


Probabilistic Inverse Consistent Image Registration Using Sparse Bayesian Network

Shenglong Yang^{1,2}, Kangrong Xu^{1,2}, Zefeng He^{1,2}, Tianchao Feng^{1,2}, and Xuan Yang^{1,2} (✉) 

¹ College of Computer Science and Software Engineering, Shenzhen University, Shenzhen 518060, Guangdong, China

yangxuan@szu.edu.cn

² Guangdong Provincial Key Laboratory of Popular High Performance Computers, Shenzhen, Guangdong, China

Abstract. This paper proposes a probabilistic inverse consistency image registration network using a sparse BNN for cardiac motion estimation, aiming to simultaneously measure aleatoric and epistemic uncertainty. We construct a sparse BNN to predict the distribution parameters of the inverse consistency transformations between two images. Two symmetric Variational Autoencoders (VAEs) are constructed to predict the distribution parameters of latent variables in deformation space. The posterior distribution parameters of network weights are estimated during optimization, and only important weights are updated. Our sparse BNNs significantly reduce the computational cost and improve the registration accuracy by Bayesian model averaging (BMA). Experiments on a public cardiac MR dataset show that our sparse BNNs significantly improve the accuracy of the bidirectional registration for small datasets. It also provides aleatoric and epistemic uncertainty of registration results.

Keywords: Image registration · Inverse consistent · Bayesian Neural Network · Variational Auto-Encoder.

1 Introduction

Cardiac motion is essential to evaluate cardiac function, detect dysfunction such as cardiomyopathy, and study the heart bio-mechanic to understand cardiac physiology [28]. Deformable image registration (DIR) is the key technique in cardiac motion estimation. Unsupervised learning-based DIR networks have recently been proposed [10,13] with great potential due to rapid inference performance. There are two issues with existing deep-learning-based DIR models. First, most existing DIR networks focus on single-directional registration. The inverse-consistent DIR encourages bi-directional DVFs symmetrically deformed toward each other. It has specific advantages for cyclic cardiac motion estimation because it provides more accurate motion whether for end-systolic (ES) to end-diastolic (ED) phase or ED to ES. Studies have focused on introducing

inverse consistent constraints in image registration [7,17,18,27]. The inverse consistency (IC) regularization is commonly used to make the deformations consistent in both forward and backward directions. IC is performed for images [7,17] or DVFs [24]. Moreover, inverse consistency can be used to estimate the DVF uncertainty and, subsequently, the dose mapping uncertainty or the assessment of the deformable image registration quality [16,25].

The other issue in DIR models is the inability to provide reliable uncertainty estimates for the network’s decision and frequently occurring overconfident predictions [12]. Probabilistic image registration models predict the distribution of DVFs, which is powerful enough to provide confidence in predictive results at the pixel level [21,11]. There are two main types of uncertainty in image registration models: aleatoric and epistemic. Aleatoric uncertainty captures the inherent noise in registration images caused by anatomical changes, image artifacts, variations across different imaging modalities, etc. Epistemic uncertainty accounts for uncertainty in the model parameters generally due to a lack of training data. To simultaneously deal with aleatoric and epistemic uncertainty, we propose a probabilistic inverse consistent image registration model using variational autoencoders (VAEs) and sparse Bayesian Neural Networks (BNNs). VAE [19] is a probabilistic model that encodes image pairs into latent variables in DVF space, which can be used to measure the aleatoric uncertainty of DVFs. BNNs learn a probability distribution over the weights of networks to mitigate overfitting, enable learning from small datasets, and measure the uncertainty of our predictions. However, most recent networks have complex architectures, resulting in expensive computation costs for BNN training [15]. Sparse BNNs are proposed to reduce the training and testing costs. Sparsity is induced by using a special distribution, such as the Laplace distribution, to make a higher probability of coefficients being close to zero [14]. Another approach is to selectively assign a subset of network weights as Bayesian to make BNNs sparse [1,22,20].

This paper proposes an inverse consistent image registration model, denoted as ICRnet, using symmetric VAEs and sparse BNNs. Two similar VAEs are constructed to predict the forward and backward DVFs and generate aleatoric uncertainty of DVFs. An inverse consistent network is trained to ensure the forward and backward DVFs are inverse to each other. Weights of networks are assumed following Gaussian distributions whose parameters are estimated during the SGD iteration [23], which can reduce the learning complexity of BNNs. The signal-to-noise ratio of network weights is used to select important weights to be Bayesian, leading to sparse BNNs with less computation cost. To the best of our knowledge, it is the first probabilistic model to estimate aleatoric and epistemic uncertainty for bi-directional DVFs simultaneously. Our contributions are detailed as follows:

- We propose a probabilistic inverse consistent image registration model using two symmetric VAEs with the cross-attention mechanism and an inverse consistent network. The aleatoric uncertainty is measured using the posterior of latent variables estimated by VAEs. The forward and backward transforms




Fig. 1. The architecture of our ICRnet. Two symmetric VAEs predict latent variable distributions of bi-directional DVFs and the inverse consistency network regularizes bi-directional DVFs to be inverse to each other.

are constrained to be inverse of each other, improving the prediction accuracy of the bidirectional DVFs.

- A sparse BNN is proposed to estimate the posterior distribution of network weights to provide epistemic uncertainty of predicted DVFs. No weight sampling is performed, and only important weights are updated during optimization. Registration accuracy can be preserved when only 5% of the Bayesian weights are used.
- Experiments on public datasets show that our image registration model outperforms other models in bi-directional registration, especially for small datasets. The aleatoric and epistemic uncertainty about DVFs are also provided for downstream tasks.

2 Method

Our proposed ICRnet is illustrated in Figure 1. The architecture comprises two symmetric VAEs for predicting latent variable distributions of bi-directional DVFs and an inverse consistency network to regularize bi-directional DVFs to be inverse to each other. The posterior distributions of network weights are estimated by incorporating essential parameters updated during training. The final registration results are obtained by Bayesian model averaging.

Symmetric Variational Autoencoders. In our ICRnet, the forward and backward VAE networks are identical in architecture, we take the forward VAE as an example. Given a source image M and a target image F , the forward VAE aims to approximate the true posterior $p(z_{M \rightarrow F} | F, M)$ using the variational posterior $q(z_{M \rightarrow F} | F, M)$ by maximizing the Evidence Lower Bound (ELBO). The ELBO is:

$$ELBO = \mathbb{E}_{z_{M \rightarrow F} \sim q(z_{M \rightarrow F} | F, M)} [\log p(F | z_{M \rightarrow F}, M)] - KL(q(z_{M \rightarrow F} | F, M) || p(z)). \quad (1)$$

where the generative likelihood $p(F|z_{M \rightarrow F}, M)$ is estimated based on the similarity between the warped source image and the target image. The variational posterior $q(z_{M \rightarrow F}|F, M) \sim (\mu_z, \Sigma_z)$, μ_z is the mean vector and Σ_z is the diagonal covariance matrix. $p(z)$ is the given prior of z [11].

By sampling $q(z_{M \rightarrow F}|F, M)$ using the re-parametric trick, samples of $z_{M \rightarrow F}$ is obtained. We employ the compact support radial basis functions (CSRBFs) based spatial transformation model Φ using n control points $\{p_i\}_{i=1}^n$ to produce the dense DVF $\Phi_{M \rightarrow F}(u) = u + \sum_{i=1}^n z_{i, M \rightarrow F} \psi(\frac{\|u - p_i\|_2}{r})$, where ψ is a CSRBF with support r , $\psi(\frac{\|u - p_i\|_2}{r})$ is denoted as ψ_i in the following text. $\{z_{i, M \rightarrow F}\}_{i=1}^n$ are the elements of latent variables z . The decoder of VAE warps the source image M using the forward DVF.

A diffeomorphic layer is used to fine-tune velocity field $v_{M \rightarrow F}(u) = \Phi_{M \rightarrow F}(u) - u$ to generate invertible DVFs. The diffeomorphic transformation is defined by the ordinary differential equation [3]: $\frac{\partial \Phi_{M \rightarrow F}^t}{\partial t} = v_{M \rightarrow F}(\Phi_{M \rightarrow F}^t)$. $\Phi_{M \rightarrow F}^0$ is the identity transformation at $t = 0$; $t \in [0, 1]$ is the time and $\Phi_{M \rightarrow F}^1 = \exp(v_{M \rightarrow F})$ is the final transformation the diffeomorphic layer estimated. The Euler method is used to compute $\Phi_{M \rightarrow F}^1$ with successive small time-steps $\frac{1}{2^T}$.

$$\Phi_{M \rightarrow F}^{1/2^T} = u + v_{M \rightarrow F}(u), \quad \Phi_{M \rightarrow F}^{1/2^{t-1}} = \Phi_{M \rightarrow F}^{1/2^t} \circ \Phi_{M \rightarrow F}^{1/2^t}. \quad (2)$$

In this model, the Variational Autoencoder (VAE) encoder consists of a T2T module and N layers of Transformer encoders. The T2T module preprocesses the source image M and the target image F to form feature sequences, which are input into the forward and backward VAE encoders. The Transformer utilizes cross-attention mechanisms that handle the relationships between the two images.

Inverse Consistency Network. In our ICRnet, an inverse consistency network (ICN) is constructed to ensure the forward and backward DVFs in the symmetric variational autoencoders are inverse to each other. Unlike existing methods [24], our ICN requires that diffeomorphic intermediate transition DVFs are mutually reversible instead of the final DVFs. The advantage is that the deformation of the intermediate DVF is relatively small, and it is relatively easier to train the ICN to generate mutually reversible DVFs. The architecture of our ICN is similar to two U-Nets with shared parameters. The forward transformation $\Phi_{M \rightarrow F}$ and backward transformation $\Phi_{F \rightarrow M}$ are input to ICN to generate fine-tuned transformations $\tilde{\Phi}_{M \rightarrow F}$ and $\tilde{\Phi}_{F \rightarrow M}$, which is expected to be inverses of each other further. The diffeomorphic layer is used to recursively compute $\{\tilde{\Phi}_{M \rightarrow F}^{1/2^i}\}_{i=1}^T$ and $\{\tilde{\Phi}_{F \rightarrow M}^{1/2^i}\}_{i=1}^T$. These diffeomorphic intermediate outputs compute the inverse consistency loss \mathcal{L}_{IC} .

$$\mathcal{L}_{IC} = \frac{1}{T} \sum_{i=1}^T \left[LCC(M \circ \tilde{\Phi}_{M \rightarrow F}^{1/2^i}, F \circ \tilde{\Phi}_{F \rightarrow M}^{1/2^{T-i+1}}) + LCC(F \circ \tilde{\Phi}_{F \rightarrow M}^{1/2^i}, M \circ \tilde{\Phi}_{M \rightarrow F}^{1/2^{T-i+1}}) \right] \quad (3)$$

$$\mathcal{L}_{ICE} = \left\| \tilde{\Phi}_{M \rightarrow F} \circ \tilde{\Phi}_{F \rightarrow M} - Id \right\|_1, \quad (4)$$

where Id is identity transformation. The two terms in Equ. (3) enforce intermediate warped images to be the same in two directions, implying the intermediate transformation $\tilde{\Phi}_{M \rightarrow F}^{1/2^i}$ is expected to be reversed to the backward intermediate transformation $\tilde{\Phi}_{F \rightarrow M}^{1/2^{T-i+1}}$, vice versa. The \mathcal{L}_{ICE} loss ensures the final forward and backward transformations are reversed to each other.

Sparse Bayesian Neural Network. Maddox *et al.* [23] proposed SWAG that approximated the Gaussian posterior distribution of neural network weights using the mean and low-rank covariance matrix of weights from the SGD iterates. Nevertheless, SWAG estimated posterior distributions of all weights. It has been observed that BNNs can be significantly condensed without greatly compromising performance [22, 1, 14]. Inspired by this observation, we make our ICRnet a sparse BNN by placing posterior distributions on weights and updating only important weights using the signal-to-noise ratio. Denote parameters of our registration model as θ , the posterior $p(\theta|D) \sim \mathcal{N}(\boldsymbol{\mu}_\theta, \boldsymbol{\Sigma}_\theta)$,

$$\boldsymbol{\mu}_\theta = \frac{1}{L} \sum_{i=1}^L \theta_i, \quad \boldsymbol{\Sigma}_\theta = \text{diag}(\bar{\theta}^2 - \boldsymbol{\mu}_\theta^2) + \frac{1}{K-1} \hat{\boldsymbol{D}} \hat{\boldsymbol{D}}^\top, \quad (5)$$

where θ_i is the snapshot of θ at the i th training iteration, $\bar{\theta}^2 = \frac{1}{L} \sum_{i=1}^L \theta_i^2$, L is the number of snapshot θ_i . The covariance matrix $\boldsymbol{\Sigma}_\theta$ combines a diagonal matrix whose diagonal elements are the variances of each parameter and a low-rank approximation matrix $\hat{\boldsymbol{D}}$. $\hat{\boldsymbol{D}}$ has K columns that is much less than the element number N of θ , where each column of $\hat{\boldsymbol{D}}$ is a snapshot of θ .

To make our BNNs sparse, the posterior is confined to a subspace R^s of the original space R^N , $s = \gamma N$ with the sparse ratio γ . Denote $I = \{n_1, \dots, n_s\}$ as the index of selected weights to be Bayesian, for the weight with index $i \notin I$, its posterior remains as before, and selected weights are updated. During the SGD iteration of training our registration model, I is updated based on the signal-to-noise ratio of currently estimated distribution parameters because weights with small magnitudes tend to have little influence on the output. To simplify weights selection, we use the k th element of $\bar{\theta}^2 - \boldsymbol{\mu}_\theta^2$ as the variance of the k th weight $\theta_{i,k}$. The signal-to-noise ratio of $\theta_{i,k}$ is defined as,

$$\text{SNR} = \frac{\mathbb{E}_{p(\theta|D)}(|\theta_{i,k}|)}{\sqrt{\text{VAR}_{p(\theta|D)}(|\theta_{i,k}|)}} = \frac{\mu_k(2\Gamma(\frac{\mu_k}{\sigma_k}) - 1) + \frac{2\sigma_k}{\sqrt{2\pi}} \exp(-\frac{\mu_k^2}{2\sigma_k^2})}{\sqrt{\sigma_k^2 + \mu_k^2 - \left[\mu_k(2\Gamma(\frac{\mu_k}{\sigma_k}) - 1) + \frac{2\sigma_k}{\sqrt{2\pi}} \exp(-\frac{\mu_k^2}{2\sigma_k^2}) \right]^2}}, \quad (6)$$

where Γ is the cumulative distribution function of standard normal. SNR has the advantage of preventing the model from instability because a high σ_k implies the weight has little effect on the model. After our ICRnet is trained, the inference is performed using the Bayesian moving average by sampling multiple $\{\theta_s\}_{s=1}^S$ to predict DVFs, and the average DVF is the predictive result. The training algorithm of our sparse BNNs and Bayesian model averaging procedure is shown in the supplement.

Aleatoric and Epistemic Uncertainty. Registration uncertainty includes aleatoric and epistemic uncertainty. Aleatoric uncertainty is provided by the variance of latent variables of VAEs, corresponding to the deviation of DVFs caused by data inherent noise. The aleatoric uncertainty $Var_a(u)$ of the DVF at pixel u is

$$Var_a(u) = trace [\boldsymbol{\Sigma}_z \times diag([\psi_1^2, \dots, \psi_n^2]^T)]. \quad (7)$$

The epistemic uncertainty $Var_e(u)$ of the DVF at u is found by the definition of variance caused by the posterior of weights:

$$Var_e(u) = \frac{1}{S} \sum_{s=1}^S \left(\sum_{i=1}^n \boldsymbol{\mu}_{z,i}(\hat{\theta}_s) \psi_i \right)^2 - \left(\frac{1}{S} \sum_{s=1}^S \sum_{i=1}^n \boldsymbol{\mu}_{z,i}(\hat{\theta}_s) \psi_i \right)^2. \quad (8)$$

where S is the number of Bayesian moving averages. Significant epistemic uncertainty might suggest that the registration result is made on data with which the model has less experience. Therefore, epistemic uncertainty might correspond to erroneous registration or outlier data.

3 Experiments

We evaluate our approach on four publicly available cardiac image datasets: ACDC [5], York [2], MICCAI2009 [26], and M&Ms [6]. The number of image pairs in these datasets is provided in supplements, where ACDC and M&Ms are large datasets, while York and MICCAI2009 are small datasets. The cardiac slices at the end of the diastolic (ED) and the end of the systolic (ES) phases are registered to each other. For image registration evaluation, the masks provided by experts are mapped using the estimated DVFs. The Dice score, Hausdorff distance (HD), the bending energy (BE), and the number of non-positive Jacobian determinants $|J_\phi| \leq 0$ are used to measure the registration results. The loss \mathcal{L}_{ICE} is used to measure the reversibility of bidirectional DVFs.

Sparse BNNs. At first, we evaluate the performance of ICRnet without BNNs, with BNNs, and with sparse BNNs (denoted as sBNN) using $\gamma = 20\%$. As listed in Table 1, it can be seen that whether for forward or backward registration, our ICRnet with BNNs improved accuracy by 1% and 1.1% for two small datasets, MICCAI2009 and York, respectively, while improved accuracy by 0.8% for two large datasets. It validates the effectiveness of our sparse BNNs, especially for small datasets. Note that the forward registration accuracy is better than the backward because the forward deformation is contractive, and the target objects are small, resulting in a large Dice value. We compare the performance of our model with different γ . The metric $|J_\phi| \leq 0$ is close whether for all cases. It is observed that the performance of our model with $\gamma = 5\%$ is close to that of ICRnet without BNNs, while the computation time reduces about 65%, which evaluates the effectiveness of our sparse BNNs. Details can be seen in supplements.

Table 1. Comparison of ICRnet without BNNs, with BNNs, and with sBNNs ($\gamma = 20\%$). Data format: mean

Dataset	Method	Dice(%)		BE		HD		$ J_\phi \leq 0$		ICE	
		fwd.	bwd.	fwd.	bwd.	fwd.	bwd.	fwd.	bwd.	fwd.	bwd.
M&Ms	ICRnet	86.3	89.7	14.69	12.86	5.95	5.82	0.00	0.58	0.268	0.268
	ICRnet w. BNN	87.1	90.4	10.87	9.58	4.62	4.25	0.00	0.32	0.260	0.261
	ICRnet w. sBNN	86.7	90.2	11.32	10.82	4.89	4.78	0.00	0.46	0.264	0.262
ACDC	ICRnet	86.2	88.9	13.49	12.57	5.59	4.92	0.00	0.32	0.155	0.155
	ICRnet w. BNN	86.9	89.7	8.51	7.55	4.87	3.81	0.00	0.18	0.149	0.147
	ICRnet w. sBNN	86.5	89.3	9.72	8.72	5.11	4.36	0.00	0.26	0.152	0.151
MICCAI	ICRnet	88.2	92.3	10.79	9.26	5.20	4.47	0.00	0.22	0.185	0.185
	ICRnet w. BNN	89.2	93.2	7.32	6.18	4.24	3.89	0.00	0.07	0.179	0.178
	ICRnet w. sBNN	88.7	92.8	9.26	8.37	4.86	4.23	0.00	0.14	0.182	0.184
York	ICRnet	85.2	89.8	12.75	10.37	6.72	5.38	0.00	0.67	0.156	0.156
	ICRnet w. BNN	86.3	90.9	9.43	8.38	5.48	4.26	0.00	0.48	0.149	0.152
	ICRnet w. sBNN	85.8	90.4	10.54	9.72	6.24	4.78	0.00	0.56	0.152	0.154

fwd., bwd: abbreviations of forward and backward.

Registration results. Furthermore, we compared our network with eight other approaches: KrebsDiff [21], DalcaDiff[9], VoxelMorph [4], NetGI[11], TransMorph[8], ICNet[27], CycleMorph[17], and SYMNet [24], where ICNet, CycleMorph, and SYMNet are inverse consistent registration networks. Using different models, Table 2 lists the average registration accuracy of the forward and backward on the M&Ms and MICCAI2009 datasets. Our network improves the Dice by 2.3% and 2.75% compared with the average Dice of other models on the two datasets, respectively. Other registration results are provided in the supplement. Our model ranks second in terms of ICE, only behind SYMnet, but achieves the Dice that is 1.1% and 1.9% higher than SYMnet for two datasets. Demonstration of DVFs and deformed images are visualized in the supplement.

Table 2. Comparison of forward and backward average registration results on all datasets using different models. $\gamma = 20\%$ used in our model. Data format: mean

Method	M&Ms					MICCAI 2009				
	Dice(%)	BE	HD	$ J_\phi \leq 0$	ICE	Dice(%)	BE	HD	$ J_\phi \leq 0$	ICE
KrebsDiff [21]	83.8	34.02	6.42	8.58	0.744	88.7	26.02	4.98	0.42	0.720
DalcaDiff [9]	84.9	99.83	6.78	0.12	0.739	88.2	53.97	5.92	0.00	0.449
VoxelMorph [4]	85.9	197.84	5.26	38.42	0.578	86.9	305.13	5.52	39.65	0.433
NetGI [11]	86.2	14.56	5.18	10.24	0.944	89.5	15.59	4.78	16.11	0.602
TransMorph [8]	86.8	613.19	5.31	108.81	0.573	88.6	338.48	5.28	125.53	0.485
ICnet [27]	87.5	93.06	4.97	1.45	0.312	86.9	44.56	5.45	0.41	0.141
CycleMorph [17]	86.4	223.61	5.47	24.61	0.435	86.7	89.42	5.77	13.76	0.325
SYMnet [24]	87.3	182.60	5.12	0.13	0.028	88.9	469.56	5.38	0.04	0.032
Ours	88.4	11.07	4.83	0.23	0.263	90.8	8.81	4.54	0.07	0.183

Aleatoric and epistemic uncertainty. Figure 2 visualizes the aleatoric and epistemic uncertainty of several cases. It is observed that high aleatoric uncer-

tainty appears in the inner smooth areas of objects due to inherent variability in these areas. Epistemic uncertainty maps show that high epistemic uncertainty occurs where large deformation occurs, such as the border of the right ventricle, indicating non-confidence in these areas.




Fig. 2. Illustration of aleatoric and epistemic uncertainty. First row: aleatoric uncertainty. Second row: epistemic uncertainty.

Ablation Study. We further conduct an ablation study on the ACDC dataset to verify the contributions of different parts in our ICRnet. Here, the BNNs are ignored to focus on the performance of two symmetric VAEs and the ICN. The ablation experiments include the bidirectional registration (BR) and two inverse consistent constraints \mathcal{L}_{IC} and \mathcal{L}_{ICE} . Ablation results are listed in Table 3. It shows that inverse consistent constraints improve registration accuracy. Moreover, the smoothness and topology-preservation of DVFs have also been improved at the same time.

Table 3. Performance comparison of different parts of ICRnet without BNNs on the ACDC dataset. Data format: mean

Parts	Dice(%)		BE		HD		$ J_\phi \leq 0$		ICE	
	fwd.	bwd.	fwd.	bwd.	fwd.	bwd.	fwd.	bwd.	fwd.	bwd.
BR	85.6	88.4	14.82	14.24	6.28	5.62	0.36	0.89	0.256	0.252
BR + \mathcal{L}_{IC}	85.9	88.7	13.26	11.72	5.92	5.21	0.00	0.58	0.212	0.208
BR + \mathcal{L}_{IC} + \mathcal{L}_{ICE}	86.2	88.9	13.49	12.57	5.59	4.92	0.00	0.32	0.155	0.155

4 Conclusion

Our proposed inverse consistent image registration model enhances bi-directional registration accuracy while providing aleatoric and epistemic uncertainty for registration. Inverse consistent regularization improves bi-directional registration accuracy simultaneously. Unlike commonly used variational inference for BNNs, no sampling of network weights is required during training, and important weights are selected and updated for our sparse BNNs. It has the advantages

of less computation and learning from small datasets, making our approach a valuable advancement in learning-based image registration for clinical applications.

Acknowledgments. This work is supported by the Shenzhen Fundamental Research Program (JCYJ20220531102407018), Guangdong Province Key Laboratory of Popular High-Performance Computers 2017B030314073, and National Natural Science Foundation of China (Grant No.62201355).

Disclosure of Interests. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Abboud, Z., Lombaert, H., Kadoury, S.: Sparse bayesian networks: Efficient uncertainty quantification in medical image analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 675–684. Springer (2024)
2. Andreopoulos, A., Tsotsos, J.K.: Efficient and generalizable statistical models of shape and appearance for analysis of cardiac mri. *Medical Image Analysis* **12**(3), 335–357 (2008)
3. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1), 95–113 (2007)
4. 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)
5. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
6. Campello, V.M., Gkontra, P., Izquierdo, C., Martín-Isla, C., Sojoudi, A., Full, P.M., Maier-Hein, K., Zhang, Y., He, Z., Ma, J., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: The m&ms challenge. *IEEE Transactions on Medical Imaging* p. 9458279 (2021)
7. Chatterjee, S., Bajaj, H., Siddiquee, I.H., Subbarayappa, N.B., Simon, S., Shashidhar, S.B., Speck, O., Nürnberger, A.: Micdir: Multi-scale inverse-consistent deformable image registration using unetmss with self-constructing graph latent. *Computerized Medical Imaging and Graphics* p. 102267 (2023)
8. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: Transmorph: Transformer for unsupervised medical image registration. *arXiv preprint arXiv:2111.10480* (2021)
9. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 729–738. Springer (2018)
10. Fu, Y., Lei, Y., Wang, T., Curran, W.J., Liu, T., Yang, X.: Deep learning in medical image registration: a review. *Physics in Medicine & Biology* **65**(20), 20TR01 (2020)

11. Gan, Z., Sun, W., Liao, K., Yang, X.: Probabilistic modeling for image registration using radial basis functions: Application to cardiac motion estimation. *IEEE Transactions on Neural Networks and Learning Systems* (2022)
12. Gawlikowski, J., Tassi, C.R.N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al.: A survey of uncertainty in deep neural networks. *arXiv preprint arXiv:2107.03342* (2021)
13. Haskins, G., Kruger, U., Yan, P.: Deep learning in medical image registration: a survey. *Machine Vision and Applications* **31**, 1–18 (2020)
14. Huang, J., Wu, Q., Ren, Y., Yang, F., Yang, A., Yang, Q., Pu, X.: Sparse bayesian deep learning for cross domain medical image reconstruction. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 38, pp. 2339–2347 (2024)
15. Khawaled, S., Freiman, M.: Npb-rec: A non-parametric bayesian deep-learning approach for undersampled mri reconstruction with uncertainty estimation. *Artificial Intelligence in Medicine* **149**, 102798 (2024)
16. Kierkels, R.G., den Otter, L.A., Korevaar, E.W., Langendijk, J.A., van der Schaaf, A., Knopf, A.C., Sijtsma, N.M.: An automated, quantitative, and case-specific evaluation of deformable image registration in computed tomography images. *Physics in Medicine & Biology* **63**(4), 045026 (2018)
17. Kim, B., Kim, D.H., Park, S.H., Kim, J., Lee, J.G., Ye, J.C.: Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis* **71**, 102036 (2021)
18. Kim, B., Kim, J., Lee, J.G., Kim, D.H., Park, S.H., Ye, J.C.: Unsupervised deformable image registration using cycle-consistent cnn. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI* 22. pp. 166–174. Springer (2019)
19. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *CoRR abs/1312.6114* (2013), <https://api.semanticscholar.org/CorpusID:216078090>
20. Kong, I., Yang, D., Lee, J., Ohn, I., Baek, G., Kim, Y.: Masked bayesian neural networks: Theoretical guarantee and its posterior inference. In: *International Conference on Machine Learning*. pp. 17462–17491. PMLR (2023)
21. Krebs, J., Mansi, T., Mailhé, B., Ayache, N., Delingette, H.: Unsupervised probabilistic deformation modeling for robust diffeomorphic registration. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, pp. 101–109. Springer (2018)
22. Li, J., Miao, Z., Qiu, Q., Zhang, R.: Training bayesian neural networks with sparse subspace variational inference. *arXiv preprint arXiv:2402.11025* (2024)
23. Maddox, W.J., Garipov, T., Izmailov, P., Vetrov, D., Wilson, A.G.: A simple baseline for Bayesian uncertainty in deep learning. *Curran Associates Inc., Red Hook, NY, USA* (2019)
24. Mok, T.C., Chung, A.: Fast symmetric diffeomorphic image registration with convolutional neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4644–4653 (2020)
25. Paganelli, C., Meschini, G., Molinelli, S., Riboldi, M., Baroni, G.: Patient-specific validation of deformable image registration in radiation therapy: overview and caveats. *Medical physics* **45**(10), e908–e922 (2018)
26. Radau, P., Lu, Y., Connelly, K., Paul, G., Dick, A., Wright, G.: Evaluation framework for algorithms segmenting short axis cardiac mri. *The MIDAS Journal-Cardiac MR Left Ventricle Segmentation Challenge* **49** (2009)

- 27. Zhang, J.: Inverse-consistent deep networks for unsupervised deformable image registration. arXiv preprint arXiv:1809.03443 (2018)
- 28. Zhang, X., You, C., Ahn, S., Zhuang, J., Staib, L., Duncan, J.: Learning correspondences of cardiac motion from images using biomechanics-informed modeling. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 13–25. Springer (2022)