



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

CardioSpectrum: Comprehensive Myocardium Motion Analysis with 3D Deep Learning and Geometric Insights

Shahar Zuler¹, Shai Tejman-Yarden^{2,3,4}, and Dan Raviv¹

¹ Tel Aviv University, Tel Aviv, Israel

shahar.zuler@gmail.com darav@tauex.tau.ac.il

² Faculty of Medicine, Tel Aviv University, Tel Aviv, Israel.

³The Edmond J. Safra International Congenital Heart Center, Sheba Medical Center, Ramat Gan, Israel.

⁴The Engineering Medical Research Laboratory, Sheba Medical Center, Ramat Gan, Israel.

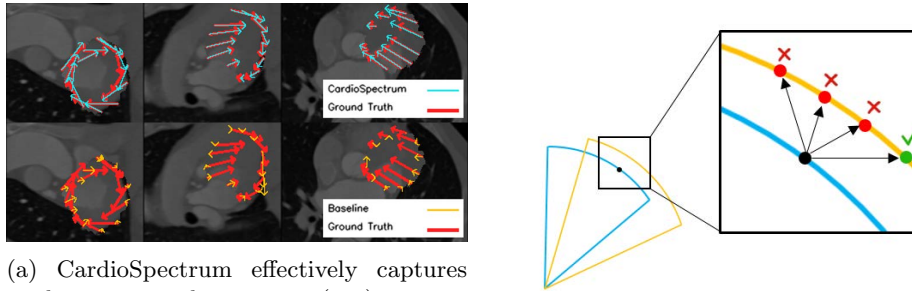
tegmnya@gmail.com

Abstract. The ability to map left ventricle (LV) myocardial motion using computed tomography angiography (CTA) is essential to diagnosing cardiovascular conditions and guiding interventional procedures. Due to their inherent locality, conventional neural networks typically have difficulty predicting subtle tangential movements, which considerably lessens the level of precision at which myocardium three-dimensional (3D) mapping can be performed. Using 3D optical flow techniques and Functional Maps (FMs), we present a comprehensive approach to address this problem. FMs are known for their capacity to capture global geometric features, thus providing a fuller understanding of 3D geometry. As an alternative to traditional segmentation-based priors, we employ surface-based two-dimensional (2D) constraints derived from spectral correspondence methods. Our 3D deep learning architecture, based on the ARFlow model, is optimized to handle complex 3D motion analysis tasks. By incorporating FMs, we can capture the subtle tangential movements of the myocardium surface precisely, hence significantly improving the accuracy of 3D mapping of the myocardium. The experimental results confirm the effectiveness of this method in enhancing myocardium motion analysis. This approach can contribute to improving cardiovascular diagnosis and treatment. Our code and additional resources are available at: <https://shaharzuler.github.io/CardioSpectrumPage/>

Keywords: Optical flow · Scene flow · Deep neural networks · Cardiac cycle.

1 Introduction

The accurate mapping of LV myocardium movement is a critical feature of cardiac CT analysis. Understanding the 3D motion of the chambers and vessels of



(a) CardioSpectrum effectively captures cardiac tangential movement (top), outperforming traditional methods that predominantly focus on local features, which often lead to inadequate results due to the aperture problem (bottom).

(b) The aperture problem: Global comprehension of geometry is essential to effectively capture tangential movements.

Fig. 1

the heart is essential to diagnosing cardiovascular conditions and guiding medical interventions. However, this task poses several challenges, including the need to capture the precise tangential movements of the myocardium.

Traditional neural networks (NNs) have shown promise in analyzing cardiac movement from imaging data [4] through the use of designated architectures [14][16][20][26], optimization schemes [3], priors and constraints [27][6][2][15]. However, in addition to the scarcity of annotated high-resolution imagery, they also often suffer from the locality problem, since they struggle to predict tangential movements accurately (as depicted in figure 1a). While they excel at minimizing photometric energy to capture radial movements of the myocardium, their inherent locality limits their ability to comprehensively model subtle tangential movements. This limitation is referred to as **the aperture problem**, as illustrated in figure 1b.

To address the challenge of capturing both radial and tangential components of myocardium motion, we propose a holistic approach that integrates DL techniques with spectral methods, and specifically FMs [17]. FMs, initially developed for non-rigid shape matching, operate by projecting functions from different shapes onto the eigenfunctions of the Laplacian Beltrami Operator (LBO) [22]. This LBO projection effectively captures the intrinsic geometry of the shapes, reducing the dimensionality of the problem and improving solution accuracy. By considering the LV surface as a manifold, we leverage FMs' global perception for understanding myocardium motion.

Our comprehensive approach as detailed in this paper includes a 3D DL architecture and the incorporation of spectral methods. We leverage 2D constraints derived from spectral correspondence methods to significantly enhance the accuracy of surface mapping. These constraints provide a more robust foundation compared to traditional segmentation priors by precisely specifying the pixel correspondences between timesteps.

The following sections detail our approach, and provide insights into each component that contributes to the accuracy and inclusiveness of our myocardium 3D mapping technique. We believe that this approach represents a significant step forward in cardiac CT analysis, by addressing the challenges associated with capturing both radial and tangential myocardium movements accurately.

2 Motivation

Understanding cardiovascular health extends beyond the study of electrical signals and anatomical structures. It hinges on comprehending the intricate dynamics of the myocardium—the muscular heart wall responsible for contractions—and its movement throughout the cardiac cycle. Accurate measurement of myocardium behavior is not just an academic aspiration, but rather is pivotal to patient well-being and the advancement of cardiovascular research.

The heart operates in a rhythmic cycle, where each beat relies on the orchestrated movements of the myocardium. This muscular layer governs the heart’s efficiency in pumping blood. Perturbations in myocardial movement can manifest as diverse cardiovascular diseases, ranging from heart attacks to cardiomyopathies.

Despite advances in technology, inaccuracies in myocardium movement extraction persist. These inaccuracies can result in incorrect diagnoses, ineffective treatment plans, and slow progress in understanding cardiac pathophysiology.

In the field of medical imaging, experts often work with 2D slices of the heart, which requires mental integration to understand its 3D structure and movements. This process can be challenging. By contrast, computer algorithms excel in handling 3D or 3D+time (3D+t) data, and can provide a holistic view of the heart’s structure and motion. When analyzing a 3D+t cardiac CT scan, for example, these algorithms generate a high-resolution 3D movement field that can be translated into deformation, velocity, and strain. These data can be represented visually, thus enhancing interpretability for medical professionals.

Accurate myocardium movement extraction and the application of 3D imaging and computer vision do not only constitute medical advances but are also holistic endeavors that can revolutionize cardiovascular health. These pursuits are not only likely to generate scientific knowhow but also tangible, life-improving benefits for individuals worldwide.

3 The Aperture Problem

In addressing 3D optical flow within a localized window, our goal is to minimize:

$$\|I_x u + I_y v + I_z w + I_t\| \quad (1)$$

where u , v , and w , the optimized parameters, correspond to the optical flow vector elements indicating motion along the x , y , and z axes, whereas I_x , I_y , I_z and I_t denote the spatial and temporal derivatives of pixel intensities.

Consider a 3D image characterized by pixel intensities, where we assume that these intensities define a level set within a small 3D window. If the level set undergoes motion, the motion vector can be decomposed into components along different axes. These include a component along the axis \mathcal{N} , normal to the level set, and components along axes t_1 and t_2 , lying in a plane perpendicular to \mathcal{N} and mutually perpendicular. The objective term to minimize in this scenario can be expressed as:

$$\|I_{\mathcal{N}}f_{\mathcal{N}} + I_{t_1}f_{t_1} + I_{t_2}f_{t_2} + I_t\| \quad (2)$$

As I_{t_1} and I_{t_2} vanish due to the motion within the level set plane, the optimization problem simplifies to:

$$\|I_{\mathcal{N}}f_{\mathcal{N}} + I_t\| \quad (3)$$

This reduction implies that the solution to the 3D optical flow problem only considers the radial component, $f_{\mathcal{N}}$, and disregards the tangential components, f_{t_1} and f_{t_2} . In essence, it exclusively addresses motion along the axis normal to the level set, leaving the tangential components unconstrained. This is called the aperture problem.

Global approaches like Horn and Schunk [7] and Convolutional Neural Networks (CNNs) have been suggested but remain curtailed by the aperture problem. Despite CNNs' broader perspective, their local context makes them prone to this issue. In cardiac CT imaging, this limitation is a major impediment because CNNs may fail to accurately estimate tangential motion on the dynamic LV myocardium surface due to their localized receptive fields.

Hypothesis: Introducing constraints on the level set may enhance optical flow evaluation around the level set, thus offering a potential approach to addressing the aperture problem.

4 Method

To address the limitations of a naive 3D optical flow approach in predicting the tangential component, we first present the global geometric solution. We then integrate the results of a global geometric solution as a constraint into the 3D voxel-based solution. This fusion yields a network that can simultaneously produce a dense flow field while correctly predicting the tangential components. Here, we define the LV surface as encompassing the outer boundary of the myocardium. In practice, this boundary constitutes the surface surrounding the union of the LV myocardium and the LV blood cavity.

4.1 Global Constraints from Functional Maps

To extract LV surface constraints, we use the ZoomOut method [13], which iteratively evaluates the FM matrix, transitioning between spatial and spectral domains. This process employs a progressive zooming-out strategy, gradually




Fig. 2: CardioSpectrum Architecture: The NN analyzes 3D image pairs of cardiac cycle timesteps, incorporating 2D constraints from ZoomOut. These constraints, derived from segmentations converted into meshes, result in a 3D optical flow.

expanding the dimensionality of the functional mapping and refining point correspondences at each step, establishing correspondences between triangle mesh representations of two timesteps, derived from their smoothed LV segmentations.

We map these surface constraints to a 3D grid by calculating an array of 3D displacements, \mathbf{D} , by subtracting corresponding vertex locations on the meshes. Using rounded vertex locations as indices, we fill a 3D grid-based flow field with sparse flow values from \mathbf{D} . Constraints are then interpolated into each voxel of the original unsmoothed segmentation map hull, providing the final constraints for the 3D optical flow network.

4.2 3D Optical Flow Network

Our 3D optical flow network is built upon the ARFlow [10] training methodology and recommended model architecture, which employs a PWC-Net-based model [23]. Adapted for 3D imaging as outlined in [9], our model integrates essential components including a feature pyramid extractor, cost volume, a flow estimator, and a refinement context network. The core objective of our 3D optical flow network is expressed in the following equation:

$$\mathcal{L} = \lambda_{\mathcal{L}_1} \cdot \mathcal{L}_1 + \lambda_{ssim} \cdot \mathcal{L}_{ssim} + \lambda_{constraints} \cdot \mathcal{L}_{constraints} \quad (4)$$

\mathcal{L}_1 and \mathcal{L}_{ssim} Losses: The \mathcal{L}_1 loss captures the pixel-wise intensity disparities, ensuring accurate matching between the target and warped source images. \mathcal{L}_{ssim} [24] focuses on preserving structural details and perceptual quality during optimization.

Constraints Loss: Our network integrates the constraints obtained from the FM pipeline operating on the LV surface. This constraints loss acts as a penalty term that incentivizes the predicted flows at constraint locations to match closely with the provided constraints:

$$\mathcal{L}_{constraints} = \sum_{i,j,k \in \text{segmentation hull}} \|\hat{\text{flow}}(i,j,k) - \text{constraint}(i,j,k)\|_1 \quad (5)$$

Here, $\hat{\text{flow}}(i, j, k)$ and $\text{constraint}(i, j, k)$ are the predicted flow and the provided constraint values at voxel locations (i, j, k) within the segmentation hull.

Thus overall, our network is trained with a balanced combination of \mathcal{L}_1 , \mathcal{L}_{ssim} , and $\mathcal{L}_{constraints}$, each weighted by the corresponding hyperparameters $\lambda_{\mathcal{L}_1}$, λ_{ssim} , and $\lambda_{constraints}$. These λ values form vectors, where each element of the vectors corresponds to a layer in the feature pyramid of our network.

Optimization Our optimization scheme involves two stages. First, we pre-train our model on a large motion-based dataset from raw cardiac 4DCT scans. Next, we fine-tune the model for each individual pair using \mathcal{L}_1 and \mathcal{L}_{ssim} , followed by further optimization with $\mathcal{L}_{constraints}$.

4.3 Accuracy and Correspondence Measures

To assess the accuracy of the predicted flow values, we compared them to the ground truth flow values using the following metrics:

Mean End-Point Error (mEPE) quantifies the distance between the end-points obtained from the predicted flow and the ground truth flow. It is defined as follows:

$$mEPE(\hat{\text{flow}}, \text{flow}) = \frac{1}{N} \sum_{i,j,k} \|\hat{\text{flow}}(i, j, k) - \text{flow}(i, j, k)\|_2 \quad (6)$$

where N denotes the number of voxels, $\hat{\text{flow}}$ represents the predicted flow, flow denotes the ground truth flow, and i, j, k are indices corresponding to voxels within a defined region.

We calculate the mEPE using flow values exclusively from voxels contained within the segmentation of the myocardium. In addition, we compute the mEPE for the radial, longitudinal, and circumferential flow components separately, as illustrated in the supplementary material. We determine the longitudinal axis using the principal component of the LV segmentation map.

For the voxels contained in the LV segmentation hull, we separately compute the mEPE for the *locally-tangential* and *locally-radial* flow components, as illustrated in the supplementary material.

Mean Angular Error Quantifies the dissimilarity between the predicted and the ground truth flow directions at each point in our analysis by computing the mean angle between the predicted and the true flow vectors.

$$\bar{E}_{angular}(\hat{\text{flow}}, \text{flow}) = \frac{1}{N} \sum_{i,j,k} \left| \cos^{-1} \left(\frac{\hat{\text{flow}}(i, j, k) \cdot \text{flow}(i, j, k)}{\|\hat{\text{flow}}(i, j, k)\|_2 \|\text{flow}(i, j, k)\|_2} \right) \right| \quad (7)$$

Here, \cdot denotes the dot product between the two vectors. The other notations are identical to those described above.

4.4 Implementation details

We optimized our PyTorch [18] network with seven feature pyramid layers, using the Adam optimizer [8] with a weight decay of 10^{-6} , $\alpha = 10^{-4}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-7}$. Training involved 2500 unconstrained iterations and an additional 2500 iterations with constraints. Losses for pyramid layers were weighted $\lambda_{ph} = 3$, $\lambda_{sim} = 1$ and $\lambda_{constraints} = 100$. For ZoomOut preprocessing, the segmentation underwent morphological closing (kernel size 5) and opening (kernel size 7), followed by conversion to a triangle mesh using marching cubes [11] from [19]. Mesh smoothing used spectral filtering via the first 120 eigenfunctions of the mesh’s LBOs [22]. ZoomOut initiated with 3 eigenvectors and 3 eigenvalues on each mesh, incorporating 50 Wave Kernel Signature (WKS) descriptors [1] over 130 iterations. The entire optimization on a single Quadro RTX 8000 GPU took approx. 240 minutes and required 1900 MB of GPU memory.

5 Experiments and Results

Our experimental evaluations consisted of in-depth assessments of the efficacy of our method through optimization on a synthetic dataset enriched with ground truth annotations for LV 3D optical flow. The dataset consisted of 300 unique scan pairs of systole and diastole, each exhibiting torsion ranging equally from 0 to approximately 35 degrees and distributed across various torsion centers. The data distribution and deformation process are described in [29]. The scans were taken from various sources [5][28][12][25][21].

We compared the performance of our method against two baseline approaches. The first baseline, here termed "3D ARFlow", corresponds to the network before optimization with additional constraint flow, akin to [9]. The second baseline, akin to [15][2], involved 3D ARFlow with an additional anatomical loss that encouraged maximal overlap between the predicted and ground truth segmentations. Both baselines were trained for 2500 iterations.

We observed that optimal outcomes were not always achieved at the point of convergence in iterations. This is a recognized issue, which is typically addressed through the optimization of hyperparameters. In our study, we halted the process based on performance against a deliberately introduced deformation. Below, we highlight the main results.

5.1 Results

Subfigure 3a shows that as the torsion angle increased, the mEPE across the myocardium region in our method consistently exhibited a reduction in comparison to the baselines. In terms of the locally-radial and locally-tangential flow components within the segmentation hull (Subfigures 3b, 3c), our method demonstrated superior performance as the torsion angle increased, particularly in capturing the intricacies of locally-tangential flow patterns compared to the baseline models. Our model exhibited significant advantages in locally-tangential components,




Fig. 3: Comprehensive comparison of proposed model and two baselines in accurately capturing cardiac deformation across a range of torsion angles. The metrics evaluated included the mEPE over the myocardium volume, the mEPE in locally-radial and locally-tangential directions within the segmentation hull. As depicted, our model exhibited strengths that were particularly evident in 3a representing the overall mEPE, and pronounced advantages in locally-tangential components 3c, which aligns with our method’s goal of addressing the aperture problem, a challenge that was difficult to face by the two baseline models. Error bars represent the standard errors (SE). See Fig. 1 in supplementary material for the component-wise mEPE in radial, circumferential, and longitudinal directions and mean angular error.

and its performance in locally-radial aspects still remained competitive, with some variations observed in lower torsion ranges. This emphasizes the ability of our approach to effectively capture locally-tangential flow patterns within the anatomical regions of interest, by addressing the challenges associated with the aperture problem. Breaking down the mEPE into radial, circumferential, and longitudinal components provides additional insights. For detailed analyses on these components, as well as the mean angular error, see the supplementary material.

The overarching trend in our results thus indicates that our method excels with increasing torsion angles compared to the baseline models.

6 Discussion and Conclusion

Our comprehensive approach that integrates a 3D DL architecture with spectral methods makes a notable contribution to LV flow analysis. We conducted a wide range of experimental evaluations that show a reduction in mEPE across myocardial regions, particularly in scenarios with increased torsion complexity, thus highlighting the superiority of our approach over baseline methods.

One of the key strengths of our method is its ability to address locally-tangential flow components, thus effectively overcoming challenges associated with the aperture problem. The constraints loss, coupled with the integration of

2D constraints from FMs, contributes to greater accuracy in predicting optical flow and has strong implications for cardiac health and disease.

Acknowledgments. This work is partially funded by the Zimin Institute for Engineering Solutions Advancing Better Lives.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops) pp. 1626–1633 (2011), <https://api.semanticscholar.org/CorpusID:17847588>
2. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
3. Dou, H., Bi, N., Han, L., Huang, Y., Mann, R., Yang, X., Ni, D., Ravikumar, N., Frangi, A.F., Huang, Y.: Gsmorph: Gradient surgery for cine-mri cardiac deformable registration. *arXiv preprint arXiv:2306.14687* (2023)
4. Duchateau, N., King, A.P., De Craene, M.: Machine learning approaches for myocardial motion and deformation analysis. *Frontiers in cardiovascular medicine* **6**, 190 (2020)
5. Gao, S., Zhou, H., Gao, Y., Zhuang, X.: Bayeseg: Bayesian modeling for medical image segmentation with interpretable generalizability. *arXiv preprint arXiv:2303.01710* (2023)
6. Hanania, E., Volovik, I., Barkat, L., Cohen, I., Freiman, M.: Pcmc-t1: Free-breathing myocardial t1 mapping with physically-constrained motion correction. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 226–235. Springer (2023)
7. Horn, B.K., Schunck, B.G.: Determining optical flow. *Artificial Intelligence* **17**(1), 185–203 (1981). [https://doi.org/https://doi.org/10.1016/0004-3702\(81\)90024-2](https://doi.org/https://doi.org/10.1016/0004-3702(81)90024-2), <https://www.sciencedirect.com/science/article/pii/0004370281900242>
8. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
9. Lifshitz, G., Raviv, D.: Unsupervised volumetric displacement fields using cost function unrolling. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 153–160. Springer (2021)
10. Liu, L., Zhang, J., He, R., Liu, Y., Wang, Y., Tai, Y., Luo, D., Wang, C., Li, J., Huang, F.: Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 6488–6497 (2020), <https://api.semanticscholar.org/CorpusID:214713881>
11. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: *Seminal graphics: pioneering efforts that shaped the field*, pp. 347–353 (1998)

12. Luo, X., Zhuang, X.: An n-dimensional information-theoretic framework for group-wise registration and deep combined computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
13. Melzi, S., Ren, J., Rodola, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: Spectral upsampling for efficient shape correspondence. arXiv preprint arXiv:1904.07865 (2019)
14. Meng, Q., Qin, C., Bai, W., Liu, T., de Marvao, A., O'Regan, D.P., Rueckert, D.: Mulvimotion: Shape-aware 3d myocardial motion tracking from multi-view cardiac mri. *IEEE Transactions on Medical Imaging* **41**, 1961–1974 (2022), <https://api.semanticscholar.org/CorpusID:247108049>
15. Morales, M.A., van den Boomen, M., Nguyen, C., Kalpathy-Cramer, J., Rosen, B.R., Stultz, C.M., Izquierdo-Garcia, D., Catana, C.: Deepstrain: A deep learning workflow for the automated characterization of cardiac mechanics. *bioRxiv* (2021). <https://doi.org/10.1101/2021.01.05.425266>, <https://www.biorxiv.org/content/early/2021/01/07/2021.01.05.425266>
16. Morales, M.A., Izquierdo-Garcia, D., Aganj, I., Kalpathy-Cramer, J., Rosen, B.R., Catana, C.: Implementation and validation of a three-dimensional cardiac motion estimation network. *Radiology: Artificial Intelligence* **1**(4), e180080 (2019)
17. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)* **31**(4), 1–11 (2012)
18. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
19. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
20. Qin, Y., Li, X.: Fsdiffreg: Feature-wise and score-wise diffusion-guided unsupervised deformable image registration for cardiac images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 655–665. Springer (2023)
21. Rosset, A., Spadola, L., Ratib, O.: Osirix: An open-source software for navigating in multidimensional dicom images. *Journal of digital imaging : the official journal of the Society for Computer Applications in Radiology* **17**, 205–16 (10 2004). <https://doi.org/10.1007/s10278-004-1014-6>
22. Rustamov, R.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. pp. 225–233 (01 2007). <https://doi.org/10.1145/1281991.1282022>
23. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8934–8943 (2018)
24. Wang, Z., Bovik, A., Sheikh, H., Simoncelli, E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004). <https://doi.org/10.1109/TIP.2003.819861>

25. Wu, F., Zhuang, X.: Minimizing estimated risks on unlabeled data: A new formulation for semi-supervised medical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(5), 6021–6036 (2022)
26. Xu, K., Huang, Q., Yang, X.: Importance weighted variational cardiac mri registration using transformer and implicit prior. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 581–591. Springer (2023)
27. Yan, W., Wang, Y., van der Geest, R.J., Tao, Q.: Cine mri analysis by deep learning of optical flow: Adding the temporal dimension. *Computers in biology and medicine* **111**, 103356 (2019)
28. Zhuang, X.: Multivariate mixture model for myocardial segmentation combining multi-source images. *IEEE transactions on pattern analysis and machine intelligence* **41**(12), 2933–2946 (2018)
29. Zuler, S., Raviv, D.: Synthetic data generation for 3d myocardium deformation analysis (2024)