



**Hôpital ophtalmique
Jules-Gonin**

Service universitaire d'ophtalmologie
Fondation Asile des aveugles

**École Polytechnique Fédérale de Lausanne
School of Life Sciences**

Automatic Grading of Vasculitis Inflammation in Fluorescein Angiography Images

Master project in Life Sciences Engineering
by Victor Amiot

Carried out in the Data Science Group at Jules-Gonin Hospital under the supervision of:
PhD. Mattia Tomasoni and Dre. Florence Hoogewoud, HOJG
Thesis Supervisors & External Experts

Under the direction of:

Prof. Jean-Marc Odobez with Dr. Andre Anjos, EPFL, IDIAP
Thesis Advisor

EPFL STI IEL LIDIAP
ELE 131 (Bâtiment ELE)
Station 14
CH-1015 Lausanne

Lausanne, 27 January 2023

ABSTRACT

Vasculitis is a disease that affects the retinal blood vessels, among other structures of the eyes. A so-called Fluorescein Angiography (FA) exam consists in capturing a time-lapse the fundus of the eye as retinal vessels are progressively being filled by a fluorescent dye, which is injected intravenously. Vasculitis weakens the wall of retinal vessels causing them to gradually leak their content into the surrounding tissues. This effect, impossible to detect at the naked eye, becomes visible during a FA exam as the dye progressively leaks causing diffused staining of the tissue surrounding the vessels.

As of today, the grading of the severity of Vasculitis signs in Fluorescein Angiography (FA) images (and of retinal vascular leakage, in particular) is performed manually. Acquiring the ability to interpret and grade such images requires years of training after a specialisation in ophthalmology.

I designed and developed an automatic pipeline for the grading of inflammation signs caused by Vasculitis, in cooperation with members of the team that hosted this MSc thesis work. Specifically, I focused on the detection of vascular leakage on late frames of the FA time-lapse. I tested several registration methods to match vessel structures across multiple frames that make up a single Fluorescein Angiography (FA) exam. I validated the output based on a novel criterion that I designed for this purpose. Thanks to a specific combination of registration techniques, the amount of frames that were successfully registered, on average, increases significantly compared to results obtained with off-the-shelf methods. I used Computer Vision methods to segment the vasculature in early frames and identify retinal background area where leakage may become visible in late frames. Our pipeline was applied on 543 patients, reading raw Fluorescein Angiography (FA) data from the clinic, preparing the dataset for expert grader annotated, and performing detection of staining due to vascular leakage. Finally, I designed strategies to grade vascular leakage according to a formal scale, recognised by Vasculitis experts, and used in the Vasculitis clinic of the Jules-Gonin hospital, where I carried out this work. Applying our pipeline to a first set of annotations performed by human experts, resulted in the detection of cases of diffuse staining with promising preliminary results.

Contents

1	Introduction	1
2	Theoretical Background	3
2.1	Vasculitis disease	3
2.1.1	Variations of the disease and affected population	3
2.1.2	Fluorescein Angiography	3
2.1.3	Associated symptoms	3
2.1.4	Tugal-Tuktun grading	4
2.2	Registration	5
2.2.1	Keypoints detection	5
2.2.2	Keypoints matching	6
2.2.3	Homography matrix	6
2.2.4	Image similarity measures	6
2.3	Computer vision methods	7
2.3.1	Sato tubeness filter	7
2.3.2	Otsu's algorithm	8
2.3.3	Morphological transformations: opening	8
3	Design	10
3.1	Dataset	10
3.1.1	Characteristics	10
3.1.2	Dataset structure	10
3.2	Registration	11
3.2.1	Standard registration	11
3.2.2	Frame reference selection and back-registration	12
3.2.3	Keypoints matching	12
3.2.4	Keypoints matches selection	12
3.2.5	Quality assessment	13
3.3	Vascular leakage detection	13
3.3.1	Generic mask creation	13
3.3.2	Vessel mask	14
3.3.3	Background mask	15
3.3.4	Optic disk detection	15
3.3.5	Hyper-fluorescence detection	16
3.4	Expert's grading	16
3.4.1	Format processing	16

3.4.2	Grading form	17
3.4.3	Hyper-parameters tuning	18
3.5	Pipeline	19
4	Results	20
4.1	I implemented a method to register Fluorescein Angiography photographs, with higher performances compared to off-the-shelf approaches, and precision that satisfied human experts	20
4.2	I implemented a method to segment the vasculature in Fluorescein Angiography photographs, using Computer Vision	22
4.3	I implemented a complete pipeline for the grading of vascular leakage in Fluorescein Angiography (FA) photographs	23
4.4	I defined a strategy to enable Fluorescein Angiography photographs grading and detect vascular leakage	24
5	Discussion	26
5.1	Interpretation of registration process	26
5.2	Limitation of current work	27
5.3	Future directions	27
6	Conclusion	29

Acronyms

AIDS Acquired ImmunoDeficiency Syndrome.

BRIEF Binary Robust Independent Elementary Features.

CHU Centre Hospitalier Universitaire.

EPFL École Polytechnique Fédérale de Lausanne.

FA Fluorescein Angiography.

FLANN Fast Library for Approximate Nearest Neighbors.

GAN Generative Adversarial Network.

HOJG Hopital Ophtalmique Jules-Gonin.

IDIAP Institut Dalle Molle d'Intelligence Artificielle Perceptive.

kNN k-Nearest-Neighbours.

LUKS Luzern Kantonsspital.

MI Mutual Information.

MSE Mean Squared Error.

NMI Normalized Mutual Information.

OCT Optical Coherence Tomography.

ORB Oriented FAST and Rotated BRIEF.

PID Patient ID.

POC Phase-Only Correlation.

RANSAC Random Sample Consensus.

SIFT Scale-Invariant Feature Transform.

SOIN Swiss Ophthalmic Imaging Network.

SSIM Structural Similarity Index Measure.

UML Unified Modeling Language.

Chapter 1

Introduction

Vasculitis is an inflammatory disease that affects notably retinal blood vessels and may lead to vision loss. The acquisition of Fluorescein Angiography (FA) photographs [1] is the best method to reveal Vasculitis. The technique allows to clearly image the retinal blood vessels thanks to a fluorescent dye injected intravenously. Retinal vessels that are compromised by the disease, progressively leak part of the dye, causing staining of the surrounding tissues, that become visible in FA images between 5 and 10 minutes after the injection. Vasculitis can manifest as several inflammatory symptoms [2] such as Macular Edema and staining of the optic disk during FA. At the Jules-Gonin ophthalmic hospital, the severity of the disease is evaluated based on the standard Tugal-Tuktun grading [3], which takes into account the many manifestations of the disease. Among them, retinal vascular leakage is graded as: focal, multifocal or diffuse.

This MSc thesis work is part of a multi-centric, multi-grant project, named *CAD4IED*. It involves Vasculitis experts and Data Science researchers from Hopital Ophtalmique Jules-Gonin (HOJG), Institut Dalle Molle d'Intelligence Artificielle Perceptive (IDIAP), Luzern Kantonsspital (LUKS) and CHU Grenoble. The overall aim is to automatically grade Vasculitis cases based on the Tugal-Tuktun method, once a Vasculitis diagnosis has been performed by a human expert. The more complex task of automatic detection of Vasculitis is beyond the scope of this project. Future phases of the project might attempt to distinguishing signs of Vasculitis in FA images from sign of other disease that cause similar injuries to the retinal structures.

My thesis work focused on a major item of the grading: retinal vascular leakage detection. The aim was to segment bright staining (hyper-fluorescence) surrounding the retinal vessels, on FA photographs. The vessels themselves also emit bright fluorescent light in FA examinations. The main challenge of my project was to differentiate pathological leakage from the vessels, from the fluorescence of other retinal structures. Related work in detection of inflammation on retinal images mostly focused on Diabetic Retinopathy [4], or global detection of hyper-fluorescence [5] in FA exams. our approach is specific to Vasculitis inflammation and grades retinal vascular leakage specifically, on FA photographs.

We devised a strategy to address the challenge of distinguishing signs of leakage from vessels. It consisted in segmenting the vasculature and automatically distinguishing the surrounding area, where staining is expected to occur, from the retinal background. We detected leakage as

pixels brighter than retinal background. A FA examination (i.e., a so-called study) consists of a time-lapse, with photographs taken at approximately regular intervals, after the dye has been injected. The examination generally lasts up to 10 minutes, with staining becoming visible only after approximately 5 minutes. In earlier frames, vessels are the only bright structure and can thus be easily segmented. Frames were registered so the vasculature detected in early frames could be aligned with vasculature detected in later frames.

In this thesis, I tested several methods to improve registration compared to off-the-shelf approaches, validating image alignment with a new specific criterion that I defined specifically for this purpose. The amount of registered frames available to segment vessels and detect vascular leakage hence increases. I then applied Computer Vision methods to build vessel and background masks, which I used to finally segment hyper-fluorescence on late frames where the masks were applied. I implemented a pipeline to run each processing step automatically, from raw machine data to segmentation of vascular leakage. Furthermore, I contributed to the preparation of a dataset to be annotated (comprising 543 patients). Finally, I designed strategies for grading of vascular leakage. I obtained encouraging preliminary results on an initial subset of images, annotated by a single grader. We expect performance to further increase once parameters will be tuned leveraging annotations on the full dataset, once they become available.

Chapter 2

Theoretical Background

2.1 Vasculitis disease

Vasculitis refers to a group of diseases characterised by inflammation of blood vessels, not only in the retina, leading to potentially long-term sequels including vision loss, aneurysm formation and kidney failure [6].

2.1.1 Variations of the disease and affected population

Incidence and prevalence are hard to determine because of the lack of reliable diagnostic criteria and the rarity and diversity of the disease [2].

Forms of Vasculitis disease vary according to age and geographical factors. Giant cell arteritis [7] occurs mainly affect persons of northern European ancestry over 50 years old, whereas Takayasu arteritis [8] occurs mainly under 40 years. Kawasaki disease [9] mainly affect children of Asian ancestry aged under 5 years old. Behçet syndrome [10] occurs most commonly in the middle east. Finally, IgA Vasculitis [11] affects children and adolescents.

2.1.2 Fluorescein Angiography

Fluorescein Angiography (FA) is an imaging technique that captures retinal vasculature. It consists on an intravenous injection of fluorescein sodium, a dye that becomes fluorescent under blue light [1]. Photographs are then taken, capturing the dye diffusion through vessels. Frame angle is the opening of the imaging device, a lower angle corresponding to a zoom on the macula (central part of the retina). Vasculitis disease is widely studied using FA [12, 13, 14]. The inflammation of blood vessels let the dye break the blood-retinal barrier and stain retinal background around vessels. This phenomenon is seen for photographs captured later (5-10 min) after the injection.

2.1.3 Associated symptoms

Vasculitis disease affects retinal vessels but may also cause Vitritis and Macular Edema [2], symptoms that are not specific to Vasculitis and that may interfere with vascular leakage detection. Macular Edema [15] is an accumulation of fluids in the extracellular space of the

neurosensory retina, causing a thickening of the retina. It can be observed in FA with bright stains caused by the inflammation. Macular Edema is a characteristic symptom of Diabetic Retinopathy [16].

Vitritis corresponds to an inflammation of cells situated in the vitreous layer of the eye, making the structure opaque and causing vision loss. Vitritis is revealed by more or less large shadows in FA photographs. It can also be an immune disease, affecting AIDS patients [17]. Vitritis is known to be strongly associated with Vasculitis, with only 4% probability of finding a patient with Vasculitis considering that this patient do not have Vitritis [18]. This rate was found by I. Tugal-Tuktun, who proposes the grading of the disease bellow.

2.1.4 Tugal-Tuktun grading

This grading addresses various aspects of the disease, with the recurrent criteria of bright stain appearance, a larger affected area inducing a higher score for the disease.

Table 1 Fluorescein angiographic scoring system		<i>Capillary leakage at 5–10 min</i>	
Angiographic sign		Posterior pole (excluding perifoveal ring of leakage)	
		Limited	1
		Diffuse	2
		For each quadrant	
<i>Optic disc hyperfluorescence at 5–10 min</i>		Foci of leakage limited in area or intensity	1
Normal staining of the scleral rim	0	Diffuse leakage	2
Staining of the disc with distinct margins		<i>Retinal capillary nonperfusion</i>	
Partial	1	Macular ischemia (enlargement of foveal avascular zone)	1
Diffuse	2	Posterior pole (excluding macular ischemia)	1
Leakage at the optic disc with blurring of margins and papillary vasculature	3	For each quadrant	1
<i>Macular edema at 10 min</i>		<i>Neovascularization of the optic disc (NVD)</i>	2
Faint hyperfluorescence	1	<i>Neovascularization elsewhere (NVE)</i>	
Incomplete ring of leakage	2	At one focus	1
Complete ring of leakage	3	Multiple	2
Pooling of dye in cystic spaces	4	<i>Pinpoint leaks</i>	
<i>Retinal vascular staining and/or leakage at 5–10 min</i>		Limited area or at one focus (≤ 3 DD)	1
Posterior pole arcades		Extensive (>3 DD)	2
Focal	1	<i>Retinal staining and/or subretinal pooling at 5–10 min</i>	
More extended or multifocal but limited area	2	Limited or at one focus (≤ 3 DD)	1
Diffuse	3	Extensive (>3 DD)	4
For each quadrant	1	DD, disc diameter	

Figure 2.1: Fluorescein angiographic Tugal-Tuktun scoring system [3]

In particular, retinal vascular leakage is distinguished between focal, multifocal and diffuse. Focal and multifocal leakage are isolated staining, intermediary stages between control FA late frame (see Figure 2.2) and diffuse staining (see Figure 2.3).

2.2 Registration

Registration between images can be done by maximising similarity via grid search directly [19], or extracting features from images with Phase-Only Correlation (POC) [20]. These methods are not suited when rotation, scaling or deformation need to be applied. Thus, medical image registration is preferably achieved by finding keypoints, matching them, and finally applying registration transformation to one image to be aligned with the other.

2.2.1 Keypoints detection

These methods can be applied to a single image to extract its main features by providing a set of keypoints with their descriptions.

Scale-Invariant Feature Transform (SIFT)

SIFT was designed to recognise features like objects [21]. It iteratively computes convolution with Gaussian function, to obtain a pyramidal structure of hierarchically smoothed images. A keypoint is detected as a local extremum at the top and is chosen comparing to less smooth images below. Gradient and orientation is computed at each level so keypoints can be robustly described by their location, scale, and orientation. Bringing adaptations and improvements to the method, SIFT have often been applied to medical data. [22, 23, 24].




Figure 2.2: Features detection with SIFT [25]

Oriented FAST and Rotated BRIEF (ORB)

ORB [26] was created as an alternative to SIFT by combining FAST detector [27] and BRIEF descriptor [28]. Binary Robust Independent Elementary Features (BRIEF) performs pixel-wise binary tests with a smoothed image patch, and is sensitive to in-plane rotation. ORB registration may reach or even slightly outperform the SIFT method [29].

2.2.2 Keypoints matching

Keypoint detectors provide descriptors that enable to cluster keypoints with a nearest-neighbours classifiers. Feature descriptors as with SIFT can be matched using euclidean distance or augmented metrics that take into account keypoint orientation. Binary descriptors obtained with ORB are matched using Hamming distance [30], evaluating the number of binary differences. Keypoints are matched minimising these distances, but may not all be relevant and disturb the last step of registration. Best keypoints can be selected by keeping a proportion that has the lowest Hamming distance (for ORB). Another strategy is to compute the distance ratio between the best and second best choice for a keypoint. This value is the Lowe's ratio [25] (design for SIFT) and 0.7 is a standard threshold to keep keypoints.

2.2.3 Homography matrix

The homography matrix represents the transformation between two 2D view of a 3D object. Let a keypoint be positioned at (x_1, y_1) and (x_2, y_2) for two views. The associated homography matrix H verifies:

$$\begin{pmatrix} x_2 \\ y_2 \\ 1 \end{pmatrix} = H \begin{pmatrix} x_1 \\ y_1 \\ 1 \end{pmatrix} \quad \text{with } H = \begin{pmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ 0 & 0 & 1 \end{pmatrix}$$

A single zoom is characterised by only $h_{00} = h_{11} \neq 0$.

For a single rotation of θ angle, $h_{00} = h_{11} = \cos\theta$ and $h_{10} = -h_{01} = \sin\theta$.

A single affine displacement will have only $h_{00} \neq 0$ and $h_{11} \neq 0$.

With 6 parameters in H , 3 pairs of positions allow to solve this equation. Random Sample Consensus (RANSAC) algorithm [31] takes randomly 4 pair of keypoints, solves the equation with 3 of them, and look if the transformation matches the last pair of positions (with a tolerance distance). Repeating this operation, RANSAC computes the homography matrix that is consistent with the most of keypoints (if enough matches where found).

2.2.4 Image similarity measures

Image similarity measures are used to compare images alignment and access the quality of registration

Mean Squared Error (MSE)

The mean of squared differences between pixels requires both images to have similar luminosity

$$MSE(x, y) = \frac{1}{N} \sum_{i=1}^N (x_i - y_i)^2 \quad \text{with } x_i, y_i \text{ the pixel values.}$$

Structural Similarity Index Measure (SSIM)

SSIM [32] measures the similarity between two images, aiming to mimic human perception combining luminance, contrast and structure:

Luminance: $l(x, y) = \frac{2\mu_x\mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1}$, where $\mu_x = \frac{1}{N} \sum_{i=1}^N x_i$

Contrast: $c(x, y) = \frac{2\sigma_x\sigma_y + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}$, where $\sigma_x^2 = \frac{1}{N-1} \sum_{i=1}^N (x_i - \mu_x)^2$

Structure: $s(x, y) = \frac{\sigma_{xy} + c_3}{\sigma_x\sigma_y + c_3}$, where $\sigma_{xy} = \frac{1}{N-1} \sum_{i=1}^N (x_i - \mu_x)(y_i - \mu_y)$

SSIM(x,y) = $l(x, y) \cdot c(x, y) \cdot s(x, y)$.

Normalized Mutual Information (NMI)

The MI [33] and NMI [34] measure the degree of dependence of variables based on their distribution. They have been proven to be good similarity measures for medical images registration [35, 36].

$$MI(x, y) = H(x) + H(y) - H(x, y), \quad NMI(x, y) = \frac{H(x) + H(y)}{H(x, y)}$$

$$\text{where } H(x) = - \sum_{i=1}^N x_i \log(x_i) \text{ is the entropy}$$

2.3 Computer vision methods

2.3.1 Sato tubeness filter

Sato tubeness filter [37] was developed for 3D medical images to detect curvilinear structures such as vessels. The method can be adapted to 2D medical images [38] and is based on the Hessian matrix:

$$\nabla^2 I(\mathbf{x}) = \begin{pmatrix} \frac{\partial^2}{\partial x^2} I(\mathbf{x}) & \frac{\partial^2}{\partial xy} I(\mathbf{x}) \\ \frac{\partial^2}{\partial yx} I(\mathbf{x}) & \frac{\partial^2}{\partial y^2} I(\mathbf{x}) \end{pmatrix} \quad \text{where } I(\mathbf{x}) \text{ is the pixel value at } (x, y)$$

An ideal bright light parallel to x axis at x_0 position is represented by:

$$I(x, y) = \frac{e^{-(x-x_0)^2}}{2\sigma^2}$$

The eigenvalues $\lambda_1(\mathbf{x}) < \lambda_2(\mathbf{x})$ of the Hessian matrix are computed. The condition $\lambda_1(\mathbf{x}) \ll \lambda_2(\mathbf{x}) \approx 0$ suggests line structure, based on this ideal example, where the constant σ controls the scale.

2.3.2 Otsu's algorithm

Otsu's algorithm [39] consists on finding a threshold to binarize a gray scale image into class 1 and 2, minimising the weighted within-class variance:

$$\sigma_{weighted}^2(t) = w_1(t)\sigma_1^2(t) + w_2(t)\sigma_2^2(t) \quad w_1(t), w_2(t) \text{ corresponding to class separation probabilities}$$

The threshold found t verifies this equation [40]:

$$t = \frac{m_1(t) + m_2(t)}{2} \quad \text{where } m_1(t), m_2(t) \text{ are the mean values in each class}$$

The method can be generalised to more classes: multi-Otsu thresholding [41].

2.3.3 Morphological transformations: opening

Morphological transformations are simple operations performed on binary images. Basic transformations are erosion and dilation, for which a kernel slides through the image and set a pixel to 0/1 if a 0/1 pixel is found under the kernel for erosion/dilatation respectively. An erosion followed by a dilation is called an opening, and removes isolated group of black pixels while preserving larger shapes.



Figure 2.3: opening: erosion followed by dilation (*OpenCV* documentation)

Binary area opening [42] refers to algebraic opening and removes group of pixels with an area smaller than a parameter λ .




Figure 2.4: Binary area opening [42]

Chapter 3

Design

3.1 Dataset

3.1.1 Characteristics

Dataset Fluorescein Angiography (FA) was extracted from Heyex database in the Hopital Ophtalmique Jules-Gonin (HOJG), using an automated interface. Patients treated in the Uveite section of HOJG were selected. Patient data often contain several visits and last examinations may be done after the patient is cured. Patients may also not have Vasculitis disease and Vasculitis rarely affects both eyes, so the dataset contains enough control cases (distribution to be determined from annotations to come).

Dataset is composed of 543 patients (295 females, 248 males), born from 1919 to 2012 (1968 ± 19 years on average,), that were examined from 2014 to 2021. Patients came to the hospital from 1 to 16 times (3 ± 2.5 visits on average), with imaging of both eyes. A visit contains 12.8 FA on average (± 5.2 frames), and may have a video of the early times after the injection for one eye. FA are mostly 55° angle images, specially for early frames, and may be 35° zoomed frames or 102° larger frames, some of them for peripheral views. Finally most images have 768x768 pixel resolution, and few others were cut to match this square shape that is kept through every processing step.

3.1.2 Dataset structure

FA saved in dicom format were selected from Heyex dataset, that contains other images modalities like Optical Coherence Tomography (OCT), and were anonymized as follow:

- Patient ID (PID) from the hospital is replaced by a new ID, with a conversion table
- Date of birth is set to January 1st, year of birth is kept
- Name and surname are removed

A study is defined by a patient visit for one eye, and is named concatenating 'Vasculitis', new PID, visit date and eye side (ex. 'Vasculitis-0017-20170429-R'). Dataset was splitted into 3303 studies, for which frames were extracted from Dicom files, and saved as JPEG images, keeping acquisition time in seconds and frame angle in the JPEG title (ex. 64-55°).

FA video were used as additional early frames. They capture ink diffusion and show vessels being gradually filled by the ink. We chose to extract 5 frames regularly out of the last half of the video, so it helped to build masks, and avoided to treat almost identical pictures of the retina.




Figure 3.1: Anonymised dataset structure

3.2 Registration

The aim was to register the maximum number of frame in a study, starting from early frames as a model, and preventing misalignments to not disturb following processing steps.

Grid search based methods were not suited because of rotations and scale variations between frames. The strategy was to find and rely on a robust quality checker so we can try any combination of the following methods to register a new frame until we succeed.

Registration was achieved with *OpenCV* Python libraries.

3.2.1 Standard registration

Registration cannot start with a bad quality frame or with the wrong angle, so only 55° frames without large black areas were automatically selected to be a reference for registration.

The standard method was defined by using Scale-Invariant Feature Transform (SIFT) key-points detector, and selecting best matches according to 0.7 Lowe's ratio, trying to register every frames to the earliest.

3.2.2 Frame reference selection and back-registration

Starting with the first good enough frames, new frames are registered to every 55° registered frame until success. This strategy led to some new registered frames as the images are progressively evolving with the dye diffusion.

Registering a new frame means transforming it to be aligned with the reference. The opposite operation was done, "back-registration", registering a frame that have already been treated, using the new one as reference. Keypoints sets were switched after matching, so the transformation was applied to the new frame.

3.2.3 Keypoints matching

The first step was to compute and save keypoints for all frames using SIFT and ORB methods.

Oriented FAST and Rotated BRIEF (ORB) registration was tested using Hamming distance to match keypoints. Scale-Invariant Feature Transform (SIFT) registration was done using Fast Library for Approximate Nearest Neighbors (FLANN). Keypoints were matched with k-Nearest-Neighbours (kNN) algorithm

3.2.4 Keypoints matches selection

The second step consisted on selecting the best keypoints to achieve computing the homography matrix for registration.

ORB best matches were first selected being the proportion with the lowest Hamming distance, as it is commonly done, but this selection led to little misalignments compare to SIFT. No matter the proportion kept and the metric used, the ORB registration performances were worse than for SIFT.

SIFT best matches were selected according to Lowe's ratio, trying values from 0.2 to 0.9. Most frames were better registered for ratio in the range 0.6 – 0.7, but in some cases, trying other ratios led to get new frames registered. Testing different ratios is not much time consuming since keypoints are saved before trying to compute the homography matrix.

Modifying the registration template, matches selection with Lowe's ratio was tested for ORB. This improved the registration success, placing ORB close or even better than SIFT for some frames, using image similarity metrics. However these improvements were rare and limited. It showed the limits of image similarity metrics, creating some examples where the score given by the metric is not conform with the observed alignment.

Moreover, ORB method was occasionally producing shifted images, specially for frames that are difficult to register. On the other side, SIFT registration had the positive behaviour to often stop computing the homography matrix, instead of returning misaligned frame.

ORB registration was not integrating to the final pipeline, with almost no improvement comparing to SIFT, and being too much challenging for the quality assessment.

3.2.5 Quality assessment

Image similarity metrics (Mean Squared Error (MSE), Structural Similarity Index Measure (SSIM), Mutual Information (MI)) were tested, allowing to set thresholds that detect most of failed SIFT registration, or ensure a good alignment for most selected results, but not both. MI performed the best to evaluate registration, but it was still feasible to better assess registration quality, and thus to get more registered frames avoiding misalignments. Moreover, similarity metrics were not always able to rank similar tries with different methods.

SIFT and ORB methods are designed to compute registration between any 2D view of a 3D object, but we needed only to register parallel views of the retina in our dataset. Thus, the computed homography matrix H may represent an undesired transformation, characterised by different values for diagonal coefficients h_{00} and h_{11} .

$$H = \begin{pmatrix} h_{00} & h_{01} & h_{02} \\ h_{10} & h_{11} & h_{12} \\ 0 & 0 & 1 \end{pmatrix}$$

Moreover, too large or too small values for diagonal coefficients correspond to a bigger scale change than possible with our data. I created the "parallel view" criterion associated to a tolerance ration r_{tol} :

$$0.5 < h_{00} < 2 \quad \text{and} \quad 1 - r_{tol} < \frac{h_{00}}{h_{11}} < 1 + r_{tol}$$

The parallel view criterion was selected with $r_{tol} = 2\%$ to assess registration quality on its own, there is no similarity measure computation in the final pipeline.

3.3 Vascular leakage detection

From a registered study, we selected early frames (acquisition time $< 2\text{min}$) and late frames (acquisition time $> 5\text{min}$). Vascular leakage is characterised by hyper-fluorescence around the vessels and only appears in the late frames. Vessels are the only bright area in early frames so the strategy was to detect them with the background area around, where vascular leakage may occur in late frames. Building vessel and background masks specific to the study, and applying them to late frames, vascular leakage was detected as fluorescence around vessels.

3.3.1 Generic mask creation

I designed a method to create both masks from early frames using hyper-parameters that regulate mask characteristics:

- σ : range used as scales of filter, controlling minimal detected vessel thickness

- p_{max} : maximal intensity percentage for contrast stretching, controlling computed vessel thickness

Masks were creating for each of the early frames and merged into final mask, adding vessels for a more complete detection. Mask creation for one early frames followed these steps in this order:

- Vessels detection, applying Sato filtering with *skimage.filter* Python library, filtering for every scale σ (higher scales implying more rejection of small vessels)
- Contrast stretching, computing intensity thresholds $t_2, t_{p_{max}}$ representing 2% and p_{max} % of frame pixels bellow this intensity threshold, and rescaling luminosity so pixels with $t_2, t_{p_{max}}$ values become respectively 0, 1
- Binarization, using Otsu's algorithm with *OpenCV* Python library
- Removal of the 10 pixels thick area at the outline of the frame, that may be detected as an edge and selected in the mask
- Removal of small vessel traces or artefacts based on binary area openings, using *skimage.filter* Python library




Figure 3.2: Example of computed vessel and background masks from the same early frames

3.3.2 Vessel mask

Vessel masks were computed with the method above using only $\sigma = 1$ to have better vessel detection by not filtering small vessels. The downside was that it left small vessel traces or detected some isolated group of point as artefact, caused by high luminosity or anatomical confounding factor. The morphological operation of removing these small areas helped getting a

clean vessel mask, while detecting more vessels using contrast stretching and small σ .

Vessel masks should segment precisely vessels so little contrast stretching was applied using $p_{max} = 98\%$, helping to detect more vessels, without overflowing around.

3.3.3 Background mask

Background masks were computed with the method above using $\sigma = (1, 10)$ to filter smallest vessels and avoid filling most of eye background. These vessels are smaller branch of main ones, so the area around is still matching with the detected background.

Background masks should be an expansion of vessel masks so more contrast stretching was applied using $p_{max} = 85\%$. This contrast transformation is not symmetric and brings vessel luminosity closer to the value, so both zones are matched by Otsu's binarization. The value of p_{max} is quite sensitive because the luminosity gradient at vessel's edges is strong even if luminosity is continuous.

An other strategy was tried to get background masks directly from the vessel masks, applying erosion then dilatation with morphological operator from *OpenCV* Python library and using a large kernel to expand detected area. The best kernel found was 5x5 pixels but this method was not better than applying strong contrast stretching, but had the defect of expanding every area the same way and creating occasionally weird shapes.

3.3.4 Optic disk detection

In order to focus on vascular leakage, the optic disk, that may be quite bright when leaking, needed to be hidden.

53 optic disks were manually annotated by doctor Florence Hoogewood on 55° angle frames. The coordinates of the optic disk centre was computed as the mean coordinates of annotated pixels (applying mirror transformation to right eyes). Then the radius of the optic disk was defined by the further annotated pixel from the centre.

These statistics for optic disk centre positions (x, y) and radius (r) were selected removing 4 outliers. The following values are in pixels, computed from 55° angle frames with 768x768 resolution:

- $x = 172 \pm 17$, (146-217) range
- $y = 362 \pm 14$, (340-396) range
- $r = 48 \pm 5$, (34-61) range

The generic optic disk was set to these measures.

3.3.5 Hyper-fluorescence detection

Relying on the registration, vessel, background, and optic disk masks were merged into the "combined" mask, which was applied to the late frames to select the area around vessels. This last step needs to be completed and tuned with annotations to come.

Hyper-fluorescence was then detected selecting pixels brighter than a threshold, an hyper-parameter of the pipeline. This threshold can be chosen using Otsu's binarization or multi-Otsu's thresholding, methods that cluster the area around vessels into two or more classes, minimising within-class variance. The image luminosity was standardised.




Figure 3.3: Combined mask with the grid to characterise vascular leakage

Finally the late frames are divided based on a squared grid, and vascular leakage was defined for hyper-fluorescence areas greater than a threshold to tune. The grading between focal, multifocal and diffuse staining was determined by the occurrence of vascular leakage in the whole frame.

3.4 Expert's grading

3.4.1 Format processing

Grading was done on Discovery Platform (tool provided by RetinAI), on Swiss Ophthalmic Imaging Network (SOIN), a secure space created in HOJG to safely share medical data.

Frames belonging to a study have been grouped into mosaic collage, with timestamps for each frame, so the grader can easily see the whole study and attach one grading form to one image in Discovery. The issue with this manipulation was that Discovery platform needs specific data

format to then extract patient metadata. We chose to upload the collage as fake Optical Coherence Tomography (OCT) images to benefit later from a tool provided by the platform, that would randomise data smartly into graders to be more robust. The format adaptation involved these steps:

- Write timestamps on the top of each frame to let the grader see easily acquisition time.
- Create one collage per study and save it as an JPEG image.
- Create Discovery specific storage files of OCT (h5 files) for each collage, generating in parallel ID to the platform.
- Fill patient and study metadata and check if h5 files content matches with collages.
- Upload h5 files to SOIN secure environment and finally to Discovery platform.

3.4.2 Grading form

The aim of the whole project is to be able to automatically grade Vasculitis based on Tugal-Tuktun table, that contains additional aspects of the disease. This is why we created a grading form to address most criterion of Tugal-Tuktun table in order to get annotations for vascular leakage detection, but also for further investigations. The thesis focused on answering 5th item, aiming to grade retinal vascular leakage between focal, multifocal, and diffuse.

1: image quality

good	medium (grading still possible)	poor (grading not doable)
------	---------------------------------	---------------------------

2: Confounding factors
Select all that apply:

atrophic lesions	vitritis or vitreous opacity	blood
other		

☐ Skip

3: Optic disc hyperfluorescence at 5–10 min

none	partial Staining of the disc with distinct margins	diffuse Staining of the disc with distinct margins	Leakage at the optic disc with blurring of margins and papillary vasculature	not assessable

4: Macular edema at 10 min

Select...

5: Retinal vascular staining and/or leakage at 5–10 min: Posterior pole arcades

none	focal	More extended or multifocal but limited area	Diffuse	not assessable
------	-------	--	---------	----------------

6: Retinal vascular staining and/or leakage at 5–10 min: periphery. Number of quadrants involved

0

7: Capillary leakage at 5–10 min: Posterior pole (excluding perifoveal ring of leakage)

none	limited	diffuse	not assessable
------	---------	---------	----------------

8: Capillary leakage at 5–10 min: for each quadrant: Foci of leakage limited in area or intensity = 1 ; Diffuse leakage = 2 (if not assessable, skip)
Select a value between 0 and 8. It must be a multiple of 1.

0

☐ Skip

9: Retinal capillary nonperfusion: posterior pole

none	Macular ischemia (enlargement of foveal avascular zone)	Posterior pole (excluding macular ischemia)	not assessable
------	---	---	----------------

10: Retinal capillary nonperfusion: periphery: number of quadrant involved

0

11: retinal neovascularization

none	of the disc (NVD)	elsewhere (NVE)	not assessable
------	-------------------	-----------------	----------------

12: Pinpoint leaks

none	Limited area or at one focus ($\leq 3DD$)	Extensive ($> 3DD$)	not assessable
------	---	-----------------------	----------------

13: Retinal staining and/or subretinal pooling at 5–10 min

none	limited or at one focus ($\leq 3DD$)	extensive ($> 3DD$)	not assessable
------	--	-----------------------	----------------

14: comments

Enter text...

☐ Skip

Figure 3.4: Grading form in Discovery platform

3.4.3 Hyper-parameters tuning

Based on annotation results from the 5th item of the grading form (Retinal vascular staining and/or leakage at 5-10min), dataset will be split into test and train set to tune hyper-parameters of the pipeline via grid search:

- r_{tol} (registration quality tolerance), x, y, r (optic disk characteristics), and σ, p_{max} (characteristics of both masks): these parameters for intermediate steps were already set, giving visual satisfaction from Vasculitis expert.
- Threshold to detect hyper-fluorescence (or Otsu's thresholding computation)

- Minimal area of detected hyper-fluorescence to characterise vascular leakage.

3.5 Pipeline

The pipeline is a Python script that apply iteratively every methods described bellow, saving for every steps JPEG images according to the designed sub-folder organisation. The pipeline can be run starting from any of these steps, ending by vascular leakage detection, but also preparing data for annotations. Input can be all dataset or any described subset.




Figure 3.5: Pipeline steps for vascular leakage detection

Chapter 4

Results

4.1 I implemented a method to register Fluorescein Angiography photographs, with higher performances compared to off-the-shelf approaches, and precision that satisfied human experts

I designed and implemented a registration method that aligns features, such as the vasculature and the optic disk, that belonged to different frames. Vasculitis expert Dre. Hoogewoud reviewed the registration results and was satisfied with its performances. This validates the SIFT method with the registration quality assessment process, namely the parallel view criterion that I created.

The process consisted of trying variations of SIFT to register a new frame (see Design 3.2):

- Use different Lowe's ratios
- Try every frame that was already registered as a reference
- Apply normal and back-registration
- All three variations combined

All tested variations of SIFT led separately to higher performance. I obtained even better results when I combined them, those methods being complementary

Thanks to the robust parallel view criterion (see Design 3.2.5), registration tries could be multiplied with no risk of alignment errors. Hence, registration success increased, especially for 102° and 35° frames. The dataset has the following distribution of angles: 13% had 35° angles; 38% had 55° angles; finally, 49% had 102°. Registration performance improved the most for 35° and 102° frames, noting that the first frame has always 55° angle.

I measured registration success by looking at the ratio between successfully registered frames and available ones. The percentage of successfully registered 102° frames increased from 15.8% to 47.8%, and overall frame registration success reached 66.8%, compared to 46.5% achieved with standard SIFT (see Design 3.2.1).

Frame Angle	Standard SIFT	+ Lowe's ratio modification	+ registration to every frames	+ back registration	All methods
35°	56.3	59.8	61.5	58.1	65.9
55°	83.2	87.3	88.1	85.3	91.8
102°	15.8	24.9	25.1	20.2	47.8
Total	46.5	53.0	53.6	49.7	66.8

Figure 4.1: Percentage of registered frames for different angles

I also measured registration success by looking at how many frames were available to be included as "early" and "late" frames in a study (i.e., a single patient visit). There are overall less early frames but they are mostly registered successfully, no matter the method variation used. Late frames were characterised by the lowest registration success rate, but also best improvements: despite only 26% percent of them being successfully registered, the average number of registered late frames per study nearly doubled, from 1.44 to 2.86.

Vascular leakage can only be detected on late frames, it is thus essential to register them. The percentage of studies in which no late frame was available after registration, decreased from 22.2% to 9.5%.

Frames	Before registration	Standard SIFT	+ Lowe's ratio modification	+ registration to every frames	+ back registration	All methods
Early	2.88	2.62	2.66	2.65	2.64	2.69
Medium	3.72	1.58	1.91	1.92	1.74	2.56
Late	5.53	1.44	1.86	1.94	1.66	2.86
Total	12.13	5.64	6.44	6.51	6.03	8.10

Figure 4.2: Averages number of registered frames (in a single study) across different acquisition times

4.2 I implemented a method to segment the vasculature in Fluorescein Angiography photographs, using Computer Vision

The results of our vasculature segmentation method were validated, via visual inspection, by Vasculitis expert Dr. Hoogewoud who was satisfied with the detection of vessels (both those of large and small diameter). Similarly, she validate the segmentation of the background , the area where vascular leakage may occur.




Figure 4.3: Top left: vessel mask computed with all processing steps; Top right: using only the first frame; Bottom left: without removing frame outline; Bottom right: without removing traces

The figure above shows on the top left a vessel mask computed with my method (see Design 3.3.1) with three kind of artifacts that were treated and removed by my method (which correspond to the three processing steps I applied in addition to standard Sato filtering and Otsu's binarization, see Design 3.3.1):

- The mask on the top right is incomplete, thus merging masks computed from each available early frame improved vessel detection. The stain close to the optic disk is an inflamed area that needed to be included in the mask to avoid its being erroneously detected as vascular leakage in later steps of our pipeline.
- The mask on the bottom left contains artefacts caused by the edges of each early frame that was combined to generate it.
- The mask at the bottom right is polluted by isolated traces, that were removed via morphological operators. Left untreated, these traces may lead to expand too much the background mask associated with the corresponding study.

Macular edema (i.e., leakage of blood into the central part of the retina) is a symptom associated with Vasculitis, and appears as a bright structure in early frames. Our pipeline incorporates the edema in the vessel mask. Despite this not being biologically correct, such behaviour is desirable, as it prevents it from being misclassified as vascular leakage in late frames, by the subsequent steps of our pipeline.

4.3 I implemented a complete pipeline for the grading of vascular leakage in Fluorescein Angiography (FA) photographs

Our pipeline is implemented as a single Python script that takes raw data in input and segments vascular leakage. The output of each step of the pipeline is stored in a corresponding image sub-folder, so the users can visualise the raw data in input, the output of the registration step and the resulting masks. Once it has been run a first time, the pipeline has been designed to be able to restart the processing from any step, by recomputing data in the image sub-folders corresponding to subsequent steps.




Figure 4.4: Unified Modeling Language (UML) of our pipeline

The script manipulates "frames", custom-made Python objects created to capture and easily access FA characteristics. This design choice makes our pipeline modular: additional processing step could be added at a future stage by adding a new function to the pipeline, using pre-built loading/saving functions that operate on frame objects.

I was faced with external issues in the upload of data to our annotation platform (i.e. the Discovery tool by RetinAI). I implemented a solution by creating a script that import the dataset from JPEG images. My script organised the data in Discovery in such a way that annotations can be performed by experts both inside our hospital, and in remote locations. This process is not specific, and will be re-used in other projects at Hopital Ophtalmique Jules-Gonin.

4.4 I defined a strategy to enable Fluorescein Angiography photographs grading and detect vascular leakage

The last step of our pipeline consists on detecting hyper-fluorescence from late frames and classifying the type of vascular leakage. The method needs expert annotations to be tuned and validated. At this early stage of the research project, vascular leakage segmentation was performed with a fixed luminosity threshold of 0.6. The figure bellow displays a late frame showing diffuse vascular leakage (left). Coloured filters were added (on the right): vessels are marked in green, hidden retinal background in purple, and vascular leakage in red. Segmentation is far from perfect, but a sufficient amount of staining was detected so Vasculitis was correctly identified, and the leakage was graded as diffuse (see Theoretical background 2.1.4), despite shadows on the bottom of the frame (Vitritis)

I obtained encouraging preliminary results. Shortly before the submission of the present manuscript, 60 images were annotated. Those comprised photographs of both the eyes of 30

patients, from their most recent visit. I tried first to distinguish diffuse leakage from other cases (focal, multifocal or control). Performance were encouraging, for a specific subset of the data: indeed, diffuse leakage was correctly detected if a late frame satisfied the following criterion: an area of hyper-fluorescence bigger than 100x100 pixels applying a fixed 0.6 luminosity threshold. Results were as follows:

- 4/6 cases (67%) of diffuse staining were correctly identified.
- 6/54 other cases (11%) were wrongly identified (namely: 1 multifocal, 1 focal, 2 with major confounding factor disease, 2 control)
- 82% accuracy with 40% precision

Chapter 5

Discussion

5.1 Interpretation of registration process

Registration performances were improved by using three variations of SIFT method, validated by the designed parallel view criterion.

ORB registration is known to be sensitive to in-plane rotation [26], a transformation often needed to register our frames. This may explain the lower performances compared to SIFT

Lowe's ratio is used to filter keypoints matches, a ratio close to 1, meaning that the distance between the best and 2nd best possibility for a match is close, so the best match does not stand out and has few chances to be correct. I discovered that the best ratio varies with images : there are few cases where registration failed with the standard 0.7 Lowe's ratio, but succeeded with a lower or higher one. The fact that a lower one works better may be caused by too many keypoints found, with matching errors being more filtered. A higher ratio may allow to keep more keypoints, which are still filtered by the SIFT keypoints selection (RANSAC).

Registration was tried on every frame that has already been registered, with improvements compare to registering only to the first frame. This means that registration success is not transitive : some frames need an intermediate images to be registered to the first one. This could be explained by the fact that frames are order by time and represent the evolution of the retina. We may think that selecting the previous registered frame as a model is better than with the first one, but it is not the case on average, because luminosity and contrast are often better for early frames. The injected dye slowly disappears over time, thus late frames are harder to be registered.

Back-registration means registering the first frame with the new frame a reference (and applying then opposite transformation). I found surprisingly that it may be complementary to normal registration, meaning that registration success is not always transitive. Its relative contribution is lower than other SIFT variations described bellow but back-registration still helped registering more frames. I honestly cannot give more interpretation on this.

The parallel view criterion can be interpreted as follow: Frames views are all perpendicular

to retina plan and SIFT registration can register any 2D view of an object, so there are very few chances that the registered frame has the exact same direction as it should, without being aligned to the reference. This would suppose that keypoints were mostly misplaced in the same way to pass the SIFT keypoints selection (RANSAC).

5.2 Limitation of current work

The aim of the thesis was to detect vascular leakage in late frames of Fluorescein Angiography (FA) studies, and registering frames was a major step to achieve it. A lot of efforts was put into improving registration, but all aspects may not be useful to help hyper-fluorescence detection. Thus, frames that are neither early nor late (31% of total) were not used to detect hyper-fluorescence. The best progression in registration performance concerned 102° angle frame, with even more improvements for peripheral views from which it should be harder to detect hyper-fluorescence. We faced limitations with images similarity measures to evaluate registration, not being able to rank or evaluate precisely registration methods. Alignment improvements are thus not quantifiable for frames that were already registered with standard SIFT method. Finally, the parallel view criterion was an efficient solution that replaced similarity measures, but did not provide quantitative registration quality metric.

The other main aspect of this thesis was the use of computer vision techniques to build masks and detect hyper-fluorescence. Here as well, mask creation is an intermediate step and efforts to get a clean mask by applying morphological operations may not all be so necessary to finally grade between vascular leakage categories. Results from first annotations led to 6/54 false positive, 2 of them being directly caused by a confounding inflammation symptom that was not segmented correctly in early frames. Vasculature is overall well segmented but vascular leakage detection seems to be sensitive to other inflammation symptoms, which are often associated to Vasculitis.

The main limitation of the project was the delay of annotations, preventing to get robust quantitative results for vascular leakage detection. With the first annotations and looking at Vasculitis examples, the method was still promising and we defined a strategy to tune hyper-parameters. Setting up data into Discovery platform was a big issue, various strategy failed to upload data properly, which delayed the project. The solution found, encapsulating collages into OCT, worked but seems obviously much harder than it should.

5.3 Future directions

The first next direction will be to get annotations and apply the defined strategy to evaluate and improve the current method, specially the last step, the hyper-fluorescence detection from late frames and computed masks, leading to vascular leakage grading.

The annotation process is ready to start, and will be carried by Vasculitis expert Dre. Hoogewoud at Hopital Ophtalmique Jules-Gonin. She will be helped by interns from Luzern Kantonsspital and Dr. Chiquet from CHU Grenoble. These hospitals also have Vasculitis patients and

data privacy agreements are in process to be able to share data. Our dataset has been transferred to the secure SOIN platform, to which some annotators have already been on-boarded and can start the grading.

The current optic disk mask creation is quite basic and not precise, but this step has been assigned to Oscar Jimenez (IDIAP) who is trying to train neural networks to detect precisely the optic disk. The strategy is to start from a model already working on coloured retinal fundus images [43] and fine-tune it with the initial 53 optic disks annotated by Dr. Hoogewoud.

Registration of peripheral frames between 2-5 min of acquisition time, that are currently not used for vascular leakage detection, may be useful to be able to detect early vascular leakage, which is a sign of the severity of the disease.

It might also be interesting to try and adapt other methods applied on similar work, as hyper-fluorescence detection from Ultra-Wide Fluorescein Angiography for diabetic retinopathy [4] or posterior uveitis [44]. The most recent study [5] is not directly linked to our grading but is able to detect inflammation in FA photographs. Their method consists on generating normal images from Vasculitis cases via a Generative Adversarial Network (GAN), computing leakage as the difference between the initial image and the generated one. Their model could be adapted, or the process could be done training the GAN with future annotations.

Finally, Vasculitis can be alternatively observed via Optical Coherence Tomography (OCT), 3D photograph of the retina, from which Macular Edema can be seen as thickening of the retinal layer in the macula [1]. Patients in our dataset may also have OCT, which can be directly uploaded to Discovery platform, that computes layer thickness and detect fluid accumulation, a characteristic of the disease. This would help to detect Macular Edema but would move a little away from the global project that aims to grade Vasculitis from FA images only.

Chapter 6

Conclusion

This MSc thesis project led me to implement a method to better register Fluorescein Angiography (FA) photographs, validating image alignment according to novel specific criteria. I applied Computer Vision methods to build vessel and background masks from early FA frames. I leveraged those to segment retinal vascular leakage from late frames. The quality of registration and of the masks was reviewed by a Vasculitis expert. I designed and implemented a modular pipeline that analyses raw data and outputs vascular leakage segmentation (detected as hyper-fluorescence), creating a collage of frames ready for expert grader annotated. Finally, I applied my pipeline to grade vascular leakage in a subset of data for which annotations are available, demonstrating positive preliminary results.

This project was started a previous intern in the lab, Pauline Eyraud, who I would like to thank for introducing me the basic concepts, and for explaining her work. She explored various methods and devised a prototype method to generate vessel and background masks which served as the basis for my work. Being able to try and evaluate her ideas and scripts, helped me in writing the pipeline I presented in this thesis. She applied standard, off-the-shelf SIFT registration method, resulting in almost 1/2 of frame being registered. The main achievements of my thesis are: the implementation of a specific registration method that reaches 2/3 success rate, using variations of SIFT; the implementation of a method to segment vasculature in FA photographs; the design and implementation of a first version of an automatic pipeline able to grade severity of vascular leakage in FA images, with encouraging preliminary results. Future work will evaluate and adapt my vascular leakage detection method based on annotations that will become available and extending the grading pipeline to cover the rest of the inflammation signs that are detailed in the Tugal-Tutkun method.

I would like to thank Dr. Jean-Marc Odobez (IDIAP/EPFL) for supervising the thesis, Dre. Florence Hoogewoud (HOJG Vasculitis expert) who helped me to validate methods and design annotation process, and Dr. André Anjos (IDIAP) for his advices. Finally, I would like to thank Dr. Mattia Tomasoni (HOJG) for supervising my work, during these last months, with a lot of interest and consideration.

Bibliography

- [1] R. F. Spaide, J. M. Klancnik, and M. J. Cooney, "Retinal vascular layers imaged by fluorescein angiography and optical coherence tomography angiography," *JAMA ophthalmology*, vol. 133, no. 1, pp. 45–50, 2015.
- [2] R. A. Watts, G. Hatemi, J. C. Burns, and A. J. Mohammad, "Global epidemiology of vasculitis," *Nature Reviews Rheumatology*, vol. 18, no. 1, pp. 22–34, 2022.
- [3] I. Tugal-Tutkun, C. P. Herbort, M. Khairallah, and A. S. for Uveitis Working Group (ASUWOG), "Scoring of dual fluorescein and icg inflammatory angiographic signs for the grading of posterior segment inflammation (dual fluorescein and icg angiographic scoring system for uveitis)," *International ophthalmology*, vol. 30, pp. 539–552, 2010.
- [4] J. P. Ehlers, K. Wang, A. Vasanji, M. Hu, and S. K. Srivastava, "Automated quantitative characterisation of retinal vascular leakage and microaneurysms in ultra-widefield fluorescein angiography," *British Journal of Ophthalmology*, vol. 101, no. 6, pp. 696–699, 2017.
- [5] W. Li, W. Fang, J. Wang, Y. He, G. Deng, H. Ye, Z. Hou, Y. Chen, C. Jiang, and G. Shi, "A weakly supervised deep learning approach for leakage detection in fluorescein angiography images," *Translational vision science & technology*, vol. 11, no. 3, pp. 9–9, 2022.
- [6] D. G. Scott and R. A. Watts, "Epidemiology and clinical features of systemic vasculitis," *Clinical and experimental nephrology*, vol. 17, no. 5, pp. 607–610, 2013.
- [7] A. Sharma, A. J. Mohammad, and C. Turesson, "Incidence and prevalence of giant cell arteritis and polymyalgia rheumatica: a systematic literature review," in *Seminars in Arthritis and Rheumatism*, vol. 50, no. 5. Elsevier, 2020, pp. 1040–1048.
- [8] W. P. Arend, B. A. Michel, D. A. Bloch, G. G. Hunder, L. H. Calabrese, S. M. Edworthy, A. S. Fauci, R. Y. Leavitt, J. Lie, R. W. Lightfoot Jr *et al.*, "The american college of rheumatology 1990 criteria for the classification of takayasu arteritis," *Arthritis & Rheumatism*, vol. 33, no. 8, pp. 1129–1134, 1990.
- [9] S. M. Skochko, S. Jain, X. Sun, N. Sivilay, J. T. Kanegaye, J. Pancheri, C. Shimizu, R. Sheets, A. H. Tremoulet, and J. C. Burns, "Kawasaki disease outcomes and response to therapy in a multiethnic community: a 10-year experience," *The Journal of pediatrics*, vol. 203, pp. 408–415, 2018.
- [10] F. Shahram, M. T. Mæhlen, M. Akhlaghi, F. Davatchi, Y. J. Liao, and C. M. Weyand, "Geographical variations in ocular and extra-ocular manifestations in behçet's disease," *European Journal of Rheumatology*, vol. 6, no. 4, p. 199, 2019.
- [11] M. Píram, C. Maldini, S. Biscardi, N. De Suremain, C. Orzechowski, E. Georget, D. Regnard, I. Koné-Paut, and A. Mahr, "Incidence of iga vasculitis in children estimated by four-source capture–recapture analysis: a population-based study," *Rheumatology*, vol. 56, no. 8, pp. 1358–1366, 2017.
- [12] A. M. Abu El-Asrar, C. P. Herbort, and K. F. Tabbara, "A clinical approach to the diagnosis of retinal vasculitis," *International ophthalmology*, vol. 30, no. 2, pp. 149–173, 2010.
- [13] A. M. A. El-Asrar, C. P. Herbort, and K. F. Tabbara, "Differential diagnosis of retinal vasculitis," *Middle East African journal of ophthalmology*, vol. 16, no. 4, p. 202, 2009.

- [14] C. D. Hart, M. Sanders, and S. Miller, "Benign retinal vasculitis. clinical and fluorescein angiographic study," *The British Journal of Ophthalmology*, vol. 55, no. 11, p. 721, 1971.
- [15] M. W. Johnson, "Etiology and treatment of macular edema," *American journal of ophthalmology*, vol. 147, no. 1, pp. 11–21, 2009.
- [16] N. Bhagat, R. A. Grigorian, A. Tutela, and M. A. Zarbin, "Diabetic macular edema: pathogenesis and treatment," *Survey of ophthalmology*, vol. 54, no. 1, pp. 1–32, 2009.
- [17] H. W. Henderson and S. M. Mitchell, "Treatment of immune recovery vitritis with local steroids," *British journal of ophthalmology*, vol. 83, no. 5, pp. 540–545, 1999.
- [18] I. Tugal-Tutkun, S. Onal, M. Stanford, M. Akman, J. W. Twisk, M. Boers, M. Oray, P. Özdal, S. Kadayifcilar, R. Amer *et al.*, "An algorithm for the diagnosis of behçet disease uveitis in adults," *Ocular Immunology and Inflammation*, vol. 29, no. 6, pp. 1154–1163, 2021.
- [19] J. Wu and S. Samant, "Su-dd-a4-02: Fast multimodal image registration using mutual information and variable step-size grid search algorithm," *Medical Physics*, vol. 32, no. 6Part2, pp. 1894–1894, 2005.
- [20] Q.-s. Chen, M. Defrise, and F. Deconinck, "Symmetric phase-only matched filtering of fourier-mellin transforms for image registration and recognition," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 16, no. 12, pp. 1156–1168, 1994.
- [21] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proceedings of the seventh IEEE international conference on computer vision*, vol. 2. Ieee, 1999, pp. 1150–1157.
- [22] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: A survey," *IEEE transactions on medical imaging*, vol. 32, no. 7, pp. 1153–1190, 2013.
- [23] M. T. Hossain, G. Lv, S. W. Teng, G. Lu, and M. Lackmann, "Improved symmetric-sift for multi-modal image registration," in *2011 international conference on digital image computing: techniques and applications*. IEEE, 2011, pp. 197–202.
- [24] M. Urschler, J. Bauer, H. Ditt, and H. Bischof, "Sift and shape context for feature-based nonlinear registration of thoracic ct images," in *International Workshop on Computer Vision Approaches to Medical Image Analysis*. Springer, 2006, pp. 73–84.
- [25] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [26] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International conference on computer vision*. Ieee, 2011, pp. 2564–2571.
- [27] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *European conference on computer vision*. Springer, 2006, pp. 430–443.
- [28] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "Brief: Binary robust independent elementary features," in *European conference on computer vision*. Springer, 2010, pp. 778–792.
- [29] Z. Yunsheng and Z. Zhengrong, "Automatic registration method for remote sensing images based on improved orb algorithm," *Remote Sensing for Natural Resources*, vol. 25, no. 3, pp. 20–24, 2013.
- [30] Y. Yu, Y. Guo, R. Wang, S. Yin, and M. Yu, "Registration based on orb and freak features for augmented reality systems," *Transactions of Tianjin University*, vol. 23, no. 2, pp. 192–200, 2017.
- [31] L.-R. Dung, C.-M. Huang, Y.-Y. Wu *et al.*, "Implementation of ransac algorithm for feature-based image registration," *J. Comput. Commun.*, vol. 1, no. 6, pp. 46–50, 2013.
- [32] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.

- [33] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, 1997.
- [34] C. Studholme, D. L. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3d medical image alignment," *Pattern recognition*, vol. 32, no. 1, pp. 71–86, 1999.
- [35] Z. F. Knops, J. A. Maintz, M. A. Viergever, and J. P. Pluim, "Normalized mutual information based registration using k-means clustering and shading correction," *Medical image analysis*, vol. 10, no. 3, pp. 432–439, 2006.
- [36] T. Yang, Q. Tang, L. Li, J. Song, C. Zhu, and L. Tang, "Nonrigid registration of medical image based on adaptive local structure tensor and normalized mutual information," *Journal of applied clinical medical physics*, vol. 20, no. 6, pp. 99–110, 2019.
- [37] Y. Sato, S. Nakajima, N. Shiraga, H. Atsumi, S. Yoshida, T. Koller, G. Gerig, and R. Kikinis, "Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images," *Medical image analysis*, vol. 2, no. 2, pp. 143–168, 1998.
- [38] J. Jin, L. Yang, X. Zhang, and M. Ding, "Vascular tree segmentation in medical images using hessian-based multiscale filtering and level set method," *Computational and mathematical methods in medicine*, vol. 2013, 2013.
- [39] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE transactions on systems, man, and cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [40] X. Xu, S. Xu, L. Jin, and E. Song, "Characteristic analysis of otsu threshold and its applications," *Pattern recognition letters*, vol. 32, no. 7, pp. 956–961, 2011.
- [41] D.-Y. Huang and C.-H. Wang, "Optimal multi-level thresholding using a two-stage otsu optimization approach," *Pattern Recognition Letters*, vol. 30, no. 3, pp. 275–284, 2009.
- [42] L. Vincent, "Grayscale area openings and closings, their efficient implementation and applications," in *First Workshop on Mathematical Morphology and its Applications to Signal Processing*, 1993, pp. 22–27.
- [43] X. Sun, Y. Xu, W. Zhao, T. You, and J. Liu, "Optic disc segmentation from retinal fundus images via deep object detection networks," in *2018 40th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 2018, pp. 5954–5957.
- [44] A. G. Venkat and S. Sharma, "Automated measurement of leakage on wide-field angiography in the assessment of retinal vasculitis," *Journal of ophthalmic inflammation and infection*, vol. 10, no. 1, pp. 1–4, 2020.