


Anatomy-Aware Cardiac Motion Estimation

Pingjun Chen¹[0000-0003-0528-1713] ^{*}, Xiao Chen²[0000-0002-7147-7311], Eric Z. Chen²[0000-0001-5002-720X], Hanchao Yu³, Terrence Chen², and Shanhui Sun²[0000-0001-9841-8592] 

¹ University of Florida, Gainesville, FL 32611, USA

² United Imaging Intelligence, Cambridge, MA 02140, USA

³ University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Corresponding author: shanhui.sun@united-imaging.com

Abstract. Cardiac motion estimation is critical to the assessment of cardiac function. Myocardium feature tracking (FT) can directly estimate cardiac motion from cine MRI, which requires no special scanning procedure. However, current deep learning-based FT methods may result in unrealistic myocardium shapes since the learning is solely guided by image intensities without considering anatomy. On the other hand, motion estimation through learning is challenging because ground-truth motion fields are almost impossible to obtain. In this study, we propose a novel Anatomy-Aware Tracker (AATracker) for cardiac motion estimation that preserves anatomy by weak supervision. A convolutional variational autoencoder (VAE) is trained to encapsulate realistic myocardium shapes. A baseline dense motion tracker is trained to approximate the motion fields and then refined to estimate anatomy-aware motion fields under the weak supervision from the VAE. We evaluate the proposed method on long-axis cardiac cine MRI, which has more complex myocardium appearances and motions than short-axis. Compared with other methods, AATracker significantly improves the tracking performance and provides visually more realistic tracking results, demonstrating the effectiveness of the proposed weakly-supervision scheme in cardiac motion estimation.

Keywords: Anatomy aware · Motion estimation · Weak supervision.

1 Introduction

Accurate cardiac motion estimation plays a critical role in cardiac function assessment, such as myocardium strain, torsion, and dyssynchrony, which have been demonstrated as sensitive and early indicators of myocardial disorders [15]. Myocardial feature tracking (FT) can provide motion estimation from breath-hold 2D cine MRI, which is recommended by the American Heart Association (AHA) for clinical routine [13]. Cardiac MRI feature tracking (CMR-FT) [17,5]

^{*} This work was carried out during the internship of the author at United Imaging Intelligence, Cambridge, MA 02140




Fig. 1. Full cardiac cycle in Cine MRI. Example frames from a long-axis cine MRI show the heart motion starting from relaxation to contraction and then back to relaxation. The left ventricular (LV) muscle (V-shape dark region) undergoes large deformation during the cycle. Papillary muscle appears similar to the LV muscle.

estimates time-varying cardiac motion from cine MRI that usually includes a complete cycle of cardiac contraction and relaxation, as shown in Fig. 1. Starting from an initial image frame (usually end-diastole ED) as a reference, the next frame in time as a source image is compared, and motion occurred in-between is estimated. Then the source image becomes a new reference frame, and the process is repeated for all the consecutive frames to obtain the motion for the full cardiac cycle.

CMR-FT is a challenging topic still under active investigation, as addressed e.g. in [17,20,11,18,24,23,22]. A conventional image registration based method was proposed in [17] to estimate left ventricle motion using 2D B-spline free form deformation (FFD). Vigneault *et al.* in [20] proposed to perform cardiac segmentation and then apply B-Spline registration on the boundary points to track myocardium. This approach requires additional segmentation work. Tracking only boundary points also limits the motion estimation accuracy since both image features inside and outside the myocardium are not considered. Recently, Krebs *et al.* in [11] presented a variational autoencoder based image registration algorithm for estimating motion field for two consecutive frames from cine CMR. The VAE encoder takes the two images and encode them into latent variables which generate cardiac deformations via decoder. Because myocardium has similar image appearances with neighboring tissues/organs such as the papillary muscle, mere image-based motion estimation will face severe ambiguity in these regions, which causes anatomically unrealistic tracked results. Qin *et al.* proposed to jointly learn motion estimation and segmentation using a supervised deep learning method [18]. Zheng *et al.* introduced a semi-supervised learning method for apparent flow estimation aiming for explainable cardiac pathology classification [24]. Both studies make use of raw MR images and cardiac segmentation together for accurate flow estimation.

In this study, we propose an end-to-end framework to estimate cardiac motion that is aware of the underlying anatomy through shape-constraints in a weakly supervision manner, coined as anatomy-aware tracker (AATracker). We first train an unsupervised CNN-based dense tracker as the baseline. We then train a convolutional VAE model that learns the latent space of realistic myocardium shapes. We apply the trained VAE model to baseline-tracked my-

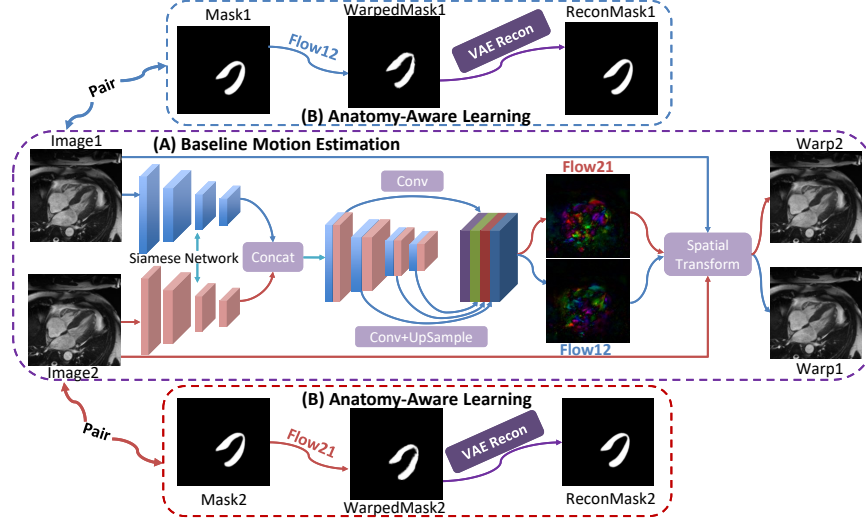


Fig. 2. The illustration of the baseline unsupervised motion estimation model and the weakly-supervised anatomy-aware model (AATracker). Module (A) presents the baseline, and Module (B) adds shape constraints to enhance the baseline motion estimation. Besides the image intensity loss and motion field regularization in the baseline model, loss between the warped mask with the VAE refined results (anatomy loss) and loss between the warped mask with its VAE reconstruction (reconstruction loss) are introduced in Module (B) to constrain the motion estimation for anatomy awareness.

ocardium and treat the anatomically reasonable myocardium masks from the VAE as self-learned shape constraints. The baseline model is further refined as AATracker using shape constraints for anatomy awareness. We evaluate the proposed method on the Kaggle Cardiovascular Disease dataset⁴.

The main contributions of this work are as follows: (1) We present AATracker, an end-to-end anatomy-aware cardiac motion estimation model via weak supervision. (2) We employ VAE to constrain motion estimation for anatomy awareness. (3) AATracker significantly improves the performance of myocardium feature tracking compared to baseline and conventional methods.

2 Method

2.1 Unsupervised Motion Estimation

The network structure for baseline cardiac motion estimation is shown in the middle part of Fig. 2. Two images, source and target at different time points of cine MRI, are inputs to the network. Three main modules include a Siamese network for mutual image feature extraction, a multi-scale decoder for flow field

⁴ <https://www.kaggle.com/c/second-annual-data-science-bowl/data>

generation, and a spatial transform that warps the source image with the flow fields [3,9]. Unlike previous studies [2,12], the motion estimation framework here is symmetrical, inspired by traditional symmetric registration algorithms [1,21].

For an input image pair $(\mathcal{I}_1, \mathcal{I}_2)$, \mathcal{F}_{12} is the flow field from \mathcal{I}_1 to \mathcal{I}_2 , and \mathcal{F}_{21} is the reverse. Using \otimes as the warping operator, $\mathcal{I}'_1 = \mathcal{F}_{12} \otimes \mathcal{I}_1$ and $\mathcal{I}'_2 = \mathcal{F}_{21} \otimes \mathcal{I}_2$ are the warped results of \mathcal{I}_1 and \mathcal{I}_2 via the spatial transform, respectively. The loss function enforcing warping consistency is defined as $\mathcal{L}_{cons} = \|\mathcal{I}_1 - \mathcal{I}'_2\| + \|\mathcal{I}_2 - \mathcal{I}'_1\|$. We add Huber loss $\mathcal{L}_H = \mathcal{H}(\mathcal{F}_{12}) + \mathcal{H}(\mathcal{F}_{21})$ on the flow fields as the regularizer for motion smoothness [8]. The loss function for the baseline model is then formulated as:

$$\mathcal{L}_{base} = \mathcal{L}_{cons} + \lambda_H \mathcal{L}_H, \quad (1)$$

where λ_H is the Huber loss weight.

2.2 Myocardium Feature Tracking

We can now perform CMR-FT for cine MRI based on the motion estimation between consecutive frames, as shown in Fig. 3(A). We denote the flow field between the $(n-1)$ -th and n -th frame as $\mathcal{F}_{(n-1)n}$. We compute the composite flow field $\hat{\mathcal{F}}_{1n}$ between the first and n -th frame using all intermediate motion fields:

$$\hat{\mathcal{F}}_{1n} = \begin{cases} \mathcal{F}_{12} & n = 2 \\ \hat{\mathcal{F}}_{1(n-1)} \oplus \mathcal{F}_{(n-1)n} & n > 2, \end{cases} \quad (2)$$

where \oplus is a flow composite operator and $\mathcal{F}_{ik} = \mathcal{F}_{ij} \oplus \mathcal{F}_{jk} = \mathcal{F}_{ij} \otimes \mathcal{F}_{jk} + \mathcal{F}_{jk}$. Note that motion is estimated only between two neighboring frames to avert large feature changes due to image intensity drifting and severe deformation seen in cine MRI. Additionally, the calculation of composite flow always refers to the ED frame because myocardium semantic information (*i.e.*, segmentation mask) is usually given at the ED frame, either by cardiologists' annotation or computer-aided algorithms. Without the loss of generality, we assume the first frame is ED. The myocardium semantic information can thus be obtained by warping the first frame with the composite flow field $\hat{\mathcal{F}}_{1n}$.

We take a refinement step to compensate for potential accumulation error through tracking. Specifically, after the warping to the n -th frame $\mathcal{I}'_n = \hat{\mathcal{F}}_{1n} \otimes \mathcal{I}_1$, motion \mathcal{F}_n^δ is estimated between \mathcal{I}'_n and \mathcal{I}_n . The final compensated motion between the first and the n -th frame is composed as $\hat{\mathcal{F}}_{1n}^* = \hat{\mathcal{F}}_{1n} \oplus \mathcal{F}_n^\delta$.

2.3 Anatomy-Aware Motion Estimation

Shape prior via VAE: The tracked myocardium using the above pipeline achieves fairly promising results. However, as the baseline model is mainly based on image intensity difference, the estimated motion can thus be severely affected by disturbances such as intensity-similar anatomies (*e.g.*, papillary muscle) and noises, which leads to tracked myocardium with unrealistic anatomy ("Tracked Masks" in Fig. 3(A)).

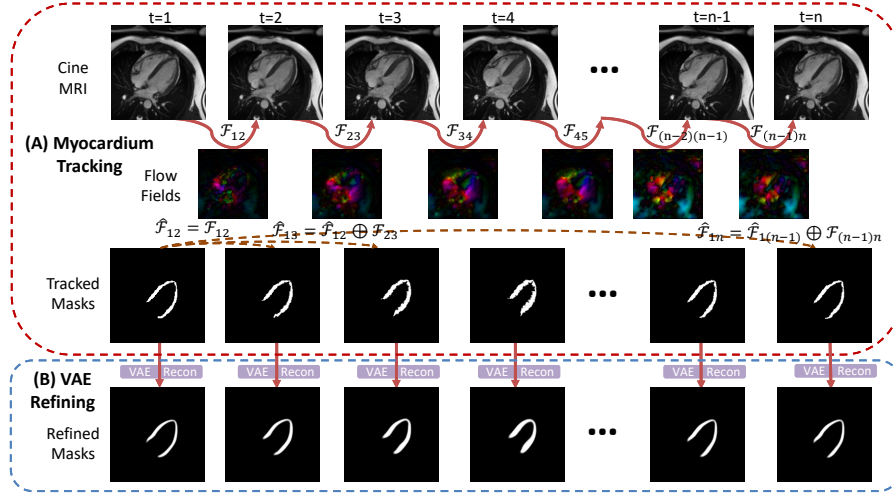


Fig. 3. Cine MRI myocardium tracking via motion estimation from consecutive frames and myocardium refining via VAE reconstruction. (A) describes the flow field composing procedure and the myocardium tracking based on the composite flow fields. (B) presents the refining of the tracked myocardium using VAE.

To solve this problem, we utilize convolutional VAE [16,7,6] to encode myocardium anatomy and enforce myocardium shape constraints (anatomy-awareness) in the motion estimation. Using available myocardium annotations, we train the VAE model to take the myocardium mask as input and reconstruct it. In addition to the reconstruction loss, the Kullback-Leibler Divergence (KLD) loss is used to enforce the latent space to conform to a standard normal distribution. Thus, the encoder’s outputs are a mean vector and a standard deviation vector. During the training, we sample from this distribution based on the mean and the standard deviation and reconstruct myocardium shape from the sample via the decoder. We used the trained VAE model (both the encoder and the decoder) to correct unrealistic myocardium masks. Specifically, the mean latent variable representing expected myocardium manifold in the latent space given the input mask is used without sampling. The decoder can decode such a latent variable into a realistic shape. Fig. 3(B) shows the reconstructed myocardium using the trained VAE model.

Weakly-supervised motion estimation using shape prior: For every cine MRI in the training dataset, we first use the baseline model to feature-track the myocardium to obtain coarse myocardium results for every image. The tracked myocardium through time, with possible unrealistic shapes, are then corrected by the VAE model, which are further used as anatomy constraints to improve the motion model. In this way, the motion estimation can mitigate the disturbances in images. Specifically, each input image pair $(\mathcal{I}_1, \mathcal{I}_2)$ now has its corresponding corrected myocardium $(\mathcal{M}_1, \mathcal{M}_2)$. We apply the flow fields

\mathcal{F}_{12} and \mathcal{F}_{21} to their corresponding masks and obtain warped masks $\mathcal{M}'_1 = \mathcal{F}_{12} \otimes \mathcal{M}_1$ and $\mathcal{M}'_2 = \mathcal{F}_{21} \otimes \mathcal{M}_2$. We expect that a plausible flow field will preserve the anatomy after warping and therefore propose the anatomy loss function $\mathcal{L}_{anat}^M = |\mathcal{M}_1 - \mathcal{M}'_2| + |\mathcal{M}_2 - \mathcal{M}'_1|$.

Furthermore, we apply the VAE model to the warped masks $(\mathcal{M}'_1, \mathcal{M}'_2)$ and obtain their reconstructed masks $(\mathcal{M}_1^{recon}, \mathcal{M}_2^{recon})$. We enforce the warped masks to be close to their VAE reconstructions to constrain the motion estimation model further using the reconstruction loss $\mathcal{L}_{recon}^M = |\mathcal{M}'_1 - \mathcal{M}_1^{recon}| + |\mathcal{M}'_2 - \mathcal{M}_2^{recon}|$. We define the anatomy-aware motion estimation loss as:

$$\mathcal{L} = \mathcal{L}_{cons} + \lambda_H \mathcal{L}_H + \lambda_{anat} \mathcal{L}_{anat}^M + \lambda_{recon} \mathcal{L}_{recon}^M, \quad (3)$$

where λ_H , λ_{anat} , and λ_{recon} are weights, and we denote this model as AATracker. Fig. 2 presents examples for the aforementioned images, flow fields, and masks. After refining the baseline model, we then apply the AATracker to the pipeline in Fig. 3(A) to track myocardium. Since only the myocardium in the first frame needs annotation, and the rest are tracked and then corrected by VAE, the whole process is weakly-supervised. It is worth pointing out that ED frame annotation is available in clinical setup for CMR-FT application. Also note that the AATracker directly estimates cardiac motions that preserve the underlying anatomy. The anatomy-aware learning module (Fig. 2B) is only performed during the training stage to infuse the anatomical knowledge into the motion estimation network.

3 Experiments and Results

3.1 Implementation Details

We benchmark the proposed method on 1,137 2-chamber and 1,111 4-chamber cine MRI from Kaggle. Each cine is from one patient with 30 frames. Data is randomly split for training, testing and validation. The VAE model is trained on fully-annotated 100 cine (3000 frames). The motion estimation refinement is trained on another 300 cine with the first frame annotated, and is tested on another 45 fully-annotated cine (1,350 frames). The remaining data, without any annotation, are used for the unsupervised baseline model training, hyper parameters tuning and model selection. All images and annotated myocardium masks are rescaled to the same resolution and cropped into 192×192 . The images are normalized into zero mean and unit standard deviation. Since no ground truth motion field is available, we evaluate the motion estimation based on the tracked myocardium with three commonly used metrics. The Dice similarity coefficient (DSC) measures the overlapping regions. The Hausdorff distance (HD) calculates the maximum distance between two boundaries, while the average symmetric surface distance (ASSD) calculates the average distance between two boundaries.

The implementation includes three aspects: baseline motion estimation, VAE, and AATracker. The baseline model is trained with λ_H as 0.02. We train the

Table 1. Comparison of FFD, diffeomorphic demons, Baseline model, baseline with anatomy loss (Baseline+anat), baseline with reconstruction loss (Baseline+recon) and the AATracker. Dice Similarity Coefficient (DSC), Hausdorff Distance (HD), and Average Symmetric Surface Distance (ASSD) on 2- and 4-chamber cine MRI are shown in mean(std).

Long-Axis	Method	DSC	HD (mm)	ASSD (mm)
2-chamber	FFD	0.768 (0.054)	6.242 (1.556)	1.480 (0.310)
	Diffeomorphic Demons	0.791 (0.051)	6.265 (1.313)	1.370 (0.238)
	Baseline	0.834 (0.039)	6.389 (1.444)	1.203 (0.216)
	Baseline+anat	0.835 (0.048)	5.659 (1.285)	1.158 (0.209)
	Baseline+recon	0.835 (0.038)	6.163 (1.507)	1.190 (0.224)
	AATracker	0.836 (0.048)	5.604 (1.252)	1.154 (0.206)
4-chamber	FFD	0.803 (0.039)	6.067 (1.339)	1.306 (0.238)
	Diffeomorphic Demons	0.813 (0.037)	6.936 (1.583)	1.274 (0.213)
	Baseline	0.861 (0.026)	6.228 (1.607)	1.062 (0.207)
	Baseline+anat	0.864 (0.026)	5.328 (1.185)	1.007 (0.163)
	Baseline+recon	0.865 (0.025)	5.936 (1.615)	1.026 (0.197)
	AATracker	0.864 (0.028)	5.303 (1.171)	0.998 (0.160)

VAE model with extensive data augmentation, including vertical and horizontal flipping, and multi-angle rotation, and set the latent space as a 32-d representation. In the AATracker training, λ_H , λ_{anat} , and λ_{recon} are set as 0.04, 6.0, 1.2, respectively. We compare the AATracker with the baseline model and the two shape-constrained models employing either anatomy loss or reconstruction loss. We also compared to two conventional registration methods, multi-scale free form deformation (FFD) [14,10] and multi-scale diffeomorphic demons [19], that have been previously used for medical image motion estimation [4]. All models are trained and tested on a standard workstation equipped with Intel Xeon Bronze 3106 CPUs and a Nvidia Titan XP GPU.

3.2 Results and Discussions

Table 1 presents the quantitative results over compared methods. Both anatomy loss and reconstruction loss can boost the performance of the baseline, while the effect of anatomy loss is more noticeable. The anatomy-aware model AATracker with both losses attains the best performance. Compared with the baseline model, AATracker reduces HD by 12.3% and 14.9% on 2- and 4-chamber cine MRI, respectively. The accuracy improvements of utilizing deep-learning-based methods over conventional methods are consistent with existing studies [2,12]. Besides, the AATracker takes much less time ($\sim 1.5s$) on average to estimate the motion than FFD ($\sim 46.1s$) and diffeomorphic demons ($\sim 25.2s$) for a cine MRI of typical size.

Fig. 4 shows a patient-wise comparison between the baseline model and AATracker. On both 2- and 4-chamber evaluation, the p-values from the Wilcoxon signed-rank test are significant for the two boundary-based metrics HD and

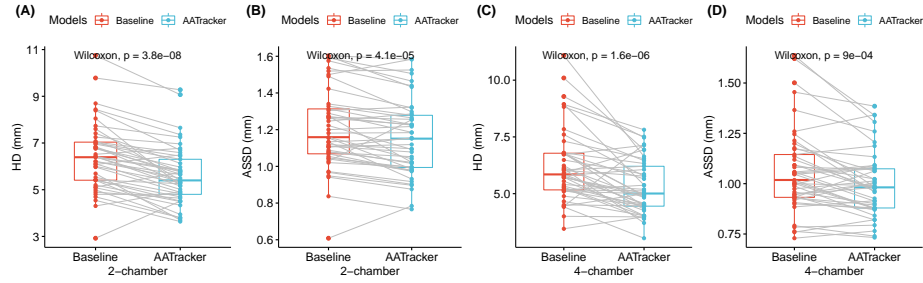


Fig. 4. Patient-wise comparison between the baseline model and AATracker on 2-chamber (A-B) and 4-chamber (C-D) cine MRI.




Fig. 5. Examples of myocardium annotation, baseline, and AATracker results.

ASSD. This result demonstrates that AATracker consistently improves the myocardium tracking results. Fig. 5 shows examples of tracked myocardium. The results of AATracker are visually more similar to the annotations. Most importantly, the anatomy-aware myocardium are more anatomically reasonable with smoother boundaries, demonstrating the effectiveness of the shape constraints.

Both Fig. 4 and Fig. 5 demonstrate the improved performance of the AATracker over baseline. These results indicate that AATracker preserves the anatomy structure during tracking. Arguably, DSC improvement is subtle after the motion estimation refinement. The main reason is that the myocardium boundary only accounts for a tiny part in the myocardium, and the refinement works mainly on the myocardium boundary without substantially affecting the overall myocardium shape. The reduction in ASSD is not as significant as HD, likely because ASSD is the average distance considering all boundary pixels while HD measures the worst error distance.

4 Conclusion

We present an end-to-end framework incorporating the anatomy prior to training for the awareness of anatomy in cardiac motion estimation. To our best knowledge, this is the first work that introduces shape constraints into the myocardium feature tracking via weak supervision. The proposed anatomy-aware method achieves consistent improvements over the baseline deep learning methods and two conventional methods. This study provides a sound basis for further cardiac function assessment, such as strain analysis.

References

1. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* **12**(1), 26–41 (2008)
2. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9252–9260 (2018)
3. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: *European conference on computer vision*. pp. 850–865. Springer (2016)
4. Haskins, G., Kruger, U., Yan, P.: Deep learning in medical image registration: a survey. *Machine Vision and Applications* **31**(1), 8 (2020)
5. Heinke, R., Pathan, F., Le, M., D’Angelo, T., Winau, L., Arendt, C., Vogl, T.J., Zeiher, A., Nagel, E., Puntmann, V.O.: Towards standardized postprocessing of global longitudinal strain by feature tracking–optistrain cmr-ft study. *BMC cardiovascular disorders* **19**(1), 267 (2019)
6. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework. *Iclr* **2**(5), 6 (2017)
7. Hou, X., Shen, L., Sun, K., Qiu, G.: Deep feature consistent variational autoencoder. In: *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. pp. 1133–1141. IEEE (2017)
8. Huber, P.J.: Robust estimation of a location parameter. In: *Breakthroughs in statistics*, pp. 492–518. Springer (1992)
9. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: *Advances in neural information processing systems*. pp. 2017–2025 (2015)
10. Joshi, S., Pizer, S., Fletcher, P.T., Thall, A., Tracton, G.: Multi-scale 3-d deformable model segmentation based on medial description. In: *Biennial International Conference on Information Processing in Medical Imaging*. pp. 64–77. Springer (2001)
11. Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T.: Learning a probabilistic model for diffeomorphic registration. *IEEE transactions on medical imaging* **38**(9), 2165–2176 (2019)
12. Mansilla, L., Milone, D.H., Ferrante, E.: Learning deformable registration of medical images with anatomical constraints. *Neural Networks* (2020)

13. MEMBERS, W.C., Hundley, W.G., Bluemke, D.A., Finn, J.P., Flamm, S.D., Fogel, M.A., Friedrich, M.G., Ho, V.B., Jerosch-Herold, M., Kramer, C.M., et al.: Accf/acr/aha/nasci/scmr 2010 expert consensus document on cardiovascular magnetic resonance: a report of the american college of cardiology foundation task force on expert consensus documents. *Circulation* **121**(22), 2462–2508 (2010)
14. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer methods and programs in biomedicine* **98**(3), 278–284 (2010)
15. Muser, D., Castro, S.A., Santangeli, P., Nucifora, G.: Clinical applications of feature-tracking cardiac magnetic resonance imaging. *World journal of cardiology* **10**(11), 210 (2018)
16. Pu, Y., Gan, Z., Henao, R., Yuan, X., Li, C., Stevens, A., Carin, L.: Variational autoencoder for deep learning of images, labels and captions. In: *Advances in neural information processing systems*. pp. 2352–2360 (2016)
17. Puyol-Antón, E., Ruijsink, B., Bai, W., Langet, H., De Craene, M., Schnabel, J.A., Piro, P., King, A.P., Sinclair, M.: Fully automated myocardial strain estimation from cine mri using convolutional neural networks. In: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. pp. 1139–1143. IEEE (2018)
18. Qin, C., Bai, W., Schlemper, J., Petersen, S.E., Piechnik, S.K., Neubauer, S., Rueckert, D.: Joint learning of motion estimation and segmentation for cardiac mr image sequences. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 472–480. Springer (2018)
19. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S72 (2009)
20. Vigneault, D.M., Xie, W., Bluemke, D.A., Noble, J.A.: Feature tracking cardiac magnetic resonance via deep learning and spline optimization. In: *International Conference on Functional Imaging and Modeling of the Heart*. pp. 183–194. Springer (2017)
21. Wu, G., Kim, M., Wang, Q., Shen, D.: S-hammer: Hierarchical attribute-guided, symmetric diffeomorphic registration for mr brain images. *Human brain mapping* **35**(3), 1044–1060 (2014)
22. Yu, H., Chen, X., Shi, H., Chen, T., Huang, T.S., Sun, S.: Motion pyramid networks for accurate and efficient cardiac motion estimation. *arXiv preprint arXiv:2006.15710* (2020)
23. Yu, H., Sun, S., Yu, H., Chen, X., Shi, H., Huang, T.S., Chen, T.: Foal: Fast online adaptive learning for cardiac motion estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4313–4323 (2020)
24. Zheng, Q., Delingette, H., Ayache, N.: Explainable cardiac pathology classification on cine mri with motion characterization by semi-supervised learning of apparent flow. *Medical image analysis* **56**, 80–95 (2019)