Recursive Uncertainty-Gated Image Registration for Learning-Based Algorithms

Clara Rodrigo González, Oscar Bates, Fu Siong Ng, Meng-Xing Tang

Imperial College London

Abstract. Typically, learning-based image registration identifies a deformation field deterministically in a single pass through the neural network, with no test-time adaption or quantification of confidence in the prediction. We propose Recursive Uncertainty-Gated Image Registration (RUGI), a model-agnostic, inference-time algorithm, to recursively refine an initial deformation field while focusing updates on regions of high uncertainty. RUGI produces uncertainty estimates which can be propagated to downstream tasks, and improves registration performance. We evaluate on two datasets spanning cardiac MRI and ultrasound, and demonstrate statistically significant improvements over single-step inference. Mean squared error (MSE) reductions exceed 20% relative to the single-step baseline, with gains in anatomical alignment confirmed by Dice similarity within labelled datasets. RUGI can also be applied directly to existing pretrained models; applied to VoxelMorph, TransMorph, and CycleMorph in a third dataset, it yields MSE reductions of 27–37% with no modification to the original training procedure. Therefore, we expect RUGI to improve performance and drive interpretability in any learning-based image registration scenario.

Keywords: Deformable Image Registration · Uncertainty Estimation · Inference-Time Adaptation · Bayesian Neural Networks · Cardiac Imaging

1 Introduction

Image registration is a fundamental task in medical image analysis, with applications in diagnostics, treatment planning, and image-guided intervention. In recent years, learning-based image registration methods have gained substantial attention due to their speed at inference and their performance. Typically, a neural network receives a fixed and moving image pair and estimates a dense displacement field that can be used to warp one image to match the other. Learning-based approaches enable fast inference at test time and can learn complex, data-driven regularisation priors, and are able to generalize across image pairs and, in some cases, across datasets without a need for further training [1, 2].

The estimation of spatial correspondence between images is inherently ill-posed, and its difficulty is often spatially varying: image quality may vary substantially across the field of view, making some correspondences easier to infer than others. Some regions may exhibit simple, smooth motion, whereas others undergo larger or more complex local deformations, making the transformation spatially heterogeneous. These factors make registration performance highly variable across datasets, with the optimal method depending strongly on the spatial heterogeneity and motion characteristics of the data [3]. This sensitivity to data characteristics limits the ability of fixed, trained models to generalize reliably to new domains, motivating inference-time strategies that can adapt to the specific properties of each image pair.

Iterative registration strategies have been applied in the past [4–8]. Repeated inference expands the transformations that can be estimated by the model, but updates are often applied uniformly across the full image domain, regardless of the spatial heterogeneity or the confidence in the predictions.

Uncertainty in learning-based registration arises from two sources: aleatoric uncertainty (due to data variability, noise, artefacts) and epistemic uncertainty (due to model structure and parameters) [2]. Probabilistic frameworks such as variational inference [9] and conditional VAEs [10] model aleatoric uncertainty by estimating a distribution over deformations, but require specialised architectures. Epistemic uncertainty is estimated via Monte Carlo dropout [11], ensembles [12], or Bayesian neural networks [13, 14]. Uncertainty has been used to weigh regularisation terms [15], and Tian et al. [16] estimate it at inference time via transformation equivariance, but no prior work combines uncertainty estimation with iterative refinement to guide focused updates at inference.

We propose a learning-based image registration framework at inference which combines focused optimisation with epistemic uncertainty estimation to improve registration accuracy and increase interpretability. Our contributions are:

1. We present an inference-time algorithm to iteratively refine image registration using uncertainty gating. Our method performs iterative inference and uses an uncertainty-based gating function to apply targeted updates to different regions of the image. Our algorithm provides uncertainty estimates that can be propagated to downstream tasks.
2. We show the method can be applied to pretrained models with no need for retraining by using an image similarity-based proxy for the uncertainty to guide updates.

2 Methods

For a pair of images I_{mov} and I_{fix}, the goal is to find a deformation field ϕ which warps I_{mov} to match the underlying structure of I_{fix}. In unsupervised motion estimation, a model is trained to predict this deformation field by minimizing a loss function between the warped moving image I_{warped} and I_{fix}. At inference, this model receives the input images and estimates the deformation field in one step. This deformation field is then used to warp the moving image.

Algorithm 1 RUGI Algorithm

Require: Fixed image I_{fix}, moving image I_{mov}, max iterations T, tolerance ϵ
Ensure: Composed deformation field $\boldsymbol{\Phi}$

```
Φ ← identity
I^(0) ← I_mov
q_0 ← Sim(I_fix, I^(0))
for t = 1, 2, ..., T do
    φ^(t) ← R(I_fix, I^(t−1))                ▷ Predict deformation
    s^(t) ← S(I_fix, I^(t−1))                ▷ Calculate gating
    if t = 1 then
        φ̂^(t) ← φ^(t)                        ▷ First step: unmodulated update
    else
        φ̂^(t) ← φ^(t) ⊙ s^(t)                ▷ Apply uncertainty gating
    end if
    Φ_new ← Φ ∘ φ̂^(t)                        ▷ Tentative composition
    I^(t) ← I_mov ∘ Φ_new                    ▷ Warp with new field
    q_t ← Sim(I_fix, I^(t))
    if q_t − q_{t−1} < ϵ then                ▷ Insufficient improvement, stop
        break
    end if
    Φ ← Φ_new                                ▷ Accept update
end for
Return Φ
```

2.1 Recursive Uncertainty-Gated Image Registration (RUGI)

Algorithm 1 describes the Recursive Uncertainty-Gated Image Registration (RUGI) Algorithm. At each iteration t, the registration model $R()$ predicts a dense deformation field over the full image domain, conditioned on the fixed image I_{fix} and the moving image warped at the previous step, $I^{(t-1)}$. The model outputs a vector flow field for that time point $\underline{\phi}^{(t)} = R(I_{fix}, I^{(t-1)})$. In parallel, the gating function $S()$ produces an uncertainty estimate for the deformation field $s^{(t)} = S(I_{\text{fix}}, I^{(t-1)})$, which is normalized for each image. The uncertainty map modulates the magnitude of the incremental deformation update at each iteration. The selective gating is anisotropic, and is calculated and applied independently along each axis $s_x^{(t)}$ and $s_z^{(t)}$. This anisotropic formulation allows the method to modulate registration confidence independently per direction, which is particularly important in settings where noise and artefact characteristics differ across axes. This process can be repeated until a loss-based or a higher-is-better similarity-based stopping criterion (as written in Algorithm 1) is satisfied.

To prevent the accumulation of interpolation artefacts that would arise from repeated resampling, the deformation fields $\hat{\underline{\phi}}^{(t)}$ at each iteration are not applied sequentially to the output of the previous step. Instead, all incremental deformation fields are composed into a single cumulative field $\boldsymbol{\Phi}$, and the original moving image I_{mov} is resampled using the composed deformation at each iteration. This reduces interpolation artefacts and preserves the image details during the optimisation.

We propose to use the uncertainty of the deformation field as the gating function. In this case, regions with low predicted uncertainty receive smaller updates and stabilize earlier in the iterative refinement, while high-uncertainty regions continue to receive updates.

While RUGI works at inference time with no modification to the registration model, the uncertainty estimation requires a lightweight network to be trained. Both networks are trained jointly, so the uncertainty network learns to predict uncertainty specific to $R()$'s predictions. During training, the registration model $R()$ is sampled N times for each image pair, producing a distribution of deformation fields $\phi_{1..N}$. We used 5 iterations. The per-pixel standard deviation of these fields σ is used to train the uncertainty model U in a supervised manner using the mean squared error (MSE) as the loss function. The mean deformation field $\bar{\phi}$ is used to warp I_{mov}, yielding I_{warped}. The similarity loss between I_{fix} and I_{warped} is used to train the registration model R. The full loss function used the MSE as the similarity term, and an added bending energy term as the regularisation loss. At inference, the uncertainty model predicts $\sigma^{(t)} = U(I_{fix}, I^{(t-1)})$. This uncertainty map is then smoothed and normalized to obtain the gating map $s^{(t)}$. While this requires an additional network, the costs are minimized by the small size of the network (465k parameters).

Since $U()$ estimates epistemic uncertainty, high values indicate regions where the model's predictions are inconsistent across multiple passes. A BNN trained on image pairs will exhibit high predictive variance in regions where the deformation is difficult to estimate consistently, due to its magnitude, complexity, or image quality. Conversely, regions of background or static anatomy are registered consistently, yielding low σ. By weighting the incremental update $\phi^{(t)}$ by the normalised uncertainty maps $s^{(t)}$, RUGI concentrates refinement effort where the initial single-step prediction is least reliable, while suppressing potentially destabilishing updates in already well-registered regions. This allows RUGI to correct the residual misalignment that single-step inference obtains.

[4, **?**,**?**,**?**] use repeated inference to expand the space of representable transformations, but apply updates uniformly across the image domain every iteration. This approach fails to account for the heterogeneous confidence of the deformations. RUGI addresses this by introducing uncertainty gating, which modulates the magnitude of the updates. Regions of low uncertainty receive attenuated updates and stabilize early, while high-uncertainty regions continue to to be refined. The non-uniform, spatially selective refinement is a fundamental departure from prior iterative approaches.

Uncertainty-aware registration methods [9, **?**,**?**,**?**] have quantified uncertainty as an additional output to their registration frameworks, or have used it to reweight regularisation terms, but do not use uncertainty to selectively refine registration at inference time. RUGI is the first method to combine epistemic uncertainty estimation with iterative inference, with spatially gated refinement that focuses updates in regions of high uncertainty. This is achieved in a model-agnostic manner with no changes to the architecture or training procedure.

2.2 Implementation details

Registration model We use a Bayesian Neural Network (BNN) as the registration model to obtain deformation estimates and train the uncertainty network $U()$.

The BNN initialization parameters and dropout rate were empirically selected through hyperparameter sweeps.

Datasets We evaluated our selective gating algorithm on two datasets. First, the Automatic Cardiac Diagnosis Challenge (ACDC) dataset, which contains cardiac magnetic resonance acquisitions from 1.5T and 3.0T MRIs [17]. The total training dataset has 100 patients with two volumes per patient, and 560 training image pairs. Cardiac Acquisitions for Multi-structure Ultrasound Segmentation (CAMUS) is an ultrasound dataset containing cardiac acquisitions from 500 healthy and anomalous patients in multiple views [18]. Each acquisition was segmented by a clinical expert, providing ground truth anatomical information. 4500 images were used for training. We trained the models using five different random seeds on each dataset, and the sample size for statistical analysis was one per training run (n=5). Train:test:validation splits were performed by dividing across patients so no patient was in both, using ratios of 70:20:10.

Pretrained models To validate the application of our algorithm without retraining, we applied the method to several publicly available pretrained models, namely TransMorph, VoxelMorph, CycleMorph, with pretrained weights extracted from [19], in an additional third dataset. Motivated by the well-established correspondence between uncertainty and prediction error, we propose an error-based gating alternative that requires no auxiliary network. The per-pixel local normalized cross-correlation (LNCC) between I_{fix} and $I^{(t-1)}$ is computed, exponentiated to amplify contrast, smoothed and normalized to produce $s^{(t)}$. This yields a spatially varying gating that acts as an uncertainty proxy and concentrates updates in regions of high error. After RUGI, the standard deviation of the similarity-based proxy is taken as an empirical measurement of uncertainty. These models were trained on the IXI dataset, a 3-D brain MRI registration dataset [20]. For each model, we compared the original one-step inference result with the output obtained when the same model was used within our iterative gating framework.

3 Results

Figure 1 shows example registrations across the four datasets for our proposed methods, as well as a baseline single-step BNN. Our method obtains improved alignment of tissue boundaries in all cases, as seen in the difference images and the masks. High-uncertainty regions correspond to areas of larger or complex deformation. In ACDC anc CAMUS, uncertainty is highest at the myocardial boundary, where motion is largest and tissue contrast is lowest. In IXI, the empirical uncertainty is highest at sulcal boundaries and ventricular margins, consistent with the large inter-subject variability in anatomical structure in these areas. Figure 2 shows the improvement according to the four metrics as a percentage change of the performance of the one-step inference. The performance improvements when applying our algorithm to the pretrained models is also shown.

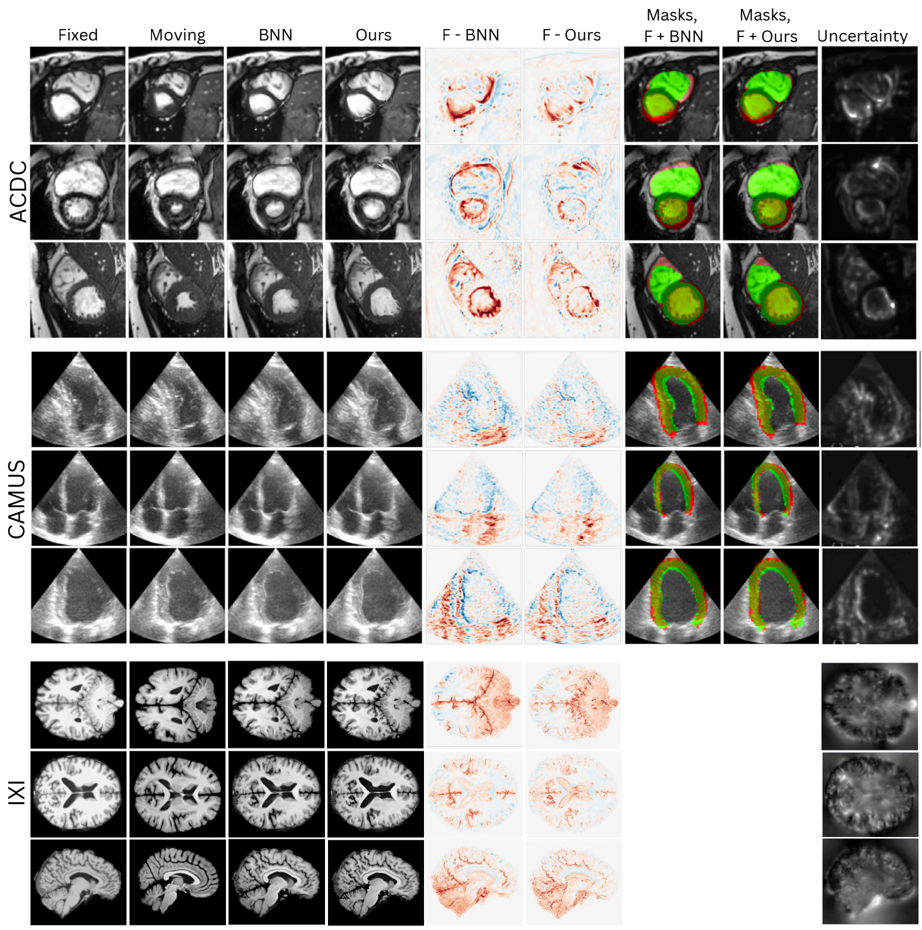

Fig. 1. Example motion corrected images using the proposed method using the uncertainty-guided iterative refinement, in comparison to the baseline BNN. Example warped images, error maps and warped masks are shown in both ACDC, CAMUS and IXI. The estimated uncertainty is shown for each example.

First, we evaluate RUGI using our proprietary models using ACDC and CAMUS. Our method obtained consistent, statistically significant improvements relative to the single-step baseline according to all four image similarity metrics. Paired one-sided Wilcoxon signed-rank tests, with Benjamini-Hochberg false discovery rate correction applied across all 40 simultaneous comparisons, yield significant results ($p_{BH} = 0.031$) in every case. We evaluated anatomical alignment directly using the Dice similarity coefficient (DSC). Our method achieved a statistically significant improvement in DSC relative to the single-step baseline (paired one-sided Wilcoxon signed-rank test, $p = 0.031$). On ACDC, the uncertainty-based method increased the median DSC across all tissues from

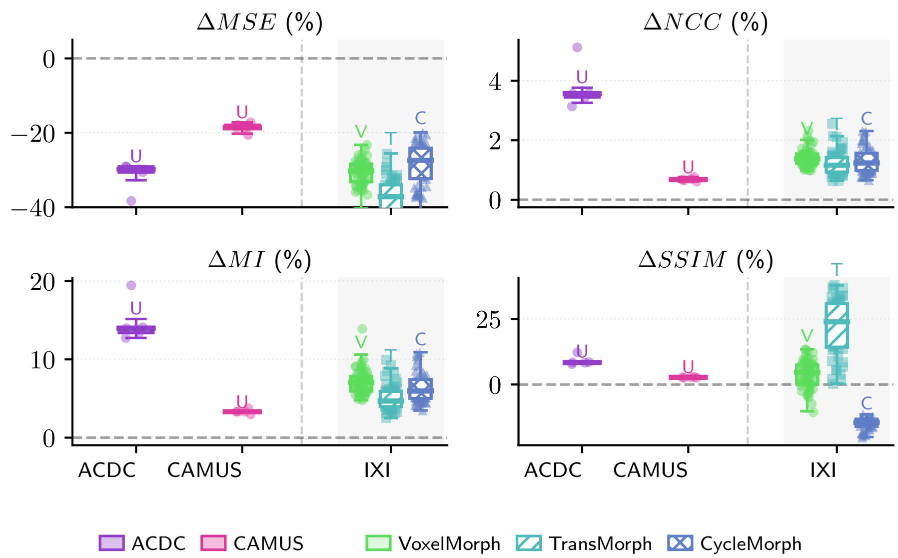

Fig. 2. Percentage change in image similarity metrics relative to one-step BNN method across the three datasets. The performance change in the IXI dataset was calculated using three pretrained models: VoxelMorph (V), TransMorph (T) and CycleMorph (C).

0.8293 to 0.8508 ($\Delta = +0.021$). On CAMUS, our method increased the median DSC from 0.813 to 0.823 ($\Delta = +0.010$). These results demonstrate that the improvements observed in intensity-based image similarity metrics correspond to a tangible gain in anatomical alignment.

Secondly, RUGI was applied at inference to three pretrained registration models, VoxelMorph, TransMorph, and CycleMorph, evaluated on 58 paired samples. Statistical comparisons used one-sided paired Wilcoxon signed-rank tests with Benjamini–Hochberg correction across all 12 simultaneous comparisons, with effect sizes signed so that positive values indicate multistep inference performing better.

Significant improvements are obtained across all three models for MSE, NCC, and MI (all $p_{\mathrm{BH}} < 0.001$, large effect sizes throughout). For MSE, median reductions of 30.2%, 37.1%, and 27.4% are observed for VoxelMorph, TransMorph, and CycleMorph respectively, with Cliff's $\delta \geq 0.84$ in all cases, indicating that the multistep approach produces lower MSE than the one-step baseline on the large majority of individual samples. NCC improvements are similarly consistent across all three models ($\delta \geq 0.94$), and MI improvements are large and significant for all three ($\delta \geq 0.80$). SSIM results are significant for VoxelMorph and TransMorph ($p_{\mathrm{BH}} < 0.001$), both showing improvements, while CycleMorph shows a significant decrease in SSIM ($p_{\mathrm{BH}} < 0.001$, $\delta = 0.72$). The divergence on SSIM for CycleMorph likely reflects differences in the spatial frequency characteristics of that model's output relative to the other two. Overall, these results demonstrate that RUGI generalises effectively to pretrained models without any retraining, producing consistent and significant improvements in the majority of metrics across all three architectures tested.

3.1 Computational Cost

Table 1 reports the median inference time per sample and median number of registration steps for our proposed method relative to the single-step BNN baseline. The multistep inference is significantly slower than the one-step baseline in all cases (BH-corrected Wilcoxon, $p < 0.001$ throughout), which is expected given that it performs multiple forward passes. The uncertainty method takes between 5.1× and 9.8× longer than the baseline depending on the dataset, with a median of 4–11 steps. The variability in the number of steps of the proposed method shows the adaptability of the method, which takes more steps when the registration problem is harder.

Dataset	One-step BNN (s)	Ours (s)	Steps	Ratio (Ours / BNN)
ACDC	0.014	0.135***	8.0	9.8×
CAMUS	0.014	0.071***	4.0	5.1×

Table 1. Computational cost of multistep inference relative to the one-step baseline. Time is median wall-clock inference time per sample (seconds). Steps is the median number of registration steps taken. Ratio = multistep / one-step time. Significance (BH-corrected Wilcoxon): $^{*}p$<0.05, $^{**}p$<0.01, $^{***}p$<0.001.

4 Discussion

This study proposed an uncertainty-gated recursive algorithm for deformation field refinement at inference. The results demonstrated that our algorithm improves learning-based image registration across diverse datasets and architectures.

RUGI produced statistically significant improvements over single-step inference across all datasets and all four image similarity metrics. This breadth of improvement suggests the method is not tailored to a specific modality or motion regime, but captures a general weakness of single-step feed-forward registration: the inability to correct residual misalignment after a single forward pass. An ablation study confirmed these improvements are not solely due to iterative refinement: uniform and intensity-gated refinement produced smaller gains than the uncertainty-gated approach, indicating spatial selectivity is a primary driver of improvement.

While there are consistent improvements in registration accuracy according to image-based and anatomical metrics, the effect sizes are modest. Further analysis should explore the potential impact of these improvements on clinical applications.

Several limitations should be noted. First, the uncertainty model requires supervised training using samples from the registration model, which adds complexity to the training pipeline and ties the uncertainty estimates to the quality of the base model. Second, the iterative inference increases wall-clock time substantially, which may be prohibitive in real-time clinical settings such as intraoperative guidance. Third, while the method was evaluated on two diverse datasets, all experiments used 2D registration; the extension to 3D (demonstrated only for the pretrained model experiments) may introduce additional computational constraints. Fourth, the SSIM inconsistency for CycleMorph and the moderate performance gap between methods on specific dataset–metric combinations suggest that the method's behaviour is not fully uniform, and failure cases remain to be characterised systematically. Finally, the stopping criterion is based on image similarity improvement, which may not always align with anatomical registration accuracy; incorporating landmark-based or segmentation-based stopping criteria could improve reliability in settings where ground truth deformation is partially available.

5 Conclusion

We proposed RUGI, an algorithm for uncertainty-gated deformation field refinement at inference-time. By focusing updates on regions of the images with high uncertainty, our method is able to adapt to spatially heterogeneous difficulty. The resulting uncertainty maps provide spatially interpretable confidence estimates that correlate with registration error and can be propagated to downstream tasks. Statistically significant improvements were demonstrated across cardiac MRI and ultrasound datasets, with improvements in intensity-based metrics and anatomical alignment. The framework generalises to existing pretrained architectures without retraining, yielding MSE reductions of 27–37% across VoxelMorph, TransMorph, and CycleMorph. This work demonstrates the value of uncertainty estimation not only to improve interpretability but also to drive performance in medical imaging tasks.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, and Adrian V. Dalca. VoxelMorph: A Learning Framework for Deformable Medical Image Registration. *IEEE Transactions on Medical Imaging*, 38(8):1788–1800, August 2019. arXiv:1809.05231 [cs].

2. Junyu Chen, Yihao Liu, Shuwen Wei, Zhangxing Bian, Shalini Subramanian, Aaron Carass, Jerry L. Prince, and Yong Du. A survey on deep learning in medical image registration: New technologies, uncertainty, evaluation metrics, and beyond. *Medical Image Analysis*, 100:103385, February 2025.
3. Clara Rodrigo González, Biao Huang, Su Yan, Rifkat Zaydullin, Marcelo Lerendegui, Cameron A. B. Smith, Matthieu Toulemonde, Megan Morris, Navita Somaiah, Oscar Bates, Fu Siong Ng, and Meng-Xing Tang. Benchmarking Image-Based Motion-Correction Methods for Ultrasound Localization Microscopy. *Ultrasound in Medicine & Biology*, May 2026.
4. Tony C. W. Mok and Albert C. S. Chung. Large Deformation Diffeomorphic Image Registration with Laplacian Pyramid Networks, June 2020.
5. Lei Xu, Ping Jiang, Tiffany Tsui, Junyan Liu, Xiping Zhang, Lequan Yu, and Tianye Niu. 4D-CT deformable image registration using unsupervised recursive cascaded full-resolution residual networks. *Bioengineering & Translational Medicine*, 8(6):e10587, August 2023.
6. Shengyu Zhao, Yue Dong, Eric I.-Chao Chang, and Yan Xu. Recursive Cascaded Networks for Unsupervised Medical Image Registration, July 2019.
7. Robin Sandkühler, Simon Andermatt, Grzegorz Bauman, Sylvia Nyilas, Christoph Jud, and Philippe C. Cattin. Recurrent Registration Neural Networks for Deformable Image Registration. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
8. Bo Hu, Shenglong Zhou, Zhiwei Xiong, and Feng Wu. Recursive Decomposition Network for Deformable Image Registration. *IEEE Journal of Biomedical and Health Informatics*, 26(10):5130–5141, October 2022.
9. Adrian V. Dalca, Guha Balakrishnan, John Guttag, and Mert R. Sabuncu. Unsupervised Learning for Fast Probabilistic Diffeomorphic Registration. volume 11070, pages 729–738. 2018. arXiv:1805.04605 [cs].
10. Julian Krebs, Hervé Delingette, Boris Mailhé, Nicholas Ayache, and Tommaso Mansi. Learning a Probabilistic Model for Diffeomorphic Registration. *IEEE Transactions on Medical Imaging*, 38(9):2165–2176, September 2019.
11. Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 1050–1059. PMLR, June 2016.
12. Sangdaow Noppitak and Olarik Surinta. dropCyclic: Snapshot Ensemble Convolutional Neural Network Based on a New Learning Rate Schedule for Land Use Classification. *IEEE Access*, 10:60725–60737, 2022.
13. Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer Science & Business Media, December 2012. Google-Books-ID: LHHrBwAAQBAJ.
14. Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight Uncertainty in Neural Network. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1613–1622. PMLR, June 2015.
15. Zhe Xu, Jie Luo, Donghuan Lu, Jiangpeng Yan, Sarah Frisken, Jayender Jagadeesan, William M. Wells, Xiu Li, Yefeng Zheng, and Raymond Kai-yu Tong. Double-Uncertainty Guided Spatial and Temporal Consistency Regularization Weighting for Learning-Based Abdominal Registration. In Linwei Wang, Qi Dou, P. Thomas Fletcher, Stefanie Speidel, and Shuo Li, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, pages 14–24, Cham, 2022. Springer Nature Switzerland.
16. Lin Tian, Xiaoling Hu, and Juan Eugenio Iglesias. Uncertainty Estimation for Pretrained Medical Image Registration Models via Transformation Equivariance, 2025. Version Number: 2.

17. Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, Gerard Sanroma, Sandy Napel, Steffen Petersen, Georgios Tziritas, Elias Grinias, Mahendra Khened, Varghese Alex Kollerathu, Ganapathy Krishnamurthi, Marc-Michel Rohe, Xavier Pennec, Maxime Sermesant, Fabian Isensee, Paul Jager, Klaus H. Maier-Hein, Peter M. Full, Ivo Wolf, Sandy Engelhardt, Christian F. Baumgartner, Lisa M. Koch, Jelmer M. Wolterink, Ivana Isgum, Yeonggul Jang, Yoonmi Hong, Jay Patravali, Shubham Jain, Olivier Humbert, and Pierre-Marc Jodoin. 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, November 2018.
18. Sarah Leclerc, Erik Smistad, Joao Pedrosa, Andreas Ostvik, Frederic Cervenansky, Florian Espinosa, Torvald Espeland, Erik Andreas Rye Berg, Pierre-Marc Jodoin, Thomas Grenier, Carole Lartizien, Jan Dhooge, Lasse Lovstakken, and Olivier Bernard. Deep Learning for Segmentation Using an Open Large-Scale Dataset in 2D Echocardiography. *IEEE transactions on medical imaging*, 38(9):2198–2210, September 2019.
19. Lin Tian, Hastings Greer, Roland Kwitt, Francois-Xavier Vialard, Raul San Jose Estepar, Sylvain Bouix, Richard Rushmore, and Marc Niethammer. uniGradICON: A Foundation Model for Medical Image Registration, March 2024. arXiv:2403.05780 [cs].
20. IXI Dataset – Brain Development.