



HHS Public Access

Author manuscript

Med Image Comput Comput Assist Interv. Author manuscript; available in PMC 2023 August 11.

Published in final edited form as:

Med Image Comput Comput Assist Interv. 2022 September ; 13436: 66–77.

doi:10.1007/978-3-031-16446-0_7

ContraReg: Contrastive Learning of Multi-modality Unsupervised Deformable Image Registration

Neel Dey¹, Jo Schlemper², Seyed Sadegh Mohseni Salehi², Bo Zhou³, Guido Gerig¹, Michal Sofka²

¹Department of Computer Science & Engineering, New York University, Brooklyn, NY, USA.

²Hyperfine Research, Guilford, CT, USA.

³Department of Biomedical Engineering, Yale University, New Haven, CT, USA.

Abstract

Establishing voxelwise semantic correspondence across distinct imaging modalities is a foundational yet formidable computer vision task. Current multi-modality registration techniques maximize hand-crafted inter-domain similarity functions, are limited in modeling nonlinear intensity-relationships and deformations, and may require significant re-engineering or underperform on new tasks, datasets, and domain pairs. This work presents ContraReg, an unsupervised contrastive representation learning approach to multi-modality deformable registration. By projecting learned multi-scale local patch features onto a jointly learned inter-domain embedding space, ContraReg obtains representations useful for non-rigid multi-modality alignment. Experimentally, ContraReg achieves accurate and robust results with smooth and invertible deformations across a series of baselines and ablations on a neonatal T1–T2 brain MRI registration task with all methods validated over a wide range of deformation regularization strengths.

1. Introduction

The spatial alignment (or *registration*) of images from sources capturing distinct anatomical characteristics enables well-informed biomedical decision-making via multi-modality information integration. For example, multi-modality registration of intra-operative to pre-operative imaging is crucial to various surgical procedures [13,24,33]. Consequently, several inter-domain image similarity functions have been developed to drive iterative or learning-based registration [11,14,38]. Yet, despite the decades-long development of multi-modality objectives, accurate deformable registration of images with highly nonlinear relationships between their appearance and shapes remains difficult.

Losses operating on intensity features (global and local 1D histograms [8, 34, 38], local descriptors [14,37,39], edge-maps [11], among others) are typically hand-crafted and do not consistently generalize outside of the domain-pair they were originally proposed for and necessitate non-trivial domain expertise to tune towards optimal results. More recent multi-modality methods based on *learned* appearance similarity [28], simulation-driven semantic similarity [15], and image translation [29] have demonstrated strong registration performance but may only apply to supervised affine registration [28], require population-

specific segmentation labels for optimal registration [15], or necessitate extensive and delicate GAN-based training frameworks [29].

To overcome these limitations, this work develops `ContraReg` (CR), an unsupervised representation learning framework for non-rigid multi-modality registration. CR requires that, once registered, corresponding positions and patches in the moved and fixed images have *high mutual information* in a jointly-learned multi-scale multi-modality embedding space. These characteristics are achieved via self-supervised contrastive training and only requires an unsupervised feature extractor (which can be pretrained or randomly-initialized and frozen). As a result, on the challenging task of neonatal inter-subject T1w–T2w diffeomorphic MRI registration, CR achieves higher anatomical overlap with comparable deformation characteristics to previous similarity metrics, validated over a wide range of regularization parameters. Finally, for experimental completeness, we then evaluate CR across a diversity of auxiliary losses and external negative sample selection and pretraining strategies.

2. Related work

Hand-crafted similarity losses.

Mutual information and its variants [19,38] typically operate on image intensity histograms and may be limited in their ability to model complex non-rigid deformations. Conversely, local mutual information methods [8,34] are spatially-aware but may not have enough samples to build accurate patch intensity histograms. Other losses operating on intensity-derived local descriptors [14,39] learn domain-invariant features and have been successful in tasks such as body cavity MR-CT registration with relatively limited adoption in neuroimaging.

Simulation.

Registration via translation approaches seek to simulate one modality from the other, such that the problem can be reduced to a mono-modality alignment [1,29,40]. While performant, recent methods can be susceptible to the hallucinatory or instability drawbacks of medical image translation leading to suboptimal alignment [20, 31]. Recently, `SynthMorph` [15] simulates both deformations and synthetic appearances for neuroimages to train a general-purpose multi-modality network. `ContraReg` instead obtains highly accurate warps for a given dataset at the cost of dataset-specific training.

Learned similarity losses.

Finally, several methods attempt to learn an inter-domain similarity metric via supervised learning with pre-aligned training images [9,18,28,36]. For example, `CoMIR` [28] uses supervised contrastive learning to learn affine registration between highly visually disparate modalities. In particular, `ContraReg` draws inspiration from `PatchNCE` [26], an image translation framework using contrastive losses on multi-scale patches. However, a straightforward extension of `PatchNCE` to registration is not possible and would lead to degenerate identity solutions to the `PatchNCE` loss as we require a warp between two distinct input images [17]. This work presents a different approach to incorporating multi-

scale patches via *externally-trained* feature extractors which enables successful registration without degenerate solutions.

3. Methods

Fig. 1 illustrates ContraReg. Architectural details are in the supplementary material.

Unsupervised pre-training.

To extract multi-scale n -dimensional features in an unsupervised manner, we first train modality-specific autoencoders A_1 and A_2 as domain-specific features can be beneficial to patchwise contrastive training [12]. Training is done on 128^3 crops with random flipping and brightness/contrast/saturation augmentation, using an \mathcal{L}_1 + Local NCC (window width = 7 voxels) reconstruction loss [2].

Registration training.

This work focuses on unsupervised learning of registration and multi-modality similarity. Given cross-modality volumes I_1 and I_2 , a VoxelMorphstyle [5] U-Net with constant channel width (ch) predicts a stationary velocity field v , which when numerically integrated with ts time steps, yields an approximately diffeomorphic displacement ϕ . We focus on bidirectional registration and obtain the inverse warp via integrating $-v$. This network is trained with a $(1 - \lambda)\mathcal{L}_{sim} + \lambda\mathcal{L}_{reg}$ objective where $\mathcal{L}_{reg} = \|v\|_2^2$ is a regularizer controlling velocity (and indirectly displacement) field smoothness, $\mathcal{L}_{sim} = \frac{1}{2}(d_{12}(I_1 \| I_2 \circ \phi) + d_{21}(I_2 \| I_1 \circ \phi^{-1}))$ is a registration term s.t. $d_{12,21}$ measures inter-domain similarity, and λ is a hyperparameter.

Here, we define d_{12} and d_{21} is defined analogously. We first extract multi-scale spatial features $A_1^k(I_1)$ and $A_2^k(I_2 \circ \phi)$ using the autoencoders, where $k = 1, \dots, L$ is the layer index and L is the number of layers. A perceptual registration loss [4] is inappropriate in this setting as these features correspond to different modality-specific spaces. Instead, we maximize a lower bound on mutual information between corresponding spatial locations in $A_1^k(I_1)$ and $A_2^k(I_2 \circ \phi)$ by minimizing a noise contrastive estimation loss [25]. As in [26], we project the channel-wise autoencoder features of size $\mathbb{R}^{N^k \times C^k}$ (where N^k is the number of spatial indices and C^k is the number of channels in layer) onto a hyperspherical representation space to obtain features $F_1^k(A_1^k(I_1))$ and $F_2^k(A_2^k(I_2 \circ \phi))$ where $F_{1,2}$ are 3-layer 256-wide trainable ReLU-MLPs [3]. In this space, indices in correspondence $f_i^k = F_1^k((A_1^k(I_1))_i)$ and $f_i^{k+} = F_2^k((A_2^k(I_2 \circ \phi))_i)$, where $i = 1, \dots, N^k$, are positive pairs. Similarly, f_i^k and $f_j^{k-} = F_2^k((A_2^k(I_2 \circ \phi))_j)$, where $j = 1, \dots, N^k$ and $j \neq i$, are negative pairs.

For optimal contrastive learning, we sample a single positive pair and $ns > 1$ negative samples, and use the following loss with τ as a temperature hyperparameter:

$$d_{I_2}(I_1 \parallel I_2 \circ \phi) = \sum_{k=1}^L \sum_{i=1}^{N^k} -\log \left(\frac{\exp(f_i^k \cdot f_i^{k+}/\tau)}{\exp(f_i^k \cdot f_i^{k+}/\tau) + \sum_{j=1, j \neq i}^{m_s} \exp(f_i^k \cdot f_j^{k-}/\tau)} \right)$$

Notably, as medical images also acquire empty space outside of the body, random patch sampling will lead to the sampling of false positive and negative pairs (e.g., background voxels sampled as both positive and negative pairs). To this end, the masked CR (mCR) model samples voxel pairs only within the union of the binary foregrounds of I_1 and I_2 and resizes this mask to the layer- k specific resolution when sampling from $A_1^k(I_1)$ and $A_2^k(I_2 \circ \phi)$. We further investigate tolerance to false positive and negative training pairs and thus also train models without masking (denoted CR only).

Hypernetwork optimization.

Registration performance strongly depends on weighing λ correctly for a given dataset and \mathcal{L}_{sim} . Therefore, for fair comparison, the entire range of λ is evaluated for all benchmarked methods using hypernetworks [10] developed for registration [16,23]. Specifically, the FiLM [27] based framework of [23] is used with a 4-layer 128-wide ReLU-MLP to generate a $\lambda \sim \mathcal{U}[0,1]$ -conditioned shared embedding, which is then linearly projected (with a weight decay of 10^{-5}) to each layer in the registration network to generate λ -conditioned scales and shifts for the network activations. At test time, we sample 17 registration networks for each method with dense λ sampling between [0,0.2) and sparse sampling between [0.2,1.0].

4. Experiments

Data.

We compare registration methods by benchmarking on inter-subject multi-modality registration. We use pre-processed T1w and T2w MRI of newborns imaged at 29–45 weeks gestational age from *dHCP* [22], a challenging dataset due to rapid temporal development in morphology and appearance alongside inter-subject variability. We follow [7] for further preprocessing to obtain $160 \times 192 \times 160$ volumes at $0.6132 \times 0.6257 \times 0.6572 \text{ mm}^3$ resolution and use 405/23/64 images for training, validation, and testing.

Evaluation methods.

Registration evaluation is non-trivial due to competing objectives. Low smoothness regularization (λ) can allow for near-exact matching of appearance but with anatomically-implausible and highly irregular deformations. Conversely, high λ enables smooth deformations with suboptimal alignment. Therefore, we evaluate registration performance and robustness as a function of λ , via Dice and Dice30 (average of 30% of lowest dice scores) scores, respectively, calculated between the target and moved label maps of the input images (segmentations provided by *dHCP* [21]). To investigate deformation smoothness, we also evaluate the percentage of folding voxels and the standard deviation of the \log Jacobian determinant of ϕ as a function of λ .

Benchmarked methods.

Using the same registration network with $ch = 64$ and $ts = 5$, we benchmark popular multi-modality metrics including Mutual Information (MI) (48 bins), local MI (48 bins, patch size = 9), MIND (distance = 2, patch size = 3), and normalized gradient fields, alongside the proposed mCR and CR models. We further compare against the general-purpose SynthMorph (SM)-shapes and brains models [15], by using their publicly released models and affine-aligning the images to their atlas. As SM uses $ch = 256$, we retrain the proposed registration models at that width. As we use public models, we cannot perform λ -conditioning and evaluation for SM. Further, as inter-subject dHCP registration can require large non-smooth deformations, we study whether a higher number of integration steps improves deformation characteristics (as in [35]) for the $ch = 256$ model, evaluating $ts = \{10,16,32\}$ with 32 as default.

To evaluate extensions of CR and *mCR*, we investigate whether combining them with a global loss (+ *MI*), incorporating more negative samples from an external randomly-selected subject (+ *ExtNegs*), or both (+ *MI* + *ExtNegs*) lead to improved results. We then evaluate the importance of feature extractor pretraining by using randomly-intialized *frozen* autoencoders instead as a worst-case feature extractor (+ *RandAE*). Finally, we evaluate whether contrastively pre-training the autoencoders and projection MLPs by using ground truth multi-modality image pairs alongside the reconstruction losses (+ *SupPretrain*) leads to improved results, with the following loss, where $I_{1,2}$ are from the same subject, $\lambda_{sp} = 0.1$, and $\hat{I}_{1,2}$ are the reconstructions,

$$\mathcal{L}_{A_1, A_2, F_1, F_2} = \lambda_{sp} d_{12}(I_1, I_2) + \sum_{i=1}^2 (\|I_{Ti} - \hat{I}_{Ti}\|_1 + NCC(I_{Ti}, \hat{I}_{Ti})).$$

Implementation details.

All models were developed in TensorFlow 2.4 and were all trained for 10^5 iterations with Adam on a V100 GPU. For stability across all methods, we use a conservative learning rate of 5×10^{-5} . For the contrastive loss, we set $n_s = 1024$ and $\tau = 0.007$. The autoencoder has *7Conv - IN - LeakyReLU(0.2)* blocks with 3 down/up sampling layers and 32-64-128-64-32-32-32-1 filters with its post-convolution features from the first 6 layers sampled for the contrastive loss. In practice, to avoid tuning the sampling strategy for λ as in [16], we add a rescaling constant $\alpha = 0.1$ to the objective function for CR and mCR with the form $\alpha(1 - \lambda)\mathcal{L}_{sim} + \lambda\mathcal{L}_{reg}$.

Results.

Sample registration visualizations are provided in Fig. 2, performance scores versus λ are plotted in Fig. 3, and a study of trading-off registration accuracy for smoothness is tabulated in Table 1. We make the following experimental observations:

(m)CR achieves higher accuracy and converges faster than baseline losses. Fig. 3 (row 1) indicates that the proposed models achieve better Dice with comparable

(mCR) or better (CR) folding and smoothness characteristics in comparison to baseline losses as a function of the 17 values of λ tested. Further, Table 1 reveals that if anatomical overlap is reduced to also achieve negligible folding (defined as folds in 0.5% of all voxels [30]), CR and mCR still achieve the optimal tradeoff.

(m)CR achieves more accurate registration than label-trained methods at the cost of lower warp regularity. While the public SM-brains model does not achieve the same Dice score as (m)CR, it achieves the third-highest performance behind (m)CR with substantially smoother deformations. We postulate that this effect stems from the intensity-invariant label-based training of SM-brains only looking at the semantics of the image, whereas our approach and other baselines are appearance based.

Masking consistently improves results. Excluding false positive and false negative pairs from the training patches yields improved registration performance across all values of λ with acceptable increases in deformation irregularities vs. λ (Fig 3 rows 1 & 3; cols 1–3). Importantly, contrastive training without foreground masks (CR) still outperforms other baseline losses and does so with smoother warps.

Pretraining has a marginal impact on (m)CR. While still outperforming other baselines, using a *randomly-initialized & frozen* feature extractor achieves marginally lower dice with longer convergence times as compared to the full pretrained model.

Using external losses or negatives with (m)CR does not improve results. Combining a global loss (MI) with CR does not improve results, which we speculate is due to the inputs already being globally affine-aligned. We also see a similar phenomenon to [26], where adding external negatives from other subjects lowers performance.

Self-supervision yields nearly the same performance as supervised pretraining. Comparing rows A5–6 and C4–5 of Table 1 reveals that utilizing supervised pairs of aligned images for pretraining $A_{1,2}$ and $F_{1,2}$ yields very similar results, indicating that supervision is not required for optimal registration in this context.

5. Discussion

Limitations and future work.

Some limitations exist in the presented material and will be addressed in subsequent work: (1) While we perform λ -conditioned hypernetwork registration to fairly compare benchmarked losses across all regularization strengths, hypernetworks may not exactly approximate all λ conditions. Further, hypernetworks were not trained for two of our baselines (SM-brains and shapes [15]) as we instead used their public models and we regularize for *velocity-smoothness* instead of *warp-smoothness* as in their work, both of which confound comparisons. (2) It is probable that combining our appearance-based approach with label-based simulation [15] would further improve results. (3) We did not explore other architectural configurations for the autoencoders and MLPs and it is plausible that there may be significant room for optimization. (4) We benchmarked the simulated inter-subject registration task and other use-cases such as pre-to-intra operative warping and preprocessing for multi-sensor fusion [6] may show different trends. (5) (m)CR currently requires ~15% more time per training iteration w.r.t. mutual information and can

be optimized. (6) Unsupervised patch sampling may introduce false negative pairs in the contrastive loss and can be avoided with unsupervised *negative-free* patch representation learning methods [32].

Conclusions.

This work presented *ContraReg*, a self-supervised contrastive representation learning approach to diffeomorphic non-rigid image registration. On the challenging task of inter-subject T1w–T2w registration with neonatal images showing high appearance and morphological variation, CR achieved high registration performance and robustness while maintaining desirable deformation qualities such as invertibility and smoothness. Finally, CR was validated across several baseline losses (including MI, MIND, NGF), training configurations, and frameworks, with results indicating that training supervision, losses, or external negative sampling strategies are not required.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Funding.

N. Dey and G. Gerig were partially supported by NIH 1R01HD088125-01A1. The dHCP data used in this study was funded by ERC Grant Agreement no. [319456].

References

1. Arar M, Ginger Y, Danon D, Bermano AH, Cohen-Or D: Unsupervised multi-modal image registration via geometry preserving image-to-image translation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2020)
2. Avants BB, Tustison NJ, Song G, Cook PA, et al. : A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage* (2011)
3. Chen T, Kornblith S, Swersky K, Norouzi M, Hinton G: Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029* (2020)
4. Czolbe S, Krause O, Feragen A: Semantic similarity metrics for learned image registration. In: *Medical Imaging with Deep Learning* (2021)
5. Dalca AV, Balakrishnan G, Guttag J, Sabuncu MR: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical Image Analysis* (2019)
6. Dey N, Li S, Bermond K, Heintzmann R, Curcio CA, Ach T, Gerig G: Multi-modal image fusion for multispectral super-resolution in microscopy. In: *Medical Imaging 2019: Image Processing*. vol. 10949, pp. 95–101. SPIE (2019)
7. Dey N, Ren M, Dalca AV, Gerig G: Generative adversarial registration for improved conditional deformable templates. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3929–3941 (October 2021)
8. Guo CK: Multi-modal image registration with unsupervised deep learning. Master's thesis, Massachusetts Institute of Technology (2019)
9. Gutierrez-Becker B, Mateus D, Peter L, Navab N: Guiding multimodal registration with learned optimization updates. *Medical image analysis* 41, 2–17 (2017) [PubMed: 28506641]
10. Ha D, Dai A, Le QV: Hypernetworks. *arXiv preprint arXiv:1609.09106* (2016)
11. Haber E, Modersitzki J: Intensity gradient based registration and fusion of multi-modal images. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 726–733. Springer (2006)

12. Han J, Shoeiby M, Petersson L, Armin MA: Dual contrastive learning for unsupervised image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 746–755 (June 2021)
13. Hata N, et al.: Multimodality deformable registration of pre-and intraoperative images for mri-guided brain surgery. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 1067–1074. Springer (1998)
14. Heinrich MP, Jenkinson M, Bhushan M, Matin T, Gleeson FV, Brady M, Schnabel JA: Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Medical image analysis* 16(7), 1423–1435 (2012) [PubMed: 22722056]
15. Hoffmann M, Billot B, Greve DN, Iglesias JE, Fischl B, Dalca AV: Synthmorph: learning contrast-invariant registration without acquired images. *IEEE Transactions on Medical Imaging* (2021)
16. Hoopes A, Hoffmann M, Fischl B, Guttag J, Dalca AV: Hypermorph: Amortized hyperparameter learning for image registration. In: International Conference on Information Processing in Medical Imaging. pp. 3–17. Springer (2021)
17. Jing L, Vincent P, LeCun Y, Tian Y: Understanding dimensional collapse in contrastive self-supervised learning. *arXiv preprint arXiv:2110.09348* (2021)
18. Lee D, Hofmann M, Steinke F, Altun Y, Cahill ND, Scholkopf B: Learning similarity measure for multi-modal 3d image registration. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 186–193. IEEE (2009)
19. Loeckx D, Slagmolen P, Maes F, Vandermeulen D, Suetens P: Nonrigid image registration using conditional mutual information. *IEEE transactions on medical imaging* (2009)
20. Lu J, Öfverstedt J, Lindblad J, Sladoje N: Is image-to-image translation the panacea for multimodal image registration? a comparative study. *arXiv preprint arXiv:2103.16262* (2021)
21. Makropoulos A, Gousias IS, Ledig C, Aljabar P, et al. : Automatic whole brain mri segmentation of the developing neonatal brain. *IEEE transactions on medical imaging* (2014)
22. Makropoulos A, et al. : The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage* 173, 88–112 (2018) [PubMed: 29409960]
23. Mok TC, Chung A: Conditional deformable image registration with convolutional neural network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 35–45. Springer (2021)
24. Nimsky C, Ganslandt O, Merhof D, et al. : Intraoperative visualization of the pyramidal tract by diffusion-tensor-imaging-based fiber tracking. *Neuroimage* (2006)
25. Oord A.v.d., Li Y, Vinyals O: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
26. Park T, Efros AA, Zhang R, Zhu JY: Contrastive learning for unpaired image-to-image translation. In: European Conference on Computer Vision. pp. 319–345. Springer (2020)
27. Perez E, Strub F, de Vries H, Dumoulin V, Courville AC: Film: Visual reasoning with a general conditioning layer. In: AAAI (2018)
28. Pielawski N, Wetzler E, Öfverstedt J, Lu J, Wählby C, Lindblad J, Sladoje N: CoMIR: Contrastive multimodal image representation for registration. In: Advances in Neural Information Processing Systems (2020)
29. Qin C, Shi B, Liao R, Mansi T, Rueckert D, Kamen A: Unsupervised deformable registration for multi-modal images via disentangled representations. In: International Conference on Information Processing in Medical Imaging. pp. 249–261. Springer (2019)
30. Qiu H, Qin C, Schuh A, et al. : Learning diffeomorphic and modality-invariant registration using b-splines. In: Medical Imaging with Deep Learning (2021)
31. Ren M, Dey N, Fishbaugh J, Gerig G: Segmentation-renormalized deep feature modulation for unpaired image harmonization. *IEEE Transactions on Medical Imaging* 40(6), 1519–1530 (2021) [PubMed: 33591913]
32. Ren M, Dey N, Styner MA, Botteron K, Gerig G: Local spatiotemporal representation learning for longitudinally-consistent neuroimage analysis. *arXiv preprint arXiv:2206.04281* (2022)
33. Risholm P, Golby AJ, Wells W: Multimodal image registration for preoperative planning and image-guided neurosurgical procedures. *Neurosurgery Clinics* 22(2), 197–206 (2011) [PubMed: 21435571]

34. Russakoff DB, Tomasi C, Rohlfing T, Maurer CR: Image similarity using mutual information of regions. In: European Conference on Computer Vision. Springer (2004)
35. Schuh A: Computational models of the morphology of the developing neonatal human brain. Ph.D. thesis, Imperial College London (2018)
36. Simonovsky M, Gutiérrez-Becker B, Mateus D, Navab N, Komodakis N: A deep metric for multimodal registration. In: International conference on medical image computing and computer-assisted intervention. pp. 10–18. Springer (2016)
37. Wachinger C, Navab N: Entropy and laplacian images: Structural representations for multi-modal registration. *Medical image analysis* 16(1), 1–17 (2012) [PubMed: 21632274]
38. Wells WM III, Viola P, Atsumi H, Nakajima S, Kikinis R: Multi-modal volume registration by maximization of mutual information. *Medical image analysis* 1(1), 35–51 (1996) [PubMed: 9873920]
39. Woo J, Stone M, Prince JL: Multimodal registration via mutual information incorporating geometric and spatial context. *IEEE Transactions on Image Processing* (2014)
40. Zhou B, Augenfeld Z, Chapiro J, Zhou SK, Liu C, Duncan JS: Anatomy-guided multimodal registration by learning segmentation without ground truth: Application to intraprocedural *cbct/mr* liver segmentation and registration. *Medical image analysis* (2021)

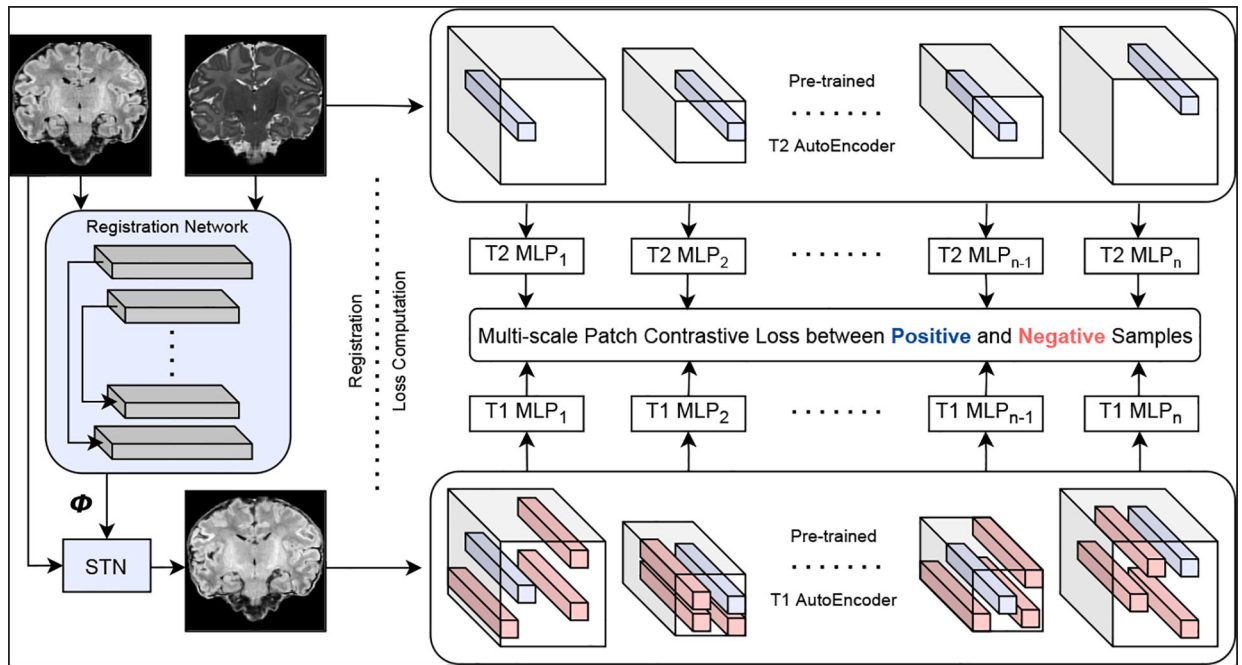


Fig. 1. ContraReg.

Given displacements and a moved image from a multi-modality registration network (**left**), a contrastive loss is calculated on multi-scale patches extracted from modality-specific networks (**right**), such that corresponding locations have high mutual information. In practice, our implementation is bidirectional s.t. the inverse modality-pair loss is also computed.

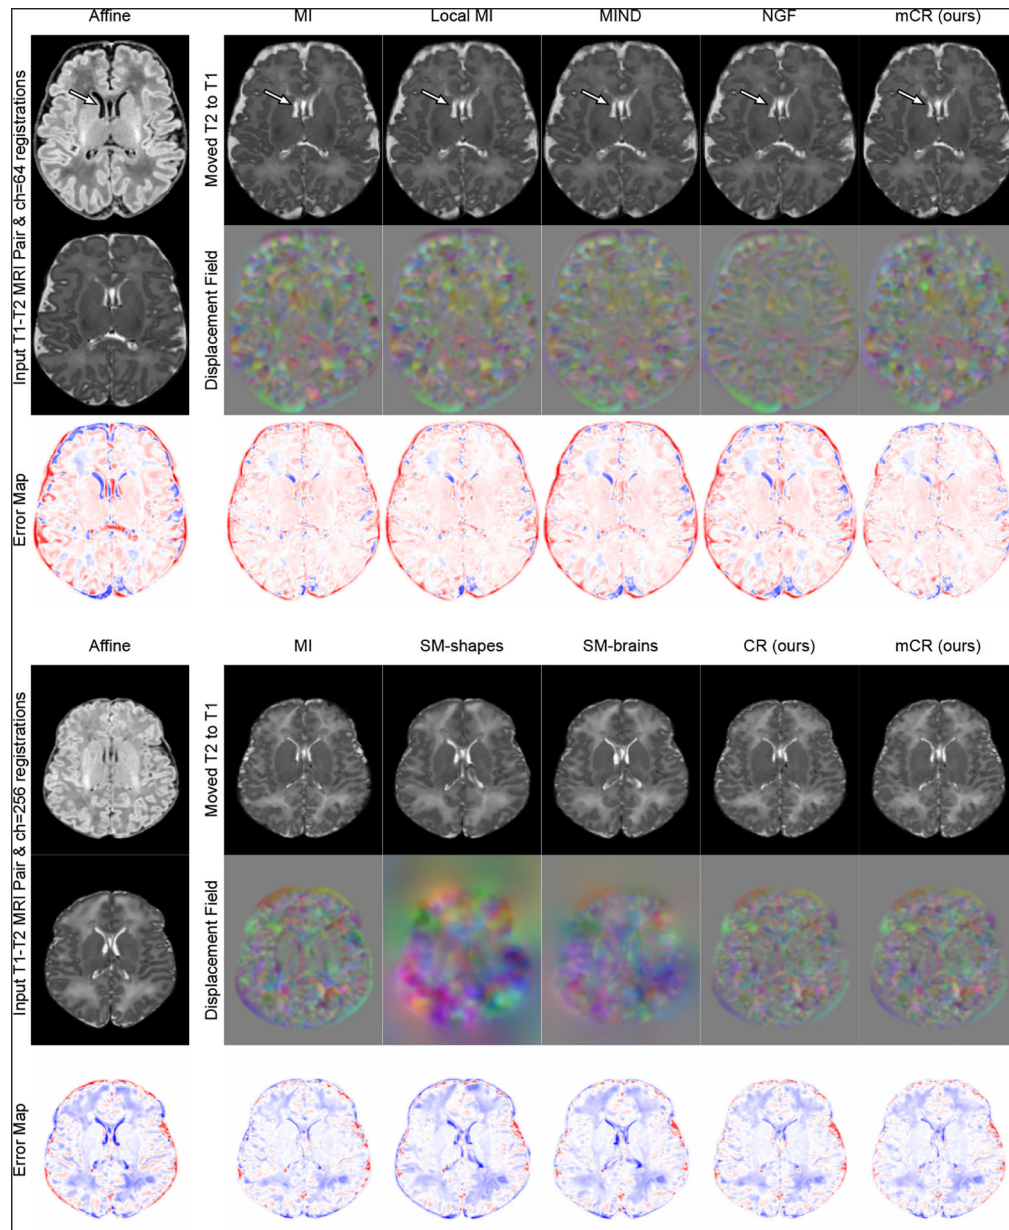


Fig. 2. T1w-T2w registration visualization between arbitrarily selected subjects for the **(top)** ch=64 and **(bottom)** ch=256 models. Error maps computed w.r.t. the T2w MRI of the target subject. Hypernetwork registration models are sampled with the same λ as Table 1.

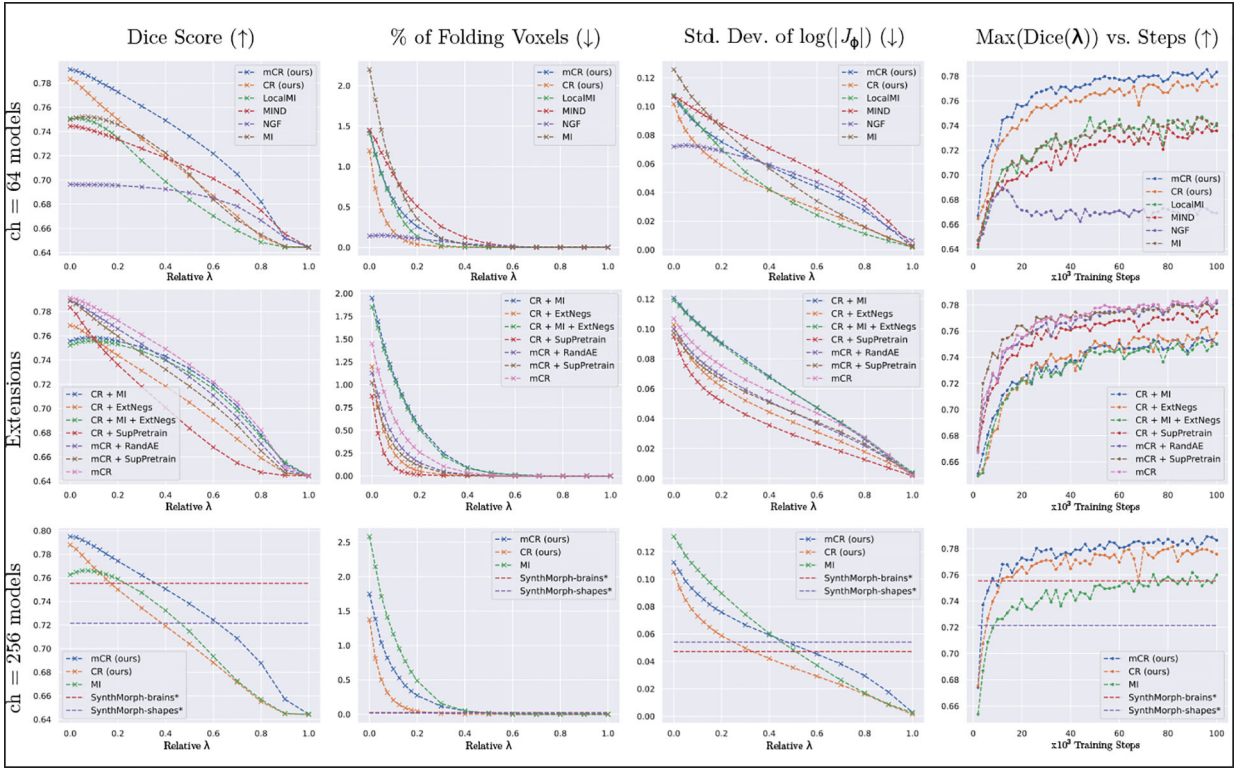


Fig. 3. Plots of registration accuracy vs. λ (**col. 1**), deformation qualities vs. λ (**cols. &3**), and accuracy vs. training steps (**col. 4**). Benchmarks are performed against commonly used multimodality losses (**row 1**), extensions of the proposed techniques (**row 2**), and recent modality-pair agnostic methods (**row 3**). Across baseline losses, CR and mCR achieve the best tradeoff between accuracy and deformation characteristics (**row 1, cols. 1–3**). Further, using external losses and/or negatives reduces performance and supervised pretraining does not yield notable improvements (**row 2, cols. 1–3**). Compared to SynthMorph-brains [15], CR and mCR obtain higher accuracy (**row 3, col. 1**) in the $\lambda = 0.0 - 0.15$ and $0.0 - 0.3$ ranges, respectively, at the cost of more irregular warps (**row 3, cols. 2–3**). See Table 1 for an analysis of trading off accuracy for smoothness.

Table 1.

Trading off performance for invertibility.

Set	Width	Method	Opt. λ	Dice (\uparrow)	Dice30 (\uparrow)	% Folds (\downarrow)	Sdlog $ J_\phi $ (\downarrow)
A	64	NGF [11]	0.0	0.696 \pm 0.023	0.686	0.141 \pm 0.043	0.072
	64	MI [38]	0.175	0.748 \pm 0.021	0.739	0.461 \pm 0.100	0.089
	64	LocalMI [8]	0.125	0.745 \pm 0.023	0.737	0.402 \pm 0.076	0.083
	64	MIND [14]	0.3	0.726 \pm 0.023	0.716	0.258 \pm 0.051	0.079
	64	CR (proposed)	0.05	0.776 \pm 0.020	0.768	0.451 \pm 0.074	0.083
	64	mCR (proposed)	0.125	0.781 \pm 0.020	0.774	0.475 \pm 0.070	0.084
B	256	SM-brains [15]	-	0.755 \pm 0.020	0.749	0.023 \pm 0.008	0.048
	256	SM-shapes [15]	-	0.721 \pm 0.021	0.715	0.017 \pm 0.011	0.056
	256	MI [38]	0.2	0.759 \pm 0.021	0.750	0.487 \pm 0.099	0.090
	256	CR (proposed)	0.075	0.774 \pm 0.020	0.765	0.315 \pm 0.0576	0.078
	256	mCR (proposed)	0.15	0.780 \pm 0.021	0.773	0.416 \pm 0.065	0.082
C	64	CR+MI	0.3	0.751 \pm 0.021	0.742	0.246 \pm 0.059	0.080
	64	CR+ExtNegs	0.05	0.764 \pm 0.020	0.756	0.489 \pm 0.073	0.085
	64	CR+MI+ExtNegs	0.3	0.747 \pm 0.021	0.739	0.214 \pm 0.056	0.078
	64	CR+SupPretrain	0.025	0.778 \pm 0.020	0.770	0.465 \pm 0.075	0.084
	64	mCR+SupPretrain	0.075	0.778 \pm 0.020	0.770	0.406 \pm 0.067	0.081
	64	mCR+RandAE	0.1	0.778 \pm 0.020	0.770	0.393 \pm 0.070	0.80
D	256	CR (10 int. steps)	0.075	0.773 \pm 0.021	0.764	0.341 \pm 0.058	0.079
	256	CR (16 int. steps)	0.05	0.779 \pm 0.020	0.772	0.462 \pm 0.071	0.083
	256	CR (32 int. steps)	0.075	0.774 \pm 0.020	0.765	0.315 \pm 0.0576	0.078

Registration accuracy (Dice), robustness (Dice30), and characteristics (% Folds, stddev. $\log|J_\phi|$) for all benchmarked methods at values of λ that maintains the percentage of folding voxels at less than 0.5% of all voxels, as in [30], s.t. high performance is achieved alongside negligible singularities. This table is best interpreted in conjunction with figure 3, where results from all λ values are visualized. **A.** CR and mCR obtain improved accuracy and robustness (A5–6) with similar deformation characteristics to baseline losses (A1–4). **B.** At larger model sizes, mCR and CR still obtain higher registration accuracy and robustness (B4–5), albeit at the cost of more irregular deformations in comparison to SM (B1). **C.** Further adding external losses, negative samples, or both to CR harms performance (C1–3), supervised pretraining (C4–5) very marginally improves results over training from scratch (A5–6), and random feature extraction only slightly reduces Dice while smoothening displacements (C6). **D.** At a given λ , increasing integration steps yields marginal Dice and smoothness improvements.