A brief introduction to Physics-Informed Neural Networks and how to solve image registration problems with WarpPINN

Pablo Arratia
Francisco Sahli, Daniel Hurtado, Sergio Uribe, Hernán Mella

UNIVERSITY OF BATH

March 30, 2023

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Introduction and Motivation

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Introduction and Motivation

- In the presence of large amounts of data, machine learning has shown great success in many areas.

Introduction and Motivation

- In the presence of large amounts of data, machine learning has shown great success in many areas.
- Do we always have this amount of data?

Introduction and Motivation

- In the presence of large amounts of data, machine learning has shown great success in many areas.
- Do we always have this amount of data? No!

Introduction and Motivation

- In the presence of large amounts of data, machine learning has shown great success in many areas.
- Do we always have this amount of data? No!
- Despite that, we have a priori information in the form of physics laws that can be incorporated into the problem.

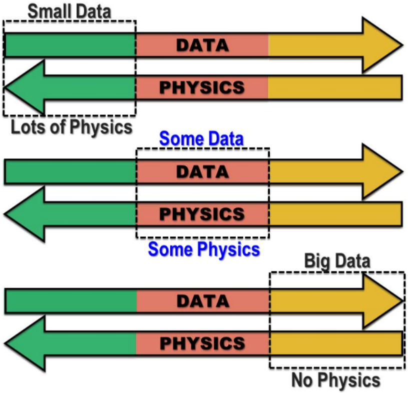

Introduction and Motivation

- In the presence of large amounts of data, machine learning has shown great success in many areas.
- Do we always have this amount of data? No!
- Despite that, we have a priori information in the form of physics laws that can be incorporated into the problem.

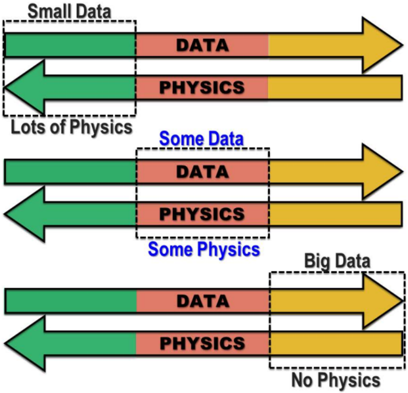

- Physics-Informed Neural Networks address this problem by leveraging the machinery of deep learning (automatic differentiation).

Physics-Informed Neural Networks and the Deep-Ritz Method

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Physics-Informed Neural Networks and the Deep-Ritz Method

- 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 (2018).
- We use Deep Learning in low-data but lots-of-physics scenarios to approximate physical quantities by leveraging the universal approximation of neural networks.

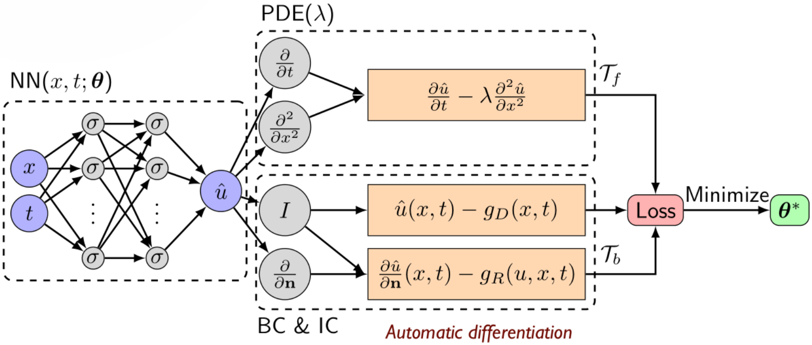

Physics-Informed Neural Networks and the Deep-Ritz Method

- Let us say we want to solve the PDE

$$\begin{cases} -\Delta u = f, & \text{in } \Omega \\ u = g, & \text{on } \partial\Omega \end{cases}$$

- The solution is approximated with a neural network u_θ that takes x as input and outputs the value $u_\theta(x)$.

$$\theta = \left\{ \left(W^i, b^i \right) \right\}_i$$

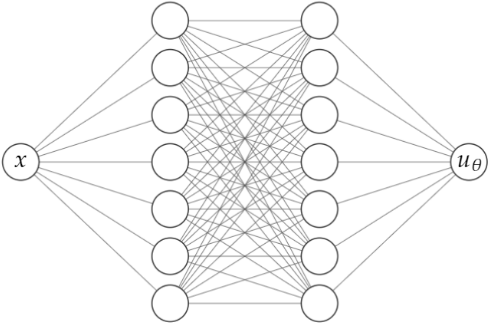

Physics-Informed Neural Networks and the Deep-Ritz Method

$$\begin{cases} -\Delta u = f, & \text{in } \Omega \\ u = g, & \text{on } \partial\Omega \end{cases}$$

- But wait, how do we train this?
- We need training points! To evaluate the boundary condition we use $\{x_i^b\}_{i=1}^{N_b} \subset \partial\Omega$. To evaluate the PDE we use collocation points $\{x_j^r\}_{j=1}^{N_r} \subset \Omega$.
- We need a loss function!
- PINNs: minimise the residual of the PDE.

$$\min_\theta \underbrace{\frac{1}{N_r}\sum_j |\Delta u_\theta(x_j^r) + f(x_j^r)|^2}_{\text{Unsupervised}} + \underbrace{\frac{1}{N_b}\sum_i |u_\theta(x_i^b) - g(x_i^b)|^2}_{\text{Supervised}}$$

Physics-Informed Neural Networks and the Deep-Ritz Method

- What about the Deep-Ritz method?

$$\begin{cases} -\Delta u = f, & \text{in } \Omega \\ u = g, & \text{on } \partial\Omega \end{cases} \iff \min_{u \in H^1_g(\Omega)} \int_\Omega \frac{1}{2} |\nabla u(x)|^2 - f(x)u(x)dx$$

- Deep-Ritz: minimise the potential of the PDE.

$$\min_\theta \underbrace{\frac{1}{N_r} \sum_j \frac{1}{2} |\nabla u_\theta(x_j^r)|^2 - f(x_j^r) u_\theta(x_j^r)}_{\text{Unsupervised}} + \underbrace{\frac{1}{N_b} \sum_i |u_\theta(x_i^b) - g(x_i^b)|^2}_{\text{Supervised}}$$

Physics-Informed Neural Networks and the Deep-Ritz Method

```
def u(x):
  u = neural_net(x, weights, biases)
  return u

def f(x):
  f = -np.pi**2 * tf.sin(np.pi*x)
  return f

def residual(x):
  u = u(x)
  u_x = tf.gradients(u,x)[0]
  u_xx = tf.gradients(u_x,x)[0]
  f = f(x)
  residual = u_xx + f
  return residual
```

Physics-Informed Neural Networks and the Deep-Ritz Method

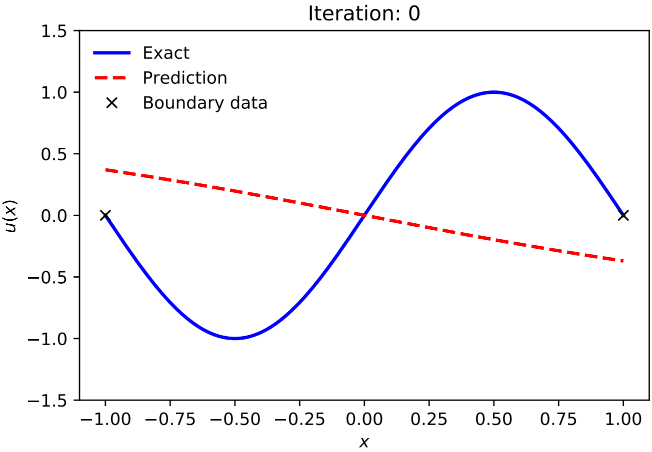

Physics-Informed Neural Networks and the Deep-Ritz Method

CAN PHYSICS-INFORMED NEURAL NETWORKS BEAT THE FINITE ELEMENT METHOD?

Tamara G. Grossmann*, Urszula Julia Komorowska†, Jonas Latz‡ and Carola-Bibiane Schönlieb*

ABSTRACT

Partial differential equations play a fundamental role in the mathematical modelling of many processes and systems in physical, biological and other sciences. To simulate such processes and systems, the solutions of PDEs often need to be approximated numerically. The finite element method, for instance, is a usual standard methodology to do so. The recent success of deep neural networks at various approximation tasks has motivated their use in the numerical solution of PDEs. These so-called physics-informed neural networks and their variants have shown to be able to successfully approximate a large range of partial differential equations. So far, physics-informed neural networks and the finite element method have mainly been studied in isolation of each other. In this work, we compare the methodologies in a systematic computational study. Indeed, we employ both methods to numerically solve various linear and nonlinear partial differential equations: Poisson in 1D, 2D, and 3D, Allen–Cahn in 1D, semilinear Schrödinger in 1D and 2D. We then compare computational costs and approximation accuracies. In terms of solution time and accuracy, physics-informed neural networks have not been able to outperform the finite element method in our study. In some experiments, they were faster at evaluating the solved PDE.

Physics-Informed Neural Networks and the Deep-Ritz Method

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations

M. Raissi [a], P. Perdikaris [b,*], G.E. Karniadakis [a]

[a] *Division of Applied Mathematics, Brown University, Providence, RI, 02912, USA*
[b] *Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA, 19104, USA*

5. Conclusions

We have introduced *physics-informed neural networks*, a new class of universal function approximators that is capable of encoding any underlying physical laws that govern a given data-set, and can be described by partial differential equations. In this work, we design data-driven algorithms for inferring solutions to general nonlinear partial differential equations, and constructing computationally efficient physics-informed surrogate models. The resulting methods showcase a series of promising results for a diverse collection of problems in computational science, and open the path for endowing deep learning with the powerful capacity of mathematical physics to model the world around us. As deep learning technology is continuing to grow rapidly both in terms of methodological and algorithmic developments, we believe that this is a timely contribution that can benefit practitioners across a wide range of scientific domains. Specific applications that can readily enjoy these benefits include, but are not limited to, data-driven forecasting of physical processes, model predictive control, multi-physics/multi-scale modeling and simulation.

We must note however that the proposed methods should not be viewed as replacements of classical numerical methods for solving partial differential equations (e.g., finite elements, spectral methods, etc.). Such methods have matured over the last 50 years and, in many cases, meet the robustness and computational efficiency standards required in practice. Our message here, as advocated in Section 3.2, is that classical methods such as the Runge–Kutta time-stepping schemes can coexist in harmony with deep neural networks, and offer invaluable intuition in constructing structured predictive

Physics-Informed Neural Networks and the Deep-Ritz Method

Main features:

- Useful in low-data regime.
- The input of the network is a point in the domain. It has a low dimensional input in comparison to typical neural networks.
- Physical constraint is imposed through PDEs or conservation laws.
- Architecture: Fully Connected. Activation function: hyperbolic tangent
- **Automatic differentiation to take derivatives of output with respect to the input.**
- **Mesh free method!** (No curse of dimensionality!)
- Easy to handle with non-linearities.
- Once trained, the solution is easy to interpolate!
- Related concepts in the literature: Neural fields, Implicit Neural Representations.

Motivation

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.
- **Cardiac Strain is an alternative metric that describes the deformation of the heart during the cardiac cycle. Simply put, it is the derivative of the deformation field.**

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.
- **Cardiac Strain is an alternative metric that describes the deformation of the heart during the cardiac cycle. Simply put, it is the derivative of the deformation field.**
- There are magnetic resonance methods used to measure cardiac deformation: tagged, phase velocity, displacement-encoded (DENSE) or strain encoded (SENC). Not suitable for clinical routines!

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.
- **Cardiac Strain is an alternative metric that describes the deformation of the heart during the cardiac cycle. Simply put, it is the derivative of the deformation field.**
- There are magnetic resonance methods used to measure cardiac deformation: tagged, phase velocity, displacement-encoded (DENSE) or strain encoded (SENC). Not suitable for clinical routines!
- **Cine SSFP MRI** is the gold standard for cardiac image acquisition but it doesn't contain information about the deformation.

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.
- **Cardiac Strain is an alternative metric that describes the deformation of the heart during the cardiac cycle. Simply put, it is the derivative of the deformation field.**
- There are magnetic resonance methods used to measure cardiac deformation: tagged, phase velocity, displacement-encoded (DENSE) or strain encoded (SENC). Not suitable for clinical routines!
- **Cine SSFP MRI** is the gold standard for cardiac image acquisition but it doesn't contain information about the deformation.
- Goal: To determine the **deformation field/cardiac strain** of the heart during the cardiac cycle from cine SSFP MRI data.

Motivation

- Cardiac diseases are one the main causes of death in the world (17.9 millions in 2016 according to the WHO).
- There are metrics used to assess cardiac activity such as Ejection Fraction.
- **Cardiac Strain is an alternative metric that describes the deformation of the heart during the cardiac cycle. Simply put, it is the derivative of the deformation field.**
- There are magnetic resonance methods used to measure cardiac deformation: tagged, phase velocity, displacement-encoded (DENSE) or strain encoded (SENC). Not suitable for clinical routines!
- **Cine SSFP MRI** is the gold standard for cardiac image acquisition but it doesn't contain information about the deformation.
- Goal: To determine the **deformation field/cardiac strain** of the heart during the cardiac cycle from cine SSFP MRI data.
- Idea: To approximate the deformation field by using a Physics-Informed Neural Network. It is trained by solving an **image registration** task on cine SSFP MR images.

Image Registration

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - **Image Registration**
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Image Registration

The image registration task

Given a reference image $R : \mathbb{R}^n \to \mathbb{R}$ and a template image $T : \mathbb{R}^n \to \mathbb{R}$, we want to find a deformation field $\varphi : \mathbb{R}^n \to \mathbb{R}^n$ such that the warped template image $T_\varphi := T \circ \varphi$ is close enough R w.r.t. some distance D:

$$\varphi^* = \operatorname{argmin}_\varphi D(R, T_\varphi)$$

.

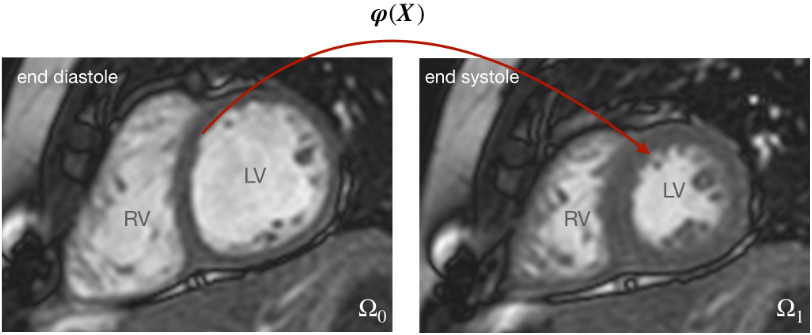

An example in 2D

- Let's take the deformation field $\varphi(X, Y) = \begin{pmatrix} X \\ Y + 0.1 \cdot X \end{pmatrix}$

Image Registration

An example in 2D

- Let's take the deformation field $\varphi(X, Y) = \begin{pmatrix} X \\ Y + 0.1 \cdot X \end{pmatrix}$
- Let's meet Bruno!

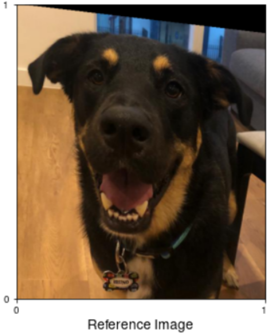

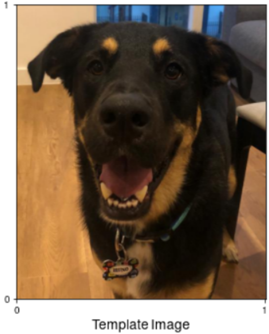

Image Registration

An example in 2D

- Let's take the deformation field $\varphi(X,Y) = \begin{pmatrix} X \\ Y + 0.1 \cdot X \end{pmatrix}$
- Let's meet Bruno!

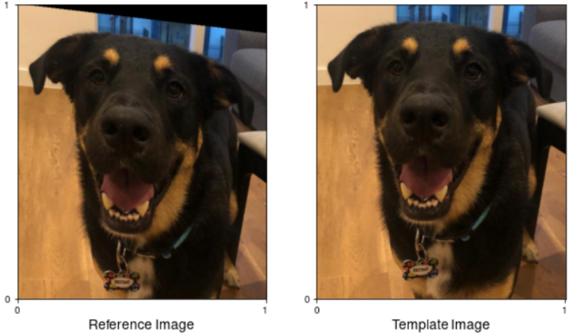

- Example: $\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} 1 \\ 0.9 \end{pmatrix}$. Then $\varphi(1, 0.9) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$.

Image Registration

An example in 2D

- Let's take the deformation field $\varphi(X, Y) = \begin{pmatrix} X \\ Y + 0.1 \cdot X \end{pmatrix}$
- Let's meet Bruno!

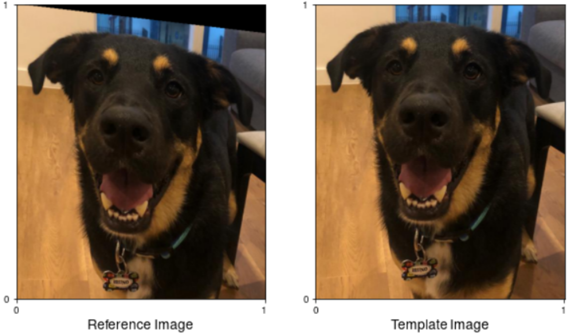

- Example: $\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} 1 \\ 0.9 \end{pmatrix}$. Then $\varphi(1, 0.9) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$.
- Finally $R(1, 0.9) = T(\varphi(1, 0.9)) = T(1, 1)$.

WarpPINN

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

WarpPINN

- We deal now with a multi-temporal image registration problem

WarpPINN

- We deal now with a multi-temporal image registration problem
- Try to approximate the deformation field $(x, y, z, t) \to \begin{pmatrix} \varphi_1(x,y,z,t) \\ \varphi_2(x,y,z,t) \\ \varphi_3(x,y,z,t) \end{pmatrix}$.

WarpPINN

- We deal now with a multi-temporal image registration problem
- Try to approximate the deformation field $(x, y, z, t) \to \begin{pmatrix} \varphi_1(x, y, z, t) \\ \varphi_2(x, y, z, t) \\ \varphi_3(x, y, z, t) \end{pmatrix}$.
- We propose *WarpPINN*: takes a point (x, y, z) in the reference frame and the time t and outputs the new position of this point at that time.

WarpPINN

- We deal now with a multi-temporal image registration problem
- Try to approximate the deformation field $(x, y, z, t) \to \begin{pmatrix} \varphi_1(x, y, z, t) \\ \varphi_2(x, y, z, t) \\ \varphi_3(x, y, z, t) \end{pmatrix}$.
- We propose *WarpPINN*: takes a point (x, y, z) in the reference frame and the time t and outputs the new position of this point at that time.

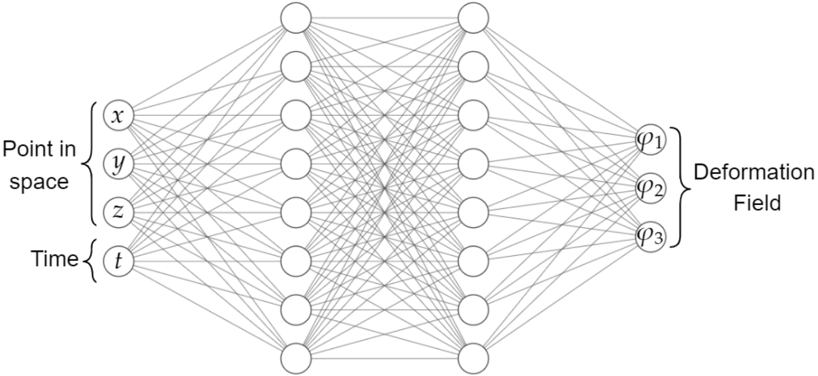

WarpPINN

- We seek to minimise

$$\min_{\varphi} \mathcal{J}(\varphi) = \|R(\boldsymbol{X}) - T \circ \varphi(\boldsymbol{X})\|^p_{L^p(\Omega)} + \alpha \mathcal{R}(\varphi).$$

WarpPINN

- We seek to minimise

$$\min_{\varphi} \mathcal{J}(\varphi) = \|R(\boldsymbol{X}) - T \circ \varphi(\boldsymbol{X})\|_{L^p(\Omega)}^p + \alpha \mathcal{R}(\varphi).$$

- Physical constraint? The cardiac tissue is nearly incompressible!

WarpPINN

- We seek to minimise

$$\min_{\varphi} \mathcal{J}(\varphi) = \|R(\boldsymbol{X}) - T \circ \varphi(\boldsymbol{X})\|^p_{L^p(\Omega)} + \alpha \mathcal{R}(\varphi).$$

- Physical constraint? The cardiac tissue is nearly incompressible!
- $J := |J_\varphi(\boldsymbol{X})| \sim 1$, for $\boldsymbol{X}$ in the cardiac tissue.

WarpPINN

- We seek to minimise

$$\min_{\varphi} \mathcal{J}(\varphi) = \|R(\boldsymbol{X}) - T \circ \varphi(\boldsymbol{X})\|_{L^p(\Omega)}^p + \alpha \mathcal{R}(\varphi).$$

- Physical constraint? The cardiac tissue is nearly incompressible!
- $J := |J_\varphi(\boldsymbol{X})| \sim 1$, for $\boldsymbol{X}$ in the cardiac tissue.
- WarpPINN minimises

$$\min_{\theta} \mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^{N} (R(\boldsymbol{X}_i) - T \circ \varphi_\theta(\boldsymbol{X}_i))^p + \alpha \mathbb{E}_{\boldsymbol{X} \in \Omega} W_{NeoHook}(\varphi_\theta)(\boldsymbol{X}),$$

where

$$W_{NeoHook}(\varphi_\theta) = \mathrm{tr}(C) - 3 - 2\log(J) + \lambda(J-1)^2; F = \frac{\partial \varphi_\theta}{\partial \boldsymbol{X}}, J = |F|, C = F^t F$$

WarpPINN

Let's see some results!

WarpPINN

A nice plot to finish

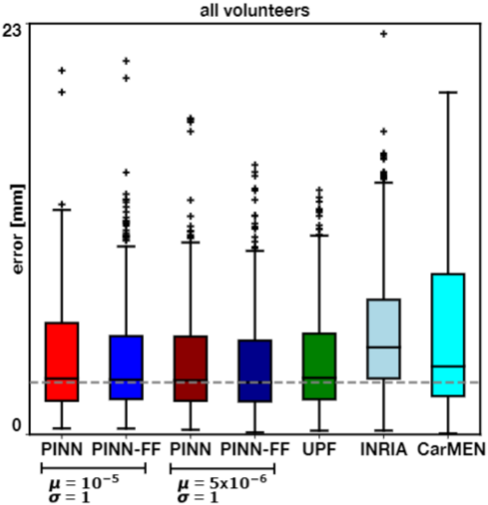

Conclusions and future work

1. Introduction
   - Introduction and Motivation
2. PINN
   - Physics-Informed Neural Networks and the Deep-Ritz Method
3. Image Registration and *WarpPINN*
   - Motivation
   - Image Registration
   - WarpPINN
4. Conclusions
   - Conclusions and future work

Conclusions and future work

- PINNs seem to be a promising new tool for the resolution of PDEs and inverse problems involving regularisers in the form of differential operators.
- We showed how to use WarpPINN to solve an image registration problem.
- WarpPINN can be employed in other image registration tasks, but how to ensure the deformation field to be a diffeomorphism?
- *Vanilla* PINNs has shown to be a versatile tool, however it has failed in some applications.

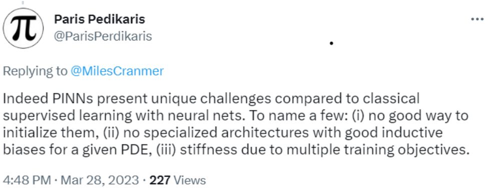

- There are too many ingredients! Initialisation of parameters, number of hidden layers, number of neurons, optimiser, sampling of collocation points, etc.
- Good news for mathematicians!

Conclusions and future work

Thanks!