

Neural-Augmented Kelvinlet for Real-Time Soft Tissue Deformation Modeling

Ashkan Shahbazi^{1*}, Kyvia Pereira^{1*}, Jon S. Heiselman¹, Elaheh Akbari¹, Annie C. Benson¹,
Sepehr Seifi¹, Xinyuan Liu¹, Garrison Lawrence Horswill Johnston¹, Jie Ying Wu¹,
Nabil Simaan¹, Michael Miga¹, Soheil Kolouri¹

¹Vanderbilt University, USA

{ashkan.shahbazi, kyvia.pereira, jon.s.heiselman, elaheh.akbari, annie.c.benson,
sepehr.seifi, xinyuan.liu, garrison.l.johnston, jieying.wu, nabil.simaan, michael.i.miga, soheil.kolouri}@vanderbilt.edu

Abstract

Accurate and efficient modeling of soft-tissue interactions is fundamental for advancing surgical simulation, surgical robotics, and model-based surgical automation. To achieve real-time latency, classical Finite Element Method (FEM) solvers are often replaced with neural approximations; however, naively training such models in a fully data-driven manner without incorporating physical priors frequently leads to poor generalization and physically implausible predictions. We present a novel physics-informed neural simulation framework that enables real-time prediction of soft-tissue deformations under complex single- and multi-grasper interactions. Our approach integrates Kelvinlet-based analytical priors with large-scale FEM data, capturing both linear and nonlinear tissue responses. This hybrid design improves predictive accuracy and physical plausibility across diverse neural architectures while maintaining the low-latency performance required for interactive applications. We validate our method on challenging surgical manipulation tasks involving standard laparoscopic grasping tools, demonstrating substantial improvements in deformation fidelity and temporal stability over existing baselines. These results establish Kelvinlet-augmented learning as a principled and computationally efficient paradigm for real-time, physics-aware soft-tissue simulation in surgical AI.

Code and data — <https://github.com/mint-vu/neural-augmented-kelvinlet-leap>

Introduction

Soft tissue deformation modeling is essential in computational biomechanics for surgical simulation, image registration, and robotic-assisted interventions. Accurate tissue simulation enables surgical planning, image-guided predictions of soft-tissue changes, and precise haptic feedback for robotics. Traditional approaches rely on physics-based methods grounded in continuum mechanics, such as the Finite Element Method (FEM), Boundary Element Method (BEM), and Tensor Mass Model (TMM). While highly accurate and physically realistic, these methods are often computationally prohibitive for real-time applications (Allard,

Courtecuisse, and Faure 2012; Zhang, Zhong, and Gu 2019). To reduce computational cost, approximate methods such as mass-spring systems (Nedel and Thalmann 1998), mesh-free methods (Belytschko et al. 1996), and particle-based methods (Liu et al. 2015) are frequently used. However, these approximations often sacrifice realism, creating a significant simulation-to-reality gap that limits their practicality in high-fidelity applications.

Recent advances in deep learning and neural operators have enabled fully data-driven methods that approximate soft tissue deformations with a single forward pass of a learned model. Convolutional neural networks (CNNs), graph neural networks (GNNs), and Transformers (Wu, Kazanzides, and Unberath 2020; Wu et al. 2021; Pfaff et al. 2021) have demonstrated the ability to capture complex deformation patterns from simulated or real-world data, enabling real-time tissue simulation. Despite these advances, fully data-driven approaches face two fundamental limitations: (a) limited generalization to unseen force distributions and boundary conditions, which results in unreliable extrapolation, and (b) a tendency to produce physically inconsistent deformations, particularly when trained on datasets of limited size or diversity. These shortcomings highlight the need for hybrid approaches that combine the expressiveness of neural models with the robustness of physics-based priors.

To address these challenges, considerable efforts have focused on integrating physics priors into data-driven machine learning models (Karniadakis et al. 2021). Broadly, these methods make learning algorithms physics-informed by introducing: (a) observational biases, such as data augmentations that embed physical constraints; (b) inductive biases in the model architecture, for example, using group-invariant or equivariant neural networks (Bronstein et al. 2017); and (c) learning biases through the design of loss functions that explicitly enforce physical laws. Despite this progress, achieving the necessary low latency for real-time simulation while maintaining a high-performing physics-informed model with appropriately chosen biases remains a challenging and largely open problem.

In this paper, we model soft tissue using a volumetric mesh and address the problem of predicting the displacement of the entire mesh given the observed displacements

*Equal contribution

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

at a limited set of interaction points on its surface. This is analogous to determining how the entire soft tissue will deform when graspers engage multiple surface locations and are moved to specified target positions. While several prior studies have explored similar setup and demonstrated that neural networks can approximate the solutions produced by high-fidelity physics-based solvers such as FEM (Pfeiffer et al. 2019), these approaches often lack strong physical priors and can suffer from reduced generalization and physical inconsistency. To overcome these limitations, we propose a physics-informed neural framework that integrates analytical elastic priors to enhance both accuracy and physical plausibility.

We introduce a physics-informed neural framework that leverages Kelvinlets (De Goes and James 2017), i.e., the closed-form Green’s function solutions to the Navier–Cauchy equations, as an analytically tractable and computationally efficient prior. Kelvinlets provide a fast, low-order approximation of elastic deformation that captures the fundamental physics of linear elasticity. We build on this prior through two complementary hybrid strategies: (a) *residual learning*: using Kelvinlets to generate a baseline displacement field and training a neural network to model the residual error; and (b) *physics regularization*: incorporating Kelvinlets as a soft constraint within the learning objective of a neural network that predicts the full deformation field. By explicitly coupling the expressiveness of neural networks with the efficiency and physical grounding of Kelvinlets, our approach achieves improved accuracy, generalization, and physical plausibility for soft tissue deformation modeling.

To establish a unified testbed for evaluating our approach alongside prior methods, we generated a finite element dataset of liver deformations using both a linear model without preload and a nonlinear model with gravity-induced preload. We show that our proposed *Neural-Augmented Kelvinlet* framework surpasses purely data-driven models in generalization, physical realism, and learning efficiency, ultimately enabling real-time, high-fidelity modeling of soft tissue deformations involving multiple graspers.

Contributions. Our key contributions are:

1. We present a dataset of 20,800 FEM solutions for data-driven soft tissue modeling, with 10,400 linear and 10,400 nonlinear simulations. It includes 5,000 single-grasper and 5,400 multi-grasper cases.
2. We provide a GPU-accelerated Kelvinlet PyTorch implementation that enables efficient training and real-time inference, making physics-informed deformation modeling accessible for interactive applications.
3. We introduce two hybrid approaches that integrate Kelvinlets into data-driven soft tissue deformation modeling and conduct extensive experiments to demonstrate their superiority over purely physics-informed and data-driven baselines across diverse neural architectures.

Related Work

Classical approaches based on the Finite Element Method (FEM) provide high-fidelity simulations of soft tissue deformation but are often computationally prohibitive for real-time surgical applications. To address this limitation, model reduction techniques such as Proper Orthogonal Decomposition (POD) have been proposed to accelerate FEM without sacrificing physical realism. For example, Lau et al. (Liu et al. 2025) employ a reduced-order linear FEM with modal warping to achieve sub-15 ms inference for liver deformation, enabling intraoperative use. Complementarily, Hu and Desai (Hu and Desai 2004) characterize nonlinear liver material behavior through indentation experiments and validate Local Effective Moduli (LEM) estimates using hyperelastic FEM models, laying foundational work for realistic soft-tissue mechanics under large strains.

More recently, physics-informed neural networks (PINNs) have emerged as a data-efficient alternative that integrates governing equations such as elasticity or dynamics directly into the learning objective. Hu et al. (Hu et al. 2025) introduce a PINN that encodes static linear elasticity to predict 3D stress and strain fields of liver tissue under surgical tool interaction, achieving millisecond-scale inference while retaining physical plausibility. Extending this paradigm, Nguyen-Le et al. (Nguyen-Le, Ballit, and Dao 2025) propose a PINODE framework that augments neural dynamics with mass-spring priors, enabling extrapolation from sparse motion data. These methods improve interpretability and generalization but often require complex training and remain constrained by solver latency or limited scalability.

In contrast, purely data-driven surrogate models have been proposed to directly regress deformation fields from control variables. Ke et al. (Camara et al. 2016) train a feed-forward network to predict liver surface displacements conditioned on retractor pose, achieving sub-millisecond inference. While fast and simple, such models lack physical constraints, which can reduce reliability and robustness in unseen conditions.

On the learning-based front, mesh-based models such as PhysGNN leverage FEM supervision to train graph neural networks for real-time deformation prediction in neurosurgery (Salehi and Giannacopoulos 2022). PEGNN (Saleh et al. 2024) extends this with contact-aware hybrid graph reasoning, enabling accurate predictions across surgical and manipulation scenarios. Meanwhile, SeeSaw (Docea et al. 2024) employs stereo video and self-supervised learning to estimate sparse 3D deformations using a GNN encoder. Zhao et al. (Gong et al. 2024) introduce a diffeomorphic mapping framework that enforces cycle consistency, capturing invertible tissue motion. These approaches show strong empirical performance but may suffer from limited interpretability or reliance on dense supervision.

Neural architectures such as U-Net (Mendizabal, Márquez-Neila, and Cotin 2020), Set Attention Block (SAB) (Lee et al. 2019), and GraphGPS (Rampásek et al. 2022) offer modular and scalable backbones for modeling deformation fields across grid-based, set-based, and graph-structured data. While effective in capturing spatial patterns

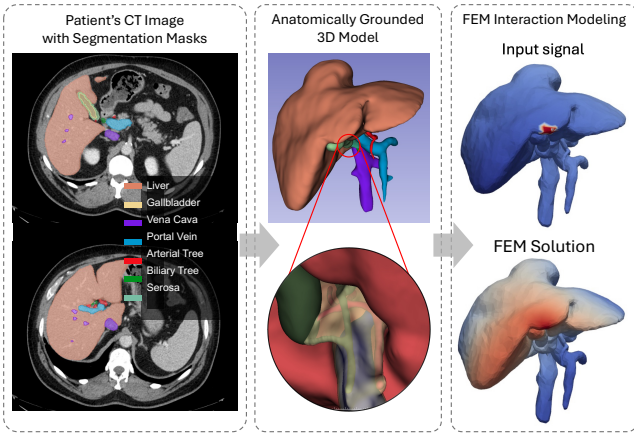


Figure 1: Overview of the large-scale FEM data generation pipeline. A patient’s abdominal CT scan is first anatomically segmented (left), followed by the generation of a 3D linear tetrahedral mesh (middle). Diverse interactions are then applied to the mesh, and FEM simulations are performed to produce training data for ML models (right). The right panel shows the norms of the input (interactions) and the resulting output displacements (predicted global effects).

and long-range dependencies, these models are generally agnostic to the underlying physics of soft tissue behavior and typically lack inductive biases grounded in biomechanics, such as force equilibrium, boundary conditions, or constitutive material priors.

Our work builds upon these foundations by integrating physically motivated priors, specifically Kelvinlet-based elastic solutions, into neural architectures through both residual learning and regularization. By embedding closed-form Green’s function responses from linear elasticity into the training process, our method encourages physically consistent deformation predictions even in data-sparse regimes. This hybrid approach improves generalization, improves physical plausibility, and remains compatible with a range of backbone architectures, while preserving the low latency performance required for real-time surgical simulation and interactive applications.

Large-Scale FEM Data

Neural network–based solutions require training on large-scale finite element (FE) simulation datasets. In this work, we focus on laparoscopic cholecystectomy and construct an FE model from an abdominal CT image and its segmentation masks, converting them into a 3D linear tetrahedral mesh with 10,400 nodes and 44,045 tetrahedra using `gmsht` (Geuzaine and Remacle 2009). As part of our contributions, we release this large-scale dataset; the details of its creation are provided below.

The bare area of the liver, where it contacts the inferior vena cava and diaphragm, and the distal ends of the ducts and veins are fixed (non-deformable). Tool interaction regions are selected to simulate the initial surgical step of lifting the liver lobes to expose the gallbladder and cystic ducts. These regions are identified by clustering nodes along the liver ridge, based on surgical video analysis and input from

our surgical collaborator. FE simulations are performed using the open-source `getFEM` package (Renard and Poullos 2020), applying tool-imposed displacements at those predefined clustered regions. This pipeline is illustrated in Figure 1.

In our FE model, nodes with restricted motion or tool-imposed displacements are flagged. These boundary conditions are encoded as binary features in the neural network (1 for constrained or controlled, 0 otherwise), allowing the model to distinguish between free and constrained regions.

Assuming a quasi-static environment, our solutions satisfy equilibrium by ensuring that the total potential energy is not only stationary but is also a minimum for permissible displacements, and is given by: $\int_{\Omega} \frac{\mathbf{W}(\mathbf{E})}{\partial \mathbf{E}} : \delta \mathbf{E} dV + \int_{\Omega} \rho \mathbf{g} \cdot \delta \mathbf{u} dV = 0$, where $\mathbf{W}(\mathbf{E})$ is the strain energy density, Ω is the volume of the model, $\delta \mathbf{u}$ is the variation of the displacement field, $\delta \mathbf{E}$ is the variation of the Green-Lagrange strain tensor, ρ is the density of the material (1000 kg/m^3), and \mathbf{g} is the acceleration due to gravity that points downward along the y -axis (9.81 m/s^2). For large displacements, the Green-Lagrange strain tensor \mathbf{E} depends on the unknown displacement field \mathbf{u} and is defined as:

$$\mathbf{E} = \frac{1}{2} (\nabla \mathbf{u} + \nabla \mathbf{u}^T + \nabla \mathbf{u}^T \nabla \mathbf{u}). \quad (1)$$

This equation relates to the Cauchy-Green deformation tensor \mathbf{C} by: $\mathbf{C} = 2\mathbf{E} - \mathbf{I}$, where \mathbf{I} is the identity tensor.

For our nonlinear simulations, we use a two-parameter Mooney-Rivlin model (Rivlin 1948), where the strain energy density can be written as functions of the Cauchy-Green deformation tensor invariants ($I_i(\mathbf{C})$, with $i = 1, 2, 3$) (Hackett 2018):

$$\mathbf{W} = C_{01} \left(I_1(\mathbf{C}) \cdot I_3(\mathbf{C})^{-1/3} - 3 \right) + C_{10} \left(I_1(\mathbf{C}) \cdot I_3(\mathbf{C})^{-2/3} - 3 \right), \quad (2)$$

where $C_{10} = 1.62 \text{ kPa}$ and $C_{01} = 1.97 \text{ kPa}$ are constitutive constants derived from ex-vivo porcine liver experiments (Benson, Pereira, and Miga 2025).

For our linear simulations, we simulate an elastic response in the absence of gravitational load ($\mathbf{g} = 0$). These solutions serve as baseline data for tissue deformation under small elastic input displacements, where the term $\nabla \mathbf{u}^T \nabla \mathbf{u}$ in Eq. (1) is neglected. Assuming isotropic material properties, the strain energy density \mathbf{W} in terms of shear modulus μ and Poisson’s ratio ν expressed as:

$$\mathbf{W} = \left(\frac{\mu\nu}{1 - 2\nu} \right) \text{Tr}(\mathbf{E})^2 + \mu \text{Tr}(\mathbf{E}^2). \quad (3)$$

In these linear simulations, we assign a Poisson’s ratio of 0.45 for all organs (Chen et al. 1996). The shear modulus values are also set as follows: 0.03 kPa for the fascia (Akhmanova et al. 2015), 0.34 kPa for the gallbladder (Nisansala et al. 2016), 0.69 kPa for the gallbladder wall and liver (Kim et al. 2013), and 1.07 kPa for the remaining organs (George et al. 2018).

Training Distribution, p . Given a mesh partitioned into regions Ω_i informed by cholecystectomy surgery, nodes \mathbf{x}_s are sampled from a mixture of Generalized von Mises-Fisher (vMF) distributions (i.e., radial densities on the surface mesh) centered in these regions:

$$p(\mathbf{x}_s) = \sum_i g(\Omega_i) f_{\text{vMF}}(\mathbf{x}_s \mid \boldsymbol{\mu}_i, \kappa_i), \quad (4)$$

where $g(\Omega_i)$ is the prior-informed probability of selecting Ω_i , $\boldsymbol{\mu}_i$ indicates the center of Ω_i , and κ_i is the spread of the vMF, which depends on the size of Ω_i . The displacements $\mathbf{u}_s \in \mathbb{R}^3$ are independently sampled from a uniform distribution over a cuboid domain $\mathcal{C} \subset \mathbb{R}^3$, defined by:

$$x \in (x_1, x_2), \quad y \in (y_1, y_2), \quad z \in (z_1, z_2),$$

with total volume $V = (x_2 - x_1)(y_2 - y_1)(z_2 - z_1)$. This ensures that node sampling is spatially structured according to surgical priors, while displacements remain isotropic within a bounded range.

For single grasping, ten regions Ω_i around the liver edge were selected, each with 250 samples. For multiple graspers, six pairwise combinations of five regions Ω_i were considered, with 30 samples per region. The displacement bounds are $x \in (-90, 50)$ mm, $y \in (-90, 50)$ mm, and $z \in (-60, 90)$ mm.

Distribution for Multiple Graspers. For multiple interactions (e.g., two graspers), we define the joint distribution as $p(\mathbf{x}_1, \mathbf{x}_2) = p(\mathbf{x}_1)p(\mathbf{x}_2 \mid \mathbf{x}_1)$, where $p(\mathbf{x}_1)$ follows Eq. (4). To ensure that \mathbf{x}_2 is not sampled from the same region Ω_k as \mathbf{x}_1 , we modify the region probabilities $g(\Omega_i)$ by setting $g(\Omega_k) = 0$, followed by renormalizing the remaining weights $g(\Omega_i)$ for $i \neq k$ based on surgical priors. The conditional distribution $p(\mathbf{x}_2 \mid \mathbf{x}_1)$ is then defined using these updated region weights.

Regularization Distribution, q . For Kelvinlet regularization, interaction points \mathbf{x}_s are sampled uniformly from the surface nodes. The corresponding displacements $\mathbf{u}_s \in \mathbb{R}^3$ are then sampled uniformly from a cone centered around the surface normal at \mathbf{x}_s , with an opening angle of α_{\max} .

Having detailed the data generation process, we now turn to our proposed neural augmented Kelvinlet solution.

Neural-Augmented Kelvinlet

In this section, we present the Kelvinlets prior and our residual learning and regularization frameworks.

Navier-Cauchy Eqs and Green's Function Solution

The displacement field $\mathbf{u} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ of an isotropic, homogeneous elastic medium under an applied body force density $\mathbf{f} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ satisfies the Navier-Cauchy equations:

$$\mu \nabla^2 \mathbf{u} + \frac{\mu}{1 - 2\nu} \nabla(\nabla \cdot \mathbf{u}) + \mathbf{f} = 0, \quad (5)$$

where $\mu > 0$ is the shear modulus, and $\nu \in (0, 0.5)$ is the Poisson's ratio, ensuring a physically consistent material response. Here, ∇ denotes the gradient operator, $\nabla \cdot \mathbf{u}$ is the

divergence of the displacement field, and $\nabla^2 \mathbf{u} = \nabla \cdot (\nabla \mathbf{u})$ is the vector Laplacian applied to \mathbf{u} .

For a concentrated force $\mathbf{f} \delta(\mathbf{x} - \mathbf{x}_s)$ applied at a source point $\mathbf{x}_s \in \mathbb{R}^3$, the Green's function solution, known as the *Kelvinlet*, provides an analytical displacement field $\mathbf{u} : \mathbb{R}^3 \setminus \{\mathbf{x}_s\} \rightarrow \mathbb{R}^3$ following:

$$\mathbf{u}(\mathbf{x}) = \frac{1}{4\pi\mu} \left(\frac{(1 - \nu)I}{r} + \nu \frac{(\mathbf{x} - \mathbf{x}_s) \otimes (\mathbf{x} - \mathbf{x}_s)}{r^3} \right) \mathbf{f} \quad (6)$$

where $r = \|\mathbf{x} - \mathbf{x}_s\|_2$ denotes the Euclidean distance from the source, and \otimes denotes the outer product operator.

The classical Kelvinlet in Eq. (6) exhibits singularities at \mathbf{x}_s . To prevent numerical instabilities, we substitute r with $r_\epsilon = \sqrt{\|\mathbf{x} - \mathbf{x}_s\|_2^2 + \epsilon^2}$ for $\epsilon > 0$, where the spatial extent of force distribution is controlled by ϵ , ensuring the solution is well-conditioned. Finally, a re-scaled solution with respect to the local displacement boundary condition $\mathbf{u}_s = \mathbf{u}(\mathbf{x}_s)$ removes the dependence on both the force \mathbf{f} and the modulus μ , yielding the regularized Kelvinlet formulation used in this paper,

$$\mathbf{u}_\epsilon(\mathbf{x}) = \frac{\epsilon}{(5 - 6\nu)} \left(\frac{(3 - 4\nu)I}{r_\epsilon} + \frac{(\mathbf{x} - \mathbf{x}_s) \otimes (\mathbf{x} - \mathbf{x}_s)}{r_\epsilon^3} \right) \mathbf{u}_s. \quad (7)$$

For multiple interaction points, we note that at each interaction point we only observe the superposition of contributions from all interaction points, without knowing the exact contribution of each individual point to this superposition. Consequently, unlike the single-grasper setting, \mathbf{u}_s in Eq. (7) is not directly available for each interaction point and must be estimated. To formalize this, let \mathbf{x}_i denote the location of the i th interaction, \mathbf{u}_i the corresponding total displacement at \mathbf{x}_i , and \mathbf{k}_i the unknown contribution of the i th interaction to this displacement. Assuming there are K interaction points, the total displacement can be expressed as:

$$\mathbf{u}_\epsilon(\mathbf{x}) = \sum_{i=1}^K \underbrace{\frac{\epsilon}{(5 - 6\nu)} \left(\frac{(3 - 4\nu)I}{r_\epsilon} + \frac{(\mathbf{x} - \mathbf{x}_i) \otimes (\mathbf{x} - \mathbf{x}_i)}{r_\epsilon^3} \right)}_{\Gamma_i(\mathbf{x})} \mathbf{k}_i, \quad (8)$$

where $\Gamma_i(\mathbf{x}) \in \mathbb{R}^{3 \times 3}$, and $r_{i,\epsilon}(\mathbf{x}) := \sqrt{\|\mathbf{x} - \mathbf{x}_i\|_2^2 + \epsilon^2}$. Note that this formulation enables the total displacement function to incorporate cross communication between interaction points.

To estimate the \mathbf{k}_i values, we solve the following optimization problem, which enforces the boundary conditions $\mathbf{u}_\epsilon(\mathbf{x}_i) = \mathbf{u}_i$ for all i :

$$\underset{\mathbf{k}}{\operatorname{argmin}} \sum_{i=1}^K \left(\left\| \sum_{j=1}^K \Gamma_j(\mathbf{x}_i) \mathbf{k}_j - \mathbf{u}_i \right\|^2 + \lambda \|\mathbf{k}_i\|^2 \right). \quad (9)$$

This formulation ensures that the individual Kelvinlet contributions collectively satisfy the prescribed displacements, while the regularization term $\lambda = 0.001$ promotes a minimum-energy solution. The optimization in Eq. (9) is essentially a regularized least squares problem, which admits a fast analytic solution:

$$\mathbf{k}^* = (\mathbf{\Gamma}^\top \mathbf{\Gamma} + \lambda \mathbf{I})^{-1} \mathbf{\Gamma}^\top \mathbf{u}, \quad (10)$$

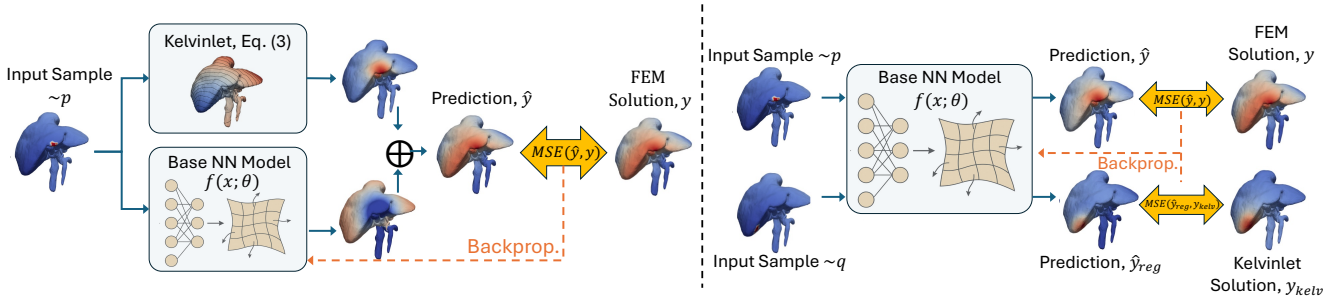


Figure 2: Overview of the two proposed training frameworks. **Method 1 (Left):** Residual learning predicts only the deviation from the Kelvinlet solution, reducing learning complexity by leveraging the physics-based analytical solution as an initialization. **Method 2 (Right):** Kelvinlet-based regularization combines FEM-supervised learning with a secondary loss enforcing consistency with Kelvinlet priors. The heatmap depicts the displacement norms.

where $\mathbf{I} \in \mathbb{R}^{3K \times 3K}$ is the identity matrix, and $\Gamma \in \mathbb{R}^{3K \times 3K}$, $\mathbf{k} \in \mathbb{R}^{3K}$, and $\mathbf{u} \in \mathbb{R}^{3K}$ are defined as:

$$\Gamma = \begin{bmatrix} \Gamma_1(\mathbf{x}_1) & \cdots & \Gamma_K(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ \Gamma_1(\mathbf{x}_K) & \cdots & \Gamma_K(\mathbf{x}_K) \end{bmatrix}, \mathbf{k} = \begin{bmatrix} \mathbf{k}_1 \\ \vdots \\ \mathbf{k}_K \end{bmatrix}, \mathbf{u} = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_K \end{bmatrix}.$$

Note that while the closed form solution requires the inversion of a $3K \times 3K$ matrix, the number of interaction points, K is often small, leading to a very fast solution for Kelvinlet. The total Kelvinlet deformation model, including the analytic solution to the optimization, can be performed in less than 1 ms for two simultaneously engaged graspers (i.e., $K \approx 20$) using native GPU acceleration in PyTorch on an NVIDIA A6000 GPU and on a mesh with $\sim 10\text{k}$ nodes.

Having established a fast physics-based solution, we now focus on integrating Kelvinlets into our models.

Method 1: Kelvinlet-Based Residual Learning

For an initial displacement \mathbf{u}_s located at \mathbf{x}_s , let $\mathbf{u}_{\text{true}}(\cdot; \mathbf{x}_s, \mathbf{u}_s)$ denote the ground-truth displacement field. Instead of directly regressing \mathbf{u}_{true} , we introduce a residual learning framework where a neural network estimates the deviation from the Kelvinlet prior, which serves as an initial approximation. Specifically, we define the residual displacement as $\mathbf{r}(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s) = \mathbf{u}_{\text{true}}(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s) - \mathbf{u}_\varepsilon(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s)$, where \mathbf{u}_ε represents the Kelvinlet solution. The neural network, parameterized by $\theta \in \mathbb{R}^d$, learns to approximate this residual via $\hat{\mathbf{r}}(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s, \theta)$. The corrected displacement field is then given by $\tilde{\mathbf{u}}(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s, \theta) = \mathbf{u}_\varepsilon(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s) + \hat{\mathbf{r}}(\mathbf{x}; \mathbf{x}_s, \mathbf{u}_s, \theta)$. Let $\|\mathbf{u}\|_{L_2(\Omega)}^2 = \int_{\Omega} \|\mathbf{u}(\mathbf{x})\|^2 d\mathbf{x}$, then to train the network, we minimize the loss:

$$\mathcal{L}_{\text{residual}}(\theta) = \mathbb{E}_{(\mathbf{x}_s, \mathbf{u}_s) \sim p} \left[\left\| \mathbf{u}_{\text{true}}(\cdot; \mathbf{x}_s, \mathbf{u}_s) - \tilde{\mathbf{u}}(\cdot; \mathbf{x}_s, \mathbf{u}_s, \theta) \right\|_{L_2(\Omega)}^2 \right], \quad (11)$$

where p represents the distribution of all possible point-wise interactions with the tissue, and Ω denotes the tissue domain. By focusing on learning only the correction beyond the Kelvinlet solution, this approach leverages the physics priors directly into its formulation.

Method 2: Kelvinlets as Regularization

Alternatively, a network can be trained to directly regress \mathbf{u}_{true} , producing an estimate $\hat{\mathbf{u}}$. To enforce physically consistent predicted deformations by the network, we leverage Kelvinlets as a regularization prior and write:

$$\mathcal{L}(\theta) = \mathbb{E}_{(\mathbf{x}_s, \mathbf{u}_s) \sim p} \left[\left\| \mathbf{u}_{\text{true}} - \hat{\mathbf{u}} \right\|_{L_2(\Omega)}^2 \right] + \lambda_{\text{reg}} \mathbb{E}_{(\mathbf{x}'_s, \mathbf{u}'_s) \sim q} \left[\left\| \mathbf{u}_\varepsilon - \hat{\mathbf{u}} \right\|_{L_2(\Omega)}^2 \right] \quad (12)$$

where $\lambda_{\text{reg}} > 0$ is the regularization coefficient. For clarity and brevity, we omit the dependencies $(\cdot; \mathbf{x}_s, \mathbf{u}_s)$ and $(\cdot; \mathbf{x}_s, \mathbf{u}_s, \theta)$ in the notation. Note that we explicitly separate $(\mathbf{x}_s, \mathbf{u}_s)$ from $(\mathbf{x}'_s, \mathbf{u}'_s)$ because, in practice, the two distributions p and q could differ, and also while Monte Carlo integration is used for both terms, the first term requires FEM simulations to obtain \mathbf{u}_{true} (which is computationally expensive), whereas the second term only involves Kelvinlet evaluations (much faster). We will describe the distributions p and q used in our experiments in the following section.

Figure 2 illustrates the proposed residual- and regularization-based learning methods, along with the displacement norm visualizations for all methods on a representative sample.

Experiments

Data. Our dataset consists of large-scale FEM simulations of soft tissue deformation. Each node in the mesh has a feature vector: $\mathbf{f} = [\mathbf{x}; \mathbf{u}_s; a] \in \mathbb{R}^7$. Where $\mathbf{x} \in \mathbb{R}^3$ is the undeformed position, $\mathbf{u}_s \in \mathbb{R}^3$ is the initial displacement, and $a \in \{0, 1\}$ indicates whether the node is actively influenced by external forces. This distinction helps differentiate constrained nodes from passively deforming ones. We create a dataset of 20,800 FEM solutions with 10,400 linear elastic simulations and 10,400 nonlinear material responses. Of these, 5,000 correspond to single manipulations and 5,400 to multiple grasper manipulations.

Neural Architectures. We evaluate two groups of models: deformation-specific FEM surrogates and generic set/graph architectures. For the former, PhysGNN (Salehi and Gianacopoulos 2022) is a mesh-based GNN that we augment




Figure 3: Qualitative comparison of deformation predictions for the best (PointNet) and worst (GraphGPS) models in linear regimes.

Model	Base	Ours	
		+Res.	+Reg.
PhysGNN(Salehi and Giannacopoulos 2022)	1.71 (0.11)	1.84 (0.13)	1.71 (0.11)
SeeSaw(Doceca et al. 2024)	1.56 (0.09)	1.69 (0.08)	1.56 (0.09)
Physics-Encoded GNN(Saleh et al. 2024)	1.65 (0.08)	1.79 (0.08)	1.65 (0.08)
Transformer-KAN(Jiang et al. 2025)	2.22 (0.17)	2.32 (0.16)	2.22 (0.17)
Diffeomorphic Deformation(Gong et al. 2024)	2.16 (0.12)	2.35 (0.11)	2.16 (0.12)
PointNet(Qi et al. 2016)	1.16 (0.10)	1.23 (0.13)	1.16 (0.10)
U-Net(Mendizabal, Márquez-Neila, and Cotin 2020)	1.55 (0.08)	1.64 (0.07)	1.55 (0.08)
GraphGPS(Rampášek et al. 2022)	2.60 (0.15)	2.78 (0.15)	2.60 (0.15)
SAB(Lee et al. 2019)	2.30 (0.02)	2.44 (0.21)	2.30 (0.20)
Kelvinlet	0.15 (0.052)		

Table 1: Inference time (ms) on an NVIDIA RTX A6000 GPU for a single grasper in linear simulations, with and without Kelvinlet priors. Values are mean (std).

with a global attention layer to counter oversmoothing (Zhao and Akoglu 2020); PEGNN (Saleh et al. 2024) is used only through its soft-tissue encoder; SeeSaw (Doceca et al. 2024) is reimplemented for our point-cloud setting; from the diffeomorphic framework of (Gong et al. 2024) we keep only the deformation network; and a Transformer-based KAN (Jiang et al. 2025) is restricted to single-step prediction. To study architecture-agnostic effects of Kelvinlet priors, we also use U-Net (Mendizabal, Márquez-Neila, and Cotin 2020), SAB (Lee et al. 2019), GraphGPS (Rampášek et al. 2022), and PointNet (Qi et al. 2016): U-Net is converted into a DeepSet-style model via shared pointwise MLPs, SAB uses two self-attention blocks with an MLP head, and GraphGPS and PointNet are shallow variants with reduced depth to avoid overfitting on high-resolution meshes.

Hyperparameters. We set shear modulus in Kelvinlet to $\mu = 0.72$ kPa, Poisson’s ratio $\nu = 0.45$, and regularization parameter $\epsilon = 0.05$, to reflect soft tissue properties and ensure numerical stability. For linear and nonlinear simulations we use $\lambda_{\text{reg}} = 1$ and $\lambda_{\text{reg}} = 0.1$, respectively, to emphasize linearity of Kelvinlet.

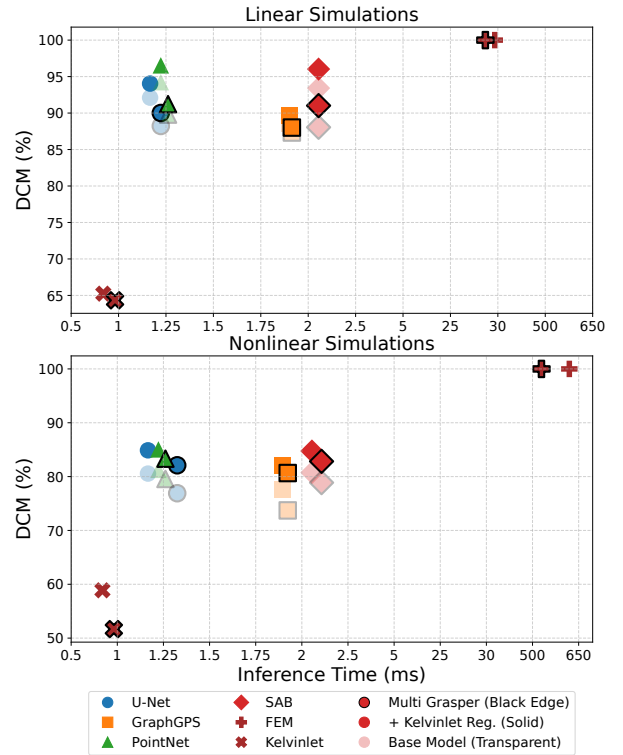


Figure 4: Log-scaled inference time vs. accuracy (DCM%) for various models in linear and nonlinear simulations. Transparent markers denote base models, while solid markers indicate Kelvinlet-regularized versions. Black-edged markers represent multi-grasper cases, while edge-free markers denote single-grasper cases. FEM and Kelvinlet baselines are included for reference.

Results and Analysis To quantitatively assess model performance, we utilize the *Deformation Capture Mean* (DCM) metric, which measures how well a predicted deformation field approximates the ground-truth displacement field. The DCM is computed as:

$$\text{DCM} = 100 \times \left(1 - \mathbb{E}_{(\mathbf{x}_s, \mathbf{u}_s) \sim p} \left[\frac{\|\mathbf{u}_{\text{true}} - \hat{\mathbf{u}}\|_{L_2(\Omega)}}{\|\mathbf{u}_{\text{true}}\|_{L_2(\Omega)}} \right] \right) \quad (13)$$

Model	Indiv. (DCM)			Comb. (DCM)			Multi (DCM)		
	Base	+Res.	+Reg.	Base	+Res.	+Reg.	Base	+Res.	+Reg.
PhysGNN(Salehi and Giannacopoulos 2022)	90.12(0.13)	91.77(0.11)	92.65(0.10)	86.38(0.13)	88.10(0.11)	88.66(0.10)	89.55(0.13)	90.96(0.11)	91.54(0.10)
SeeSaw(Doceca et al. 2024)	89.24(0.15)	91.02(0.10)	91.83(0.08)	85.50(0.15)	87.35(0.10)	87.84(0.08)	88.67(0.15)	90.21(0.10)	90.72(0.08)
Physics-Encoded GNN(Saleh et al. 2024)	88.55(0.14)	90.44(0.09)	91.58(0.12)	84.81(0.14)	86.77(0.09)	87.59(0.12)	87.98(0.14)	89.63(0.09)	90.47(0.12)
Transformer-KAN(Jiang et al. 2025)	90.89(0.12)	92.01(0.08)	93.22(0.07)	87.15(0.12)	88.34(0.08)	89.23(0.07)	90.32(0.12)	91.20(0.08)	92.11(0.07)
Diffeomorphic Deformation(Gong et al. 2024)	91.14(0.10)	92.88(0.11)	94.33(0.09)	87.40(0.10)	89.21(0.11)	90.34(0.09)	90.57(0.10)	92.07(0.11)	93.22(0.09)
PointNet(Qi et al. 2016)	94.23(0.11)	94.66(0.09)	96.54(0.09)	89.79(0.07)	90.11(0.07)	91.22(0.07)	94.00(0.08)	94.12(0.07)	95.01(0.06)
U-Net(Mendizabal, Márquez-Neila, and Cotin 2020)	92.11(0.12)	93.56(0.10)	94.04(0.11)	88.23(0.08)	89.67(0.09)	90.01(0.08)	91.20(0.09)	92.33(0.10)	92.88(0.08)
GraphGPS(Rampášek et al. 2022)	88.60(0.21)	89.45(0.15)	89.61(0.23)	87.33(0.06)	87.92(0.02)	88.01(0.04)	88.35(0.13)	88.91(0.10)	89.45(0.09)
SAB(Lee et al. 2019)	93.42(0.11)	94.60(0.04)	96.03(0.04)	88.05(0.09)	89.91(0.11)	91.02(0.04)	92.54(0.06)	93.66(0.07)	94.42(0.06)
Kelvinlet	65.24(0)			64.33(0)			64.82(0)		

Table 2: Performance of various models on **linear** simulations with individual graspers, combined graspers, and multitask learning. DCM (%) is reported with standard deviations. “Base” denotes the original model, “+Res.” augments it with a Kelvinlet-based residual, and “+Reg.” applies a Kelvinlet-inspired regularizer.

Model	Indiv. (DCM)			Comb. (DCM)			Multi (DCM)		
	Base	+Res.	+Reg.	Base	+Res.	+Reg.	Base	+Res.	+Reg.
PhysGNN(Salehi and Giannacopoulos 2022)	78.05(0.13)	82.19(0.11)	82.79(0.10)	75.30(0.13)	79.18(0.11)	80.83(0.10)	77.73(0.13)	80.58(0.11)	81.34(0.10)
SeeSaw(Doceca et al. 2024)	77.17(0.15)	81.44(0.10)	81.97(0.08)	74.42(0.15)	78.43(0.10)	80.01(0.08)	76.85(0.15)	79.83(0.10)	80.52(0.08)
Physics-Encoded GNN(Saleh et al. 2024)	76.48(0.14)	80.86(0.09)	81.72(0.12)	73.73(0.14)	77.85(0.09)	79.76(0.12)	76.16(0.14)	79.25(0.09)	80.27(0.12)
Transformer-KAN(Jiang et al. 2025)	78.82(0.12)	82.43(0.08)	83.36(0.07)	76.07(0.12)	79.42(0.08)	81.40(0.07)	78.50(0.12)	80.82(0.08)	81.91(0.07)
Diffeomorphic Deformation(Gong et al. 2024)	79.07(0.10)	83.30(0.11)	84.47(0.09)	76.32(0.10)	80.29(0.11)	82.51(0.09)	78.75(0.10)	81.69(0.11)	83.02(0.09)
PointNet(Qi et al. 2016)	81.25(0.22)	84.65(0.19)	85.05(0.15)	79.56(0.18)	82.66(0.16)	83.33(0.11)	83.25(0.17)	84.75(0.13)	85.21(0.09)
U-Net(Mendizabal, Márquez-Neila, and Cotin 2020)	80.56(0.10)	84.22(0.09)	84.88(0.11)	76.91(0.05)	79.83(0.08)	82.10(0.21)	78.10(0.12)	80.45(0.09)	81.33(0.14)
GraphGPS(Rampášek et al. 2022)	77.52(0.17)	80.63(0.13)	82.11(0.09)	73.72(0.15)	77.92(0.10)	80.66(0.16)	76.24(0.18)	79.31(0.15)	80.42(0.10)
SAB(Lee et al. 2019)	80.75(0.19)	84.44(0.07)	84.75(0.10)	78.88(0.11)	81.50(0.15)	82.83(0.16)	81.22(0.10)	83.01(0.08)	84.00(0.07)
Kelvinlet	58.88(0)			51.67(0)			55.56(0)		

Table 3: Performance of various models on **nonlinear** simulations with individual graspers, combined graspers, and multitask learning. DCM (%) is reported with standard deviations. “Base” denotes the original model, “+Res.” augments it with a Kelvinlet-based residual, and “+Reg.” applies a Kelvinlet-inspired regularizer.

Table 2 shows that Kelvinlet priors consistently improve deformation accuracy in the linear regime. Across all architectures, residual learning yields roughly 1 to 2 percentage points of DCM improvement, while Kelvinlet regularization adds a further 1 to 2 points, with the largest gain of **+3.19%** achieved by the diffeomorphic deformation network. PointNet achieves the highest absolute performance, reaching **96.54%** DCM in the single grasper setting. In contrast, the analytical Kelvinlet alone remains substantially less accurate at **65.24%** DCM, highlighting the importance of learned residual corrections.

As shown in Table 1, the proposed methods incur negligible runtime overhead. Regularized models match the base inference time, while residual variants increase latency by only about 0.1ms. Although standalone Kelvinlet inference is extremely fast at **0.15ms**, its accuracy is insufficient, supporting its role as a prior rather than a predictor. Overall, accuracy gains are achieved without sacrificing efficiency.

Figure 4 illustrates the accuracy versus runtime trade off across linear and nonlinear simulations. Kelvinlet regularization consistently improves accuracy relative to the base

models, as indicated by the upward shift of solid markers, while maintaining comparable inference times. Models trained on multi grasper data follow the same trend, indicating robust generalization across grasping conditions. All runtime measurements were obtained on a single NVIDIA A6000 GPU.

Conclusion

We present Neural Augmented Kelvinlets, a physics informed framework for real time soft tissue deformation modeling that improves accuracy, physical plausibility, and generalization while preserving real time performance. Kelvinlets are incorporated both as a residual correction mechanism and as a regularization prior that enforces physically consistent deformations. Extensive FEM simulations demonstrate clear performance gains, particularly in multi grasper scenarios. Our GPU accelerated PyTorch implementation enables real time inference and realistic visualization, making the method suitable for surgical robotics and medical training applications.

Acknowledgments

This work was fully supported by the Wellcome-Leap SAVE program.

References

- Akhmanova, M.; Osidak, E.; Domogatsky, S.; Rodin, S.; and Domogatskaya, A. 2015. Physical, Spatial, and Molecular Aspects of Extracellular Matrix of In Vivo Niches and Artificial Scaffolds Relevant to Stem Cells Research. *Stem Cells International*, 2015: 167025.
- Allard, J.; Courtecuisse, H.; and Faure, F. 2012. Chapter 21 - Implicit FEM Solver on GPU for Interactive Deformation Simulation. In mei W. Hwu, W., ed., *GPU Computing Gems Jade Edition*, Applications of GPU Computing Series, 281–294. Boston: Morgan Kaufmann. ISBN 978-0-12-385963-1.
- Belytschko, T.; Krongauz, Y.; Organ, D.; Fleming, M.; and Krysl, P. 1996. Meshless methods: An overview and recent developments. *Computer Methods in Applied Mechanics and Engineering*, 139(1): 3–47.
- Benson, A.; Pereira, K.; and Miga, M. I. 2025. Gravity: weighty attributes in surgical simulation. In *Medical Imaging 2025: Image-Guided Procedures, Robotic Interventions, and Modeling*. International Society for Optics and Photonics, SPIE.
- Bronstein, M. M.; Bruna, J.; LeCun, Y.; Szlam, A.; and Vandergheynst, P. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4): 18–42.
- Camara, M.; Mayer, E.; Darzi, A.; and Pratt, P. 2016. Soft tissue deformation for surgical simulation: a position-based dynamics approach. *International Journal of Computer Assisted Radiology and Surgery*, 11(6): 919–928.
- Chen, E. J.; Novakofski, J.; Jenkins, W. K.; and O'Brien, W. D. 1996. Young's Modulus Measurements of Soft Tissues with Application to Elasticity Imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 43(1): 191–194.
- De Goes, F.; and James, D. L. 2017. Regularized kelvinlets: sculpting brushes based on fundamental solutions of elasticity. *ACM Trans. Graph.*, 36(4).
- Docea, R.; Xu, J.; Ling, W.; Jenke, A. C.; Kolbinger, F. R.; Distler, M.; Riediger, C.; Weitz, J.; Speidel, S.; and Pfeiffer, M. 2024. SeeSaw: Learning Soft Tissue Deformation From Laparoscopy Videos With GNNs. *IEEE Transactions on Biomedical Engineering*, 71(12): 3432–3445.
- George, D.; Baniassadi, M.; Hoarau, Y.; Kugler, M.; and Rémond, Y. 2018. Influence of the liver vascular distribution on its overall mechanical behavior: A first approach to multiscale fluid-structure homogenization. *Journal of Cellular Immunotherapy*, 4(1): 35–37.
- Geuzaine, C.; and Remacle, J.-F. 2009. Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11): 1309–1331.
- Gong, S.; Long, Y.; Chen, K.; Liu, J.; Xiao, Y.; Cheng, A.; Wang, Z.; and Dou, Q. 2024. Self-Supervised Cyclic Deformation Mapping for Soft Tissue Deformation Recovery in Robotic Surgery Scenes. *IEEE Transactions on Medical Imaging*, 43(12): 4356–4367.
- Hackett, R. M. 2018. *Strain Measures*, 5–8. Cham: Springer International Publishing. ISBN 978-3-319-73201-5.
- Hu, T.; and Desai, J. P. 2004. Characterization of Soft-Tissue Material Properties: Large Deformation Analysis. In Cotin, S.; and Metaxas, D., eds., *Medical Simulation*, 28–37. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Hu, Z.; Liao, S.; Kui, X.; Wu, R.; Yuan, F.; and Chen, Q. 2025. Real-time integrated modeling of soft tissue deformation and stress based on deep learning. *Physics in Medicine and Biology*, 70(12).
- Jiang, J.; Fu, T.; Liu, J.; Fan, J.; Song, H.; Ai, D.; Xiao, D.; Wang, Y.; and Yang, J. 2025. Incremental energy-based recurrent transformer-KAN for time series deformation simulation of soft tissue. *Expert Systems with Applications*, 271: 126619.
- Karniadakis, G. E.; Kevrekidis, I. G.; Lu, L.; Perdikaris, P.; Wang, S.; and Yang, L. 2021. Physics-informed machine learning. *Nature Reviews Physics*, 3(6): 422–440.
- Kim, Y.; Chang, D.; Kim, J.; and Park, S. 2013. Gallbladder Removal Simulation for Laparoscopic Surgery Training: A Hybrid Modeling Method. *Journal of Computer Science and Technology*, 28(3): 499–507.
- Lee, J.; Lee, Y.; Kim, J.; Kosiosek, A.; Choi, S.; and Teh, Y. W. 2019. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning*, 3744–3753.
- Liu, H.; Han, Y.; Emerson, D.; Rabin, Y.; and Kara, L. B. 2025. A data-driven approach for real-time soft tissue deformation prediction using nonlinear presurgical simulations. *PLOS ONE*, 20(4): e0319196.
- Liu, X.; Wang, R.; Li, Y.; and Song, D. 2015. Deformation of Soft Tissue and Force Feedback Using the Smoothed Particle Hydrodynamics. *Computational and Mathematical Methods in Medicine*, 2015(1): 598415.
- Mendizabal, A.; Márquez-Neila, P.; and Cotin, S. 2020. Simulation of hyperelastic materials in real-time using deep learning. *Medical image analysis*, 59: 101569.
- Nedel, L. P.; and Thalmann, D. 1998. Real time muscle deformations using mass-spring systems. In *Proceedings. Computer Graphics International (Cat. No.98EX149)*, 156–165.
- Nguyen-Le, H.-D.; Ballit, A.; and Dao, T.-T. 2025. Physics-informed neural ordinary differential equations and mass-spring system modeling framework for predicting the real-time soft-tissue deformations. *Engineering Applications of Artificial Intelligence*, 156: 111217.
- Nisansala, A.; Dias, G. K. A.; Kodikara, N.; Weerasinghe, M.; Sandaruwan, D.; Keppitiyagama, C.; and Dammika, N. 2016. Stiffness Parameter Evaluation for Graphical and

Haptic Gallbladder Model. In *Proceedings of the 2nd International Conference on Communication and Information Processing (ICCIP '16)*, ICCIP '16, 249–254. New York, NY, USA: Association for Computing Machinery. ISBN 978-1-4503-4819-5.

Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; and Battaglia, P. 2021. Learning Mesh-Based Simulation with Graph Networks. In *International Conference on Learning Representations*.

Pfeiffer, M.; Riediger, C.; Weitz, J.; and Speidel, S. 2019. Learning soft tissue behavior of organs for surgical navigation with convolutional neural networks. *International journal of computer assisted radiology and surgery*, 14(7): 1147–1155.

Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2016. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *arXiv preprint arXiv:1612.00593*.

Rampásek, L.; Galkin, M.; Dwivedi, V. P.; Luu, A. T.; Wolf, G.; and Beaini, D. 2022. Recipe for a General, Powerful, Scalable Graph Transformer. *Advances in Neural Information Processing Systems*, 35.

Renard, Y.; and Poullos, K. 2020. GetFEM: Automated FE Modeling of Multiphysics Problems Based on a Generic Weak Form Language. *ACM Trans. Math. Softw.*, 47(1).

Rivlin, R. S. 1948. Large elastic deformations of isotropic materials IV. Further developments of the general theory. *Philosophical transactions of the royal society of London. Series A, Mathematical and physical sciences*, 241(835): 379–397.

Saleh, M.; Sommersperger, M.; Navab, N.; and Tombari, F. 2024. Physics-Encoded Graph Neural Networks for Deformation Prediction under Contact. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 17160–17166. IEEE.

Salehi, Y.; and Giannacopoulos, D. 2022. PhysGNN: A Physics-Driven Graph Neural Network Based Model for Predicting Soft Tissue Deformation in Image-Guided Neurosurgery. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.

Wu, J. Y.; Kazanzides, P.; and Unberath, M. 2020. Leveraging vision and kinematics data to improve realism of biomechanic soft tissue simulation for robotic surgery. *International journal of computer assisted radiology and surgery*, 15(5): 811–818.

Wu, J. Y.; Munawar, A.; Unberath, M.; and Kazanzides, P. 2021. Learning soft-tissue simulation from models and observation. In *2021 International Symposium on Medical Robotics (ISMR)*, 1–7. IEEE.

Zhang, J.; Zhong, Y.; and Gu, C. 2019. Neural network modelling of soft tissue deformation for surgical simulation. *Artificial Intelligence in Medicine*, 97: 61–70.

Zhao, L.; and Akoglu, L. 2020. PairNorm: Tackling Over-smoothing in GNNs. In *International Conference on Learning Representations (ICLR)*.