Utilizing Radiomic Features for Automated MRI Keypoint Detection: Enhancing Graph Applications

Sahar Almahfouz Nasser[a], Shashwat Pathak[b], Keshav Singhal[c], Mohit Meena[d], Nihar Gupte[e], Ananya Chinmaya[f], Prateek Garg[g] and Amit Sethi[h]

Department of Electrical Engineering, Indian Institute of Technology Bombay, Mumbai, India

Keywords: Image Matching, Image Registration, Kepoint Detection, Radiomic Features, Brain MRI, GNN.

Abstract: Graph neural networks (GNNs) present a promising alternative to CNNs and transformers for certain image processing applications due to their parameter-efficiency in modeling spatial relationships. Currently, an active area of research is to convert image data into graph data as input for GNN-based models. A natural choice for graph vertices, for instance, are keypoints in images. SuperRetina is a promising semi-supervised technique for detecting keypoints in retinal images. However, its limitations lie in the dependency on a small initial set of ground truth keypoints, which is progressively expanded to detect more keypoints. We encountered difficulties in detecting a consistent set of initial keypoints in brain images using traditional keypoint detection techniques, such as SIFT and LoFTR. Therefore, we propose a new approach for detecting the initial keypoints for SuperRetina, which is based on radiomic features. We demonstrate the anatomical significance of the detected keypoints by showcasing their efficacy in improving image registration guided by these keypoints. We also employed these keypoints as ground truth for a modified keypoint detection method known as LK-SuperRetina for improved image matching in terms of both the number of matches and their confidence scores.

1 INTRODUCTION

Graph neural networks (GNNs) have shown promising results for reducing the computational requirements for certain image processing tasks as shown in (He et al., 2023). However, converting images into graphs is an active area of research. One approach is to break down an image into patches and treat each patch as a node in a graph while another is to segment an image into super-pixels and use each super-pixel as a graph vertex (He et al., 2023). Our method relies on detecting important keypoints in the images along with their features and making graphs.

Detecting important keypoints in certain types of images, such as magnetic resonance images of the brain is not straightforward due to the lack of well-defined landmarks. In various registration competitions, such as the BraTSReg challenge (Baheti et al., 2021), experts have had to mark specific landmark points on 3D brain images to help the participants align images and evaluate how well their registration algorithms worked. But this marking is neither easy nor fast. Also, only a few points – between 6 and 50 per set of images – get marked, which can be inadequate for good registration.

Traditional keypoint detection algorithms such as SIFT (Lindeberg, 2012) fall behind deep learning-based keypoint detection algorithms which have different types include: supervised, unsupervised, and semi-supervised methods. Some examples are UnsuperPoint (Christiansen et al., 2019), SuperPoint (DeTone et al., 2018), GLAMpoints (Truong et al., 2019), and SuperRetina (Liu et al., 2022).

After trying and failing to find reliable keypoints in brain images using methods such as SIFT (Lindeberg, 2012) and LoFTR (Sun et al., 2021), we proposed a new algorithm. Our proposed method finds keypoints and extract their descriptors using radiomic features in brain MR images. We assessed the utility of these keypoints in brain images by using them for image registration. Then we used a dataset that

[a] https://orcid.org/0000-0002-5063-9211
[b] https://orcid.org/0000-0002-4508-8531
[c] https://orcid.org/0009-0003-2290-4924
[d] https://orcid.org/0009-0008-1641-2727
[e] https://orcid.org/0009-0004-5972-2088
[f] https://orcid.org/0009-0008-5844-5498
[g] https://orcid.org/0009-0005-5348-8905
[h] https://orcid.org/0000-0002-8634-1804

Nasser, S., Pathak, S., Singhal, K., Meena, M., Gupte, N., Chinmaya, A., Garg, P. and Sethi, A.
Utilizing Radiomic Features for Automated MRI Keypoint Detection: Enhancing Graph Applications.
DOI: 10.5220/0012568800003657

In *Proceedings of the 17th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2024) - Volume 1*, pages 319-325
ISBN: 978-989-758-688-0; ISSN: 2184-4305

we prepared, which consists of MRI images from OASIS dataset (Marcus et al., 2007) and the keypoints we detected, to train the LK-SuperRetina algorithm (Almahfouz Nasser et al., 2023) to detect new keypoints. We showed that using a GNN, such as SuperGlue (Sarlin et al., 2020), for image matching improves the features of detected keypoints. This method of detecting keypoints opens doors for using graph-based neural networks for different tasks, including brain classification, segmentation, and registration.

Our contribution can be outlined as follows:

- Introduction of an innovative, fully automated approach for keypoint detection in brain MRI images.
- Development of a two-stage method: the first stage involves detecting initial keypoints by utilizing segmentations of regions of interest in the brain, employing radiomic features as keypoint descriptors, and establishing keypoints at the centers of mass of the segmentations. The second stage employs our previously developed algorithm, LK-SuperRetina, gradually increasing the number of detected keypoints, commencing from the initial set.
- Enhancement of various tasks, including image registration and image matching, achieved through the incorporation of detected keypoints.

2 RELATED WORK

Traditional approaches to keypoint detection have long been prominent in the field of computer vision, with a focus on identifying keypoints in images that remain consistent across scaling, rotation, and lighting variations. These techniques involve characterizing the local image patch around these keypoints using a set of features, allowing for the matching of keypoints between different images and object recognition. However, these methods come with their limitations, including high computational complexity, reduced accuracy in the presence of extreme lighting and viewpoint changes, and challenges in handling occlusions and cluttered backgrounds.

In recent years, deep learning-based keypoint detection algorithms have emerged as a promising alternative. These algorithms have the ability to autonomously learn robust and discriminative features directly from data, making them better suited to handle complex and diverse image variations. As a result, they have found application in various domains, including object detection, semantic segmentation, and image retrieval.

Within the realm of deep learning, keypoint detection algorithms come in various forms, including supervised, semi-supervised, self-supervised, and unsupervised techniques. Supervised techniques require annotated data, where keypoints are manually labeled in training images, proving advantageous in scenarios with a substantial volume of labeled data, such as facial recognition or object detection. Conversely, unsupervised techniques operate independently of labeled data, with the network learning to identify keypoints by maximizing specific objectives, like information preservation during feature extraction. These methods are particularly valuable in contexts where obtaining labeled data is challenging or expensive, such as in medical imaging or remote sensing.

Prominent deep learning-based keypoint detection methods include UnsuperPoint, SuperPoint, GLAMpoints, and SuperRetina. UnsuperPoint (Christiansen et al., 2019) introduces an innovative unsupervised training approach, utilizing a combination of differentiable soft nearest neighbor loss and unsupervised clustering loss. SuperPoint (DeTone et al., 2018), on the other hand, represents a self-supervised deep learning-based algorithm for keypoint detection and description, relying on a novel loss function for training on unlabeled images, which makes it adaptable and scalable across various applications.

The primary distinction between UnsuperPoint and SuperPoint lies in their training methodologies. While SuperPoint follows a self-supervised approach, UnsuperPoint takes an unsupervised path. Additionally, UnsuperPoint achieves state-of-the-art performance across various benchmarks, even outperforming SuperPoint in challenging scenarios characterized by significant viewpoint alterations and illumination shifts. GLAMpoints (Truong et al., 2019) emerges as a semi-supervised deep learning-based algorithm for interest point detection and description, employing a unique greedy training strategy for end-to-end learning of keypoint detection and description. This strategy involves learning to select the most precise keypoints and their descriptors, resulting in higher accuracy and efficiency, especially in demanding scenarios involving significant viewpoint changes, scaling, rotation, and benchmarks like HPatches (Balntas et al., 2017). Furthermore, GLAMpoints is well-suited for accommodating multiple object instances within the same image, making it suitable for multi-object tracking and matching. SuperRetina (Liu et al., 2022) represents a semi-supervised approach for keypoint detection and description in retinal images, leveraging both labeled and unlabeled data to enhance the performance of the keypoint detector and descriptor.

In the following section, we introduce our proposed approach for detecting keypoints in brain MRI images. As we will demonstrate, this task proves to be highly challenging, and traditional keypoint detection methods like SIFT, as well as well-known deep learning techniques such as LoFTR, struggle to identify robust and repeatable keypoints. Our solution involves leveraging radiomic features for initial keypoint detection in the brain. Subsequently, we employ LK-SuperRetina to increase the number of identified keypoints in the images. Finally, we showcase the effectiveness of these keypoints in applications like image matching and registration, highlighting their positive impact on the overall performance of these tasks.

3 PROPOSED METHOD

We now introduce our approach for keypoint detection using radiomic features. Given an image alongside its segmentation labels, our method identifies radiomic keypoints as the centers of radiomic segmentation labels within the image. These radiomic features encompass a range of intensity characteristics (such as mean intensity, standard deviation, skewness and kurtosis), shape characteristics (such as volume, surface area, and compactness), texture characteristics (such as gray level co-occurrence matrix (GLCM), gray level run length matrix (GLRLM), and gray level run length matrix)), wavelet-Based characteristics, spatial and statistical characteristics (such as centroid position, eccentricity, and entropy), and fractal characteristics. These radiomic features are specific to the regions defined by the segmentation map. To compute these features, we utilized the Pyradiomic library (Van Griethuysen et al., 2017).

Radiomic keypoints are closely tied to segmentation regions predicted by neural networks. We trained Swin UNetR (Hatamizadeh et al., 2021) to predict the segmentation maps of brain images and used them as masks to extract the keypoints. Each segmentation mask yields a keypoint location, accompanied by 53 descriptive radiomic features, see figures 1 and 2 for more details. Radiomic keypoints exhibit repeatability across various brain samples even in the presence of varying intensity and non-rigid deformations which make our proposed method a more robust alternative to methods, such as SIFT (Lindeberg, 2012). These detected keypoints can serve as a initial keypoints for detecting additional keypoints through deep learning techniques as we will show in our results, in which we trained LK-SuperRetina (Almahfouz Nasser et al., 2023) to automatically detect these ground truth keypoints and extra keypoints. This underscores the significance of these keypoints as essential landmarks within brain images.

4 EXPERIMENTS AND RESULTS

This section begins with an introduction to the dataset used in all our experiments. Following that, we explore various approaches applied for keypoint detection, explaining how the detected keypoints were utilized to improve registration and identify additional keypoints. Finally, we will conclude with our results on GNN-based image matching.

4.1 OASIS Dataset

The OASIS dataset, as described in (Marcus et al., 2007), contains MRI data obtained from 414 subjects. This dataset has been divided into three separate subsets for training, validation, and testing, following a ratio of 314:50:50, respectively. Each subject in the dataset has T1-weighted scan, as well as segmentation masks of various regions of the brain. This dataset incorporates three distinct types of brain segmentation: a four-label mask, a thirty-five label mask, and a twenty-four label mask.The 3D T1-weighted scans and their corresponding masks are of resolution (160, 192, 224). Additionally, the 2D T1-weighted scans and their corresponding masks have a resolution of (160, 192).

4.2 SIFT

SIFT (Lindeberg, 2012) was initially chosen for experimentation due to its reputation in keypoint detection. Utilizing gradients, SIFT excels in identifying scale and rotation invariant keypoints. The method calculates orientation gradients across scales, creating a distinctive 128-dimensional vector for each keypoint.

Figure 3 shows significant issues we encountered with keypoint detection using SIFT. Notably, there was inconsistency in keypoint locations across different MRI slices within the brain, reflecting a lack of repeatability. Our experiments highlighted SIFT's ineffectiveness under conditions involving large deformations.

4.3 LoFTR

Due to SIFT's inability to establish correspondences, we opted to explore a deep learning method namely LoFTR (Sun et al., 2021), which stands for detector-free local feature matching transformer.

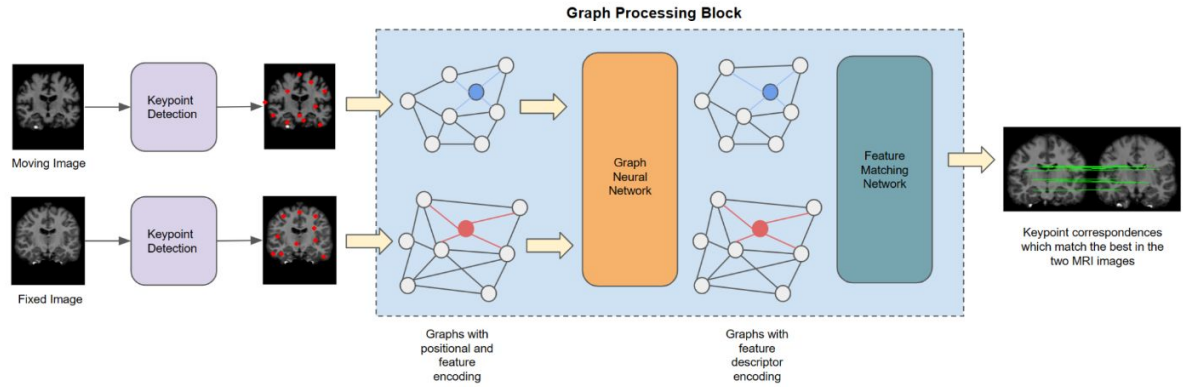

Figure 1: The image matching pipeline, encompassing keypoint detection using neural networks, graph formation from detected keypoints, graph neural network (GNN) processing to enhance keypoint features, and a dedicated head for keypoint matching.

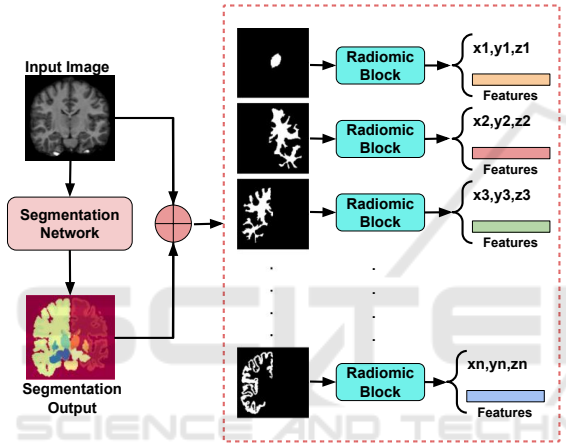

Figure 2: Radiomic features-based Keypoint detection method.

In the LoFTR approach, a convolutional neural network (CNN) with an encoder-decoder architecture extracts both low-resolution and high-level features. The LoFTR module, incorporating self-attention and cross-attention blocks, transforms these features, and a differential matching layer offers two methods: optimal transport and dual softmax (Bridle, 1989) . Our implementation used a pretrained model from the Aachen Day Night Dataset (Sattler et al., 2018). LoFTR succeeded in generating keypoint matches for similar intensity profiles but faced challenges in other cases similar to SIFT.These findings emphasize the need for keypoints exhibiting consistency, accurate matching, and intensity profile invariance across the dataset.

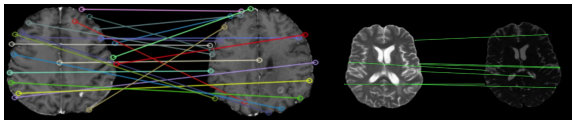

Figure 3: SIFT (Lindeberg, 2012) and LoFTR (Sun et al., 2021) performances in keypoint detection. The left pair of images depicts results for SIFT, while the right one illustrates the performance of LoFTR. It is evident that LoFTR outperforms SIFT in keypoint detection and matching. Nevertheless, LoFTR encounters challenges when the source and target images exhibit varying intensity distributions.

4.4 Radiomic Features-Based Keypoint Detection

Having faced challenges with keypoint detection using SIFT and LoFTR, which proved sensitive to intensity variations, we shifted to our proposed radiomic features-based method. In this section, we prove the importance of the detected keypoints by showing our results on registering brain images among different subjects. we also showcase the applications of the dataset containing the original OASIS scans and the corresponding detected radiomic keypoints on automatic keypoint detection. And finally train a GNN-based image matcher on the same dataset.

4.4.1 Image Registration

To assess the significance of radiomics keypoints as landmarks, we integrated them into the loss function of a registration network (TransMorph (Chen et al., 2022)). Our findings revealed that incorporating the keypoints' loss led to a notable 3% enhancement in the registration performance.

Vision transformers, excelling in capturing long-range spatial relationships, prove effective in medical image tasks due to their large receptive fields. TransMorph (Chen et al., 2022), a hybrid Transformer-

ConvNet model, utilizes these advantages for volumetric medical image registration. The encoder divides input volumes into 3D patches, projecting them to feature representations through linear layers. Sequential patch merging and Swin Transformer blocks follow. The decoder, with upsampling and convolutional layers, connects to the encoder stages via skip connections, producing the deformation field. We contributed by designing a customized loss function for keypoints, utilizing Gaussian-blurred keypoints to create a ground truth heatmap. Combining Dice and inverted Dice losses addressed imbalanced masks, resulting in a **3%** Dice score improvement over the OASIS test dataset. TransMorph achieved a dice score of **0.89** with keypoint loss, compared to 0.86 without, underscoring keypoints' role in enhancing registration performance. A potential avenue for future research involves developing a loss function that considers both the feature descriptors of keypoints and the disparity between the locations of registered keypoints and their counterparts in the target image.

4.4.2 Automated Keypoint Detection

SuperRetina, introduced in (Liu et al., 2022) is an apdaptive version of SuperPoint model (DeTone et al., 2018) for identifying important keypoints in retinal images. Utilizing a semi-supervised learning framework, SuperRetina maximizes the utility of limited labeled retinal image data by combining both supervised and unsupervised techniques. Yet, its utilization requires an initial set of ground truth keypoints to initiate the process, subsequently increasing the detected keypoints iteratively. In our approach, we use our radiomic keypoints as the initial sets for OASIS images.

LK-SuperRetina (Almahfouz Nasser et al., 2023) which is a modified version of SuperRetina consists of an encoder for downsampling, along with two decoders—one for keypoint detection and another for descriptor generation. Keypoint detection utilizes a mix of labeled and unlabeled data, while descriptor training employs self-supervised learning.

Following the U-Net (Ronneberger et al., 2015) design, LK-SuperRetina's shallow encoder begins with a single convolutional layer, followed by three blocks containing two convolutional layers, a 2×2 max-pooling layer, and ReLU activation. The keypoint decoder has three blocks with two convolutional layers, ReLU activation, and concatenation block. The detection map (P) is generated through a convolutional block with three convolutional layers and a sigmoid activation.

The loss function combines the detector and the descriptor losses. The detection loss consists of the classification loss and the geometric loss as shown in Equation 1.

$$l_{det} = l_{clf} + l_{geo} \tag{1}$$

The classification loss component (l_{clf}) is defined in Equation 2, where $\tilde{Y}$ represents the smoothed version of the binary ground truth labels Y of the keypoints after blurring them with a 2D Gaussian.

$$l_{clf}(I;Y) = 1 - \frac{2.\sum_{i,j}(P \circ \tilde{Y})_{i,j}}{\sum_{i,j}(P \circ P)_{i,j} + \sum_{i,j}(\tilde{Y} \circ \tilde{Y})_{i,j}} \tag{2}$$

Where $\circ$ denotes element-wise multiplication.

The detector generates a heatmap as its output. Coordinates where the intensity value exceeds a specified threshold (subject to non-maximum suppression) are regarded as the keypoint coordinates. When feeding both the image I and its augmented version I' to the network, two tensors for the descriptors D and D' are obtained. For each keypoint (i, j) in the non-maximum suppressed keypoint set $\tilde{P}$, two distances are computed: $\Phi^{rand}_{i,j}$ between the descriptors of (i, j) in the set $\tilde{P}$ and a random point from the registered heatmap $H(\tilde{P})$, and $\Phi^{hard}_{i,j}$ representing the minimal distance, as depicted in Equation 3.

$$l_{des}(I,H) = \sum_{(i,j)\in\tilde{P}} max\left(0, m + \Phi_{i,j} - \frac{1}{2}\left(\Phi^{rand}_{i,j} + \Phi^{hard}_{i,j}\right)\right) \tag{3}$$

For more in-depth information on the loss function, please refer to the SuperRetina paper (Liu et al., 2022). Figure 4 shows the results obtained from LK-SuperRetina. As demonstrated, the number of additionally identified keypoints meets expectations, showcasing the network's proficiency in capturing good new keypoints. The network successfully detects both the ground truth keypoints and extra keypoints during the testing phase.

Figure 4 presents two instances demonstrating the resilience of the keypoint detection model against deformations. The images were deformed randomly by an affine deformation. We passed the original image (target) and the deformed image (reference) seperately to LK-SuperRetina. The model successfully identified corresponding keypoints in both images, as indicated by the number of good matches. We adjusted the thresholds of LK-SuperRetina to detect a smaller set of keypoints for the clarity of the visualization, but in practice, the model can detect over 300 good keypoints.

4.4.3 Image Matching

Following the detection of keypoints within the brain images, we proceed to construct graphs to be used as

Table 1: A comparison between brute force matcher and SuperGlue. SuperGlue outperforms the brute force matcher in terms of both evaluation metrics: the average number of good matches and the average confidence score across the entire test dataset.

Method	Avg. No. Good Matches	Confidence Score
BF (Lowe, 2004)	7	0.449 ± 0.007
SuperGlue	**20**	$\mathbf{0.988 \pm 0.010}$

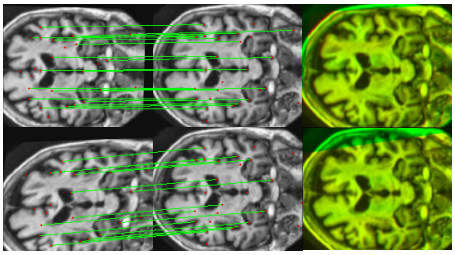

Figure 4: Two examples show the robustness of the keypoint detection model. In both rows, the sequence of images, from left to right, includes: the reference image, the target image, and the registration output. The good matches contribute to aligning the images effectively.

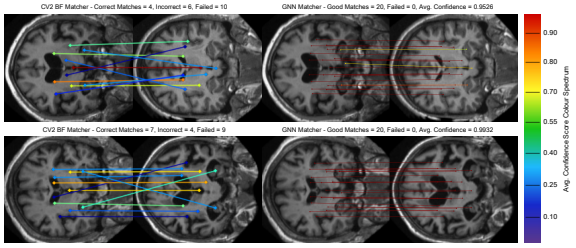

Figure 5: A comparison of the SuperGlue and brute force matcher performance in matching detected keypoints on brain images.

inputs of the GNN, for accomplishing specific tasks such as matching in our study. Within this paper, we demonstrate our success in training a GNN-based matcher (SuperGlue (Sarlin et al., 2020)) using the graphs formed from the detected keypoints.

SuperGlue designed for matching two sets of local features by identifying correspondences and filtering non-matchable points. Using attention-based graph neural networks, it integrates context aggregation, matching, and filtering within a unified architecture. SuperGlue employs self-attention to enhance the receptive field of local descriptors and cross-attention for cross-image communication. The network handles partial assignments and occluded points by solving an optimal transport problem. With superior performance over other learned approaches, SuperGlue achieves state-of-the-art results in pose estimation for challenging real-world indoor and outdoor environments.

Figure 5 and Table 1 show a performance comparison between the brute force matcher and SuperGlue across the test dataset. SuperGlue enhances the features of the detected keypoints which improves the matching performance. The data presented in the table suggests that exploring the application of GNNs for keypoint matching is a promising endeavor. To pursue this, it is essential to identify keypoints in images and format the data to serve as appropriate input for GNNs. Having introduced a technique for detecting significant keypoints in brain MRI images, these keypoints can now function as nodes in graphs, providing data for GNNs. This development paves the way for a multitude of graph-based applications in the analysis of brain MRI images.

5 CONCLUSION

To sum up, our radiomic keypoint detection algorithm provides a solution for automated keypoint detection in MRI scans, overcoming challenges encountered by traditional and other deep learning methods. The limited set of radiomic keypoints facilitates training SuperRetina for increased keypoint detection. These keypoints are consistent and deformation resilient.

Our approach paves the way for the application of GNN-based models on brain images, offering a faster and more parameter-efficient alternative compared to CNNs and transformers. Moreover, the detection of keypoints contributes to various tasks, including registration as justified in this work.

The limitation of this study is the necessity for segmenting regions of interest to acquire radiomic features, which are subsequently utilized as the initial keypoints for the keypoint detection algorithm, LK-SuperRetina.

REFERENCES

Almahfouz Nasser, S., Gupte, N., and Sethi, A. (2023). Reverse knowledge distillation: Training a large model using a small one for retinal image matching on limited data. *arXiv e-prints*, pages arXiv–2307.

Baheti, B., Waldmannstetter, D., Chakrabarty, S., Akbari, H., Bilello, M., Wiestler, B., Schwarting, J., Calabrese, E., Rudie, J., Abidi, S., et al. (2021). The brain tumor sequence registration challenge: establishing correspondence between pre-operative and follow-up mri scans of diffuse glioma patients. *arXiv preprint arXiv:2112.06979*.

Balntas, V., Lenc, K., Vedaldi, A., and Mikolajczyk, K.

(2017). Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5173–5182.

Bridle, J. (1989). Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. *Advances in neural information processing systems*, 2.

Chen, J., Frey, E. C., He, Y., Segars, W. P., Li, Y., and Du, Y. (2022). Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis*, 82:102615.

Christiansen, P. H., Kragh, M. F., Brodskiy, Y., and Karstoft, H. (2019). Unsuperpoint: End-to-end unsupervised interest point detector and descriptor. *arXiv preprint arXiv:1907.04011*.

DeTone, D., Malisiewicz, T., and Rabinovich, A. (2018). Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236.

Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H. R., and Xu, D. (2021). Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In *International MICCAI Brainlesion Workshop*, pages 272–284. Springer.

He, X., Hooi, B., Laurent, T., Perold, A., LeCun, Y., and Bresson, X. (2023). A generalization of vit/mlp-mixer to graphs. In *International Conference on Machine Learning*, pages 12724–12745. PMLR.

Lindeberg, T. (2012). Scale invariant feature transform.

Liu, J., Li, X., Wei, Q., Xu, J., and Ding, D. (2022). Semi-supervised keypoint detector and descriptor for retinal image matching. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*, pages 593–609. Springer.

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110.

Marcus, D. S., Wang, T. H., Parker, J., Csernansky, J. G., Morris, J. C., and Buckner, R. L. (2007). Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience*, 19(9):1498–1507.

Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer.

Sarlin, P.-E., DeTone, D., Malisiewicz, T., and Rabinovich, A. (2020). Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947.

Sattler, T., Maddern, W., Toft, C., Torii, A., Hammarstrand, L., Stenborg, E., Safari, D., Okutomi, M., Pollefeys, M., Sivic, J., et al. (2018). Benchmarking 6dof outdoor visual localization in changing conditions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8601–8610.

Sun, J., Shen, Z., Wang, Y., Bao, H., and Zhou, X. (2021). Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931.

Truong, P., Apostolopoulos, S., Mosinska, A., Stucky, S., Ciller, C., and Zanet, S. D. (2019). Glampoints: Greedily learned accurate match points. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10732–10741.

Van Griethuysen, J. J., Fedorov, A., Parmar, C., Hosny, A., Aucoin, N., Narayan, V., Beets-Tan, R. G., Fillion-Robin, J.-C., Pieper, S., and Aerts, H. J. (2017). Computational radiomics system to decode the radiographic phenotype. *Cancer research*, 77(21):e104–e107.