Otto-von-Guericke Universität Magdeburg

Fakultät für Informatik

Multi-modal co-registration of high-resolution 7T MRI vessel data

Masterarbeit

Autor:

Lena Spitz

Erstprüferin:

PD Dr.-Ing. habil. Sylvia Saalfeld

Zweitprüfer:

apl. Prof. Dr.-Ing. Gábor Janiga

Magdeburg, 11.11.2020

Lena Spitz
lena.spitz@ovgu.de
Matrikelnummer: 207425
Studiengang: Informatik

Otto-von-Guericke-Universität
Fakultät für Informatik
Postfach 4120
39016 Magdeburg

Acknowledgements

I want to thank my thesis advisor and examiner PD Dr.-Ing. habil. Sylvia Saalfeld for her unfailing support and continuous encouragement throughout the process of developing this thesis. She always supported and encouraged me and readily gave advice.
I also want to thank M.Sc. Franziska Gaidzik for her help and patience in answering questions.
Furthermore I want to thank apl. Prof. Dr.-Ing. Gábor Janiga for accrediting and examining this thesis.
Lastly a most heartfelt thanks to my loved ones who always had an open ear for my worries and who provided much needed stress relief and moral support.

Abstract

Cerebral blood flow gives insight into workings of the human brain, including the development and progression of neurodegenerative diseases. To explore the causality relationships between such diseases and blood flow, simulations can be used. Computational fluid dynamics (CFD) are an appropriate method, but they need a suitable model of the vasculature and patient-specific boundary conditions in order to be realistic. Spatially limited phase-contrast magnetic resonance imaging (PC-MRI) measures blood flow and can thus give the necessary boundary parameters, while time-of-flight (TOF) MRI is high-resolution and can give a detailed vascular model. To make these two sets of data voxel-wise comparable, a co-registration is needed.

This thesis proposes a non-rigid, model-based co-registration of TOF-MRI data to PC-MRI data. The hybrid approach fits the segmented centerline of the TOF data to the intensity ridges that make up the vessels within the PC-MRI volume. A multi-start, coarse-to-fine optimisation strategy based on Powell's method with a scaled sum of weighted intensities metric is used. Global-to-local affine transformations with decoupled parameters are applied following the inherent hierarchical structure of the vessel tree.

Landmark evaluation yielded an average mean squared error of $134.63mm$ and an average Hausdorff distance of $10.08mm$. Among clinical and synthetic test sets, it became apparent that both handling of vessel tree parent-child hierarchy as well as treatment of small vessels that are not represented in the PC-MRI data influenced the registration result.

The presented implementation offers a starting point for further development and could be extended through refining hierarchy treatment and additional non-linear transformations.

Contents

1 Introduction

Understanding blood flow within the cerebral vasculature is a crucial part of understanding the function of the human brain. Realising how external and internal stimuli and relationships influence blood flow and how blood flow influences behaviour is an interdisciplinary concern. Especially the reach of illnesses and diseases that affect the brain are relevant and often related to blood flow. Research into the development and progression of neurodegenerative diseases and cognitive decline becomes more and more important in particular, seeing as the age of the human population is steadily on the rise.
One such disease is cerebral small vessel disease (CSVD), which has been linked to dementia and Alzheimer's disease [1]. Studies have provided evidence that cerebral blood flow declines with advance of CSVD, and that cognitive decay is related to blood flow [2]. Despite progress in this field, more research is necessary, as it is still unclear whether it is the advanced stages of CSVD that lead to blood flow reduction, or if blood flow reduction is what causes CSVD to advance [1, 2].

One way to examine blood flow is via computational fluid dynamics (CFD). They simulate blood flow within a vessel model based on a variety of boundary conditions [3, 4]. For a realistic simulation, patient-specific parameters for these conditions are needed, as well as a model of the patient's vasculature.
One option of measuring blood flow and acquiring patient-specific parameters is with 4D phase-contrast magnetic resonance imaging (PC-MRI), which yields information about velocities within the image volume over time. However, PC-MRI is limited by its temporal and spatial resolution, giving good results for diagnosis of large vessels like the aorta [5, 6], but not so much for small vessels like small cerebral arteries. Instead, state-of-the-art 7Tesla (T) time-of-flight magnetic resonance imaging (TOF-MRI) is high resolution and can depict arterioles down to a diameter of $40\mu M$ [1]. Using the high-resolution vascular structure of the TOF-MRI data and the numerical boundary conditions from the PC-MRI data, the blood flow of the specific patient for whom these images were acquired can be simulated via CFD. For such a simulation, first a co-registration of the PC-MRI and the TOF-MRI data is necessary to establish exact correspondence between the two sets of data. This is because CFD is not just dependent on boundary parameters, but also on vessel structure. Two imaging sessions on the same patient, even if done right after another, will not be perfectly aligned

due to movement artefacts and differences in the modalities. Therefore, a transformation that fits the data of one modality to the other in order to make them voxel-wise comparable is of essence. What this means is illustrated in Figure 1.1.

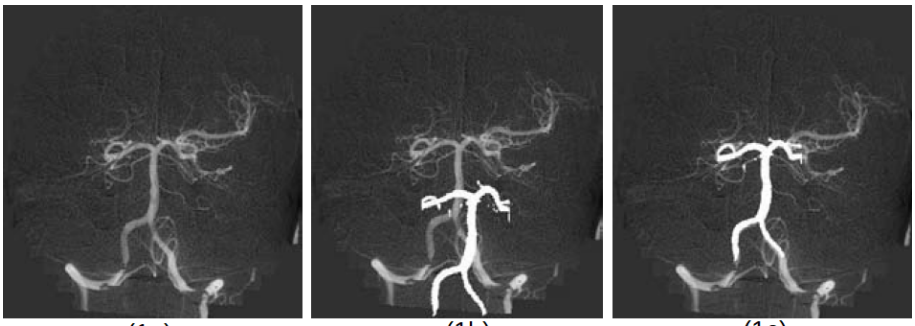

(a) Two-dimensional intensity-based co-registration example. (1a) shows base data, (1b) adds second set of data in white before co-registration, and (1c) shows result after co-registration [7].

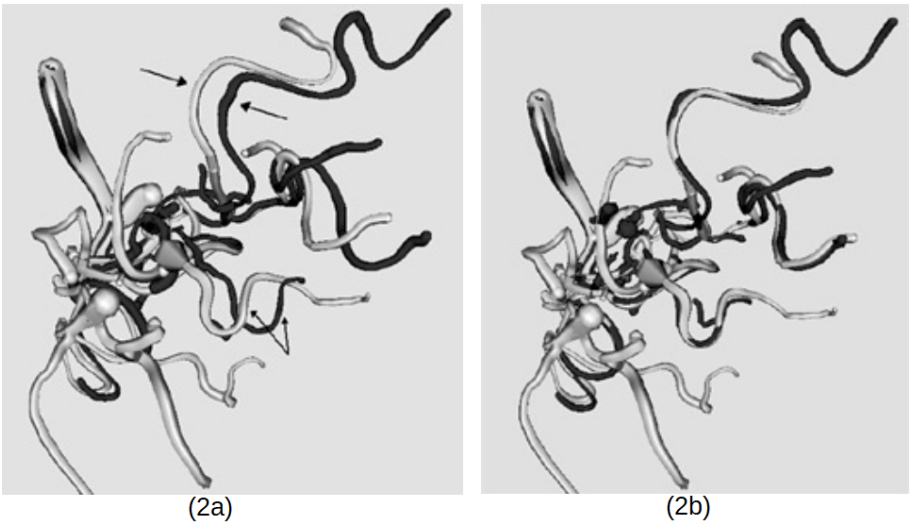

(b) Three-dimensional segmentation-based co-registration example. Black vessels shall be registered to light-grey vessels. (2a) shows data prior to co-registration, (2b) after co-registration [8].

Figure 1.1. Two examples of medical co-registrations, (a) shows a two-dimensional example based on image intensities, (b) a three-dimensional example based on vessel segmentations.

A correct co-registration is a necessary step for enabling CFD in the TOF-MRI model using numerical parameters from the PC-MRI data. Through combination of state-of-the-art 7T TOF-MRI data with PC-MRI patient-specific real parameters for CFD, new insights into cerebral blood flow of even the smallest vessels can be gained. This is valuable not just for further research into CSVD, but for any and all fields exploring cerebral vasculature and its blood flow.

This thesis proposes a co-registration of high-resolution 7T TOF-MRI data to 4D PC-MRI data of the brain's vasculature, specifically of the Circle of Willis. The registration shall

be non-rigid, as a rigid transformation has proven to not be sufficient for this subject, and model-based, making use of the 7T TOF-datas' high resolution and to make it easier to subsequently facilitate CFD application.

2 Prerequisites and Related Work

First, this thesis's background is elaborated. Insight into the modalities used is given, as well as into the medical background. The basics of and related work in co-registration are outlined, specifically for the use-case of vascular structures.

2.1 Medical Background and Modalities

Cerebral small vessel disease (CSVD) means the remodelling of small vessels in the brain that results in stiffening as well as thickening of the vessel wall. It is closely connected to and seen as cause of strokes, white matter hyperintensities, as well as vascular dementia and Alzheimer's disease [1, 2]. At the core of CSVD is its effect on cerebral blood flow and cerebrovascular reactivity, specifically its pulsatility and ability to dilate vessels when blood flow shall be increased [2, 9, 10, 11].
Advanced CSVD concurs with lowered blood flow in the cerebral vasculature, and rate of cognitive decay is dependent on how much the blood flow is reduced [2, 12, 13, 14]. Despite research in this area, the exact causality relation of CSVD and blood flow reduction has not yet been determined [1, 2]. This is partially due to lack of resolution in state-of-the-art imaging methods [2].
Phase-contrast magnetic resonance imaging (PC-MRI) is used to detect cerebral blood flow and pulsatility. However, due to the limited resolution of the modality, detailed results are often only acquired for major vessels with a diameter of a few millimeters, while for other areas it only gives information about their overall blood supply. However, per definition CSVD is a disorder of the small arteries and arterioles with a diameter smaller than $1mm$. Detailed information for the individual small vessels, like velocity, wall shear stress, and other quantitative measures, can not yet be acquired with current imaging modalities [2], and is a relevant research field for many diseases and exploration of cognitive function.

MRI is based on the magnetic spin momentum of hydrogen atoms‘ nuclei. Hydrogen occurs in the human body in abundance, particularly in water and fat, which is why it can be

used to visualise blood and its flow. The basic concept of MRI is using a strong magnetic field to align the axes of the hydrogen atoms‘ spin. The frequency of that spin corresponds to the strength of the magnetic field, which is determined by its number of Tesla (T). The spins are then disturbed, or in technical terms „excited“, using radiofrequency pulses to change the axes of the spins. The receiver coils of the MRI apparatus can then measure the realignment, which differs based on tissue properties. Various parameters of the MRI can be adjusted depending on which kind of tissue shall be highlighted [15, 16]. Figure 2.1 shows two time-of-flight MRI images, which highlights blood vessels, of the anterior Circle of Willis, which is part of the brain's vasculature, and depicts the influence of an MRI apparatus's Tesla on the resulting image.

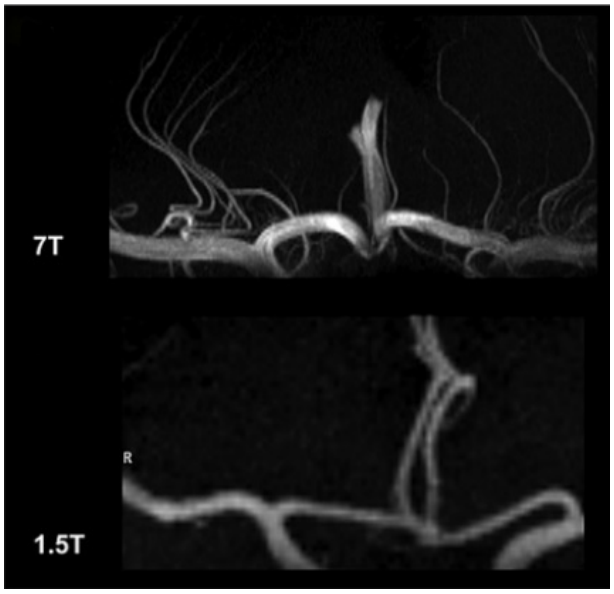

Figure 2.1. Comparison of anterior Circle of Willis acquired with 7T (above) and 1.5T (below) 3D time-of-flight MRI [17].

Each voxel in an MRI dataset represents transverse magnetisation, which is a vector containing magnitude and direction. The latter can also be referred to as phase. Most MRI applications only consider the magnitude to calculate image intensities, but the phase, being direction, can give insights about motion within the acquired volume [18].

PC-MRI is based on that phase shift of moving atoms, since stationary tissue does not have a phase change. This phase shift is measured via a bipolar gradient that encodes the velocity of moving atoms. The gradient is applied two times in opposing directions during the radiofrequency pulses. A gradient measures a difference, so if there is a change in position, meaning something is moving, that gradient will not be zero, but instead gives information about the movement's velocity. Since blood moves through the otherwise static tissues of the body, PC-MRI can thus be used to measure not only the vessel structure but also direction and velocity of blood flow as well as other hemodynamic parameters [19, 20].

PC-MRI usage is limited by its long acquisition time and dependency on imaging parameters. The bipolar gradient can only be measured one direction at a time, so in order to visualise blood flow in all directions, the imaging procedure has to be done three times, for x-, y- and z-direction, and a fourth time for the reference image [19]. Additionally, the strength of the bipolar gradient has to be set previously to imaging and dictates the maximum velocity of blood flow that can be detected. This value is called velocity encoding sensitivity (VENC) and it determines the kind of aliasing, noise and other artefacts in the final imagine volume. Selection of VENC can be complex, as a low value leads to aliasing of vessels with a faster blood flow, while a high value creates more noise, though an increase of Tesla used with the MRI scanner generally leads to less noise [20]. Different vessel sizes having different blood flow velocities, together with the spatial limitations of the modality, make PC-MRI usable for big vessels like the aorta, but thus far not adequate for smaller vessels, even when using a high Tesla MRI [5, 6].
Initially, PC-MRI was applied in 2D, where only one slice at a time was acquired. By now, 4D PC-MRI is standard, hence not only acquiring a 3D volume data but also including a time component and visualising blood flow over time. VENC should be chosen with the blood flow velocity of the vessel of interest in mind. This makes finding the best VENC value a challenge when an area shall be scanned that contains vessels of various different sizes, as vessel size affects blood flow velocity [19].
After image acquisition, a magnitude volume can be generated, combining the phase volumes and their information about blood flow in one direction into one set where overall blood flow properties and anatomy can be examined [19].

Figure 2.2. Comparison of the same patient's Circle of Willis acquired with PC-MRI (left) and TOF-MRI (right).

Rather than PC-MRI, **Time-of-flight MRI (TOF-MRI)** is the commonly used modality for acquisition of vessel images [20, 21]. It uses high-frequency radiofrequency waves to oversaturate static tissue, meaning that fresh blood that moves into the imaging volume will not yet have been excited by as many pulses, thus not being as saturated. This gives incoming flow a higher signal and highlights it in the final image [20, 21], like in Figure 2.1. In order to only highlight arteries and suppress the signal from veins, so-called saturation bands are used in the areas of the veins, as well as other techniques to further clear up the image [20, 21]. As with PC-MRI, a higher Tesla MRI scanner will lead to even better results [17]. A 7T TOF-MRI scanner can depict vessels with a diameter as small as $40\mu M$ [1]. A comparison of vasculature segmented from the PC-MR and TOF-MR images of the same patient can be seen in Figure 2.2.

2.2 Co-Registration of Vascular Image Data

Co-registration, also simply called registration, is particularly important in a medical context. Due to various artefacts during image acquisition like geometric distortions, partial volume effects, gradient inhomogeneities or motion artefacts, even two images of the same patient do not perfectly align and have to undergo registration [22]. Registration can make images acquired with various modalities comparable on a voxel level, which can then be used used for diagnosis, assessment, treatment and therapy planning as well as research. It enables comparison of the same patient at different times, and even direct comparison of two different patients with individual anatomy by creating an atlas [23, 24, 25]. Another use case is the registration of high-resolution 3D images taken pre-operatively to real-time 2D images taken during the intervention, thus aiding the surgeon without the drawbacks of often space-related problems of such higher resolution 3D imaging modalities [26].

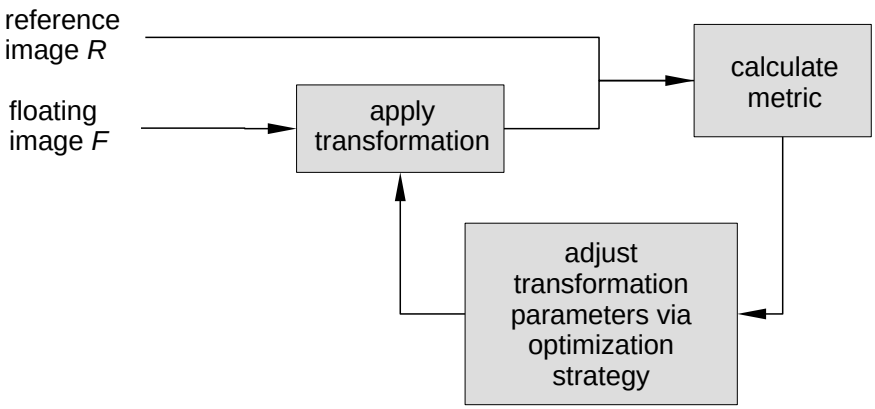

Figure 2.3. The basic workflow of a co-registration.

In general, co-registration is the process of finding a transformation of an image that shall be moved, floating image F, that makes it voxelwise comparable to a reference image R [27, 28]. This basic workflow of a co-registration is illustrated in Figure 2.3.
First, a feature room is defined, meaning the image features via which the two images shall be compared. This feature room can be based on appearance, meaning the raw or filtered image data, or on extracted data, like models or segmentations, which can be referred to as geometry-based approaches [29]. Within this feature room, a similarity criterion has to be determined to quantify the level of match between the two images. This similarity criterion, also called metric or cost function [30], then has to be optimised, which means its value shall be maximised or minimised (depending on definition.) This is done by adjusting the transformation parameters defined by the feature room [27, 28, 31, 32], thus transforming the floating image.

There is a huge amount of co-registration approaches, as even just in the medical field there is a variety of different application scenarios, all of which benefit from different registration methods. All of the approaches are based on the presented workflow, and there are different ways to classify them [30, 32]. One way to classify co-registration techniques for vascular structures is based on the following 13 categories [30]:

1. Application
2. Modality
3. Dimensionality
4. Subject
5. Encoding function
6. Geometry representation
7. Cost function
8. Image representation
9. Registration bias
10. Optimisation
11. Global geometric transformation
12. Local deformation model
13. Validation

2.2.1 Application, Modality, Dimensionality and Subject

The first set of classification categories are self-explanatory: **Application** refers to the anatomical application area, for example *head, liver,* or more specifically the exact vascular networks that shall be registered, like *Circle of Willis.* **Modality** refers to the imaging modality, for example which MRI or computer tomography techniques were used to acquire the images. **Dimensionality** refers not only to how many dimensions the used modalities cover, but also to whether more than one modality was used, thus separating the category in mono-modal and multi-modal approaches. It also specifically denotes if one of the dimensions is time [30]. Counterintuitively, a registration technique does not always have the same dimensionality as the modalities used. This can be achieved by averaging the time-dimension or only choosing the timepoint with the highest signal [33], or higher dimensions can be projected into lower ones [7, 34, 35, 36]. For example, Chan et al. [34] reduce a 3D-3D registration problem to a 2D-3D one by generating three 2D projections along the orthogonal axes from the pre-processed binary 3D volume.
When a registration co-registers images from different patients, the approach is tagged *inter-subject*, while when it uses images from the same patient it is tagged *intra-subject*. These tags describe the category **subject**.

2.2.2 Encoding Function and Geometry Representation

Encoding function is nothing but the above mentioned feature room that defines what sort of features are used for registration. Figure 2.4 shows how the feature room overall is classified and how encoding function and geometry representation come into play.

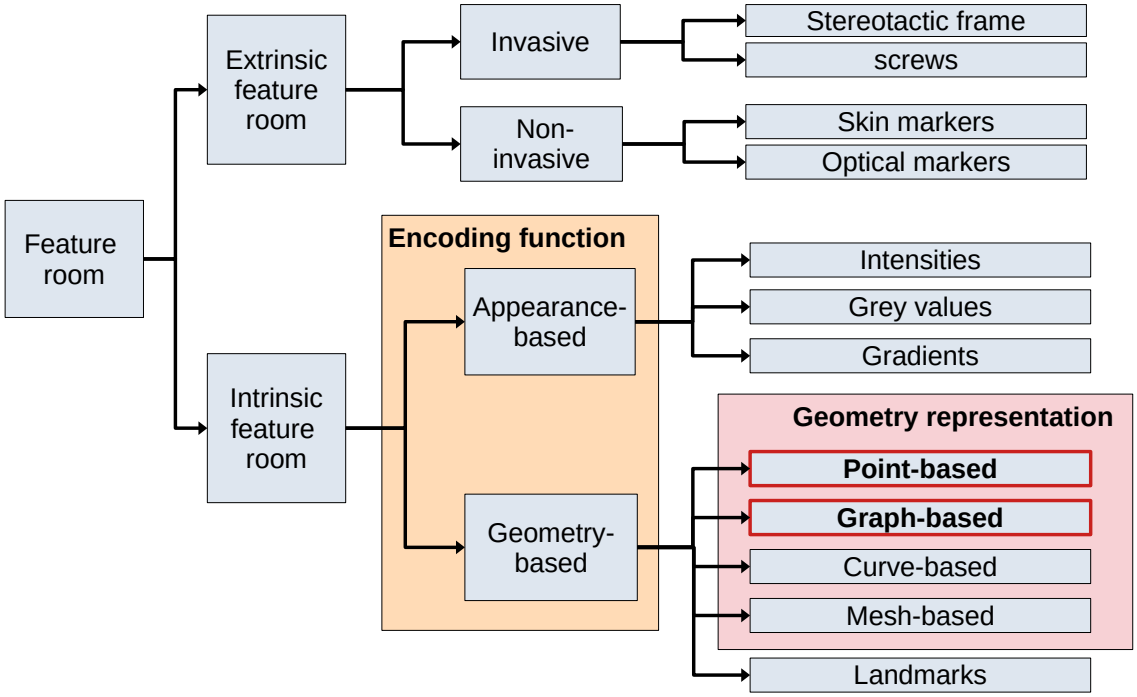

Figure 2.4. Feature room classification and where encoding function and geometry representation come into play. Point- and graph-based geometry representation, which infer a geometry-based encoding function, are particularly relevant for registration of vascular structures and are highlighted with red borders.

Generally, the feature room can be separated into extrinsic and intrinsic measurements [27, 32], though in this classification only intrinsic measurements are considered [30]. This is because extrinsic features are entirely based on external markers, which can be invasive like stereotactic frames or screws, or non-invasive like optical or skin markers [16, 27, 32]. Intrinsic markers can be separated into *appearance-based* and *geometry-based* [28, 30, 37]. They are not added externally, but their separation is determined by whether they are inherent in the raw or filtered image data, like voxel features, or derived from models that are extracted from the image. Voxel features can be *intensities, grey values, or gradients* [34, 38, 39]. They are easy to obtain and inherently included in the raw image data, no extra steps are necessary to use them, and they are usually obtained globally, not only in specific areas [27, 32]. Such *appearance-based* features, also called *intensity-based* features, are usually compared based on their similarity [30]. *Geometry-* or *model-based* features rely on a model of the image data. A common approach is a segmentation of the relevant structures, which can be *rigid point clouds, curves* or *surfaces* or even *deformable meshes* [27], a distinction that is defined by the category **geometry representation** [30]. A vast amount of registration algorithms meant for vascular images are *geometry-based*, more specifically making use of the vessel structures in form of centerlines and bifurcation points, thus making them *point-based* [36, 40, 41, 42, 43, 44]. Another approach to vessel

segmentations is *graph-based*, making use of the inherent tree structure of vasculature, with an inlet vessel as root, bifurcations as nodes with the vessels connecting them being the edges, direction being towards the smallest vessels which represent the leaves [8, 44, 45]. Segmentations provide a good basis for registration, but the result is highly depended on the segmentation method, many of which are error prone [34, 41, 46]. This is why appearance-based approaches are considered to be more accurate, yet also computationally less efficient [34].
Another option for geometry-based techniques are *landmarks*, which function similarly to extrinsic markers, but instead of using external markers the position of the markers are part of the image. Landmarks can still be external, though in this case it means external features of the patient like the tip of the nose or the navel, as opposed to internal markers like vessel bifurcations or the border of an organ [31]. Using landmarks has the advantage of simple implementation and only using data that is already in the image, but the same drawbacks as for extrinsic markers remain: the quality of a marker-based registration will always rely on the quality of the marker placement and selection. Additionally, a minimum of four markers are needed, and finding enough unambiguous landmarks in the data can be a challenge, especially since determination of exact positions can be difficult due to patient specific anatomy [31]. Nevertheless, they can and are being used for registration of vascular structures [47, 48], set either manually [39, 43] or automatically [24]. Combining them with centerline and bifurcation segmentations, which give a good basis for landmark extraction, is also possible [39, 41, 42, 43].

2.2.3 Cost Function

Using these appearance- or geometry-based features, be it intensities, landmarks or point clouds, a similarity- and/or distance-based metric, the **cost function**, is defined [30, 31, 32]. There is a variety of approaches, like *sum or mean of squared or scaled differences* or *absolute differences* based on Euclidean or other distance measures [34, 35, 36, 45, 49], various *gradient-* or *correlation-based* methods [43, 50], using *entropy, patches, mutual information* [24, 38, 51], or even taking a *sum of weighted intensities* [8, 25, 46, 52]. *Intensity-based* cost functions like *mutual information* or *correlation ratios* are considered hard to optimise, as there are "many non-linear, potentially discontinuous terms involved that result in functions that are non-smooth and irregular" [51]. *Distance-based* cost functions are generally based on segmentations, and thus suffer the same drawbacks as geometry-based registration techniques.

2.2.4 Image Representation and Registration Bias

Related to encoding function and cost function is **image representation**, which categorises if both floating and reference image are represented as images (meaning via appearance or intensities), as models (meaning via geometry), or one as image and one as model [30], thus tagging the algorithm as *image-image, model-model or model-image*. **Registration bias** also refers to the differentiation between *model-* and *image-based* representations, though it more specifically categorises the values that are compared. *Intensity-based* means intensities from both images are used, *model-based* means models from both images are used, and a *hybrid* algorithm compares models to intensities [30].
Such *hybrid* techniques are particularly suited to registration of vascular structures, as centerlines can be matched to intensities. This has the added benefit of not being as error prone to segmentation errors, as only one image is segmented, yet still making use of the entire image instead just one extracted model [46]. Performing just one segmentation also means that more computational effort can be put into it rather than splitting computation costs between two sets of data [46]. Aylward et al. [46, 52] did a lot of ground work in this area and presented an approach that matches extracted centerlines from one dataset to the intensity ridges formed by the vessels in the other image, presented as intensity volume. Cost function is made up of *sum of scaled and weighted intensities* corresponding to the *centerline points* after transformation. Centerline points are weighted according to their radius since smaller vessels are more likely to be affected by noise [52]. This algorithm is further used and adapted in other registration methods [8, 25], among them Jomier et al. [8], who make use of the inherent hierarchical tree-like nature of vasculature to apply local transformations, traversing from root, meaning vessel inlet, to leaves, meaning child vessels.

2.2.5 Global Geometric Transformation and Local Deformation Model

The nature of the transformation, also called search room, is a crucial part of the registration, and can be split into **global geometric transformation** and **local deformation model** [30]. It determines how the best transformation is found for the given feature room, meaning it determines what sort of transformations are valid [27, 30, 32].
The simplest transformations are *rigid* transformations, which include rotation, translation and scaling, and preserve angle and distance between objects [32]. Rigid transformations are a subset of *affine* transformations, which add shearing parameters and preserve parallelism but not angles. They are considered the simplest, most commonly used transformations in registration [51]. *Projective* transformations are like higher order *affine* transforma-

tions, but do not preserve parallelism, and are a subset of *elastic* transformations, which, unlike *affine,* map lines to curves instead of lines [30, 32]. All these transformations can be used globally, while *deformable* transformations are applied locally. Those transformations include *B-Splines, Thin-Plate-Splines* and *displacement fields* in addition to the above mentioned transformations [28, 30]. Free-form deformations based on one-dimensional cubic *B-Splines* are one approach used by Rueckert et al. [38, 53]. They shift the points of an underlying mesh, only ever affecting points in the immediate neighborhood rather than the entire image, leading to "a smooth and continuous transformation" [38]. The resolution of such a deformable mesh determines the degrees of freedom (DOF) of the transformation, as each point adds one DOF. A $10 \times 10 \times 10$ deformable mesh describing a transformation can thus have 3000 DOF, while a rigid transformation has 6 DOF and an affine has 12 DOF. However, a higher DOF transformation does not equal a better transformation [54]. Local transformations might have higher DOF, but they do not take into account the entire information given by the image and thus cannot reliably find the best global fit [51]. There are also studies arguing that locally rigid transformations can achieve comparable results to deformable transformations [41, 46], which are more computationally complex, though there are also studies arguing the opposite, that higher order transformations are needed to accurately register vasculature [51, 53]. The latter is mentioned particularly in inter-subject applications due to patient specific anatomy [53].

2.2.6 Optimisation

While all the previously introduced categories are important, at the core of any registration is the **optimisation.** It determines the strategy by which the best transformation, meaning the transformation that results in the best metric via the defined features, is found. There is an abundance of optimisation algorithms, but they can be divided into *derivative-based* and *derivative-free* [30], as the computation of derivatives can be complex. Well known derivative-free algorithms include *Brent's method, Powell's method, greedy search* and *best neighbour,* while *Newton's method, nonlinear conjugate gradient* and *gradient descent* represent some of the most widely used gradient-based ones [30]. Other optimisers include *probabilistic/stochastic* approaches like *Gaussian models* [30], *evolutionary algorithms* [25], *least squares* [42], or in the case of *point-based* registration, algorithms like *iterative closest point (ICP)* techniques can be used [36, 39, 41, 42].
Powell's method [55] is considered to be efficient [51] and one of the most commonly used optimisation methods [7, 34, 35, 36, 51]. It is gradient-free and makes use of the concept of conjugate directions [55].

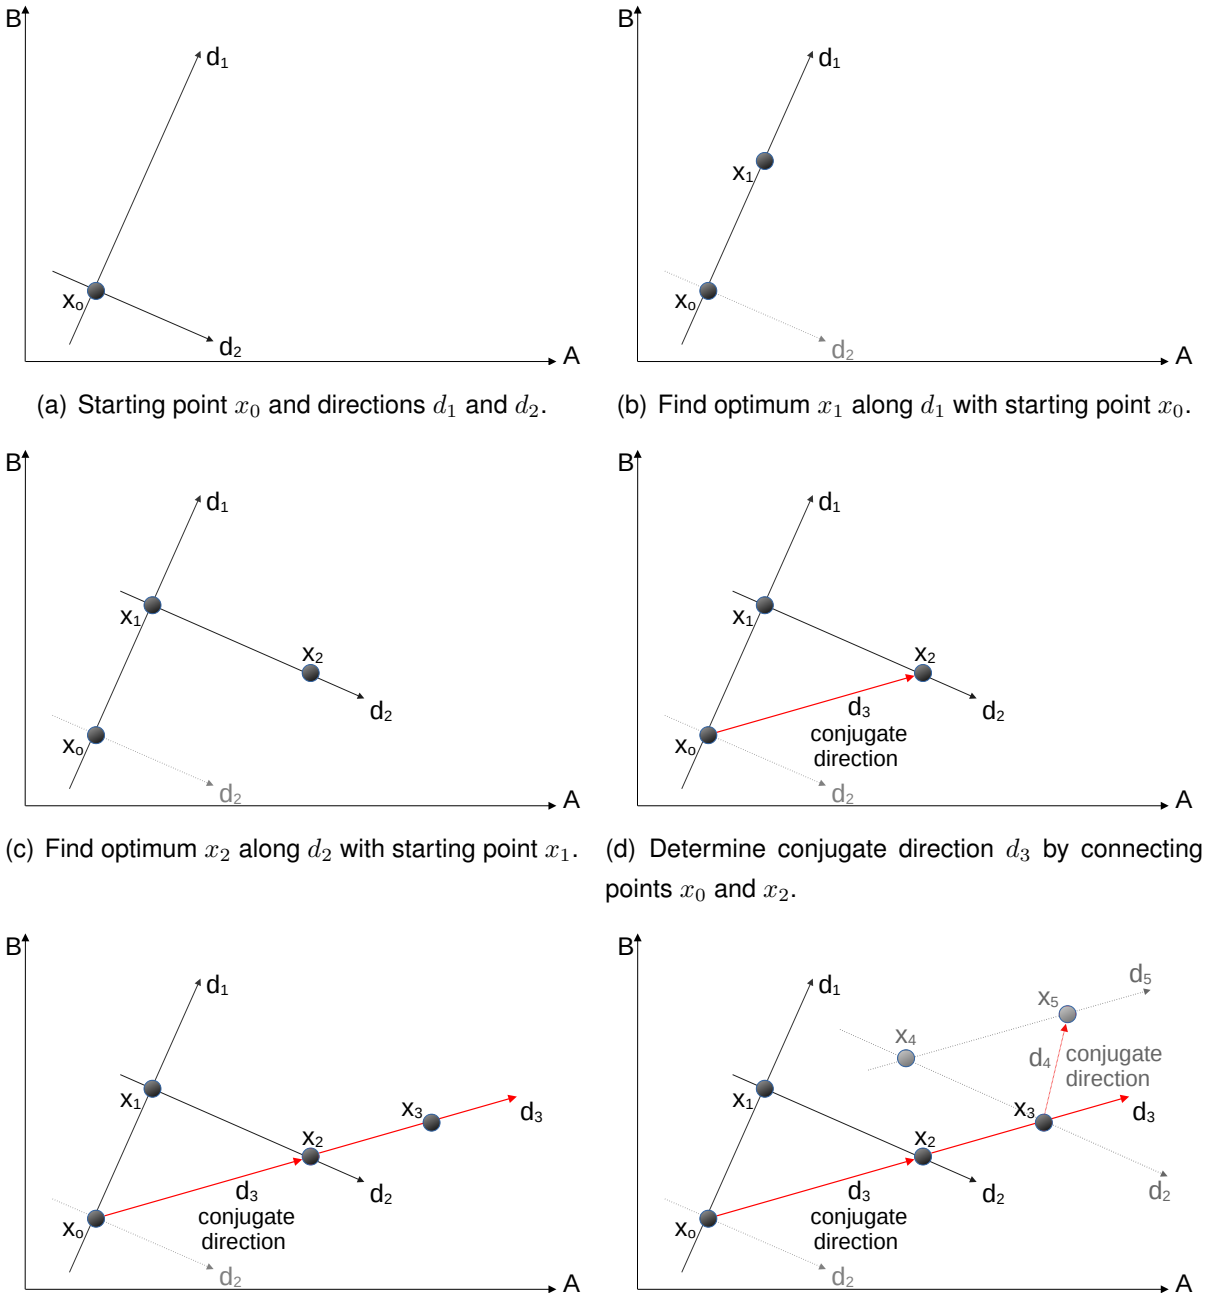

(a) Starting point x_0 and directions d_1 and d_2.

(b) Find optimum x_1 along d_1 with starting point x_0.

(c) Find optimum x_2 along d_2 with starting point x_1.

(d) Determine conjugate direction d_3 by connecting points x_0 and x_2.

(e) Find optimum x_3 along conjugate direction d_3 with starting point x_2. This concludes first iteration.

(f) Start second iteration with starting point x_3, using d_2 as first and conjugate direction d_3 as second search direction.

Figure 2.5. Illustrating first and beginnings of second iteration of Powell's method within parameter space A and B.

It first selects a starting point $x0$ in the parameter space and two search directions $d1$ and $d2$ (Figure 2.5(a)). Starting at $x0$, the best metric along $d1$ shall be found via one-dimensional optimisation, and the parameter set found is denoted as $x1$ (Figure 2.5(b)). Then, using $x1$ as new starting point and $d2$ as search direction, $x2$ is found, again as the

best set of parameters in that one-dimensional optimisation (Figure 2.5(c)). Connecting the starting point $x0$ and the just found second best point $x2$ yields the conjugate direction $d3$ which points towards the optimum metric in that parameter space (Figure 2.5(d)). From $x2$ along the conjugate direction $d3$ again the best point $x3$ shall be found, meaning the parameter set with the best metric result from a 1D optimisation (Figure 2.5(e)). This concludes the first iteration. $x3$ is used as the starting point for the second iteration, with $d2$ becoming the first search direction and conjugate direction $d3$ becoming the second (Figure 2.5(f)) [56]. For quadratic functions, this second iteration already finds the global optimum, but even for non-quadratic function the found best set of parameters on the conjugate direction is closer to the global optimum that the initial search point [55, 56].

When gradient computation is not an obstacle, gradient descent or ascent are commonly used in a variety of registrations [8, 24, 46]. There, gradients of the cost function are computed and followed in direction of the desired optimum, where the best parameter set can be found.

The goal of any optimisation is to find the global optimum of the cost function. This is often a minimisation problem, but for some metrics the optimisation finds the maximum instead of the minimum. The biggest issue for any optimisation method are local optima, where the cost function has the smallest (or highest) value in its neighbourhood (Figure 2.6(a)), leading the optimisation towards that value when beyond the neighbourhood an even smaller (or higher) value might exist as the global optimum (Figure 2.6(b)).

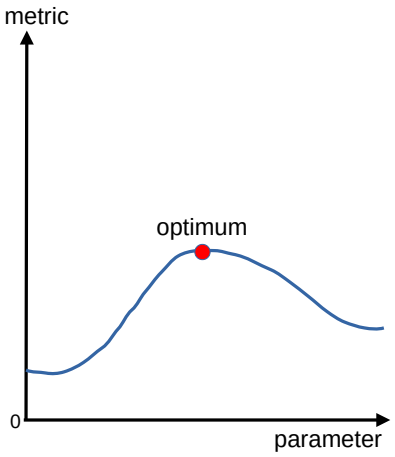

(a) Optimum within local neighbourhood.

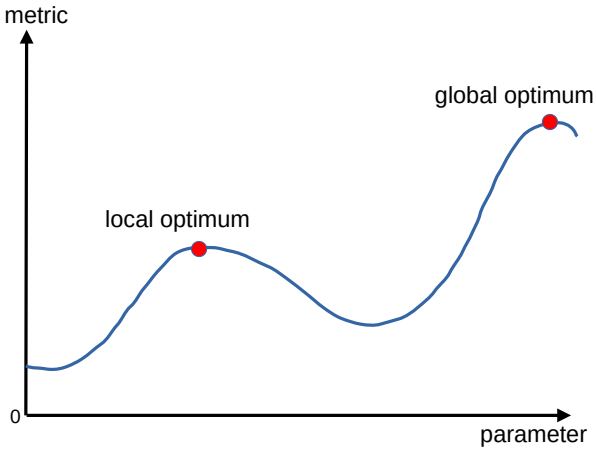

(b) That same optimum turns out to only be a local optimum when viewed on a global scale.

Figure 2.6. Example of local vs global optimum in a simple maximisation example.

One measure to prevent termination in a local optimum is a *multi-start* method, meaning the registration is not started once with one set of parameters, but several times, each time using a different starting point spread over the feature room [30, 35, 39, 51]. Another option

is a *multi-level* approach, meaning a coarse-to-fine strategy is used. One common way of doing this is first running a low-computation-effort global registration, often using only rigid transformations, which is then used as initialisation for a more complex local non-rigid registration, thus effectively combining the advantages of global and local registration to eliminate their disadvantages [8, 24, 42, 43]. This approach is often combined with even more multi-level measures. The assumption generally is that an optimiser is less likely to slip into a local optimum when bigger, meaning coarser, transformation steps are used first, covering more ground within the feature space. This also tends to make registration more efficient, as the finer steps where high computational costs arise, can be limited to the previously established area from the coarser, less costly steps [51].
For intensity-based registration, multi-level often means optimisation is done on different resolutions, leading to registration starting on a down-sampled resolution. Meanwhile model-based registrations often refine the sampling frequency of the model [30]. Aylward et al. [46, 52], who did a hybrid approach, make use of a variety of these measures. They map the segmented centerlines of one dataset to the intensities of another, using intensity ridges making up the vessels in the raw image. Their metric therefore maximises the scaled sum of weighted intensities:

$$M(T) = \frac{1}{\sum_{i=1}^{n} w_i} \sum_{i=1}^{n} w_i I_{\kappa\sigma_i}(x_i T) \tag{2.1}$$

T denotes the current transformation for which the metric is evaluated, $I_{\kappa\sigma_i}$ is the intensity at the current centerline point x_i, while n represent the total number of centerline points, which can be sampled at a high or lower frequency to control the coarse-to-fine approach. w_i are the weights, which are determined by the radius around the centerline at point x_i and can be used to first heighten the influence of big vessels before focussing on smaller ones. Another multi-level adjusting parameter is $\kappa\sigma_i$, which is the standard deviation of a Gaussian kernel that is used to blur the data. A higher $\kappa\sigma_i$ leads to a stronger blurring and thus eliminates noise, but also smaller vessels. Having a higher blurring parameter at the start thus lowers the chance of misregistration due to noise, and lowering it will lead to smaller vessels later being considered in the registration [46, 52].

2.2.7 Validation

Last category for registration classification is **validation** [30]. While there are a lot of different ways to validate, this category does not refer to the validation method, of which there are many, for example landmark-based or indirect validation [27]. Instead, the data used for validation determines the classification, and the three options are *synthetic, phantom*

and clinical. Synthetic data is artificially created, has the advantage of a definite ground truth and can be used when clinical data is not available. Clinical and phantom data can be considered more realistic, but often have noise and artefacts, and a ground truth can not always be determined. Clinical data used for validation are the best option, but not always feasible due to lack of access to patients. Phantom data are a good middle way between synthetic and clinical, as they can be created without patients but are acquired with the relevant modality used for clinical data [30].

2.2.8 Application-Specific Co-Registrations

Through all these different categories, registration techniques are highly specialised for their respective use case, and a method that works well for one application area, modality or subject might not work for another. In regards to methods best for inter- or intra-subject registration, Charnoz et al. [44] note that *centerline tree structures* and their *tree matching* algorithm [45] is especially good for *intra-subject* registration, as it is the same tree that shall be registered. Crum et al.'s [54] comparison of three different non-rigid registration algorithms for example notes that *B-splines* are widely used for *inter-subject brain* registration, while *fluid* methods, which are "mathematical model[s] of a compressible viscous fluid [that] model[...] the transformation between images" [54] are more often used for *intra-subject* registration [54].
For registration of vessel structures, a *mutual information* cost function has the drawback of being error prone [7, 34] because it takes surrounding anatomical structures into account, too, rather than just the vessels. This affects vessels alignment because they tend to make up only a small portion of the entire image, mreaning the registration will rely more on other tissues than the vessels, resulting in lowered accuracy [46, 52].
Another aspect of registration of vasculatures are vascular network changes. Zana et al. [40] specifically account for that in the context of temporal *intra-subject* changes, though their algorithm, which uses a *Bayesian Hough Transform* to estimate the best *global affine* transformation based on *point-matching* of *segmented bifurcation points*, is robust for *multi-modal* registration too. Aylward et al. [46, 52], too, specifically take care to account for only partially overlapping vasculature, meaning only part of the vessel structure in both images is the same. They are also motivated by registering high-resolution data to one of much lower detail in particular.

An example for a specific use case outside of medical context would be Sablatnig et al.'s [57] work, who developed a model-based registration for shards of archaeological finds, making use of the fact that shards that belong to the same vase have the same axis of rotation due

to the nature of pottery, and that rigid body transforms are sufficent for non-deformable clay surfaces.

3 Method

Due to the high amount of methods available for co-registration, conceptualisation was an important step before development began. To determine the basic setup and which methods have to be selected and which are fix, the co-registration of this work was sorted into the described registration classification system [30].

The parameters for this thesis were a multi-modal TOF-MRI to PC-MRI, intra-subject registration of the Circle of Willis. Technically, dimensionality could have been 3D-3D$+t$, as PC-MRI data has been acquired over time, but the time dimension was excluded to fit the scope of this thesis.
When working with the presented data, other teams [33] were not successful when using automatic registration solutions that used rigid transformations like Slicer 3D (National Institutes of Health, 2020). A non-rigid approach was therefore a prerequisite.
Furthermore, model-based registration has proven a wide application area and has previously been used for co-registrations using these modalities [7, 34], as well as for the registration of vascular structures in general, as has been established in Section 2.2.
This was why originally a true model-model registration was aimed for, meaning extraction of centerlines in both modality data sets and registering based on those sets of centerline points. This approach was however quickly changed to a model-image hybrid strategy, as during pre-processing it became apparent that centerline extraction was not easily possible for the PC-MRI data. This did not pose a problem, as there are established strategies for hybrid registration strategies for vascular structures. Aylward et al.'s [46, 52], as well as Jomier et al.'s [8] were suitable inspirations, using the fact that vessels form intensity ridges in image data for a centerline point-to-intensity registration. They also takes into account the inherent graph-structure of a vascular network. Aylward et al.'s [46, 52] work further is suitable as basis for this thesis's goals, as it specifically considers a registration of high-resolution data to low-resolution ones, as well as data that only partially overlaps, which is the case for the TOF- to PC-MRI co-registration, as many smaller vessels that are visible in the TOF-MRI data are not represented in the PC-MRI data.
A multi-level coarse-to-fine approach is almost a given for any complex co-registration, especially one using both local and global transformations. The overall optimisation strategy could be based on Jenkinson et al.'s [51] approach which in turn uses Powell's method [55]

like many other registrations do, and adds a multi-start approach to the overall strategy.
Based on these related works, an affine transformation seemed fitting, as it naturally extends rigid registration as well as the above mentioned works, and is commonly used in registrations and is argued to replace non-rigid transformations on a local level in the presented intra-subject application [41, 46, 51].
Table 3.1 presents the classification of the described concept.

Category	Tag
Application	Circle of Willis
Modality	TOF-MRI, PC-MRI
Dimensionality	multi-modal, 3D-3D
Cost function	sum of weighted intensities
Encoding function	geometry-based
Geometric representation	point-based, graph-based
Image representation	model-image
Registration bias	hybrid
Global geometric transformation	affine
Local deformation model	affine
Optimisation	multi-level, multi-start, derivative-free (Powell's method)
Subject	intra-subject
Validation	clinical and synthetic

Table 3.1. The presented TOF-MRI to PC-MRI co-registration as sorted into Matl et al.'s [30] classification of vessel-based registration algorithms.

3.1 Data Acquisition

The bulk of this work was developed using data from one healthy volunteer.
The 4D PC-MRI data was acquired on a 7T whole-body MRI system (Siemens Healthineers, Erlangen, Germany) using a 32-channel head coil (Nova Medical, Wilmington, MA, USA). A radiofrequency-spoiled gradient echo with quantitative flow encoding in all three spatial dimensions [58, 59] was used as a basis for the image acquisition sequence [33]. It resulted in 17 time steps, for each of which there were three velocity maps that measure velocity in x-, y- and z-direction, and one magnitude image. The represented voxel size is $0.640625mm$ and has a temporal resolution of $54.4ms$. The VENC value, which sets the highest measured velocity, was set to $0.9m/s$.
The TOF-MRI data is of the same healthy volunteer and MRI system, which was set to high-resolution, giving an isotropic voxel size of $0.32mm$.
Both sets of data underwent a prospective motion correction system in-scan [33].

3.2 Pre-Processing

The PC-MRI data were first processed with the automated tool described by Bock et al. [60] in MeVisLab 2.3.1 (MeVis Medical Solutions AG, 2020), which includes noise masking, antialiasing and conversion to EnSight file format (ANSYS Inc, 2020). The data were then further segmented according to Gaidzik et al.'s [33] proposed workflow. This entailed temporal averaging of all three phase as well as the magnitude images, resulting in a single 3D volume each. The phase images were then multiplied with the original temporally averaged ones, after which the vessel surface was extracted via thresholding [33]. For registration, only one time point was used, namely the one with the highest signal peak, meaning the one with the most flow information, which entails the most vessel information [33].

To prepare the 7T TOF-MRI data for centerline extraction, its DICOM data files were preprocessed in MeVisLab 3.3 (MeVis Medical Solutions AG, 2020). First multiple vesselness filters were used: one with six scales with $\sigma = \{1, 2, 3, 4, 5, 6\}$ acted as base, but was additionally summed up with two one scale vesselness filters with $\sigma = \{1, 6\}$, respectively, to further emphasise the thickest and thinnest vessels.
A mask for the vessels was then created via region growing with a threshold of 1.25. The result was morphologically dilated with a $5 \times 5 \times 5$ kernel in case small vessels were missed. This mask was then applied to the original data, leaving only the vessels, and acted as a base for segmentation.

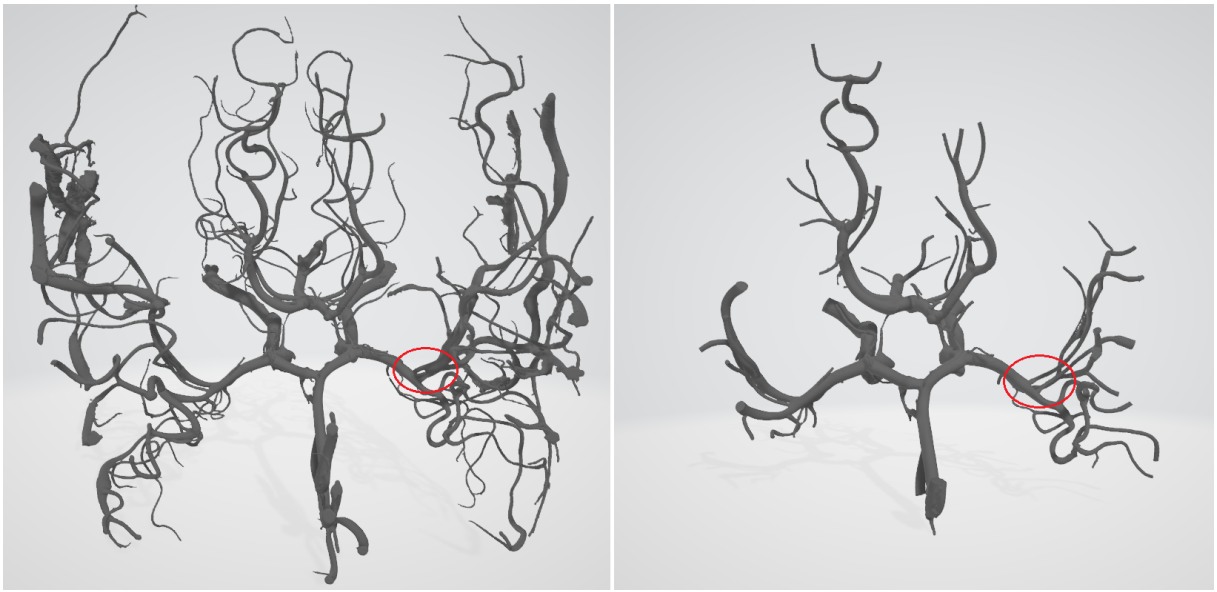

Figure 3.1. Comparison of the TOF-MRI segmentation before and after manual editing, with an area where fused vessels were separated highlighted in a red ellipses.

For segmentation of the vessel tree, the tool and pipeline proposed by Saalfeld et al. [61] was used. Parameters $center$ and $width$ for windowing and thresholding were set to 0.035 and 0.05, respectively. The $center$ value was placed so low in order to keep small vessels in the center of the Circle of Willis, particularly the posterior communicating arteries (PCOM) which connect the vessel trees from the three inlet vessels. This however also meant that there were artefacts in the periphery, and several other vessels were melted together. To ensure more accurate centerline extraction, the vessel structure was next edited using Blender 2.78 (The Blender Foundation, 2002) and Sculptris Alpha 6 (Pixologic Sculptris, 2018). There, artefacts as well as the smallest vessels of the periphery, where exact segmentation would not be possible due to noise, and vessels that were fused together due to low $center$ value during segmentation were separated. Figure 3.1 shows the TOF-MRI data segmentation before and after manual editing and highlights an area where heavy edits due to fused vessels were necessary.
Next, the centerline was automatically extracted using the Vascular Modeling Toolkit VTMK 1.4.0 (Orobix, 2020). Finally, the .vtp file yielding from VMTK was converted into a .vtk file using ASCII with ParaView 5.4.1 (Kitware Inc., 2020).
After pre-processing, the PC-MRI data is represented as a volume of intensity values, and the TOF-MRI data as a centerline. The resulting registration problem is depicted in Figure 3.2. The TOF-MRI centerlines will be fitted to the intensity ridges in the PC-MRI data. The picture also shows the mismatch of corresponding vessels that the transformation resulting from the registration shall fix.

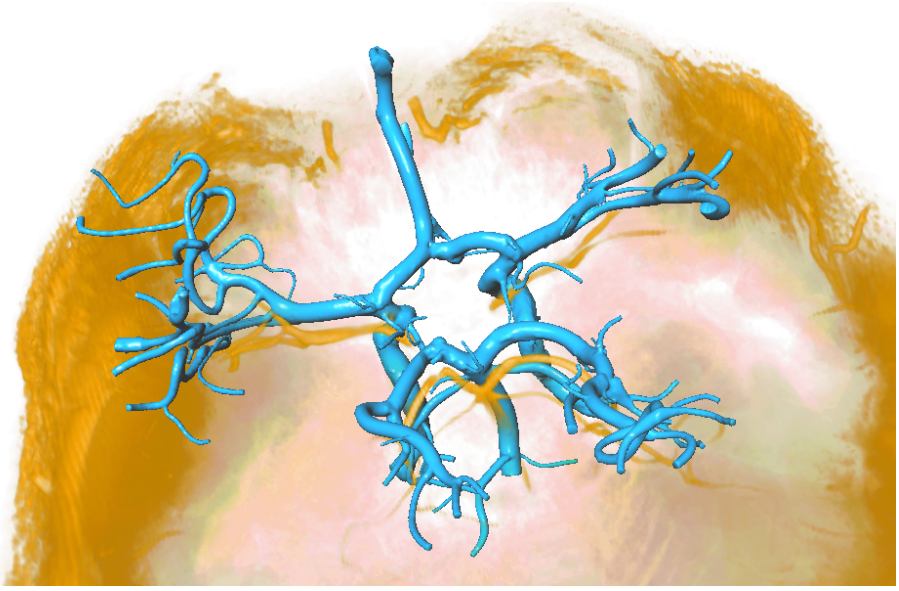

Figure 3.2. PC-MRI data and TOF-MRI data. The PC-MRI data is a volume of intensities. High intensities are drawn in orange-yellow, highlighting what it means that vessels are shown as intensity ridges in the intensity volume. The TOF-MRI data is shown here as segmentation in blue, though in the registration it is represented by its vessels' centerlines

3.3 Data Structure

The registration was entirely developed in Matlab R2020a (The MathWorks Inc. 2020).

Since the PC-MRI data was already given as magnitude data, its handling was trivial. The intensity values of the DICOM data are saved in a matrix of the same resolution size as the dataset per dimension. However, drawing the PC-MRI data was not trivial. For one, drawing it as a solid 3D volume would of course obstruct anything but the outside edges. This was why only one axial slice was drawn at a time, and a slider was added to draw additional slices if desired. The way Matlab internally saves images was another issue, which led to the x- and y-dimensions being switched.

The TOF-MRI data is represented in Matlab as a polygon mesh and the extracted centerline. The mesh .obj file is loaded via the function proposed by Harwin et al. [62] and saved as two arrays, one for the vertices and one for the faces.
The centerline is split into three subtrees. This was done for the previously mentioned centerline extraction. The Circle of Willis has multiple inlets and can contain circles, which led to some vessels not being assigned a centerline in VMTK. Since there are three inlets, the basiliar artery and the left and right internal carotid arteries, the vessel structure was split into three subtrees. They were cut apart at the posterior communicating arteries and the anterior communicating artery, as well as one cut at the anterior cerebral artery, as seen in Figure 3.3.

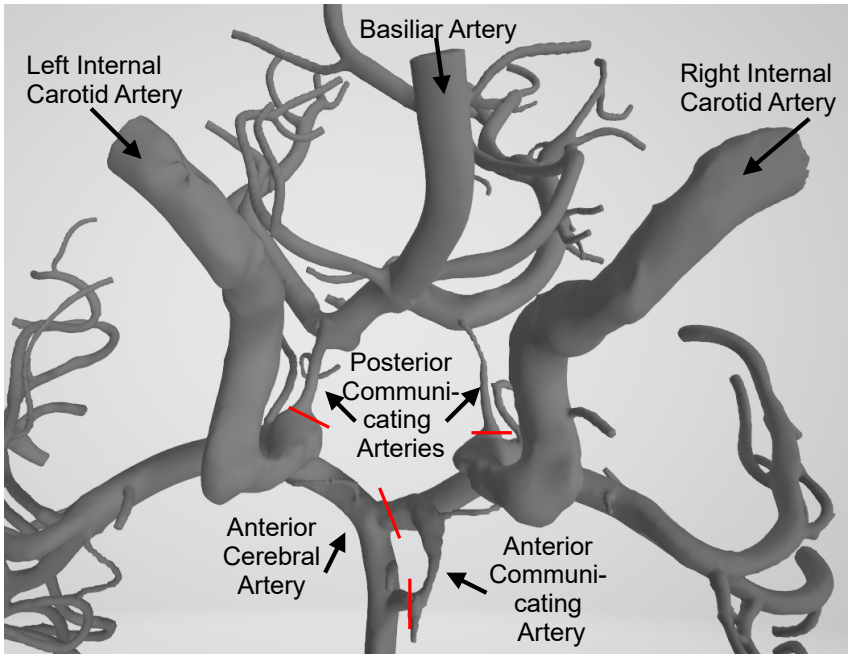

Figure 3.3. The extracted and segmented Circle of Willis from the TOF-MRI data and selected annotated arteries. Red lines highlight where the vessels were cut.

The registration algorithm was written in such a way that the user only has to select the TOF-MRI data's mesh. It will then automatically load the three centerline segments. The centerline is saved as an array of its vertices' coordinates. Using Saalfeld et al.'s [63] flow split functionality, the centerline is furthermore split into segments, which are the equivalent to the edges if the bifurcations are considered as nodes of the graph that represents the entire vessel structure. Additionally, the hierarchical relationship of the tree structure is saved, meaning for each segment its parent and children segments are noted.

3.4 Hierarchy Reconstruction

Here it is important to mention that these hierarchical structures only apply to the three subtrees of the centerline, not the vasculature as a whole, due to it having been cut apart and the centerlines having been calculated separately. The algorithm combines the three subtrees' information - points array, segments, and hierarchical relationships - as one structure, but internally the three are still separate. This is however only relevant for the hierarchical relationships, where the segments that have been cut apart do not have a parent-child relationship.

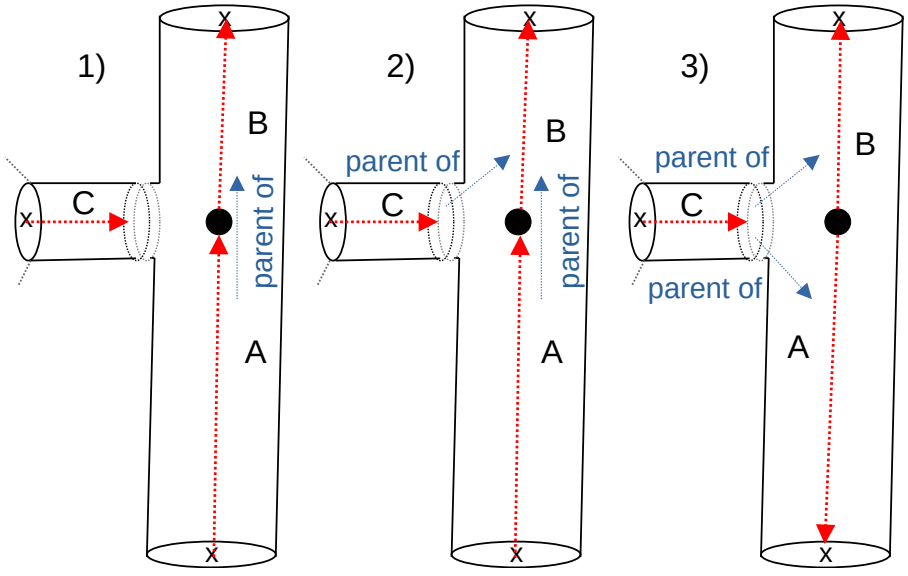

Figure 3.4. Illustration of hierarchy reconstruction. Red lines are centerlines, blue ones show parent-child relationships. C has been cut away from vessels A and B for centerline extraction and is a leaf of another vessel tree. $1)$ shows no reconstruction, C has no hierarchical relationship to A and B, while A is parent of B. $2)$ shows simple reconstruction, where B now has both A and C as parents. $3)$ shows reconstruction with flipped hierarchy, where C is now parent of A and B, and B no longer has a relationship to A.

The option to remedy the cuts and reintroduce a whole parent-child hierarchy for the entire Circle of Willis was developed. This was done by adding an additional parent and child at the respective segments where they were cut apart. An additional hierarchy reconstruction method was devised, where following the added parent-child hierarchies the hierarchies in direction of the inlet were flipped. The purpose of this second strategy is elaborated in Section 3.6.5. All reconstruction options are illustrated in Figure 3.4.

3.5 Initialisation

Before the actual co-registration can begin, the data have to be initialised. This means ensuring that both TOF- and PC-MRI data are comparable on a voxel level, and that they are as close to global maximum as possible from the get go. The first step of this is plotting them in the same coordinate system where the transformations then can be applied.
Initially, both sets of data - meaning the mesh and centerlines of the TOF-MRI data and the DICOM slices of the PC-MRI data - are plotted in a plain coordinate system as is given by Matlab.
The plotting of the PC-MRI data in the corner of the coordinate system, meaning in a positive range, is not trivial. As described above, the x- and y-dimensions have to be switched. For visibility reasons, only one axial slice is shown at first, though the others can be made visible.

In order to make the TOF-MRI data directly comparable to the PC-MRI data, its voxels have to be resized to the voxel size of the PC-MRI data. This is done by dividing each coordinate of the TOF-MRI data by the corresponding voxel size, meaning the x-coordinates are divided by the length of a PC-RMI voxel in x-dimension, yielding:

$$\frac{xCoordinate_{TOF-data}}{xDimensionVoxelSize_{PC-MRI-data}}$$

for x-dimension, and for y- and z-dimensions in the same way.
All centerline and mesh points are then translated into the middle of the PC-MRI volume. This is achieved by first translating the center point of the TOF-MRI data into the origin of the coordinate system, and then translating it into the middle of the PC-MRI data.
All needed scalars, be it voxel sizes, coordinates of the center TOF-MRI point, or center PC-MRI point, were extracted from the raw data in MeVisLab using the WorldVoxelConvert and Info panels, or they were read from the DICOM headers of the data.

Since this does not automatically yield a sufficient initial alignment for further registration, a simple 4 DOF optimisation is done for translation in x-, y- and z-direction as well as isotropic

scaling. This can be viewed as a restricted rigid registration for initialisation, and the begin of the multi-level strategy.

3.6 Optimisation

The optimisation drives the process of maximising the metric through iterative transformations, thus finding the transformation that best fits the TOF-MRI data to the PC-MRI data. The strategy used for this problem is coarse-to-fine, as well as multi-start and multi-level. It decouples the transformation parameters, starting with different parameters and a low amount of degrees of freedom and increases them. It also has a hierarchical approach that makes use of the inherent tree structure of the vessel network, and helps ensure local optimisation in addition to global.

3.6.1 Metric

The metric quantifies if a given transformation of the floating data is appropriately fitted to the reference volume data on their corresponding structures. In this case, the metric reflects how well the centerline points of the TOF-MRI data lay within the vessels of the PC-MRI data. This is executed by calculating the normalised sum of the intensities of the PC-MRI voxels in which the centerline points lie after a transformation has been applied. The PC-MRI data are pre-processed to give vessel voxels a high intensity value, therefore an optimal transformation returns a high metric and the optimisation is a maximisation problem. Furthermore, the metric is specialised through blur and weight parameters. The blurring can add to the coarse-to-fine strategy, with an initially high blur factor to eliminate termination in a false local optimum due to noise, and to then be lowered later on to increase the influence of details and further perfect the fit. A Gaussian blur was used with its kernel set to $\kappa\sigma = 2$, though due to the PC-MRI's pre-processing its influence is only mildly significant.

The weight has a similar functionality. It depends on vessel thickness and can be used to give thicker or thinner vessels a stronger impact on the metric. Thus, thicker vessels can be prioritised when the optimisation is at its almost global stage, and then heighten the import of thinner vessels when the smaller vessels are locally optimised. However, since mainly bigger vessels are visible in the PC-MRI data, in this case the weight exclusively serves to

give vessels with a higher radius a bigger impact.
First, the weight proposed by Ayward et al. [46, 52] was used:

$$w_i = \frac{2}{1 + e^{-2 \cdot r_i}} - 1 \tag{3.1}$$

where w_i is the weight of centerline point i based on its radius r. It assigns a point with radius 0 the weight 0, and a maximum weight of 1 to points with radius bigger than 3. To make the assigned weights more flexible to the whichever radii are currently given to the metric, rather than the same radius always getting the same weight regardless of overall radii distribution at the given time, and to give thicker vessels even more import, a second weighting strategy was devised. Especially in the case of highly varying radii the thicker vessels shall have a bigger influence on the metric. Therefore, using the maximum and minimum radii, and the average and standard deviation of the radii, the metric of Aylward et al. [46, 52] is only used if all radii lie within the range of 0 and 3 and have a very small standard deviation, meaning only similar radii are used given to the metric. If the standard variation is bigger or radii exceed the range of 0 to 3, the weight will simply be a rescaling of all given radii to a range of 0.1 to 10. This gives the biggest radii a significantly bigger impact on the metric than the smaller radii and is relevant for the overall global-to-local optimisation strategy, more on which follows in Section 3.6.5. Should none of the conditions be met for weights to be assigned according to the previously described ones, the raw radius will be used as weight. Exact conditions were determined empirically.
Remembering Equation 2.1, w gets replaced with Equation 3.1, T with the transformations described in Section 3.6.3 and $\kappa\sigma$ is set to 2. This yields:

$$M(T) = \frac{1}{\sum_{i=1}^{n} \frac{2}{1+e^{-2 \cdot r_i}} - 1} \sum_{i=1}^{n} \left(\frac{2}{1 + e^{-2 \cdot r_i}} - 1\right) I_{\kappa\sigma=2_i}(x_i T) \tag{3.2}$$

where x_i is a centerline point, T the transformation, I the interpolated intensity value in which the centerline point lies, n the number of centerlines points and r the radius at the given point.
The metric always traverses each centerline point it is given to evaluate. Before the corresponding interpolated intensity value in the PC-MRI data is weighted and processed by the metric, it first checks whether the centerline point lies within the boundaries of the PC-MRI data. At first, points outside of the data were dismissed and had a negative impact on the metric. This punishment was achieved by setting the corresponding intensity values to zero, so nothing would be added to the sum of intensities, but at the same time multiplying the weight by factor 10. This had a significant impact on the metric, since the sum of all weighted corresponding intensities are divided by the sum of all weights. This division is done to normalise the metric and make it comparable regardless of the number of points it

is given to evaluate.
This punishment was later removed from the algorithm, since it was determined that the TOF-MRI data covers a bigger area of the patient's brain than the PC-MRI data, meaning that some of the vessel structures are not only not visible in the PC-MRI data because they are too small, but also because they lie in an area that was not covered in the imaging process. The weight multiplication factor was thus removed from the metric. A centerline point laying outside of the PC-MRI area will have a more neutral impact on the metric, not adding to the sum of intensities but still adding to the sum of weights. This would still lead to a negative impact, however, if all centerline points from vessels not covered in the PC-MRI data were given into the metric. Measures to account for this are described in Section 3.6.6.

When talking about the centerline points the metric is given to evaluate, it is always a subsampled subset of the centerline points. A percentage parameter determines exactly how many points of the centerline is given to the metric. The subsampling strategy always takes the first and last point of a vessel segment, and then selects the rest spread evenly between them. A constraint of four centerline points minimum after subsampling has been set, which is only relevant for very short or very roughly extracted centerlines and acts as a failsafe.

3.6.2 Powell's Optimiser

For the optimiser, meaning the algorithm that finds the transformation parameters that yield the best metric for the current points, Powell's method [55] was chosen, which works according to Source Code 3.1

Source Code 3.1. Pseudo Code of the functionality of Powell's method.

```
1  while optimumNotFound
2      x_0 << starting point
3      d_1, d_2 << starting directions
4      x_1 = oneDimensionalOptimisation(x_0, d_1);
5      x_2 = oneDimensionalOptimisation(x_1, d_2);
6      conjugateDirection = directionConntecting(x_0, x_2);
7      x_3 = oneDimensionalOptimisation(x_2, conjugateDirection);
8      if x_3 == optimum
9          optimumNotFound = FALSE;
10     else
```

```
11            x_0 = x_3;
12            d_1 = d_2;
13            d_2 = conjugateDirection;
14        end
15 end
```

In Matlab, Tonel's [64] proposed code was used for Powell's optimiser, with some adaptions to fit the current use case and enviroment. It uses Coggings for the one dimensional optimisation along the search directions.

3.6.3 Transformation

Affine transformations were chosen as they are commonly used, extend Jomier's [8] and Aylward et al.'s [46, 52] approach and are used by Jenkison et al. [51]. They have 12 DOF based on 15 parameters: three translation parameters, three scaling parameters, three rotation parameters, and six shearing parameters.
The data can be translated, scaled and rotated separately in x-, y- and z- direction. For shearing, it can be skewed parallel to the x-axis in y- and z-direction, parallel to the y-axis in x- and z-direction, and parallel to the z-axis in x- and y-direction. Since for each axis shearing is possible in two directions, each have two parameters, despite only one degree-of-freedom being influenced.
Each kind of transformation has its own matrix, which are provided in Table 3.2, with x, y and z denoting the transformation parameters for the respective dimensions:

Translation	Scaling	Shearing
$\begin{pmatrix} 1 & 0 & 0 & x \\ 0 & 1 & 0 & y \\ 0 & 0 & 1 & z \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} x & 0 & 0 & 0 \\ 0 & y & 0 & 0 \\ 0 & 0 & z & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} 1 & y_1 & z_1 & 0 \\ x_1 & 1 & z_2 & 0 \\ x_2 & y_2 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$

Rotation		
x-Axis	y-Axis	z-Axis
$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(x) & -\sin(x) & 0 \\ 0 & \sin(x) & \cos(x) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} \cos(y) & 0 & \sin(y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(y) & 0 & \cos(y) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$	$\begin{pmatrix} \cos(z) & -\sin(z) & 0 & 0 \\ \sin(z) & \cos(z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$

Table 3.2. Translation, scaling, shearing and rotation matrices.

Extension from $\mathbb{R}^3$ to $\mathbb{R}^4$ by homogeneous coordinates is done in order to write translations, which are not linear transformations, as a matrix too [65]. This also means that any point $p = (x, y, z)$ that shall be transformed via the matrices is extended by a fourth coordinate, the ω-coordinate, as well:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} \Rightarrow \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

The ω-coordinate is set to 1, as the coordinates represent points, not directions, in which case ω would be set to 0.

Rotation is expressed as three separate rotation matrices, one around each axis. Rotation parameters x, y and z are Euler angles and are always multiplied in the same order, as a different order of rotation matrices would change the transformation result [65]. This also applies to transformations in general, so the affine transformation matrices are always used in the same order.

All transformations are executed around the origin of the world coordinate system. That means that simply applying the transformations to local coordinate system of the TOF-MRI data it would be transformed entirely out of frame of the PC-MRI data, as it is necessary to express any transformation in world coordinates. In order for the transformation to be correctly applied to the data that shall be transformed, is has to first be translated into world coordinate origin. There the transformations will be applied before it shall be translated back to its original position. Figure 3.5 illustrates how the result of a rotation can differ if the object that shall be rotated is moved into origin first.

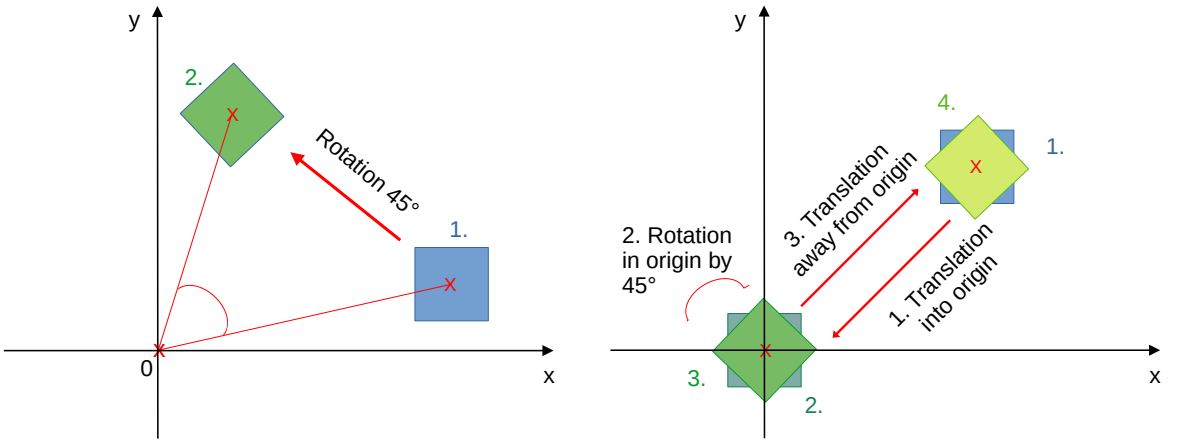

(a) Simple 45° application of the rotation.

(b) The same 45° rotation after translating into origin and then translating back.

Figure 3.5. Difference of translating data into world coordinate origin before rotation.

The point that lies in the origin is the one that after all transformation are applied will not have moved, and therefore has to be chosen carefully. For global transformation, the center point of mass can be chosen. For the later described Hierarchical Tree Traversal strategy in Section 3.6.5, the point that connects parent and child segment makes more sense.
Transformations were constrained to only be valid if their parameters lay within a certain range. This was done to prevent errors like scaling the entire dataset so small that it fits within one vessel, thus maximising the metric, or rotating a vessel to lay parallel to its parent vessel and thus be optimised to the same intensity ridge. Transformation parameters were thus constrained according to Table 3.3. Constraints were determined by empirical observations.

Transformation	Range
Scaling	0.8 - 1.3
Translation	-30 - 30
Rotation	$-0.5 dimensionSize$ - $0.5 dimensionSize$
Shearing	-0.5 - 0.5

Table 3.3. Range wherein the transformation parameters must lie to be valid.

3.6.4 Optimisation Strategy

For data as complex and high-resolution as the 7T TOF-MRI data of the intracranial vessels to be co-registered to the far less detailed 7T PC-MRI data, one time optimisation over all transformation parameters using Powell's optimiser does not yield sufficient results. A more complex strategy is needed.

One of the key issues to avoid in all optimisation problems is the avoidance of local extrema, which might cause the strategy to terminate prematurely on the way to the global optimum. One strategy to minimise the possibility of optimising towards a local extremum is a coarse-to-fine approach. This means to first resolve for a rougher version of the input data, using less DOF and less precise parameters, and refining these and others over the course of the optimisation process. Each step of refinement therefore already solves for an optimum, so by the time the most precise step is reached, the parameters are assumed to already be close to the global optimum. This strategy has the additional benefit of being more efficient, since fewer and less precise parameters are faster to solve, and by the time the more expensive steps are reached, their closeness to the global optimum ensures that resolving them will not be as expensive as starting from figurative zero.
This coarse-to-fine approach is already partly realised through the number of points given to the metric. Before optimisation, the centerline points are resized, or rather, subsampled.

A certain percentage of them is chosen, rather than taken all of them, effectively speeding up computation.
The basic concept of the core strategy chosen comes from the related paper of Jenkinson et al. [51]. The key concept of their idea is based on the fact that (affine and rigid) transformations are heavily dependent on rotations. Even a slight difference in rotation can change the result of the entire transformation. Therefore, the rotation angles are decoupled from the rest of the parameters and resolved first.
A rough grid over a number of rotation angles is created for a multi-start approach. The angles are determined by parameters $stepsize$ and $steps$. The number of angles to resolve for in each dimension is determined by $steps$, while $stepsize$ denotes the interval between the $steps$. The $steps$ with a distance $stepsize$ between them form a grid of points. Each point will be assigned three rotation angles based on the $steps$ and $stepsize$ parameters. After applying the rotation to the data, a 4 DOF optimisation is done to find the three best translation and one best isotropic scaling parameters for each point in the grid. After each grid point has a set of seven parameters - the rotation, translation, and scaling ones - as well as the resulting metric value, the grid is refined: one new grid point between each two old ones, as well as in the center of each cell of four points. Where the initial rough grid had a size of $steps^3$, the fine grid has a size of $((2 \cdot steps) - 1)^3$.
For each new grid point, the rotation and translation parameters are interpolated from the surrounding rough grid points, with the exception of the scaling parameter: it is universally set to the median of the rough grid point results. This is sufficient as scaling is decoupled from rotation [51]. Using these interpolated parameters, the data is transformed and the metric evaluated. Note that there is no optimisation done in this step.
From this refined grid, the three parameter sets that have yielded the best metric value are chosen. Each of those three sets is then perturbated. Perturbation parameters are dependent on resolution of the PC-MRI data and voxel size. This results in seven new sets of parameters per best metric value. These total of 24 sets of parameters are then used as starting transformations for 7 DOF optimisations: three translation, three rotation and one isotropic scaling parameter. The set with the best metric result is then chosen as starting parameters for a final 12 DOF optimisation, with three translation, three rotation, three scaling and six shearing parameters.

This strategy has the additional benefit that when the optimisation returns an invalid metric, it can still take the best result from the previous steps to transform the data. An invalid metric can be reached when transformation is too strong, transforming the data out of range and thus eliminating all overlap, or when transformation parameters become too small or infinite or outside of the constraint range.

3.6.5 Hierarchical Vessel Tree Traversal

Given the stark difference in detail between TOF- and PC-MRI data, a global co-registration, even with the decoupling of the parameters and a coarse-to-fine strategy, is not applicable. The big center vessels might be registered well, but the further away from the center the vessels are, the worse the quality in the PC-MRI data, and consequently the quality of registration among the smaller vessels decreases. Therefore a hierarchical vessel tree traversal strategy was devised based on related work by Jomier et al. [8], who in turn based their work in Ayward et al.'s [46, 52] works, who also inspired the previously explained metric in Section 3.6.1. The entire registration pipeline is illustrated in Figure 3.6.

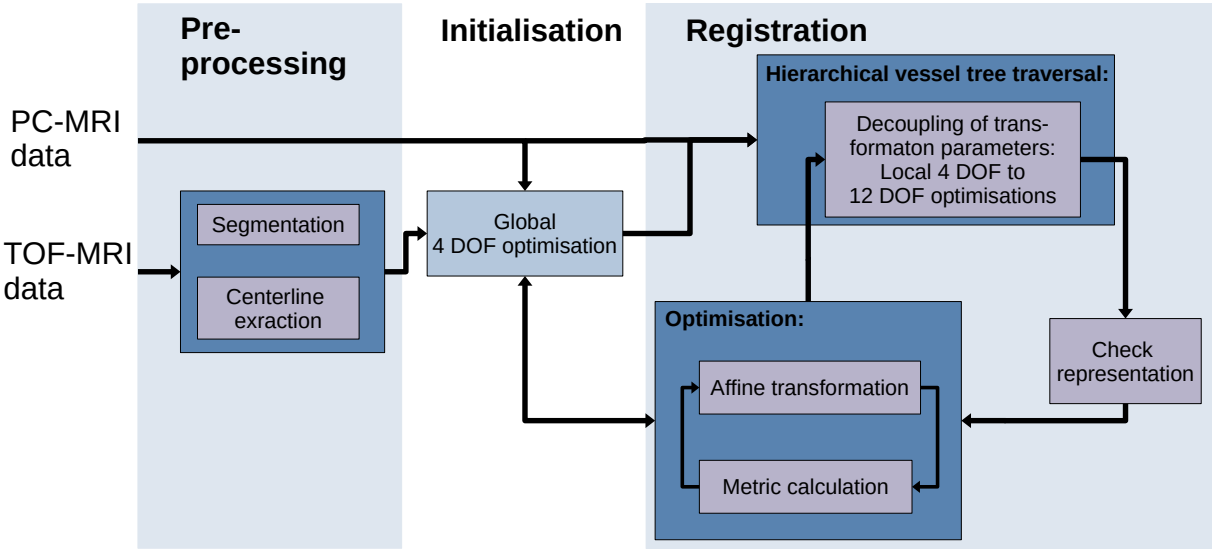

Figure 3.6. Registration pipeline.

Making use of the inherent children and parent tree graph structure of the vessel network, the root segments of the vessel tree are found. All children and children's children to the leaves are added, and for the entire tree from root to leaves the optimisation strategy as described in Section 3.6.4 is applied, and the entire tree is transformed accordingly. After that, each of the immediate children of the root is treated as a new root of a tree spanning its children and children's children to the leaves, and optimised and transformed again, but this time excluding the parent of the new root.

Here the weighting according to radius becomes important. It can be assumed that each of the root segments is bigger than its children, so the weight, which gives vessels with a bigger radius more impact on the metric, assures that the optimisation is always in favour of the current root segment. This way, the vessel structure is optimised from its inlets to the leaves, with the larger root vessels influencing the transformation of its smaller children,

but not the other way around. Here the assumption is that if a parent vessel is correctly aligned, its children vessels will be closer to their corresponding vessels as well, as they are all transformed according to the metric on which the root vessel has the biggest impact. If the assumption holds true, this way even smaller vessels that are not visible in the PC-MRI data will be aligned.

Important here is that the vessels have to stay together - the tree structure has to be preserved, there shall be no gaps or disconnections in the vessel structure that has not been there before. This means that aside from the first root segments, meaning the inlets of the observed vessel structure, no translations are allowed. The point that connects the vessel to its parent must remain in the same position, even as the subsequent points shift. To achieve this, that connecting point becomes the central origin point for the transformations, meaning the point that will be moved into origin for transformations, as described in Section 3.6.3 and illustrated in Figure 3.5. Therefore all branch points that are not inlets can only be rotated, stretched and scaled, and all transformations of the branch points lose the 3 DOF from the translation.

Also important to keep in mind is that in order for centerline to be correctly extracted, the Circle of Willis was split into three subtrees to eliminate cycles. This yielded three separate trees with separate roots and no connections between children. Without constraints, the above described optimisation strategy could lead to these three trees to become separated, which is also not wanted. One measure to pre-empt this was to reconnect the three centerline structures by adding the missing parent-child relationships where the vessels were cut. This however leads to some of the segments to be optimised twice, since they will have two parents. A constraint for each segment to only be optimised once can be added, but it will thus only be optimised according to the one parent vessel through which it is reached first. Furthermore, this does not entirely solve the problem, as this can also lead to the tree to be torn apart, meaning that if the subtree is transformed according to one parent vessel, it might be transformed away from the second parent vessel.
Ultimately the issue is that there are three separate trees, meaning three separate inlets or roots from which optimisation can start. To completely combine the tree, the third reconnecting strategy from Figure 3.4 was devised. There the two other roots can be turned into leaves by flipping the parent-child hierarchy after closing the gaps in the hierarchy created by the cuts for centerline extraction. This does not render the roots, which are now leaves, to have a minor influence, since they still have big radii and thus a bigger influence in the metric.

During testing, it was observed that this strategy can lead to rather strong distortion of the smallest vessels, presented in Figure 3.7. This is assumed to be because the smallest vessels will be transformed in almost every step, since a lot of the bigger vessels will eventually

lead to them via the parent-child hierarchy. So even if transformation parameters for each subtree are small, they can add up in each iteration. To counteract this, it was implemented that the part of the subtree that is not represented will only be transformed in a limited manner, excluding shearing parameters which are suspected to have the biggest effect on the observed distortion. What it means for a segment to not be represented is elaborated in the following Section 3.6.6.

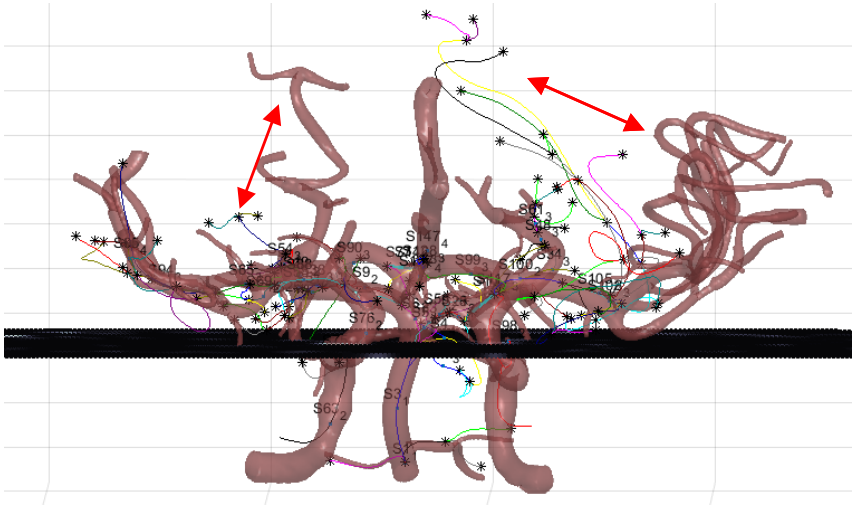

Figure 3.7. The mesh shows TOF-MRI data before registration, the colourful centerlines TOF-MRI data after registration. Only centerline segments with a label are represented. Red arrows highlight the strong transformations the smaller vessels underwent. The black volume is one slice of the PC-MRI intensity volume.

3.6.6 Representation

Up until now, the optimisation process was described as being based on a subgroup of all centerline points of whichever segments are handed into the optimiser. However, another aspect to take into consideration is the lower grade of detail of the PC-MRI data, which results in significantly less vessels being visible in comparison to the TOF-MRI data. This means that for a lot of the centerline segments, particularly smaller ones in the periphery, there is no corresponding intensity ridge visible in the PC-MRI data. Therefore those centerline segments will always be registered to a structure that does not actually correspond to it.

This is why the concept of representation was introduced. Each segment is flagged as either being represented or not, meaning if the corresponding vessel is visible in the PC-MRI data or not. This flag was called representation and was constructed as an array set to 0 or

1 at the segment index based on the respective segment's representation status. This array can then be used during the optimisation, where a segment is henceforth only included in the metric calculation steps if it is represented.
Non-represented segments will still be transformed according to their parent vessels, but they themselves and their children will not undergo the optimisation steps and be assigned their own transformation parameters, and they will not be considered during metric calculation. Since the TOF- and PC-MRI data are derived from the same patient, meaning both show the same vessels, it is assumed that if the visible parent vessels are aligned, the smaller child vessels, who will have been transformed along with their parents, are aligned as well, or at least closer to be aligned than previously.

To calculate representation of the vessels, first the Strahler order of the segments was determined [66]. Strahler order can be used to not only assess hierarchical structure of asymmetrical trees like vasculature, but also to assign vessels of the same approximate size the same level, despite solely relying on the topology of a given tree [67, 68]. This is done bottom up, as illustrated in Figure 3.8 and in Source Code 3.2. Leaves are assigned level one, and from then upwards at each bifuraction the order of all children of the parent are compared. Do all children have the same level, the parent gets a higher level. If children do not all have the same level, the parent gets the highest level among the children.
In the present case, the assigned levels of the Strahler order were then flipped, so that the leaves get the highest level and the previously highest level segments, which presumably would be the inlets, will be level one.

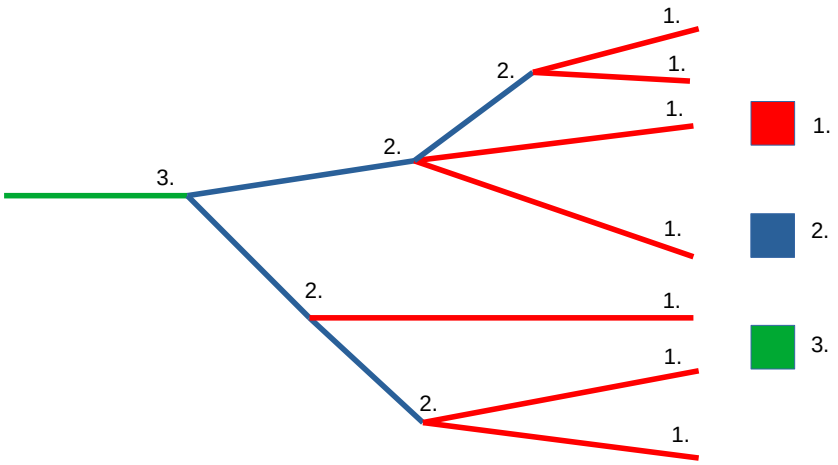

Figure 3.8. Example tree structure with each segment's Strahler order level.

Based on that level as well as vessel radius, the representation is determined. Vessel radius is needed in additon to Strahler Order, as vessels with a big radius are visible in the PC-MRI data and thus are definitely represented, regardless of the level according to Strahler order. The cutoff radius was determined empirically and finally set to 1. Throughout

testing, it also became apparent that the use of the vessel radius rendered Strahler order redundant, as all vessels with the needed radius to be represented also had a low Strahler order level, but not the other way around, which was why the radius was required in the first place.

Source Code 3.2. Pseudo Code of assignment of Strahler order.

```
1  leaves << leaves of a tree
2  level(leaves) = 1;
3  currentParents = parents(leaves)
4  while rootIsNotReached
5      for i = 1:number(currentParents);
6          currentChildren = children(currentParent_i);
7          if all currentChildren levels are the same
8              level(currentParent_i) = level(currentChildren) + 1;
9          else
10             level(currentParent_i) = max(level(currentChildren));
11         end
12     end
13     currentParents = parents(currentParents);
14 end
```

3.7 Mesh Transform

After the centerline is transformed according to the optimised metric, the mesh shall be transformed as well. In the initialisation step, correspondences to the centerline for each mesh point was found.

Testing has determined that segment correspondence is sufficient, since all centerline points within a segment get the same transformation matrix. Therefore, each mesh point needs to be assigned its closest segment.

To find the closest segment, each segment is represented through its first, last and middle point (Figure 3.9(a)). For each mesh vertex, the three closest of those representative segment points are found (Figure 3.9(b) and Figure 3.9(c)). Then the mesh vertex is compared to every centerline point within those three closest segments (Figure 3.9(d)), and the segment to which the closest centerline point belongs will be assigned to the mesh vertex for transformation (Figure 3.9(e)).

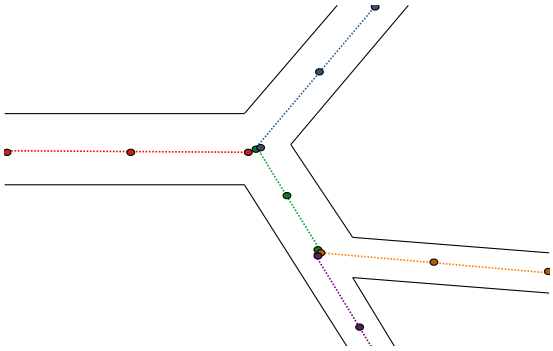

(a) Every centerline is represented by its start, end and middle point.

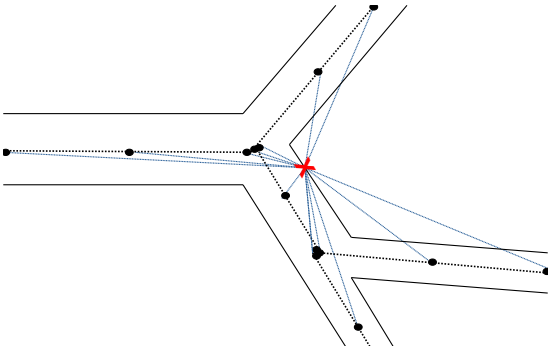

(b) Distance of a mesh vertex (red X) to each centerline representation point is calculated.

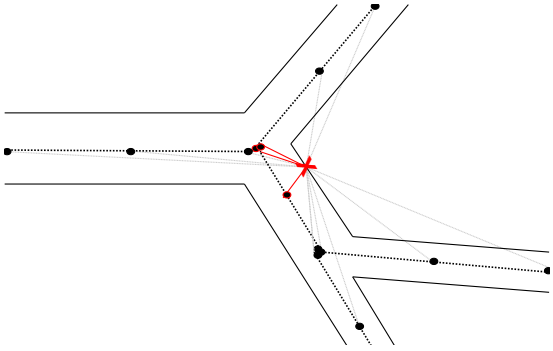

(c) The three closest centerline representation points to the mesh vertex (red X) are found.

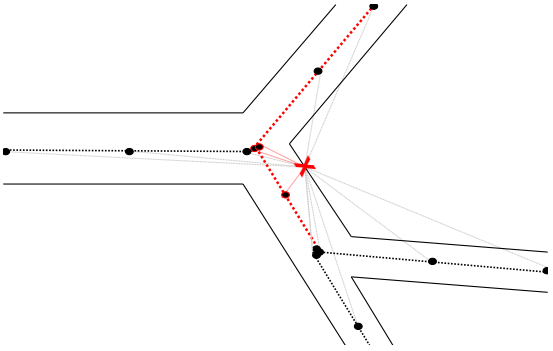

(d) The corresponding centerlines are found, and the distance of the mesh vertex to every centerline point among the closest ones is caclulated.

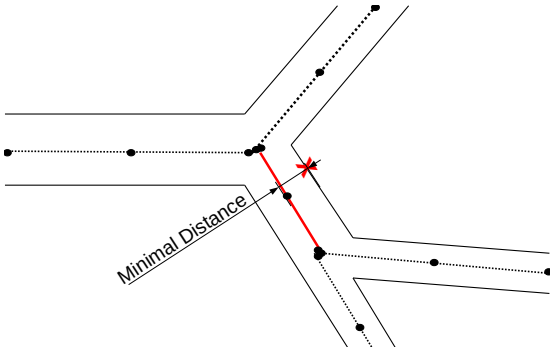

(e) Mesh vertex (red X) is assigned to the segment on which's centerline the closest point with the minimal distance lies.

Figure 3.9. Illustration of mesh vertex assignment to vessel segments according to closest centerline points, for later mesh transform.

This is significantly faster than naively calculating the distance between every mesh vertex and every centerline point. While distances between mesh vertex and each point within the three closest segments are computed, in practice a lot of vertices only have one or two closest segments, since every segment is represented by three of its points. It is thus

possible that the three closest points to a mesh vertex all belong to the same segment.
To find closest points, their distance was calculated via Euclidean distance.
Despite the precaution of taking the assigned segment from the closest point among the three closest segments, some mesh vertices were misassigned. To prevent this, after assigning closest points each point's assigned segment is compared to the surrounding point's assigned segments. The surrounding points are found by finding all triangles that contain the current vertex. All other points that make up those triangles are considered surrounding points. The assumption is made that points that are close to each other have the same assigned segment, and that within a neighbourhood the same assigned segments are found. Therefore, if a point has an assigned segment that none of its surrounding points have assigned, it will be assigned the same segment as the majority of its neighbours. It shall be noted that the number of points that do not share their assigned segments with at least one neighbouring point is smaller than 1%, and goes as low as 0.07% and 0.03% in test data runs.

4 Evaluation

The presented co-registration approach was tested on three sets of data: the clinical data described in Section 3.1, and two sets of synthetic data.
The first set of synthetic data is based on the original clinical TOF-MRI data. First, Gaussian noise was added before artificially downsampling the DICOM slices. The data were then slightly rotated and shifted. These distorted TOF-MR images were segmented in a similar manner of the clinical data described in Section 3.2, just without the additional masking, and for validation they were registered to the clinical PC-MR images. This set of data is compared to the clinical data in Figure 4.1. Due to the downsampling, the synthetic data has less small vessels. The posterior communicating arteries for examples were completely missing.

(a) Clinical data.

(b) Synthetic data based on clinical.

Figure 4.1. The segmented clinical dataset compared to the segmentation of the synthetic data that were based on the clinical one.

The other synthetic dataset is a segmented partial Circle of Willis with a simpler vessel network, as well as an aneurysm. Centerlines were extracted using VMTK, and a binary DICOM dataset was created to act as replacement of the PC-MRI data and for it to be co-registered to. Afterwards, the segmented vessels were shifted slightly, with single segments

and or subtrees being rotated or scaled. Of the two thusly created sets of data, one was slightly transformed (Figure 4.2(b)), and the other more heavily (Figure 4.2(c)), to examine the limits of the proposed registration algorithm.

(a) Original. (b) Slightly distorted. (c) Strongly distorted.

Figure 4.2. The three sets of synthetic data, (a) is the original data, (b) is slightly distorted and (c) is strongly distorted.

4.1 Landmark-based Evaluation

The proposed registration was evaluated via a landmark evaluation based on clinical and synthetic data.

In the clinical data set and the synthetic one created based on it, thirteen landmarks were set manually in both PC-MRI and TOF-MRI data. Bifurcations were selected as anchors for the landmarks, as they are distinct. Furthermore, landmarks were only set in represented vessels, meaning they are visible in the PC-MRI data. The other set of synthetic data received seven landmarks. Landmark positions were discussed with a second reviewer with a broader experience in segmentation and editing of vascular structures.

Based on the landmarks, the mean squared error (MSE) and Hausdorff distance (HD) before and after registration were used as evaluation measures. MSE is the average of squared distances of all landmark pairs:

$$MSE = \frac{1}{n}\sum_{i}^{n} distance(landmarkPCMRI_i, landmarkTOF_i)^2 \tag{4.1}$$

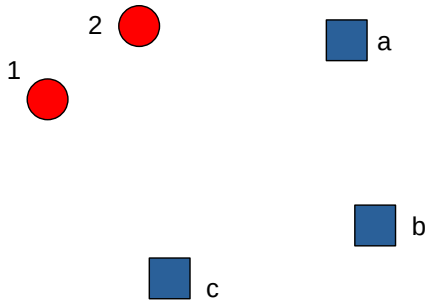

(a) Two sets of landmarks, red circles 1 and 2, and blue squares a, b and c. Find minimum distances between each pair of landmarks

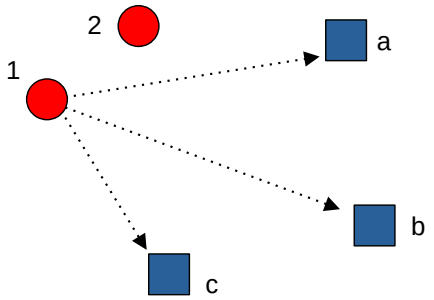

(b) Find nearest blue square for circle 1.

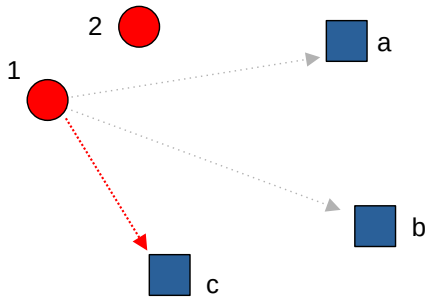

(c) Blue square c is closest to red circle 1.

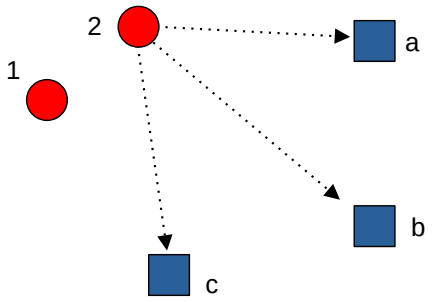

(d) Find nearest blue square for circle 2.

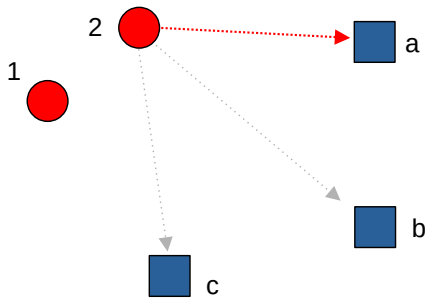

(e) Blue square a is closest to red circle 2.

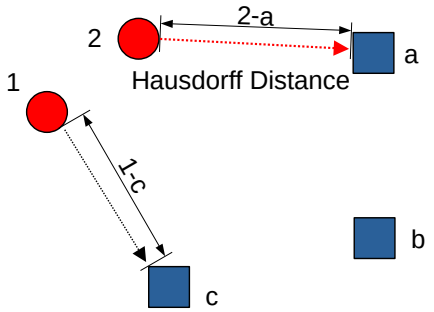

(f) Find maximum distance among the minimal distances between landmark pairs. Distance 2-a is larger than distance 1-c, therefore the former is the Hausdorff distance.

Figure 4.3. Illustration of Hausdorff distance calculation between red circle set of landmarks and blue square set of landmarks.

HD is the maximum distance of all minimal distances between landmark pairs:

$$HD = max(min(distance(landmarkPCMRI_i, landmarkTOF_j))) \tag{4.2}$$

To calculate HD, for each landmark the closest landmark from the other set is found (Figure 4.3(a) to Figure 4.3(e)). Among these closest distances, the maximum is found and used as return value (Figure 4.3(f)). This value can be seen as the amount of mismatch between the landmarks and is considered to be robust to small errors of positions [69]. For any distances that needed to be calculated, Euclidean distance was used.
These results were then compared to the MSE of the same landmarks, but resolved via Horn's [70] "closed form solution of absolute orientation using unit quaternions" (CFS), implemented by Wengert et al. [71]. CFS gives a rigid transform to fit one set of points to a second set, without assumptions about correspondences. The global rotation is solved via unit quaternions and represented via eigenvectors. After rotation is found, global scale and translation can easily be determined, the latter via the difference in the point sets' centroids. Since CFS solves in one step, it is not iterative and therefore not dependent on a good initial alignment [70]. CFS was applied once with the landmarks before registration as input, and once with them after registration as input, which was done for each dataset.

These error measures are all given in voxel lengths of the PC-MRI dataset, where in the clinical dataset one voxel side equals $0.640625mm$.

For evaluation, sampling size was chosen to be set to 20%, meaning for metric evaluation 20% of any given centerline points are used. For the optimisation strategy, parameter $steps$ was set to 11 in order to achieve minimum resolution for good results according to Jenkinson et al. [51], and $stepsize$ was set to 3.

The proposed algorithm's runtime depends highly on the selected data and parameters. For the clinical data described in Section 3.1, co-registration takes approximately $1087seconds$ on average with a standard deviation of $152.7s$. The second set of synthetic data took approximately $246s$ to register, while the first set of synthetic data that was based on the clinical data took $670s$. In comparison, when choosing $steps = 5$, clinical data takes only $267s$ on average and the first set of synthetic data only $53s$.
All tests were run on a computer with an Intel core i5-3470 CPU @ 3.20GHz processor, 8GB working RAM and an AMD Radeon HD 7800 series graphics card.

4.2 Results

The evaluation measures are first compared based on method, meaning the value before registration will be compared to the value after registration. This is first done separately for raw landmarks and for CFS, and then those two methods are compared.

4.2.1 Intra-Method Evaluation

The clinical data was evaluated once for each of the three different states of hierarchy reconstruction: one where the three subtrees remained separate, one where the hierarchy was reconstructed by adding parent-child relationships where the original vessel tree was cut apart before centerline extraction, and one where hierarchy was reconstructed and then flipped after the reconnection so that there is only one inlet, turning the two other inlets into leaves, as elaborated earlier. Each of these three reconstruction states was also evaluated once with non-represented segments being transformed same as the root segment, and once where non-represented segments were only transformed with a limited transformation, as elaborated in Section 3.6.5.

All results of the clinical data evaluation can be seen in Figure 4.4.
Figure 4.4(a) shows that the metric of the clinical data was always maximised. The factor by which it increased ranges from 2.25 in the no-hierarchy-reconstruction and non-represented vessels only receiving limited transformation dataset to 11.79 in the flipped hierarchy with all segments receiving same transformation dataset. On average, the metric increased by a factor of 5.95, so almost six times the original metric, which was 1.71.
MSE before registration was 309.07 and was decreased by registration to about 75% of the original value, meaning around 232.7.
HD meanwhile was lowered in all datasets, except the one where hierarchy was reconstructed with non-represented vessels receiving limited transformation. On average, the original value of 19.27 before registration was decreased to 15.6 after registration with a standard variation of 3.5. This means an average change to 80.9%, ranging between 55.7% and 112.4%. In the one case where HD was increased, it increased to 21.66.

Results of CFS using non-registered landmarks as input were always better than using the landmarks after registration as input, though both results were always better than the values based on the landmarks. MSE increase factor lies between 8.98 for datasets with no reconstruction and 14.51 for datasets with reconstruction, increasing the value of 16.93 using initial landmarks to 187.05 on average using registered landmarks, with a standard deviation of 46.75. HD looks similar, with the initial 8.91 being increased to 16.13 on average, with a standard deviation of 2.56. Average increase was to 181%.

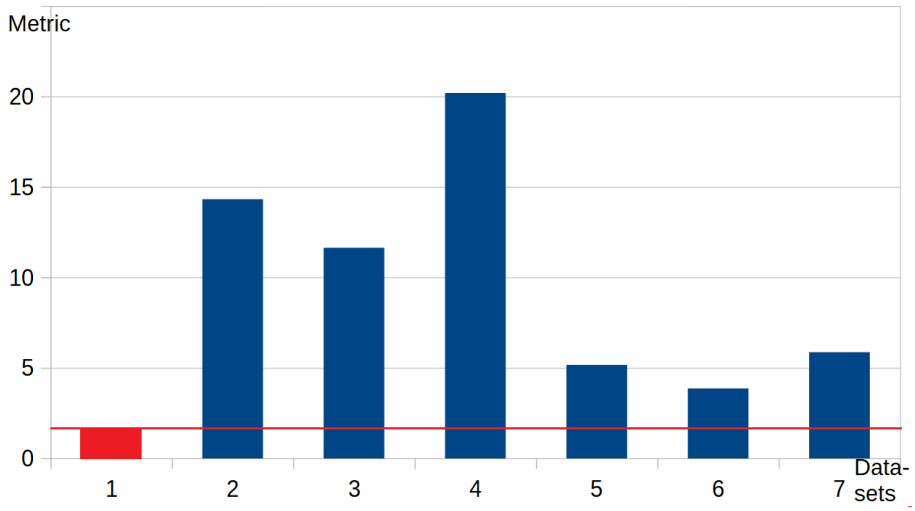

(a) Metric before and after registration of clinical data. Values bigger than the red line are wanted.

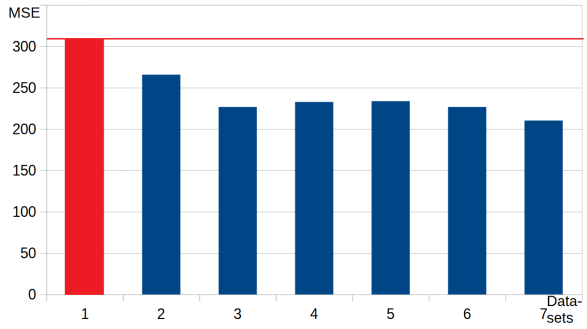

(b) MSE before and after registration of clinical data. Values smaller than the red line are wanted.

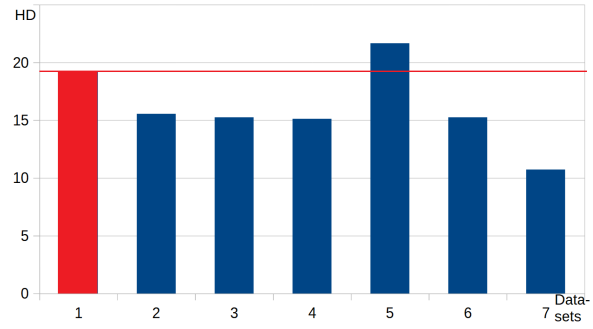

(c) HD before and after registration of clinical data. Values smaller than the red line are wanted.

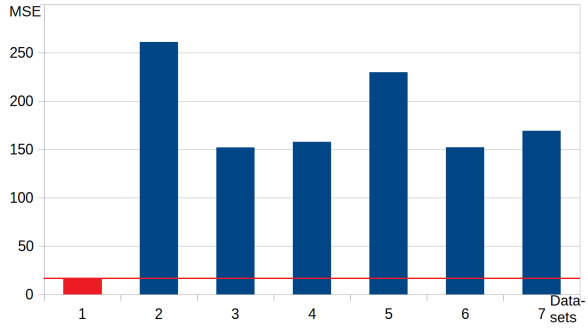

(d) MSE of CFS before and after registration of clinical data. Values smaller than the red line are wanted.

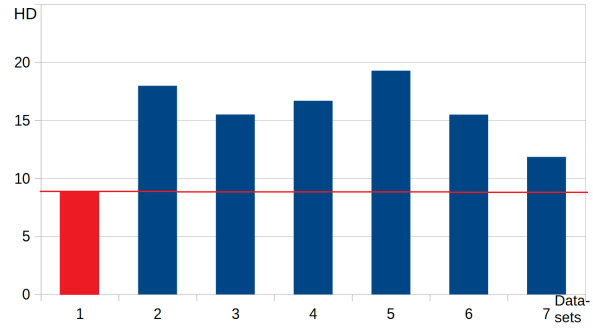

(e) HD of CFS before and after registration of clinical data. Values smaller than the red line are wanted.

Figure 4.4. Results of the evaluation of the clinical data. 1, in red while the rest is blue, always represents the result of the measures before registration, with a red line highlighting its result for comparison. 2 had vessel tree hierarchy reconstructed. 3 had no reconstruction of hierarchy. 4 had vessel tree reconstructed, as well as hierarchy flipped. 5 has reconstructed hierarchy, and non-represented segments had limited transformation. 6 had no reconstructed hierarchy and non-represented segments had limited transformation. 7 had reconstructed hierarchy as well as flipped hierarchy, and non-represented segments only received limited transformation.

The synthetic data based on the clinical data was evaluated once, as seen in Figure 4.5. MSE decreased from 306.4 to 74.89, which is 24.4% of the value before registration. HD decreased to 85.9%, from 19.27 to 16.56.
Here, too, CFS increased MSE by a factor of 3.18 when comparing landmarks before and after registration used as input. Initial MSE of CFS was higher than that of the clinical data with 17.34, but increased to 55.2 when using the registered landmarks as input. HD of CFS also increased to 139.2% of the original value, raising 8.96 to 12.47.
On the other hand, registration increased the metric significantly by a factor of almost 27, from 1.97 to 52.9.

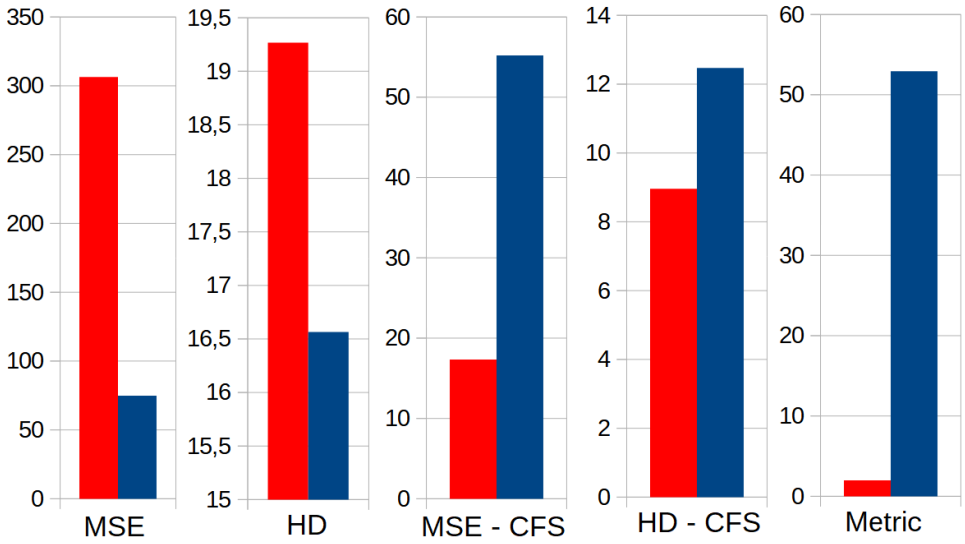

Figure 4.5. Results of MSE and HD for the synthetic data that is based on the clinical data, as well as MSE and HD for the CFS of that data, and the metric results. The red left column shows result before registration, the right blue column shows result after registration.

There were three sets of synthetic data not based on the clinical data that were evaluated: one was the original, unedited synthetic data, one was slightly distorted, and one was more strongly distorted. The results are illustrated in Figure 4.6.
Since here the basic centerlines were different rather than the workflows of how the registration handled them, the values of the evaluation measures before registration differ along with the ones after registration. Overall, results for these three synthetic sets of data were quite different. The metric increased by a factor of 3.84 on average, with a standard deviation of 3.82, factor thus ranging between 1.38 from the original data and 8.24 from the strongly distorted dataset.
Counterintuitively, only the data with stronger distortions had positive MSE and HD results. It decreased the MSE from 4394 to 1510.59, a decrease to 34.38%, and HD by a factor of 0.78, from 87.41 to 68.31. For the other two datasets, registration increased both MSE and

HD. HD increase for both was an approximate factor of 3.03, while MSE increased by 3.35 for the mildly distorted dataset, and by 4.1 for the original dataset.

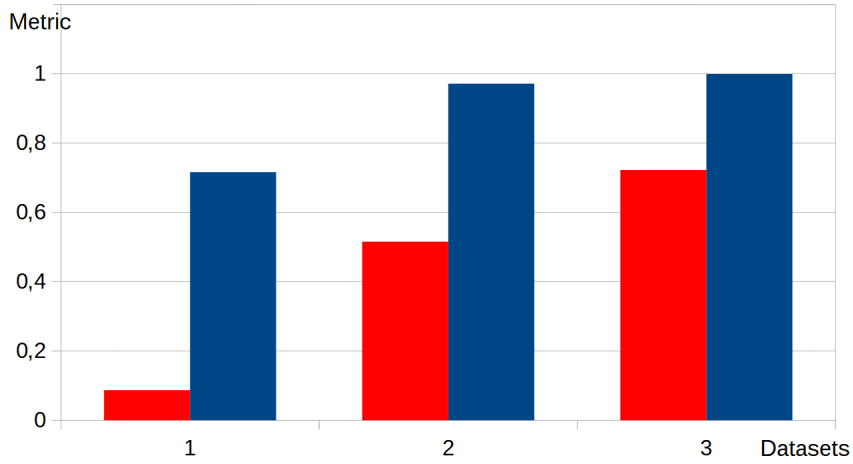

(a) Metric before and after registration of synthetic data. Blue values bigger than their red counterparts are wanted.

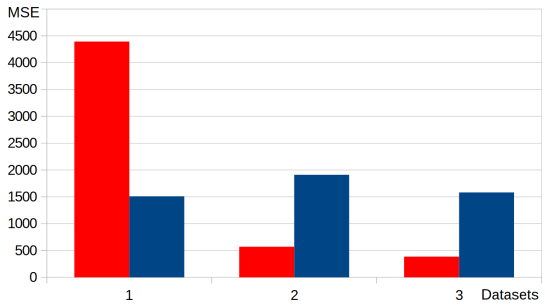

(b) MSE before and after registration of synthetic data. Blue values smaller than their red counterparts are wanted.

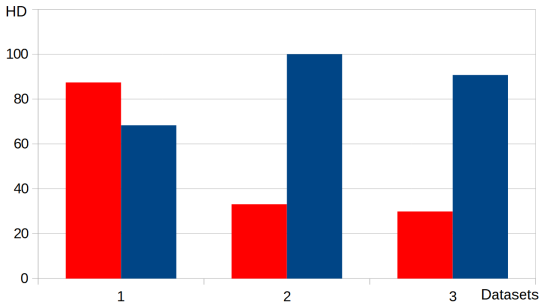

(c) HD before and after registration of synthetic data. Blue values smaller than their red counterparts are wanted.

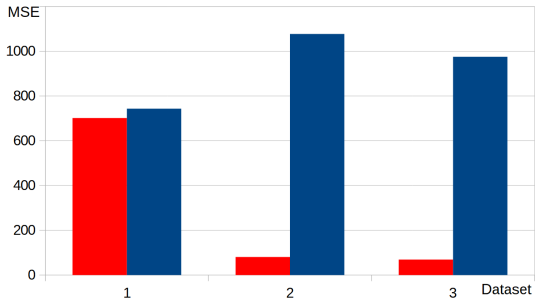

(d) MSE of CFS before and after registration of synthetic data. Blue values smaller than their red counterparts are wanted.

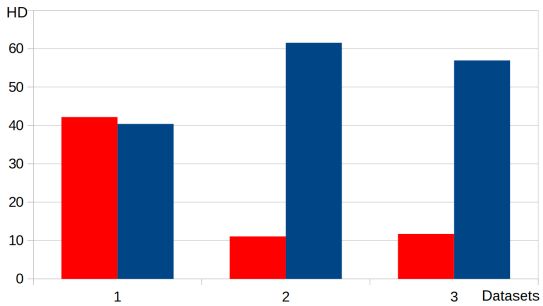

(e) HD of CFS before and after registration of synthetic data. Blue values smaller than their red counterparts are wanted.

Figure 4.6. Results of the evaluation of synthetic data not based on the clinical data. Red column represents the result before registration, blue column the result after registration. 1 is the strongly distorted data. 2 is the slightly distorted data. 3 shows the result for the non-distorted data.

CFS has negative results when comparing landmark input prior to and after registration, except for the strongly distorted dataset, where using the registered landmarks improved MSE by a factor of 0.96, lowering it from 42.17 to 40.37. On average, MSE increased to 951.8% and HD to 379.2%, though both CFS results improved the ones of the landmarks themselves, both before and after registration.

4.2.2 Inter-Method Evaluation

Thus far the results of the evaluation measures were only compared within method before and after registration. Now the results from the landmarks themselves will be compared to the results of the CFS.

For the clinical data, before registration CFS naturally achieved much better results. Initial MSE of the landmarks was 309.07, which CFS reduced to 16.96 on average, a decrease to 5.5%. HD was not decreased quite as dramatically, lowering to 46.26%, from 19.27 to 8.91 on average.
Using the landmarks after registration as input resulted in more mixed numbers. For MSE, CFS achieved better results for all data lowering it to 79.83% of that of the landmarks. When CFS used the registered landmarks as input, it often resulted in a worse HD than the one from the registered landmarks themselves, increasing it by an average of 1.05. The one exception is the dataset without hierarchical reconstruction and limited transformation of non-represented vessels. This, however, is only because this was also the only dataset where registration itself resulted in a heightened HD of the landmarks. CFS there still has a minimally higher HD than the initial landmarks that had not been registered yet. So while for all datasets the HD of CFS using the registered landmarks was worse than HD of the landmarks themselves, both HDs are still lower than those of the landmarks without registration.

For the synthetic data based on clinical data, CFS performed better than registration for the landmarks. Upon initialisation, landmarks show an MSE of 306.4. Registration reduces this to 74.89, while CFS further reduces it to 17.34. Using the position of the landmarks after registration as initial positions, CFS results in an MSE of 55.2, which is approximately 73.7% of the registration landmark MSE. HD is reduced to 46.5% before registration, and 75.28% after registration.

When converting the results for which clinical data was used into millimeters, average MSE of the landmarks is $134.63mm$ and average HD is $10.08mm$. For MSE, CFS achieved a better average result, though it varies wildly if landmarks before or after registration are used as input. Using landmarks before registration, average MSE for CFS is $10.88mm$,

while using landmarks after registration yields $107.763mm$ MSE. The resulting HD using landmarks before registration as input for CFS is $5.71mm$, while using landmarks after registration gives a bigger HD, namely $9.99mm$. The results are graphically compared in Figure 4.7.

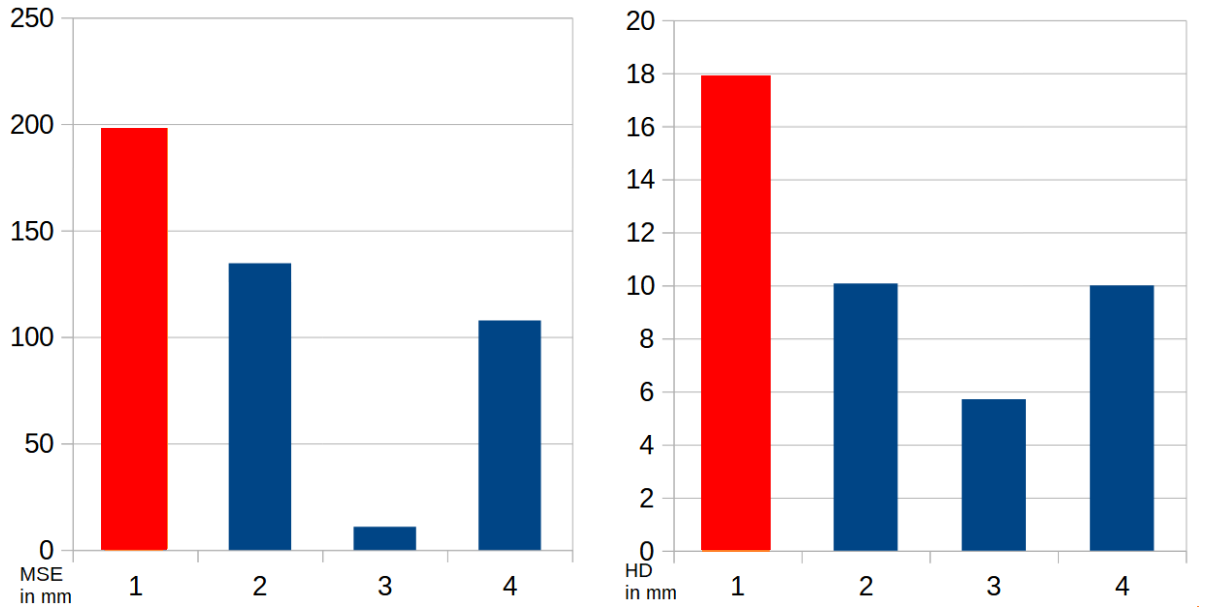

(a) MSE before and after registration of clinical data. Blue columns with smaller values than the red column are wanted.

(b) HD before and after registration of clinical data. Blue columns with smaller values than the red column are wanted.

Figure 4.7. Average error measures in millimeters. Red column 1 represents the result before registration. Blue columns represent the result after registration. 2 is the result of the landmarks after registration. 3 is results of CFS with landmarks before registration as input. 4 is results of CFS with landmarks after registration as input.

CFS resolved the landmarks better than the registration could in all cases for the synthetic data not based on clinical data, however, it once again differs in whether results are better using landmarks' initial position or the ones after registration as input.
Using the inital positions as input, CFS achieved a MSE that is smaller than initial landmark MSE by factor 0.16 on average. Using the registered landmark positions as input, MSE was smaller by a factor of 0.56. HD was decreased to 40.3% on average using landmarks before registration and to 61.1% using landmarks after registration as input when comparing CFS result to landmark result. Curiously, the more strongly distorted dataset performed best in all areas, even better than the original dataset.

4.2.3 Evaluation Overview

Of the clinical data, the hierarchy reconstruction with flipped hierarchy had the best results for all methods and error measures, with limited transformation of the non-represented segments having a positive influence. The data that only reconstructed hierarchy without flipping it had the worst results, particularly when applying limited transformation to the non-represented segments, where HD was worse than before registration. The non-reconstructed hierarchy datasets had the smallest metrics all around, but no significant differences to the reconstructed ones in other areas. When comparing all sets that fully transformed non-represented segments to the ones that applied limited transformations, there are no distinct differences on average, though the latter set had considerably smaller metrics. When looking at CFS results, using the landmarks registered with simple hierarchy reconstruction as input yielded the worst results for both MSE and HD by a significant amount.
Within the synthetic data, the one based on the clinical data had the best results, maximising its metric with by highest factor while also minimising its MSE the most.
Among the purely synthetic data, the strongly distorted one stood out with the best results, not the original one. It was the only one that had a better MSE and HD after registration than before, as well as maximising its metric the most, making it almost 10 times higher than upon initialisation, though the initial metric was particularly bad due to the strong distortion in the first place, much worse than that of the other two datasets.

For an overview of all results of MSE, HD and CSF, please refer to Figure 4.8,

4.3 Discussion

The results of the evaluation as well as implementation overall lead to various conclusions. Addressing the evaluation first, it is important to keep in mind that the landmarks used for evaluation and calculating the error measures were set manually. Despite being reviewed twice, they will inherently be flawed, resulting in an error above zero even for perfect co-registration due to micro misplacements that are not visible to the human eye. Conversely, for an automatic and accurate placement of landmarks, a previous co-registration would be required.

During evaluation it can be observed and reconfirmed that a high metric does not automatically lead to a good co-registration. This becomes apparent when looking at the clinical dataset that had hierarchical reconstruction and flipped hierarchy and where non-represented segments were fully transformed.

	Clinical Data						Synthetic Data			
				Partial transformation for non-represented segments			based on clinical	not based on clinical		
	hierarchy reconstructed	not reconstructed	flipped hierarchy	hierarchy reconstructed	not reconstructed	flipped hierarchy		distorted	mildly distorted	not distorted
MSE										
initial	309,0692	309,0692	309,0692	309,0692	309,0692	309,0692	306,3962	4394	570,4286	385,8571
after registration	265,7767	226,7575	232,786	233,795	226,7533	**210,3207**	74,8929	**1510,5878**	1911,5772	1582,5462
HD										
initial	19,2678	19,2678	19,2678	19,2678	19,2678	19,2678	19,2678	87,4128	33,1059	29,8831
after registration	15,5571	15,254	15,1234	21,6634	15,254	**10,7277**	16,5648	**68,3143**	100	90,7562
CSF										
MSE										
initial	16,9256	16,9256	16,9256	16,9256	16,9256	16,9256	17,3401	701,5117	80,3727	69,2327
after registration	261,0956	**152,0557**	157,8781	229,8694	152,132	169,2561	55,2016	**743,5559**	1077,3954	975,3514
HD										
initial	8,9132	8,9132	8,9132	8,9132	8,9132	8,9132	8,9593	42,1673	11,0724	11,7127
after registration	17,975	15,5062	16,6879	19,2797	15,4942	**11,8443**	12,4706	**40,3746**	61,5374	56,9302
Metric										
initial	1,7125	1,7125	1,7125	1,7125	1,7125	1,7125	1,9651	0,086823	0,51537	0,72203
after registration	14,3231	11,6421	**20,1899**	5,16	3,8602	5,8646	52,9315	0,71582	0,97148	0,99933

Figure 4.8. Comprehensive results of the evaluation. Best results across comparable datasets are highlighted in green.

It had the highest metric of all clinical datasets, which have a comparable metric. Yet the dataset that also had flipped hierarchy but did not fully transform non-represented segments had a better MSE and HD, despite having a much lower metric.
Treatment of the non-represented vessels seemed to have made a bigger influence than the reconstruction of the centerline, though among those, the one with reconstructed and flipped hierarchy performed best. Therefore it can be concluded that having only one vessel tree is best for registration. When multiple roots/inlets occur, it can lead to ambiguities, which express themselves in cases like the one illustrated in in Figure 4.9, where two inlet vessels are fitted to the same intensity ridge.

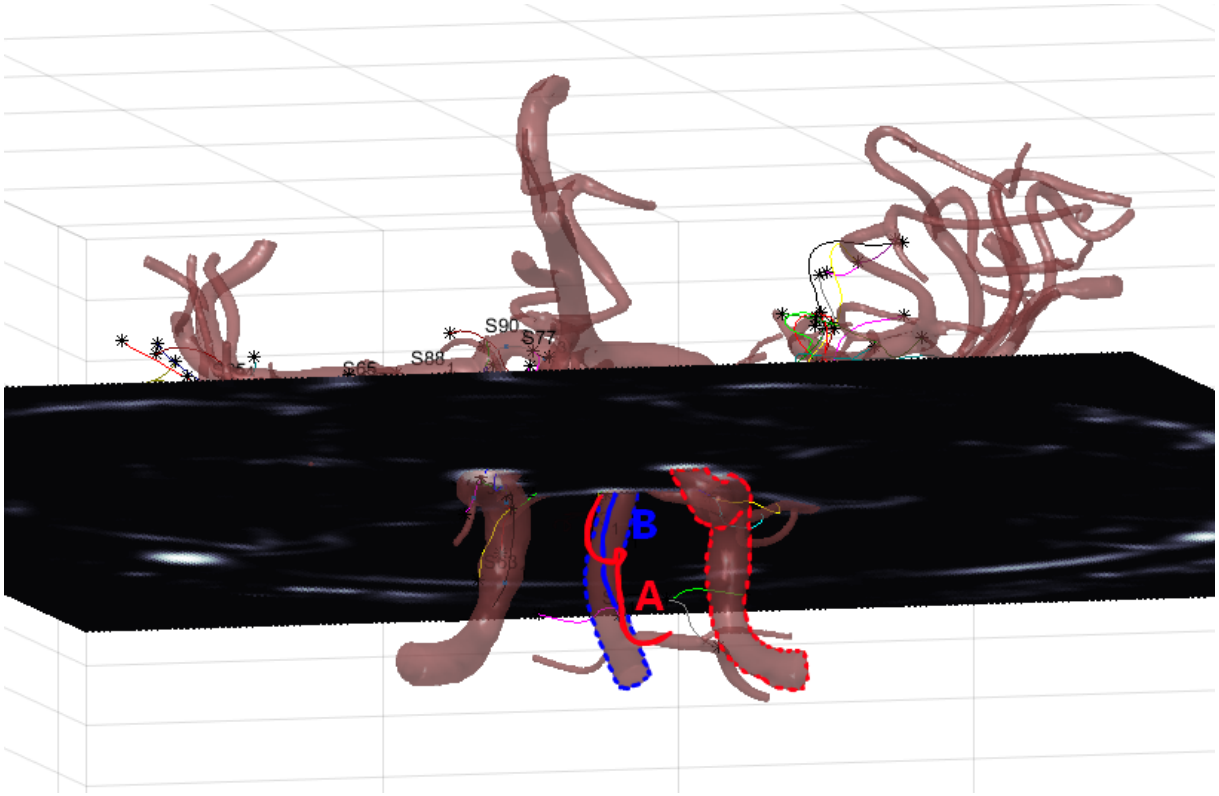

Figure 4.9. Example of two centerlines being registered to the same intensity ridge. Red centerline A originally comes from the vessel hatched in red, but was registered to the vessel hatched in blue that belongs to the blue centerline B.

Another observation is that while CFS almost always achieved better results than error measures prior to registration, those results varied wildly depending on whether landmarks before or after registration were used as input. Counterintuitively, the landmarks that have not been registered yet caused CFS to result in a lower error measure than when the registered landmarks were used as result. While using initial landmarks often meant better results than the registration itself, in case of the hierarchically reconstructed set where non-represented vessels were not transformed fully, using registered landmarks as input even caused the HD CFS result to be worse than the registration result.

When looking at CFS results, it is important to consider that CFS does not make use of any underlying image or model information which are assumed to play a significant role in registration, and arguably are what this registration is based on. CFS merely takes two sets of landmarks, the positions of which will be the only thing it considers when calculating the transformation that yields the best match. CFS also only does a global rigid transformation, finding the best fit for the landmark set by applying an isotropic scaling factor, rotation and translation matrix. There is therefore no local transformation happening, only a global 7 DOF transformation. The landmarks after registration, however, will have been not only transformed with a global affine transformation that includes five more DOF (anisotropic scaling, giving two more scaling DOF, and an additional three from shearing), but will furthermore have been locally transformed, skewing them further apart. It is possible that CFS was not able to account for these local affine 12 DOF transformations when using the registered landmarks as input, which resulted in worse results than if no affine transformations were used on the landmarks

Regardless of CFS, an average MSE of $134.63mm$ and average HD of $10.08mm$ is not sufficient for a co-registration result. Additionally, it has to be considered that the landmarks used to calculate these measures are mainly at bifurcations in the center of the Circle of Willis, where most of the bigger vessels are located, since landmarks can only be placed where the bifurcations are visible in the PC-MRI dataset. Therefore, the errors only minimally consider the fit of the smallest vessels. This is particularly troubling, as visually examining the registration result shows that the smallest vessels underwent a very strong transformation (recall Figure 3.7). A visible translation in z-direction is to be expected, as the PC-MRI data covers less area in that dimension. However, given that this is the same patient, such a strong transformation is likely incorrect. Those strong transformations of the small vessels were exactly why the the option to only partially transform non-represented vessels, which those small vessels are, was implemented. This did not solve this problem though, as results of the not fully transformed non-represented vessels were only minimally better than those of fully transformed non-represented vessels, especially considering that the landmarks were not influenced by the transformations of the non-represented vessels in the first place.

There are several possible reasons for the insufficiency of the registration.

4.3.1 Landmark Placement

Landmark placement refers to the previously mentioned manual placement of the landmarks, which will inherently be flawed. Even with best placement possible, a small distance

will have a considerable effect, especially with error measurements such as MSE. It can be assumed that every corresponding pair of landmarks has an inherent distance between them even in perfect registration due to human error.

4.3.2 Pre-Processing Errors

As mentioned in Section 2.2, one of the problems with model-based registration is that any error during segmentation will affect the registration. It is therefore possible that the pre-processing steps described in Section 3.2 added to the registration results. Especially the segmentation of the TOF-MRI data has to be scrutinised for this.
The first area where key information might have gotten lost or erroneous steps might have been taken is the pre-processing for segmentation in MeVisLab via vesselness filters and masking. The second is the segmentation and conversion into .obj file via windowing and thresholding. A different choice of parameters in any of these steps might have led to different results The third area where pre-processing might have had a negative effect on registration is the manual editing in Blender and Sculptris, where arguably the biggest changes to the raw data were made. However, this was a necessary step to enable centerline extraction, as fused vessels would have had one wrong centerline instead of the actual two. This editing, cutting and smoothing was done to the best of ability and under discussion with a second reviewer with more experience with cerebral vessel networks, but it is possible that some vessels were incorrectly reconstructed.
The next possibly error-prone area is the cutting apart of the centerline and subsequent optional reconstruction. Not only is the cutting apart itself a possible source for errors, the attempted hierarchy reconstruction might be even more so, as it is not clear what the exact hierarchical relationship of the blood flow is, particularly where two feeding vessels meet. This is also affected by the choice of where to place the cut. Had the cuts for example been placed at the other end of the posterior communicating arteries, the hierarchy of them would have been completely different.
The last area where pre-processing might have led to errors happens in Matlab, where despite the manual editing of the segmented data for centerline extraction the centerline itself had to be manually edited to avoid artefacts. An example of what kind of changes had to be made is illustrated in Figure 4.10. Distinct parent vessels were often split apart into small segments when several children occured shortly after another. This meant that the vessel was treated as several vessels in the registration instead of one, which leads to each segment being treated separately, thus causing errors.

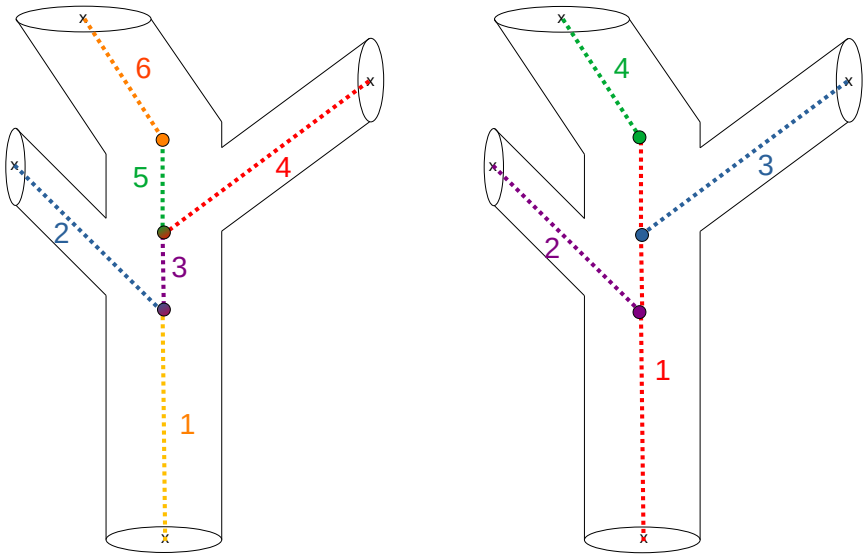

Figure 4.10. Illustrates how the centerline was edited. Left image shows before the edit, right image after the edit, simplifying the centerline to better represent the vessels as one parent vessel with multiple child vessels, rather than splitting the parent vessel.

Despite these many areas where TOF-MRI pre-processing might originate errors, the PC-MRI data might also be a source, starting with only the peak time point that the data is registered too. As mentioned before, extraction of a vessel segmentation in the PC-MRI data was not feasible within the scope of this work, but highlighting of the vessels with the described method might also harbour opportunity for later errors.

4.3.3 Choice of Optimiser Algorithm

Choice of optimiser algorithm could be another reason for the presented co-registration results. Powell's method might be commonly used in the field of medical co-registration for its relatively simple approach and lack of having to compute gradients, but it is possible that exactly those gradients might have been of benefit. Furthermore, the presented registration utilises a pre-existing implementation of Powell's method from the Matlab File Exchange (`https://www.mathworks.com/matlabcentral/fileexchange/`, accessed 23.10.2020) where anyone can upload algorithms. The file chosen had several comments and ratings from other users about its correct function, but it is possible that the used algorithm [64] does not work correctly, despite having reviewed the code, as code comments are lacking.

4.3.4 Choice of Transformations

What is assumed to be the most likely reason for the presented results is the transformations chosen, namely affine transformations. They were selected due to their frequent and established usage in related work, and particularly in the works that inspired the presented implementation. However, it has also been stated that they might not be able to account for the non-linear transformations vessels might undergo, though this referred to inter-subject registration [51, 53]. Due to the limitations of affine transformations mainly being related to inter-subject issues and the presented problem being intra-subject, it was assumed that affine transformations would be sufficient due to their frequent use. This might have been a miscalculation.

4.3.5 Constraints

Lastly, what might have had a significant effect on registration are the various constraints that guide it. These constraints encompass the entire registration workflow, from structuring the data over metric and optimisation to transformation and traversal of the vessel tree. Constraints include:

- Range of valid transformation parameters
- Treatment of out-of-range centerline points in metric calculation
- Representation
- Treatment of non-represented vessels
- Resolution of rotation grid calculated for optimisation
- Treatment of already optimised vessels
- Reconstructed hierarchy
- Enforcing of constraints

Range of Valid Transformation Parameters The first constraint affects the linear transformations that make up the local and global affine transformation. Thereby the data can only be transformed with parameters that lie within a certain range, meaning it can, for example, not be doubled in size or scaled down to a fraction of its original size, or rotated by 180°. A setting of such constraints is reasonable, as extreme transformations are not feasible, especially not in an intra-subject co-registration scenario. Furthermore, it prevents

misregistrations that maximise the metric but are evidently wrong, like the previously mentioned example of scaling the centerlines to be so small they all fit within one voxel within one of the intensity ridges. However, both the exact values of these constraints and their enforcing is open to interpretation. For the presented implementation these values were determined empirically, and they were enforced by severely punishing the metric result if one of the parameters fell out-of-range.

Treatment of out-of-Range Centerline Points in Metric Calculation Centerline points that fall outside of the PC-MRI range were originally punished in a similar way, with having a negative impact on the metric, if not one quite as extreme as the out-of-range transformation parameters. It was ultimately decided to *not* punish out-of-range centerline points, as the PC-MRI data covers a smaller area than the TOF-MRI data, particularly along the longitudinal axis, as shown in Figure 4.11. Therefore a perfect co-registration would fit several of the vessels out of range. However, the concept of representation was introduced to specifically note which vessels were, in fact, represented in the PC-MRI data, with the non-represented ones not influencing the metric as they would have no intensities to be optimised to in the PC-MRI data. It could therefore be reasoned that punishing represented centerline points out of range would still make sense.

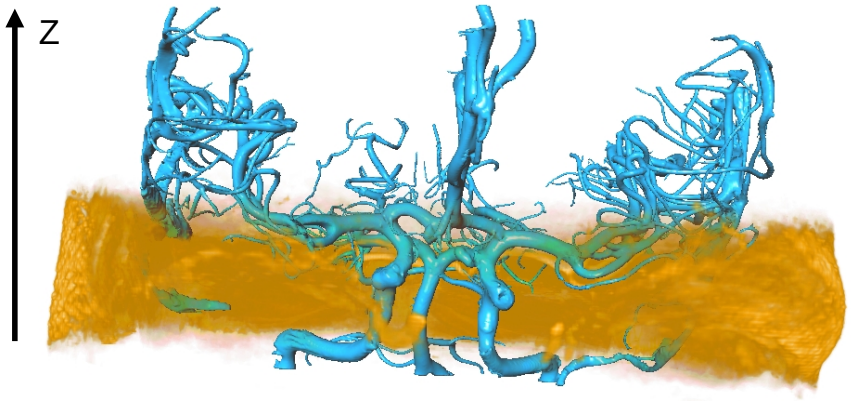

Figure 4.11. Shows mismatch in longitudinal range of TOF- vs PC-MRI data. Longitudinal range is shown via the z-direction and arrow. PC-MRI data is depicted as yellow intensity volume. TOF-MRI is depicted as blue segmentation.

Representation But here arises the questions of how to quantify representation. This was done via vessel radius, another empirically estimated constraint. This does not entirely solve the out-of-range problem from the previous paragraph, though, due to the mentioned smaller area that the PC-MRI data covers. Looking only at the z-axis, the longitudinal axis, the basiliar artery for example is definitely represented in the PC-MRI data. However, due to PC-MRI data covering a smaller area, the lower part of the basiliar artery is not represented, while the upper part is. Figure 4.12 illustrates this problem. While the vessel

network is split into segments, these segments generally cover a vessel up to a bifurcation and in some cases past it. A parent vessel as prominent and high in radius as the basiliar artery would not be split into several segments, and splitting it might actually lead to further registration errors. Still, the presented algorithm thus far does not specifically account for vessels that are tagged as represented but are actually only partially represented in the data - the centerline points that thus lie out of range will not contribute an intensity value to the sum of weighted intensities, though its weight will be added to the normalising sum, meaning the metric will ultimately be penalised for such out-of-range centerline points.

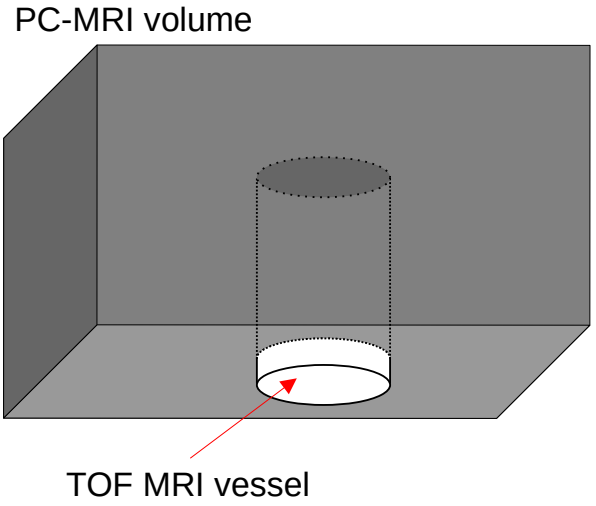

Figure 4.12. Illustration of how a represented TOF-MRI vessel might correctly lie in the PC-MRI data, yet be partially out of range.

Treatment of Non-Represented Vessels While on the topic of representation, the treatment of non-represented vessels is another crucial part of the co-registration, as evidenced by the multiple approaches the presented algorithm implemented. Ultimately they were not counted towards the metric, as they would falsify the results due to their lack of appearance in the PC-MRI data. The way they are transformed seemed to be a more relevant problem to solved, which was attempted by giving them only a partial transformation, so the skewing parameters, despite being very small in a single transformation, would not add up and wildly transform the non-represented vessels to be almost unrecognisable. This, however, barely had an impact on the evaluation measures, due to the landmarks necessarily being affected by represented vessels only, as they had to be placed in the PC-MRI data as well. Evaluating the treatment of the non-represented vessels in a quantitative way therefore is nigh impossible. It shall be noted that a different treatment of those vessels could severely influence the overall registration result.

Resolution of Rotation Grid Calculated for Optimisation Looking at the optimisation step, the resolution of the rotation grid used for multi-start coarse-to-fine approach can be counted as a relevant constraint. Jenkison et al. [51] states that a coarse grid resolution of 6 and fine grid resolution of 20 is sufficient for good results, which in the presented implementation leads to a coarse resolution of 11, which makes fine resolution 21. These resolution values denote the value $steps$ in the algorithm, meaning the number of steps per dimension for which initial rotation is solved for, as elaborated in Section 3.6.4. Along with changing the resolution $steps$, changing $stepsize$, meaning the interval between steps, might also have a considerable effect. Here $stepsize$ was chosen to always lead to whole values, which is relevant for the calculation of the grid back-end. It is important to consider both parameters with the previous constraint of the optimisation parameters in mind, which also constraint the rotation, since $steps$ and $stepsize$ are what ultimately determines the starting parameters for the rotation. Therefore, giving a high $stepsize$ together with a high $steps$ would waste a lot of computing time, as it would lead to most starting positions returning penalised metrics. Feasible parameters are consequently limited, yet their selection might impact the registration. A different implementation that allows for non-whole values could also have lead to different results, though given minimum $steps$ in accordance with the constraints, $stepsize$ is chosen small enough that it is assumed to cover whatever benefit a rational $stepsize$ might have.

Treatment of Already Optimised Vessels and Reconstructed Hierarchy Closely related to the previously mentioned ambiguities around the hierarchical reconstruction of the vessel network is the question of how to handle the segments that after reconstruction have two parents. The base idea of the optimisation strategy, which makes use of the inherent tree structure of the vasculature, is that parent vessel transformations also affect the child vessels, and that if parent vessels are aligned, children vessels are more likely to be aligned too. But if a vessel has two parents, meaning two vessels that feed it, which parent's transformation should be applied? And furthermore, if it is transformed according to one parent, would that not sever its connection to the other parent? Part of this problem was solved by only allowing translations for the absolute root vessels, meaning vessels that have no parents, which are the inlet vessels. Every subsequent vessel will only be allowed to rotate, skew, and scale, while the point connecting it to its parent remains the same. This still means that any vessel with two parents will have one parent that is a child of one inlet and another that is a child of another inlet. When these inlet vessels are optimised, translations are valid, meaning the vessel with two parents will be moved, and whichever parent was not the one being translated might now be further away. This distance will be kept in subsequent transformations. This exact problem is what the flipped hierarchy is trying to account for. This however did not necessarily lead to significantly better results in the evaluation, arguably because it models the hierarchy against the blood flow of the

vessels where the hierarchy was flipped, since using the tree structure usually follows the blood flow via the hierarchy. Still, the different ways to reconstruct the vessel structure where it was cut had an impact on the results, meaning that a different handling, be it reconstructing, not cutting apart for centerline extraction at all, or not chosing an hierarchical approach would likely change the results.

Enforcing of Constraints Lastly, the way these constraints are enforced has to be considered. While some are a simple question of which method is applied to for example reconstruct the tree or which branches are traversed, others have direct influences on the metric. This is often managed by assigning a very low metric value to the metric, often a negative one, which will return to Powell's optimiser and signal not to optimise further in that direction. As Powell's optimiser searches through the transformation parameter space, one could argue that this works against the methodology of the optimiser, as the constraints are not directly calculated in the metric, but rather return their negative value without even calculating the metric. On the other hand, these constraints merely set limits to the valid search space, giving Powell's optimiser a window in which the parameters must lie to achieve a positive metric.

Changing any of these constraints or how they are enforced might have a significant impact on the registration. The presented results from Section 4.2 are likely due to a combination of all the above mentioned reasons. Possible solutions and extensions that might account for some of the theorised error sources are proposed in Section 5.2 at the end of this work.

4.3.6 Mesh Transform Errors

The mesh transform after registration has not been evaluated or discussed thus far. This is because for further use of the data the transformed centerline is sufficient, so a good result of the mesh transform is not relevant. Nevertheless, the mesh transform implemented shall be addressed.
The two big issues the presented approach from Section 3.7 has are artefacts around bifurcations and misassigned vertices. The former is not crucial, as there are various mesh deformation programs that specifically have functions to smooth vertices around bifurcations, where a simple rotation of the centerline can lead to the surrounding mesh cutting itself or stretching apart.
The misassignment of vertices also happens around bifurcations, where a vertex has multiple close centerlines, but also in cases where assigning which centerline a vertex belongs to is obvious to the human observer, but not so to an algorithm based on closest

centerline. These cases are mostly due to the closeness of very thin to very thick vessels. This leads to vertices of the thick vessel being closer to the centerline of the nearby thin vessel than to that of the thick vessel due to its large radius, as illustrated in Figure 4.13.

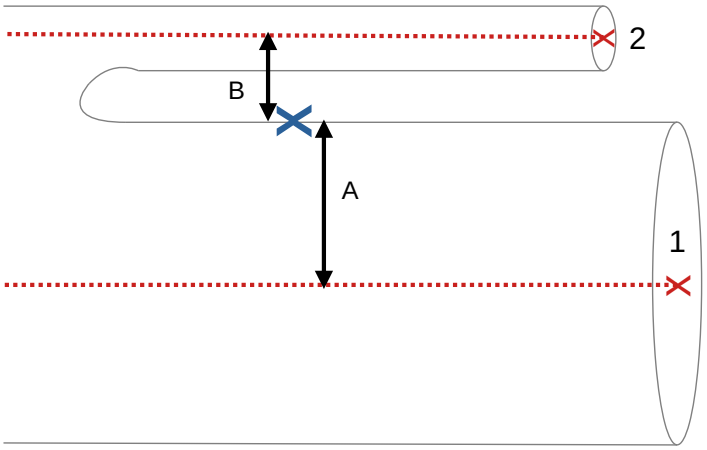

Figure 4.13. Illustration of a common problem of mesh-to-centerline correspondence. Blue X is a mesh vertex of vessel 1. A is the distance from X to the centerline of thick vessel 1 and B is distance of X to centerline of thin vessel 2. Since thin vessel 2 is very close to thick vessel 1, the distance B is shorter than distance A. X will therefore be assigned to vessel 2, despite belonging to mesh vertex of thick vessel 2.

5 Conclusion

The goal of this work was to co-register 7T TOF-MR images to PC-MR images of the Circle of Willis in a non-rigid, model-based way. The proposed vascular registration method is a hybrid model-to-image approach with a multi-start, multi-level strategy that fits the segmented centerlines of the TOF-MRI data to the intensity ridges of the PC-MRI intensity volume. It does so via a scaled sum of weighted intensities, which is optimised with a hierarchical optimisation strategy based on Powell's method that utilises the inherent tree structure of the cerebral vasculature. At the start, a global affine transformation roughly fits the centerlines into the PC-MRI data, beginning the coarse-to-fine process. Starting at the vessel inlets, first the entire vessel tree is optimised. The optimisation decouples the transformation parameters, using rotation parameters on a grid as basis, thus implementing the multi-start approach. The parameters are then iteratively refined until a full 12 DOF affine transformation is used on the vessel tree. The optimisation traverses the tree, iteratively following the parent-child relationships between the vessel segments to consider a smaller subtree in every step, thus slowly switching from global to local transformations.
Due to the considerably lower resolution of the PC-MRI data, not all small vessels that are part of the TOF-MRI segmentation are visible in the DICOM image. Therefore the concept of representation was developed based on vessel radius, assigning a flag to each vessel segment which denotes if a vessel is represented in PC-MRI data or not. If not, this means there is no corresponding intensity ridge of the vessel to be fitted to, so non-represented centerline points are not used for metric calculation.
While all transformation parameters are constrained to only lie within a feasible range, the smallest vessels will be transformed along with their parent vessels. The assumption behind this is that if parent vessels are aligned, child vessels are more likely to be aligned too, meaning this can be used to find the most likely best fit for non-represented vessels, too. However, due to the strategy traversing the vessel tree from root to leaves, the smallest leaf vessels will be transformed multiple times, with any vessel that is a parent's parent. Therefore, even the smallest skewing parameters can add up over time and heavily distort the small vessels, which is not wanted. To prevent this, an option to only apply limited transformations to the non-represented vessels, which are the smallest vessels, was added.
For centerline extraction, it was necessary to cut apart the Circle of Willis to prevent circu-

lar networks. Before registration, the vasculature was reconnected via adding the missing parent-child relationships. This still lead to three inlets, which confused the hierarchical relationship, and the option to flip hierarchy after reconstruction, turning two of the inlets into outlets, was added. Since the metric weights intensities based on the radius of the centerline points, this does not diminish those previous inlets' influence on the registration fit.
In fact, when evaluating the presented registration with manually set landmarks and mean squared error and Hausdorff distance measures, the datasets with the flipped hierarchy had the best results, allowing the conclusion that while reconstructing hierarchy is beneficial, circular structures are confusing to the hierarchical traversal. Furthermore, applying only limited transformation of the non-represented vessels seemed to also yield better results, though this was not as conclusive, as landmarks had to be placed in represented areas.
Registration with the proposed method always leads to metric maximisation, though not always to lowered MSE and HD. The average MSE was $134.63mm$ and average HD $10.08mm$.

5.1 Limitations

Even considering the inherent misalignment of the landmarks due to manual placement, according to the results of the evaluation the proposed implementation is not sufficient for co-registration.
This is likely due to the combination of reasons listed in Section 4.3, but the two that are assumed to have the most impact are the transformations and the tree structure.
With transformation the affine transformations are meant. Replacing the local affine transformations with higher-order ones might result in a more precise fit for the vessels. Another option would be to add a higher-order transformation after the affine ones instead of replacing them. This was considered, but the scope of this work did not allow for this addition. It must be noted that a higher-order transformation could never entirely replace the affine one, as a rough initial fit is needed for most local and higher-order transformations to work. For this, the global affine transformation, and possibly the local one too, is sufficient.
Tree structure refers to how the vessel connections are treated. Evaluation showed that the way the centerlines were reconstructed had an influence on the result. Cyclical relationships in particular posed a problem, but as the hierarchical approach was chosen early on, a switch away from the tree structure, or to look for a different centerline extraction that did not require removal of the cyclic structures and thus cutting apart the vessel tree was not possible.

5.2 Future Work

Due to the scope of this work being limited, some concessions had to be made during development. There are various ways in which the proposed registration could be extended and potentially changed to yield better results.

One thing to change applies to what happens before the actual registration, namely the centerline extraction. With the current method, cutting apart the Circle of Willis to be three separate vessel trees is necessary because circular structures led to errors. A different method might be able to extract the centerline without the cutting apart. A skeletonisation algorithm would be thinkable. The current method also has some issues around bifurcations, making additional editing necessary.
Related to the centerline extraction itself and the fact that it currently has to get cut apart is the subsequent hierarchy reconstruction, which has a significant impact on the registration results. Not cutting the vessel structure apart of course does not require a reconstruction, though ambiguities in the hierarchy might still occur. This is why completely abandoning a hierarchical approach might also work, instead looking at the vessels on their own, or only at immediate neighborhood without considering which is a feeding vessel or child.

A reoccurring problem during implementation was that occasionally two centerlines would be transformed to run along the same intensity ridge. A way to solve this would be to enforce a minimum distance between two centerlines determined by the radii of the centerlines in question.

When talking about constraints in Section 4, something that came up repeatedly was the optimisation. One part of this was the Powell's optimiser used, which could either be replaced with another optimiser, possibly a gradient based one, or to implement it from scratch instead of using a third party code from the Matlab File Exchange.
Another constraint mentioned was the rotation grid resolution, which was limited to whole numbers in the presented implementations. Allowing for rational numbers here could possibly change the registration results, as would loosening up the constraints of the transformation parameters to not apply while the rotation grid is being set up. Another option would be to implement adaptable parameter ranges based on the step of the optimisation, adding to the coarse-to-fine approach.
Generally, implementing various constraints to be flexible could be beneficial. Constraints could change based on in which step of the optimisation it is applied, either increasing or decreasing dynamically, or being removed entirely.

Possibly easiest to change in future work would be the transformations themselves. A global affine transformation registration to start with is a commonly used idea, but to locally refine the transformation and heighten the DOF could be highly beneficial. The local affine transformations of the smaller vessel subtrees could either antecede a non-linear transformation or be entirely replaced by it. A suitable transformation for this would be free deformation models via B-Splines [8, 39] or thin plate splines [42, 43].
If replacing the affine transformation does not seem feasible, it could still be further refined, for example by limiting local transformations of vessels to each centerline's normal direction [46, 52]. This way, centerlines are not transformed along a vessel, but perpendicular to it. This can be further extended by a more selective subsampling of the centerline, excluding outliers via medialness and rigidness [46, 52].

Going back to methodology, another option for this registration would be to completely change the registration bias: model-based seems like best approach for vessel co-registration, but a true model-model approach rather than a hybrid one, as was originally considered, could still be achieved, even if not possible within the scope of this work. If a centerline could be extracted for PC-MRI, active shape models or graph matching like Charnoz et al.'s [44, 45] technique could be used, which would also work with the abandoning of hierarchical structures.

Another concession that had to be made was using only one time point of the PC-MRI data to register to. The other timepoints could however still be used for registration in theory. When scrolling through timepoints, it can be observed that different timepoints have different blood flow, often leading different vessels being more or less visible at different timepoints. This could be used for a more accurate registration, either registering to different timepoints and selecting from those transformations, or pre-processing the PC-MRI data so that each vessel is selected from the timepoint during which it was displayed best.

If a better mesh transform after registration is desired, the assigning of which segment a mesh vertex belongs to could be extended by taking the vessel radii into account. After the closest centerline point has been determined, the distance between that centerline point and the mesh vertex could be compared to the centerline point's radius, and if the distance exceeds the radius, a centerline point from the second closest vessel could be chosen instead. This could potentially solve the problem illustrated in Figure 4.13, where vertices of very big vessels are assigned to nearby small vessels.

Lastly, to fully round out a registration for TOF- and PC-MRI data it would be beneficial to find a way to evaluate the fit of the non-represented vessels after registration in a quan-

titative way. This thesis merely compared the represented centerlines after registration to where they lay before registration, which is not sufficient.

Bibliography

[1] J.M. Wardlaw, C. Smith, and M. Dichgans. “Small Vessel Disease: Mechanisms and Clinical Implications”. In: *The Lancet Neurology* 18.7 (2019), pp. 684–696. DOI: 10.1016/S1474-4422(19)30079-1.

[2] Y. Shi, M.J. Thrippleton, G.W. Blair, D.A. Dickie, I. Marshall, I. Hamilton, F.N. Doubal, F. Chappell, and J.M. Wardlaw. “Small Vessel Disease is Associated with Altered Cerebrovascular Pulsatility but not Resting Cerebral Blood Flow”. In: *Journal of Cerebral Blood Flow & Metabolism* 40.1 (2020), pp. 85–99. DOI: 10.1177/0271678X18803956.

[3] P. Berg, D. Stucht, G. Janiga, O. Beuing, O. Speck, and D. Thévenin. “Cerebral Blood Flow in a Healthy Circle of Willis and Two Intracranial Aneurysms: Computational Fluid Dynamics Versus Four-Dimensional Phase-Contrast Magnetic Resonance Imaging”. In: *Journal of Biomechanical Engineering* 136.4 (2014), pp. 55–66. DOI: 10.1115/1.4026108.

[4] J.R. Cebral, M.A. Castro, O. Soto, R. Löhner, and N. Alperin. “Blood-Flow Models of the Circle of Willis from Magnetic Resonance Data”. In: *Journal of Engineering Mathematics* 47.3 (2003), pp. 369–386. DOI: 10.1023/B:ENGI.0000007977.02652.02.

[5] J. Garcia, A.J. Barker, I. Murphy, K. Jarvis, S. Schnell, J.D. Collins, J.C. Carr, S.C. Malaisrie, and M. Markl. “Four-Dimensional Flow Magnetic Resonance Imaging-Based Characterization of Aortic Morphometry and Haemodynamics: Impact of Age, Aortic Diameter, and Valve Morphology”. In: *European Heart Journal - Cardiovascular Imaging* 17.8 (2016), pp. 877–884. DOI: 10.1093/ehjci/jev228.

[6] A. Harloff, F. Albrecht, J. Spreer, A. Stalder, J. Bock, A. Frydrychowicz, J. Schöllhorn, A. Hetzel, M. Schumacher, and J. Hennig. “3D Blood Flow Characteristics in the Carotid Artery Bifurcation Assessed by Flow-Sensitive 4D MRI at 3T”. In: *Magnetic Resonance in Medicine* 61.1 (2009), pp. 65–74. DOI: 10.1002/mrm.21774.

[7] H.M. Chan, A.C.S. Chung, S.C.H. Yu, and W.M. Wells. “2D-3D Vascular Registration between Digital Subtraction Angiographic (DSA) and Magnetic Resonance Angiographic (MRA) Images”. In: *Proceedings of the 2nd IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. 2004, pp. 708–711. DOI: 10.1109/ISBI.2004.1398636.

[8] J. Jomier and S.R. Aylward. “Rigid and Deformable Vasculature-to-Image Registration: A Hierarchical Approach”. In: *Proceedings of the 7th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2004, pp. 829–836. DOI: 10.1007/978-3-540-30135-6_101.

[9] S. Jandke, C. Garz, D. Schwanke, M. Sendtner, H.-J. Heinze, R.O. Carare, and S. Schreiber. “The Association between Hypertensive Arteriopathy and Cerebral Amyloid Angiopathy in Spontaneously Hypertensive Stroke-Prone Rats”. In: *Brain Pathology* 28.6 (2018), pp. 844–859. DOI: 10.1111/bpa.12629.

[10] S. Schreiber, C.Z. Bueche, C. Garz, S. Kropf, F. Angenstein, J. Goldschmidt, J. Neumann, H.-J. Heinze, M. Goertler, K.G. Reymann, and H. Braun. “The Pathologic Cascade of Cerebrovascular Lesions in SHRSP: Is Erythrocyte Accumulation an Early Phase?” In: *Journal of Cerebral Blood Flow & Metabolism* 32.2 (2011), pp. 278–290. DOI: 10.1038/jcbfm.2011.122.

[11] S.M. Wong, J.F.A. Jansen, C.E. Zhang, E.I. Hoff, J. Staals, R.J. van Oostenbrugge, and W.H. Backes. “Blood-Brain Barrier Impairment and Hypoperfusion are Linked in Cerebral Small Vessel Disease”. In: *Neurology* 92.15 (2019), pp. 1669–1677. DOI: 10.1212/WNL.0000000000007263.

[12] K. Kitagawa, N. Oku, Y. Kimura, Y. Yagita, M. Sakaguchi, J. Hatazawa, and S. Sakoda. “Relationship between Cerebral Blood Flow and Later Cognitive Decline in Hypertensive Patients with Cerebral Small Vessel Disease”. In: *Hypertension Research* 32.9 (2009), pp. 816–820. DOI: 10.1038/hr.2009.100.

[13] F.J. Wolters, H.I. Zonneveld, A. Hofman, A. van der Lugt, P.J. Koudstaal, M.W. Vernooij, and M.A. Ikram. “Cerebral Perfusion and the Risk of Dementia”. In: *Circulation* 136.8 (2017), pp. 719–728. DOI: 10.1161/CIRCULATIONAHA.117.027448.

[14] C.E. Zhang, S.M. Wong, R. Uiterwijk, J. Staals, W.H. Backes, E.I. Hoff, T. Schreuder, C.R.L.P.N. Jeukens, J.F.A. Jansen, and R.J. van Oostenbrugge. “Intravoxel Incoherent Motion Imaging in Small Vessel Disease”. In: *Stroke* 48.3 (2017), pp. 658–663. DOI: 10.1161/STROKEAHA.116.015084.

[15] M. Reiser and W. Semmler. *Magnetresonanztomographie*. Second Edition. Springer Verlag, 2013. ISBN: 3642979629.

[16] B. Preim and C. Botha. *Visual Computing for Medicine - Theory, Algorithms, and Applications*. Second Edition. Morgan Kaufmann, 2014. ISBN: 978-0-12-415873-3.

[17] N.C. Campeau and J. Huston III. “Vascular Disorders – Magnetic Resonance Angiography: Brain Vessels”. In: *Neuroimaging Clinics of North America* 22.2 (2012), pp. 207–233. DOI: 10.1016/j.nic.2012.02.006.

[18] N.J. Pelc, R.J. Herfkens, A. Shimakawa, and D.R. Enzmann. “Phase Contrast Cine Magnetic Resonance Imaging”. In: *Magnetic Resonance Quarterly* 7.4 (1991). PMID: 1790111, pp. 229–154.

[19] Z. Stankovic, B.D. Allen, J. Garcia, K.B. Jarvis, and M. Markl. “4D Flow Imaging with MRI”. In: *Cardiovascular Diagnosis & Therapy* 4.2 (2014), pp. 173–192. DOI: 10.3978/j.issn.2223-3652.2014.01.02.

[20] M.P. Hartung, T.M. Grist, and C.J. François. “Magnetic Resonance Angiography: Current Status and Future Directions”. In: *Journal of Cardiovascular Magnetic Resonance* 13.1 (2011), p. 19. DOI: 10.1186/1532-429X-13-19.

[21] M. Miyazaki and V.S. Lee. “Nonenhanced MR Angiography”. In: *Radiology* 248.1 (2008), pp. 20–43. DOI: 10.1148/radiol.2481071497.

[22] F. Gaidzik, C. Stucht D. Roloff, O. Speck, D. Thévenin, and G. Janiga. “Transient flow prediction in an idealized aneurysm geometry using data assimilation”. In: *Computers in Biology and Medicine* 115.103507 (2019). DOI: 10.1016/j.compbiomed.2019.103507.

[23] F. Maes, D. Vandermeulen, and P. Suetens. “Medical Image Registration Using Mutual Information”. In: *Proceedings of the IEEE* 91.10 (2003), pp. 1699–1722. DOI: 10.1109/JPROC.2003.817864.

[24] D. Rueckert, A. Frangi, and J. Schnabel. “Automatic Construction of 3D Statistical Deformation Models of the Brain Using Non-Rigid Registration”. In: *IEEE Transactions on Medical Imaging* 22 (2003), pp. 1014–1025. DOI: 10.1109/TMI.2003.815865.

[25] D. Chillet, J. Jomier, D. Cool, and S.R. Aylward. “Vascular Atlas Formation Using a Vessel-to-Image Affine Registration Method”. In: *Proceedings of the 6th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2003, pp. 335–342. DOI: 10.1007/978-3-540-39899-8_42.

[26] J. Yao and R. Taylor. “Assessing Accuracy Factors in Deformable 2D/3D Medical Image Registration Using a Statistical Pelvis Model”. In: *Proceedings of the 9th IEEE International Conference on Computer Vision*. 2003, pp. 1329–1334. DOI: 10.1109/ICCV.2003.1238644.

[27] T. König and K.D. Tönnies. *7. Registrierung I und 8. Registrierung II*. Veranstaltung Medizinsche Bildverarbeitung, Otto-von-Guericke Universität Magdeburg. 2016.

[28] A.A. Goshtasby. *2-D and 3-D Image Registration for Medical, Remote Sensing, and Industrial Applications*. Wiley Press, 2005. ISBN: 978-0-471-64954-0.

[29] D. Lesage, E.D. Angelini, I. Bloch, and G. Funka-Lea. "A Review of 3D Vessel Lumen Segmentation Techniques: Models, Features and Extraction Schemes". In: *Medical Image Analysis* 13.6 (2009), pp. 819–845. DOI: https://doi.org/10.1016/j.media.2009.07.011.

[30] S. Matl, R. Brosig, M. Baust, N. Navab, and S. Demirci. "Vascular Image Registration Techniques: A Living Review". In: *Medical Image Analysis* 35.1 (2017), pp. 1–17. DOI: 10.1016/j.media.2016.05.005.

[31] S. Engelhardt and C. Hansen. *9. Registration*. Veranstaltung Computer-Assisted Surgery, Otto-von-Guericke Universität Magdeburg. 2017.

[32] S. Loncaric. *Introduction to Image Registration*. Image Processing Group at the Faculty of Electrical Engineering and Computing, University of Zagreb. 2004. URL: http://www.inf.u-szeged.hu/ssip/2004/proceedings/lectures/Lecture_21.pdf (visited on 04/14/2020).

[33] F. Gaidzik, S. Pathiraja, S. Saalfeld, D. Stucht, O. Speck, D. Thévenin, and G. Janiga. "Hemodynamic Data Assimilation in a Subject-Specific Circle of Willis Geometry". In: *Clinical Neuroradiology* (2020). DOI: 10.1007/s00062-020-00959-2.

[34] H.-M. Chan and A.C.S. Chung. "Efficient 3D-3D Vascular Registration Based on Multiple Orthogonal 2D Projections". In: *Proceedings of the 2nd International Workshop on Biomedical Image Registration (WBIR)*. 2003, pp. 301–310. DOI: 10.1007/978-3-540-39701-4_32.

[35] K. Lau and A.C.S. Chung. "A Global Optimization Strategy for 3D-2D Registration of Vascular Images". In: *Proceedings of the British Machine Vision Conference*. 2006, pp. 489–498. DOI: 10.5244/C.20.51.

[36] H. Sundar, A. Khamene, C. Xu, F. Sauer, and C. Davatzikos. "A Novel 2D-3D Registration Algorithm for Aligning Fluoro Images with 3D Pre-OP CT/MR Images". In: *Proceedings of the International Society for Optical Engineering (SPIE) - Medical Imaging: Visualization, Image-Guided Procedures, and Display*. Vol. 6141. 2006, pp. 760–766. DOI: 10.1117/12.654251.

[37] G. Sharma and A. Thé. *Automating Image Registration with MATLAB*. 2013. URL: https://de.mathworks.com/company/newsletters/articles/automating-image-registration-with-matlab.html (visited on 04/16/2020).

[38] D. Rueckert, L.I. Sonoda, C. Hayes, D.L.G. Hill, M.O. Leach, and D.J. Hawkes. “Non-rigid Registration Using Free-Form Deformations: Application to Breast MR Images”. In: *IEEE Transactions on Medical Imaging* 18.8 (1999), pp. 712–721. DOI: 10.1109/42.796284.

[39] T. Lange, S. Eulenstein, M. Hünerbein, H. Lamecker, and P.M. Schlag. “Augmenting Intraoperative 3D Ultrasound with Preoperative Models for Navigation in Liver Surgery”. In: *Proceedings of the 7th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2004, pp. 534–541. DOI: 10.1007/978-3-540-30136-3_66.

[40] F. Zana and J.C. Klein. “A Multimodal Registration Algorithm of Eye Fundus Images Using Vessels Detection and Hough Transform”. In: *IEEE on Transactions on Medical Imaging* 18.5 (1999), pp. 419–428. DOI: 10.1109/42.774169.

[41] Y. Song, J. Totz, S. Thompson, S. Johnsen, D. Barratt, C. Schneider, K. Gurusamy, B. Davidson, S. Ourselin, D. Hawkes, and M.J. Clarkson. “Locally Rigid, Vessel-Based Registration for Laparoscopic Liver Surgery”. In: *International Journal of Computer Assisted Radiology and Surgery* 10.12 (2015), pp. 1951–1961. DOI: 10.1007/s11548-015-1236-8.

[42] I. Reinertsen, M. Descoteaux, K. Siddiqi, and D.L. Collins. “Validation of Vessel-Based Registration for Correction of Brain Shift”. In: *Medical Image Analysis* 11.4 (2007), pp. 374–388. DOI: 10.1016/j.media.2007.04.002.

[43] T. Lange, N. Papenberg, S. Heldmann, J. Modersitzki, B. Fischer, H. Lamecker, and P.M. Schlag. “3D Ultrasound-CT Registration of the Liver Using Combined Landmark-Intensity Information”. In: *International Journal of Computer Assisted Radiology and Surgery* 4.1 (2009), pp. 79–88. DOI: 10.1007/s11548-008-0270-1.

[44] A. Charnoz, V. Agnus, G. Malandain, C. Forest, M. Tajine, and L. Soler. “Liver Registration for the Follow-Up of Hepatic Tumors”. In: *Proceedings of the 8th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2005, pp. 155–62. DOI: 10.1007/11566489_20.

[45] A. Charnoz, V. Agnus, G. Malandain, S. Nicolau, M. Tajine, and L. Soler. “Design of Robust Vascular Tree Matching: Validation on Liver”. In: *Proceedings of the 19th International Conference on Information Processing in Medical Imaging*. 2005, pp. 443–55. DOI: 10.1007/11505730_37.

[46] S.R. Aylward, J. Jomier, S. Weeks, and E. Bullitt. “Registration and Analysis of Vascular Images”. In: *International Journal of Computer Vision* 55.2 (2003), pp. 123–138. DOI: 10.1023/A:1026126900358.

[47] E.Y. Kim, H. Johnson, and N. Williams. “Affine Transformation for Landmark Based Registration Initializer in ITK”. In: *The MIDAS Journal - Medical Imaging and Computing* (2011). URL: http://hdl.handle.net/10380/3299 (visited on 08/20/2020).

[48] H. Späth. “Fitting Affine and Orthogonal Transformations between Two Sets of Points”. In: *Mathematical Communications* 9.1 (2004), pp. 27–34. URL: https://hrcak.srce.hr/712.

[49] G. Adluru, E.V.R. DiBella, and M.C. Schnabel. “Model-Based Registration for Dynamic Cardiac Perfusion MRI”. In: *Journal of Magnetic Resonance Imaging* 24.5 (2006), pp. 1062–1070. DOI: 10.1002/jmri.20756.

[50] B.B. Avants, C.L. Epstein, M. Grossman, and J.C. Gee. “Symmetric Diffeomorphic Image Registration with Cross-Correlation: Evaluating Automated Labeling of Elderly and Neurodegenerative Brain”. In: *Medical Image Analysis* 12.1 (2008), pp. 26–41. DOI: 10.1016/j.media.2007.06.004.

[51] M. Jenkinson and S. Smith. “A Global Optimisation Method for Robust Affine Registration of Brain Images”. In: *Medical Image Analysis* 5.2 (2001), pp. 143–156. DOI: 10.1016/s1361-8415(01)00036-6.

[52] S.R. Aylward, S. Weeks, and E. Bullitt. “Analysis of the Parameter Space of a Metric for Registering 3D Vascular Images”. In: *Proceedings of the 4th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2001, pp. 932–939. DOI: 10.1007/3-540-45468-3_111.

[53] D. Rueckert, P. Aljabar, R. Heckemann, J. Hajnal, and A. Hammers. “Diffeomorphic Registration Using B-Splines”. In: *Proceedings of the 9th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2006, pp. 702–709. DOI: 10.1007/11866763_86.

[54] W. Crum, D. Rueckert, M. Jenkinson, D. Kennedy, and S. Smith. “A Framework for Detailed Objective Comparison of Non-Rigid Registration Algorithms in Neuroimaging”. In: *Proceedings of the 7th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2004, pp. 679–686. DOI: 10.1007/978-3-540-30135-6_83.

[55] M.J.D. Powell. “An Efficient Method for Finding the Minimum of a Function of Several Variables Without Calculating Derivatives”. In: *The Computer Journal* 7.2 (1964), pp. 155–162. DOI: 10.1093/comjnl/7.2.155.

[56] CEM Lectures. *Topic 8c – Multivariable Optimization*. 2018. URL: https://www.youtube.com/watch?v=ep4C_R7dUMs&t=1579s (visited on 10/20/2020).

[57] R. Sablatnig and M. Kampel. “Model-Based Registration of Front- and Backviews of Rotationally Symmetric Objects”. In: *Computer Vision and Image Understanding* 87.1 (2002), pp. 90–103. DOI: 10.1006/cviu.2002.0985.

[58] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben. “4D Flow MRI”. In: *Journal of Magnetic Resonance Imaging* 36.5 (2012), pp. 1015–1036. DOI: 10.1002/jmri.23632.

[59] M. Markl, A. Harloff, T.A. Bley, M. Zaitsev, B. Jung, E. Weigang, M. Langer, J. Hennig, and A. Frydrychowicz. “Time-Resolved 3D MR Velocity Mapping at 3T: Improved Navigator-Gated Assessment of Vascular Anatomy and Blood Flow”. In: *Journal of Magnetic Resonance Imaging* 25.4 (2007), pp. 824–831. DOI: 10.1002/jmri.20871.

[60] J. Bock, B. Kreher, J. Hennig, and M. Markl. “Optimized Pre-Processing of Time-Resolved 2D and 3D Phase-Contrast MRI Data”. In: *Proceedings of the 15th Annual Meeting of the International Society for Magnetic Resonance in Medicine (ISMRM)*. 2007, p. 3138. URL: https://cds.ismrm.org/ismrm-2007/files/03138.pdf (visited on 08/19/2020).

[61] S. Saalfeld, P. Berg, M. Neugebauer, and B. Preim. “Reconstruction of 3D Surface Meshes for Blood Flow Simulations of Intracranial Aneurysms”. In: *Proceedings of the 14th Jahrestagung der deutschen Gesellschaft für Computer- und Roboterassistierte Chirurgie (CURAC)*. 2015, pp. 163–168. URL: http://www.vismd.de/lib/exe/fetch.php?media=files:misc:glasser_2015_curac.pdf (visited on 08/18/2020).

[62] W.S. Harwin and D. Hackett. *loadOBJ_externBSD*. University Reading, Matlab BSD license. 2010.

[63] S. Saalfeld, S. Voß, O. Beuing, and P. Berg. “Flow-Splitting-Based Computation of Outlet Boundary Conditions for Improved Cerebrovascular Simulation in Multiple Intracranial Aneurysms”. In: *International Journal of Computer Assisted Radiology and Surgery* 14.10 (2019), pp. 1805–1813. DOI: 10.1007/s11548-019-02036-7.

[64] G. Tonel. *Unconstrained Optimization Using Powell*. 2007. URL: https://www.mathworks.com/matlabcentral/fileexchange/15072-unconstrained-optimization-using-powell (visited on 08/23/2020).

[65] D.H. House and J.C. Keyser. *Foundations of Physically Based Modeling and Animation*. A K Peters/CRC Press, 2006. ISBN: 9781482234602.

[66] A.N. Strahler. “Quantitative Analysis of Watershed Geomorphology”. In: *American Geophysical Union Transactions* 8.6 (1957), pp. 913–920. DOI: 10.1029/TR038i006p00913.

[67] H. Hahn, B. Preim, D. Selle, and H.-O. Peitgen. “Visualization and Interaction Techniques for the Exploration of Vascular Structures”. In: *Proceedings of the IEEE Conference on Visualization*. 2001, pp. 395–578. DOI: 10.1109/VISUAL.2001.964538.

[68] H. Hahn, M. Georg, and H.-O. Peitgen. “Fractal Aspects of Three-Dimensional Vascular Constructive Optimization”. In: *Proceedings of Conference on Fractals in Biology and Medicine*. 2005, pp. 55–66. DOI: 10.1007/3-7643-7412-8_5.

[69] D.P. Huttenlocher, G.A. Klanderman, and W.J. Rucklidge. “Comparing Images Using the Hausdorff Distance”. In: 15.9 (1993), pp. 850–863. DOI: 10.1109/34.232073.

[70] B. Horn. “Closed-Form Solution of Absolute Orientation Using Unit Quaternions”. In: *Journal of the Optical Society of America* 4.4 (1987), pp. 629–642. DOI: 10.1364/JOSAA.4.000629.

[71] C. Wengert and G. Bianchi. *Absolute Orientation Quaternion*. ETH Zürich, Computer Vision Laboratory. 2010. URL: https://www.mathworks.com/matlabcentral/fileexchange/22422-absolute-orientation (visited on 08/18/2020).

Selbstständigkeitserklärung

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe. Andere als die angegebenen Literatur und Hilfsmittel wurden nicht benutzt und ich habe mich sonst keiner unerlaubten Hilfe bedient. Die aus fremden Quellen direkt oder indirekt übernommenen Stellen sind als solche kenntlich gemacht.
Dieses Arbeit wurde bisher weder in gleicher noch ähnlicher Form im In- oder Ausland einer anderen Prüfungsbehörde zur Beurteilung vorgelegt.

Datum: _____ Unterschrift: _____