


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Computer Methods and Programs in Biomedicine

journal homepage: www.elsevier.com/locate/cmpb

RegWSI: Whole slide image registration using combined deep feature- and intensity-based methods: Winner of the ACROBAT 2023 challenge

Marek Wodzinski ^{a,b,*}, Niccolò Marini ^a, Manfredo Atzori ^{a,c}, Henning Müller ^{a,d}^a Institute of Informatics, University of Applied Sciences Western Switzerland, Sierre, Switzerland^b Department of Measurement and Electronics, AGH University of Kraków, Krakow, Poland^c Department of Neuroscience, University of Padova, Padova, Italy^d Medical Faculty, University of Geneva, Geneva, Switzerland

ARTICLE INFO

Keywords:

Digital pathology
 Deep learning
 WSI registration
 Image registration
 Whole slide images
 ACROBAT
 ANHIR
 HyReCo

ABSTRACT

Background and objective: The automatic registration of differently stained whole slide images (WSIs) is crucial for improving diagnosis and prognosis by fusing complementary information emerging from different visible structures. It is also useful to quickly transfer annotations between consecutive or restained slides, thus significantly reducing the annotation time and associated costs. Nevertheless, the slide preparation is different for each stain and the tissue undergoes complex and large deformations. Therefore, a robust, efficient, and accurate registration method is highly desired by the scientific community and hospitals specializing in digital pathology. **Methods:** We propose a two-step hybrid method consisting of (i) deep learning- and feature-based initial alignment algorithm, and (ii) intensity-based nonrigid registration using the instance optimization. The proposed method does not require any fine-tuning to a particular dataset and can be used directly for any desired tissue type and stain. The registration time is low, allowing one to perform efficient registration even for large datasets. The method was proposed for the ACROBAT 2023 challenge organized during the MICCAI 2023 conference and scored 1st place. The method is released as open-source software.

Results: The proposed method is evaluated using three open datasets: (i) Automatic Nonrigid Histological Image Registration Dataset (ANHIR), (ii) Automatic Registration of Breast Cancer Tissue Dataset (ACROBAT), and (iii) Hybrid Restained and Consecutive Histological Serial Sections Dataset (HyReCo). The target registration error (TRE) is used as the evaluation metric. We compare the proposed algorithm to other state-of-the-art solutions, showing considerable improvement. Additionally, we perform several ablation studies concerning the resolution used for registration and the initial alignment robustness and stability. The method achieves the most accurate results for the ACROBAT dataset, the cell-level registration accuracy for the restained slides from the HyReCo dataset, and is among the best methods evaluated on the ANHIR dataset.

Conclusions: The article presents an automatic and robust registration method that outperforms other state-of-the-art solutions. The method does not require any fine-tuning to a particular dataset and can be used out-of-the-box for numerous types of microscopic images. The method is incorporated into the DeeperHistReg framework, allowing others to directly use it to register, transform, and save the WSIs at any desired pyramid level (resolution up to 220k x 220k). We provide free access to the software. The results are fully and easily reproducible. The proposed method is a significant contribution to improving the WSI registration quality, thus advancing the field of digital pathology.

1. Introduction

The automatic registration of whole slide images (WSIs) is crucial for several downstream applications in digital pathology: (i) the multimodal diagnosis and prediction using artificial intelligence (AI) techniques [1], (ii) the annotation transfer between consecutive or restained

slides to decrease the annotation time, and thus the related costs [2], (iii) the 3-D reconstruction of consecutive slides to create the volumetric representation [3].

It is challenging to propose an automatic, robust, and efficient WSI registration method that does not require fine-tuning or hyperparameter adaptation to a particular dataset. There are several associated chal-

* Corresponding author.

E-mail address: wodzinski@agh.edu.pl (M. Wodzinski).

<https://doi.org/10.1016/j.cmpb.2024.108187>

Received 5 February 2024; Received in revised form 5 April 2024; Accepted 17 April 2024

Available online 22 April 2024

0169-2607/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).




Fig. 1. Exemplary registration pairs from the ACROBAT, ANHIR, and HyReCo datasets. Note the clinical quality of the ACROBAT samples (without any initial preprocessing and numerous artifacts), large initial misalignment in the ANHIR dataset, and the good quality of the HyReCo samples.

lenges: (i) the slides are acquired in various medical centers worldwide using different equipment, stains, and acquisition protocols, resulting in high heterogeneity and varying quality, (ii) no assumption about the relative orientation of the registered WSIs can be made leading to the necessity of large receptive fields of the learning-based methods, (iii) the requirement of relatively low computational complexity which is in contrast to the resolution of the WSI images (that is usually at the level of more than $200k \times 200k$ pixels).

The heterogeneity of data in digital pathology is enormous. The images are acquired in biopsy or resection procedures from nearly any part of not only the human body, but also other species. The samples are being prepared according to standards from different acquisition centers and digitized using various available scanners. The quality of the samples varies strongly with a large number of potential artifacts [4,5]. All these factors lead to the requirement of high generalizability of the digital pathology algorithms. A visualization of registration pairs acquired in different medical centers is shown in Fig. 1. Unfortunately, the current state-of-the-art methods in WSI registration still struggle with the necessity of fine-tuning for learning-based solutions or the requirement of hyperparameters adjustment for the classical, iterative methods.

In digital pathology, in contrast to radiology, no assumption about the initial orientation of the registered images can be made. The same observation applies to the initial spatial offset. Therefore, traditional intensity-based methods usually fail making the initial alignment of WSIs considerably challenging. The problem is even more impactful for the learning-based methods resulting in the requirement of large receptive fields, limiting the usefulness of CNN-based architectures.

The resolution of WSIs makes it difficult to perform efficient registration. The WSIs are usually saved in pyramidal format with the highest resolution above $200k \times 200k$ pixels. That makes the WSIs orders of magnitude larger than the high-resolution computed tomography (CT) or magnetic resonance (MR) 3-D volumes, leading to significant computational complexity. Considering that the registration is usually the initial processing step performed before the downstream task, the registration time should be relatively low. For the consecutive slides the problem can be mitigated by reasonable downsampling, however, for

the restained slides, the high resolution should be maintained to enable the cell-level registration accuracy.

Contribution: In this article, we propose an automatic, robust, efficient, and accurate method for the registration of whole slide images (WSIs). We evaluate the method using three open datasets and show a significant improvement with respect to the current state-of-the-art in WSIs registration. We openly release the method source code and incorporate it in the DeeperHistReg framework allowing other researchers to directly use it in their research [6,7]. The method has superior generalizability, it does not require any fine-tuning or retraining to a particular dataset. Moreover, it can be used to register other kinds of microscopic images, without any changes to the parameters setup.

2. Related work

There are numerous contributions in the field of automatic WSI registration, both in terms of open data availability, as well as methodological advances.

Currently there are three major open datasets dedicated to evaluating advances in WSI registration. The first and probably the most recognized is the ANHIR dataset introduced during a challenge organized jointly with the IEEE ISBI 2019 conference [1]. The dataset is heterogeneous, containing several tissue types stained with various dyes. However, the dataset is of relatively good quality, without significant artifacts or distortions, which is rare in practical scenarios. The second dataset was released during the ACROBAT challenge organized jointly with the MICCAI conference [8,9]. The dataset consists of several hundred breast cancer clinical slides acquired using various stains. The quality of the slides varies strongly, from good quality ones to cases with foreign bodies, low contrast, and missing or additional tissues. The third notable dataset is called HyReCo [10,11] and contains both consecutive and restained slides. The dataset consists of high-resolution and good-quality sections and enables researchers to evaluate the influence of the registration resolution on the restained samples. Other big datasets could be useful for evaluating WSI registration, like MSHIR [12] or COMET [13], however, up-to-date they remain closed to the scientific community.

Medical image registration is a research area with a significant number of notable contributions [14–24]. However, due to the challenges specific to the WSIs, the general state-of-the-art registration methods usually fail [25]. Therefore, contributions proposed specifically to WSIs are required. The two most successful methods are named HistokatFusion [26] and VALIS [27]. Both of them are based on classical, iterative image registration. The HistokatFusion is ready-to-use closed software implementing both initial and nonrigid intensity-based iterative procedures with normalized gradient fields (NGF) as a similarity measure and B-Splines as the deformation model. The method was the winner of the ANHIR challenge and was among the best-performing ones during the ACROBAT challenge [8,9]. The VALIS library is fully operational software using SimpleITK as the registration backend. The main limitation of VALIS is a significant computational complexity limiting the usefulness to consecutive slides, with limited applicability to registering the restained ones. There are also numerous other scientific contributions, however, not directly useful in practice. Among the classical ones, it is worth mentioning: (i) the UPENN method based on the “Greedy” tool dedicated to diffeomorphic registration [28,29], (ii) the multistep AGH method based on the Demons algorithm with modality independent neighborhood descriptor [30], (iii) the intensity-based registration driven by unsupervised classification of structural similarity [31], (iv) the patch-based nonrigid registration method enabling registration of WSIs at any resolution [32], (v) a combined feature- and quad-tree based method for robust affine alignment [33].

There are several deep learning-based contributions to the WSI registration. The TUB method based on a modified volume tweening network was among the first ones [34] and scored considerably well in the ANHIR challenge. The method was not on the podium, however, confirmed that the learning-based solutions can significantly speed up the registration process. Another worth mentioning contribution is the DeepHistReg method that introduced a patch-based approach to the learning-based WSI registration, allowing one to perform the learning-based registration at any desired resolution [35,36]. One of the most accurate solutions is based on CNN guided by structural features, with a fully unsupervised training pipeline [37]. Another work attempted to use data-driven descriptors learned by CNN architecture and compared them to the hand-crafted descriptors [13]. Additionally, the proposed method was combined with the HistokatFusion deformable registration resulting in significant improvement. Recently, several researchers attempted to address the initial alignment of WSIs. Several works used transformers to increase the receptive field [38,39], thus solving one of the limitations of CNN-based affine registration networks [40]. Another work proposed an end-to-end affine registration method by utilizing weak annotations [12]. The researchers reformulated the problem into the classification of the initial misalignment to reduce the requirement of a large receptive field of the following affine registration network.

There are several limitations of the current state-of-the-art registration algorithms: (i) the majority of the methods are evaluated using just a single dataset, rarely coming from different acquisition centers, (ii) the algorithms require fine-tuning or hyperparameters adjusting whenever samples from new datasets are going to be registered, (iii) the initial alignment is not robust resulting in numerous failures, and (iv) the source code, which is also rarely released, usually do not allow other researchers to directly replicate the experiments or use the methods in practical applications.

In this work, we solve the limitations: (i) we evaluate the proposed method on several datasets acquired in different medical centers and show good stability and generalizability, (ii) the proposed method does not require any fine-tuning, retraining or parameter adjustment to new datasets, (iii) the initial alignment method works correctly for almost any case, and (iv) we openly release the source code and incorporate it in the DeeperHistReg framework allowing others to use the algorithm in research and practical applications.

3. Methods

3.1. Overview

The method consists of three steps: (i) the preprocessing, (ii) the initial alignment, and (iii) the deformable registration. The goal of preprocessing is to convert the input images to the same color space and to resample them to resolutions desired at subsequent registration steps. The initial alignment finds the global affine transformation to roughly match the images. The deformable registration calculates the dense displacement field to recover more complex and nonrigid deformations. All these steps are described in the following sections and presented in Fig. 2.

3.2. Preprocessing

The preprocessing starts with an automatic estimation of the pyramid level that is then loaded and resampled based on the resolution desired at the nonrigid registration step. The resolution level with the closest higher resolution to the desired one is chosen and then the image is loaded and resampled. Such an approach limits the computational complexity since the initial downsampling is performed using a relatively similar resolution. The initial loading and resampling are implemented using the PyVips [41] patch-based approach. The initial downsampling is necessary because the input images may be stored in a format not supporting pyramidal representation. For such cases performing the patch-based downsampling to each pipeline step would significantly increase the registration time.

Next, the images are converted to grayscale, normalized to the [0-1] range, and both the images are processed with the CLAHE algorithm. Finally, the images are resampled again to the resolutions defined in the configuration of the initial alignment step, preceded by Gaussian smoothing to avoid aliasing. Since the interpolation performs only downsampling and is applied at most two times, the interpolation error is negligible.

3.3. Initial alignment

The initial alignment is based on pretrained SuperPoint and SuperGlue methods [42,43]. The SuperPoint is responsible for the feature extraction and the SuperGlue for the feature pairing. The paired features are then used to calculate the affine transformation matrix using the least squares method. Importantly, the SuperPoint and SuperGlue algorithms are not fine-tuned to the histology data. The trained weights based on self-supervised learning are directly used to calculate and match the features. It confirms the superior generalizability of the SuperPoint feature extractor.

Nevertheless, the SuperPoint and SuperGlue methods are not orientation and scale invariant. Therefore, we extend these methods by embedding them in a pipeline that firstly resamples the source and target images to a given resolution (to extract features at different scales) and secondly rotates the source image with given angles to improve the robustness with respect to the initial orientation. The feature extraction and pairing are then performed repeatedly for each scale and orientation. The result with the highest number of matches calculated by the SuperGlue is chosen as the final outcome. The exhaustive search for the best scale and initial orientation is crucial to increase the robustness of the initial alignment and to make it useful for various datasets. Even though for high-quality datasets a single scale and orientation may be enough, for cases with poor quality and artifacts the exhaustive approach positively impacts the robustness. The initial alignment procedure is shown in Fig. 3.

3.4. Nonrigid registration

The nonrigid registration is implemented as a multilevel instance optimization-based iterative algorithm. The objective function is de-

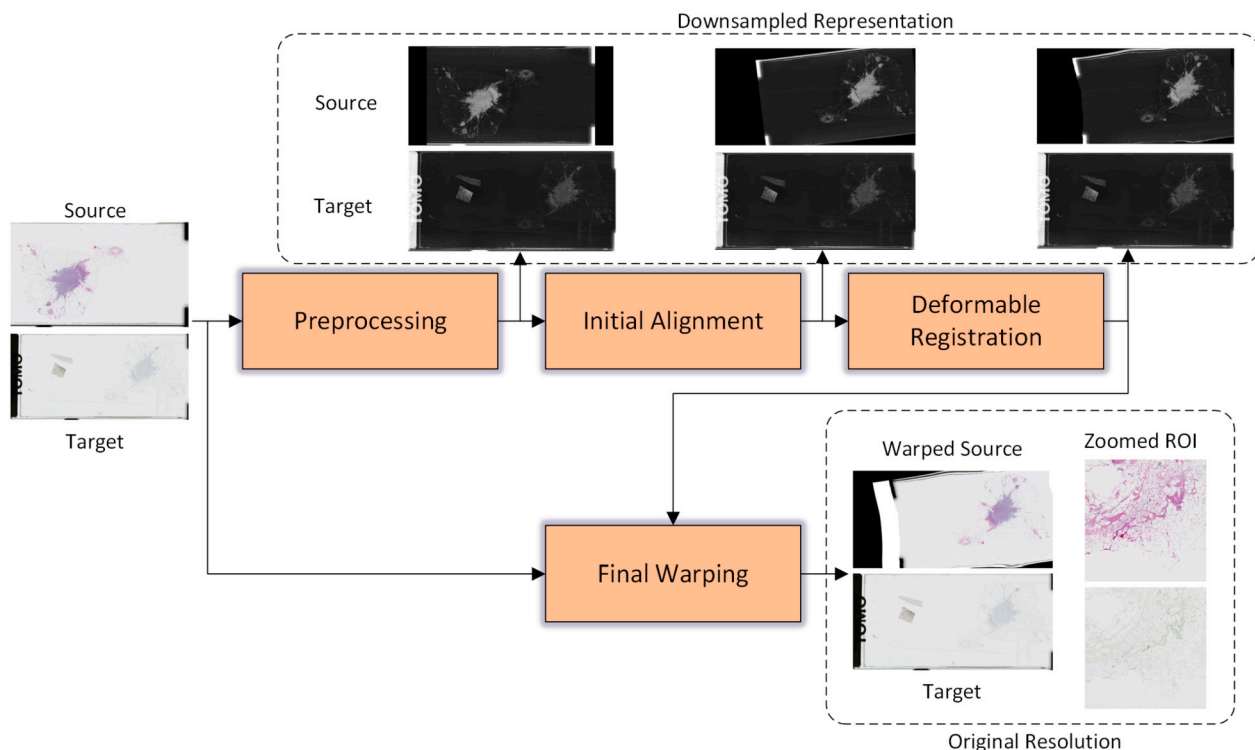


Fig. 2. The pipeline of the proposed method. The processing starts with the preprocessing consisting of resampling and color normalization, followed by the robust initial alignment, and finally the accurate deformable registration. The final warping is performed in a separate module at the original resolution.

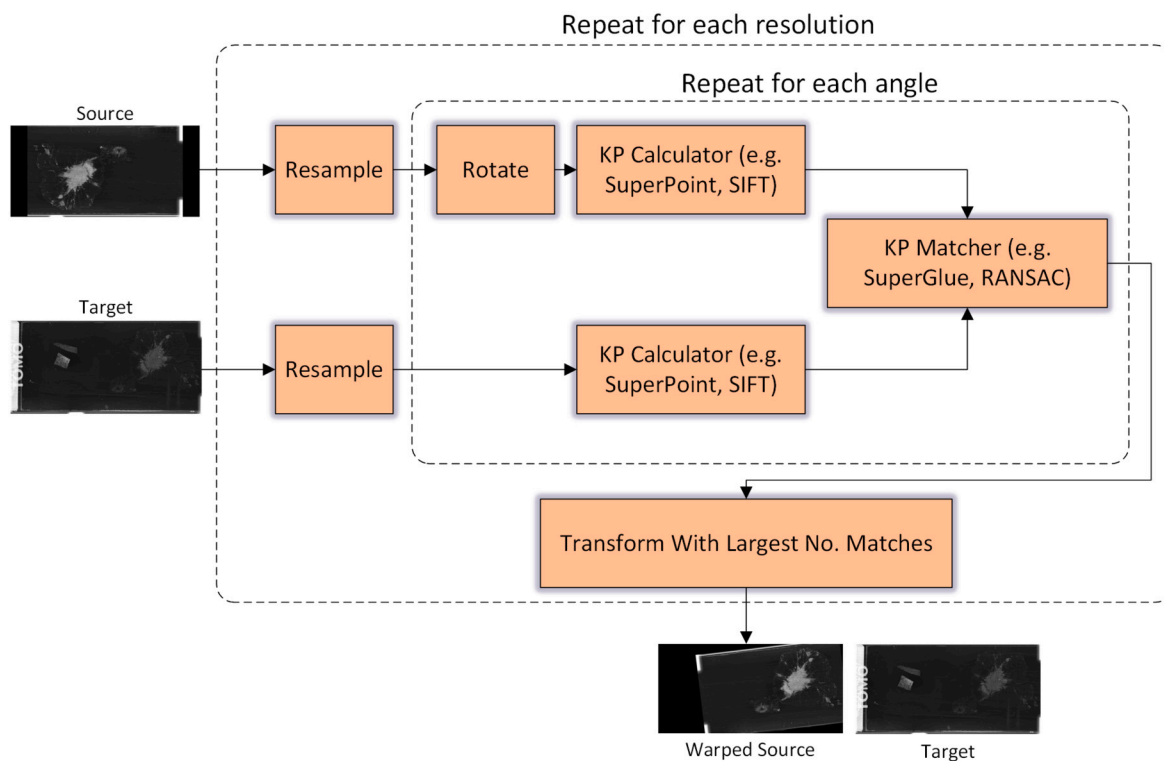


Fig. 3. The multi-scale and multi-angle initial alignment pipeline. Since the keypoint extractors are never perfectly scale- and rotation-invariant, such an exhaustive procedure improves the alignment robustness. The transform with the largest number of matches is chosen as the final one, choosing from all transformations calculated across the predefined scales and rotation angles. The procedure improves the alignment quality especially for outliers which would be otherwise incorrectly registered, potentially even decreasing the results quality with respect to the initial misalignment.

defined as the weighted sum of the local normalized cross-correlation (NCC) as the similarity measure, and the diffusive regularization as the regularizer, defined as:

$$O_{REG}(S, T, u) = NCC(S_i \circ u_i, T_i) + \theta_i Reg(u_i), \quad (1)$$

where S_i, T_i are the source and target slides at the i -th resolution level respectively, u_i is the calculated displacement field, θ_i denotes the regularization coefficient at i -th resolution level, NCC denotes the normalized cross-correlation, Reg is the diffusive regularization, \circ denotes the warping operation, and N is the number of resolution levels.

We decided to use the local NCC as the similarity measure because the local contrast changes can be accurately captured by the cross-correlation, unlike as in e.g. CT/MRI registration. Moreover, our recent studies confirmed the great potential of NCC for histology registration, outperforming multimodal metrics like Modality Independent Neighbourhood Descriptor, Mutual Information, or Normalized Gradient Fields [44,8,30,35].

The nonrigid registration starts with warping the source image using the transformation found in the initial alignment step. The warping is applied by transforming the affine matrix into displacement field and performing bicubic interpolation. We have chosen this approach because initializing the displacement field with the affine transformation impacts the diffusive regularization term and expects larger receptive field from the optimizer. In theory, we could avoid that by using the curvature regularization, however, the initial ablations confirmed that the curvature regularization decreases the results quality and is more prone to folding than the diffusive one. Moreover, the proposed initial alignment step is robust and effective, therefore, it can be enforced as the initial status of the deformable registration. The result of the nonrigid registration is saved and can be used to transform the source image at any desired resolution level by using the B-Splines-based upsampling.

All the hyperparameters are adjustable to each resolution level separately. The relative weight of the similarity measure and regularizer can be set differently for each resolution which makes sense since at higher resolutions the weight of regularization should be increased to avoid tissue folding. The same applies to the learning rate that is higher at coarse resolutions and then decreases when the resolution increases. The number of iterations is predefined in the registration settings, without any early stopping mechanism. We decided to use fixed number of iterations because tuning this parameter would require combined grid search together with an optimal regularization coefficient and learning rate which is beyond the computational power of our infrastructure and evaluation possibilities of the Grand-Challenge platform. Moreover, such hyperparameter tuning would require to be performed separately for each dataset. Therefore, we decided that it is better idea to set the parameter fixed and show the robustness of the proposed method because its accuracy outperforms other methods without any dedicated hyperparameter adjustment.

3.5. Datasets

The evaluation is performed using three open datasets: (i) Automatic Nonrigid Histological Image Registration Dataset (ANHIR), (ii) Automatic Registration of Breast Cancer Tissue Dataset (ACROBAT), and (iii) Hybrid Restained and Consecutive Histological Serial Sections Dataset (HyReCo). Exemplary registration pairs from each dataset are shown in Fig. 1.

The ANHIR dataset consists of 481 image pairs and is split into 230 training and 251 evaluation pairs [1,45–48]. Noteworthy, the ANHIR dataset and the evaluation platform are biased and almost half of the dataset has to be excluded from the evaluation. Even though the results in the system are aggregated for both training and evaluation subsets, the annotations for the training set are openly released. Therefore, we exclude all the training pairs ($N=230$) from the ANHIR dataset in our evaluation to make a fair comparison. The evaluation is performed

server-side, automatically using scripts provided by the ANHIR organizers [44]. The inter-observer variability is 0.05% of the image size which can be used as the indication of the best performance possible for the automatic registration methods, below which the methods become indistinguishable [1]. The images are provided in .jpg and .png format, without the resolution pyramid and metadata. There are 8 different tissue types which were stained using 10 different dyes. A full description of the dataset can be found in [1].

The ACROBAT dataset consists of consecutive female breast cancer slides. There are 750, 100, and 303 training, validation, and test cases respectively. Since the cases are stained using Hematoxylin&Eosin (H&E), ER, KI67, PGR, and HER2, there are 3406 training, validation, and test WSIs respectively (the training cases are stained using multiple dyes). The WSIs were acquired using three scanners, one NanoZoomer S360, and two different NanoZoomer XRs. The images are saved and released in .tiff format using 10X magnification, downscaled from the original 40X. The dataset was annotated by 13 experienced individuals on the original 40X scale. In this paper, the evaluation is performed on the 100 hidden validation pairs using the ACROBAT evaluation platform [49]. Unfortunately, the 303 test cases are hidden and not released publicly. The description of the ACROBAT dataset is available in [8,9].

The HyReCo dataset, released in 2021, is unique by providing both consecutive and restained sections. The dataset was acquired at the Radboud University Medical Center and annotated by experienced researchers from Fraunhofer MEVIS. The dataset consists of 9 high-resolution consecutive sections, each stained with H&E, CD8, CD45, and KI67. For each section, there are between 11 and 19 annotated landmarks, resulting in 138 annotations per stain and 690 annotations in total. The landmarks were placed and verified by two experienced researchers. Additional slides were restained using the PHH3, resulting in 54 restained H&E/PHH3 pairs with about 43 annotations per restained pair. The slides are saved and released in BigTIFF format with corresponding landmarks released in a CSV file containing world coordinates in mm. The HyReCo dataset, together with its full description, is available in [10,11].

3.6. Experimental setup

We compare the proposed method to the state-of-the-art solutions. The target registration error (TRE) is used as the evaluation metric. The TRE calculates the Euclidean distance between corresponding landmarks, manually annotated by experts. For the ACROBAT and HyReCo datasets, we report the TRE in physical units while for the ANHIR dataset, we normalize it by dividing by the image diagonal (rTRE), following the convention introduced by the ANHIR dataset creators and other methods reporting the results [1,25].

We perform several ablation studies to confirm the method's robustness and the impact of the most crucial hyperparameters. The experiments verify (i) the influence of each registration step, (ii) the difference between the SuperPoint/SuperGlue method and the traditional SIFT/RANSAC approach [30], (iii) the influence of the number of rotation angles and resolutions on the initial alignment stability, (iv) the impact of the resolution on the nonrigid registration quality, both for the consecutive and restained slides. Any claim in the Results or Discussion sections about significant improvement means that the claim is supported by Wilcoxon signed-rank test with p-value lower than 0.01.

All the algorithms and experiments are implemented using the PyTorch library with extensions based on PyVips [41]. The experiments were performed using NVIDIA RTX 3090 GPU with 24 GB of dedicated memory. The proposed method is integrated into the DeeperHistReg framework and openly released [7]. The method does not require any fine-tuning or retraining to be used for other datasets.

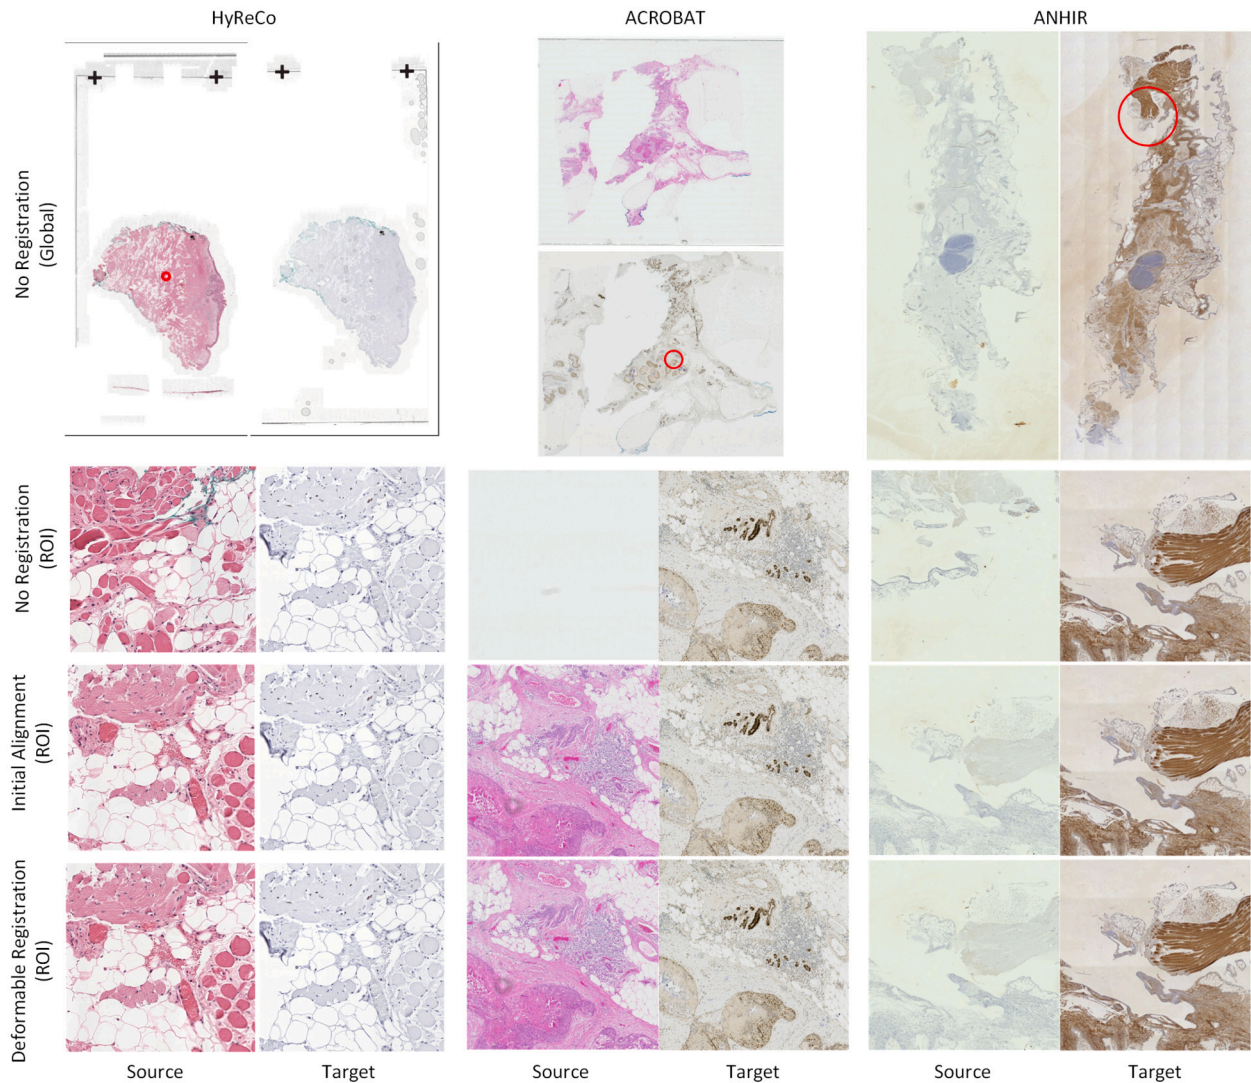


Fig. 4. Visual registration results for exemplary cases from all three datasets. The results are presented on a chosen ROI after subsequent registration steps. The global overview for ANHIR case is rotated, and the ACROBAT source image is cropped, for the presentation clarity. Note the difference between the registration quality of restrained (HyReCo) and consecutive slides (ACROBAT/ANHIR).

4. Results

4.1. Affine vs nonrigid registration

Fig. 5 presents the mean and median TRE distribution after each registration step. The ticks in violin plots present the mean and median value respectively. It can be noted that the registration improves the alignment for majority of registration pairs, without a single case that could be considered a failure (for which mean/median TRE significantly increases after the registration). The Median-Average rTRE before the registration, after the affine alignment, and after the nonrigid registration are 0.0683, 0.0050, 0.0029 for the ANHIR dataset. For the ACROBAT dataset, the Median TRE in μm are 14424.35, 409.86, 137.27 after subsequent registration steps. Finally, for the HyReCo datasets, the Median TRE in μm are 174.90, 4.96 (consecutive), and 12.24, 0.59 (restrained) after the initial alignment and the deformable registration, respectively. It confirms significant improvement after each registration step ($p\text{-value} < 0.01$). Exemplary visualization of the registration after each algorithm step is shown in Fig. 4. Importantly, the percentage of foldings, defined as the ratio of a number of pixels with negative Jacobian to all pixels is below 0.1% for all nonrigid experiments.

4.2. Comparison between SuperPoint-SuperGlue and SIFT-RANSAC

One of the ablations confirms the impact of the novel SuperPoint/SuperGlue initial registration to the state-of-the-art SIFT/RANSAC methods. The results are presented in Fig. 6. One can note that the learning-based methods are more robust without failures which are quite common for the SIFT/RANSAC method. The Median-Average rTRE for SuperPoint/SuperGlue, and SIFT + RANSAC are equal to 0.0050, 0.0059 for the ANHIR dataset. For the ACROBAT dataset, the Median-Average TRE in μm for those methods are equal to 409.86, 852.43. The comparison is skipped for the HyReCo dataset due to large number of failures of SIFT/RANSAC method accounting for more than 10% of the registered pairs. Moreover, the proposed learning-based method is computationally more efficient (~ 15 vs ~ 27 seconds), however, the effect is negligible when considering the time required for further nonrigid registration and final image warping.

4.3. Stability concerning the number of angles and scales

The outcomes of the experiment related to the influence of the number of scales are shown in Table 1, and the impact of the rotation angle step is shown in Table 2. The difference between using a single scale and using eight scales is notable (0.0128 vs 0.0057 in terms of the rTRE




Fig. 5. The quantitative results in terms of the target registration error (TRE) after the subsequent registration steps. The HyReCo results are reported for the consecutive slides only. Note the significant improvement (p -value < 0.01) after each registration step.

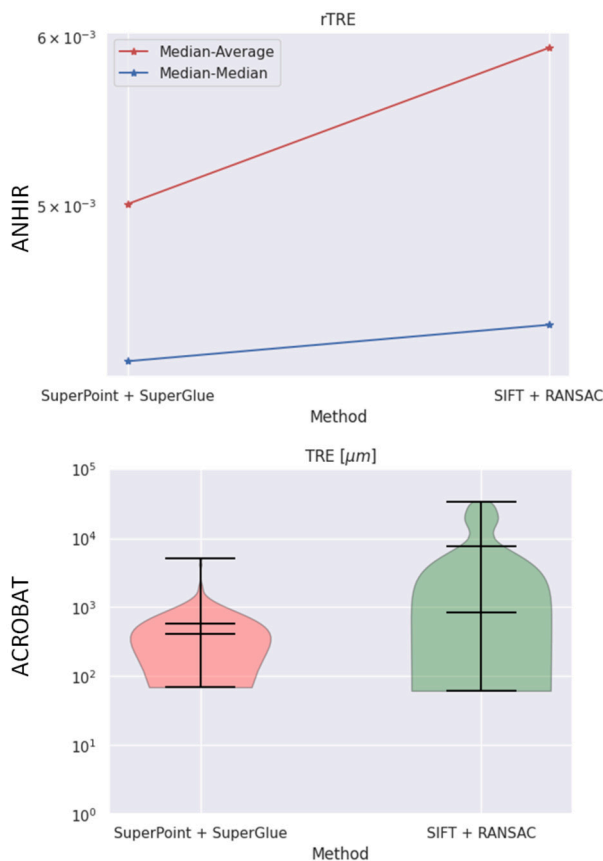


Fig. 6. Comparison of the SuperPoint/SuperGlue-based approach to the SIFT/RANSAC method. The Average-Median and Average-Average for ANHIR dataset are skipped due to large number of failures of the SIFT/RANSAC method (more than 10%). The comparison is skipped for the HyReCo dataset due to large number of failures of SIFT/RANSAC method (more than 10%).

for the ANHIR dataset, 470.98 vs 396.38 μm for the ACROBAT dataset, and 288.47 vs 232.64 μm for the HyReCo dataset). The small angle step (30 degrees vs 180 degrees) is also influential (0.0057 vs 0.0110 in $r\text{TRE}$ for the ANHIR dataset, 396.38 vs 417.30 μm for the ACROBAT dataset), with the exception of HyReCo datasets where the best alignment was achieved for angle step equal to 120. The differences in terms of both Average-Median and Average-Average TRE are significant for both angle- and scale-related ablations (p -value < 0.01) for all the datasets. It can be noted that the multi-scale and multi-angle extension is not crucial to achieving considerable performance, better than

Table 1

The effect of the number of scales on the stability of the initial alignment based on the SuperPoint and SuperGlue descriptors. In case of the HyReCo dataset, the results are reported for the consecutive slides.

ANHIR		
#Scales	Average-Median $r\text{TRE}$	Average-Average $r\text{TRE}$
1	0.0128	0.0140
2	0.0092	0.0104
3	0.0083	0.0095
4	0.0068	0.0080
5	0.0063	0.0074
6	0.0061	0.0072
7	0.0058	0.0070
8	0.0057	0.0071
ACROBAT		
#Scales	Average-Median TRE [μm]	Average-Average TRE [μm]
1	470.98	656.73
2	452.78	946.40
3	434.55	629.83
4	466.00	614.22
5	446.39	562.28
6	410.68	541.34
7	404.90	537.69
8	396.38	540.87
HyReCo		
#Scales	Average-Median TRE [μm]	Average-Average TRE [μm]
1	288.47	321.82
2	275.28	312.13
3	289.18	323.75
4	255.69	291.04
5	263.44	298.31
6	238.13	268.99
7	234.92	266.39
8	234.64	263.08

the multi-scale/multi-angle SIFT/RANSAC method. However, it ensures as high robustness as possible, without failures of the initial alignment.

4.4. Impact of the resolution on consecutive and retained slices

The influence of the registration resolution on the nonrigid registration step is shown in Fig. 7 for each dataset. The Fig. 7 also presents the comparison between the consecutive and the retained slices from the HyReCo dataset. The comparison between the consecutive and retained slides is reported only for the HyReCo cases because other datasets consist of only consecutive slides. It can be noted that the resolution used for registration is less impactful for the consecutive slides and easily saturates (insignificant changes), however, for the retained ones further increase of the resolution continues to significantly (p -

Table 2

The effect of the rotation angle step on the stability of the initial alignment based on the SuperPoint and SuperGlue descriptors. In case of the HyReCo dataset, the results are reported for the consecutive slides. Note that the results for the HyReCo dataset are strongly influenced by the fact that in several slides there are more than one tissue.

ANHIR		
Angle Step	Average-Median rTRE	Average-Average rTRE
30	0.0057	0.0071
60	0.0079	0.0093
90	0.0079	0.0093
120	> 0.1	> 0.1
150	0.0084	0.0098
180	0.0110	0.0124
ACROBAT		
Angle Step	Average-Median TRE [μm]	Average-Average TRE [μm]
30	396.38	540.87
60	409.86	574.50
90	401.58	546.31
120	490.42	993.30
150	401.587	588.47
180	417.30	561.56
HyReCo		
Angle Step	Average-Median TRE [μm]	Average-Average TRE [μm]
30	234.64	263.08
60	223.58	253.28
90	223.58	253.28
120	52.45	65.18
150	138.85	156.22
180	223.58	253.28

value < 0.01) improve the registration quality (MTRE equal to 2.20, 1.33, 1.01, 0.59 μm for 1024², 2048², 4096², 8192² respectively). Therefore, the increase of the registration resolution beyond 8192² would be beneficial only for the restained slides, however, such a change would require a patch-based approach to the registration due to the GPU memory limitations.

4.5. Comparison to the state-of-the-art

The comparison to the state-of-the-art methods using the ANHIR and ACROBAT datasets are shown in Table 3 and Figs. 8, 9 respectively. The results for the state-of-the-art methods are acquired by downloading the best submission for each method from the ANHIR/ACROBAT challenge platforms. For the ANHIR dataset, we report results only using the evaluation cases because the annotations for the training set are openly available which biases towards the learning-based methods utilizing the annotations during training. The proposed method achieves only slightly worse results compared to the best-performing ones (0.0017 vs 0.0015 for Median-Median rTRE and 0.0044 vs 0.0034 for the Average-Average rTRE). For the ACROBAT dataset, we report the results using the closed validation set. The method has the lowest Average-Median TRE among all the proposed methods (331.01 μm vs 385.58 μm achieved by HistoKat Fusion, the second most accurate registration method.). For the HyReCo dataset, we compare the proposed method to the HistoKatFusion in Table 4 by using the results reported in the respective publication [10]. The method outperforms the HistoKatFusion in the deformable registration (0.90 μm , 4.96 μm for the restained, and consecutive slides respectively, versus 0.90 μm and 5.30 μm achieved by the HistoKat Fusion).

5. Discussion

The proposed method produces accurate registrations for all three datasets, as shown in Fig. 5. Importantly, the method uses the same configuration for all registration pairs and does not require adjusting

the parameters to a particular dataset. The method is the most accurate on the ACROBAT benchmark, among the best-performing ones for the ANHIR dataset, and comparable to the HistoKatFusion on the HyReCo slides. The average registration time, for the 8192² resolution is below 2 minutes.

Interesting observations, confirming the results obtained by the creators of the HyReCo dataset [10], can be made when comparing the registration accuracy between the consecutive and restained slides (Fig. 7). Increasing the registration resolution substantially decreases the TRE for the restained slides, however, for the consecutive ones the impact is less significant, and further increase of the resolution does not improve the results. Even though the nonrigid registration still significantly improves the accuracy when compared to the initial affine alignment, one should keep in mind that achieving the cell-level registration accuracy for the consecutive slides may be challenging if not impossible. The registration accuracy strongly depends on the quality of the slides, therefore to achieve as accurate registration as possible, it is important to prepare the slides carefully.

The initial alignment method is stable and robust, as presented in Fig. 5 and Tables 1, 2. The failures are not significant for the HyReCo dataset, and there are no failures for the ANHIR and ACROBAT datasets. The initial alignment decreases the initial TRE for all registration pairs in ANHIR and ACROBAT datasets and increases the TRE only those pairs in the HyReCo dataset that contain multiple tissue sections. Tables 1, 2 confirm the impact of the multi-scale and multi-angle approach. Even though it is not necessary for high-quality slides, it improves the stability of the low-quality ones. The results also show that the tissues are usually rotated by 0, 45, 90, 135, or 180 degrees. The experiments with the angle step equal to 60 or 120 degrees resulted in several outliers. One could argue that the performance of the initial alignment for the HyReCo dataset is suboptimal, however, it turns out that it does not impact the following deformable step for both consecutive and restained slides. Moreover, the suboptimal performance is caused by factors that should be addressed during the quality control step.

The algorithm has several advantages. The background segmentation [50,51], which is a preprocessing step for several other contributions, is not necessary. The need for background segmentation usually decreases the registration methods' robustness and generalizability because the segmentation is nowadays based on learning-based techniques. As far as Hematoxylin&Eosin (H&E) images are concerned, the problem may be considered solved, however, the current state-of-the-art methods still struggle with other stains and need further improvement. Therefore, eliminating the necessity of background segmentation by using the feature-based approach to the initial alignment increases the robustness and eliminates the error-prone part of the intensity-based pipeline.

Another advantage is connected with the method's high robustness and generalizability. Even though parts of the method are learning-based, they do not require fine-tuning to new datasets. The instance optimization approach by design captures the heterogeneity of WSIs by using the inference time optimization and the SuperPoint/SuperGlue methods are characterized by high stability and reliability.

We openly release the method's source code and incorporate it in the DeeperHistReg framework allowing one to reproduce the results and use the method in their research. The method is ready to use, one needs to only clone the repository, set up paths to the images to be registered, and run the registration. This is a significant contribution to the field of digital pathology since other registration methods are often closed source or have a high entry barrier requiring the registration expertise from the end-user. Importantly, the DeeperHistReg framework allows one to not only perform the registration but also to load any OpenSlide compatible format, transform the WSIs, and save the output at any desired resolution level. Our other experiments confirmed also the usability to other microscopic images, e.g. unstained, DAPI-stained, or ExM images, without the need for the parameters adjustment.

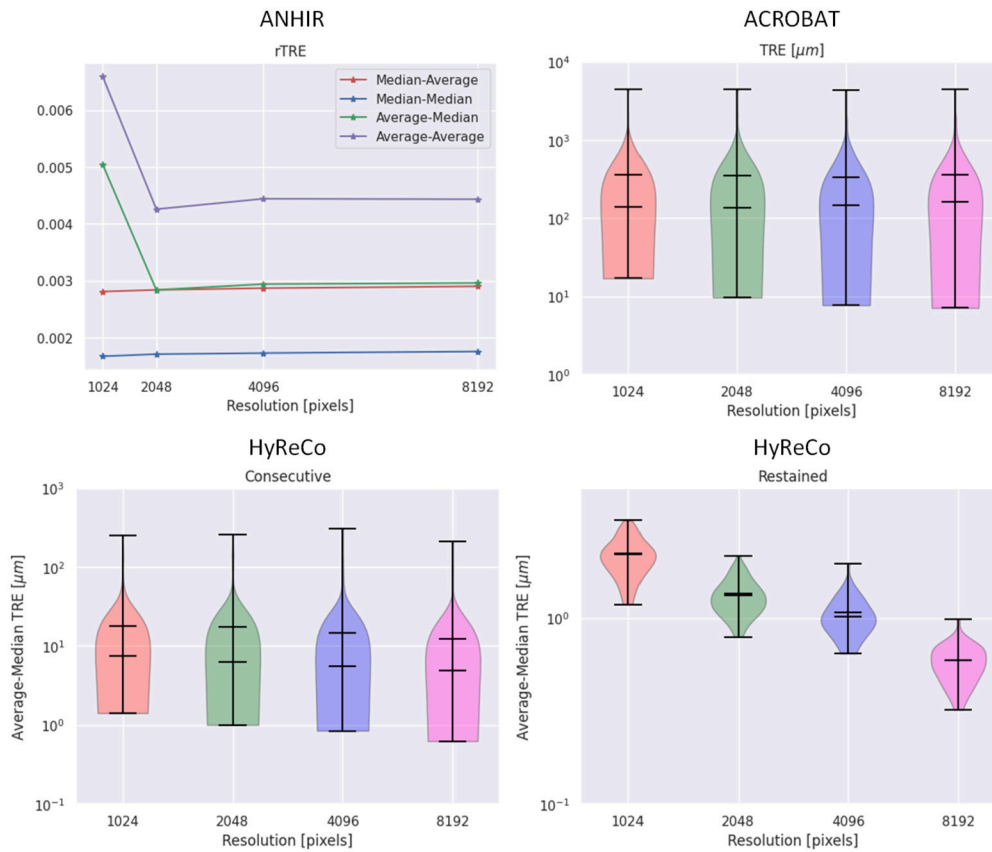


Fig. 7. The influence of the image resolution on the deformable registration quality. Note that for the consecutive slides increasing the resolution has negligible positive effect (ACROBAT, HyReCo consecutive) or even decreases the performance (ANHIR), however, for the restrained slides, further increase of the resolution has a significant impact (p-value < 0.01) on the registration quality.

Table 3

Comparison of the proposed method to other best-performing methods evaluated on the ANHIR dataset. The results are taken directly from the ANHIR evaluation platform for the best submission from each participant. The results are reported for the evaluation subset of the ANHIR dataset. The Med-Med and Med-Avg show the ability of the methods to recover fine details while the Avg-Med and Avg-Avg present the robustness with respect to difficult cases.

Method	Med-Med rTRE ↓	Med-Avg rTRE ↓	Avg-Med rTRE ↓	Avg-Avg rTRE ↓
Proposed	0.0017	0.0029	0.0029	0.0044
DDL	0.0015	0.0023	0.0039	0.0051
SFG	0.0016	0.0025	0.0070	0.0083
SIAT	0.0016	0.0024	0.0022	0.0034
VALIS	0.0017	0.0025	0.0059	0.0073
MEVIS	0.0018	0.0027	0.0029	0.0043
HistoReg	0.0019	0.0029	0.0029	0.0042
MASI	0.0021	0.0031	0.0048	0.0064
AGH	0.0017	0.0031	0.0057	0.0071
USTC	0.0024	0.0037	0.0104	0.0117
CASIA	0.0026	0.0037	0.0041	0.0054

Table 4

Comparison of the proposed method to the HistokatFusion on the HyReCo consecutive and restrained slides in terms of the median TRE (MTRE). Note that even though the initial alignment is worse than for the HistokatFusion, the following deformable refinement still significantly improves the results.

Method	MTRE Affine [μm] ↓		MTRE Deformable [μm] ↓	
	Restained	Consecutive	Restained	Consecutive
Proposed	12.24	174.90	0.59	4.96
HistokatFusion [10]	1.70	20.20	0.90	5.30

The method has also several limitations. The disadvantage of the deformable step is its limited ability to recover large non-rigid displacements, as confirmed in Table 3. It is connected with the limitations of instance optimization that inherently has a narrow receptive field, even in multilevel scenarios. As far as good quality samples are concerned, especially restrained ones, the effect is negligible, as presented in Table 4, where the proposed method even outperformed the HistokatFusion. However, for low-quality consecutive slides, the power to recover large deformations could further improve the outcomes, thus further increasing the registration quality for ANHIR and ACROBAT datasets. Another limitation of the method is lack of downstream validation on classification or segmentation tasks, however, this is a topic for sepa-




Fig. 8. Comparison to the best performing methods on the evaluation subset of the ANHIR dataset. The results are taken directly from the ANHIR evaluation platform for the best submission from each participant. The challenge platform reports only the aggregates.




Fig. 9. Comparison to the best performing methods on the hidden validation subset of the ACROBAT dataset. The results are taken directly from the ACROBAT evaluation platform for the best submission from each participant.

rate studies. Additionally, the method could benefit from a complete GPU-based implementation, starting from loading the images directly to the GPU memory and then performing all the steps using just one device. Such an approach would not change the registration accuracy, however, would be beneficial from the practical point of view. Nevertheless, the transition to a GPU-only implementation is a pure technical contribution with limited scientific value.

One could argue that the proposed modifications to the initial alignment based on the SuperPoint and SuperGlue architectures add additional computational complexity. It is indeed true, however, the time required for the initial alignment is negligible when compared to the time required to perform the following nonrigid registration or to warp the source image after the registration. The initial alignment takes no more than 20 seconds while warping the source image after the registration at the highest resolution level requires between 2 to 10 minutes.

In future work, we plan to extend the method in several ways. First, we forecast that fine-tuning the SuperPoint and SuperGlue methods using self-supervised learning on large histology datasets, e.g. from the BigPicture consortium [52] could further improve the robustness of the initial alignment, and allow one to skip the multi-scale and multi-orientation steps. Additionally, it would be interesting to propose variants of the SuperPoint and SuperGlue methods for volumetric data which could be beneficial for the initial alignment in radiology [15].

Another direction we plan to explore is the use of stain normalization and transfer to improve the preprocessing step. The normalization of color spaces of both the source and target images could strongly influence the quality of the deformable registration. Finally, the difficulties of the nonrigid registration to handle complex deformations could be solved by the neural instance optimization of deep models dedicated to recovering large displacements, e.g. LapIRN [16,53].

To conclude, we proposed an automatic, robust, and accurate method dedicated to WSI registration. We evaluated the method using three open datasets: ANHIR, ACROBAT, and HyReCo, performed several ablation studies, and compared it to other state-of-the-art methods. We freely released the source code and incorporated it in the DeeperHistReg framework, allowing other researchers to use the method in their research and to easily reproduce the results. The proposed method is a significant contribution to the field of digital pathology.

CRedit authorship contribution statement

Marek Wodzinski: Writing – review & editing, Writing – original draft, Visualization, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Niccolò Marini:** Writing – review & editing, Writing – original draft. **Manfredo Atzori:** Writing – review & editing, Writing – original draft, Supervision. **Henning Müller:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare no conflict of interest.

Acknowledgements

This project has received funding from the Innovative Medicines Initiative 2 Joint Undertaking under grant agreement No 945358. This Joint Undertaking receives support from the European Union's Horizon 2020 research and innovation program and EFPIA, Belgium (www.imi.europe.eu). The research reflects only the author's view and the Joint Undertaking is not responsible for any use that may be made of the information it contains. Additionally, the research was supported in part by PLGrid Infrastructure. We gratefully acknowledge Poland's high-performance computing infrastructure PLGrid (HPC Centers: ACK Cyfronet AGH) for providing computer facilities and support within computational grant no. PLG/2023/016239.

References

- [1] Jiří Borovec, Jan Kybic, Ignacio Arganda-Carreras, Dmitry V. Sorokin, Gloria Bueno, Alexander V. Khvostikov, Spyridon Bakas, I. Eric, Chao Chang, Stefan Heldmann, et al., ANHIR: automatic non-rigid histological image registration challenge, *IEEE Trans. Med. Imaging* 39 (10) (2020) 3042–3052.
- [2] Philippe Weitz, Viktoria Sartor, Balazs Acs, Stephanie Robertson, Daniel Budelmann, Johan Hartman, Mattias Rantalainen, Increasing the usefulness of already existing annotations through wsi registration, arXiv preprint, arXiv:2303.06727, 2023.
- [3] Yi Song, Darren Treanor, Andrew J. Bulpitt, Derek R. Magee, 3d reconstruction of multiple stained histology images, *J. Pathol. Inform.* 4 (2) (2013) 7.
- [4] Syed Ahmed Taqi, Syed Abdus Sami, Lateef Begum Sami, Syed Ahmed Zaki, A review of artifacts in histopathology, *J. Oral Maxillofac. Pathol.* 22 (2) (2018) 279.
- [5] Andrew Janowczyk, Ren Zuo, Hannah Gilmore, Michael Feldman, Anant Madabhushi, HistoQC: an open-source quality control tool for digital pathology slides, *JCO Clin. Cancer Inform.* 3 (2019) 1–7.
- [6] M. Wodzinski, DeeperHistReg library, <https://github.com/MWod/DeeperHistReg/>. (Accessed 14 December 2023), 2023.
- [7] Marek Wodzinski, Niccolò Marini, Manfredo Atzori, Henning Müller, DeeperHistReg: robust whole slide images registration framework, in: Review, 2024.
- [8] Philippe Weitz, Masi Valkonen, Leslie Solorzano, Circe Carr, Kimmo Kartasalo, Constance Boissin, Sonja Koivukoski, Aino Kuusela, Dusan Rasic, Yanbo Feng, et al., The ACROBAT 2022 challenge: automatic registration of breast cancer tissue, arXiv preprint, arXiv:2305.18033, 2023.

- [9] Philippe Weitz, Masi Valkonen, Leslie Solorzano, Circe Carr, Kimmo Kartasalo, Constance Boissin, Sonja Koivukoski, Aino Kuusela, Dusan Rasic, Yanbo Feng, et al., ACROBAT - a multi-stain breast cancer histological whole-slide-image data set from routine diagnostics for computational pathology, arXiv preprint, arXiv:2211.13621, 2022.
- [10] Johannes Lotz, Nick Weiss, Jeroen van der Laak, Stefan Heldmann, Comparison of consecutive and restained sections for image registration in histopathology, *J. Med. Imag.* (2023).
- [11] Johannes Lotz, Nick Weiss, Jeroen van der Laak, Stefan Heldmann, HyReCo - hybrid re-stained and consecutive histological serial sections, *IEEE Data Port* (2021).
- [12] Yuanhua Lin, Zhendong Liang, Yonghong He, Wenting Huang, Tian Guan, End-to-end affine registration framework for histopathological images with weak annotations, *Comput. Methods Programs Biomed.* 241 (2023) 107763.
- [13] Ruqayya Awan, Shan E. Ahmed Raza, Johannes Lotz, Nick Weiss, Nasir Rajpoot, Deep feature based cross-slide registration, *Comput. Med. Imaging Graph.* 104 (2023) 102162.
- [14] Aristeidis Sotiras, Christos Davatzikos, Nikos Paragios, Deformable medical image registration: a survey, *IEEE Trans. Med. Imaging* 32 (7) (2013) 1153–1190.
- [15] Alessa Hering, Lasse Hansen, Tony C.W. Mok, Albert C.S. Chung, Hanna Siebert, Stephanie Häger, Annkristin Lange, Sven Kuckertz, Stefan Heldmann, Wei Shao, et al., Learn2Reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning, *IEEE Trans. Med. Imaging* 42 (3) (2022) 697–712.
- [16] Tony C.W. Mok, Albert C.S. Chung, Large deformation diffeomorphic image registration with Laplacian pyramid networks, in: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23*, Springer, 2020, pp. 211–221.
- [17] Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John Guttag, Adrian V. Dalca, Voxelmorph: a learning framework for deformable medical image registration, *IEEE Trans. Med. Imaging* 38 (8) (2019) 1788–1800.
- [18] Stefan Klein, Marius Staring, Keelin Murphy, Max A. Viergever, Josien P.W. Pluim, Elastix: a toolbox for intensity-based medical image registration, *IEEE Trans. Med. Imaging* 29 (1) (2009) 196–205.
- [19] Grant Haskins, Uwe Kruger, Pingkun Yan, Deep learning in medical image registration: a survey, *Mach. Vis. Appl.* 31 (1) (2020) 8.
- [20] Dwarikanath Mahapatra, Bhavna Antony, Suman Sedai, Rahil Garnavi, Deformable medical image registration using generative adversarial networks, in: *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, IEEE, 2018, pp. 1449–1453.
- [21] Kasper Marstal, Floris Berendsen, Marius Staring, Stefan Klein, Simpleelastix: a user-friendly, multi-lingual library for medical image registration, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2016*, pp. 134–142.
- [22] Ben Glocker, Aristeidis Sotiras, Nikos Komodakis, Nikos Paragios, Deformable medical image registration: setting the state of the art with discrete methods, *Annu. Rev. Biomed. Eng.* 13 (2011) 219–244.
- [23] Mattias P. Heinrich, Mark Jenkinson, Manav Bhushan, Tahreema Matin, Fergus V. Gleeson, Michael Brady, Julia A. Schnabel, Mind: modality independent neighbourhood descriptor for multi-modal deformable registration, *Med. Image Anal.* 16 (7) (2012) 1423–1435.
- [24] Frederik Maes, Dirk Vandermeulen, Paul Suetens, Medical image registration using mutual information, *Proc. IEEE* 91 (10) (2003) 1699–1722.
- [25] Jiří Borovec, Arrate Munoz-Barrutia, Jan Kybic, Benchmarking of image registration methods for differently stained histological slides, in: *2018 25th IEEE International Conference on Image Processing (ICIP)*, IEEE, 2018, pp. 3368–3372.
- [26] Johannes Lotz, Nick Weiss, Stefan Heldmann, Robust, fast and accurate: a 3-step method for automatic histological image registration, arXiv preprint, arXiv:1903.12063, 2019.
- [27] Chandler D. Gatenbee, Ann-Marie Baker, Sandhya Prabhakaran, Otilie Swinyard, Robert J.C. Slebos, Gunjan Mandal, Eoghan Mulholland, Noemi Andor, Andriy Marusyk, Simon Leedham, et al., Virtual alignment of pathology image series for multi-gigapixel whole slide images, *Nat. Commun.* 14 (1) (2023) 4502.
- [28] Ludovic Venet, Sarthak Pati, Paul Yushkevich, Spyridon Bakas, Accurate and robust alignment of variable-stained histologic images using a general-purpose greedy diffeomorphic registration tool, arXiv preprint, arXiv:1904.11929, 2019.
- [29] Ludovic Venet, Sarthak Pati, Michael D. Feldman, MacLean P. Nasrallah, Paul Yushkevich, Spyridon Bakas, Accurate and robust alignment of differently stained histologic images based on greedy diffeomorphic registration, *Appl. Sci.* 11 (4) (2021) 1892.
- [30] Marek Wodzinski, Andrzej Skalski, Multistep, automatic and nonrigid image registration method for histology samples acquired using multiple stains, *Phys. Med. Biol.* 66 (2) (2021) 025006.
- [31] Yi Song, Darren Treanor, Andrew J. Bulpitt, Nagitha Wijayathunga, Nick Roberts, Ruth Wilcox, Derek R. Magee, Unsupervised content classification based nonrigid registration of differently stained histology images, *IEEE Trans. Biomed. Eng.* 61 (1) (2013) 96–108.
- [32] Johannes Lotz, Janine Olesch, Benedikt Müller, Thomas Polzin, P. Galuschka, J.M. Lotz, Stefan Heldmann, Hendrik Laue, Margarita González-Vallinas, Arne Warth, et al., Patch-based nonlinear image registration for gigapixel whole slide images, *IEEE Trans. Biomed. Eng.* 63 (9) (2015) 1812–1819.
- [33] Christian Marzahl, Frauke Wilm, Lars Tharun, Sven Perner, Christof A. Bertram, Christine Kröger, Jörn Voigt, Robert Klopffleisch, Andreas Maier, Marc Aubreville, et al., Robust quad-tree based registration on whole slide images, in: *MICCAI Workshop on Computational Pathology*, PMLR, 2021, pp. 181–190.
- [34] Shengyu Zhao, Tingfung Lau, Ji Luo, I. Eric, Chao Chang, Yan Xu, Unsupervised 3d end-to-end medical image registration with volume tweening network, *IEEE J. Biomed. Health Inform.* 24 (5) (2019) 1394–1404.
- [35] Marek Wodzinski, Henning Müller, DeepHistReg: unsupervised deep learning registration framework for differently stained histology samples, *Comput. Methods Programs Biomed.* 198 (2021) 105799.
- [36] Marek Wodzinski, Henning Müller, Unsupervised learning-based nonrigid registration of high resolution histology images, in: *Machine Learning in Medical Imaging: 11th International Workshop, MLMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 11*, Springer, 2020, pp. 484–493.
- [37] Lin Ge, Xingyue Wei, Yayu Hao, Jianwen Luo, Yan Xu, Unsupervised histological image registration using structural feature guided convolutional neural network, *IEEE Trans. Med. Imaging* 41 (9) (2022) 2414–2431.
- [38] V.A. Pyatov, D.V. Sorokin, Affine registration of histological images using transformer-based feature matching, *Pattern Recognit. Image Anal.* 32 (3) (2022) 626–630.
- [39] Vladislav A. Pyatov, Dmitry V. Sorokin Tahir, Transformer-based affine histological image registration, in: *International Conference on Pattern Recognition*, Springer, 2022, pp. 541–552.
- [40] Marek Wodzinski, Henning Müller, Learning-based affine registration of histological images, in: *Biomedical Image Registration: 9th International Workshop, WBIR 2020, Portorož, Slovenia, December 1–2, 2020, Proceedings 9*, Springer, 2020, pp. 12–22.
- [41] PyVips, PyVips library, <https://www.libvips.org/>. (Accessed 14 December 2023), 2023.
- [42] Daniel DeTone, Tomasz Malisiewicz, Andrew Rabinovich Superpoint, Self-supervised interest point detection and description, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2018*, pp. 224–236.
- [43] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, Andrew Rabinovich, SuperGlue: learning feature matching with graph neural networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020*, pp. 4938–4947.
- [44] ANHIR Evaluation Platform, <https://anhir.grand-challenge.org/evaluation/challenge/leaderboard/>. (Accessed 14 December 2023).
- [45] R. Fernandez-Gonzalez, et al., System for combined three-dimensional morphological and molecular analysis of thick tissue specimens, *Microsc. Res. Tech.* 59 (6) (2002) 522–530.
- [46] L. Gupta, et al., Stain independent segmentation of whole slide images: a case study in renal histology, in: *Proceedings - International Symposium on Biomedical Imaging, 2018*, pp. 1360–1364.
- [47] I. Mikhailov, N. Danilova, P. Malkov, The immune microenvironment of various histological types of ebv-associated gastric cancer, *Virchows Arch.* (2018).
- [48] G. Bueno, O. Deniz, AIDPATH: academia and industry collaboration for digital pathology, <http://aidpath.eu>.
- [49] ACROBAT Evaluation Platform, <https://acrobat.grand-challenge.org/evaluation/model-development/leaderboard/>. (Accessed 14 December 2023).
- [50] Péter Bándi, Maschenka Balkenhol, Bram van Ginneken, Jeroen van der Laak, Geert Litjens, Resolution-agnostic tissue segmentation in whole-slide histopathology images with convolutional neural networks, *PeerJ* 7 (2019) e8242.
- [51] Artur Jurgas, Marek Wodzinski, Manfredo Atzori, Henning Müller, Robust multiresolution and multistain background segmentation in whole slide images, in: *Polish Conference on Biocybernetics and Biomedical Engineering*, Springer, 2023, pp. 29–40.
- [52] BigPicture project, <https://bigpicture.eu/>. (Accessed 14 December 2023).
- [53] Tony C.W. Mok, Zi Li, Yingda Xia, Jiawen Yao, Ling Zhang, Jingren Zhou, Le Lu, Deformable medical image registration under distribution shifts with neural instance optimization, in: *International Workshop on Machine Learning in Medical Imaging*, Springer, 2023, pp. 126–136.