Physics-Informed Neural Operators for Tissue Elasticity Reconstruction

Youjin Kim[1][0000−0001−6260−0237], Jae Yong Lee[†1][0000−0003−0193−545X], and Junseok Kwon[†1][0000−0001−9526−7549]

Chung-Ang University, 84, Heukseok-ro, Dongjak-gu, Seoul, Republic of Korea
{ju9073, jaeyong, jskwon}@cau.ac.kr

Abstract. Magnetic Resonance Elastography (MRE) is a non-invasive imaging technique that estimates tissue elasticity using Magnetic Resonance Imaging. The conventional approach for elasticity reconstruction in MRE involves solving an inverse problem through numerical methods such as Helmholtz inversion and the finite element method. However, these techniques suffer from noise sensitivity and high computational costs due to iterative optimization. Recently, Physics-Informed Neural Networks (PINNs) have been studied for tissue elasticity reconstruction, integrating physical constraints into deep learning models. While PINNs improve noise resistance, they require a separate network to be trained for each instance, resulting in a computationally inefficient training. In this study, we introduce an operator learning-based approach to tissue elasticity reconstruction, which learns a generalized mapping from input measurements to tissue elasticity. This method enables simultaneous learning across multiple instances, significantly improving computational efficiency. Experimental results using box and abdomen simulation data show that our approach achieves superior reconstruction performance and robustness to noise.

Keywords: Operator Learning · Neural Operator · Magnetic resonance elastography · Elasticity reconstruction.

1 Introduction

Magnetic Resonance Elastography (MRE) is a non-invasive imaging technique that utilizes Magnetic Resonance Imaging (MRI) to estimate tissue elasticity [16]. Quantitative assessment of tissue elasticity has been applied in various clinical domains, including the diagnosis of pathological lesions, disease progression evaluation, and treatment monitoring for conditions such as liver fibrosis, breast cancer, neurological disorders, and cardiovascular diseases [19, 24, 22, 9].

MRE estimates tissue elasticity through three stages. First, a shear wave is induced within the tissue, generating shear deformation. Next, the resulting wave displacement is measured using MRI and captured as a wave image. Finally, tissue elasticity is reconstructed by solving an inverse problem formulated from

[0] [†]Corresponding author

the elasticity equation, which defines the relationship between displacement and elasticity. Traditionally, this inverse problem has been solved using numerical methods, with Helmholtz inversion [18] and finite element method (FEM) [6] being the most common approaches. However, these methods suffer from noise sensitivity and high computational cost due to iterative optimization.

To overcome these challenges, Physics-Informed Neural Networks (PINNs) have been applied to tissue elasticity reconstruction. PINNs approximate the solution of partial differential equations (PDEs) by embedding physical constraints into the loss function, allowing them to enforce underlying physical laws while leveraging available data. In particular, MRE-PINN [20] demonstrates robustness to noise and produces smoother elastograms, highlighting its potential as an alternative approach.

However, a key limitation of PINNs is that they require retraining for each new PDE instance, making them computationally expensive. In contrast, operator learning via neural operators [10] learns a general mapping from inputs (e.g., initial/boundary conditions) to outputs, allowing multiple instances to be learned efficiently with a single model.

In this study, we present the first operator learning-based method for tissue elasticity reconstruction that directly maps diverse wave images to elasticity predictions. In contrast to PINN-based approaches, which typically require separate training for each wave image, our method eliminates this overhead and significantly improves training efficiency. Experiments with box and abdomen simulation data demonstrate superior reconstruction performance and noise robustness, suggesting its effectiveness in elasticity reconstruction tasks.

2 Related works

PINNs. PINNs incorporate physical laws into neural networks to solve PDEs, addressing the limitations of traditional solvers and data-driven methods [21, 8]. By enforcing PDE constraints in the loss function, they provide noise-robust solutions and handle complex or incomplete data effectively. However, PINNs must be retrained for each new set of initial and boundary conditions, leading to high computational costs and limited scalability for large or diverse datasets.
Operator Learning. Operator learning [10] has been developed to overcome the limitations of numerical methods and neural PDE solvers, offering a scalable approach to learning mappings between function spaces. As the primary implementation of operator learning, neural operators generalize across multiple problem instances without requiring separate training, making them more computationally efficient than PINNs. Among various neural operator architectures, DeepONet [14] is one of the most well-known approaches. It employs a branch-trunk architecture to approximate nonlinear operators, and several extensions have further enhanced its capabilities [15, 11, 2, 3]. However, although neural operators like DeepONet exhibit strong generalization capabilities across different conditions, they struggle to enforce physical constraints, often resulting in physically inconsistent solutions. To overcome this, physics-informed neural

operators have been developed, integrating PDE constraints into the learning process to ensure adherence to physical laws while maintaining generalization ability [13, 23, 12].

Deep Learning Approaches for Inverse Problems. Many deep learning methods have been proposed for solving inverse problems, utilizing different techniques to tackle key challenges [17, 5, 4, 7]. MRE-PINN [20] was the first to apply PINNs to tissue elasticity reconstruction, using two PINN modules—one for displacement and another for elasticity—trained with a PDE-based loss. While MRE-PINN improves noise robustness and produces smoother elastograms compared to numerical methods, its per-instance training requirement results in high computational costs.

3 Proposed Method

We propose **MRE-Hyper**, a HyperDeepONet-based model specifically designed for tissue elasticity reconstruction. In this section, we first present our baseline models, DeepONet [14] and HyperDeepONet [11], followed by a detailed explanation of the proposed MRE-Hyper model and its training strategy.

Notations. The goal is to learn the operator $\mathcal{G}$, which maps an input function u to its corresponding output function: $\mathcal{G} : u \mapsto \mathcal{G}(u)$. The input function $u(x)$ is defined on the domain $x \in X \subset \mathbb{R}^{d_x}$. The output function is evaluated on the domain $y \in Y \subset \mathbb{R}^{d_y}$, where the output function is denoted as $\mathcal{G}(u)(y)$.

Baselines. DeepONet approximates $\mathcal{G}$ using a branch-trunk structure. The branch network $\mathcal{N}_{Branch}$ maps the input u to a feature vector, encoding the input function into a finite-dimensional feature space. The trunk network $\mathcal{N}_{Trunk}$ maps y to a basis function representation, encoding the query point in terms of the learned basis functions. The operator is then computed as follows:

$$\mathcal{G}_\theta(u)(y) = \sum_{k=1}^{p} \mathcal{N}_{Branch}(u;\theta)_k \mathcal{N}_{Trunk}(y;\theta)_k, \tag{1}$$

where p is the number of basis functions used in the approximation. However, the expressivity of the model is limited due to the linear combination of the branch and trunk network outputs. To overcome this limitation, HyperDeepONet replaces the branch network with a hypernetwork $\mathcal{H}$ that generates the parameters Θ of a target network $\mathcal{T}$:

$$\Theta = \mathcal{H}(u;\theta), \quad \mathcal{G}_\theta(u)(y) = \mathcal{T}(y;\Theta). \tag{2}$$

Here, θ is the learnable parameters of networks in general, while Θ refers to the parameters generated by $\mathcal{H}$ which are then used in $\mathcal{T}$. This allows HyperDeepONet to capture complex operators with greater expressivity than DeepONet, making it particularly effective for tissue elasticity reconstruction, where sharp variations in the elasticity solution require a more sophisticated approach.

Model. The proposed MRE-Hyper comprises two HyperDeepONet networks, Hyper-u and Hyper-μ, as shown in Fig. 1. Both networks take as input u_{sample},

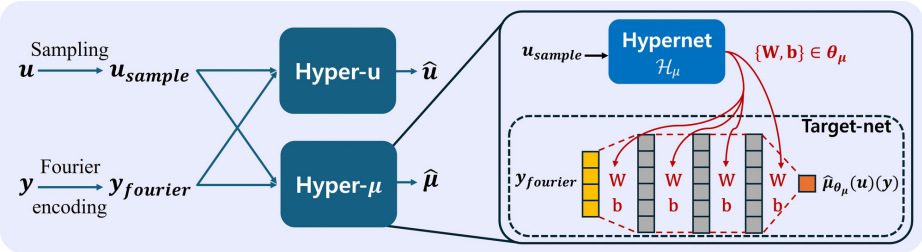

Fig. 1: **Overall architecture of the proposed MRE-Hyper.**

subsampled from the wave image u, and y_{fourier}, the Fourier-encoded spatial coordinates of the target points. Each HyperDeepONet consists of a hypernetwork $\mathcal{H}$ and a target network $\mathcal{T}$. The hypernetwork encodes u_{sample} and generates a set of parameters Θ, which the target network then applies to y_{fourier} to predict the corresponding function values.

For Hyper-u, the hypernetwork $\mathcal{H}_u$ generates parameters Θ_u, which the target network $\mathcal{T}_u$ uses to predict the continuous displacement function:

$$\Theta_u = \mathcal{H}_u(u_{\text{sample}}; \theta_u), \quad \hat{u}_{\theta_u}(y) = \mathcal{T}_u(y_{\text{fourier}}; \Theta_u). \tag{3}$$

Similarly, for Hyper-μ, the hypernetwork $\mathcal{H}_\mu$ generates parameters Θ_μ, enabling the target network $\mathcal{T}_\mu$ to predict the elasticity function:

$$\Theta_\mu = \mathcal{H}_\mu(u_{\text{sample}}; \theta_\mu), \quad \hat{\mu}_{\theta_\mu}(y) = \mathcal{T}_\mu(y_{\text{fourier}}; \Theta_\mu). \tag{4}$$

Inspired by [23], we incorporate Fourier encoding to enhance function approximation. By employing two HyperDeepONet modules, MRE-Hyper achieves efficient and expressive reconstruction of tissue elasticity.

Training. The model adopts a physics-informed neural operator framework, where the physical knowledge of linear elasticity is incorporated through a PDE loss to enforce physical consistency. Unlike conventional operator learning, which relies on data-driven supervision, this approach leverages physics-based constraints, enabling operator learning even in the absence of explicit training data. The training process follows a sequential learning strategy. First, Hyper-u is trained using a reconstruction loss $\mathcal{L}_{\text{recon}}$:

$$\mathcal{L}_{\text{recon}} = \frac{1}{mn} \sum_{j=1}^{m} \sum_{i=1}^{n} \left\| \hat{u}_{\theta_u}(u_j)(y_i) - u_j(y_i) \right\|^2. \tag{5}$$

where n is the number of query points used to evaluate the target function and m is the number of wave images. Subsequently, Hyper-u and Hyper-μ are jointly optimized with a PDE residual loss $\mathcal{L}_{\text{pde}}$ derived from linear elasticity theory. For simplicity, we denote $\hat{\mu}_{\theta_\mu}(u_j)(y_i)$ and $\hat{u}_{\theta_u}(u_j)(y_i)$ as $\hat{\mu}_{ji}$ and $\hat{u}_{ji}$, respectively, leading to the following form:

$$\mathcal{L}_{\text{pde}} = \frac{1}{mn} \sum_{j=1}^{m} \sum_{i=1}^{n} \left\| \hat{\mu}_{ji} \nabla^2 \hat{u}_{ji} + (\nabla \hat{u}_{ji} + \nabla \hat{u}_{ji}^T) \nabla \hat{\mu}_{ji} + \rho \omega^2 \hat{u}_{ji} \right\|^2. \tag{6}$$

Here, ρ is the mass density of the material, and ω is the actuator frequency in MRE. This formulation allows Hyper-μ to learn the mapping to elasticity without requiring ground-truth μ supervision.

4 Experiments

We conducted experiments on two datasets, comparing the proposed MRE-Hyper with both PINN-based methods and numerical approaches.

Datasets. We used two publicly available FEM-based simulation datasets from the BIOQIC research group [1]. FEM-Box dataset consists of numerical FEM simulations of elastic wave propagation in an incompressible rectangular domain containing four stiff inclusions of decreasing size. It provides six wave fields of dimensions $80 \times 100 \times 10 \times 3$ and corresponding ground-truth (GT) elasticity maps, generated at frequencies from 50 to 100 Hz in 10 Hz increments. The FEM-Abdomen dataset comprises numerical FEM simulations of elastic wave propagation in a 3D human abdomen model with known tissue properties. It includes four wave fields with dimensions $174 \times 136 \times 20 \times 3$ and corresponding GT elasticity, generated at frequencies ranging from 30 to 48 Hz in 6 Hz increments.

Compared Methods. We compared our MRE-Hyper model with several existing methods. One of them is Algebraic Helmholtz Inversion (AHI) [18], which estimates the shear modulus by solving the Helmholtz equation as a linear system after computing the Laplacian of the wave image using finite differences. We also evaluated FEM-HH [6], which directly inverts the Helmholtz equation using the Finite Element Method (FEM), and FEM-het [6], which incorporates heterogeneous material assumptions into the PDE inversion. Additionally, we compared MRE-PINN, a PINN-based model for elasticity reconstruction. Standard HyperDeepONet and DeepONet could not be directly applied to this task without modifications, so they were excluded from the comparison.

Evaluation Metrics. We evaluated the Pearson correlation between the reconstructed elasticity and the GT elasticity, referred to as the μ correlation. For the FEM-Box dataset, we also computed the Contrast Transfer Efficiency (CTE), which measures the ratio of the contrast in predicted elasticity to that in GT elasticity. Following [20], CTE was computed for each of the four target regions, and the CTE error was defined as the mean absolute deviation of these values from the ideal value of 100%.

Noise Robustness. Since real MRE data is often noisy, robustness to noise is a crucial factor in evaluating elasticity reconstruction methods. Following MRE-PINN [20], we added Gaussian noise at five different levels to the wave images and evaluated the elasticity reconstruction performance of each method.

Implementation Details. Each reconstruction method was applied to individual single-frequency wave images. We trained MRE-PINN with independent networks for each wave image, whereas our MRE-Hyper was trained on all wave images simultaneously. All networks were implemented using PyTorch 2.1.2. Our code is available for reproducibility. For details, please refer to the repository.[1]

[1] https://github.com/youjin-DDAI/MRE-Hyper

Table 1: **Quantitative comparisons in terms of μ correlations and CTE error.**

Freq	AHI		FEM-HH		FEM-het		MRE-PINN		MRE-Hyper (Ours)	
(Hz)	$\mu\uparrow$	error $\downarrow$	$\mu\uparrow$	error $\downarrow$	$\mu\uparrow$	error $\downarrow$	$\mu\uparrow$	error $\downarrow$	$\mu\uparrow$	error $\downarrow$
50	0.61	66	0.38	87	0.32	81	**0.79**	**59**	0.76	60
60	0.54	93	0.48	75	0.37	78	**0.84**	53	0.81	**48**
70	0.56	60	0.54	53	0.38	51	0.75	56	**0.81**	**50**
80	0.58	73	0.57	54	0.46	**41**	**0.76**	67	0.74	66
90	0.44	79	0.55	54	0.36	65	0.72	**49**	**0.82**	**49**
100	0.46	59	0.63	57	0.53	**44**	0.71	80	**0.80**	53
Mean	0.532	71.7	0.525	63.3	0.423	60.0	0.762	60.7	**0.790**	**54**

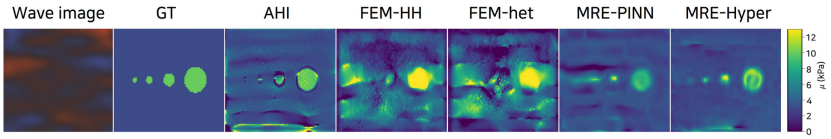

Fig. 2: **Qualitative comparison of μ reconstructions at 90 Hz.**

4.1 Experiments on FEM-Box Data

Experimental Results. Table 1 presents the μ correlations and CTE errors across different frequencies, where the best results are highlighted in bold. The results indicate that deep learning-based methods, MRE-PINN and MRE-Hyper, achieve higher correlations and lower CTE errors compared to numerical approaches. Notably, MRE-Hyper outperformed MRE-PINN, achieving the highest mean correlation of 0.790 while maintaining consistently strong correlations across all frequencies. The proposed MRE-Hyper also achieved the lowest mean CTE error, surpassing all other methods. It consistently produced lower errors, confirming its effectiveness in reconstruction. These findings are further supported qualitatively in Fig. 2.

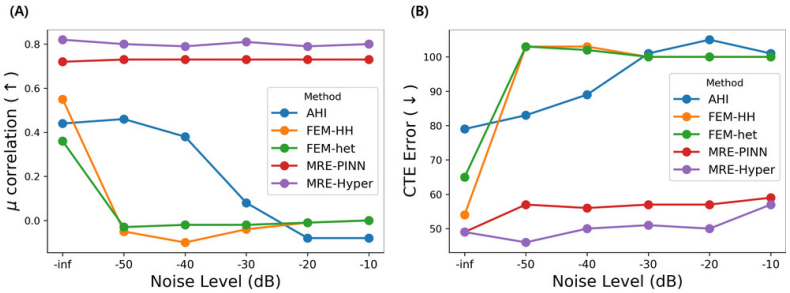

(a) μ correlation (b) CTE error

Fig. 3: **Performance comparisons under increasing noise levels.**

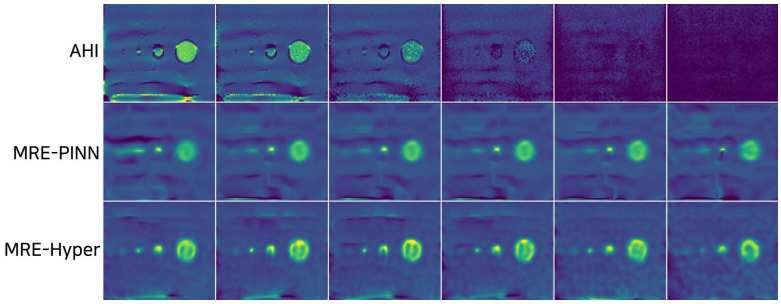

Fig. 4: **Qualitative comparison of μ reconstructions at 90 Hz under increasing noise levels.** Columns correspond to increasing noise levels: -inf, -50, -40, -30, -20, -10 dB.

Robustness to Noise. Fig. 3(a) illustrates the μ correlations of each method at 90 Hz as noise levels increase. While the deep learning-based methods remained largely unaffected by increasing noise, the numerical methods, AHI, FEM-HH, and FEM-het, experienced a significant performance drop. Among them, FEM-based methods exhibited particularly severe degradation even with a small amount of noise. Fig. 3(b) shows the CTE error at 90 Hz under increasing noise levels. Similar to the μ correlation results, the deep learning-based methods maintained consistently low errors with minimal variation, whereas numerical methods exhibited clear sensitivity to noise. The robustness of the proposed MRE-Hyper method against noise is further demonstrated in Fig. 4.

4.2 Experiments on FEM-Abdomen Data

Experimental Results. Table 2 presents the μ correlation for each frequency, computed only for soft tissues by excluding bones with elasticity below 10 kPa. Correlation values for FEM-HH and FEM-het were NaN and thus omitted. The results indicate that MRE-Hyper consistently achieves the highest μ correlation across all frequencies. Additionally, as shown in Fig. 6, MRE-Hyper closely follows the GT elasticity trend.

Robustness to Noise. Fig. 5 visualizes how μ correlation varies with noise levels at 36 Hz. FEM-HH and FEM-Het were excluded from the graph as their correlation values were NaN. The results suggest that deep learning-based methods are more robust to noise than the numerical method AHI, with MRE-Hyper maintaining high performance even at high noise levels.

4.3 Discussion on Computational Time

Table 3 presents the computational times and requirements of each method for elasticity reconstruction on the FEM-Abdomen dataset. AHI and FEM require no training but rely on iterative solvers, leading to higher inference times. MRE-PINN and MRE-Hyper, on the other hand, require an initial training phase but

Freq (Hz)	AHI	MRE-PINN	MRE-Hyper
30	0.57	0.55	**0.79**
36	0.67	0.68	**0.81**
42	0.66	0.70	**0.83**
48	0.68	0.77	**0.84**
Mean	0.645	0.675	**0.818**

Table 2: μ **correlations.**

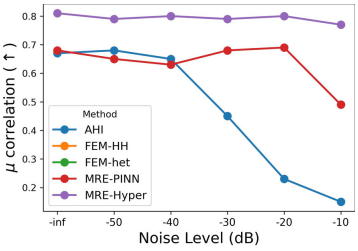

Fig. 5: μ **correlations at 36 Hz under increasing noise levels.**

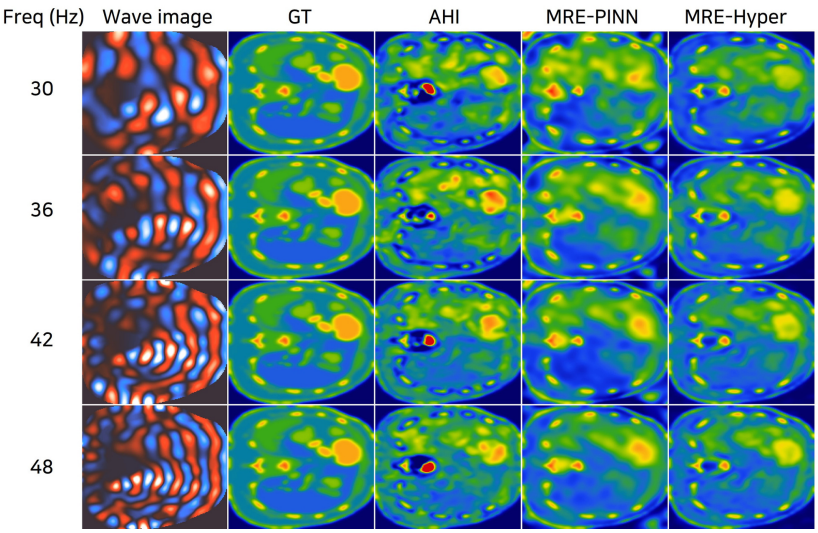

Fig. 6: **Qualitative comparisons of μ reconstructions.**

Table 3: **Comparison of computational times and requirements on FEM-Abdomen data**

	AHI	FEM	MRE-PINN	MRE-Hyper
Requires Initial Training	No	No	Yes	Yes
Inference Time (s)	1.1	20.03	0.04	0.18
Generalization to New Instances	Yes	Yes	No	Yes

enable faster inference. Notably, MRE-PINN necessitates retraining for each new instance, whereas MRE-Hyper generalizes across different wave images without additional training. Although MRE-Hyper's per-inference time is slightly longer

than MRE-PINN's, its scalability for new instances provides a clear computational advantage in real-world scenarios.

5 Conclusion

In this study, we introduced MRE-Hyper, a physics-informed neural operator framework that learns a generalized mapping from wave images to elasticity by incorporating PDE constraints. Our approach overcomes the limitations of conventional methods by enabling efficient learning across multiple instances while eliminating susceptibility to noise. Experimental results support the superiority of the proposed method, highlighting operator learning as a promising approach for tissue elasticity reconstruction.

Acknowledgments. This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. https://bioqic-apps.charite.de/downloads
2. Azizzadenesheli, K., Kovachki, N., Li, Z., Liu-Schiaffini, M., Kossaifi, J., Anandkumar, A.: Neural operators for accelerating scientific simulations and design. Nature Reviews Physics **6**(5), 320–328 (2024)
3. Cao, S.: Choose a transformer: Fourier or galerkin. In: NeurIPS (2021)
4. Cho, D., Yun, H., Lee, J., Lim, M.: Conformal mapping coordinates physics-informed neural networks (coco-pinns): learning neural networks for designing neutral inclusions. arXiv preprint arXiv:2501.07809 (2025)
5. Cho, S.W., Son, H.: Physics-informed deep inverse operator networks for solving PDE inverse problems. In: ICLR (2025)
6. Honarvar, M.: Dynamic elastography with finite element-based inversion. Ph.D. thesis, University of British Columbia (2015)
7. Hwang, R., Lee, J.Y., Shin, J.Y., Hwang, H.J.: Solving pde-constrained control problems using operator learning. In: AAAI (2022)
8. Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L.: Physics-informed machine learning. Nature Reviews Physics **3**(6), 422–440 (2021)
9. Kolipaka, A., Araoz, P.A., McGee, K.P., Manduca, A., Ehman, R.L.: Magnetic resonance elastography as a method for the assessment of effective myocardial stiffness throughout the cardiac cycle. Magnetic resonance in medicine **64**(3), 862–870 (2010)
10. Kovachki, N.B., Lanthaler, S., Stuart, A.M.: Operator learning: Algorithms and analysis. arXiv preprint arXiv:2402.15715 (2024)
11. Lee, J.Y., Cho, S., Hwang, H.J.: HyperDeepONet: learning operator with complex target function space using the limited resources via hypernetwork. In: ICLR (2023)

12. Lee, J.Y., Ko, S., Hong, Y.: Finite Element Operator Network for Solving Elliptic-Type Parametric PDEs. SIAM J. Sci. Comput. **47**(2), C501–C528 (2025). https://doi.org/10.1137/23M1623707
13. Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. ACM/JMS Journal of Data Science **1**(3), 1–27 (2024)
14. Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. Nature Machine Intelligence **3**(3), 218–229 (2021)
15. Lu, L., Meng, X., Cai, S., Mao, Z., Goswami, S., Zhang, Z., Karniadakis, G.E.: A comprehensive and fair comparison of two neural operators (with practical extensions) based on fair data. Computer Methods in Applied Mechanics and Engineering **393**, 114778 (2022)
16. Manduca, A., Oliphant, T.E., Dresner, M.A., Mahowald, J., Kruse, S.A., Amromin, E., Felmlee, J.P., Greenleaf, J.F., Ehman, R.L.: Magnetic resonance elastography: non-invasive mapping of tissue elasticity. Medical image analysis **5**(4), 237–254 (2001)
17. Molinaro, R., Yang, Y., Engquist, B., Mishra, S.: Neural inverse operators for solving PDE inverse problems. In: ICML (2023)
18. Papazoglou, S., Hamhaber, U., Braun, J., Sack, I.: Algebraic helmholtz inversion in planar magnetic resonance elastography. Physics in Medicine & Biology **53**(12), 3147 (2008)
19. Petitclerc, L., Sebastiani, G., Gilbert, G., Cloutier, G., Tang, A.: Liver fibrosis: Review of current imaging and mri quantification techniques. Journal of Magnetic Resonance Imaging **45**(5), 1276–1295 (2017)
20. Ragoza, M., Batmanghelich, K.: Physics-informed neural networks for tissue elasticity reconstruction in magnetic resonance elastography. In: MICCAI (2023)
21. Raissi, M., Perdikaris, P., Karniadakis, G.E.: 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 (2019)
22. Sinkus, R., Siegmann, K., Xydeas, T., Tanter, M., Claussen, C., Fink, M.: Mr elastography of breast lesions: understanding the solid/liquid duality can improve the specificity of contrast-enhanced mr mammography. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine **58**(6), 1135–1144 (2007)
23. Wang, S., Wang, H., Perdikaris, P.: Learning the solution operator of parametric partial differential equations with physics-informed deeponets. Science advances **7**(40), eabi8605 (2021)
24. Wuerfel, J., Paul, F., Beierbach, B., Hamhaber, U., Klatt, D., Papazoglou, S., Zipp, F., Martus, P., Braun, J., Sack, I.: Mr-elastography reveals degradation of tissue integrity in multiple sclerosis. Neuroimage **49**(3), 2520–2525 (2010)