

Intraoperative Liver Surface Completion with Graph Convolutional VAE

Simone Foti^{1,2}, Bongjin Koo^{1,2}, Thomas Dowrick^{1,2}, João Ramalhinho^{1,2},
Moustafa Allam³, Brian Davidson³, Danail Stoyanov^{1,2}, and
Matthew J. Clarkson^{1,2}

¹ Wellcome/EPSRC Centre for Interventional and Surgical Sciences,
University College London, London, UK

² Centre For Medical Image Computing, University College London, London, UK

³ Division of Surgery and Interventional Science, University College London, London,
UK

`s.foti@cs.ucl.ac.uk`

Abstract. In this work we propose a method based on geometric deep learning to predict the complete surface of the liver, given a partial point cloud of the organ obtained during the surgical laparoscopic procedure. We introduce a new data augmentation technique that randomly perturbs shapes in their frequency domain to compensate the limited size of our dataset. The core of our method is a variational autoencoder (VAE) that is trained to learn a latent space for complete shapes of the liver. At inference time, the generative part of the model is embedded in an optimisation procedure where the latent representation is iteratively updated to generate a model that matches the intraoperative partial point cloud. The effect of this optimisation is a progressive non-rigid deformation of the initially generated shape. Our method is qualitatively evaluated on real data and quantitatively evaluated on synthetic data. We compared with a state-of-the-art rigid registration algorithm, that our method outperformed in visible areas.

Keywords: Laparoscopic Liver Surgery · Geometric Deep Learning
· Graph Convolution · Surface Completion · Variational Autoencoder.

1 Introduction

The loss of direct vision and tactile feedback in laparoscopic procedures introduces an additional level of complexity for surgeons. Augmented reality (AR) is a promising approach to alleviate these limitations and provide guidance during the procedure. However, it remains an open challenge for laparoscopic surgery of the liver, which is one of the largest and most deformable organs. AR is usually achieved by registering a preoperative 3D model to the intraoperative laparoscopic view. Clinically available state-of-the-art commercial systems use manual point-based rigid registration [14], while recent research works focus on either rigid [12,16] or deformable [3,6,10,13] registration techniques requiring different

amounts of manual interactions and computations on the preoperative data. In contrast, we formulate the deformable registration as a shape completion problem that does not rely on patient specific preoperative computations.

Even though the underlying techniques are different, the common application and presence of an optimisation procedure make our method similar to registration. We believe that our method has the potential to become a precursor to a new approach for registration. Thus, we directly compare our method with a rigid registration algorithm (Go-ICP [20]) that aligns two point clouds. This algorithm was successfully used for laparoscopic liver applications in [12], where the preoperative model was registered onto the intraoperative point cloud obtained using an unsupervised neural network for depth estimation. Our method is similar, but relying on a manual interaction it is also able to predict a deformed model that better fits the point cloud. Other methods, such as [10,13] attempt the registration of preoperative models directly on the intraoperative images requiring manual image annotations. Even though they still show high errors in areas not visible from the camera, these methods showed extremely good performances in coping with deformations. Both use biomechanical models to simulate deformations, and [13] requires an additional preoperative step where multiple possible patient-specific simulations have to be performed.

The most similar works to ours are [1] and [11]. The former leverages a voxel-based conditional variational autoencoder (VAE) to complete missing segments of bone and plan jaw reconstructive surgical procedures. Not only the anatomical structures considered in their work are not deformable and the missing segments are small compared to the complete shape, but also their solution is constrained by the remaining healthy structures. On the other hand, our problem is more ill-posed because the liver is highly deformable and the missing parts are much larger than the partial intraoperative shape. In addition, voxel-based representations of shapes are inefficient volume representations that struggle to achieve high resolutions and to handle deformations. Therefore, we represent 3D shapes as Riemannian manifolds discretised into meshes and use geometric deep learning techniques to process these data. In particular, our work adapts [11] to achieve shape completion in laparoscopic liver surgery by *(i)* overcoming the shortage of data; *(ii)* compensating the lack of point correspondences between partial and complete shapes; *(iii)* redefining the VAE training loss to deal with non-uniformly sampled meshes; and *(iv)* leveraging preoperative data for the initialisation. The optimisation process for shape completion, makes the methodology suitable for registration, but there are a few key steps that need innovating. We believe this is the first attempt to bring geometric deep learning in to computer assisted interventions.

2 Methods

The proposed method (Fig. 1) estimates the complete mesh of a liver given a partial point cloud of its surface. A graph convolutional variational autoencoder

is trained to generate complete shapes (Sec. 2.1) and a distinct optimisation procedure non-rigidly deforms them to fit the partial point cloud (Sec. 2.2).

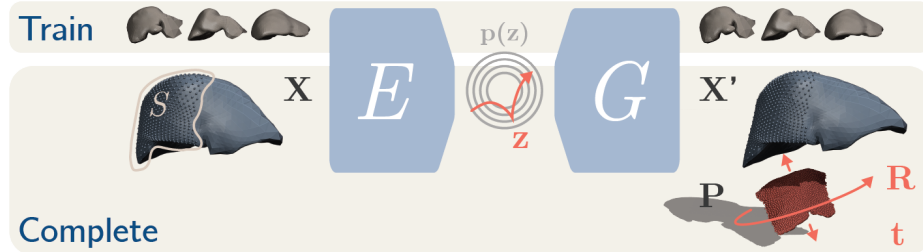


Fig. 1. Schematic Description of the Proposed Method. *Top:* a VAE ($\{E, G\}$) is trained on complete preoperative meshes of the liver. *Bottom:* the shape completion starts with a manual selection S on the preoperative mesh \mathbf{X} . The latent representation obtained encoding \mathbf{X} is used to initialise \mathbf{z} . The error between the selection on the generated mesh ($\mathbf{X}'_G = S \circ G(\mathbf{z})$) and the partial intraoperative point cloud \mathbf{P} is minimised optimising over $\mathbf{z}, \mathbf{R}, \mathbf{t}$, which control the shape of new generated meshes, the rotation of \mathbf{P} and its translation respectively.

2.1 Shape Generator

A 3D mesh can be represented as a graph $\mathcal{M} = \{\mathbf{X}, \varepsilon\}$, where $\mathbf{X} \in \mathbb{R}^{N \times 3}$ is its vertex embedding and $\varepsilon \in \mathbb{N}^{\varepsilon \times 2}$ is the edge connectivity that defines its topology. Traditional convolutional operators, well suited for grid data such as images and voxelizations, are generally incompatible with the non-Euclidean domain of graphs. Following [11], we chose to build our generative model with the Feature-Steered graph convolutions defined in [18]. This operator dynamically assigns filter weights to k-ring neighbourhoods according to the features learned by the network. In particular, given a generic feature vector field where for each vertex i we have a feature vector \mathbf{x}_i , we can define the output of the convolutional operator as

$$\mathbf{y}_i = \mathbf{b} + \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} \sum_{m=1}^M q_m(\mathbf{x}_i, \mathbf{x}_j) \mathbf{W}_m \mathbf{x}_j \quad (1)$$

where \mathbf{b} is a learnable bias, $q_m(\mathbf{x}_i, \mathbf{x}_j)$ is a translation-invariant assignment operator that, using a soft-max over a linear transformation of the local feature vectors, learns to dynamically assign \mathbf{x}_j to the m -th learnable weight matrix \mathbf{W}_m , and \mathcal{N}_i is the neighbour of the i -th vertex with cardinality $|\mathcal{N}_i|$.

Every VAE is made of an encoder-decoder pair, where the decoder is used as a generative model and is usually referred to as generator. Following this convention, we define our architecture as a pair $\{E, G\}$. Let \mathcal{X}_c be the vertex embedding

domain of complete shapes and \mathcal{Z} the latent distribution domain. Then, the two networks are defined as two non-linear functions such that $E : \mathcal{X}_c \rightarrow \mathcal{Z}$ and $G : \mathcal{Z} \rightarrow \mathcal{X}_c$. With $\mathbf{X} \in \mathcal{X}_c$ and $\mathbf{z} \sim \mathcal{Z}$, the generator is described by the likelihood $p(\mathbf{X}|\mathbf{z})$ while the encoder is defined as a variational distribution $q(\mathbf{z}|\mathbf{X})$ that approximates the intractable model posterior distribution. Both E and G are parametrised by neural networks whose building blocks are the Feature-Steered graph convolutions. During training, a reconstruction loss (\mathcal{L}_{recon}) encourages the output of the VAE to be as close as possible to its input and a regularisation term (\mathcal{L}_{KL}) pushes the variational distribution towards the prior distribution $p(\mathbf{z})$, which is defined as a standard spherical Gaussian distribution. While we set \mathcal{L}_{KL} to be the Kullback–Leibler (KL) divergence, we define \mathcal{L}_{recon} as a vertex-density-weighted mean-squared-errors loss. Let \mathbf{x}_i be the i -th position vector (i.e. a feature vector of size 3) and \mathbf{x}'_i its corresponding point in $\mathbf{X}' = G(E(\mathbf{X}))$. We have:

$$\mathcal{L}_{recon} = \frac{1}{N} \sum_{i=0}^{N-1} \gamma \|\mathbf{x}'_i - \mathbf{x}_i\|_2^2 \quad \text{with } \gamma \propto \frac{1}{N_i} \sum_{j \in \mathcal{N}_i} \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \quad (2)$$

where γ is a vertex-wise weight that increases the contribution of the errors in low vertex-density regions, thus preventing the generated mesh from fitting only densely sampled areas. The total loss is then computed as linear combination of the two terms: $\mathcal{L}_{tot} = \mathcal{L}_{recon} + \alpha \mathcal{L}_{KL}$.

Data Preparation Though the chosen graph convolution was proven effective also on datasets with different graph topologies [18], we decided to remesh our data in order to have the same topology and known point correspondences across all the preoperative meshes and all the generated shapes. This accelerates and eases the training procedure, allowing us to define a simple and computationally-efficient loss function (Eq. 2). In addition, thanks to the consistent vertex indexing it is possible to easily perform the initial manual selection described in Sec. 2.2.

In order to consistently remesh our dataset, we run an optimisation procedure that iteratively deforms an ico-sphere with a predefined topology and fixed number of vertices. Following [19], the loss function is given by $\mathcal{L}_{remesh} = \lambda_0 \mathcal{L}_{Ch} + \lambda_1 \mathcal{L}_n + \lambda_2 \mathcal{L}_L + \lambda_3 \mathcal{L}_E$. \mathcal{L}_{Ch} is the Chamfer distance that averages the distances between each point in a mesh and the closest point in the other mesh, and vice versa. \mathcal{L}_n is the normal loss that requires the edge between a vertex and its neighbours to be perpendicular to the normal of the closest vertex in the ground truth. \mathcal{L}_L is the mesh Laplacian regularisation loss that avoids self-intersections, acting as a surface smoothness term. \mathcal{L}_e is an edge regularisation that reduces flying vertices by penalising long edges. $\lambda_{0,1,2,3}$ are the weights of each loss term.

Spectral Augmentation The small size of our dataset makes it difficult to train a generative model that can generalise to new shapes. Simple shape aug-

mentation techniques such as random rotations, translations and scalings can be used to augment the dataset, but shapes are not deformed and the performance gain is therefore limited. Instead of attempting an augmentation in the spatial domain we propose a data augmentation technique that operates in the frequency domain, which is a known concept in the literature [17]. However, we simplify and randomise the spectral deformation making it suitable for data augmentation. We thus compute the un-normalized graph Laplacian operator $\mathbf{L} = \mathbf{D} - \mathbf{A}$, where $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix and $\mathbf{D} \in \mathbb{R}^{N \times N}$ is the diagonal degree matrix with $D_{ii} = \sum_j A_{ij}$. Computing the eigenvalue decomposition of the Laplacian, $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$, we obtain a set of orthonormal eigenvectors (columns of \mathbf{U}) which are the Fourier bases of the mesh, and a series of eigenvalues (diagonal values of $\mathbf{\Lambda}$) that are its frequencies. The Fourier transform of the vertices can be computed as $\tilde{\mathbf{X}} = \mathbf{U}^T \mathbf{X}$ and the inverse Fourier transform as $\mathbf{X} = \mathbf{U}\tilde{\mathbf{X}}$ [5].

Using these operators, each mesh is transformed into its spectral domain, perturbed, and transformed back to the spatial domain. Hence, the spectral augmented mesh \mathbf{X}^\dagger is computed as $\mathbf{X}^\dagger = \mathbf{U}\boldsymbol{\xi}\mathbf{U}^T\mathbf{X}$, where $\boldsymbol{\xi}$ is a vector that randomly perturbs four mesh frequencies. In particular, the first frequency is never modified because, playing the role of a direct current component [2], it would not deform the shape. One of the following three frequencies, responsible for low frequency variations similar to scalings along the three major axes of the mesh, is arbitrarily perturbed. The remaining three perturbations are applied to randomly selected higher frequencies with the effect of affecting the fine details of the shape.

It is worth noting that the remeshed data share the same topology, thus the set of orthonormal eigenvectors used to compute the direct and inverse Fourier transforms can be computed one time and then used for every mesh.

2.2 Shape Completion

This section illustrates how a complete shape is obtained from a partial intraoperative point cloud $\mathbf{P} \in \mathbb{R}^{P \times 3}$. In contrast to [11], we do not have known (or easily computable) point correspondences between intraoperative point clouds and the generated meshes. Therefore, we relax this assumption at the expense of the introduction of a manual step in the procedure. In fact, the surgeon is asked to roughly select from the preoperative 3D model $\mathbf{X} \in \mathcal{X}_c$ a region of interest that corresponds to the visible surface in the intraoperative image. To reduce computational time and increase robustness against the errors in manual region selection and varying vertex density in the region, we sample the selected vertices with an iterative farthest point sampling [15], obtaining a selection operator S that gives sparser and uniformly sampled vertices. Since mesh topology consistency is guaranteed by construction, the selected vertices will always have the same indexing for each mesh $\mathbf{X}' \in \mathcal{X}_c$ generated with the model discussed in Sec. 2.1. The shape completion problem is thus formulated as finding the best latent variable \mathbf{z}^* that generates a complete shape $\mathbf{X}^{/*}$ plausibly fitting \mathbf{P} . Given

$\mathbf{X}'_S = S \circ \mathbf{X}' = S \circ G(\mathbf{z}) \subset \mathbf{X}'$ the subset of selected and sampled vertices from a generated shape, we optimise

$$\min_{\mathbf{z}, \mathbf{R}, \mathbf{t}} \mathcal{L}_{Ch}(S \circ G(\mathbf{z}), \mathbf{R}\mathbf{P} + \mathbf{t}). \quad (3)$$

It is worth mentioning that not having point correspondences between \mathbf{P} and \mathbf{X}'_S we cannot compute the rotation \mathbf{R} and translation \mathbf{t} in a closed form solution as in [11]. Therefore, they are iteratively updated alongside \mathbf{z} in the same optimisation procedure. The gradient of the loss in Eq. 3 directly influences \mathbf{R} and \mathbf{t} , but it needs to be back-propagated through the generator network G , without updating the network’s weights, to update \mathbf{z} . The completion procedure is initialised by centering \mathbf{P} and \mathbf{X}'_S , and by setting an initial $\mathbf{z} = \mathbf{z}_0 = E(\mathbf{X})$, thus using as prior the latent representation of the preoperative mesh. The initialisation \mathbf{z}_0 can be further refined to \mathbf{z}_0^* by running a few iterations of a second optimisation $\mathbf{z}_0^* \leftarrow \arg \min_{\mathbf{z}_0} (\max_i \|\mathbf{x}_i - \mathbf{x}'_i\|_2^2)$. Finally, by adding to the latent initialisation a small Gaussian noise $\boldsymbol{\eta} \sim \mathcal{G}(\mathbf{0}, \boldsymbol{\Sigma})$ with $\Sigma_{ii} \ll I_{ii}$, we can generate multiple complete shapes conditioned on the preoperative data and that plausibly fit the intraoperative point cloud \mathbf{P} .

3 Results

Our dataset consists of 50 meshes of livers which were segmented and reconstructed from preoperative CT scans of different patients. The segmentation and initial mesh generation was performed by Visible Patient. 45 meshes were used to train the VAE, while the remaining 5 meshes were used as a test set to evaluate the network, data augmentation, and shape completion. Given the limited size of the dataset, to not bias results toward the test set, hyperparameters were tuned on the training set. The study was approved by the local research ethics committee (Ref: 14/LO/1264) and written consent obtained from each patient.

The remeshing was performed by deforming an ico-sphere with 2564 vertices. For this, and all the other optimisations described in this paper, we used the ADAM optimiser [9]. We remeshed every model with 500 iterations at a fixed learning rate of $lr = 5e^{-3}$. The weights of the loss function \mathcal{L}_{remesh} were $\lambda_0 = 5$, $\lambda_1 = 0.2$, $\lambda_2 = 0.3$, and $\lambda_3 = 15$.

The VAE was built using $M = 8$ weight matrices, batch size 20 and latent size 128. LeakyReLU and batch normalisation were used after every layer. The network was implemented in PyTorch Geometric [7] and trained for 200 epochs with $lr = 1e^{-3}$ and a KL divergence weight $\alpha = 1e^{-6}$. The training was performed on an NVIDIA Quadro P5000 and took approximately 9 hours.

We evaluated the reconstruction performance of the VAE with and without data augmentation while fixing the number of iterations. Applying the spectral augmentation (Fig. 2A) we generated 100 new meshes for each model in the training set, thus obtaining 4500 models. An additional online data augmentation composed of random rotations, scalings, and translations was applied. We

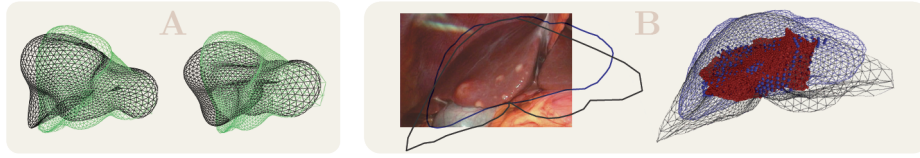


Fig. 2. Augmentation and Qualitative Results *A*: effects of the spectral augmentation where a real liver (green) is subject to two different random augmentation (black). *B*: laparoscopic image and comparison between the proposed shape completion (blue) and the Go-ICP registration (black). The intraoperative point cloud is shown in red and the selected point in blue. The contours of the silhouettes are overlaid also on the laparoscopic image.

obtain a mean-squared testing error of (0.28 ± 0.04) mm when both augmentations are applied, of (0.45 ± 0.18) mm with the online augmentation alone, of (0.50 ± 0.03) mm with the spectral augmentation only, and of (0.92 ± 0.22) mm without any augmentation. We then evaluated the computational cost of the spectral augmentation, finding that when it is performed by computing the Fourier operators for each mesh it takes 0.4532 ± 0.0568 seconds per mesh, while, when the operators are precomputed (Sec. 2.1) the computational time is reduced by one order of magnitude to 0.0487 ± 0.0092 seconds per mesh.

Given the lack of intraoperative 3D ground truths for registration in laparoscopic liver surgery, the evaluation of our method on real data is purely qualitative. In a real operative scenario the computation of a dense and reliable point cloud is still a major challenge. To obtain \mathbf{P} from rectified images of a calibrated laparoscope we used an off-the-shelf depth reconstruction network [4]. Given the predicted depth map and a manual segmentation of the liver, we first compute \mathbf{P} and then estimate the complete shape \mathbf{X}^* (Fig. 2B).

The quantitative assessment of the shape completion is performed on synthetic data. The five meshes in the test set were manually deformed, trying to reproduce intraoperative liver deformations similar to those expected in a laparoscopic procedure and characterised in [8]. To obtain intraoperative partial point clouds, the surface of the deformed models was sampled with vertex selections on three regions: entire front surface \mathbf{P}_F , left lobe \mathbf{P}_L , and right lobe \mathbf{P}_R . Each deformed model is considered the intraoperative ground truth \mathbf{X}_{GT} that we want to infer given a partial point cloud. To maintain a higher \mathbf{P} density, \mathbf{X}_{GT} was not remeshed. Eq. 3 is optimised for 100 iterations using ADAM with a different learning rate for each term. To encourage the optimisation over \mathbf{z} and thus the generation of more diverse, progressively deformed meshes \mathbf{X}' , we set $lr_z = 5e^{-2}$. The learning rates for \mathbf{R} and \mathbf{t} are empirically set to $lr_R = 1e^{-2}$ and $lr_t = 5e^{-5}$. Rotations are regressed faster because the two point clouds were initially centred. In case \mathbf{z}_0 is further refined to \mathbf{z}_0^* , the same learning rate $lr_z = 5e^{-2}$ is used for 20 iterations. In case multiple complete shape proposals are desired, Eq. 3 can be generalised to process batches with a refined initialisation perturbed by $\boldsymbol{\eta}$ with $\Sigma_{ii} = \frac{1}{10}$.

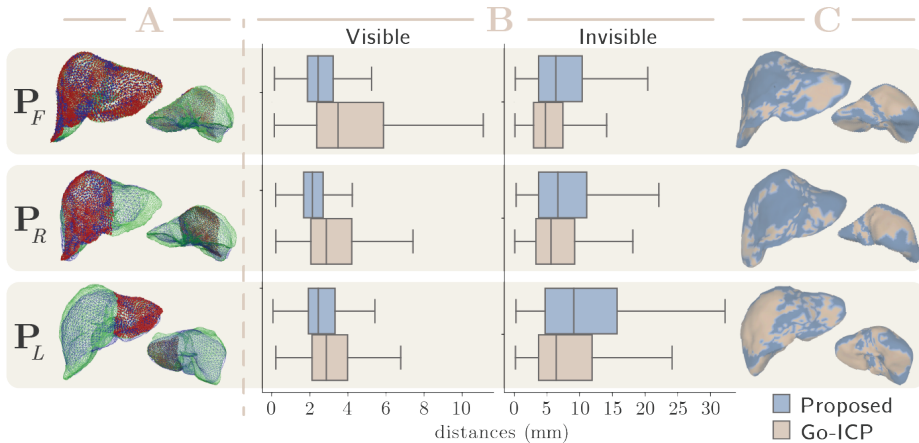


Fig. 3. Quantitative Results. Rows: results for intraoperative point clouds of front surface \mathbf{P}_F , right lobe \mathbf{P}_R , and left lobe \mathbf{P}_L . Columns: A) front and back view of intraoperative point cloud \mathbf{P} (red), intraoperative ground truth \mathbf{X}_{GT} (green), and prediction \mathbf{X}^* (blue). B) Comparison of vertex-wise distances (mm) for the proposed method (blue) and Go-ICP (pink). We compute errors separately in the visible and invisible parts of the liver in the camera’s field of view. C) Front and back view of \mathbf{X}^* . Colours represent the algorithm with a smaller vertex error.

The shape completion was evaluated for each partial shape without η perturbation. Since the procedure requires a manual step currently performed with a lasso selection that might affect results, we repeated the evaluation 3 times, for a total of 45 experiments. Selections could be refined and took approximately one minute each. We compared our method with the rigid registration using Go-ICP [12] which has comparable computational time to ours. The lack of point correspondences between \mathbf{X}_{GT} and \mathbf{X}^* does not allow us to evaluate our method using mean-squared errors. Therefore, we define a variation of \mathcal{L}_{Ch} that allows us to compute vertex-wise errors on \mathbf{X}^* . For each vertex of one mesh we compute the distance to the closest point on the other mesh. All the distances are assigned to the vertices of \mathbf{X}^* from which they were computed and are locally averaged. Results are reported in Fig. 3.

4 Conclusion

While this work is about shape completion, we believe it could become an alternative to registration or provide a better initialisation for such algorithms. From Fig. 2 we notice that the proposed method seems to fit better the point cloud especially on the left lobe. The lack of a ground truth makes impossible to draw further conclusions from this result. However, observing Fig. 3 (columns B and C), we can conclude that our method outperforms Go-ICP in visible areas, and, despite performing worse in invisible areas, it predicts a realistic looking

model of the liver. Of particular importance is the improvement over visible areas, because these regions are the only ones in the narrow field of interest of the surgeon, where an accurate deformation is required. Since the manual selection of the visible area on the preoperative model affects the quality of the results, as future work not only we aim at quantifying the uncertainty involved in the manual interaction, but also at avoiding it by predicting point correspondences between partial and complete shapes. We also believe that the use of biomechanical constraints for deformation could reduce errors in invisible areas. In fact, the unconstrained deformations operated by our method through the optimisation of \mathbf{z} , despite generating plausible livers fitting the partial intraoperative point cloud, often downscale invisible areas. Even though our method can propose multiple solutions (Sec. 2.2), identifying the correct complete shape is essential to improve our method and outperform Go-ICP everywhere. Thus, we shall also research the introduction of biomechanical constraints while avoiding any patient specific training or simulation. Finally, we are planning to incorporate the liver’s internal structure in our method in order to overlay them on a laparoscopic video during surgery.

Acknowledgments This work was supported by the the Wellcome Trust/EPSRC [203145Z/16/Z], and Wellcome Trust / Department of Health [HICF-T4-317]

References

1. Abdi, A.H., Pesteie, M., Prisman, E., Abolmaesumi, P., Fels, S.: Variational Shape Completion for Virtual Planning of Jaw Reconstructive Surgery. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 227–235. Springer (2019)
2. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* **34**(4), 18–42 (2017)
3. Brunet, J.N., Mendizabal, A., Petit, A., Golse, N., Vibert, E., Cotin, S.: Physics-Based Deep Neural Network for Augmented Reality during Liver Surgery. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 137–145. Springer (2019)
4. Chang, J.R., Chen, Y.S.: Pyramid Stereo Matching Network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5410–5418 (2018)
5. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In: Advances in Neural Information Processing Systems (NIPS). pp. 3844–3852 (2016)
6. Espinel, Y., Ozgür, E., Calvet, L., Le Roy, B., Buc, E., Bartoli, A.: Combining Visual Cues and Interactions for 3D-2D Registration in Liver Laparoscopy. *Annals of Biomedical Engineering* (2020)
7. Fey, M., Lenssen, J.E.: Fast Graph Representation Learning with PyTorch Geometric. In: ICLR Workshop on Representation Learning on Graphs and Manifolds (2019)

8. Heiselman, J.S., Clements, L.W., Collins, J.A., Weis, J.A., Simpson, A.L., Geevarghese, S.K., Kingham, T.P., Jarnagin, W.R., Miga, M.I.: Characterization and Correction of Intraoperative Soft Tissue Deformation in Image-Guided Laparoscopic Liver Surgery. *Journal of Medical Imaging* **5**(2), 021203 (2017)
9. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980 (2014)
10. Koo, B., Özgür, E., Le Roy, B., Buc, E., Bartoli, A.: Deformable Registration of a Preoperative 3D Liver Volume to a Laparoscopy Image using Contour and Shading Cues. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. pp. 326–334. Springer (2017)
11. Litany, O., Bronstein, A., Bronstein, M., Makadia, A.: Deformable Shape Completion with Graph Convolutional Autoencoders. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1886–1895 (2018)
12. Luo, H., Yin, D., Zhang, S., Xiao, D., He, B., Meng, F., Zhang, Y., Cai, W., He, S., Zhang, W., et al.: Augmented Reality Navigation for Liver Resection with a Stereoscopic Laparoscope. *Computer Methods and Programs in Biomedicine* p. 105099 (2019)
13. Özgür, E., Koo, B., Le Roy, B., Buc, E., Bartoli, A.: Preoperative Liver Registration for Augmented Monocular Laparoscopy using Backward–Forward Biomechanical Simulation. *International Journal of Computer Assisted Radiology and Surgery (IJCARS)* **13**(10), 1629–1640 (2018)
14. Prevost, G.A., Eigl, B., Paolucci, I., Rudolph, T., Peterhans, M., Weber, S., Beldi, G., Candinas, D., Lachenmayer, A.: Efficiency, Accuracy and Clinical Applicability of a New Image-Guided Surgery System in 3D Laparoscopic Liver Surgery. *Journal of Gastrointestinal Surgery* pp. 1–8 (2019)
15. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In: *Advances in Neural Information Processing Systems (NIPS)*. pp. 5099–5108 (2017)
16. Robu, M.R., Ramalhinho, J., Thompson, S., Gurusamy, K., Davidson, B., Hawkes, D., Stoyanov, D., Clarkson, M.J.: Global Rigid Registration of CT to Video in Laparoscopic Liver Surgery. *International Journal of Computer Assisted Radiology and Surgery (IJCARS)* **13**(6), 947–956 (2018)
17. Rong, G., Cao, Y., Guo, X.: Spectral Mesh Deformation. *The Visual Computer* **24**(7-9), 787–796 (2008)
18. Verma, N., Boyer, E., Verbeek, J.: FeaStNet: Feature-Steered Graph Convolutions for 3D Shape Analysis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2598–2606 (2018)
19. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 52–67 (2018)
20. Yang, J., Li, H., Campbell, D., Jia, Y.: Go-ICP: A Globally Optimal Solution to 3D ICP Point-Set Registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **38**(11), 2241–2254 (2015)