

Convolutional Deep Operator Networks for Learning Nonlinear Focused Ultrasound Wave Propagation in Heterogeneous Spinal Cord Anatomy

Avisha Kumar¹, Xuzhe Zhi¹, Zan Ahmad¹, Minglang Yin², Amir Manbachi^{1,2}

¹Johns Hopkins University, Baltimore, MD, USA

²Johns Hopkins University School of Medicine, Baltimore, MD, USA

Abstract

Focused ultrasound (FUS) therapy is a promising tool for optimally targeted treatment of spinal cord injuries (SCI), offering submillimeter precision to enhance blood flow at injury sites while minimizing impact on surrounding tissues. However, its efficacy is highly sensitive to the placement of the ultrasound source, as the spinal cord’s complex geometry and acoustic heterogeneity distort and attenuate the FUS signal. Current approaches rely on computer simulations to solve the governing wave propagation equations and compute patient-specific pressure maps using ultrasound images of the spinal cord anatomy. While accurate, these high-fidelity simulations are computationally intensive, taking up to hours to complete parameter sweeps, which is impractical for real-time surgical decision-making. To address this bottleneck, we propose a convolutional deep operator network (DeepONet) to rapidly predict FUS pressure fields in patient spinal cords. Unlike conventional neural networks, DeepONets are well equipped to approximate the solution operator of the parametric partial differential equations (PDEs) that govern the behavior of FUS waves with varying initial and boundary conditions (i.e., new transducer locations or spinal cord geometries) without requiring extensive simulations. Trained on simulated pressure maps across diverse patient anatomies, this surrogate model achieves real-time predictions with only a 2% loss on the test set, significantly accelerating the modeling of nonlinear physical systems in heterogeneous domains. By facilitating rapid parameter sweeps in surgical settings, this work provides a crucial step toward precise and individualized solutions in neurosurgical treatments.

Code and Dataset —

<https://github.com/avishakumar21/nonlinear-fus-with-neural-operators>

Introduction

Spinal cord injury (SCI) often leads to significant physiological issues following the initial trauma, including prolonged reduced blood flow to the injury site (hematoma) (Quadri et al. 2020). Current treatments, like surgical decompression, cannot precisely restore blood flow to the hematoma for optimal healing (Lonjaret et al. 2014; Ahuja et al. 2017). Focused ultrasound (FUS) therapy offers a targeted solution by generating acoustic pressure to enhance localized blood flow directly to a focal region (Hwang et al. 2021; Tshay

et al. 2022; Morishita et al. 2014; Hong et al. 2022). When administering FUS therapy, clinicians must ensure that the focal region of the resulting FUS pressure distribution maximally overlaps with the hematoma with minimal exposure to surrounding tissue for optimal therapeutic benefit. However, predicting FUS pressure fields in the spinal cord relative to transducer positions is difficult due to its complex geometric variations and acoustic heterogeneity that cause beam distortion (Kumar et al. 2023b).

Computational Bottleneck

Computer simulations that numerically solve the governing FUS wave propagation equations can determine the pressure distribution through patient-specific spinal cord anatomy using ultrasound images (Treeby et al. 2018; Kumar et al. 2023a). This pressure distribution is crucial for non-invasively inferring the biological effects of a specific transducer placement, especially when treating a sensitive organ such as the spinal cord. To achieve optimal therapeutic results, all possible source locations along the region of interest in the spine must be evaluated. However, these predictions cannot be made prior to surgery, as a laminectomy is required to remove surrounding vertebral bone and create an acoustic window for ultrasound imaging. Other imaging methods, such as magnetic resonance imaging (MRI) and computed tomography (CT), do not adequately capture all the soft tissues within the spinal cord that have varying acoustic properties, which are essential for accurate simulations. Since these simulations can take several minutes to hours per patient, it is important to develop a faster method for solving the FUS wave equations in patient-specific spinal cord to support timely intraoperative decision-making.

Previous works in predicting ultrasound behavior includes physics-informed neural networks (PINNs) on ultrasound acoustic wave data to simulate regions with surface-breaking cracks and PINNs for transcranial ultrasound wave propagation (Wang et al. 2023; Shukla et al. 2020). These methods embed physical constraints of the acoustic wave equation into the loss equation of their models to ensure that predictions align closely with the underlying physics, improving accuracy in scenarios where direct measurements or high-quality imaging data may be limited (Raissi, Perdikaris, and Karniadakis 2019). While these approaches address the limitations of data-driven models that require extensive, diverse

training datasets and overcome the challenges of traditional mesh-based methods—such as time-intensive calculations for high-dimensional problems—they must be retrained for each new input domain. This retraining is impractical for FUS treatment planning, as each patient’s spinal cord geometry is unique, especially after sustaining an injury.

Operator Learning for PDEs

Neural operator learning is a framework for learning mappings between infinite-dimensional function spaces, such as the solution generators for systems of partial differential equations (PDEs). This approach differs from traditional neural networks, which are designed to learn mappings between finite-dimensional vector spaces. By learning the underlying relationships between input functions (initial and boundary conditions) and their corresponding PDE solutions, neural operators aim to generalize across families of PDEs without requiring retraining for each new configuration (Kovachki et al. 2023).

Mathematically, a neural operator \mathcal{G}_θ maps an input function space $\mathcal{U}(\Omega_\alpha; \mathbb{R}^{d_u})$ to an output function space $\mathcal{V}(\Omega_\alpha; \mathbb{R}^{d_v})$, as follows:

$$\mathcal{G}_\theta : \mathcal{U}(\Omega_\alpha; \mathbb{R}^{d_u}) \rightarrow \mathcal{V}(\Omega_\alpha; \mathbb{R}^{d_v}), \quad (1)$$

where $\Omega_\alpha \subset \mathbb{R}^d$ represents the spatial domain and $\alpha \in \mathcal{A}$ parametrizes its shape. The parameters of the neural operator, $\theta \in \Theta$, are optimized during training. The dimensions d_u and d_v refer to the sizes of the input and output function spaces, which may be subspaces of Sobolev spaces or spaces of continuous functions. These function spaces provide a broad mathematical setting for applications where the PDE solutions vary smoothly or have certain regularity properties.

The true solution operator \mathcal{G} is approximated by \mathcal{G}_θ using training data in the form of input-output pairs, $\{u_i, v_i\}_{i=1}^N$, where $u_i \in \mathcal{U}$ represents an input function, and $v_i = \mathcal{G}(u_i) \in \mathcal{V}$ is the corresponding PDE solution. These pairs are typically generated from high-fidelity numerical simulations of the governing equations. For example, in modeling focused ultrasound (FUS) pressure fields in the spinal cord, \mathcal{U} may represent the initial and boundary conditions of the acoustic wave equation in a patient-specific spinal cord geometry, while \mathcal{V} corresponds to the resulting pressure distribution in the tissue.

By learning a direct mapping from inputs to solutions, neural operators can predict the solution to a PDE without solving it numerically, offering significant computational advantages. This makes them well-suited for applications requiring rapid predictions, such as real-time surgical planning for FUS treatment.

Proposed Framework

The proposed network architecture used in this work combines convolutional neural networks (CNNs) (Anwar et al. 2018) and deep operator networks (DeepONets) (Lu et al. 2021) to learn a generalizable mapping between images of the spinal cord anatomy and transducer/source location (input functions) and the resulting pressure fields across

the spatial domain after FUS sonication (output functions). Trained on simulated pressure maps from diverse patient anatomies, the proposed model predicts pressure distributions across varying spinal cord geometries in real-time, achieving just 2% error on the test set. While operator learning has been explored in medical contexts (Loeffler et al. 2024; Zhou et al. 2024; Yin et al. 2022), this paper introduces its application for real-time optimization of neurosurgical interventions.

Methods

Data Generation

Learning a surrogate model to accelerate simulations requires an abundantly diverse and expressive training dataset, with both input images and corresponding ground truth simulation results. To train our neural operator, we generated simulated pressure maps using 1,000 sagittal B-mode ultrasound images of the thoracic spinal cord from 25 porcine subjects, captured both before and after contusion injuries (Kumar et al. 2024). Each image ($25.6\text{mm} \times 8.1\text{mm}$) displayed anatomical structures with varying acoustic properties, including the dorsal space, dura, cerebrospinal fluid (CSF), pia, spinal cord, ventral space, and injury site (Figure 1A). Ground truth masks for these soft-tissue regions were used to define the computational domain to generate patient-specific acoustic phantoms of the spinal cords (Figure 1B). For some images, due to the acquisition angle or the severity of injury, the anatomical boundaries between the dura, CSF, and pia or between the ventral dura and ventral space were indistinct. In these cases, the regions were grouped and labeled as the dura/pia complex or dura/ventral complex, respectively.

Computer Simulations

We use `k-wave`, a MATLAB toolbox for modeling acoustic wave propagation, to generate the simulated pressure maps (Treeby et al. 2018). By solving a system of first-order partial differential equations (PDEs) computationally equivalent to the generalized Westervelt equation, `k-wave` is designed to account for nonlinearities introduced to wave propagation caused by high magnitude acoustic waves common in biomedical ultrasonics. The framework also models acoustically heterogeneous media, such as the spinal cord, where sound speed and ambient density vary spatially. The governing PDEs (Equations 2 – 4) are solved at each time step via pseudo-spectral methods to obtain the pressure of the system at each grid location,

$$\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{\rho_0} \nabla p, \quad (2)$$

$$\frac{\partial \rho}{\partial t} = -(2\rho + \rho_0) \nabla \cdot \mathbf{u} - \mathbf{u} \cdot \nabla \rho_0, \quad (3)$$

$$p = c_0^2 \left(\rho + \mathbf{d} \cdot \nabla \rho_0 + \frac{B}{2A} \frac{\rho^2}{\rho_0} - L\rho \right), \quad (4)$$

where \mathbf{u} is the acoustic particle velocity, ρ is the acoustic density, p is the acoustic pressure, ρ_0 is ambient density, c_0 is the sound speed, and d is the acoustic particle displacement.

Table 1: Acoustic properties of spinal cord anatomies of interest (Hasgall et al. 2022)

Material	Sound Speed (m/s)	Density (kg/m ³)	Attenuation Constant (dB/MHz ^y cm)	Acoustic Nonlinearity (B/A)
Dorsal Space	1578.2	1050	0.2152	6.11
Dura	1500.0	1174	1.1641	6.72
CSF	1504.5	1007	0.0087	4.96
Pia	1500.0	1174	1.1641	6.72
Dura/Pia Complex	1500.0	1174	1.1641	6.72
Spinal Cord	1542.0	1075	0.778	6.72
Hematoma	1560.1	1062.5	0.4968	6.42
Dura/Ventral Complex	2064.0	1339	1.77642	7.00
Ventral Space	2577.2	1504	2.38875	7.28

L is a linear integro-differential operator used in k -Wave that accounts for acoustic absorption and dispersion.

The acoustic properties for each anatomical region were obtained with a literature search for physiological accuracy and are summarized in Table 1 (Hasgall et al. 2022). Due to limited data in the literature, the following assumptions were made. The hematoma properties were estimated by averaging the acoustic values of the spinal cord and blood to represent the effects of ruptured blood vessels during injury. The pia is assumed to have the same acoustic properties as the dura because of their similar echogenicity. Acoustic nonlinearity parameters for the dura, CSF, and spinal cord were derived from analogous biological tissues. Specifically, the dura, pia, and spinal cord were given the nonlinearity parameter of brain tissue, while the CSF was modeled using the nonlinearity parameter of water. The ventral space itself was assumed to exhibit properties averaging those of bone and cartilage, and the dura/ventral complex was assumed to be average values of dura and ventral space. The attenuation constant for all these media was defined by a power-law absorption exponent (y) of 1.05, indicating that dispersion effects were insignificant in the simulations.

After calibrating our patient-specific computational grids (mesh size: 512×162 grid points, element spacing: $5e-5$) to exhibit acoustic properties of the spinal cord, a single element transducer emitting a continuous sine wave was configured based on characteristics of typical FUS transducers. The operating frequency of the transducer was set to 2.5 MHz, with a sonication time of $1e-4$ seconds and a focal length of 5 mm. The source location was systematically varied across 8 positions, spaced 3 mm apart along the sagittal axis of the spine. The corresponding pressure distributions were recorded at each position, providing a comprehensive discretization of possible pressure maps across the entire imaged anatomy.

These pressure maps (Figure 1C) serve as the ground truth simulation dataset for the proposed neural operator model, which was divided into training, validation, and test sets with an 80-10-10 split. There was no overlap of patient spinal cord images between the sets to prevent data leakage and overfitting.

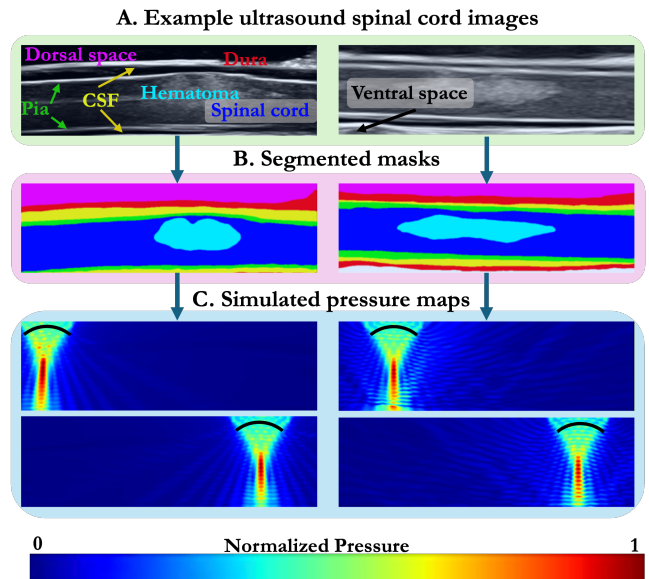


Figure 1: Example ultrasound spinal cord images in the dataset along with their segmented masks and visualizations of corresponding pressure maps at various source locations.

Network Architecture

Our proposed network architecture features a convolutional DeepONet, which is well suited for our problem given that we have a fixed rectangular domain Ω for all training samples. In this context, we allow the pressure distributions from the FUS simulations to depend on 2 critical parameters:

1. The unique spinal cord geometry of each patient which prescribes the acoustic heterogeneity in the parametric PDE.
2. The specific transducer locations at which the pressure field is evaluated.

To account for these dependencies, our model incorporates two branch networks to encode the input function space and one trunk network to encode the output domain. The branch networks process patient-specific geometric and




Figure 2: Visualization of proposed model architecture. This model consists of two branch nets for encoding the input function space (spinal cord image and transducer location) and a trunk net for encoding the domain of the output functions (discrete coordinates at which they are evaluated). The outputs of these deep neural networks are then merged with a dot product to approximate the true solution generator \mathcal{G} with the neural operator \mathcal{G}_θ . The loss is minimized to obtain the optimal set of parameters (θ) for the solution operator during the training process.

source-specific information, while the trunk network encodes the grid locations of the output pressure distribution.

Let the governing wave equation be represented succinctly by:

$$\mathcal{L}(p; \mathbf{u}, \mathbf{z}) = f, \quad \text{in } \Omega,$$

where:

- $p \in \mathcal{V}(\Omega)$ is the acoustic pressure distribution,
- $\mathbf{u} \in \mathcal{U}_1$ represents the patient-specific spinal cord geometry (e.g., anatomical masks),
- $\mathbf{z} \in \mathcal{U}_2$ represents the transducer location parameters (e.g., (x, y)),
- $\Omega \subset \mathbb{R}^2$ is the rectangular computational domain.

The neural operator approximates the solution operator \mathcal{G} : $\mathcal{U}_1 \times \mathcal{U}_2 \rightarrow \mathcal{V}$, such that:

$$p = \mathcal{G}(\mathbf{u}, \mathbf{z}).$$

The complete network architecture (depicted in Figure 2) consists of the following components:

- **Branch Network 1:** A CNN with three convolutional layers, which encodes the patient-specific spinal cord geometry \mathbf{u} . The CNN processes segmented anatomical masks from ultrasound images and outputs a latent representation $\mathbf{g}_{\text{geo}} \in \mathbb{R}^q$.
- **Branch Network 2:** A fully connected neural network (FCNN) with three hidden layers, which encodes the transducer location \mathbf{z} . This network outputs a latent representation $\mathbf{g}_{\text{src}} \in \mathbb{R}^r$.
- **Trunk Network:** A fully connected neural network (FCNN) with three hidden layers, which encodes the grid

locations $\mathbf{x} \in \Omega$. The output of the trunk network is a spatial representation $\mathbf{h}_{\text{spatial}} \in \mathbb{R}^s$.

The outputs of the branch and trunk networks are combined using a Hadamard product, resulting in a continuous representation of the predicted pressure map:

$$p(\mathbf{x}) \approx \mathcal{G}_\theta(\mathbf{u}, \mathbf{z})(\mathbf{x}) = \sum_{i=1}^n (\mathbf{g}_{\text{src}, i} \cdot \mathbf{g}_{\text{geo}, i}) \cdot \mathbf{h}_{\text{spatial}, i}(\mathbf{x}).$$

The model is trained using simulation data generated by *k-Wave*, with the goal of minimizing the relative L_2 loss:

$$\mathcal{L} = \frac{\|p_{\text{pred}} - p_{\text{true}}\|_2}{\|p_{\text{true}}\|_2}.$$

The Adam optimizer with a step-based learning rate schedule is employed for training. By efficiently approximating the operator \mathcal{G} , this architecture maps patient-specific anatomical features and transducer parameters to acoustic pressure distributions. This enables rapid parameter sweeps, facilitating optimal source placement for treatment.

We compare the proposed operator network against two baseline models to evaluate performance in predicting pressure maps in heterogeneous spinal cord anatomy. The first baseline is a CNN model, similar to the branch net in the DeepONet architecture, consisting of three convolutional layers with additional layers for upsampling the output to match the dimensions of the expected simulation. The second baseline is a Fully Convolutional Network (FCN) with regression, a deep learning architecture typically designed for dense, pixel-wise prediction tasks, such as semantic segmentation, where continuous values are predicted

Table 2: Characteristics of deep learning models for approximating focused ultrasound pressure distribution in patient-specific spinal cord anatomy.

Model	Trainable Parameters	Learning Rate	Batch Size	Training Epochs	Optimizer	Loss function
Proposed Neural Operator	5,286,656	0.001	4	1000	Adam	Relative L_2
Baseline CNN	23,649	0.01	32	100	Adam	Relative L_2
Baseline FCN	22,833,473	0.01	32	100	Adam	Relative L_2

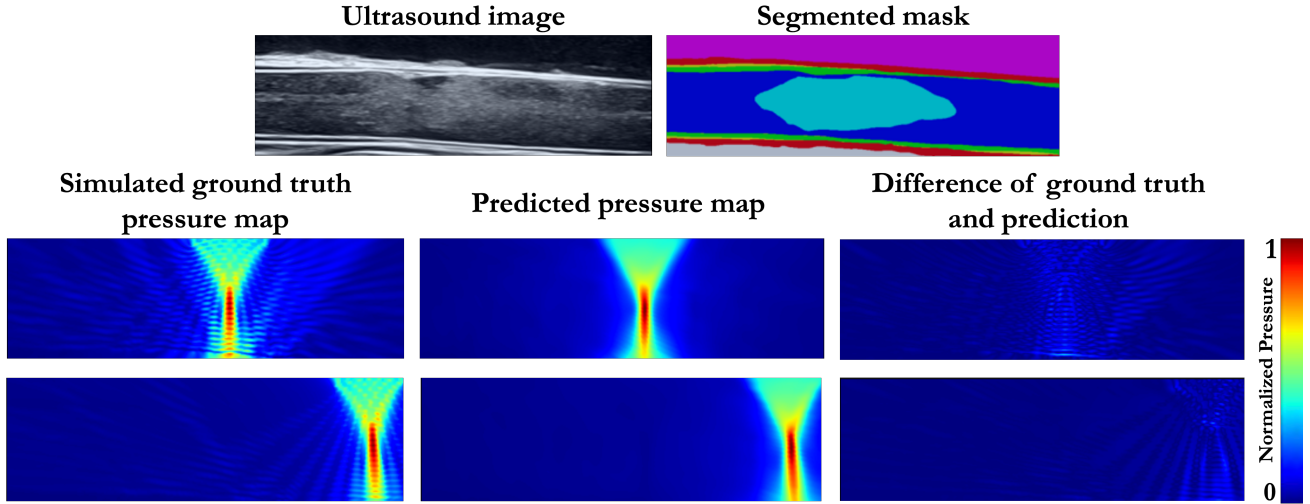


Figure 3: Visualization of results from the proposed neural operator model, including the ultrasound image of the injured spinal cord, the corresponding segmented mask used as input for the neural operator, high-fidelity ground truth pressure maps at source locations 4 and 8 obtained from k -Wave, the model’s predicted pressure maps for those locations, and the difference between the predictions and the ground truth.

for each pixel (Long, Shelhamer, and Darrell 2015; Sofka et al. 2017). Our FCN is a pretrained ResNet34 with up-sampling layers for dimension matching. These models were trained on a Windows 11 Machine (8 GB RAM) with 24 GB NVIDIA GeForce RTX 3090 graphics unit and 14th Gen Intel Core i5-14600 processor (14 cores, 20 threads, 2.7GHZ to 5.2GHZ turbo frequency). The details of these models are provided in Table 2, and the hyperparameters are tuned to optimize performance for each model.

Results

After developing the simulation training data and the task-specific deep learning models for predicting FUS pressure maps, each model was evaluated on unseen patient spinal cord geometries in the test set. The results are summarized in Table 3, with our proposed neural operator model outperforming the other deep learning models with only a 2% loss in the test set evaluated against all 8 locations (Figure 4). The visualization of the predictions made by the proposed neural operator model is presented in Figure 3. We also assessed the performance of CNN and FCN models trained on a single source location as a baseline comparison against the DeepONet model. From the test loss and prediction re-




Figure 4: The loss for the training and validation set recorded at each training epoch for the proposed neural operator model.

Table 3: Performance comparison of models trained on either all source locations or on a single source location (position 4) for generating pressure maps in heterogeneous patient-specific spinal cords.

Model and Task	Relative L2 Loss on Validation Set
Proposed neural operator - All source location	0.024
Baseline CNN - Source location 4	0.647
Baseline FCN - Source location 4	0.024
Baseline FCN - All source locations	0.687

Table 4: Time comparison between methods for generating 8 pressure maps across a heterogeneous patient-specific spinal cord.

Method	Time for generating 8 pressure maps
Proposed neural operator	0.05 seconds
k-Wave (numerical solver)	76.1 minutes

sults (Figure 5), it is evident that the baseline CNN model fails to learn FUS pressure distributions in new spinal cord anatomy, with a loss of 64% even when trained on a single source location. Although the baseline FCN model trained on a single source location achieves a low loss of 2%, it proves impractical for location sweeps without embedding the source location into its architecture (69% loss).

Additionally, we compare the time required to generate 8 pressure maps using predictions from the proposed model versus high-fidelity numerical solutions computed with *k-Wave* on a macOS system equipped with an Apple M1 chip, featuring an 8-core CPU with 4 performance cores and 4 efficiency cores, 8-core GPU, and 8 GB of unified memory. This benchmarking provides insights into the computational efficiency of the model on consumer-grade hardware. Our findings indicate that the inference on our model takes 0.05 seconds, while obtaining the pressure maps using the numerical solver *k-Wave*, takes 76.1 total minutes (Table 4).

Discussion

Numerical modeling is a promising tool for approximating the behavior of physical systems in complex media by solving the governing physical equations iteratively in a specified computational grid. These simulations allow clinicians to infer the biological effects of FUS in patient-specific anatomy, providing valuable insight in individualized treatment planning (i.e., where to place the source). However, the computational demands of these simulations limit their utility for real-time parameter sweeps needed to optimize in-

Simulated ground truth pressure map




Figure 5: Visualization of the predicted pressure maps from the baseline comparison models for source location 4, including results from a convolutional neural network (CNN) trained only on source location 4, and fully convolutional networks (FCNs) trained on source location 4 and all source locations.

traoperative decision-making. High-fidelity ultrasound simulations using traditional finite difference and finite element methods require a resolution of 10 mesh points per acoustic wavelength (Zingg, Lomax, and Jurgens 1996). This results in prohibitively large computational grids, especially when modeling high-frequency biomedical transducers operating in the MHz range. This challenge is particularly significant in the context of spinal cord injury, where preoperative ultrasound imaging is not feasible until vertebral bone removal, requiring simulations to be performed intraoperatively. While MRI and CT imaging are available, they fail to provide adequate soft-tissue delineation required for the acoustic solvers. Ultrasound imaging, on the other hand, inherently captures the acoustic differences within the spinal cord, which are crucial for accurately predicting the propagation of the nonlinear FUS beam.

Our proposed neural operator architecture, which integrates CNNs and DeepONets, demonstrates the ability to learn the solution operator for nonlinear FUS wave propagation in heterogeneous, patient-specific spinal cord anatomy, enabling real-time prediction of pressure maps. Unlike traditional neural networks which excel at finding patterns in data, operator learning frameworks provide a more principled approach that results in a model which respects the underlying continuous structure of the true PDE solution generator. This allows for predictions to be made based on various input functions without need for retraining or running expensive simulations repeatedly. In this case, the nonlinear

operator is the mapping from a space of functions (patient spinal cord geometry) to another space of functions (acoustic pressure distribution). Accelerating FUS simulations can be transformative across several medical domains, such as glioblastoma or deep vein thrombosis, offering a novel approach to optimize ultrasound transducer placements for targeting tumors or blood clots while minimizing exposure to adjacent healthy tissue (Zhou 2011).

While a FCN can learn pressure distributions across the spinal cord for a single source location with model loss comparable to the proposed approach, it lacks the ability to generalize effectively across multiple source locations. As a result, it is unsuitable for parameter sweeps to determine the optimal source location. To achieve comparable generalization capabilities with this architecture, it would require training with source location embeddings, likely increasing the computational and resource burden. This limitation is further exacerbated by the fact that the baseline FCN model is more than 4 times larger than the proposed neural operator (Table 2).

Our proposed model efficiently and accurately learns the relationship between patient-specific anatomy and the corresponding *in vivo* pressure distribution for various transducer placements. The online computing time of this network is over **91,000** times faster than traditional simulations (Table 4), demonstrating the practical applicability of operator learning in surgical medicine, particularly for optimizing the use of therapeutic ultrasound.

Given the similarities in soft-tissue morphology, vasculature, and immune response after injury between porcine and human spinal cord anatomy, we anticipate that this model can generalize effectively to the human spinal cord (Toossi et al. 2021). Fine-tuning the final layers with human data would further enhance its performance and applicability. Moving forward, we aim to enhance this model by integrating a segmentation model into the preprocessing pipeline, replacing the use of the soft-tissue masks as inputs. This improvement will streamline intraoperative decision-making by eliminating the need for the tedious and manual segmentation currently required. Future efforts will focus on adapting the model to accept raw ultrasound images as inputs to the branch network instead of segmented masks, while maintaining high performance. Additionally, we plan to expand the training and evaluation process to include human spinal cord images, enhancing the clinical utility of this tool.

Conclusion

The deployment of advanced computational models in predictive ultrasound modeling holds potential to revolutionize personalized treatment planning by dramatically enhancing time efficiency and enabling scalable parameter sweeps at a fraction of the computational cost. Traditional numerical models, while highly accurate and physics-based, face a critical limitation: each new input, such as transducer location or patient anatomy, requires a complete recomputation of the solution, a computational luxury that cannot be afforded in time-sensitive environments like the operating room. In this paper, we introduced a deep operator network capable of predicting nonlinear focused ultrasound pressure maps in

patient-specific, heterogeneous spinal cord anatomy, offering a paradigm-shifting approach to accelerating intraoperative treatment planning. With the advent of high-resolution ultrasound imaging and the increasing interest in focused ultrasound therapy for noninvasive treatments, the application of deep learning to develop efficient surrogate models can have far-reaching implications in healthcare, including tumor ablation and blood clot removal. This work underscores the immense promise of AI-driven solutions to address critical challenges in clinical decision-making and treatment optimization.

Acknowledgments

Amir Manbachi acknowledges funding support from the Defense Advanced Research Projects Agency (DARPA) (award #N660012024075), along with the Johns Hopkins Institute for Clinical and Translational Research (ICTR)’s Clinical Research Scholars Program (KL2), administered by the National Center for Advancing Translational Sciences (NCATS), National Institute of Health (NIH), as well as the National Science Foundation (NSF)’s Small Businesses Technology Transfer (STTR) (award #1938939).

References

- Ahuja, C. S.; Wilson, J. R.; Nori, S.; Kotter, M.; Druschel, C.; Curt, A.; and Fehlings, M. G. 2017. Traumatic spinal cord injury. *Nature reviews Disease primers*, 3(1): 1–21.
- Anwar, S. M.; Majid, M.; Qayyum, A.; Awais, M.; Alnowami, M.; and Khan, M. K. 2018. Medical image analysis using convolutional neural networks: a review. *Journal of medical systems*, 42: 1–13.
- Hasgall, P.; Di Gennaro, F.; Baumgartner, C.; Neufeld, E.; Lloyd, B.; Gosselin, M.; Payne, D.; Klingensböck, A.; and Kuster, N. 2022. IT’IS Database for thermal and electromagnetic parameters of biological tissues.
- Hong, Y.-r.; Lee, E.-h.; Park, K.-s.; Han, M.; Kim, K.-T.; and Park, J. 2022. Ultrasound stimulation improves inflammatory resolution, neuroprotection, and functional recovery after spinal cord injury. *Scientific Reports*, 12(1): 3636.
- Hwang, B. Y.; Mampre, D.; Ahmed, A. K.; Suk, I.; Anderson, W. S.; Manbachi, A.; and Theodore, N. 2021. Ultrasound in traumatic spinal cord injury: a wide-open field. *Neurosurgery*, 89(3): 372–382.
- Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2023. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89): 1–97.
- Kumar, A.; Kotkar, K.; Jiang, K.; Bhimreddy, M.; Davidar, D.; Weber-Levine, C.; Krishnan, S.; Kerensky, M. J.; Liang, R.; Leadingham, K. K.; et al. 2024. A novel open-source ultrasound dataset with deep learning benchmarks for spinal cord injury localization and anatomical segmentation. *arXiv preprint arXiv:2409.16441*.
- Kumar, A.; Punnoose, J.; Leadingham, K. M. K.; Kerensky, M. J.; Theodore, N.; Thakor, N. V.; and Manbachi, A. 2023a.

- A Patient-specific Preplanning Treatment Algorithm for Focused Ultrasound Therapy of Spinal Cord Injury. In *2023 11th International IEEE/EMBS Conference on Neural Engineering (NER)*, 1–4. IEEE.
- Kumar, A.; Tsehay, Y.; Gonzalez, E.; Kerensky, M. J.; Bell, M. A. L.; Theodore, N.; Thakor, N. V.; and Manbachi, A. 2023b. Computational modeling towards focused ultrasound therapy for spinal cord injury: visualization of beam propagation through patient-specific anatomy. In *Medical Imaging 2023: Image-Guided Procedures, Robotic Interventions, and Modeling*, volume 12466, 276–282. SPIE.
- Loeffler, S. E.; Ahmad, Z.; Ali, S. Y.; Yamamoto, C.; Popescu, D. M.; Yee, A.; Lal, Y.; Trayanova, N.; and Maggioni, M. 2024. Graph Fourier Neural Kernels (G-FuNK): Learning Solutions of Nonlinear Diffusive Parametric PDEs on Multiple Domains. *arXiv preprint arXiv:2410.04655*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3431–3440.
- Lonjaret, L.; Lairez, O.; Minville, V.; and Geeraerts, T. 2014. Optimal perioperative management of arterial blood pressure. *Integrated blood pressure control*, 49–59.
- Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; and Karniadakis, G. E. 2021. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3): 218–229.
- Morishita, K.; Karasuno, H.; Yokoi, Y.; Morozumi, K.; Ogi-hara, H.; Ito, T.; Fujiwara, T.; Fujimoto, T.; and Abe, K. 2014. Effects of therapeutic ultrasound on intramuscular blood circulation and oxygen dynamics. *Journal of the Japanese Physical Therapy Association*, 17(1): 1–7.
- Quadri, S. A.; Farooqui, M.; Ikram, A.; Zafar, A.; Khan, M. A.; Suriya, S. S.; Claus, C. F.; Fiani, B.; Rahman, M.; Ramachandran, A.; et al. 2020. Recent update on basic mechanisms of spinal cord injury. *Neurosurgical Review*, 43: 425–441.
- 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.
- Shukla, K.; Di Leoni, P. C.; Blackshire, J.; Sparkman, D.; and Karniadakis, G. E. 2020. Physics-informed neural network for ultrasound nondestructive quantification of surface breaking cracks. *Journal of Nondestructive Evaluation*, 39: 1–20.
- Sofka, M.; Milletari, F.; Jia, J.; and Rothberg, A. 2017. Fully convolutional regression network for accurate detection of measurement points. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*, 258–266. Springer.
- Toossi, A.; Bergin, B.; Marefatallah, M.; Parhizi, B.; Tyreman, N.; Everaert, D. G.; Rezaei, S.; Seres, P.; Gatenby, J. C.; Perlmutter, S. I.; et al. 2021. Comparative neuroanatomy of the lumbosacral spinal cord of the rat, cat, pig, monkey, and human. *Scientific Reports*, 11(1): 1955.
- Treeby, B. E.; Budisky, J.; Wise, E. S.; Jaros, J.; and Cox, B. 2018. Rapid calculation of acoustic fields from arbitrary continuous-wave sources. *The Journal of the Acoustical Society of America*, 143(1): 529–537.
- Tsehay, Y.; Weber-Levine, C.; Kim, T.; Chara, A.; Alomari, S.; Awosika, T.; Liu, A.; Ehresman, J.; Lehner, K.; Hwang, B.; et al. 2022. Advances in monitoring for acute spinal cord injury: a narrative review of current literature. *The Spine Journal*, 22(8): 1372–1387.
- Wang, L.; Wang, H.; Liang, L.; Li, J.; Zeng, Z.; and Liu, Y. 2023. Physics-informed neural networks for transcranial ultrasound wave propagation. *Ultrasonics*, 132: 107026.
- Yin, M.; Ban, E.; Rego, B. V.; Zhang, E.; Cavinato, C.; Humphrey, J. D.; and Em Karniadakis, G. 2022. Simulating progressive intramural damage leading to aortic dissection using DeepONet: an operator–regression neural network. *Journal of the Royal Society Interface*, 19(187): 20210670.
- Zhou, T.; Wan, X.; Huang, D. Z.; Li, Z.; Peng, Z.; Anandkumar, A.; Brady, J. F.; Sternberg, P. W.; and Daraio, C. 2024. AI-aided geometric design of anti-infection catheters. *Science Advances*, 10(1): eadj1741.
- Zhou, Y.-F. 2011. High intensity focused ultrasound in clinical tumor ablation. *World journal of clinical oncology*, 2(1): 8.
- Zingg, D. W.; Lomax, H.; and Jurgens, H. 1996. High-accuracy finite-difference schemes for linear wave propagation. *SIAM Journal on Scientific Computing*, 17(2): 328–346.