



ConKeD: multiview contrastive descriptor learning for keypoint-based retinal image registration

David Rivas-Villar^{1,2} · Álvaro S. Hervella^{1,2} · José Rouco^{1,2} · Jorge Novo^{1,2}

Received: 15 February 2024 / Accepted: 25 June 2024
© The Author(s) 2024

Abstract

Retinal image registration is of utmost importance due to its wide applications in medical practice. In this context, we propose ConKeD, a novel deep learning approach to learn descriptors for retinal image registration. In contrast to current registration methods, our approach employs a novel multi-positive multi-negative contrastive learning strategy that enables the utilization of additional information from the available training samples. This makes it possible to learn high-quality descriptors from limited training data. To train and evaluate ConKeD, we combine these descriptors with domain-specific keypoints, particularly blood vessel bifurcations and crossovers, that are detected using a deep neural network. Our experimental results demonstrate the benefits of the novel multi-positive multi-negative strategy, as it outperforms the widely used triplet loss technique (single-positive and single-negative) as well as the single-positive multi-negative alternative. Additionally, the combination of ConKeD with the domain-specific keypoints produces comparable results to the state-of-the-art methods for retinal image registration, while offering important advantages such as avoiding pre-processing, utilizing fewer training samples, and requiring fewer detected keypoints, among others. Therefore, ConKeD shows a promising potential towards facilitating the development and application of deep learning-based methods for retinal image registration.

Keywords Self-supervised learning · Feature-based registration · Image registration · Retinal image registration · Medical imaging

1 Introduction

Image registration is the process of spatially aligning a pair of images. In this process, one image is used as the reference, the fixed image, while the second image, the moving image, is spatially transformed to match the first one. These images are

usually from the same subject, and they differ in their point of view or in the instant they were captured in. In medicine, image registration is particularly important due to its multiple applications on clinical practice [1]. Image registration enables the simultaneous analysis of multiple images which, in turn, allows clinicians to draw more informed conclusions [2]. Additionally, medical image registration enables longitudinal studies which are useful to monitor disease progression or remission [3].

Retinal image registration (RIR) is a highly relevant task as the eyes are the only organs in the human body that allow non-invasive in vivo observation of the blood vessels and neuronal tissue [4]. In particular, there is a high interest in the registration of color fundus (CF) images due to the widespread use of this image modality in clinical practice [5]. CF images are cost-effective due to the relatively low price of CF cameras combined with their efficacy in the diagnosis of numerous diseases [6, 7]. However, CF images also present particular characteristics that complicate the registration. For instance, their photographic nature can cause several imperfections

✉ David Rivas-Villar
david.rivas.villar@udc.es

Álvaro S. Hervella
a.suarez@udc.es

José Rouco
jrouco@udc.es

Jorge Novo
jnov@udc.es

¹ Grupo VARPA, Instituto de Investigación Biomédica de A Coruña (INIBIC), Universidade da Coruña, A Coruña 15006, A Coruña, Spain

² Departamento de Ciencias de la Computación y Tecnologías de la Información, Universidade da Coruña, A Coruña 15071, A Coruña, Spain

due to incorrect placement of the device, incorrect capture by the clinician, or movements by the patient. Some of these imperfections are color or illuminance variation, changes in focus, halos, etc. Moreover, certain diseases can drastically alter the retinal appearance. These morphological changes complicate the registration process in the presence of pathology. The combination of all of these characteristics makes the CF registration process a challenging task.

In broad terms, automatic registration approaches can be classified according to their algorithmic basis as either classical approaches or deep learning approaches. Furthermore, registration methods can also be divided into three main groups: feature-based registration (FBR), which uses keypoints and descriptors; intensity-based registration (IBR), which uses the intensity values of the image; and direct parameter regression (DPR), which predicts a deformation field or matrix directly from the input images [8, 9].

Classical registration methods are still widely used nowadays. However, novel deep learning approaches have several advantages that usually make them more desirable over the classical methods. For instance, deep learning methods can be trained end-to-end from raw data to the expected result. This eliminates the need for ad hoc feature engineering which allows deep learning methods to be more flexible and adaptable to changes in the input data.

FBR methods use keypoints to drive the registration process. These keypoints are distinctive spatial locations that can be detected in the different images to register. The keypoints are matched among the images and used to compute the geometric transformation that aligns them. FBR methods are inherently explainable as the keypoints used to calculate the transformation are known and can be visualized. Commonly, to aid in the matching process, point descriptors are also computed. These descriptors are unique feature representations that characterize each keypoint and facilitate their distinction. FBR methods are divided into methods that use generic keypoints (valid for any domain) and domain-specific keypoints (generally more accurate but only usable in a single domain). Generic methods are potentially less accurate and more computationally expensive due to the lower detection specificity that makes necessary the detection and processing of a higher number of keypoints.

IBR approaches work by iteratively maximizing a similarity metric between the intensity values of the images to be registered across transformation parameter space. There are multiple suitable similarity metrics [10, 11], including deep learning-based ones [12].

DPR methods allow direct transformation prediction or deformation field prediction which, in turn, allows image alignment [13, 14]. To achieve this, a deep neural network is trained such that its output is used to transform the moving

image. In this case, the training objective is to maximize the similarity between the transformed moving image and the fixed image or, equivalently, between two simplified representations of these images (such as, e.g., blood vessel segmentation maps for CF images [15]). The key difference between IBR and DPR is how they use similarity metrics. IBR explicitly maximizes a similarity metric while DPR uses the similarity metric to guide the learning of neural network. This network learns to create suitable transformations directly from the input data, only relying on the similarity metric as guidance during its training step.

Finally, there are hybrid methods joining several of these registration paradigms. In this case, FBR methods usually provide an initial alignment which is then refined by other IBR or DPR methods that are capable of deformable registration thus improving the results [16, 17]. However, while these approaches might be successful in terms of results, they are not desirable since they increase the complexity and the computational cost and may introduce extra hyperparameters.

Currently, most state-of-the-art medical registration methodologies are based on IBR or DPR methods [14, 18, 19]. However, RIR tasks, and particularly CF images, have certain characteristics that limit the applicability of these methods. In particular, the patterns that are useful for the registration of CF images (i.e., blood vessel bifurcations and crossovers, optic disc) are relatively small and are scattered over a homogeneous background. Furthermore, CF images can present large displacements and low degree of overlapping between the images in each pair. Additionally, there can be multiple intra-pair differences produced by the photographic nature of the capture process (i.e., light halos, blurriness due to suboptimal focus). These changes are specifically detrimental to IBR methods as creating a robust similarity metric is challenging under these conditions. Similarly, as the transformations on the retina are almost always rigid, deformable registration methods like DPR using deformation fields are not suitable as they pose the risk of overfitting these fields creating unrealistic transformations. Therefore, commonly used methods in other medical image modalities such as [14, 18] are unsuitable for color fundus registration.

Traditionally, the CF registration field has been dominated by classical FBR methods [20, 21]. Some of these methods, such as REMPE and VOTUS, still obtain the best results in most of the categories of the well-known FIRE benchmark dataset [20, 21]. REMPE [20] combines generic points (SIFT) and domain-specific keypoints (blood vessel bifurcations). Then, using RANSAC (Random Sample Consensus) [22] and Particle Swarm Optimization, the transformation is computed according to an ellipsoidal model of the eye that has been created specifically for this task. Differently,

VOTUS [21] creates graphs for the arterio-venous tree. Then, the graphs from the corresponding images are matched using a novel algorithm (Vessel Optimal Transform) that relies on classical image features. The transformation matrix is computed using DeSAC (Deterministic Sample And Consensus).

Recently, some novel FBR deep learning-based methods have been able to compete with and even outperform the classical approaches in some cases [8, 23, 24], whereas previous DPR deep learning methods could not [25]. Some of these methods are based on domain-specific keypoints [23, 24] while others employ generic keypoint detectors [8].

The work of Rivas-Villar et al. [23] was the first FBR deep learning method for CF image registration. It uses a CNN to detect blood vessel crossovers and bifurcations which are then used to infer a transformation using RANSAC without the need for descriptor computation. However, the lack of descriptors limits the complexity of the transformation as calculating all the possible combinations of keypoints with RANSAC by brute-force increases in computational cost with the degrees of freedom in the geometric transformation. SuperRetina [24] adapted SuperPoint [26], a natural image registration method, to RIR. They propose to train this network using keypoints labeled by an expert manually or by a classical approach, known as PBO [27]. To perform the matching of keypoints, they use descriptors that are learned using a triplet loss approach. Their method requires extensive ad hoc image pre-processing, firstly taking only the green channel from the RGB images, normalizing it, applying CLAHE contrast enhancement [28], and finally using gamma adjustment. These pre-processing steps simplify the task but require multiple parameters that may need to be tuned on a per-dataset or per-device basis. After this, the keypoints and descriptors are jointly trained starting from a set of labeled keypoints which are progressively expanded using the more confident detections of the network outside of the ground truth. In order to complete the registration pipeline, SuperRetina uses a double inference step. This means that the registration is firstly done normally and, after transforming the moving image, the process is repeated again to improve the results at the expense of more computation cost. Similarly, in Rivas-Villar et al. [8], R2D2 [29] is adapted for the retinal image registration task. R2D2 is a natural image registration method capable of jointly learning keypoints and descriptors without ground truth data, making this method fully unsupervised. Unlike SuperRetina, this method does not require pre-processing or double inference. However, it obtains worse results. In this method, the descriptors are learned using a multi-negative loss, known as AP loss.

Overall, using domain-specific keypoints, such as blood vessel crossovers and bifurcations, involves the use of supervised learning and manually labeled data [23, 24], whereas

methods which use generic keypoint detectors do not require labeled data [8]. However, due to the lack of precise and highly distinctive keypoints, generic detectors usually carry an increased computational cost. In terms of descriptor learning, deep learning methods use contrastive learning, a technique based on comparing samples. For a given keypoint, which is used as a reference (anchor keypoint), there are positive samples (the same point in another view of the same image) and negative samples (any other point in the same or another view). The main difference among methods is how they use the different samples and how many samples they are able to process on each contrastive learning step. Currently, in natural image registration as well as in RIR, most approaches use triplet loss [24, 26] which exploits a single positive and negative sample per anchor. Moreover, in RIR, there are some approaches that employ multi-negative single-positive methods [8] or do not use descriptors at all [23]. In this regard, none of the state-of-the-art methods use novel multi-positive multi-negative approaches for image registration, neither in natural images nor in medical images. Multiple positive contrastive learning has been successfully used in other tasks, such as image classification, where it has demonstrated improved performance over just multi-negative approaches [30]. These methods employ multiviewed batches to create multiple positives for each anchor, which in turn, allows to leverage more information from the input images. However, no multi-positive multi-negative framework has been developed for keypoint-based registration despite the success and benefits that this self-supervised approach demonstrated in other areas [31, 32].

In this work, we propose ConKeD, **C**ontrastive **K**eypoint **D**escriptors, a novel deep learning method to learn descriptors for domain-specific keypoint-based RIR. In contrast to previous methods, we use a multi-positive multi-negative contrastive learning approach that allows to leverage additional information from the input images, hence improving the learning of descriptors. Our multi-positive framework represents the first approach of this kind for image registration. In that regard, we propose a complete methodology for the registration of CF images, using vessel crossovers and bifurcations as keypoints together with the descriptors learned using ConKeD. These keypoints are highly specific and unique which helps reducing the number of required detections to perform the registration. The main contributions of the paper are the following:

- We propose the first multi-positive multi-negative contrastive learning approach for keypoint-based image registration.
- We propose a complete RIR deep learning methodology that takes advantage of domain-specific keypoints and our novel descriptors.

- We provide extensive experiments in the public FIRE dataset comparing ConKeD against established baselines and analyzing different factors of the method.
- Our experimentation reveals that the novel multi-positive multi-negative approach improves the performance of the commonly used triplet loss as well as multi-negative approaches.
- Our proposal achieves results comparable to those in the state-of-the-art while avoiding image pre-processing, requiring less training samples, and also requiring less detected keypoints.

2 Materials and methods

2.1 Methodology

The proposed methodology for CF registration is described in Fig. 1. This methodology is based on ConKeD, our novel framework for learning descriptors for keypoint-based image registration. These keypoints are detected using a state-of-the-art neural network [23]. Meanwhile, another network is used to compute the descriptors of these keypoints. The descriptor network is trained with a multiviewed batch containing one original image and N views or augmentations of the original image. Therefore, each keypoint (i.e., anchor) will have N positive samples while all the rest of detected

keypoints act as negatives, including the points in the own image. These samples coupled with a suitable loss allow the keypoint descriptor to be trained, learning to match the positive samples and differentiate the negatives. At inference time, the keypoints are detected, described, and matched, and, finally, RANSAC is used to infer a projective transform (homography) to align both images.

2.1.1 Keypoint detection

The first step of the proposed methodology is to detect blood vessel crossovers and bifurcations. These keypoints are highly specific and distinctive, allowing for accurate matching. However, the number of these keypoints in each image is unknown a priori and can even change between images of the same individual (i.e., in the same pair). This is due to the difference in viewpoint as well as pathology progression, which can occlude or banish some keypoints.

In order to detect these keypoints, we train a CNN to generate heatmaps depicting all the possible keypoints in the input image [33]. Using heatmaps instead of binary maps is beneficial due to the heavy imbalance between the positive class (the keypoint pixels) and the negative class (the background pixels) [33]. The produced heatmaps have maximum values at the location of the detected keypoints and progressively decreasing values in the neighboring pixels.

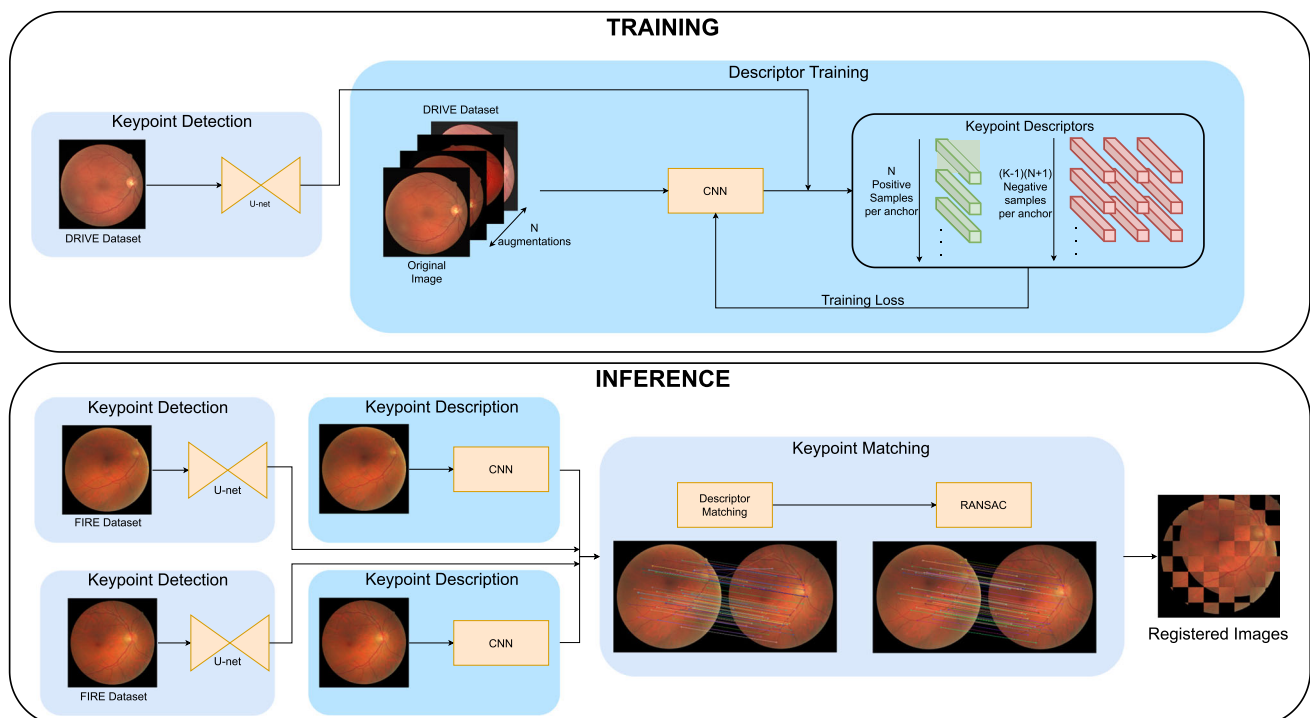


Fig. 1 Overview of the proposed method for both training and inference

The heatmaps increase the information available to the network during training, improving the learning process [33].

To distinguish between crossovers and bifurcations, two separate heatmaps are generated, one for each type of keypoint. Moreover, to incentivize the network to detect all the keypoints even if their type could not be inferred, a third output heatmap containing both types of keypoints is also generated. The network is trained using the mean squared error (MSE) between the predicted heatmaps and the target heatmaps as the loss function. In order to create the target heatmaps, the original binary ground truth is convolved using a Gaussian kernel.

Finally, during inference, in order to extract the keypoints, we compute the heatmaps using this network. From the heatmaps, we can obtain the discrete locations of the keypoints by using a local maxima filter and an intensity threshold [33].

2.1.2 Keypoint description

The second step is to generate the descriptors for the detected keypoints. For this step, we propose a novel framework, ConKeD, which is schematically described in Fig. 2.

ConKeD follows a novel multi-positive multi-negative contrastive learning approach in order to learn descriptors for image registration. Typically, deep learning-based FBR methods learn descriptors using triplet loss approaches. These approaches compare each given point against both a positive and a negative sample. Some recent methods also explored multi-negative approaches which, instead of selecting a particular negative sample, use of all of them. In contrast, our method leverages multiple samples for both

positive and negative examples. This avoids the need for explicit mining and increases the amount of contrastive comparisons, facilitating training.

The fundamental idea behind our proposal is to compare the representations (i.e., descriptors) of different keypoints. In particular, the key concept is to pull together the anchor (a particular point used as reference) and all the positive samples in the descriptor space while pushing apart the anchor from all the different negative samples.

In order to process multiple positive and multiple negative samples during training, we create a multiviewed batch for each input image during training. This batch includes an image from the dataset, serving as base, and N views or image augmentations. Thus, our multiviewed batch is of size $1 + N$. The image augmentations (views) are generated from the base image using spatial and color transformations. Therefore, all of the positive samples are augmentations of the anchor keypoint while negative samples are the rest of the detected keypoints, belonging to any image in the batch. In any case, it should be noted that ConKeD does not require multiple images from the same individual and, therefore, it could be trained using unlabeled data.

Additionally, in order to maximize the number of comparisons, our novel approach is trained with the whole images (as opposed to patches). For each image, the CNN produces a dense block of descriptors, each descriptor corresponding to one pixel of the input image. Therefore, for each keypoint location detected in the previous step, it is possible to directly select its specific descriptors. Considering that we have an average of K keypoints per image and a multiviewed batch size $1 + N$, there are a total of $K(1 + N)$ keypoint samples per batch. In particular, each one of the K keypoints will have

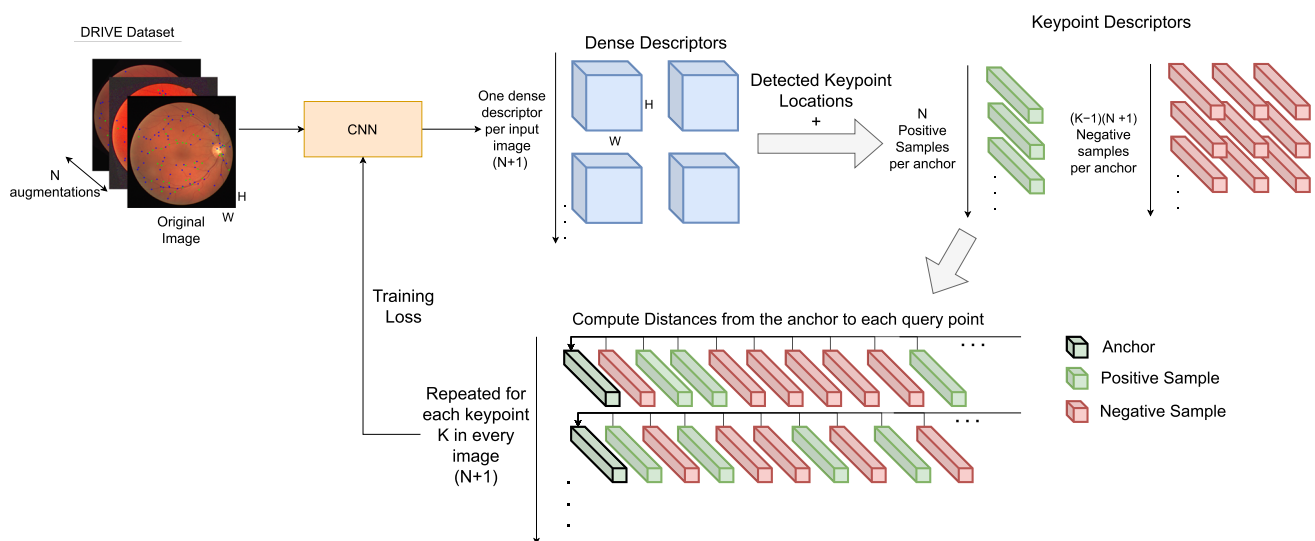


Fig. 2 ConKeD methodology for description learning

N positive samples, and $(K - 1)(N + 1)$ negative samples. As we do not perform any explicit negative point mining, for each keypoint, every other keypoint that is not a positive sample is considered a negative sample, including the ones on the same image where this keypoint was detected. However, in this case, all of the negative samples can be considered hard-negatives. This is because they are all from the set of detected keypoints and, therefore, they are either vessel crossovers or bifurcations (i.e., closely related points). Finally, given the multiple positives and negatives for each anchor point, all the possible pairings are created and all the possible comparisons are computed between the descriptors.

In order to compare the descriptors, we use cosine similarity as the metric. Similarly to other works [30, 32, 34], we apply L2 normalization to the output of the neural network which in combination with the dot product used in the losses described below is equivalent to cosine similarity. Thus, the cosine similarity between two descriptors can be directly computed as the dot product. Moreover, this normalization helps contrastive losses to perform intrinsic hard positive/negative mining [30].

As training loss for ConKeD, we propose two different alternatives: SupCon Loss [30] and MP-InfoNCE [31, 32, 34]. A schematic comparison between both losses can be seen in Fig. 3.

SupCon loss SupCon Loss [30] is a loss function specifically designed for multi-positive multi-negative contrastive

learning in the context of supervised image classification. This loss function can be directly applied to our proposal. For our framework, which uses a multiviewed batch of $N + 1$ images, each of them containing K keypoints, k is the unique index of an arbitrary keypoint such that $k \in [1, K]$. Therefore, the SupCon Loss function is defined as follows:

$$\mathcal{L}_{sup} = \frac{1}{N} \sum_{i=1}^{N+1} \sum_{\substack{j=1, \\ j \neq i}}^{N+1} \sum_{k=1}^K \exp(z_{ik} \cdot z_{jk} / \tau),$$

$$- \log \frac{\exp(z_{ik} \cdot z_{jk} / \tau)}{\sum_{\substack{c=1, \\ c \neq k}}^K \exp(z_{ik} \cdot z_{ic} / \tau) + \sum_{\substack{l=1, \\ l \neq i}}^{N+1} \sum_{c=1}^K \exp(z_{ik} \cdot z_{lc} / \tau)}, \quad (1)$$

where we represent the network output for any keypoint as z . This way, z_{ik} , z_{jk} , and z_{lk} represent the output of network for keypoint k in the views of index i , j , and l , respectively. We use the symbol \cdot to denote the dot product. Finally, $\tau \in \mathbb{R}$ is the scalar temperature parameter.

MP-InfoNCE loss InfoNCE [34], also known as NT-Xent, is a loss function commonly used in many metric learning methods [31, 32, 34]. However, it is only suitable for the paradigm of a single positive and many negatives. Therefore we propose to modify it so that it can be applied in our framework.

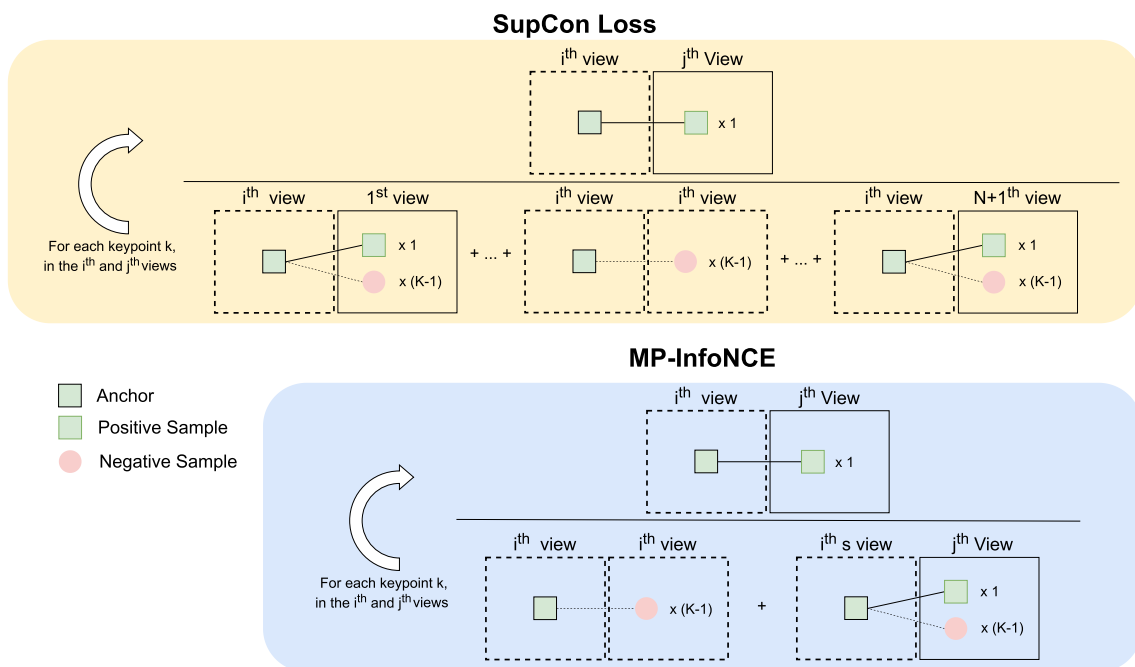


Fig. 3 Schematic comparison between SupCon Loss and MP-InfoNCE, considering $N + 1$ views, where N is the number of augmented versions of the original (1) image, and K keypoints per view, each of which matches with one keypoint on each other view (i.e., N positives per keypoint/sample)

We define our version of InfoNCE, MP-InfoNCE as follows.

$$\mathcal{L}_{NCE} = \frac{1}{\binom{N+1}{2}K} \sum_{i=1}^{N+1} \sum_{\substack{j=1, \\ j>i}}^{N+1} \sum_{k=1}^K \exp(z_{ik} \cdot z_{jk}/\tau) - \log \frac{\exp(z_{ik} \cdot z_{jk}/\tau)}{\sum_{\substack{c=1, \\ c \neq k}}^K \exp(z_{ic} \cdot z_{ic}/\tau) + \sum_{c=1}^K \exp(z_{ic} \cdot z_{jc}/\tau)}, \quad (2)$$

where $N + 1$ is the number of views in our multiviewed batch (1 original image, N augmented views). Since we compute the loss pairwise, the total number of combinations amounts to $\binom{N+1}{2}$. i and j , such that $j > i$, are the indexes of two different images from the multiviewed batch forming a pair. k represents the unique index of a keypoint such that $k \in [1, K]$. This way, z_{ik} represents the output of network for keypoint k in the view of index i . Likewise, z_{jk} represents the output of network for keypoint k in the view of index j .

Overall, the main differences between SupCon Loss and MP-Info-NCE lie in the denominator of the innermost summation term. The denominator in SupCon Loss contains the matching similarities among all keypoints across all views, involving as many positives as augmented views N . Conversely, the denominator of MP-Info-NCE only contrasts against the matching similarities within the views involved in the numerator, and only involving one positive from one augmented view. Note that for $N = 1$, i.e., 2 images, both losses would be equivalent.

2.1.3 Keypoint matching and transformation computation

After the keypoint extraction, the keypoint descriptors are computed by the descriptor network trained using ConKeD, which produces a dense block of descriptors for each image. Using the keypoints' locations, the descriptors for each keypoint are selected. Next, these descriptors are matched among the images in each registration pair. It should be noted that, for a descriptor to match another one, this match should be bidirectional. That is, a given descriptor A matches with another descriptor B, if both A is the closest to B and B is the closest to A. As in the training stage, the distance between descriptors is measured using cosine similarity. Additionally, as the keypoint detection network can accurately classify the keypoints in crossovers or bifurcations, we leverage this capability to ease and speed up the matching. In particular, we only match the descriptors of points within each class, that is, crossovers with crossovers and bifurcations with bifurcations. This drastically reduces the number of required computations, which was already low due to the high specificity of these keypoints.

After the descriptor matching step, the paired keypoints are used to estimate the geometric transformation using RANSAC [22]. This is a well-known algorithm that is commonly used to estimate the parameters of mathematical models when there may be outliers in the data. In this case, the model is the transformation matrix between the images and the data are the matched keypoints. We employ a projective (or homographical) transformation, as previous works in the state of the art [8, 24].

2.2 Experimental setup

2.2.1 Dataset

For the training phase, we employ the public DRIVE dataset. The images were obtained from a diabetic retinopathy screening program using a Canon CR5 non-mydratic 3CCD camera with a FOV of 45°. Each image has a resolution of 584 × 565 pixels. This dataset of 40 images is equally divided in the training and test sets, each of 20 images. This dataset has a ground truth of vessel crossovers and bifurcations that we use to train our keypoint detector [35]. The original binary labeling from DRIVE is converted to heatmaps, as explained in Hervella et al. [33].

For the testing phase, we use FIRE, which is the only public CF dataset which has ground truth for registration. This ground truth consists of a set of labeled control points which can be used to compute the distance between the optimal and the produced registration. This dataset provides 134 image pairs, which can be divided into three categories. Category S has high overlapping between the images that compose each pair. On the contrary, category P has low overlapping. Finally, category A has high overlapping but the images show pathology progression within each pair. This complicates the registration but makes this category the most relevant one for clinical practice, where the registration is usually needed to monitor diseases. The total 134 image pairs are divided into 71 from category S, 49 from P, and 14 from A. Finally, it should be noted that there is an error in the ground truth provided by FIRE. In particular, a single point in an image in category P is incorrectly tagged. We opted to simply discard that point and evaluate the image with one less control point.

2.2.2 Keypoint detection

We train the keypoint detector network as in Hervella et al. [33]. That is, we use the training set of the DRIVE dataset (i.e., 20 images), while reserving 25% of the images for the validation set. The remaining 20 images compose the test set. For this task, we use the U-Net [36] network architecture as in Hervella et al. [33]. To train the network, we use Adam [37] with learning rate decay. The network was trained from scratch. The learning rate is originally set to $1e - 4$ with

a patience of 2500 batches before reducing it. Each batch contains a single image, and each time the learning rate is reduced, it is by a factor of 0.1. The training stops after the learning rate reaches $1e - 7$. We use spatial augmentation consisting of random affine transformations with the following parameters: random rotations of $\pm 90^\circ$, random scaling between 0.9 and $1.1 \times imageSize$, and random shearing of $\pm 20^\circ$. Finally, color augmentation is also used randomly changing image components in the HSV color space [38].

The value for the intensity threshold was selected by testing multiple values and finding the one that provided the best F1-score in the test set of the DRIVE dataset. In this case, 0.35 was found to be the optimal value [33].

2.2.3 Keypoint description

For the training, we use the whole DRIVE dataset (i.e., all the 40 images) in combination with the keypoint output produced by the detection network. Therefore, this network is trained using the DRIVE dataset resolution. As network architecture, we use the modified L2-net [39] proposed in R2D2 [29]. This network produces a descriptor per input pixel without the need to upscale the descriptor block. Moreover, it has been demonstrated to produce accurate descriptors in other works [29]. We train the network from scratch for 15,000 epochs using Adam [37] with a constant learning rate of $1e - 4$. For both losses, SupCon Loss and MP-InfoNCE, we set the temperature parameter $\tau = 0.1$. As augmentations, we use random affine transformations with rotations of $\pm 60^\circ$, translations of $0.25 \times imageSize$ in each axis, scaling between 0.75 and $1.25 \times imageSize$, and shearing of $\pm 30^\circ$. Moreover, we also include color augmentation in the form of random changes in the HSV color space [40] as well as random Gaussian noise with a mean of 0 and a standard deviation of 0.05. The geometric and color augmentations are applied to every augmented image while the random Gaussian noise is applied with a probability of 0.25. During training, the detection of keypoints is performed only on the original images. Therefore, the geometric augmentations are applied to both images and keypoints. Finally, regarding the multiviewed batch sizes, we perform experiments with three different sizes. Using a single original image, we try the multiviewed batch sizes 1 + 1 (i.e., 1 original image and 1 augmentation), 1 + 9, and 1 + 19. Moreover, we also train a descriptor network using the triplet loss, to use as baseline. In order for the comparison to be completely fair, we use cosine similarity as the distance metric. Moreover, we use a margin of value 0.05, which we empirically found to be the most adequate. The training is done using the exact same settings described above except that, in this case, we use a learning rate of $1e - 5$ and train the network during 75,000 epochs.

2.2.4 Keypoint matching

After the description step, we match the descriptors for both images composing a registration pair. The resulting keypoint matches are upscaled from the detection resolution (DRIVE) to the test resolution (FIRE). Once the keypoints are upscaled, they are used in combination with RANSAC [22] to create a projective transformation that aligns the image pair. Therefore, the test images are aligned using their full size which is required to compare among methods.

2.3 Evaluation methodology

To validate our approach, first, we use the evaluation metric proposed by the authors of the FIRE dataset [41]. This metric, called registration score, measures the average Euclidean distance between the control points of each image pair after the registration process. Depending on whether the error of an image pair falls below a threshold, the pair is classified as either successful or unsuccessful. By plotting the ratio of successful registrations on the Y axis and progressively increasing the error threshold along the X axis (defined between 1 and 25 pixels for FIRE), we can create a 2D graph and compute its area under curve (AUC), allowing for quick and easy comparison among methods [20, 21, 23, 24].

We calculate this metric for each category separately as well as for the whole dataset as proposed in FIRE [41]. Moreover, in order to compare our results to Liu et al. [24], we also compute the average (Avg.) of the registration score among the three categories. We also propose to compute a weighted average (W. Avg.) among the three categories, using the number of images per category as the weighting factor. This metric should provide a more representative picture of the performance through the whole dataset.

The registration score only considers the error produced using all the detected keypoints. Thus, this metric may encourage brute-forcing both detection and description to improve the results, without considering the efficiency or computation time. Given an accurate enough description, simply detecting more points (i.e., higher keypoint frequency) improves the chance that each keypoint matches with another one in its approximate corresponding space, even if the detector is pseudo-random. However, this entails a much higher computational cost. Performing the registration with a smaller number of keypoints (i.e., more spaced out) puts more emphasis on both steps, as the detection needs to be repeatable enough so that the keypoints are in approximately the same locations in both image of a pair.

In order to perform a more complete evaluation, we propose a new metric based on the registration score that takes into account the number of matching keypoints. By limiting

the amount of paired keypoints (i.e., after descriptor matching) used in RANSAC, we can better evaluate the quality of the detected keypoints and their descriptors. For this metric, we propose to use the top N keypoints, as ordered by the distance of their descriptors. That is, we take the N keypoints whose descriptors are closest, i.e., which offer a higher quality match. By increasing the value of N (i.e., the number of keypoint pairs) and computing the registration score for each one of the values, we can plot another curve and calculate its AUC. This way, we can evaluate the performance of the network across a different number of keypoints. It should be noted that we cannot have less keypoints than the required by the transformation model; therefore, $N > 4$. Moreover, as we match the keypoints independently per class (i.e., crossovers and bifurcations), we compute the top N points for each class. In particular, we start the evaluation with the top 3 keypoints per class, which is a total of 6 keypoints (i.e., the best 3 crossovers and the best 3 bifurcations). Next, we increase the number to top 4, which is 8 points in total and so on, up to top 25 (with a total of 50 keypoints). In this evaluation, we set the RANSAC budget according to the amount of keypoints such that all RANSAC possibilities are computed. It should be noted that, in some cases, the value of N might be bigger than the number of available or matched keypoints for each class (either crossovers or bifurcations). In this case, no more points are added to that specific class. We name this metric Variable Top Keypoint Registration Score (VTKRS). Similarly to the registration score, we also compute the VTKRS per category and overall for the whole dataset.

3 Results and discussion

This section is structured as follows: first, in Sect. 3.1, we evaluate our method. This includes a comparison with a triplet baseline, as well as analyzing the effect of including multiple positives as well as the effect of increasing the number of positives, by increasing the multiviewed batch size. Next, in Sect. 3.2, we compare both of the proposed loss functions using the best multiviewed batch size. Finally, in Sect. 3.3, we compare our approach with the different state-of-the-art methods.

3.1 Evaluation of the proposed approach

In order to evaluate the performance of ConKeD, we compare our approach to the Triplet Loss, used in state-of-the-art approaches [24]. Additionally, we compare three different variants of our proposal, the first one using a batch of 1 + 1 views such that our method is multi-negative and single-positive and the other two using batches of 1 + 9 and 1 + 19 views each such that our method is both multi-negative and multi-positive. Therefore, we are comparing a single-negative single-positive (triplet) against a multi-negative single-positive (1 + 1 views) to the novel multi-positive multi-negative methods (1 + 9, 1 + 19). For this comparison, we use SupCon Loss [30]. Table 1 shows the results for this experiment.

Firstly, the results show that our proposal clearly outperforms the triplet loss approach, regardless of the multiviewed batch size. Moreover, we have empirically found that training with multi-positive multi-negative losses (either SupCon Loss or MP-InfoNCE) drastically reduces the training time, requiring five times less iterations to converge than the triplet loss. In terms of categories in the FIRE dataset, the difference between methods is less notorious in the category S, which is more forgiving with inaccurate descriptor matches due to the high amount of overlapping. On the other hand, the difference in categories A and P is notable due to the disease progression and low overlapping. These characteristics make accurate descriptor matching more relevant, due to the lower amount of available keypoints. In the category P, there are less keypoints in the overlapping zone and, therefore, less usable keypoints in the registration. Thus, it is more important to correctly match the few usable keypoints, in order to be able to produce an accurate alignment. In category A, due to the progression of pathologies, the amount of keypoints is also relevant, as some may be vanished or obstructed. However, given the higher amount of overlapping, the number of keypoints is less limiting, putting more weight on the robust description that allows to recognize keypoints under different conditions (i.e., pathological progression). Therefore, to evaluate the descriptors, category A is the most relevant one. It should be noted that, as the detected keypoints are exactly the same in all the approaches, the differences are purely due to the improved descriptors.

Table 1 Results measured in registration score AUC for the different number of image view sizes using SupCon Loss as well as the baseline with triplet loss

	FIRE	A	P	S	Avg	W. Avg
Triplet (SP-SN)	0.696	0.683	0.349	0.939	0.657	0.697
Proposed 1 + 1 (SP-MN)	0.748	0.743	0.469	0.942	0.718	0.748
Proposed 1 + 9 (MP-MN)	0.755	0.76	0.477	0.946	0.728	0.755
Proposed 1 + 19 (MP-MN)	0.754	0.757	0.477	0.944	0.726	0.753

Best results in bold. *SP* means single-positive, *SN* means single-negative, *MP* means multi-positive, and *MN* means multi-negative

Table 2 Results measured in VTKRS for the different number of image views using SupCon Loss as well as the baseline with triplet loss

	FIRE	A	P	S	Avg	W. Avg
Triplet (SP-SN)	0.602	0.613	0.231	0.857	0.567	0.603
Proposed 1 + 1 (SP-MN)	0.626	0.676	0.324	0.824	0.623	0.642
Proposed 1 + 9 (MP-MN)	0.656	0.685	0.368	0.849	0.634	0.656
Proposed 1 + 19 (MP-MN)	0.654	0.696	0.344	0.862	0.631	0.650

Best results in bold. *SP* means single-positive, *SN* means single-negative, *MP* means multi-positive, and *MN* means multi-negative

Regarding the different multiviewed batch sizes, overall, the results are accurate and the differences between the different number of views are small but highly relevant. These differences in the results are constant across all the FIRE categories, creating a clear trend. In particular, the most notable difference is between the batch size 1 + 1 and the other two (1 + 9 and 1 + 19). In this case, the results on FIRE are around 1% less with the 1 + 1 batch size than with bigger batch sizes. Similarly, in category P, the difference among the 1 + 1 batch and the rest is around 1% as well. For category A, the difference increases to almost 2% while on category S, it is less than 1%. This exemplifies the different levels of difficulty in learning adequate descriptors for each category, as previously stated. Category S, having the most overlapping, is the easiest, while A and P are more complex due to the pathologies and low overlapping.

Table 2 shows the results of our evaluation using VTKRS. These results are shown graphically in Fig. 4. Similarly to

the previous analyses, the difference between our proposal and the triplet loss baseline is clear. In category S, where the descriptor matching is less relevant, the results for all the approaches are equivalent. However, in categories A and P, the differences are notable and can be clearly seen in Fig. 4, where the triplet VTKRS AUC is significantly lower than in our proposals. Overall, in FIRE, due to the prevalence of the category S, the difference between our method and the triplet baseline is less notable but still highly significant. Regarding the different variants of our proposal, it can be seen that adding multiple positives (i.e., 1 + 9 and 1 + 19) consistently improves the performance across all categories and, especially, in the more complicated ones (A and P).

Regarding the number of views in a multi-positive setting (i.e., 1 + 9 and 1 + 19), the results are equivalent with both metrics. Moreover, given the extended training in our experimentation, i.e., 15,000 epochs, we can ensure that the difference in performance between 1 + 1 and other

Fig. 4 AUC curves for the different losses measured in VTKRS. In the legend, “Max.” represents the maximum registration score AUC produced using all the keypoints

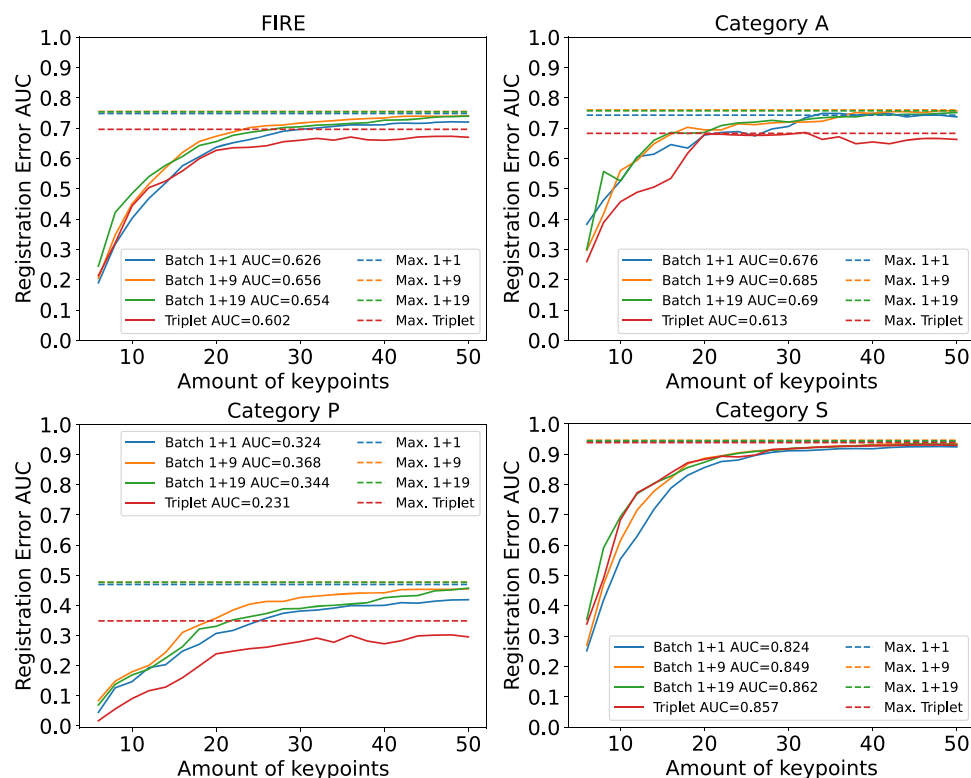


Table 3 Results for the different losses measured in registration score

	FIRE	A	P	S	Avg	W. Avg
SupCon loss	0.755	0.76	0.477	0.946	0.728	0.755
MP-InfoNCE	0.758	0.749	0.489	0.945	0.728	0.758

Best results highlighted in bold

Table 4 Results for the different losses measured in VTKRS

	FIRE	A	P	S	Avg	W. Avg
SupCon loss	0.656	0.685	0.368	0.849	0.634	0.656
MP-InfoNCE	0.656	0.692	0.362	0.852	0.635	0.656

Best results highlighted in bold

two approaches is not due to the different number of seen image views. All three approaches are trained well beyond their convergence point, such that there are no more meaningful improvements possible. This serves to isolate the improvements of the multi-positive approach from the potential improvements of using a higher number of views per batch and thus allowing the network to see more images per iteration. Therefore, the difference between adding multiple positives (1 + 1 vs the rest) and just increasing the number of views in the batch (1 + 9 vs 1 + 19) is clear. Given that the results of batches 1 + 9 and 1 + 19 are equivalent, we choose 1 + 9 as the reference batch to compare losses and models in the following sections.

3.2 Comparison of loss functions

In this section, we compare SupCon Loss against MP-InfoNCE. The experiments are performed using the optimal multiviewed batch size found in the previous experiments (i.e., 1 + 9). The results are detailed in Tables 3 and 4 as well

as in Fig. 5. Representative examples of registered images from the FIRE dataset using our both losses can be seen in the top rows of Fig. 6.

In terms of registration score over the FIRE dataset, detailed in Table 3, the results for both losses are very similar. The biggest differences appear in category A (+1.1% for SupCon) and category P (+1.2% for MP-InfoNCE). However, these differences are small due to the limited nature of the registration score evaluation over FIRE. The results in terms of VTKRS can be seen in Table 4 and in Fig. 5. The final numeric differences in the AUC values for both losses are small, like in the previous evaluation. However, looking at the curves, the differences are much more evident. In particular, for category A, we can see that MP-InfoNCE provides results that are more accurate with a low number of keypoints. In category P, both losses perform similarly, and MP-InfoNCE obtains better results with less points while SupCon Loss improves more given more points. However, in the end, the two approaches become virtually even when many or all of the keypoints are used.

Fig. 5 AUC curves for the different losses measured in VTKRS. In the legend, “Max.” represents the maximum registration score AUC produced using all the keypoints

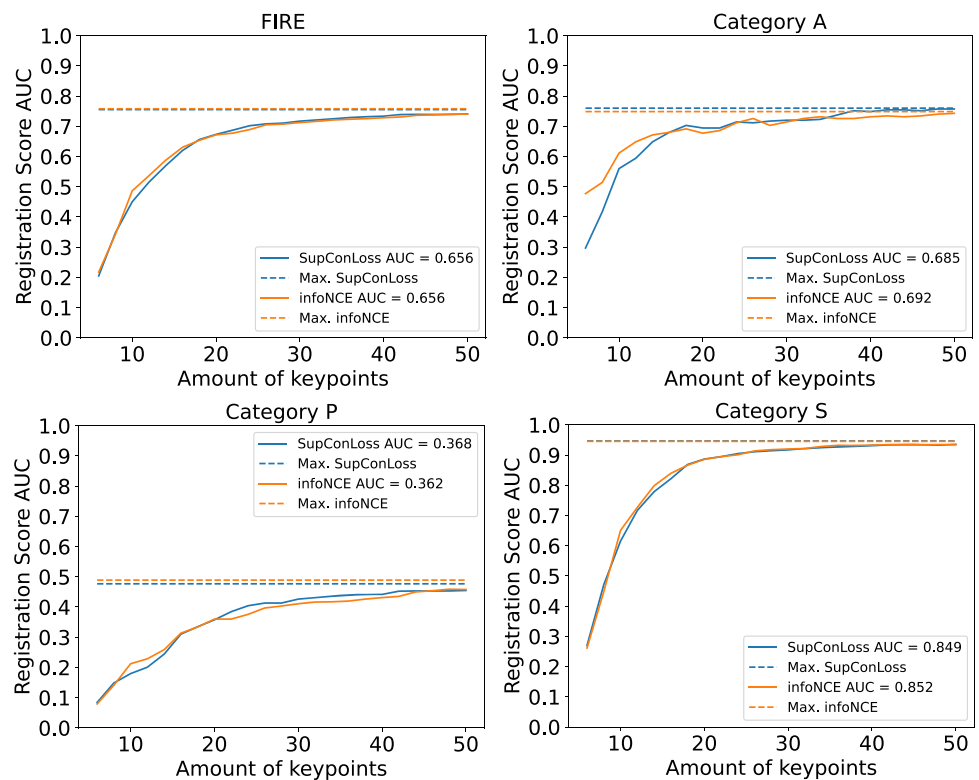
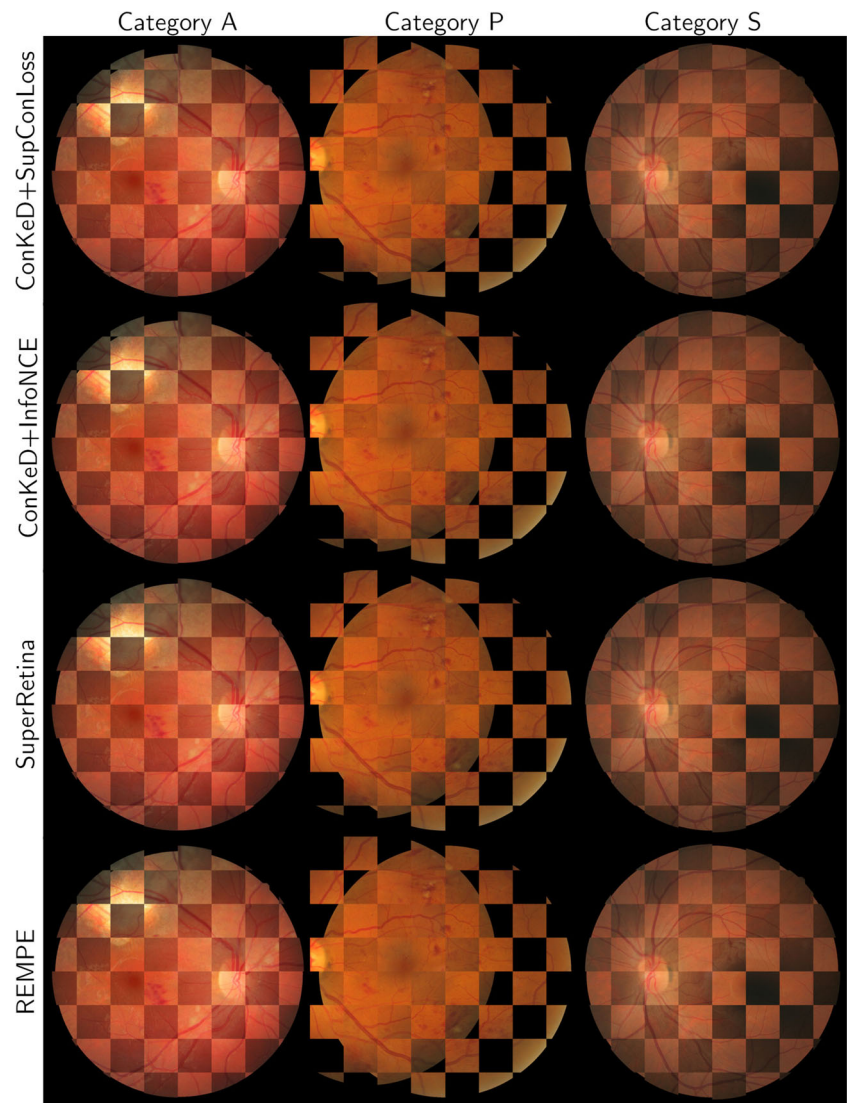


Fig. 6 Representative registration examples from the FIRE dataset for our proposals and different state-of-the-art methods



The difference in performance in the category A (the only one where the difference is notable) is especially relevant as this is the category with the most clinical relevance. RIR is commonly used for disease monitoring and study follow-up; therefore, any automated method should desirably be robust to pathological lesions. Additionally, category A can be viewed as the hardest category to align due to the number and variety of lesions and diseases that severely alter the morphology of the ocular fundus. While category P obtains the lowest scores, this is due to the lower overlapping and thus the lower amount of common keypoints on both images of the registration pair. However, in category A, the detected keypoints can disappear or be altered by pathology progression. Therefore, improving the results in this category is desirable and even more so with a low number of keypoints as this makes the network more robust to pathology progression which can completely occlude the keypoints. Given this benefit of MP-InfoNCE over SupCon Loss, we will use

the MP-InfoNCE as the method to compare our approach against the state of the art. Nevertheless, while there are differences between MP-InfoNCE and SupCon Loss, both loss terms produce satisfactory performance. Therefore, we can conclude that our multi-negative multi-positive approach is robust to the specific loss term used.

3.3 Comparison with state-of-the-art registration methods

The comparison with the state of the art is depicted in Table 5. In this comparison, we include the three best-performing approaches on the FIRE dataset as well as previous approaches using deep learning. The results show that the best-performing method overall is VOTUS [21] and the best deep learning method is SuperRetina [24]. However, it should be noted that SuperRetina does not provide the AUC for the whole FIRE dataset as it is done in every other work

Table 5 Comparison between our approach and state-of-the-art methods, sorted by average

	FIRE	A	P	S	Avg	W. Avg	Keypoints	Training images
VOTUS [21]	0.812	0.681	0.672	0.934	0.762	0.811	N/A	N/A
SR Manual [24]	–	0.783	0.542*	0.94	0.755*	0.780*	740	97 + 844
SR PBO [24]	–	0.789	0.516*	0.944	0.750*	0.773*	740	97 + 844
<i>Proposed</i>	0.758	0.749	0.489	0.945	0.728	0.758	115	15 and 40
REMPE [20]	0.773	0.66	0.542	0.958	0.72	0.774	–	N/A
Retina-R2D2 [8]	0.695	0.726	0.352	0.925	0.645	0.575	5000	1748
Rivas-Villar [23]	0.657	0.660	0.293	0.908	0.620	0.552	115	20

The asterisk symbol (*) indicates that the results were calculated using one image less. *N/A* indicates that the method does not use keypoints or does not require training samples. The hyphen symbol (-) indicates that the method is not publicly available and does not specify the number of keypoints used. $k + d$ denotes that k images are used for the detector training and $k + d$ images are used for the descriptor training. Meanwhile, k and d denote that k images are used for the detector training and d images are used for the descriptor training. The best results are in bold. Results sourced from the referenced papers

in the state of the art [20, 21, 23]. Analyzing each category separately, we can see that in A, our proposal obtains the second-best results only behind SuperRetina by around 4%. In category P, both our method and SuperRetina produce significantly worse results than VOTUS [21], a classical method that uses a quadratic transformation to help improving the alignment in these low overlapping cases. Finally, in category S, which is the easiest to register, the results of most methods are within a margin of around 1%; therefore, the differences are not notable.

In terms of registration score, our method is the second-best deep learning approach and overall provides results that are close to those reported by SuperRetina. Regarding the different categories, the most relevant difference is in category P where our method obtains 2.7% less and 5.3% less AUC than both SuperRetina methods. As explained, this is due to the lower number of keypoints detected per image. In the other categories, where the amount of keypoints is less relevant, our results are even closer, even surpassing SuperRetina in category S. Importantly, our method provides several advantages, and there are also some important experimental differences between both works that should be considered. SuperRetina trains its detector using a private dataset containing 97 images whereas we train the detector with just 15. Additionally, SuperRetina uses an extra set of 844 images to train the descriptor, for a total of 941, whereas we use just 40 images for our descriptor. Furthermore, SuperRetina requires extensive ad hoc image pre-processing. Moreover, the descriptor computed by SuperRetina has 256 channels compared to 128 in our descriptor. Additionally, the results of SuperRetina depend greatly on the number of detected keypoints. As shown in Table 5, SuperRetina detects more than $6\times$ more keypoints than our approach. In that regard, the authors of SuperRetina also tested a variant of their method that detects approximately the same amount as our work (using just the labeled keypoints in their training set),

obtaining a notable performance drop, from 0.755 to 0.685 (Avg.) in which case our method, which produces approximately the same number of keypoints, would obtain better results. Using more keypoints usually increases performance as it provides more chances for the descriptors to match and more paired keypoints to compute a transformation, at the cost of more computation. This is especially notable in category P where the small overlapping limits the number of keypoints usable for registration. Additionally, SuperRetina requires double inference (i.e., double run), including double detection, description, descriptor matching, and transformation calculation. In contrast, our method does not require any of these extra steps making it more efficient.

Our method was designed with efficiency in mind as our multi-positive multi-negative training approach allows us to leverage more information during training, allowing for increased data efficiency. Additionally, through the use of domain-specific keypoints, we can limit the cost of matching and transformation estimation. Overall, our proposed approach demonstrates execution times that are suitable for use in clinical day-to-day practice, providing fast image registrations. In particular, the average execution time for our whole methodology including each step in the process is, approximately, 0.088 s per image pair. This time can be itself divided into the different steps of the method: 0.03 s for keypoint detection inference, 0.01 s for descriptor inference, 0.0006 for descriptor matching, and 0.047 for transformation computation through RANSAC. In relation to this, our method detects, on average, 115 keypoints throughout the FIRE dataset. Given that the complexity of the descriptor matching stage is $O(n^2)$, this means that $115^2 = 13225$ comparisons are performed for each image pair. However, if our approach required the same amount of keypoints as SuperRetina (740 keypoints) or Retina-R2D2 (5000 keypoints), the number of comparisons required for the matching process would increase to 547,600 or 25,000,000, respectively.

This would increase the matching time by factors of $41\times$ and $1890\times$. Moreover, SuperRetina uses heavier descriptors and double inference, which would increase execution time and computational cost further. Our experiments revealed that running our method a second time does not improve the results. This evidences that our descriptors and keypoints are more robust to transformations since they are correctly calculated in a single pass, despite the larger transformations.

Finally, representative examples of registered images from the FIRE dataset using our approach as well as available state-of-the-art methods can be seen in Fig. 6. While there are slight numerical differences between our approach and others, such as SuperRetina, they can hardly be visually spotted in the images. This is due to the high resolution of the images in FIRE. Small registration errors (i.e., 1 or 2 pixels) become really hard to spot on images of such size while significantly impacting numerical results.

In conclusion, overall, our method, with the novel ConKeD description framework, provides very similar results to SuperRetina while requiring significantly less training samples, avoiding the need for ad hoc pre-processing, using far less keypoints, and just requiring a single execution.

4 Conclusion

In this work, we proposed ConKeD, the first multi-positive multi-negative contrastive descriptor learning approach for keypoint-based image registration. As it uses multiple positives and negatives, it can leverage additional information from the input images, which improves the resulting descriptors. We propose a complete methodology for the registration of CF images, using vessel crossovers and bifurcations as keypoints together with the descriptors learned using ConKeD.

In order to validate our method, we perform several experiments in the well-known FIRE dataset and analyze several factors in the methodology. The results show that our proposal clearly outperforms triplet loss, the alternative typically used to learn descriptors in the state of the art. In terms of the losses, we found that SupCon Loss and a multi-positive InfoNCE performed similarly with a slight advantage to the latter evidencing that our framework is robust to specific loss terms. Moreover, our method provides results competitive with the best deep learning state-of-the-art method while using less training samples and less keypoints and not requiring double inference or pre-processing.

As future work, it would be interesting to test more losses in combination with our ConKeD framework. In this regard, developing specific losses for the multi-positive paradigm is also desirable to try and improve the results. Additionally, a limitation of our current approach and of others in the state of the art is the reliance on inherently limited domain-specific

keypoints, such as crossovers and bifurcations. These keypoints are not distributed evenly across the retina and may be completely absent in some parts. This can reduce the overall registration performance since there are no points to accurately drive the transformation in every part of the fundus. The arbitrary nature of these domain-specific keypoints may also reduce the performance of these methods on patients with less crossovers and bifurcations visible in the eye fundus. Therefore, future works could focus on developing methods to increase the number of keypoints while keeping a high specificity. This could potentially improve the results, especially in category P, which has the lowest overlapping. In that regard, it could even be convenient to study the inclusion of generic keypoints. These, despite their generally lower specificity, could help regularize the registration if there are zones of the image that lack domain-specific keypoints. Therefore, the inclusion of generic keypoints in a measured way may be beneficial without a significant reduction in the efficiency of the method.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature. This work is supported by Ministerio de Ciencia e Innovación, Government of Spain, through the PID2019-108435RB-I00, TED2021-131201B-I00, and PDC2022-133132-I00 research projects; Consellería de Cultura, Educación e Universidade, Xunta de Galicia, through Grupos de Referencia Competitiva ref. ED431C 2020/24, predoctoral fellowship ref. ED481A 2021/147 and the postdoctoral fellowship ref. ED481B-2022-025; and Instituto de Salud Carlos III (ISCIII) under the grant FORT23/00010 as part of the Programa FORTALECE of Ministerio de Ciencia e Innovación. Funding for open access charge: Universidade da Coruña/CISUG.

Declarations

Conflict of Interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Viergever MA, Maintz JBA, Klein S, Murphy K, Staring M, Pluim JPW (2016) A survey of medical image registration-under review. *Med Image Anal* 33:140–144. <https://doi.org/10.1016/j.media.2016.06.030>. 20th anniversary of the Medical Image Analysis journal (MedIA)

2. Hajnal J, Hill D, Hawkes DJ (2001) Medical image registration. Biomedical engineering series, pp 1–383. CRC Press, Boca Raton, FL
3. Narasimha-Iyer H, Can A, Roysam B, Tanenbaum HL, Majerovics A (2007) Integrated analysis of vascular and nonvascular changes from color retinal fundus image sequences. *IEEE Trans Biomed Eng* 54(8):1436–1445. <https://doi.org/10.1109/TBME.2007.900807>
4. Forrester JV, Dick AD, McMenamin PG, Roberts F, Pearlman E (2020) *The Eye E-book: Basic Sciences in Practice*. Elsevier Health Sciences, United Kingdom
5. Hernandez-Matas C, Zabulis X, Argyros AA (2021) Retinal image registration as a tool for supporting clinical applications. *Comput Methods Programs Biomed* 199:105900. <https://doi.org/10.1016/j.cmpb.2020.105900>
6. Ho R, Song LD, Choi JA, Jee D (2018) The cost-effectiveness of systematic screening for age-related macular degeneration in South Korea. *PLoS ONE* 13(10):1–14. <https://doi.org/10.1371/journal.pone.0206690>
7. Salmon JF (2020) *Kanski's clinical ophthalmology: a systematic approach*. Elsevier, Edinburgh
8. Rivas-Villar D, Hervella S, Rouco J, Novo J. Joint keypoint detection and description network for color fundus image registration. *Quantitative Imaging in Medicine and Surgery*
9. Rivas-Villar D, Motschi AR, Pircher M, Hitzemberger CK, Schranz M, Roberts PK, Schmidt-Erfurth U, Bogunović H (2023) Automated inter-device 3D oct image registration using deep learning and retinal layer segmentation. *Biomed Opt Express* 14(7):3726–3747. <https://doi.org/10.1364/BOE.493047>
10. Pluim JPW, Maintz JBA, Viergever MA (2000) Image registration by maximization of combined mutual information and gradient information. In: Delp SL, DiGoia AM, Jaramaz B (eds) *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2000*. Springer, Berlin, Heidelberg, pp 452–461
11. Balakrishnan G, Zhao A, Sabuncu MR, Dalca AV, Guttag J (2018) An unsupervised learning model for deformable medical image registration. In: 2018 IEEE/CVF conference on computer vision and pattern recognition, pp 9252–9260. <https://doi.org/10.1109/CVPR.2018.00964>
12. Cheng X, Zhang L, Zheng Y (2018) Deep similarity learning for multimodal medical images. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization* 6(3):248–252. <https://doi.org/10.1080/21681163.2015.1135299>
13. Haskins G, Kruger U, Yan P (2020) Deep learning in medical image registration: a survey. *Mach Vis Appl* 31(1):8. <https://doi.org/10.1007/s00138-020-01060-x>
14. Balakrishnan G, Zhao A, Sabuncu MR, Guttag J, Dalca AV (2019) VoxelMorph: a learning framework for deformable medical image registration. *IEEE Trans Med Imaging* 38(8):1788–1800. <https://doi.org/10.1109/TMI.2019.2897538>
15. Benvenuto GA, Colnago M, Dias MA, Negri RG, Silva EA, Casaca W (2022) A fully unsupervised deep learning framework for non-rigid fundus image registration. *Bioengineering* 9(8). <https://doi.org/10.3390/bioengineering9080369>
16. Hervella S, Rouco J, Novo J, Ortega M (2018) Multimodal registration of retinal images using domain-specific landmarks and vessel enhancement. *Procedia Computer Science* 126, 97–104. <https://doi.org/10.1016/j.procs.2018.07.213>. Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 22nd International Conference, KES-2018, Belgrade, Serbia
17. Rivas-Villar D, Motschi AR, Pircher M, Hitzemberger CK, Schranz M, Roberts PK, Schmidt-Erfurth U, Bogunović H (2023) Automated inter-device 3D OCT image registration using deep learning and retinal layer segmentation. *Biomed Opt Express* 14(7):3726–3747. <https://doi.org/10.1364/BOE.493047>
18. de Vos BD, Berendsen FF, Viergever MA, Sokooti H, Staring M, Išgum I (2019) A deep learning framework for unsupervised affine and deformable image registration. *Med Image Anal* 52:128–143. <https://doi.org/10.1016/j.media.2018.11.010>
19. Xiao H, Teng X, Liu C, Li T, Ren G, Yang R, Shen D, Cai J (2021) A review of deep learning-based three-dimensional medical image registration methods. *Quantitative Imaging in Medicine and Surgery* 11(12)
20. Hernandez-Matas C, Zabulis X, Argyros AA (2020) REMPE: registration of retinal images through eye modelling and pose estimation. *IEEE J Biomed Health Inform* 24(12):3362–3373. <https://doi.org/10.1109/JBHI.2020.2984483>
21. Motta D, Casaca W, Paiva A (2019) Vessel optimal transport for automated alignment of retinal fundus images. *IEEE Trans Image Process* 28(12):6154–6168. <https://doi.org/10.1109/TIP.2019.2925287>
22. Fischler M, Bolles R (1981) Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun ACM* 24:381–395
23. Rivas-Villar D, Hervella S, Rouco J, Novo J (2022) Color fundus image registration using a learning-based domain-specific landmark detection methodology. *Comput Biol Med* 140:105101. <https://doi.org/10.1016/j.compbiomed.2021.105101>
24. Liu J, Li X, Wei Q, Xu J, Ding D (2022) Semi-supervised keypoint detector and descriptor for retinal image matching. In: Avidan S, Brostow G, Cissé M, Farinella GM, Hassner T (eds) *Computer Vision - ECCV 2022*. Springer, Cham, pp 593–609
25. Zou B, He Z, Zhao R, Zhu C, Liao W, Li S (2020) Non-rigid retinal image registration using an unsupervised structure-driven regression network. *Neurocomputing* 404:14–25. <https://doi.org/10.1016/j.neucom.2020.04.122>
26. DeTone D, Malisiewicz T, Rabinovich A (2018) SuperPoint: self-supervised interest point detection and description. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops
27. Oinonen H, Forsvik H, Ruusuvoori P, Yli-Harja O, Voipio V, Huttunen H (2010) Identity verification based on vessel matching from fundus images. In: 2010 IEEE International conference on image processing, pp 4089–4092. <https://doi.org/10.1109/ICIP.2010.5650657>
28. Pizer SM, Amburn EP, Austin JD, Cromartie R, Geselowitz A, Greer T, ter Haar Romeny B, Zimmerman JB, Zuiderveld K (1987) Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing* 39(3):355–368. [https://doi.org/10.1016/S0734-189X\(87\)80186-X](https://doi.org/10.1016/S0734-189X(87)80186-X)
29. Revaud J, De Souza C, Humenberger M, Weinzaepfel P (2019) R2D2: reliable and repeatable detector and descriptor. In: Wallach H, Larochelle H, Beygelzimer A, Alché-Buc F, Fox E, Garnett R (eds) *Advances in neural information processing systems*, vol 32. Curran Associates Inc, Vancouver, Canada
30. Khosla P, Teterwak P, Wang C, Sarna A, Tian Y, Isola P, Maschinot A, Liu C, Krishnan D (2020) Supervised contrastive learning. In: Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H (eds) *Advances in neural information processing systems*, vol 33. Curran Associates Inc, Vancouver, Canada, pp 18661–18673
31. He K, Fan H, Wu Y, Xie S, Girshick R (2020) Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)
32. Chen T, Kornblith S, Norouzi M, Hinton G (2020) A simple framework for contrastive learning of visual representations. In: III, H.D., Singh, A. (eds.) Proceedings of the 37th international conference on machine learning. Proceedings of machine learning research, vol 119, pp 1597–1607. PMLR, Vienna, Austria
33. Hervella S, Rouco J, Novo J, Penedo MG, Ortega M (2020) Deep multi-instance heatmap regression for the detection of reti-

nal vessel crossings and bifurcations in eye fundus images. *Comput Methods Programs Biomed* 186:105201. <https://doi.org/10.1016/j.cmpb.2019.105201>

34. Oord A, Li Y, Vinyals O (2018) Representation learning with contrastive predictive coding. [arXiv:1807.03748](https://arxiv.org/abs/1807.03748)
35. Abbasi-Sureshjani S, Smit-Ockeloen I, Bekkers E, Dashtbozorg B, Romeny BTH (2016) Automatic detection of vascular bifurcations and crossings in retinal images using orientation scores. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI), pp. 189–192. <https://doi.org/10.1109/ISBI.2016.7493241>
36. Ronneberger O, Fischer P, Brox T (2015) U-Net: convolutional networks for biomedical image segmentation. In: Navab N, Hornegger J, Wells WM, Frangi AF (eds) *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015*. Springer, Cham, pp 234–241
37. Kingma D, Ba J (2015) Adam: a method for stochastic optimization. In: *International Conference on Learning Representations (ICLR)*
38. Liskowski P, Krawiec K (2016) Segmenting retinal blood vessels with deep neural networks. *IEEE Trans Med Imaging* 35(11):2369–2380. <https://doi.org/10.1109/TMI.2016.2546227>
39. Tian Y, Fan B, Wu F (2017) L2-Net: deep learning of discriminative patch descriptor in Euclidean space. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp 6128–6136. <https://doi.org/10.1109/CVPR.2017.649>
40. Hervella S, Rouco J, Novo J, Ortega M (2020) Self-supervised multimodal reconstruction of retinal images over paired datasets. *Expert Syst Appl* 161:113674. <https://doi.org/10.1016/j.eswa.2020.113674>
41. Hernandez-Matas C, Zabulis X, Triantafyllou A, Anyfanti P, Douma S, Argyros A (2017) FIRE: fundus image registration dataset. *Journal for Modeling in Ophthalmology* (to appear)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



David Rivas-Villar received his degree in Computer Engineering in 2019 from the University of A Coruña (Spain), and his M.Sc. degree in Bioinformatics for Health Sciences in 2021. Since 2020, he is pursuing his Ph.D. degree in Computer Science and Artificial Intelligence. He has also been visiting researcher at the Medical University of Vienna. His main research interests lie in computer vision, self-supervised learning, medical image analysis and machine learning.



computer vision, and medical image processing and analysis.

Álvaro S. Hervella received his Ph.D. degree (cum laude) in Computer Science from the University of A Coruña (UDC), Spain, in 2022. Currently, he is a postdoctoral research fellow at UDC. He is also affiliated at the Institute of Biomedical Research of A Coruña (INIBIC), in Spain. Previously, he has also been research assistant at CEIT technology center in Spain and visiting researcher at INESC TEC research center in Portugal. His research interests lie in the areas of machine learning, com-



where he currently serves as Assistant Professor since 2019. Since 2016 he is also a researcher at INIBIC, UDC, Spain. His research interests lie in the areas of computer vision, pattern recognition, machine learning and biomedical image processing and analysis.

José Rouco received his M.Sc. degree in Computer Science from the University of A Coruña (UDC), Spain, in 2005, and his Ph.D. degree (cum laude) in Computer Science and Artificial Intelligence from UDC in 2011. He worked as Postdoctoral Research Fellow at INEB (2012–2014) and INESC-TEC (2014–2016), University of Porto, Portugal. In 2016 he joined as non-tenured Professor at the Dept. of Computer Science and Information Technologies of UDC,



professor in the Department of Computer Science and Information Technologies at the University of A Coruña, and his main research interests lie in the fields of computer vision, pattern recognition and biomedical image processing.

Jorge Novo received his M.Sc. and Ph.D. degrees (cum Laude) in Computer Science from the University of A Coruña in 2007 and 2012, respectively. He has also worked as a visiting researcher with CMR images in the detection of landmark points at Imperial College London, and as a postdoctoral research fellow at the INEB and INESC-TEC research institutes in the development of CAD systems for lung cancer diagnosis with chest CT images. He is currently an Associate Pro-