



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

One registration is worth two segmentations

Shiqi Huang^{1,2}, Tingfa Xu², Ziyi Shen¹, Shaheer Ullah Saeed¹, Wen Yan¹,
Dean Barratt¹, and Yipeng Hu¹

¹ University College London, London, UK

² Beijing Institute of Technology, Beijing, China
{huangsq, ciom_xtf1}@bit.edu.cn
yipeng.hu@ucl.ac.uk

Abstract. The goal of image registration is to establish spatial correspondence between two or more images, traditionally through dense displacement fields (DDFs) or parametric transformations (e.g., rigid, affine, and splines). Rethinking the existing paradigms of achieving alignment via spatial transformations, we uncover an alternative but more intuitive correspondence representation: a set of corresponding regions-of-interest (ROI) pairs, which we demonstrate to have sufficient representational capability as other correspondence representation methods. Further, it is neither necessary nor sufficient for these ROIs to hold specific anatomical or semantic significance. In turn, we formulate image registration as searching for the same set of corresponding ROIs from both moving and fixed images - in other words, two multi-class segmentation tasks on a pair of images. For a general-purpose and practical implementation, we integrate the segment anything model (SAM) into our proposed algorithms, resulting in a SAM-enabled registration (SAMReg) that does not require any training data, gradient-based fine-tuning or engineered prompts. We experimentally show that the proposed SAMReg is capable of segmenting and matching multiple ROI pairs, which establish sufficiently accurate correspondences, in three clinical applications of registering prostate MR, cardiac MR and abdominal CT images. Based on metrics including Dice and target registration errors on anatomical structures, the proposed registration outperforms both intensity-based iterative algorithms and DDF-predicting learning-based networks, even yielding competitive performance with weakly-supervised registration which requires fully-segmented training data.

Keywords: Image Registration · Correspondence Representation · Segment Anything Model (SAM).

1 Introduction

Registering a pair of moving and fixed images is one of the basic tasks in many medical image computing applications [8,10,22], among which the registration-estimated spatial transformation can warp the moving image to be *aligned* with the fixed image. Common transformation functions may be parameterized by translation, rotation, scaling, spline coefficients, spatially-sampled control point

coordinates, and, most generally, displacement vectors at all voxel locations - a dense displacement field (DDF). Like their *classical* interactive counterparts, recent learning-based registration algorithms have also been formulated to predict these transformation parameters [9].

Warping images is only one of many possible use cases of the output transformation that represents spatial correspondence. For example, registration-enabled atlas-based label fusion [30], morphological analysis on predefined anatomical structures [5], surgical navigation [17] and target tracking [25], require accurate correspondence to be estimated on only one or a small number of regions-of-interest (ROIs). In fact, few clinical applications require ubiquitous correspondence and, arguably, the registration becomes an *ill-posed* problem when estimating $N \times 3$ displacements at all voxel locations, given $N \times 2$ intensity values from two images, N being the average number of voxels in one image.

Representing correspondence with the above-discussed transformation functions is sufficient for most downstream tasks, but their necessity in representing correspondence (therefore efficiency) is application-dependent. This motivated this study to propose a different type of correspondence representation, using a set of corresponding ROIs. This is considered flexible to allow correspondence to be represented with only a few paired ROI masks, yet without losing generality (e.g. more and smaller ROI pairs) if denser local correspondence is required.

For designing a learning-based registration algorithm to estimate corresponding ROI pairs, it is intuitive to use predefined ROI types and their segmented labels for training, akin to a weakly-supervised algorithm [13], with or without intensity-based unsupervised loss [2]. However, requiring ROI-labelled training data limits the general applicability of the algorithm.

In this work, we propose a practical and general-purpose registration algorithm to overcome the limitation due to training data. Specifically, the registration process is streamlined using Segment Anything Model (SAM) [16] to automatically generate prototypes from the multiple segmented ROIs on both images, and match them in pairs based on similarity. The presented experiments show that neither the definition nor the quantity of these SAM-segmented ROIs needs to be prescribed (or kept the same), for estimating sufficiently useful correspondences, in three different clinical image registration tasks.

In summary, our main contributions are: 1) We pioneer a new correspondence representation for image registration using intuitive multiple paired-ROIs. 2) We introduce the general-purpose SAMReg by segmenting corresponding ROIs from both images, without any labels or fine-tuning. To our knowledge, it is the first application trial of SAM in image registration. 3) Extensive experiments show that our method outperforms the common-used registration tool and unsupervised registration, and is competitive with weakly-supervised registration that requires additional labeled training images, across three clinical applications.

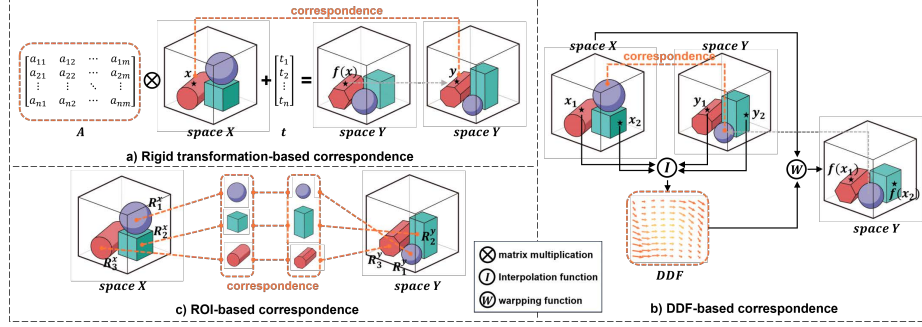


Fig. 1. The correspondence representation diagram.

2 ROI-based Correspondence Representation

This section considers spatial correspondence to be a mapping $f : \mathbf{X} \rightarrow \mathbf{Y}$ from spatial locations \mathbf{X} in a moving image coordinate system to locations \mathbf{Y} in the fixed image coordinate system. For example in 3D, $\mathbf{X} \in \mathbb{R}^3$ and $\mathbf{Y} \in \mathbb{R}^3$ are random vectors containing Euclidean coordinates, in x-, y- and z-directions.

Existing methods for representing spatial correspondence, f , fall into two main types: transformation functions of spatial locations or paired samples. The first type includes parametric functions like rigid, affine, and spline-based transformations (e.g. thin-plate-splines [23]) that map input vectors \mathbf{x} to coordinate vectors \mathbf{y} through $\mathbf{y} = f(\mathbf{x}) : \mathbf{X} \rightarrow \mathbf{Y}$. An example (shown in Fig. 1a) is the rigid transformation $f^{rigid}(\mathbf{x}) = A\mathbf{x} + \mathbf{t}$, where A is a rotation matrix and \mathbf{t} is a translation vector. The second type represents correspondence through N pairs of corresponding locations $(\mathbf{x}_n, \mathbf{y}_n)$, where $n = 1, \dots, N$. For instance, DDF (shown in Fig. 1b) is represented by $\{(\mathbf{x}_n, \mathbf{x}_n + \mathbf{d}_n)\}$ where \mathbf{d}_n is a set of vectors defined on all voxel locations \mathbf{x}_n . Interpolation is applied to estimate \mathbf{y} for \mathbf{x} at locations beyond the defined voxels \mathbf{x}_n . In this context, we formulate a more intuitive ROI-based correspondence representation as an alternative:

Theorem 1. Denoting \mathbf{x} and \mathbf{y} as spatial locations at respective moving- and fixed image spaces (i.e. coordinate systems), it is sufficient for K pairs of regions-of-interest (ROIs), denoted as $\{(R_k^x, R_k^y)\}_{k=1}^K$, to indicate any spatial correspondence between the two image spaces, if K is sufficiently large, where $R_k^x = \{\mathbf{x}_l\}_{l=1}^{L_k^x}$ and $R_k^y = \{\mathbf{y}_l\}_{l=1}^{L_k^y}$ are two sets of L_k^x and L_k^y spatial locations that represent the same (corresponding) ROIs, in the moving- and fixed image spaces, respectively.

This is intuitive in an extreme case, in which a large number ($K = N$) of single-voxel (i.e. $L_k^x = L_k^y = 1$) ROIs are used, the ROI pairs can be sampled at N voxels to an equivalent DDF representation $\{(R_k^x = \mathbf{x}_k, R_k^y = \mathbf{y}_k)\}_{k=1}^N$.

Therefore, an informal proof for Theorem 1 can be understood by iterating the following process until sufficiency is reached: When a point-to-point correspondence $\mathbf{y}^* = f(\mathbf{x}^*)$ that has NOT been represented by current ROI

pairs $\{(R_k^x, R_k^y)\}_{k=1}^K$, one can always sample an additional ROI pair $\{(R_{K+1}^x = \mathbf{x}^*, R_{K+1}^y = \mathbf{y}^*)\}$, such that the resulting $(K + 1)$ ROI pairs $\{(R_k^x, R_k^y)\}_{k=1}^{K+1} = \{(R_k^x, R_k^y)\}_{k=1}^K \cup \{(R_{K+1}^x, R_{K+1}^y)\}$ suffice.

Further, properties of the proposed ROI-based correspondence are summarized: 1) individual ROIs can be practically depicted as binary pair masks; 2) paired ROIs can represent both local and global correspondence; 3) Overlapping ROIs facilitate one-to-many correspondence; 4) if dense-level correspondence is required, the correspondence at any spatial location can be inferred. We discuss specific algorithms in Sec. 4.2; 5) if dense-level correspondence is not required, omitting smoothness regularization (e.g., rigidity or bending energy) during estimation enhances registration efficiency with ROI-based representation.

3 One Registration versus Two Segmentations

Given a single-image I^x , a multi-class segmentation is to estimate the joint probability of a random vector $\mathbf{C}^x = [C_1^x, \dots, C_{K^x}^x]^\top \in \mathbb{R}^{K^x}$, representing K^x -class probabilities over N voxel locations, $\prod_{n=1}^N p_n(\mathbf{C}^x | I^x)$. The segmented ROI set $\{\hat{R}_k^x\}_{k=1}^{K^x}$ can be computed by binarising the estimated class probability vectors $\{\hat{\mathbf{c}}_n^x\}_{n=1}^N$, where $\hat{\mathbf{c}}_n^x \sim \mathbf{C}^x$. For a second image I^y , ROIs \hat{R}_k^y are segmented by estimating K^y -class probabilities $\hat{\mathbf{c}}_n^y$, where $\hat{\mathbf{c}}_n^y$ is the class probabilities for K^y classes. In this section, we consider the following two cases for registering I^x and I^y , by estimating ROI-based correspondence described in Sec. 2.

Segmenting the same ROIs from two images. In this case, the segmentation algorithm (e.g. one or two neural networks) is capable of segmenting the same classes of ROIs, i.e. $\mathbf{C} = \mathbf{C}^x = \mathbf{C}^y$ and $K = K^x = K^y$, from both images. Therefore, the obtained ROIs pairs can directly represent the correspondence, described in Sec. 2, $\{(R_k^x = \hat{R}_k^x, R_k^y = \hat{R}_k^y)\}_{k=1}^K$, that is to estimate the monotonic joint probability given two images $p_n(\mathbf{C} | I^x, I^y) \propto p_n(\mathbf{C} | I^x) \cdot p_n(\mathbf{C} | I^y)$, subject to normalization constants. This approach usually requires predefined ROI types and segmentation networks trained using these ROI-labelled data. Interestingly, these segmented training data are exactly the same data required for training weakly-supervised registration algorithms [13,12].

Matching the segmented ROI candidates. Assume an alternative "unsupervised" segmentation can be used to obtain two sets of candidate ROIs, $\{\hat{R}_k^x\}_{k=1}^{K^x}$ and $\{\hat{R}_k^y\}_{k=1}^{K^y}$, from the fixed and moving- images, respectively. Search for a subset of common ROI classes, $\mathbf{C} = [C_1, \dots, C_K]^\top$, where $\{C_k\}_{k=1}^K = \{C_k^x\}_{k=1}^{K^x} \cap \{C_k^y\}_{k=1}^{K^y}$. Thus, the corresponding ROI pairs $\{(R_k^x, R_k^y)\}_{k=1}^K$ derived from the estimated probability vectors $\hat{\mathbf{c}}_n \sim \mathbf{C}$, that is the joint probability between the two conditionally independent "candidate probabilities", $p_n(\mathbf{C} | I^x, I^y) = p_n(\mathbf{C}^x, \mathbf{C}^y | I^x, I^y) = p_n(\mathbf{C}^x | I^x, I^y) \cdot p_n(\mathbf{C}^y | I^x, I^y) \propto p_n(\mathbf{C}^x | I^x) \cdot p_n(\mathbf{C}^y | I^y)$. In the following Sec. 4, we propose a specific algorithm using this strategy.

4 Segment Everything and Match: A Training-Free Registration Algorithm

4.1 Preliminary: SAM Architecture

SAM [16] consists of an image encoder, a prompt encoder and a mask decoder, is adept at image segmentation without task-specific training, and has recently been adapted for medical imaging. The adaptation includes training on large medical datasets [18,6], thus establishing a foundation model, and fine-tuning for targeted medical datasets using, for example, additional adaptation layers [29,11,28]. Despite the advancements, as highlighted by these authors, SAM may yield limited sensitivity to medical images compared to natural images [14,19], resulting in mixed results in direct applications for specific anatomy and pathology segmentation without further training.

It is interesting to argue that requiring known anatomical significance in segmentation may not limit medical image registration. Although correspondence between medical images may be application-dependent to some extent [7], the main objective is to find *similar* regions across images, a task that suits SAM’s strong feature and pixel classification capabilities. For example, SAM can effectively segment ROIs with clear edges such as the iliac internal, which are not typical targets but useful for establishing local correspondence. By deploying SAM to segment extensive foreground areas, we aim to discover corresponding ROIs that capture diverse anatomical structures and areas that may lack precise anatomical definitions (*e.g.*, common cavities between structures).

4.2 The SAMReg Algorithm

Since the multiple ROIs generated by SAM segmentation lack explicit one-to-one correspondence, our strategy involves embedding these ROIs before matching them in feature space, with the proposed pipeline illustrated in Fig. 2.

ROI Embedding: Leveraging the SAM, we independently encode the input images I^x and I^y into features, F^x and F^y respectively. By setting the SAM decoder to the *everything* mode with a simple outlier-removing filter (further details in Sec. 5), we exploit its generality to produce a comprehensive array of segmentation masks for each image, obtaining mask sets $\{M_k^x\}_{k=1}^{K^x}$ and $\{M_k^y\}_{k=1}^{K^y}$, representing ROIs $\{R_k^x\}_{k=1}^{K^x}$ and $\{R_k^y\}_{k=1}^{K^y}$, respectively.

For each binary mask, we derive moving prototypes $\{p_k^x\}_{k=1}^{K^x}$ by element-wise multiplying the mask with their corresponding feature map, element being indexed by n : $p_k^x = \frac{\sum_n M_k^{x,res}(n) \cdot F^x(n)}{\sum_n M_k^{x,res}(n)}$, where *res* indicates resizing M_k^x to the size of F^x as $M_k^{x,res}$. This also applies on the fixed prototypes $\{p_k^y\}_{k=1}^{K^y}$.

ROI Matching: Assume a similarity matrix $S \in \mathbb{R}^{K^x \times K^y}$ to measure the cosine similarity between the moving- and fixed prototypes: $S(i, j) = \frac{p_i^x \cdot p_j^y}{\|p_i^x\| \cdot \|p_j^y\|}$, where $\|\cdot\|$ normalises each element $S(i, j) \in [0, 1]$, $i \in [1, K^x]$ and $j \in [1, K^y]$.

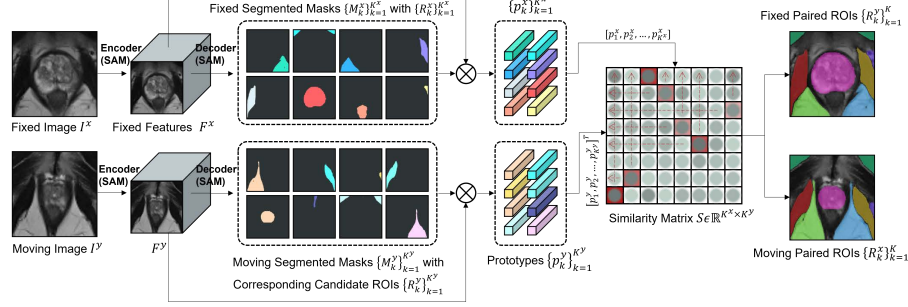


Fig. 2. The pipeline of the proposed SAM-Reg method.

A set of index pairs P then can be identified: $P = \arg \max_{i,j} (S(i,j)), s.t. S(i,j) > \epsilon \forall (i,j)$, corresponding to pairings with maximum similarities. A further constraint is added to ensure each i, j is chosen at most once, if overlapping ROIs segmented from SAM are filtered. ϵ is the similarity threshold and the set length $K = |P| \leq \min(K^x, K^y)$. Using the selected index pairs P , we construct two new sets of masks, $\{M_k^{x,cor}\}_{k=1}^K$ and $\{M_k^{y,cor}\}_{k=1}^K$:

$$(M_k^x, M_k^y)^{cor} = \{(M_i^x, M_j^y) | (i, j) \in P, k = 1, 2, \dots, K. \quad (1)$$

This leads to the final paired ROIs indicating the correspondence between I^x and I^y . This matching mechanism can be applied on either 2D images or 3D images, in the latter of which, SAM-segmented 2D slices can be stacked to provide the candidate masks, $\{M_k^x\}_{k=1}^{K^x}$ and $\{M_k^y\}_{k=1}^{K^y}$ in 3D, before computing prototypes.

Optional ROI-to-Dense Correspondence: Further, as discussed in Sec. 3, we explore the capability of converting estimated ROI-based correspondence to its dense counterpart DDF, useful for applications such as full-image alignment. In this section, we describe a general method that iteratively refines the DDF, to optimize an objective function \mathcal{L} that combines a region-specific alignment measure \mathcal{L}_{roi} with a regularization term \mathcal{L}_{ddf} to ensure smooth interpolation:

$$\mathcal{L}_\Theta(\{(M_k^x, M_k^y)\}) = \sum_{k=1}^K \mathcal{L}_{roi}(M_k^x, \mathcal{T}(M_k^y, \Theta)) + \lambda \mathcal{L}_{ddf}(\Theta) \quad (2)$$

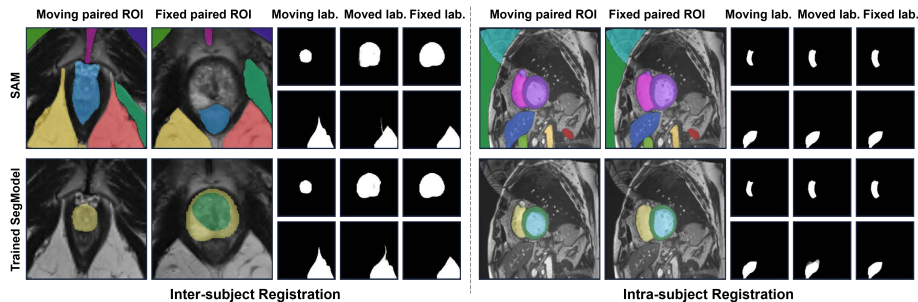
where Θ represents parameters of the transformation function \mathcal{T} and λ is a regularization parameter that balances the alignment accuracy with the deformation smoothness. This may itself be considered a multi-ROI alignment algorithm but with known ROI correspondence. In this work, we use an equally-weighted MSE and Dice as \mathcal{L}_{roi} and a L^2 -norm of DDF gradient as \mathcal{L}_{ddf} .

Table 1. Comparison with the-state-of-the-art medical image registration methods, where the * indicates methods specifically trained on the dataset.

Method	MR-Prostate [1]		MR-Cardiac [3]		CT-Abdomen [15]	
	Dice	TRE	Dice	TRE	Dice	TRE
NiftyReg [21]	7.68±3.98	4.67±3.48	9.93±2.21	3.29±2.89	6.81±3.02	5.86±3.64
VoxelMorph* [2]	56.84±3.41	3.68±1.92	60.10±3.95	3.03±2.41	58.34±3.72	4.13±2.23
LabelReg* [13]	77.32±3.56	2.72±1.23	78.97±2.42	1.73±1.34	74.53±3.43	2.58±1.18
SAMReg(Ours)	75.67±3.81	2.09±1.35	80.28±3.67	1.43±1.13	71.74±3.72	2.48±1.01

Table 2. Ablation study of segmentation performance: SAM versus leading dataset-specific segmentation models. *Lab-ROI* represents the dataset’s original labeled ROIs, while *Pse-ROI* refers to pseudo random ROIs identified by SAM.

Dataset	Seg.Model	Lab-ROI		Pse-ROI	
		Dice	TRE	Dice	TRE
Prostate [1]	SAM [16]	75.67±3.81	2.72±1.23	97.73±3.15	0.82±0.23
	Vnet [20]	80.35±3.14	2.13±1.24	43.56±3.54	5.32±2.51
Cardiac [3]	SAM [16]	80.28±3.67	1.73±1.34	98.87±3.61	1.62±0.67
	FCT [27]	83.54±3.46	1.31±1.25	31.35±3.74	7.24±2.52
Abdomen [15]	SAM [16]	71.74±3.72	2.58±1.18	98.24±3.17	1.45±0.75
	MedNeXt [24]	78.26±3.34	1.84±1.54	45.73±3.73	5.37±2.12

**Fig. 3.** The qualitative comparisons of SAM and fully-supervised segmentation results, including ROI-label(top) and pseudo-label(bottom)wrapping outcomes.

5 Experiments and Results

Datasets and Implementation Details: Our method is evaluated on MR-Prostate [1], MR-Cardiac [3], and CT-Abdomen [15] datasets, highlighting diverse anatomical areas and modalities. We apply inter-subject registration for Prostate and Abdomen datasets to align images to a standard reference, and intra-subject registration for Cardiac images, aligning time-variant image pairs. Performance is measured using the Dice score for overlap and target registration error (TRE) for landmark alignment, where TRE is defined as the root-mean-square (RMS) distance. For non-rigid registration [10,26], we assess accu-

racy by the mean transformation accuracy of critical anatomical structures [13], with TRE calculated from the centroids of the relevant ROIs. In configuring the SAM parameters, beyond the standard settings, we adjusted the pred-iou-thresh and stability-score-thresh to 0.90. Spatial size of input images are normalized to 256×256 and the ROI filtering strategy for output masks is based on area and overlapping ratio, setting a minimum of 200 and a maximum of 7000 for the ROI sizes and a maximum ratio of 0.8 empirically. Spatial size of $M_k^{x,res}$ is 32×32 , and the similarity threshold ϵ is set to 0.8. Our network is implemented with pytorch and MONAI [4]. Code demo has been released at: <https://github.com/squuang0103/SAMReg.git>.

Table 3. Ablation of paired ROI quantity on MR-prostate dataset.

PairNum	Dice	TRE
1	69.59±3.51	3.65±1.53
2	74.19±3.41	2.63±1.76
3	76.00±3.23	2.02±1.35
4	75.83±3.43	2.42±1.65
All	74.67±3.31	2.72±1.53

Table 4. Ablation of minimum similarity threshold ϵ on MR-prostate dataset.

Threshold ϵ	Dice	TRE
0.2	72.76±3.63	3.65±1.38
0.4	74.43±3.40	2.46±1.56
0.6	75.78±3.24	2.20±1.43
0.8	70.43±3.18	3.13±1.53

Comparing SAMReg with SOTA Methods: In our evaluation, we fairly compare SAMReg with NiftyReg [21], a non-learning tool, and also contrast it, less equitably, with tailored methods like VoxelMorph [2] and LabelReg [13]. VoxelMorph and LabelReg are fine-tuned on the research datasets, employing unsupervised and weakly-supervised learning respectively, with LabelReg also depending on anatomical annotations. Despite this, SAMReg demonstrates outstanding performance without dataset-specific training, particularly in the intra-subject cardiac dataset due to consistent patterns, as illustrated in Table 1. For the inter-subject datasets of MR-Prostate and MR-Cardiac, SAMReg slightly lags behind in Dice scores but showcases competitive TRE performance, highlighting its interesting local alignment ability over shape conformity, perhaps due to the involvement of multiple ROIs. Furthermore, the ability to achieve such results without dataset-specific training underscores the robust generalization capacity of SAMReg, suggesting its wide applicability in medical tasks, and reducing the time and resources needed for clinical deployment.

Ablative Study:

ROI Segmentation Comparison: Table 2 and Fig. 3 present ROI segmentation and the registration alignment performance comparisons between SAM and top segmentation models. Despite an unfair advantage for benchmark models designed and trained for groundtruth annotations (Label-ROI), SAMReg demonstrates robust performance and effectiveness, without ground-truth reliance, even outperforming these models with SAM-generated pseudo ROIs (Pse-ROI).

Optimal Number K of ROI Pairs: Table 3 shows that using up to three ROI pairs optimizes registration, improving Dice and TRE scores. This

is an interestingly small number, suggesting future work for even furthering registration performance, by more effectively utilizing smaller paired ROIs.

Threshold ϵ for ROI Correspondence: As illustrated in Table 4, a 0.6 similarity threshold enhances Dice scores and reduces TRE, while a higher threshold (0.8) reduces ROI selection and performance, underscoring the importance of a balanced threshold for optimal registration results.

6 Conclusion

In this study, We introduce a novel ROI-based correspondence representation. With this representation, the image registration is reformalized as two multi-class segmentation tasks, with a proposed, general-purpose and practical implementation, SAMReg. The comprehensive experiments show competitive registration performance that promises a new research direction in image registration.

Acknowledgments. This work was supported by the International Alliance for Cancer Early Detection, a partnership between Cancer Research UK [C28070/A30912; C73666/A31378], Canary Center at Stanford University, the University of Cambridge, OHSU Knight Cancer Institute, University College London and the University of Manchester.

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

References

1. Ahmed, H.U., Bosaily, A.E.S., Brown, L.C., Gabe, R., Kaplan, R., Parmar, M.K., Collaco-Moraes, Y., Ward, K., Hindley, R.G., Freeman, A., et al.: Diagnostic accuracy of multi-parametric mri and trus biopsy in prostate cancer (promis): a paired validating confirmatory study. *The Lancet* **389**(10071), 815–822 (2017)
2. 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 (2019)
3. 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)
4. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. *arXiv:2211.02701* (2022)
5. Cerveri, P., Belfatto, A., Manzotti, A.: Pair-wise vs group-wise registration in statistical shape model construction: representation of physiological and pathological variability of bony surface morphology. *Computer Methods in Biomechanics and Biomedical Engineering* **22**(7), 772–787 (2019)
6. Cheng, J., Ye, J., Deng, Z., Chen, J., Li, T., Wang, H., Su, Y., Huang, Z., Chen, J., Jiang, L., et al.: Sam-med2d. *arXiv:2308.16184* (2023)

7. Crum, W.R., Griffin, L.D., Hill, D.L., Hawkes, D.J.: Zen and the art of medical image registration: correspondence, homology, and quality. *NeuroImage* **20**(3), 1425–1437 (2003)
8. De Vos, B.D., Berendsen, F.F., Viergever, M.A., Sokooti, H., Staring, M., Išgum, I.: A deep learning framework for unsupervised affine and deformable image registration. *Medical image analysis* **52**, 128–143 (2019)
9. DeTone, D., Malisiewicz, T., Rabinovich, A.: Deep image homography estimation. arXiv:1606.03798 (2016)
10. Eppenhof, K.A., Lafarge, M.W., Moeskops, P., Veta, M., Pluim, J.P.: Deformable image registration using convolutional neural networks. In: *Medical Imaging 2018: Image Processing*. vol. 10574, pp. 192–197. SPIE (2018)
11. Gong, S., Zhong, Y., Ma, W., Li, J., Wang, Z., Zhang, J., Heng, P.A., Dou, Q.: 3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable medical image segmentation. arXiv:2306.13465 (2023)
12. Hu, Y., Gibson, E., Barratt, D.C., Emberton, M., Noble, J.A., Vercauteren, T.: Conditional segmentation in lieu of image registration. In: *MICCAI 2019*. pp. 401–409. Springer (2019)
13. Hu, Y., Modat, M., Gibson, E., Li, W., Ghavami, N., Bonmati, E., Wang, G., Bandula, S., Moore, C.M., Emberton, M., et al.: Weakly-supervised convolutional neural networks for multimodal image registration. *Medical image analysis* **49**, 1–13 (2018)
14. Huang, Y., Yang, X., Liu, L., Zhou, H., Chang, A., Zhou, X., Chen, R., Yu, J., Chen, J., Chen, C., et al.: Segment anything model for medical images? *Medical Image Analysis* **92**, 103061 (2024)
15. Ji, Y., Bai, H., Ge, C., Yang, J., Zhu, Y., Zhang, R., Li, Z., Zhanng, L., Ma, W., Wan, X., et al.: Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. In: *NeurIPS 2022*. pp. 36722–36732 (2022)
16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv:2304.02643 (2023)
17. Luebbbers, H.T., Messmer, P., Obwegeser, J.A., Zwahlen, R.A., Kikinis, R., Graetz, K.W., Matthews, F.: Comparison of different registration methods for surgical navigation in cranio-maxillofacial surgery. *Journal of Cranio-Maxillofacial Surgery* **36**(2), 109–116 (2008)
18. Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. *Nature Communications* **15**(1), 654 (2024)
19. Mazurowski, M.A., Dong, H., Gu, H., Yang, J., Konz, N., Zhang, Y.: Segment anything model for medical image analysis: an experimental study. *Medical Image Analysis* **89**, 102918 (2023)
20. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *3DV 2016*. pp. 565–571. Ieee (2016)
21. Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. *Journal of medical imaging* **1**(2), 024003–024003 (2014)
22. Rohé, M.M., Datar, M., Heimann, T., Sermesant, M., Pennec, X.: Svf-net: learning deformable image registration using shape matching. In: *MICCAI 2017*. pp. 266–274. Springer (2017)
23. Rohr, K., Stiehl, H.S., Sprengel, R., Buzug, T.M., Weese, J., Kuhn, M.: Landmark-based elastic registration using approximating thin-plate splines. *IEEE Transactions on medical imaging* **20**(6), 526–534 (2001)

24. Roy, S., Koehler, G., Ulrich, C., Baumgartner, M., Petersen, J., Isensee, F., Jaeger, P.F., Maier-Hein, K.H.: Mednext: transformer-driven scaling of convnets for medical image segmentation. In: MICCAI 2023. pp. 405–415. Springer (2023)
25. Seregini, M., Paganelli, C., Summers, P., Bellomi, M., Baroni, G., Riboldi, M.: A hybrid image registration and matching framework for real-time motion tracking in mri-guided radiotherapy. *IEEE Transactions on Biomedical Engineering* **65**(1), 131–139 (2017)
26. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: A survey. *IEEE transactions on medical imaging* **32**(7), 1153–1190 (2013)
27. Tragakis, A., Kaul, C., Murray-Smith, R., Husmeier, D.: The fully convolutional transformer for medical image segmentation. In: WCACV 2023. pp. 3660–3669 (2023)
28. Wang, H., Guo, S., Ye, J., Deng, Z., Cheng, J., Li, T., Chen, J., Su, Y., Huang, Z., Shen, Y., Fu, B., Zhang, S., He, J., Qiao, Y.: Sam-med3d. arXiv:2310.15161 (2023)
29. Wu, J., Fu, R., Fang, H., Liu, Y., Wang, Z., Xu, Y., Jin, Y., Arbel, T.: Medical sam adapter: Adapting segment anything model for medical image segmentation. arXiv:2304.12620 (2023)
30. Yang, H., Sun, J., Li, H., Wang, L., Xu, Z.: Neural multi-atlas label fusion: Application to cardiac mr images. *Medical image analysis* **49**, 60–75 (2018)