

Heteroscedastic Uncertainty Estimation Framework for Unsupervised Registration

Xiaoran Zhang^{*1(✉)}, Daniel H. Pak^{*1}, Shawn S. Ahn^{1,7},
Xiaoxiao Li^{3,4,5}, Chenyu You⁶, Lawrence H. Staib^{1,5},
Albert J. Sinusas^{1,5,8}, Alex Wong², and James S. Duncan^{1,5}

¹ Biomedical Engineering, Yale University, New Haven, USA
xiaoran.zhang@yale.edu

² Computer Science, Yale University, New Haven, USA

³ Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, Canada

⁴ Vector Institute, Toronto, Canada

⁵ Radiology & Biomedical Imaging, Yale University, New Haven, USA

⁶ Electrical Engineering, Yale University, New Haven, USA

⁷ Department of Surgery, University of Pennsylvania, Philadelphia, USA

⁸ Department of Internal Medicine (Cardiology), Yale University, New Haven, USA

Abstract. Deep learning methods for unsupervised registration often rely on objectives that assume a uniform noise level across the spatial domain (e.g. mean-squared error loss), but noise distributions are often heteroscedastic and input-dependent in real-world medical images. Thus, this assumption often leads to degradation in registration performance, mainly due to the undesired influence of noise-induced outliers. To mitigate this, we propose a framework for heteroscedastic image uncertainty estimation that can adaptively reduce the influence of regions with high uncertainty during unsupervised registration. The framework consists of a collaborative training strategy for the displacement and variance estimators, and a novel image fidelity weighting scheme utilizing signal-to-noise ratios. Our approach prevents the model from being driven away by spurious gradients caused by the simplified homoscedastic assumption, leading to more accurate displacement estimation. To illustrate its versatility and effectiveness, we tested our framework on two representative registration architectures across three medical image datasets. Our method consistently outperforms baselines and produces sensible uncertainty estimates. The code is publicly available at https://voldemort108x.github.io/hetero_uncertainty/.

1 Introduction

Deformable image registration solves for a dense pixel-wise displacement map that aligns one image to another. It is a key medical image analysis technique that enables disease progression monitoring and surgical guidance [8, 18, 24, 31]. Classical methods (e.g. elastic-type models [13], free-form deformation with b-splines

* Equal contribution.

[21], diffeomorphic models [1]) are often computationally expensive and impractical for large-scale analyses. Thus, numerous works have proposed training a neural network to accelerate the test-time prediction of the displacement map [3, 5, 9, 11, 32, 30]. Such frameworks commonly utilize an unsupervised objective by minimizing the mean-square error (MSE) between the warped and target images. By using this objective, the existing frameworks assume an additive *homoscedastic* Gaussian image noise, i.e. $\epsilon \sim \mathcal{N}(0, \sigma^2)$, with a constant noise variance across the spatial domain. However, this assumption is problematic for medical images, where the noise is intrinsically *heteroscedastic* and *input-dependent* (e.g. MRI [7, 26] or ultrasound [19, 29]). The non-uniform noise variance across the image space (Fig. 1, left) is due to multiple factors including changes in anatomical structures or patient motion artifacts. The simplified homoscedastic assumption disregards such variations in noise levels and results in undesired penalization of false noise-induced outliers, causing unnatural and inaccurate deformations.

To mitigate this, we propose a probabilistic heteroscedastic noise modeling framework for unsupervised image registration (Fig. 1, right). Our method involves two modules: (1) a standard displacement estimator and (2) a variance estimator that predicts the input-specific heteroscedastic noise for each image pair. For proper utilization of the predicted noise, we introduce a novel adaptive weighting strategy based on the relative γ -exponentiated signal-to-noise ratio (SNR). This enables the successful convergence of our collaborative training strategy, where the two estimators improve upon each other via information exchange and loss calibration. We validated the effectiveness and versatility of our proposed framework with extensive experiments using two representative neural network architectures and three cardiac datasets. Our proposed framework can be operated in a plug-and-play manner, and provides sensible heteroscedastic uncertainty measures to reflect spatially varying noise. This work may help provide a new perspective in data-driven input noise modeling for deep learning-based unsupervised registration.

Our contributions include (1) We first analyze a naive approach of applying heteroscedastic noise modeling to unsupervised registration and identify the pitfalls of such an approach. (2) From which, we propose a probabilistic framework for data-driven estimation of heteroscedastic uncertainty for unsupervised registration. (3) We introduce an adaptive γ -exponentiated relative SNR weighting strategy for proper loss calibration. (4) We demonstrate the effectiveness and versatility of our method on two representative registration architectures across three datasets, demonstrating consistent statistically significant improvements over the baselines while providing sensible uncertainty maps.

2 Related works

2.1 Unsupervised image registration

Voxelmorph is a popular convolutional neural network (CNN) model for unsupervised registration that was trained with the MSE loss, which assumes homoscedastic input noise [3]. Extensions of this work include Voxelmorph-diff [6],

Hypermorph [11], and Synthmorph [9] with similar assumptions. Transformers have also been explored for unsupervised registration by replacing the convolutional encoders in the Voxelmorph U-Net with swin transformer layers [5]. Transformers can increase the receptive field of the encoding branch, potentially improving the long-range modeling capability [15]. Further extensions include replacing the convolutional decoders with transformer layers [23] and introducing multi-scale pyramids [16].

In this paper, we selected Voxelmorph [3] and Transmorph [5] as the representative registration architectures and tested our proposed framework and other baselines on both across different datasets. We also compared Voxelmorph-diff [6] with a direct extension of our framework that incorporates the *displacement* uncertainty, in addition to our main contribution of modeling *image* uncertainty. The detailed analysis can be found in the Supp. Mat.

2.2 Heteroscedastic uncertainty estimation

To model the heteroscedastic uncertainty in imaging problems, Kendall et al. [12] proposed a joint mean-and-variance optimization strategy by minimizing the negative log-likelihood (NLL) under the maximum likelihood estimation (MLE) framework. This formulation has been effective in many applications, such as surface normal estimation [2] and image segmentation [17]. Seitzer et al. further advanced the formulation with β -NLL to address the undesired undersampling effects of the inverse variance weighting of the data fidelity term [22]. In this paper, we selected both NLL [12] and β -NLL [22] as our baselines for heteroscedastic uncertainty estimation.

2.3 Adaptive weighting schemes

A wide range of imaging problems is cast into optimizing an energy function, which contains a data-fidelity term and a regularization term. The relative importance between the two terms is usually tuned via hyperparameter optimization. However, such a strategy disregards the heteroscedastic nature of the error residuals [27]. Several adaptive weighting schemes have been proposed to address this issue by adjusting the weighting throughout optimization [10, 28]. Wong et al. [27] extended the formulation to multi-frame setting. In this paper, we selected AdaReg [28] and AdaFrame [27] as baselines to compare with our proposed adaptive signal-to-noise weighting scheme.

3 Analysis of a naive approach

Let $I : \Omega \rightarrow \mathbb{R}$ be an image and $\Omega \subseteq \mathbb{R}^d$ its spatial domain. Unsupervised image registration aims to find the alignment between a moving image I_m and a fixed image I_f , without ground truth information about warping displacement field $z : \Omega \rightarrow \mathbb{R}^d$. For each image pair (I_m, I_f) , we assume I_m is a fixed parameter

and I_f is a noisy observation of the warped image $\hat{I}_f = I_m \circ z$. Then, z is a sample from the posterior distribution $p(z|I_f; I_m)$.

With this formulation, the main task is finding the displacement posterior $p(z|I_f; I_m)$ that will maximize the data likelihood $p(I_f; I_m)$. Since directly solving for $p(z|I_f; I_m)$ is intractable, we apply the variational approach to approximate it using a neural network $q_\theta(z|I_f; I_m)$. To minimize the discrepancy between approximate and real posterior distributions, we compute the KL-divergence and arrive at the negative evidence lower bound:

$$\begin{aligned} \text{KL}(q_\theta(z|I_f; I_m)||p(z|I_f; I_m)) &= \text{KL}(q_\theta(z|I_f; I_m)||p(z)) \\ &\quad - \mathbb{E}_{q_\theta}(\log p(I_f|z; I_m)) + \text{const.} \end{aligned} \quad (1)$$

For the warped image distribution $p(I_f|z; I_m)$, we assume an additive zero-mean heteroscedastic Gaussian noise with spatially varying variance $\sigma_I^2 : \Omega \rightarrow \mathbb{R}^+$ to incorporate input-dependent heteroscedasticity

$$p(I_f|z; I_m) = \mathcal{N}(I_f; \hat{I}_f, \sigma_I^2). \quad (2)$$

Similar to [3, 12], we assume pixel-wise independence, isotropic Gaussian for the displacement posterior $q_\theta(z|I_f; I_m) = \mathcal{N}(z; \mu_z, \sigma_z^2 \mathbb{I})$, and uniform Gaussian for the displacement prior $p(z) = \mathcal{N}(0, \mathbb{I})$. Then, we arrive at the preliminary loss to be minimized for each image pair

$$\mathcal{L} = \mathbb{E}_\Omega \left[\underbrace{\frac{1}{\sigma_I^2} \|I_f - \hat{I}_f\|^2}_{\mathcal{L}_{\text{data}}} + \log \sigma_I^2 + \lambda \|\nabla z\|^2 \right], \quad (3)$$

where $\hat{I}_f = I_m \circ z$. σ_I^2 is estimated by our variance estimator as $\hat{\sigma}_I^2 = h_\phi(I_f, \hat{I}_f)$, using the fixed and moved image pair as input. z is estimated by our displacement estimator as $\hat{z} = q_\theta(I_m, I_f)$, using the moving and fixed pair as input.

4 Our approach

4.1 Collaborative heteroscedastic uncertainty estimation

Thus far, the formulation implies a single objective (Eq. 3) for the optimization of both displacement and variance estimators. Even with perfectly estimated variance, such naïve application may result in undersampling of higher intensity regions [22], as their noise levels tend to be elevated due to the relatively preserved SNR throughout the image space. To mitigate this, we propose a learning strategy that utilizes separate objectives for the two modules, but with collaborative information exchange (Fig. 1). Specifically, the displacement estimator is optimized using Eq. 4, which largely depends on the quality of the predicted variance. In parallel, the variance estimator directly uses the displacement estimator output (i.e. the warped image) as a part of its input to predict the variance.




Fig. 1: Left: We propose a heteroscedastic uncertainty estimation scheme to adaptively weight the data-fidelity term accounting for the non-uniform variations of noise across the image. Right: Overview of our proposed method. The noise variance estimator uses a U-Net backbone that takes reconstructed frame $\hat{I}_f$ along with frame I_f to predict the heteroscedastic variance for the noise in [Eq. 2](#).

4.2 Adaptive weighting using exponentiated relative SNR

Our proposed formulation above offers the flexibility of designing separate objectives for the displacement and variance estimators. For the displacement loss, we addressed the undersampling issue by modifying the image fidelity weighting term from the inverse of *absolute* variance to the γ -exponentiated *relative* SNR. Our proposed displacement loss is defined as

$$\mathcal{L}_\theta = \mathbb{E}_\Omega \left[\mathcal{T} \left[\left(\frac{I_f}{\lfloor \hat{\sigma}_I \rfloor} \right)^{2\gamma} \right] \|I_f - \hat{I}_f\|_2^2 + \lambda \|\nabla \hat{z}\|^2 \right], \quad (4)$$

where $\mathcal{T}(\cdot)$ is the sigmoid activation function. The hyperparameter γ indicates the confidence of the displacement model on the estimated variance; when $\gamma = 0$ the image fidelity term reduces to MSE. Gradient stopping (denoted as $\lfloor \cdot \rfloor$) was necessary to treat the weighting term as a scalar map and to prevent duplicate back-propagation by the two estimators.

For optimizing the variance estimator, we used the β -NLL objective [\[22\]](#)

$$\mathcal{L}_\phi = \mathbb{E}_\Omega \left[\lfloor \hat{\sigma}_I^{2\beta} \rfloor \left(\frac{1}{\hat{\sigma}_I^2} \|I_f - \lfloor \hat{I}_f \rfloor\|_2^2 + \log \hat{\sigma}_I^2 \right) \right]. \quad (5)$$

Our network was implemented to output $\log \hat{\sigma}_I^2$ for numerical stability. We updated each estimator in an alternating fashion using two separate optimizers.

5 Datasets and results

5.1 Datasets

We tested our method on three distinctive cardiac datasets: (1) **ACDC** [\[4\]](#): Human MRI, multiple 2D slices + time, 150 total sequences. (2) **CAMUS** [\[14\]](#):

Table 1: Contour-based metrics compared against baselines. Units: DSC (%) HD (px) ASD (px). Our method consistently improves across different architectures and datasets. The second-best method is highlighted with †.

		ACDC [4]			CAMUS [14]		
		DSC ↑	HD ↓	ASD ↓	DSC ↑	HD ↓	ASD ↓
CNN	Undeformed	47.98	7.91	2.32	66.77	10.87	2.61
	Elastix [13]	77.26	4.95	1.28	80.18	10.02	1.81
	vxm (NCC) [3]	78.55	4.94	1.29	77.01	10.23	1.89
	vxm (MI) [3]	78.04	5.25	1.35	78.18	9.83	1.99
	vxm (MSE) [3] †	80.20	4.64	1.24	81.76	8.93	1.70
	NLL [12]	76.49	5.46	1.45	75.24	11.05	2.20
	β-NLL [22]	78.74	5.07	1.33	79.75	9.39	1.93
	AdaFrame [27]	66.38	5.80	1.67	77.88	10.54	1.93
	AdaReg [28]	78.75	5.13	1.33	79.31	9.78	1.88
	Ours	80.73	4.57	1.21	81.96	8.80	1.66
Transformer	tsm (NCC) [5]	73.77	6.64	1.12	73.03	11.87	1.70
	tsm (MI) [5]	73.57	6.57	1.11	74.83	11.94	1.83
	tsm (MSE) [5] †	76.94	5.51	1.30	79.24	10.30	1.79
	NLL [12]	73.12	7.22	1.27	75.08	11.60	1.79
	β-NLL [22]	75.74	6.12	1.29	77.39	10.99	1.86
	AdaFrame [27]	67.95	5.72	1.59	78.06	9.86	1.91
	AdaReg [28]	76.22	5.68	1.29	78.12	10.62	1.84
	Ours	78.12	5.04	1.26	80.38	9.86	1.72

	Voxelmorph	Transmorph
ACDC	$p = 0.02$	$p = 1.02 \times 10^{-41}$
CAMUS	$p = 1.36 \times 10^{-6}$	$p = 0.005$

Table 3: Quantitative evaluation on 3D private Echo dataset using vxm-based architecture.

Private 3D Echo			
	DSC ↑	HD ↓	ASD ↓
Voxelmorph [3]	74.61	5.59	0.94
NLL [12]	72.59	6.35	0.97
β-NLL [22]	73.81	5.77	0.96
AdaFrame [27]	73.01	5.74	0.94
AdaReg [28]	73.41	5.91	0.95
Ours	75.04	5.55	0.93

Table 4: Effect of γ in Eq. 4 for $\hat{\sigma}_I^2$ estimation.

	ACDC [4]			CAMUS [14]		
	DSC ↑	HD ↓	ASD ↓	DSC ↑	HD ↓	ASD ↓
Ours ($\gamma = 0.25$)	79.74	4.74	1.26	82.07	8.53	1.65
Ours ($\gamma = 0.5$)	80.73	4.57	1.21	81.96	8.80	1.66
Ours ($\gamma = 0.75$)	80.00	4.69	1.24	81.82	8.45	1.66
Ours ($\gamma = 1$)	79.78	4.71	1.25	81.31	9.08	1.69

Human echocardiography, 2D + time, 1000 total sequences from 500 patients, (3) **Private 3D Echo**: *in vivo* Porcine, *in vivo* canine, and synthetic echocardiography, 3D + time, 99 total sequences. All datasets included segmentations of the left ventricular myocardium at end-diastole (ED) and end-systole (ES). For all of our experiments, we warped the ED frame to reconstruct the ES frame. The training, validation, and testing splits were 60/20/20(%). Images were pre-processed with resizing and normalization (Supp. Mat.). Implementation details (hyperparameters, compute, etc.) are in the Supp. Mat.

5.2 Registration accuracy

Quantitative and qualitative evaluations. We present extensive quantitative evaluations of our method performance, where we show a consistent improvement from both classical and deep learning baselines across all datasets (Table 1, 3) with statistical significance (Table 2). Qualitative visualizations also demonstrate improved registration performance of our approach with smoother contour edges and more locally consistent myocardium (Fig. 2).

Our consistent improvements from the vanilla version of Voxelmorph (vxm) [3] and Transmorph (tsm) [5], which both utilize homoscedastic noise assumptions, especially confirm the benefit of heteroscedastic noise modeling (Table 1).




Fig. 2: Qualitative evaluation of the registration accuracy via segmentation warping for all datasets (top two rows: Voxelmorph architecture [3], bottom two rows: Transmorph architecture [5]). Our method in the last column (overlayed with ground truth (GT) ES myocardium label in yellow) predicts more natural and accurate deformations compared to baselines, evidenced by better matching with the GT, smoother contour edges, and locally consistent myocardial region.

We note that both NLL [12] and β -NLL failed to improve upon the vanilla baselines (Table 1). This validates our analysis that the joint training of displacement and variance estimators would degrade registration performance due to the undersampling from inverse absolute variance weighting (Section 4). This also shows the advantage of our proposed collaborative learning framework, which provides the flexibility of designing separate objectives for the two estimators. Furthermore, we observe that existing adaptive weighting schemes of AdaReg [28] and AdaFrame [27] are ineffective, potentially due to their assumptions on error residuals failing to represent the complicated real-world data distribution; this is in contrast to our proposed data-driven SNR-based weighting scheme. For 3D echo data, we show quantitative results in Table 3 and qualitative results in the Supp. Mat. Note: Improvements come solely from smoother optimization and does not incur additional complexity during inference.

Effect of γ . The exponentiated hyperparameter γ provides the flexibility to adjust the estimated variance’s degree of influence on the displacement estimator during training. This can be viewed as a trade-off in the predicted uncertainty es-




Fig. 3: Left: Estimated $\hat{\sigma}_f^2$ and the corresponding weighting map of (top row: ACDC [4]; bottom row: CAMUS [14]). Right: Sparsification error plots of $\hat{\sigma}_f^2$. Both plots are from our proposed framework under Voxelmorph architecture [3]

timates, where $\gamma = 0$ means no confidence, reducing the adaptive weighting map to a uniform scalar map, and $\gamma = 1$ indicating full confidence. We empirically chose $\gamma = 0.5$ for the best performance across all datasets (Table 4).

5.3 Evaluation on heteroscedastic uncertainty

We utilize sparsification error plots [20] to provide quantitative measure of accuracy for our estimated variance $\hat{\sigma}_f^2$. To obtain the plots, we removed one pixel at a time from largest to smallest uncertainty magnitudes, and measured the MSE of the remaining pixels ($\mathbb{E}_{\Omega - \Omega_x} [I_f - \hat{I}_f]^2$) [25]. An ideal sparsification error plot should monotonically decrease, which would indicate that the estimated uncertainty map is able to correctly identify pixels with the largest errors. Fig. 3 shows our estimated uncertainty is sensible across different datasets – based on the overall shape of the plots and the Area Under the Sparsification Error (AUSE) metrics – with similar calibration levels compared to β -NLL [22] and better calibration than NLL [12]. This illustrates the effectiveness of our variance estimator and optimization framework.

We further provide a qualitative visualization on both datasets, where we demonstrate that our estimated $\hat{\sigma}_f$ provides a sensible heteroscedastic noise variance map between the fixed image and reconstructed image according to our assumption (Fig. 3). From the figure, our estimated noise variance shown in the third column reflects the intensity mismatches of the corresponding regions between the fixed image I_f in the first column and our reconstructed/moved image ($\hat{I}_f$) in the second column. This corroborates the validity of our heteroscedastic noise assumption (Eq. 2) and shows the effectiveness of our proposed variance estimator. Furthermore, our computed adaptive weighting maps (last column, Fig. 3) accurately reflect the relative importance based on SNR, as evidenced by the highlighting of regions with potential imaging artifacts. Nonetheless, we do have failure cases, as shown in the Supp. Mat., where the myocardium (typically challenging due to irregular volumes) is considerably thin.

6 Conclusion

We propose a heteroscedastic uncertainty estimation framework for probabilistic unsupervised image registration to adaptively weight the displacement estimation with *relative* γ -exponentiated signal-to-noise, which improves registration performance from the commonly used homoscedastic assumption while also providing accurate and sensible uncertainty measures. Our proposed framework consists of a displacement and a variance estimator, optimized under an alternating collaborative strategy. We demonstrate the effectiveness and versatility of our proposed framework on two representative registration architectures across diverse cardiac datasets and show consistent statistically significant improvements over other baselines. Though our proposed framework is promising, it still relies on a manually crafted adaptive map on the data fidelity term, which might not be able to fully reflect the complicated characteristics of large-scale real-world data. Future work will aim to explore a more data-driven objective with further validation of clinical datasets for more potential impact.

Acknowledgments. This work is supported by NIH grant R01HL121226.

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

References

1. Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95–113 (Oct 2007)
2. Bae, G., Budvytis, I., Cipolla, R.: Estimating and Exploiting the Aleatoric Uncertainty in Surface Normal Estimation. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 13117–13126. IEEE, Montreal, QC, Canada (Oct 2021)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Transactions on Medical Imaging* **38**(8), 1788–1800 (Aug 2019)
4. 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)
5. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: TransMorph: Transformer for unsupervised medical image registration. *Medical Image Analysis* **82**, 102615 (Nov 2022)
6. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical Image Analysis* **57**, 226–236 (Oct 2019)
7. Eklund, A., Lindquist, M.A., Villani, M.: A Bayesian heteroscedastic GLM with application to fMRI data with motion spikes. *NeuroImage* **155**, 354–369 (Jul 2017)
8. Hill, D.L.G., Batchelor, P.G., Holden, M., Hawkes, D.J.: Medical image registration

9. Hoffmann, M., Billot, B., Greve, D.N., Iglesias, J.E., Fischl, B., Dalca, A.V.: SynthMorph: learning contrast-invariant registration without acquired images. *IEEE Transactions on Medical Imaging* **41**(3), 543–558 (Mar 2022), arXiv:2004.10282 [cs, eess, q-bio]
10. Hong, B.W., Koo, J.K., Burger, M., Soatto, S.: Adaptive Regularization of Some Inverse Problems in Image Analysis (May 2017), arXiv:1705.03350 [cs]
11. Hoopes, A., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.V.: HyperMorph: Amortized Hyperparameter Learning for Image Registration (May 2021), arXiv:2101.01035 [cs, eess]
12. Kendall, A., Gal, Y.: What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? (Oct 2017), arXiv:1703.04977 [cs]
13. Klein, S., Staring, M., Murphy, K., Viergever, M., Pluim, J.: elastix: A Toolbox for Intensity-Based Medical Image Registration. *IEEE Transactions on Medical Imaging* **29**(1), 196–205 (Jan 2010)
14. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., et al.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE transactions on medical imaging* **38**(9), 2198–2210 (2019)
15. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. pp. 10012–10022 (2021)
16. Ma, T., Dai, X., Zhang, S., Wen, Y.: Pivit: Large deformation image registration with pyramid-iterative vision transformer. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 602–612. Springer (2023)
17. Monteiro, M., Le Folgoc, L., Coelho de Castro, D., Pawlowski, N., Marques, B., Kamnitsas, K., van der Wilk, M., Glocker, B.: Stochastic Segmentation Networks: Modelling Spatially Correlated Aleatoric Uncertainty. In: *Advances in Neural Information Processing Systems*. vol. 33, pp. 12756–12767. Curran Associates, Inc. (2020)
18. Oliveira, F.P.M.: Medical Image Registration: a Review
19. Ouzir, N., Ollila, E., Vorobyov, S.A.: Data-Adaptive Similarity Measures for B-mode Ultrasound Images Using Robust Noise Models. *IEEE Journal of Selected Topics in Signal Processing* **14**(6), 1244–1254 (Oct 2020), conference Name: IEEE Journal of Selected Topics in Signal Processing
20. Poggi, M., Aleotti, F., Tosi, F., Mattoccia, S.: On the Uncertainty of Self-Supervised Monocular Depth Estimation. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3224–3234. IEEE, Seattle, WA, USA (Jun 2020)
21. Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D.: Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging* **18**(8), 712–721 (Aug 1999)
22. Seitzer, M., Tavakoli, A., Antic, D., Martius, G.: On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks (Apr 2022), arXiv:2203.09168 [cs, stat]
23. Shi, J., He, Y., Kong, Y., Coatrieux, J.L., Shu, H., Yang, G., Li, S.: XMorpher: Full Transformer for Deformable Medical Image Registration via Cross Attention (Jun 2022), arXiv:2206.07349 [cs]
24. Ta, K., Ahn, S.S., Thorn, S.L., Stendahl, J.C., Zhang, X., Langdon, J., Staib, L.H., Sinusas, A.J., Duncan, J.S.: Multi-task learning for motion analysis and segmentation in 3d echocardiography. *IEEE Transactions on Medical Imaging* (2024)

25. Truong, P., Danelljan, M., Van Gool, L., Timofte, R.: Learning Accurate Dense Correspondences and When to Trust Them (Apr 2021), arXiv:2101.01710 [cs]
26. Wegmann, B., Eklund, A., Villani, M.: Bayesian heteroscedastic regression for diffusion tensor imaging. In: *Modeling, Analysis, and Visualization of Anisotropy*. pp. 257–282. Springer (2017)
27. Wong, A., Fei, X., Hong, B.W., Soatto, S.: An Adaptive Framework for Learning Unsupervised Depth Completion. *IEEE Robotics and Automation Letters* **6**(2), 3120–3127 (Apr 2021)
28. Wong, A., Soatto, S.: Bilateral Cyclic Constraint and Adaptive Regularization for Unsupervised Monocular Depth Prediction. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5637–5646. IEEE, Long Beach, CA, USA (Jun 2019)
29. Zhang, J., He, Q., Xiao, Y., Zheng, H., Wang, C., Luo, J.: Ultrasound image reconstruction from plane wave radio-frequency data by self-supervised deep neural network. *Medical Image Analysis* **70**, 102018 (May 2021)
30. Zhang, X., Dong, H., Gao, D., Zhao, X.: A comparative study for non-rigid image registration and rigid image registration. arXiv preprint arXiv:2001.03831 (2020)
31. Zhang, X., Noga, M., Martin, D.G., Punithakumar, K.: Fully automated left atrium segmentation from anatomical cine long-axis mri sequences using deep convolutional neural network with unscented kalman filter. *Medical image analysis* **68**, 101916 (2021)
32. Zhang, X., You, C., Ahn, S., Zhuang, J., Staib, L., Duncan, J.: Learning Correspondences of Cardiac Motion from Images Using Biomechanics-Informed Modeling. In: Camara, O., Puyol-Antón, E., Qin, C., Sermesant, M., Suinesiaputra, A., Wang, S., Young, A. (eds.) *Statistical Atlases and Computational Models of the Heart. Regular and CMRxMotion Challenge Papers*. pp. 13–25. Lecture Notes in Computer Science, Springer Nature Switzerland, Cham (2022)