

Generalizing Soft Tissue Deformation and Force Prediction Across Material Stiffness and Geometry

Madina Kojanazarova, Sidaty El Hadramy, and Philippe C. Cattin

University of Basel, Department of Biomedical Engineering, Allschwil, Switzerland
M.Kojanazarova@unibas.ch

Abstract. Accurate soft tissue simulation is essential for surgical training, pre-operative planning, and haptic feedback systems. While learning-based surrogate models trained on data using the finite element method (FEM) offer a promising path to real-time inference, their reliability depends on well-calibrated constitutive models. Existing approaches neither provide systematic guidance on model selection across stiffness levels, nor generalize across different tissue stiffnesses or geometries. We perform a comprehensive calibration of hyperelastic constitutive models in the SOFA Framework using gravity-loaded silicone beams with different stiffnesses. Using calibrated simulations as training data, we use a softness conditioned equivariant graph neural network, enabling deformation and force prediction across multiple tissue types and unseen geometries. Our model achieves sub-millimeter mean deformation accuracy at 0.010s inference time, while showing that force prediction quality is directly tied to upstream calibration consistency.

Keywords: soft tissue simulation, finite element method, graph neural networks

1 Introduction

Minimally invasive surgery and robotic-assisted interventions increasingly rely on accurate models of soft tissue behavior to support surgical training, pre-operative planning, and intraoperative guidance [4]. A core challenge in all of these applications is predicting how tissue deforms under tool contact in real time: deformation fields are needed to update augmented reality overlays [10], and contact forces must be rendered haptically to give trainees realistic tactile feedback [17]. Physics-based simulation using the finite element method (FEM) can reproduce soft tissue deformation with high fidelity by solving the equations of continuum mechanics on a volumetric mesh [21], but the computational cost of solving large nonlinear systems at each time step typically precludes real-time use, with simulation times orders of magnitude above the millisecond budgets required for interactive applications [18].

A widely adopted strategy to overcome this bottleneck is to use FEM offline to generate large bodies of high-fidelity training data, and then train a fast surrogate model that approximates the FEM response at inference time [15,5]. Once

trained, such surrogates can deliver predictions in milliseconds, enabling integration into real-time surgical simulators without sacrificing the physical realism of the underlying simulation. Among the available surrogate architectures, graph neural networks (GNNs) are particularly well-suited to this task. They operate natively on unstructured meshes and point clouds, naturally handle irregular and patient-specific geometries, and can propagate local deformation cues globally through message passing [19,5].

However, existing surrogates face key limitations. First, generalization across unseen geometries and varying material stiffness simultaneously remains an open challenge. Existing approaches typically fix one or both of these properties at training time [5,7,13,12,15]. Second, any surrogate trained on FEM data is only as reliable as the underlying simulation. A critical and often overlooked prerequisite is the selection and calibration of the constitutive model [6,14] and the mathematical description of how stress relates to strain in the tissue [3]. Simulation frameworks such as the SOFA Framework [8], widely used in medical simulation research, provide several hyperelastic constitutive models, each encoding different assumptions about material microstructure. Yet there is no systematic guidance on which model best captures the behavior of silicone-like soft materials across different stiffness levels, and parameters overfit to a single configuration will directly corrupt the FEM training data and, by extension, the surrogate trained on it.

In this work, we address these limitations in a unified pipeline, illustrated in Fig. 1. First, we fabricate silicone beams at three stiffness levels and use gravity-loading experiments across multiple lengths and boundary conditions as a controlled calibration benchmark (Section 3.1). We then systematically optimize the parameters of all applicable hyperelastic models in SOFA [8] and evaluate which models reproduce the physical deformation most accurately across stiffness levels (Section 3.2). Using the best-performing models and calibrated parameters, we generate a synthetic dataset of poking interactions on soft volumes with varied embedded geometries spanning all three stiffness levels (Section 3.3). Finally, this dataset is used to train a softness-conditioned extension of the equivariant GNN of [13], enabling a single network to predict deformation and contact force across multiple material classes and generalize to unseen geometries, bringing us closer to a fast, generalizable surrogate for real-time soft tissue simulation.




Fig. 1: Overview of the proposed pipeline for experimental material calibration, synthetic data generation, and surrogate learning.

2 Background

2.1 Physics-Based Soft Tissue Simulations

Soft tissues exhibit hyperelastic behavior where, under large deformations, the stress-strain relationship is nonlinear. This behavior is commonly modeled in continuum mechanics as a strain energy density function Ψ , defined as a function of the deformation gradient $\mathbf{F}$, its determinant $J = \det \mathbf{F}$ (the local volume ratio), and the right Cauchy-Green tensor $\mathbf{C} = \mathbf{F}^T \mathbf{F}$, whose invariants are

$$I_1 = \text{tr}(\mathbf{C}), \quad I_2 = \frac{1}{2} [(\text{tr} \mathbf{C})^2 - \text{tr}(\mathbf{C}^2)]. \quad (1)$$

Hyperelastic models differ in how Ψ is formulated, reflecting different assumptions about the underlying material microstructure and trade-offs between accuracy and numerical stability.

Saint Venant-Kirchhoff The simplest of the models considered, extends the classical linear (Hookean) stress-strain relationship to large displacements by replacing the linear strain tensor with the nonlinear Green-Lagrange strain tensor $\mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I})$:

$$\Psi_{\text{StVK}} = \frac{\lambda}{2} \text{tr}(\mathbf{E})^2 + \mu \text{tr}(\mathbf{E}^2), \quad (2)$$

with λ and μ being the Lamé parameters. Because the stress remains a linear function of $\mathbf{E}$, the model is computationally efficient but physically unrealistic at large compressive strains, where it can produce non-physical, non-monotonic stress responses [2].

Neo-Hookean Derived from the statistical mechanics of polymer chains under the assumption of Gaussian chain-length statistics, the compressible Neo-Hookean model is given by:

$$\Psi_{\text{NH}} = \frac{\mu}{2}(I_1 - 3) - \mu \ln J + \frac{\lambda}{2}(\ln J)^2. \quad (3)$$

With only two material parameters, it is the simplest true hyperelastic model and captures the qualitative nonlinearity of rubber-like materials, but it does not reproduce the pronounced strain-stiffening observed in real elastomers at large stretches [16].

Stable Neo-Hookean A more recent reformulation of the Neo-Hookean model designed for robustness in simulation rather than fidelity to a specific microstructural derivation:

$$\Psi_{\text{SNH}} = \frac{\mu}{2}(I_1 - 3) - \mu(J - 1) + \frac{\lambda + \mu}{2}(J - 1)^2. \quad (4)$$

By modifying the volumetric term relative to the standard compressible Neo-Hookean model, this formulation guarantees a positive semi-definite Hessian even under element inversion or extreme compression, eliminating the numerical instabilities that can otherwise arise during large or fast deformations [20].

Mooney-Rivlin Proposed as a generalization of the Neo-Hookean model. The Mooney-Rivlin model introduces a second invariant term to better fit experimental stress-strain data over a wider strain range:

$$\Psi_{\text{MR}} = C_{10}(\bar{I}_1 - 3) + C_{01}\bar{I}_2 + \frac{K}{2}(\ln J)^2, \quad (5)$$

where $\bar{I}_1, \bar{I}_2$ are the volume-normalized invariants and K is the bulk modulus. The Neo-Hookean model is recovered as the special case when $C_{01} = 0$ [16].

Ogden Rather than expressing Ψ in terms of invariants, the Ogden model is formulated directly in terms of the principal stretches $\lambda_1, \lambda_2, \lambda_3$:

$$\Psi_{\text{Ogden}} = \sum_{p=1}^N \frac{\mu_p}{\alpha_p} (\lambda_1^{\alpha_p} + \lambda_2^{\alpha_p} + \lambda_3^{\alpha_p} - 3) + \frac{K}{2}(\ln J)^2. \quad (6)$$

This formulation offers greater flexibility in fitting complex, large-strain stress-stretch behavior, but with the need for an eigen-decomposition of the deformation tensor at each step, increasing computational cost [11].

Arruda-Boyce Motivated by a non-Gaussian statistical model of polymer chains arranged on an eight-chain unit cell, the Arruda-Boyce model captures the limiting chain extensibility responsible for strain-stiffening at large stretches:

$$\Psi_{\text{AB}} = \mu \sum_{i=1}^5 \frac{C_i}{\lambda_m^{2i-2}} (I_1^i - 3^i) + \frac{K}{2}(\ln J)^2, \quad (7)$$

where $C_1, \dots, C_5$ are fixed coefficients and λ_m is the locking stretch, the stretch at which the polymer network becomes fully extended, and it captures strain-stiffening behavior that the Neo-Hookean and Mooney-Rivlin models cannot [1].

2.2 Graph Neural Network for Soft Tissue Simulation

Graph neural networks propagate information through message-passing layers that update node coordinates and features based on local neighborhoods, making them well-suited for unstructured mesh and point cloud data [9]. Equivariance to rotations and translations allows the network to generalize across orientations without explicit data augmentation [19]. The conditional equivariant GNN of [13] builds on this principle, taking as input a surface point cloud $\mathbf{x}^0 \in \mathbb{R}^{N \times d}$ (N surface points in 3D), per-point tissue features $\mathbf{h} \in \mathbb{R}^{N \times d}$ encoding subsurface material composition sampled at 128 depth intervals beneath each point, and condition vectors $\mathbf{C}_s, \mathbf{C}_e \in \mathbb{R}^3$ representing the start and end coordinates of the applied force. Information is propagated through L equivariant message-passing layers, and predictions are produced by two dedicated branches: one regressing per-point displacements $\mathbf{u} \in \mathbb{R}^{N \times 3}$ and one estimating scalar contact force F , with the deformed point cloud recovered as $\hat{\mathbf{y}} = \mathbf{u} + \mathbf{x}^0$. We adopt this architecture without modification; full details are given in [13]. The sole adaptation concerns the tissue feature encoding $\mathbf{h}$, described in Section 3.4.

3 Methods

3.1 Physical Setup

Silicone beams of three different stiffnesses were fabricated by varying the softener concentration in the mixture. Each beam was cast using a two-part silicone rubber (Eurosil 4 $A + B$, Schouten SynTec, the Netherlands) combined with an oil-based softener (Eurosil Softener, Schouten SynTec, the Netherlands), mixed in an $A:B:S$ ratio by weight, with $A:B$ fixed at 1:1 and S varied at 0, 0.5, and 1. Throughout this paper, each mixture is referred to by its softener ratio, e.g. $S0$, $S0.5$, and $S1$ (corresponding to $A:B:S = 1:1:0$, $1:1:0.5$, and $1:1:1$).

The cylindrical molds were 3D-printed with a length of 200 mm and a radius of 10 mm. Due to minor silicone leakage during casting, the beams varied slightly in final length. To obtain consistent and comparable deflection measurements, each beam was tested at standardized lengths: 50 mm and 100 mm in a cantilevered configuration (clamped at one end, free at the other), and 150 mm in a clamped-clamped (fixed-fixed) configuration. In each case, deformation was recorded after the beam had fully stabilized under gravity. The deformed profiles were captured using a smartphone camera with a checkerboard reference pattern included in each frame (Fig. 2). We used OpenCV to compute a homography correcting perspective distortion and establishing an image resolution of 0.086 ± 0.006 mm/pixel. The corrected images were imported into CAD software and scaled using both the reference grid and the known diameter of an undeformed beam, minimizing scaling inaccuracies and compensating for perspective error from camera-to-specimen distance variation. The outer contour of the deformed specimen was manually traced from the calibrated images, and a centerline corresponding to the neutral axis was generated from this contour using the CAD software’s centerline function. The known cross-sectional diameter was then swept along this centerline, constrained by the traced outer contour, to reconstruct the 3D deformed geometry, serving as the reference for parameter calibration.




Fig. 2: Silicone beams at different lengths for the softener ratio $S0.5$.




Fig. 3: Optimization results for the Ogden model for softness $S0.5$.

3.2 Biomechanical Model & Parameter Optimization

For each silicone mixture, we simulated the gravity-induced deformation of a cylindrical mesh with 7257 nodes. The beam was fixed in SOFA, setting the boundary condition at the ends of the beam to match the physical experiments (Section 3.1). We rigidly constrained all translational degrees of freedom of every mesh node within a finite-length region at each clamped end (rather than a single end cross-section), matching the finite grip length of a physical clamp; no rotation or sliding was permitted within the clamped regions. The 50 and 100 mm beams were clamped at one end only (cantilevered, free at the other), while the 150 mm beam was clamped identically at both ends. The mass density was set to 1000 kg/m^3 for all beams as the density variation due to softener content was considered negligible.

We performed a parameter search over all hyperelastic constitutive models available in SOFA that are applicable to incompressible or near-incompressible soft materials. Each SOFA simulation ran for up to 10 000 steps at a 0.001 s time step. The simulation was run until equilibrium, defined as the root-mean-square (RMS) nodal displacement, measured as the change in node position over consecutive 50-step windows, remaining below 1 mm, averaged across three consecutive windows (150 steps total). For each model and each softener ratio S , parameters were optimized by maximizing the average volumetric overlap error across all three beam lengths, measured using the Dice score between the simulated and ground-truth 3D surface meshes. Optimizing across lengths simultaneously ensures the recovered parameters are consistent across boundary conditions rather than tuned to a single configuration. The optimization results and metrics for each model and softness level are reported in Table 1. We observe that most models could replicate the beam behavior with Dice scores > 0.90 for the $S0$ and $S0.5$ softnesses. Although the Ogden model was the slowest, it outperforms the other models in the $S1$ case and manages to reflect the softer tissue best. Based on these results, Ogden and Mooney-Rivlin were selected for data generation, as they achieved the most reliable fits across all three softness levels.

Table 1: Hyperelastic constitutive model calibration results across three softness levels ($S0$, $S0.5$, $S1$), optimized by maximizing volumetric overlap (Dice) between simulated and ground-truth beam deformations. For each softness level, metrics are averaged over the three tested beam lengths (50 mm, 100 mm, and 150 mm). Surface distance, 95th percentile Hausdorff distance (HD95), Dice score, volume difference, and simulation time are reported per softness level and as an average.

Model	Softness	Surface dist.↓ [mm]	HD95↓ [mm]	Dice↑	Vol. diff.↓ [%]	Time↓ [s]
Arruda- Boyce	$S0$	1.697	4.333	0.955	0.732	125.70
	$S0.5$	1.765	4.521	0.936	2.606	83.04
	$S1$	10.281	22.661	0.571	2.802	92.30
	Average	4.581	10.505	0.821	2.047	100.35
Mooney- Rivlin	$S0$	1.706	4.352	0.953	0.731	66.09
	$S0.5$	1.753	4.502	0.943	2.605	42.18
	$S1$	2.008	5.074	0.889	2.274	41.62
	Average	1.822	4.643	0.923	1.870	49.96
Neo-Hookean	$S0$	1.691	4.330	0.957	0.735	47.29
	$S0.5$	1.838	4.792	0.916	2.617	34.09
	$S1$	1.993	5.032	0.891	2.302	26.89
	Average	1.841	4.718	0.921	1.885	36.09
Stable Neo-Hookean	$S0$	1.718	4.349	0.951	0.738	45.06
	$S0.5$	1.842	4.727	0.914	2.625	28.50
	$S1$	3.474	9.625	0.728	2.232	30.40
	Average	2.345	6.234	0.868	1.865	34.65
Ogden	$S0$	1.773	4.420	0.934	0.710	245.48
	$S0.5$	1.677	4.380	0.967	2.583	220.20
	$S1$	1.877	4.822	0.926	4.068	153.31
	Average	1.776	4.541	0.942	2.454	203.33
St Venant- Kirchhoff	$S0$	1.700	4.319	0.952	0.948	20.12
	$S0.5$	1.994	5.054	0.872	3.507	35.89
	$S1$	2.631	6.696	0.739	3.628	18.82
	Average	2.108	5.356	0.854	2.694	24.95

3.3 Data Generation

Training data was generated using the two best-performing models from the parameter optimization study, Ogden and Mooney-Rivlin, applied across all three softness levels. We constructed a rectangular soft tissue volume of $100 \times 100 \times 32 \text{ mm}^3$ with an embedded rigid inclusion. Eleven distinct inclusion shapes were designed, comprising cubes, spheres and cylinders of varying sizes, randomly positioned at the base of the volume (Fig. 4). The inclusion surfaces and the bottom face together defined the rigid boundary, while the remaining volume was discretized into tetrahedral elements representing the soft tissue.

For each shape, we sampled two sets of poking locations on the top surface: 20 fixed locations shared across all three softness levels, enabling direct comparison of the tissue response to softness alone, and 20 randomly sampled locations




Fig. 4: Examples of generated binary shapes of soft tissues with rigid inclusions.

drawn independently per softness level to increase contact configuration diversity. In both cases, the poking angle was drawn up to 45° from the surface normal, with a maximum penetration depth of 30 mm.

The mesh was locally refined around each poking site and poking was performed using a cylindrical rod with a radius of 2.5 mm and length of 40 mm. Mesh resolution was determined via a convergence study on a homogeneous volume without rigid inclusion, comparing output forces at a penetration depth of 10 mm, chosen because not all models are guaranteed to reach the full 30 mm across all softness levels. The resulting dataset meshes ranged from 13 297 to 17 732 nodes (mean $15\,954 \pm 1254$) and 58 330 to 86 130 elements (mean $75\,792 \pm 8116$) across all shapes and softness levels.

Each poking simulation was prescribed over 2 s with a time step of 0.002 s. The recorded outputs at each time step were the positions of all surface nodes and the total reaction force at the boundary nodes. To remove near-static timesteps, frames were retained only when the surface displacement at the contact tip exceeded 1 mm relative to the previous retained frame, yielding a dataset of $11 \times 3 \times 40 = 1320$ pokings, with 14305 snapshots in the Ogden dataset and 13672 snapshots in the Mooney-Rivlin dataset. The complete features of the generated dataset can be seen in Table 2.

3.4 Training Details

We sampled 32×32 surface points on a regular grid from the FEM simulations to use as the input point cloud. In the original work [13], the depth profiles $\mathbf{h}$

Table 2: Features of the simulated dataset. Sim_{max} and F_{max} indicate the average maximum depth and force reached per poke. Tip_{mean} and F_{mean} indicates the average depth and force in each snapshot.

Model	Softness	Sim_{max} [mm]	F_{max} [N]	Tip_{mean} [mm]	F_{mean} [N]
Mooney-Rivlin	S0	13.92 ± 3.88	12.92 ± 5.32	7.91 ± 4.82	4.80 ± 4.59
	S0.5	13.78 ± 4.14	6.19 ± 2.69	7.94 ± 4.89	2.33 ± 2.30
	S1	13.46 ± 4.24	3.87 ± 2.43	7.82 ± 4.87	1.42 ± 1.54
Ogden	S0	16.32 ± 6.53	29.71 ± 25.31	9.88 ± 6.54	6.94 ± 11.44
	S0.5	16.40 ± 6.41	18.35 ± 18.67	9.89 ± 6.54	3.94 ± 7.47
	S1	11.20 ± 4.82	2.34 ± 1.98	7.04 ± 4.87	1.01 ± 1.19

were assigned a binary label: 0 for soft tissue and 1 for bone, reflecting the binary nature of the tissue. Since our dataset spans three silicone mixtures with distinct stiffnesses, we assigned each material a continuous stiffness value $s \in [0, 1]$, where higher values indicate stiffer material. The rigid inclusions retain a value of 1.0, while each silicone formulation is assigned a value derived from its softener ratio S , normalized such that the stiffest silicone $S0$ maps to 0.8, $S0.5$ maps to 0.4 and the softest $S1$ maps to 0.0. We split the data into training:validation:test in the order of shapes in an 7:2:2 ratio with 5-fold cross-validation, and trained the two datasets (Ogden and Mooney-Rivlin) separately for comparison, both including all softnesses.

The models were trained for 200 epochs using the Adam optimizer with a learning rate initialized to 10^{-4} and a batch size of 64. The loss function was defined the same as in [13], minimizing the weighted Euclidean distance error $\mathcal{L}_d$ and MSE for the force predictions $\mathcal{L}_f$. We report all deformation errors computed using the standard (unweighted) Euclidean distances.

The FEM simulations were performed on an AMD Epyc 7742 64-core CPU, and the GNN models were trained in PyTorch using the PyTorch Geometric (v2.8.0) package and run on a single NVIDIA RTX 2080 GPU.

4 Results & Discussion

Table 3 reports deformation and force prediction errors per softness level and per constitutive model used to generate the training data (Mooney-Rivlin and Ogden), alongside the total across all conditions.

Deformation accuracy Mean Euclidean distances are low and consistent across all conditions, averaging 0.095 ± 0.155 mm for Mooney-Rivlin and 0.114 ± 0.198 mm for Ogden, both below the 0.156 ± 0.298 mm reported in [13] for a single-material and shape model. This suggests that extending training to multiple softness levels, combined with the continuous stiffness encoding, does not degrade deformation accuracy relative to the single-material baseline. The lower errors are

Table 3: GNN prediction results for the different datasets and softness levels.

Dataset	Softness	Force Absolute Error [N]	Mean Euclidean Distance [mm]	Max Euclidean Distance [mm]	Mean Relative Tip Error [%]
Mooney- Rivlin	$S0$	0.537 ± 1.001	0.097 ± 0.152	1.205 ± 0.854	2.3 ± 2.5
	$S0.5$	0.310 ± 0.540	0.097 ± 0.155	1.228 ± 0.924	2.1 ± 2.3
	$S1$	0.258 ± 0.650	0.092 ± 0.157	1.268 ± 0.944	2.2 ± 2.6
	Average	0.337 ± 0.708	0.095 ± 0.155	1.239 ± 0.918	2.2 ± 2.4
Ogden	$S0$	2.269 ± 6.358	0.127 ± 0.209	1.496 ± 1.084	2.5 ± 2.6
	$S0.5$	1.407 ± 5.155	0.132 ± 0.219	1.517 ± 1.104	2.4 ± 2.5
	$S1$	0.222 ± 0.776	0.079 ± 0.145	1.237 ± 1.013	2.1 ± 2.7
	Average	1.253 ± 4.751	0.114 ± 0.198	1.427 ± 1.080	2.3 ± 2.6

likely due to the increased diversity of training configurations and shapes. The mean relative tip error remains consistently around 2–2.5% across all conditions for both models, further confirming reliable deformation prediction across the stiffness range.

Force prediction Force errors are more variable and show a clear dependence on both the constitutive model and softness level. For Mooney-Rivlin, force errors are low and consistent across softness levels (0.258 N to 0.537 N), compared to those reported in [13] (0.701 ± 1.111 N). For Ogden, force errors are substantially higher and more variable for the stiffer mixture ($S0$: 2.269 ± 6.358 N, before dropping sharply for the softest level ($S1$: 0.222 ± 0.776 N). This pattern directly reflects the characteristics of the training data: as seen in the data features in Table 2, the Ogden model produces large and highly variable forces for stiffer materials, particularly $S0$ and $S0.5$, where the simulation reaches greater penetration depths before stabilizing. The resulting high variance in training targets makes force regression harder, and more training examples or a stronger tissue stiffness encoding in the model would likely reduce these errors. This high variance also partly reflects the imperfect calibration fit for stiffer Ogden conditions: parameters that do not perfectly reproduce physical behavior produce inconsistent force outputs across similar configurations, which directly increases the difficulty of the regression task.

Effect of constitutive model The Mooney-Rivlin dataset produces more consistent GNN predictions overall, which we attribute to its more uniform force distributions across softness levels rather than superior physical accuracy. The Mooney-Rivlin simulations reach shallower maximum depths on average (Table 2), which limits force variability and results in a more learnable dataset. Ogden, by contrast, sustains larger deformations and generates physically richer force profiles, particularly for stiffer tissues, but at the cost of higher variance that the current model struggles to fit. This distinction is practically important, as Mooney-Rivlin may be preferable when training data consistency is the priority, while Ogden may be more appropriate when the goal is to cover the full range of tissue responses, provided sufficient training data and a more expressive tissue encoding are available. In both cases, deformation accuracy remains comparable, suggesting that the GNN’s displacement prediction is robust to the choice of constitutive model, while force prediction is more sensitive to training data characteristics.

Across both datasets, a single inference took 0.010 ± 0.004 s (100 Hz), providing close to real-time predictions required for interactive surgical simulations. Additionally, as demonstrated in [13], the underlying architecture generalizes across point cloud densities and structures, a property inherited by our extension without additional training.

5 Conclusions

We presented a pipeline for generalizable soft tissue simulation combining systematic calibration of SOFA hyperelastic constitutive models with a softness-conditioned graph neural network surrogate. Our calibration study identified Ogden and Mooney-Rivlin as the most reliable models across tested stiffness levels, while revealing the limitations of current constitutive models for very compliant materials. Conditioning the GNN on material stiffness enabled accurate deformation and force across three trained material softness levels and unseen inclusion geometries, maintaining sub-millimeter mean deformation accuracy at close to real-time inference speeds. These findings support the use of carefully calibrated FEM simulations as a basis for fast, generalizable data-driven surrogates in surgical simulation. The close to real-time inference speed of 0.010 s demonstrates the practical viability of such surrogates for integration into interactive surgical training systems with haptic feedback.

References

1. Arruda, E.M., Boyce, M.C.: A three-dimensional constitutive model for the large stretch behavior of rubber elastic materials. *Journal of the Mechanics and Physics of Solids* **41**(2), 389–412 (1993). DOI [https://doi.org/10.1016/0022-5096\(93\)90013-6](https://doi.org/10.1016/0022-5096(93)90013-6)
2. Barbič, J., James, D.L.: Real-time subspace integration for st. venant-kirchhoff deformable models. *ACM Trans. Graph.* **24**(3), 982–990 (2005). DOI 10.1145/1073204.1073300
3. Bertram, A.: Hyperelasticity. In: *Elasticity and Plasticity of Large Deformations*, pp. 211–232. Springer International Publishing, Cham (2021)
4. Cotin, S., Delingette, H., Ayache, N.: Real-time elastic deformations of soft tissues for surgery simulation. *IEEE Trans. Vis. Comput. Graph.* **5**(1), 62–73 (1999). DOI <https://doi.org/10.1109/2945.764872>
5. Deshpande, S., Bordas, S.P., Lengiewicz, J.: Magnet: A graph u-net architecture for mesh-based simulations. *Engineering Applications of Artificial Intelligence* **133**, 108,055 (2024). DOI <https://doi.org/10.1016/j.engappai.2024.108055>
6. El Hadramy, S., Acidi, B., Padoy, N., Cotin, S.: Optimization in latent space for real-time intraoperative characterization of digital twins. *Eng. with Comput.* **42**(2) (2026). DOI 10.1007/s00366-025-02268-z
7. El Hadramy, S., Padoy, N., Cotin, S.: Hyperu-mesh: Real-time deformation of soft-tissues across variable patient-specific parameters. In: *Computational Biomechanics for Medicine*, pp. 129–139. Springer Nature Switzerland, Cham (2025). DOI https://doi.org/10.1007/978-3-031-94128-3_12
8. Faure, F., Duriez, C., Delingette, H., Allard, J., Gilles, B., Marchesseau, S., Talbot, H., Courtecuisse, H., Bousquet, G., Peterlik, I., Cotin, S.: SOFA: A Multi-Model Framework for Interactive Physical Simulation. In: Y. Payan (ed.) *Soft Tissue Biomechanical Modeling for Computer Assisted Surgery, Studies in Mechanobiology, Tissue Engineering and Biomaterials*, vol. 11, pp. 283–321. Springer (2012). DOI 10.1007/8415\2012\125
9. Hamilton, W.L.: Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* **14**(3), 1–159 (2020)

10. Kenngott, H.G., Wagner, M., Gondan, M., Nickel, F., Nolden, M., Fetzner, A., Weitz, J., Fischer, L., Speidel, S., Meinzer, H.P., Böckler, D., Büchler, M.W., Müller-Stich, B.P.: Real-time image guidance in laparoscopic liver surgery: first clinical experience with a guidance system based on intraoperative ct imaging. *Surgical Endoscopy* **28**(3), 933–940 (2014). DOI 10.1007/s00464-013-3249-0
11. Kim, B., Lee, S.B., Lee, J., Cho, S., Park, H., Yeom, S., Park, S.: A comparison among neo-hookean model, mooney-rivlin model, and ogden model for chloroprene rubber. *International Journal of Precision Engineering and Manufacturing* **13** (2012). DOI 10.1007/s12541-012-0099-y
12. Kojanazarova, M., Bieder, F., Sandkühler, R., Cattin, P.C.: Conditional graph neural network for predicting soft tissue deformation and forces. In: *Computational Biomechanics for Medicine*, pp. 62–73. Springer Nature Switzerland, Cham (2025). DOI https://doi.org/10.1007/978-3-031-94128-3_7
13. Kojanazarova, M., Hadramy, S.E., Wilkie, J., Rauter, G., Cattin, P.C.: Soft tissue simulation and force estimation from heterogeneous structures using equivariant graph neural networks. *Healthcare Technology Letters* **12**(1), e70,042 (2025). DOI <https://doi.org/10.1049/htl2.70042>
14. Mazier, A., Bilger, A., Forte, A.E., Peterlik, I., Hale, J.S., Bordas, S.P.A.: Inverse deformation analysis: an experimental and numerical assessment using the fenics project. *Engineering with Computers* **38**(5), 4099–4113 (2022). DOI 10.1007/s00366-021-01597-z
15. Mendizabal, A., Márquez-Neila, P., Cotin, S.: Simulation of hyperelastic materials in real-time using deep learning. *Med. Image Anal.* **59**, 101,569 (2020). DOI <https://doi.org/10.1016/j.media.2019.101569>
16. Mihai, L.A., Chin, L., Janmey, P.A., Goriely, A.: A comparison of hyperelastic constitutive models applicable to brain and fat tissues. *Journal of The Royal Society Interface* **12**(110), 20150,486 (2015). DOI 10.1098/rsif.2015.0486
17. Nguyen, C.C., Wong, S., Thai, M.T., Hoang, T.T., Phan, P.T., Davies, J., Wu, L., Tsai, D., Phan, H.P., Lovell, N.H., Do, T.N.: Advanced user interfaces for teleoperated surgical robotic systems. *Advanced Sensor Research* **2**(4), 2200,036 (2023). DOI <https://doi.org/10.1002/adsr.202200036>
18. Nguyen, T.N., Ho Ba Tho, M.C., Dao, T.T.: A systematic review of real-time medical simulations with soft-tissue deformation: Computational approaches, interaction devices, system architectures, and clinical validations. *Applied Bionics and Biomechanics* **2020**(1), 5039,329 (2020). DOI <https://doi.org/10.1155/2020/5039329>
19. Satorras, V.G., Hoogeboom, E., Welling, M.: E(n) equivariant graph neural networks. In: M. Meila, T. Zhang (eds.) *Proceedings of the 38th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol. 139, pp. 9323–9332. PMLR (2021). DOI <https://doi.org/10.1021/acs.jpcb.3c07304>
20. Smith, B., Goes, F.D., Kim, T.: Stable neo-hookean flesh simulation. *ACM Trans. Graph.* **37**(2) (2018). DOI 10.1145/3180491
21. Taylor, Z.A., Cheng, M., Ourselin, S.: High-speed nonlinear finite element analysis for surgical simulation using graphics processing units. *IEEE Trans. Med. Imaging* **27**(5), 650–663 (2008)