Deform Any Liver In Real-time

Sidaty EL HADRAMY, Nicolas Padoy, Stéphane Cotin

DEFORM ANY LIVER IN REAL-TIME

Sidaty El Hadramy[⋆†] *Nicolas Padoy*[†‡] *Stéphane Cotin*[⋆†]

[⋆] Inria, Strasbourg, France
[†] University of Strasbourg, CNRS, INSERM, ICube, UMR7357, Strasbourg, France
[‡] IHU Strasbourg, France

ABSTRACT

This paper introduces Deform Any Liver (DAL), a real-time surrogate model designed to accurately predict any liver deformation, given a set of applied forces. DAL leverages a hypernetwork to condition a U-Net architecture on the liver shape, with the U-Net receiving the applied external forces as input to predict the resulting displacement field. This design enables the model to generalize across various patient geometries, overcoming the limitations of existing patient-specific approaches. The DAL model is trained, validated, and tested using real human liver geometries, employing finite element method (FEM) simulations to generate realistic deformation data. Our results demonstrate the model's effectiveness and adaptability in accurately predicting the deformation of various liver shapes.

***Index Terms*—** HyperNetworks, Liver geometric variation, Real-time simulation.

1. INTRODUCTION

Soft-tissue simulations can advance healthcare in various critical areas, including preoperative planning, medical training, and intraoperative guidance [1, 2]. Simulations enable surgeons and medical trainees to visualize, manipulate, and anticipate tissue behavior in a controlled environment by providing a realistic and dynamic representation of soft tissues. Numerous studies have highlighted the advantages of physics-based simulations due to their accuracy and predictive capacity when dealing with sparse or noisy intraoperative data [3, 4]. These simulations typically rely on continuum mechanics principles and frequently use finite element methods (FEM) to solve numerical systems that describe the physical behavior of organs. Although these models enable highly accurate simulations of organ and tissue dynamics, their computational demands often lead to prolonged calculation times, mainly when modeling the hyper-elastic behavior characteristic of soft tissues [4]. This limits their applicability in scenarios requiring real-time processing, such as intraoperative guidance, where delays can impact the precision and effectiveness of surgical interventions.

Several works have proposed a trade-off between accuracy and computational expense, balancing the degree of accuracy loss to suit specific application needs. One strategy is model order reduction (MOR) [5], which simplifies simulations by approximating the solution space and thus accelerates computations. MOR methods include proper orthogonal decomposition (POD) [6] and Proper Generalized Decomposition (PGD) [7] reduced basis approaches, often used in FEM-based simulations. These snapshot-based methods, however, come at the price of long preprocessing times and their accuracy is generally inversely proportional to their computation time. Deep neural network architectures have shown remarkable effectiveness in capturing complex, nonlinear relationships between diverse sets of input and output data. A key advantage of these networks is their capacity for accurate real-time inference once trained with sufficient data. Building on this strength, numerous studies have proposed neural network-based surrogate models trained on FEM simulation data [8, 9, 10]. These models replicate complex FEM behaviors while reducing computational demands, enabling fast, accurate predictions suitable for real-time applications.

Various neural network architectures have been explored for FEM-based surrogate modeling, where the goal is to predict a dense volumetric displacement field given a set of external forces while incorporating the hyperelastic behavior of materials as prior knowledge. To cite a few, Odot *et al.* [11] proposed a multi-layer perceptron for this purpose. However, their approach loses spatial information between finite element mesh nodes. To address this limitation, Mendizabal *et al.* [8] introduced U-Mesh, a U-Net architecture designed around a structured finite-element mesh to align with the CNN architecture. Yet, structured meshes can introduce errors when applied to complex geometries like the liver. To overcome this, Deshpande *et al.* [9] proposed MagNet, a U-Net architecture based on Graph Neural Networks, which preserves the unstructured nature of the geometry mesh. El Hadramy *et al.* [12] extended Mendizabal's approach with HyperU-Mesh, which uses a hypernetwork to condition the U-Mesh on patient-specific material properties, making it possible to predict over a range of material properties. While each method shows promising results, they all share a limitation: they are patient-specific, meaning they are trained

on the exact geometry intended for use at inference. This limits flexibility and transferability to new patients, as each geometry requires a separate and time-intensive training.

In this work, we propose the *Deform Any Liver* (DAL) model, which builds upon the HyperU-Mesh [12] framework by conditioning the U-Mesh architecture on liver geometry. This geometric conditioning enables the network to generalize across different liver shapes, accurately predicting deformations for various liver geometries. By directly incorporating the liver shape into the model, DAL overcomes the limitations of patient-specific approaches, offering a more flexible and efficient solution to handle anatomical variations.

2. METHOD

We aim to provide a versatile surrogate model capable of deforming any liver geometry given any set of external forces. The core of our method is a HyperNetwork architecture, where the primary network, a U-Net, receives the applied forces as input and predicts the resulting displacement field (see Fig. 1). The U-Net weights are dynamically updated by the hypernetwork, which takes as input a 3D mask representing the liver shape in its rest state. This conditioning mechanism allows the model to adapt to different liver geometries, enabling accurate deformation predictions across varying anatomical shapes.

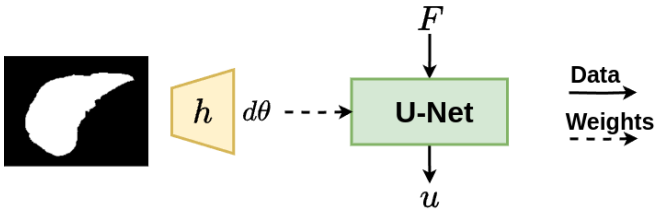

Fig. 1: Overview of Deform Any Liver (DAL). The model uses a U-Net architecture that inputs external forces (F) and predicts the displacement field (u). It is conditioned on liver geometries provided as a 3D mask which is fed into a hypernetwork. The hypernetwork computes additive weights ($d\theta$) for the U-Net. DAL is trained, validated, and tested using real human liver geometries, with FEM simulations used to generate the corresponding deformations.

2.1. DAL architecture

2.1.1. Main network

The architecture of the main network in the DAL model builds upon the U-Mesh framework proposed by Mendizabal *et al.* [8], with specific optimizations that reduce model complexity without compromising performance. By adjusting the number of convolutional layers and channels, we reduced the total weight count to 167563 parameters, significantly lower than the million order weights in the original U-Mesh. This weights optimization ensures that the hypernetwork h can efficiently predict the weights of the U-Net, making the model more computationally feasible. Like U-Mesh, the DAL network operates on a regular grid constructed from the liver shape axis-aligned bounding box. Applied forces are encoded at the regular grid nodes nearest to the force application points on the liver's surface, and the network predicts the displacement field in this same structured format. For further details on U-Mesh, see Mendizabal *et al.* [8].

2.1.2. Hypernetwork

The hypernetwork in the DAL model is composed of a CNN followed by an MLP. It inputs a 3D liver segmentation mask to condition the main network according to the specific liver geometry. The mask undergoes processing through three 3D convolutional layers, featuring channel sizes of $[1, 4, 4]$ and a kernel size of 3, with ReLU activation functions applied after each convolutional layer. A max-pooling layer follows each convolutional layer to reduce the spatial dimensions. The output from the final convolutional layer is then flattened and fed into a fully connected layer, which produces a vector of size 167563, corresponding to the number of weights needed for the main network. We employ the hypernetwork strategy from Ortiz et *al.* [13], which treats the predicted weights as additive modifications to the trainable weights of the main network. This approach enhances model flexibility and stability, leading to faster convergence of the loss function during training while improving performance in predicting liver deformations. We employ the hypernetwork strategy from Ortiz *et al.* [13], which processes the predicted weights as additive changes to the trainable weights of the main network. This allows for fast and stable training.

2.2. Liver geometries dataset

To obtain a diverse range of liver shapes, we collected a dataset of liver geometries from two public datasets. The first source is the Liver Tumor Segmentation (LiTS) [14] challenge from MICCAI 2017 and ISBI 2017, which includes 92 labeled liver segmentations. The closed-surface meshes were automatically extracted from the 3D segmentation masks using the Marching Cubes algorithm [15]. The resulting meshes were then automatically cleaned and uniformly resampled to guarantee a uniform distribution of the vertices on the surface. Our dataset comprises 92 liver meshes, enabling robust training and validation for our model. We use 78 geometries for training and 14 for testing. In Figure 2, we illustrate five shapes from the dataset.

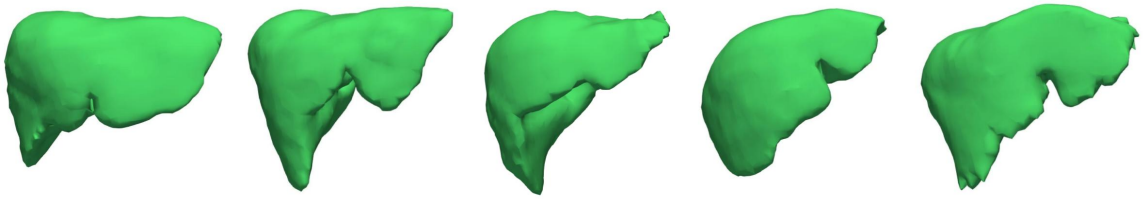

Fig. 2: Five arbitrarily chosen examples of liver shapes from the dataset, showing variations in local geometry, particularly in the size of the left lobe.

2.3. Data generation

The training of the DAL model necessitates the acquisition of applied forces as inputs for each liver shape, along with the corresponding ground truth displacement field for supervision. To obtain these quantities, we utilize physics simulations that accurately model the mechanical behavior of the liver based on its real anatomical shapes. By simulating the effects of various applied forces, we generate the required displacement fields, which serve as the ground truth for training the main network. For a given liver shape, we formulate a boundary value problem to compute the deformation of a hyperelastic material under both Dirichlet and Neumann boundary conditions. The liver occupies a volume Ω, and the Dirichlet boundary conditions, denoted Γ_D, are applied to a manually selected region of each liver shape, specifically targeting the posterior part of the liver, which corresponds to the parenchyma's connection with the inferior vena cava. While Neuman boundary conditions, denoted Γ_N, are everywhere else on the liver surface. The material behavior is approximated with a Saint-Venant Kirchhoff model with a Young modulus of 7 kPa. The boundary value problem is formulated as follows:

$$\begin{cases} \nabla(FS) = b \quad \text{on} \quad \Omega \\ u(X) = 0 \quad \text{on} \quad \Gamma_D \\ (FS) \cdot n = t \quad \text{on} \quad \Gamma_N \end{cases} \tag{1}$$

Where F represents the deformation gradient tensor and S is the second Piola-Kirchoff stress tensor. b is the body forces, n the unit normal to Γ_N, u the displacement field and X is the material coordinates. The weak form of Equation 1 brings forward the boundary term and is expressed as:

$$\int_\Omega (\boldsymbol{FS}) : \boldsymbol{\delta E}\, d\Omega = \int_\Omega \boldsymbol{b\eta}\, \boldsymbol{d\Omega} + \int_{\Gamma_N} \boldsymbol{t\eta}\, \boldsymbol{d\Gamma} \tag{2}$$

Where $\boldsymbol{\delta E} = \frac{1}{2}(\boldsymbol{F^T \nabla \eta + \nabla^T \eta F})$ is the variation of the strain, and $\boldsymbol{\eta} = \{\boldsymbol{\eta} \in \boldsymbol{H^1(\Omega)} \mid \eta = 0 \; on \; \Gamma_D\}$ is any vector-valued test function in an Hilbert space $\boldsymbol{H^1(\Omega)}$. Equation 2 is solved by employing the finite element method with the domain Ω discretized using an Immersed Boundary method (IBM) [16]. This allows for a regular structure of the liver shape, which makes it compatible with the inputs and outputs of the CNN-based U-Net architectures. We have generated $8,000$ deformations from the 78 training geometries. To create each deformation, four traction forces are sampled from a uniform distribution ranging from 1 N to 8 N for their magnitudes, with their application locations randomly selected on the liver surface. This process is illustrated in Figure 3, where the gray liver represents the initial state, the blue spheres indicate the areas where the forces are applied, and the red liver depicts the resulting deformation. Moreover, we have generated 1000 deformations from the 14 test geometries. These samples are used for network testing.

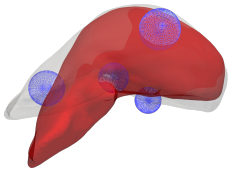

Fig. 3: Example of a training set generation sample. Spheres (in blue) centers are randomly chosen on the surface's initial state (gray). Forces are applied at their intersection with the surface. In red, the resulting deformation is shown.

2.4. DAL training

The network was supervised during 400 epochs with a batch size of 1 using $8,000$ training samples. Each sample comprises the applied forces, the corresponding liver geometry mask, and the ground truth displacement field. During training, we optimize the network's weights by minimizing the mean squared error (MSE) between the predicted and the ground truth displacement fields, allowing it to learn the relationship between the input forces and the resulting liver deformations for any liver geometry expressed as a mask, input to the hypernetwork.

3. RESULTS

3.1. Implementation details

In this section, we present the results of our method across two scenarios. We evaluate DAL by deforming the 14 ge-

ometries from our testing dataset, which were not seen during training. We compare the deformation results with those obtained from finite element method (FEM) solutions. Next, we show that DAL can be fine-tuned rapidly on a specific liver shape to become patient-specific. The implementation of DAL is carried out in Python using the PyTorch[1] library, while the simulations used for data generation and testing are implemented within the SOFA framework [17, 18].

3.2. Evaluation over the test set

We evaluated DAL on 14 liver shapes from the test set, which were not seen during training. Each liver was subjected to 100 deformations, resulting from external forces randomly sampled from a uniform distribution between 1 and 8 N and applied on the liver surface at random locations. We compared DAL's deformation results to those obtained with FEM, using the mean absolute error (MAE) between the node positions in the DAL and FEM-deformed livers. DAL achieved an average MAE of 5.9 ± 2.7 mm across the 14 shapes. The shape with the lowest error had an MAE of 3.8 ± 1.8 mm, while the shape with the highest error reached an MAE of 9.3 ± 4.4 mm. Moreover, DAL's prediction takes 4 ms, compared to 3000 ms, when computing the deformation with the FEM. We have observed a high correlation between the errors on test shapes and their distance from the distribution of liver shapes used during training. Specifically, after visual inspection, we observed that the liver shape with the highest errors (red liver in Figure 4) differs significantly from all training shapes in size and overall geometry. Conversely, the shape with the lowest errors (green in Figure 4) is visually similar to those seen during training.

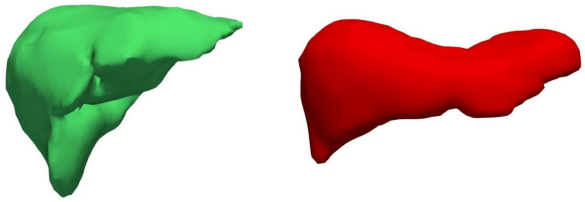

Fig. 4: DAL's performance on two test liver shapes. The shape with the highest error is shown in red (worst result), while the shape with the lowest error is depicted in green (best result). This illustrates the variability in DAL's accuracy depending on the similarity of test shapes to the training dataset.

3.3. Patient-specific DAL

Since DAL is trained on a diverse dataset of liver geometries and has already learned the general force-displacement field relationship for various liver shapes, this foundational knowledge can be effectively transferred through additional, shape-specific training. Fine-tuning DAL on a complex liver shape, like the one associated with the worst results, would enable DAL to adapt to unique anatomical variations, enhancing its performance for individualized, patient-specific applications. To demonstrate this, we generated 700 deformations of the previously identified worst-case liver shape and re-trained DAL on 500 of these generated deformations. We then tested the fine-tuned model on the other 200 samples. The results were compared to the deformations obtained using FEM, and DAL achieved a mean absolute error (MAE) of 2.5 ± 0.8 mm, a significant improvement over the previous MAE of 9.3 ± 4.4 mm. Notably, this fine-tuning process only took 30 minutes, compared to the 12 hours required for training a patient-specific U-Mesh model from scratch, as proposed by Mendizabal *et al.* [8]. With fine-tuned DAL demonstrating similar performance to U-Mesh, less than 3 mm of errors. This highlights DAL's potential for rapid adaptation to complex, patient-specific liver geometries.

4. CONCLUSION

To provide real-time and accurate simulations of the liver, we introduce the Deform Any Liver (DAL) model, a surrogate approach for predicting liver deformations across various geometries. By utilizing a hypernetwork to condition a U-Net architecture, DAL generalizes beyond individual patient shapes while maintaining good accuracy. The hypernetwork inputs a segmentation mask that delineates the liver's geometry, providing information about its shape. Meanwhile, the U-Net processes the applied forces to predict the resulting displacement field based on this geometric input, allowing DAL to capture the relationship between the applied forces and liver deformations. Our evaluations demonstrate that DAL achieves good results compared to the finite element method (FEM) simulations while delivering real-time performance. Trained on 78 liver shapes and tested on 14, DAL achieves a mean absolute error of 5.9 ± 2.7 mm. Additionally, DAL can be fine-tuned for patient-specific complex geometries in just half an hour, achieving less than 3 mm of error, similar to U-Mesh [8].

DAL was evaluated on liver shapes, which happen to vary significantly from one patient to another, as illustrated in Fig 2. This large variability would certainly require more samples to improve the prediction accuracy. On the other hand, we believe that DAL could perform better on other organs, such as the kidney or brain, for which anatomical variations are less important. Future work will focus on expanding the dataset and enhancing the model's robustness to accommodate complex anatomical variations. We will also explore alternative methods for describing the geometry as input for the hypernetwork. A signed distance function could be a promising approach, as it not only represents the shape of the liver but also encodes geometric information about the surface.

[1] https://pytorch.org/

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study was conducted retrospectively using human subject data made available in open access. Ethical approval was *not* required as confirmed by the license attached with the open access data.

6. ACKNOWLEDGEMENTS

This work was partially supported by French state funds managed by the ANR under reference ANR-10-IAHU-02 (IHU Strasbourg).

7. REFERENCES

[1] Suwelack S, Röhl S, Bodenstedt S, Reichard D, Dillmann R, dos Santos T, Maier-Hein L, Wagner M, Wünscher J, Kenngott H, Müller BP, and Speidel S, "Physics-based shape matching for intraoperative image guidance," *Med Phys. 41(11):111901*, 2014.

[2] Delingette H. and N. Ayache, "Soft tissue modeling for surgery simulation. handbook of numerical analysis," *IEEE transactions on medical imaging*, 2009.

[3] Alvarez P., Chabanas M., Rouzé S., Castro M., Payan Y., and J. L. Dillenseger, "Lung deformation between preoperative ct and intraoperative cbct for thoracoscopic surgery: a case study," *Medical Imaging, Vol. 10576D*, 2018.

[4] Haouchine N., Dequidt J., Peterlik I., Kerrien E., Berger M O., and Cotin S., "Image-guided simulation of heterogeneous tissue deformation for augmented reality during hepatic surgery," in *2013 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. Oct. 2013, pp. 199–208, IEEE.

[5] Koutsovasilis P. and Beitelschmidt M., "Model order reduction of finite element models: improved component mode synthesis," *Math. Comput. Model. Dyn. Syst.*, vol. 16, no. 1, pp. 57–73, Mar. 2010.

[6] Niroomandi S., Alfaro I., Cueto E., and Chinesta F., "Real-time deformable models of non-linear tissues by model reduction techniques," *Comput. Methods Programs Biomed.*, pp. 223–231, Sept. 2008.

[7] Niroomandi S., Gonzalez D., Alfaro I., Bordeu F., Leygue A., Cueto E., and Chinesta F., "Real-time simulation of biological soft tissues: a PGD approach," *Int. J. Numer. Method. Biomed. Eng.*, vol. 29, no. 5, pp. 586–600, May 2013.

[8] Brunet J. N., Mendizabal A., Petit A., Golse N., Vibert E., and Cotin S., "Physics-based deep neural network for augmented reality during liver surgery," *Medical Image Computing and Computer Assisted Intervention. Cham: Springer International Publishing*, 2019.

[9] Deshpande S, Bordas S P A, and Lengiewicz J., "MAgNET: A graph U-Net architecture for mesh-based simulations," *Eng. Appl. Artif. Intell.*, vol. 133, no. 108055, pp. 108055, July 2024.

[10] El Hadramy S., Verde J., Padoy N., and Cotin S., "Towards real-time vessel guided augmented reality for liver surgery," in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. May 2024, vol. 1, pp. 1–5, IEEE.

[11] Odot A., Haferssas R., and Cotin S., "DeepPhysics: a physics aware deep learning framework for real-time simulation," Sept. 2021.

[12] El Hadramy S., Padoy N., and Cotin S., "HyperU-Mesh: Real-time deformation of soft-tissues across variable patient-specific parameters," in *Computational Biomechanics Workshop at MICCAI 2024*. 2024.

[13] Ortiz J. J. G., Guttag J., and Dalca A., "Magnitude invariant parametrizations improve hypernetwork learning," 2023.

[14] Bilic P et al., "The liver tumor segmentation benchmark (LiTS)," *Med. Image Anal.*, vol. 84, no. 102680, pp. 102680, Feb. 2023.

[15] Lorensen W. E. and Cline H. E., "Marching cubes: A high resolution 3D surface construction algorithm," *Comput. Graph. (ACM)*, vol. 21, no. 4, pp. 163–169, Aug. 1987.

[16] Duprez M., Lleras V., and Lozinski A., "-FEM: an optimally convergent and easily implementable immersed boundary method for particulate flows and stokes equations," *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 57, no. 3, pp. 1111–1142, May 2023.

[17] Faure F., Duriez C., Delingette H., Allard J., Gilles B., Marchesseau S., and Cotin S., "Sofa: A multi-model framework for interactive physical simulation," *Soft tissue biomechanical modeling for computer assisted surgery, 283-321*, 2012.

[18] Mazier A., El Hadramy S., Brunet J-N., Hale J S., Cotin S, and Bordas S P A., "Sonics: develop intuition on biomechanical systems through interactive error controlled simulations," *Eng. Comput.*, vol. 40, no. 3, pp. 1857–1876, June 2024.