



Abstract: Physics Informed Neural Networks (PINNs) have gained significant interest for solving PDE-driven systems, particularly for data assimilation, but still face many challenges. This paper introduces an approach that enhances the speed and accuracy of PINNs training with a new natural gradient scaling as $\min(P^2S, S^2P)$, where P is the number of parameters and S is the batch size. In addition, it shows an explicit connection with Green's function theory.

Physics Informed Neural Networks

Given a Partial Differential and Boundary Operator:

$$D : \begin{cases} \mathcal{H} \rightarrow L^2(\Omega \rightarrow \mathbb{R}, \mu) \\ u \mapsto D[u] \end{cases}, B : \begin{cases} \mathcal{H} \rightarrow L^2(\partial\Omega \rightarrow \mathbb{R}, \sigma) \\ u \mapsto B[u] \end{cases}$$

we aim to solve: $\begin{cases} D(u) = f \in L^2(\Omega \rightarrow \mathbb{R}, \mu) \text{ in } \Omega \\ B(u) = g \in L^2(\partial\Omega \rightarrow \mathbb{R}, \sigma) \text{ on } \partial\Omega \end{cases}$

PINNs key idea: Optimize a neural network u_θ on the loss

$$\ell(\theta) := \frac{1}{2S_D} \sum_{i=1}^{S_D} (D[u_\theta](x_i^D) - f(x_i^D))^2 + \frac{1}{2S_B} \sum_{i=1}^{S_B} (B[u_\theta](x_i^B) - g(x_i^B))^2,$$

Problem: This leads to low accuracy when using SGD.

Neural Tangent Kernel (NTK)

Jacot et al. show that for an empirical quadratic loss: $\ell(\theta) := \frac{1}{2} \sum_{i=1}^S (u_\theta(x_i) - y_i)^2$, the functional dynamic of the gradient descent on ℓ can be described by:

$$\frac{du_{\theta_t}}{dt}(x) = - \sum_{i=1}^S NTK_{\theta_t}(x, x_i) (u_{\theta_t}(x_i) - y_i),$$

with: $NTK_{\theta_t}(x, y) := \sum_{p=1}^P (\partial_p u_{\theta_t}(x)) (\partial_p u_{\theta_t}(y))^T$.

Natural Gradient

Given the functional quadratic loss: $\mathcal{L} : v \in L^2(\Omega) \mapsto \frac{1}{2} \|v - f\|_{L^2(\Omega)}^2$,

whose gradient is: $\nabla \mathcal{L}|_v = v - f$, the natural gradient update is given by:

$$\theta_{t+1} \leftarrow \theta_t - \eta du_{\theta_t}^\dagger \left(\Pi_{T_{\theta_t} \mathcal{M}}^\perp (\nabla \mathcal{L}|_{u_{\theta_t}}) \right)$$

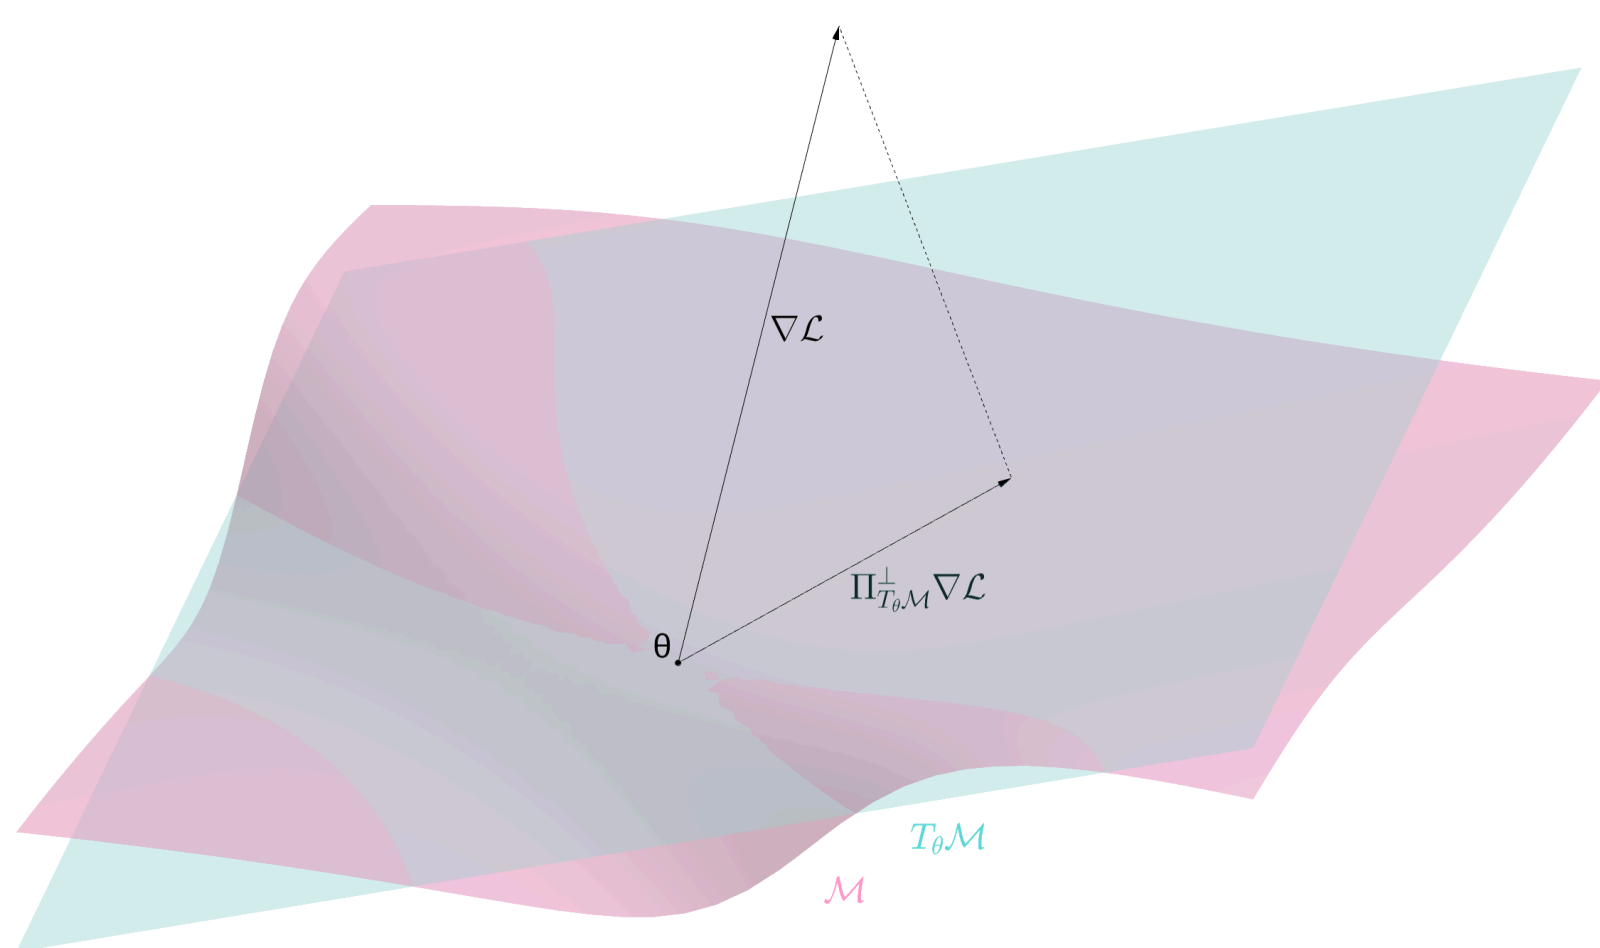


Illustration of Natural Gradient in functional space

Rudner et al. show that functional dynamics then follow the Natural NTK (NNTK):


$$\frac{du_{\theta_t}}{dt}(x) = - \sum_{i=1}^S NNTK_{\theta_t}(x, x_i) (u_{\theta_t}(x_i) - y_i), \text{ with:}$$

$$NNTK_{\theta_t}(x, y) := \sum_{1 \leq p, q \leq P} (\partial_p u_{\theta_t}(x)) G_{\theta_t, p, q}^\dagger (\partial_q u_{\theta_t}(y))^t; \quad G_{\theta_t, p, q} := \langle \partial_p u_{\theta_t}, \partial_q u_{\theta_t} \rangle_{\mathcal{H}}.$$

Lemma $\Pi_{T_{\theta_t} \mathcal{M}}^\perp$ is an integral operator whose kernel is $NNTK_{\theta_t}$.


In particular the natural gradient update rewrites:

$$\theta_{t+1} \leftarrow \theta_t - \eta G_{\theta_t}^\dagger \nabla \mathcal{L}|_{u_{\theta_t}}$$



Comparison of NTK and NNTK at initialization for Heat equation in (1+1)D.

Reading: the red cross being x_i , the plot represents the function $(N)NTK_{\theta_0}(x_i, \cdot)$



Comparison of final NTK and NNTK for Heat equation in (1+1)D.

Reading: the red cross being x_i , the plot represents the function $(N)NTK_{\theta_{\text{end}}}(x_i, \cdot)$

Empirical Natural Gradient

We introduce the **empirical Natural Gradient** as the projection of $\nabla \mathcal{L}|_{u_{\theta_t}}$ onto the **empirical Tangent Space**:

$$\hat{T}_{\theta_t}^{NNTK} \mathcal{M} := \text{Span}(NNTK_{\theta_t}(\cdot, x_i) : (x_i)_{1 \leq i \leq S}) \subset T_{\theta_t} \mathcal{M},$$

where $(x_i)_{1 \leq i \leq S}$ are given collocation points.

Theorem empirical Natural Gradient is accurately approximated by:

$$\theta_{t+1} \leftarrow \theta_t - \eta \hat{\phi}_\theta^\dagger \nabla \mathcal{L}|_{u_{\theta_t}}; \quad \hat{\phi}_{\theta_t, p} := \partial_p u_{\theta_t}(x_i); \quad \nabla \mathcal{L}|_{u_{\theta_t}} := \nabla \mathcal{L}|_{u_{\theta_t}}(x_i),$$

resulting in $O(\min(P^2S, S^2P))$ computational complexity.

PINNs are a quadratic regression

Applying our algorithm to PINNs only requires to notice that PINNs are a quadratic regression problem, using the model:

$$(D, B) \circ u : \begin{cases} \mathbb{R}^P \rightarrow \mathcal{H} \rightarrow L^2(\Omega) \times L^2(\partial\Omega) \\ \theta \mapsto u_\theta \mapsto (D[u_\theta], B[u_\theta]) \end{cases}$$




Illustration of Natural Gradient of PINNs

Connection to Green's function theory

Definition A generalized Green's function is any kernel function $g : \Omega \times \Omega \rightarrow \mathbb{R}$ such that the operator:

$$R : f \in L^2(\Omega \rightarrow \mathbb{R}, \mu) \mapsto \left(x \in \Omega \mapsto \int_{\Omega} g(x, s) f(s) \mu(ds) \right) \in \mathcal{H}$$

verifies the equation: $D \circ R = \Pi_{D[\mathcal{H}]}$

Theorem Let $D : \mathcal{H} \rightarrow L^2(\Omega \rightarrow \mathbb{R}, \mu)$ be a linear differential operator and $u : \mathbb{R}^P \rightarrow \mathcal{H}$ a parametric model. Then for all $\theta \in \mathbb{R}^P$ the generalized Green's function of D on $T_\theta \mathcal{M} = \text{Im } du_\theta$ is given by: for all $x, y \in \Omega$

$$g_{T_\theta \mathcal{M}}(x, y) := \sum_{1 \leq p, q \leq P} \partial_p u_\theta(x) G_{p, q}^\dagger \partial_q D[u_\theta](y),$$


with for all $1 \leq p, q \leq P$

$$G_{p, q} := \langle \partial_p D[u_\theta], \partial_q D[u_\theta] \rangle_{L^2(\Omega \rightarrow \mathbb{R}, \mu)}.$$

In particular, the natural gradient of PINNs can be rewritten:


$$\theta_{t+1} \leftarrow \theta_t - \eta du_{\theta_t}^\dagger \left(x \in \Omega \mapsto \int_{\Omega} g_{T_{\theta_t} \mathcal{M}}(x, y) \nabla \mathcal{L}|_{u_{\theta_t}}(y) \mu(dy) \right),$$

Interpretation: Following the natural gradient for PINNs boils down to infinitesimally moving in the direction of the solution to the linearized equation.



(a) Initialization (b) Final
Generalized Green's functions for Laplace equation in 2D.

Reading: the red cross being x_i , the plot represents the function $g_{T_{\theta_0} \mathcal{M}}(x_i, \cdot)$



(a) Initialization (b) Final
Generalized Green's functions for Heat equation in (1+1)D

Reading: the red cross being x_i , the plot represents the function $g_{T_{\theta_{\text{end}}} \mathcal{M}}(x_i, \cdot)$


Collocation points selection criterion

One of the pleasant byproducts of the empirical natural gradient formulation is that it gives a straightforward criterion for collocation points selection:


$$(x_i^*) := \underset{(x_i) \in \Omega^S}{\text{argmin}} \left\| \Pi_{\text{Span}(NNTK_{\theta_t}(x_i, \cdot) : 1 \leq i \leq S)}^\perp (\nabla \mathcal{L}|_{u_{\theta_t}}) - \nabla \mathcal{L}|_{u_{\theta_t}} \right\|_{\mathcal{H}}^2.$$

Making the most of this criterion is currently the subject of an in-depth study.


Results




Median L^2 errors and Test losses for Laplace equation in 2D



Median L^2 errors and Test losses for Heat equation in (1+1)D



Median L^2 errors and Test losses for Laplace equation in 5D



Median L^2 errors and Test losses for Allen-Cahn equation in (1+1)D

References

- Müller, J., and Zeinhofer, M. (2024). Position: Optimization in SciML Should Employ the Function Space Geometry. In Forty-First International Conference on Machine Learning.
- Müller, Johannes, and Marius Zeinhofer. "Achieving high accuracy with PINNs via energy natural gradient descent." International Conference on Machine Learning. PMLR, 2023.
- Rudner, Tim GJ, et al. "The natural neural tangent kernel: Neural network training dynamics under natural gradient descent." 4th workshop on Bayesian Deep Learning (NeurIPS 2019), 2019.
- Bai, Qinxun, Steven Rosenberg, and Wei Xu. "A geometric understanding of natural gradient." arXiv preprint arXiv:2202.06232 (2022).
- Cuomo, Salvatore, et al. "Scientific machine learning through physics-informed neural networks: Where we are and what's next." Journal of Scientific Computing 92.3 (2022): 88.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. Journal of Computational Physics, 378:686–707.
- Jacot, Arthur, Franck Gabriel, and Cl m nt Hongler. "Neural tangent kernel: Convergence and generalization in neural networks." Advances in neural information processing systems 31 (2018).
- Wang, Sifan, Xinling Yu, and Paris Perdikaris. "When and why PINNs fail to train: A neural tangent kernel perspective." Journal of Computational Physics 449 (2022): 110768.
- Paulsen, Vern I., and Mrinal Raghupathi. An introduction to the theory of reproducing kernel Hilbert spaces. Vol. 152. Cambridge university press, 2016.