

Blood Flow Velocity Estimation from Cardiac Angiography Images

Xiangyu Lu¹

¹MSc in IT and Cognition, University of Copenhagen
xilu@di.ku.dk

Abstract

The prediction of blood flow velocity is a crucial and continually evolving research area in the fields of medicine and biology. Accurate predictions of blood flow velocity provide more precise information for disease diagnosis and treatment. In this paper, we have designed a framework that includes a distance calculation algorithm based on A* and a conversion algorithm based on distance lines to predict the projected velocity and 3D velocity of blood flow in vessel branches depicted in angiographic images. A series of experiments were conducted on X-ray angiography image sequences from the Department of Cardiology at the University Hospital of Stavanger to assess the performance of our framework.

1 Introduction

1.1 Research background

The human heart is a muscular organ located in the chest, responsible for pumping blood throughout the body, ensuring the delivery of oxygen and nutrients to various tissues and organs. It consists of four chambers that work together to facilitate efficient blood circulation, and its rhythmic contractions are controlled by electrical impulses, making it a vital component of the circulatory system [1].

The coronary arteries are a crucial part of the human heart's blood supply system. They are a group of blood vessels located on the surface of the heart and are responsible for delivering oxygen-rich blood and nutrients to the heart muscles to maintain normal cardiac function [3]. When these vessels are affected by disease or narrow, it can lead to inadequate blood supply to the heart muscles, a condition known as Coronary Artery Disease (CAD). CAD is a leading cause of heart disease and can result in serious consequences such as angina and myocardial infarction (heart attack) [4].

In order to diagnose and treat CAD, as well as assess the condition of a patient's coronary arteries, coronary angiography has become an indispensable diagnostic procedure. The history of coronary angiography dates back to the early 20th century when medical professionals used X-rays and contrast agents to study the heart's blood vessels. Over time, this technology has been continuously improved to become safer and more precise. In coronary angiography, doctors typically make a small incision in the patient's wrist or thigh and then inject a special contrast agent into the coronary arteries using a catheter. Simultaneously, X-ray equipment (typically a C-arm machine) takes continuous images to display the distribution of the contrast agent within




Figure 1. Conducting system of human heart from [2].




Figure 2. An example of a C-arm system setup from [6].

the coronary arteries [5]. These images provide precise information about the arteries, including the degree and location of any narrowings.

If narrowing or blockage is detected during coronary angiography, doctors can take immediate intervention measures, such as angioplasty (coronary balloon dilation) and stent placement [7]. These treatments can restore blood flow, alleviate symptoms, and improve cardiac perfusion. Furthermore, coronary angiography allows healthcare professionals to assess the overall health of the coronary arteries, determining whether multiple vessels are affected and whether the disease is limited to specific segments or more widespread. This information is vital for developing a comprehensive treatment plan tailored to the patient’s condition [8].

The technology of coronary angiography has continually improved. Modern medicine has developed more precise and safer contrast agents and catheters, as well as advanced X-ray imaging equipment. Additionally, non-invasive imaging techniques like CT angiography and MRI angiography have been developed, providing detailed vascular images without the need for invasive catheter-based procedures, making this diagnostic procedure even more accurate and reliable [9].

1.2 Problem description

In this paper, we aim to develop a comprehensive framework for the estimation of blood flow velocity in angiography image sequences. As vascular angiographic images are typically obtained through X-ray projection, our task is divided into two main components. The first is to estimate the projected blood flow velocity on the imaging plane through these projection images. The second is to estimate the 3D blood flow velocity based on the estimated blood flow velocity combined with the real surface model of the heart.

This requires extracting valid blood flow information from angiography image sequences. Since the images often contain intrinsic movements from the patient (such as heartbeat and respiration) and other noise, it is essential to design algorithms to estimate blood flow in various vessel branches after eliminating these interferences. In the estimation of 3D velocity, we also need to consider the velocity loss along the vertical axis caused by the projection transformation. We will provide more details in the Data Section and Methods Section.

2 Related works

2.1 Image registration

Image Registration plays a crucial role in medical image analysis. It is a key technique used to precisely align and overlay medical images from different times, different devices, or different imaging modalities and it enables doctors, researchers, and computer algorithms to more effectively analyze, diagnose, and treat diseases. Image registration can be employed to track the growth of lesions, navigate surgical procedures, compute patient anatomies, and fuse multiple images, thus providing more comprehensive and accurate medical information.

The early stages of image registration were characterized by the use of traditional techniques, primarily relying on manual operations and visual observation. This period of development can be traced back to the early 20th century [10]. Medical imaging was one of the pioneering fields to employ these techniques for aligning medical images taken at different times or using different imaging modalities. Methods during this era were mainly based on traditional computer vision techniques such as feature points, edges, and histograms [11]. During this time, professionals manually identified corresponding points or features in images and used them as references to align the images. This process was time-consuming, and the accuracy depended on the operator’s skill. Despite its limitations, these early methods laid the foundation for more advanced image registration techniques.

With advancements in computer hardware and mathematical algorithms, feature-based image

registration methods gained prominence. Feature-based methods involve the detection of salient points or features in images [12]. These features could be corners, edges, blobs, or other distinctive patterns. Common algorithms used for feature detection include Harris corners [13] and Canny edge detection [14]. Once features are detected in both images, the next step is to find correspondences between these features. Various matching algorithms, such as nearest-neighbor matching or Random Sample Consensus (RANSAC) [15], are used to establish correspondences. After establishing feature correspondences, the transformation needed to align one image with the other is estimated, which is often represented by a mathematical model, such as an affine or projective transformation. Notable feature-based techniques like Scale-Invariant Feature Transform (SIFT) [16], Speeded-Up Robust Features (SURF) [17] and Oriented FAST and Rotated BRIEF (ORB) [18] contributed to the success of feature-based registration methods.

Deep learning-based image registration leverages neural networks, particularly convolutional neural networks (CNNs) [19], to directly learn the transformation between two images, making it a powerful and end-to-end approach for image alignment. Unlike traditional methods that involve multiple stages like feature extraction and matching, deep learning-based methods learn the registration transformation in an end-to-end manner [20]. This means that the network takes raw images as input and outputs the transformation needed to align them optimally. These methods often employ spatial transformer networks (STNs) within the neural architecture to perform the image warping necessary for alignment [21]. STNs allow the network to learn how to transform one image to match the other. Deep learning-based registration can handle a wide range of image deformations, including nonlinear deformations and complex transformations. It is capable of aligning images even when traditional methods struggle due to these deformations.

2.2 Medical image segmentation

Medical image segmentation is a key research area within the field of medical image processing, aiming to accurately extract and identify regions of interest within medical images. Significant progress has been made in this field over the past few decades, providing robust support for clinical diagnosis, treatment planning, and medical research.

In the early stages of medical image segmentation, methods were primarily based on traditional image processing techniques such as thresholding and edge detection [22]. However, these methods struggled when faced with complex anatomical structures and irregularly shaped lesions, limiting their effectiveness in practical clinical applications.

The advent of deep learning technologies, particularly the success of convolutional neural networks (CNNs), has marked a significant breakthrough in medical image segmentation. By learning from vast amounts of medical image data, CNN can automatically extract and learn features [23], effectively overcoming the limitations of traditional methods. Deep learning applications in medical image segmentation, such as U-Net [24], SegNet [25], and other network architectures, have become mainstream methods. Furthermore, the development of medical image segmentation has benefited from advancements in multimodal image fusion [26] and cross-modal learning [27]. Integrating information from different medical imaging devices enhances the robustness and accuracy of segmentation algorithms. Cross-modal learning enables models to effectively learn representations from one modality's data for improved segmentation performance on another modality's data.

2.3 Vascular Flow Prediction

The prediction of vascular flow speed is an evolving area of research in the fields of medicine and biology. With advancements in technology and continuous innovation in research methods, scientists are able to more accurately predict and measure the velocity of blood flow in vessels, providing more precise information for disease diagnosis and treatment.

Initially, the estimation of blood flow speed relied heavily on fundamental hemodynamic theories, such as the Poisson equation [28] and the Navier-Stokes equation [29]. However, these theories

had limitations when faced with the complex physiological environment, such as the nonlinearity and non-uniformity of blood vessel walls.

With the advancement of imaging technologies, such as ultrasound, magnetic resonance imaging (MRI), and computed tomography (CT) [30], scientists can directly observe and measure blood flow velocity without relying on theoretical estimates. These high-resolution imaging techniques allow physicians and researchers to monitor changes in blood flow velocity in real-time [31] within living organisms and provide powerful tools for diagnosing cardiovascular diseases, assessing hemodynamics, and guiding surgical procedures.

On another front, the development of computational simulations and numerical models has also played a crucial role in predicting blood flow velocity. By using tools like computational fluid dynamics (CFD) [32], scientists can simulate and analyze the hemodynamic behavior under different conditions, predicting changes in blood flow velocity and providing in-depth insights for medical research and treatment.

3 Data

The dataset used in this paper is sourced from the University Hospital of Stavanger’s Department of Cardiology. It comprises seven sets of image sequences, with each sequence containing a variable number of time-series X-ray angiography images of the coronary arteries. The specific number of images in each sequence varies and is detailed in Table 1.

Set number	00	01	03	04	06	07	08
Number of images	42	32	47	47	58	57	58

Table 1. Dataset size

All these images were acquired using the same C-arm machine, although the exact model is undisclosed. Different machine parameters were set to capture images from different angles. Notably, specific information about the individuals being examined, such as age and gender, is unknown.

Each image sequence depicts the dispersion of a contrast agent within the coronary arteries. In adjacent images, variations can be attributed to the diffusion of the contrast fluid (Figure 3) and intrinsic bodily movements (Figure 4), such as heartbeats or respiration, as the machine’s capture angle remains fixed.




Figure 3. A selection of angiograms from set 07. The main differences observed between the images are primarily due to the diffusion of the contrast fluid.

Within this dataset, all images have are single-channel grayscale images with a specification of 736x736 pixels, and pixel values range from 0 to 255. Additionally, the dataset comes with annotations that correspond one-to-one with each image and were generated through a random forest segmentation method. These annotations are binary and also have dimensions of 736x736




Figure 4. A selection of angiograms from set 08. The main differences observed between the images are primarily due to intrinsic bodily movements, in this case, the heartbeat.

pixels, with pixel values indicating whether a particular area is part of a blood vessel or not. It's important to note that during the process of vascular angiography, some contrast agent may leak outside the blood vessels as it diffuses, resulting in a cloud of contrast fluid in the image. In the annotations, these contrast fluid clouds are also labeled as blood vessels, although this labeling may not be entirely accurate (Figure 5). A noteworthy point is that, in set 07, images ranging from frame 16 to 57, and in set 08, images from frame 16 to 58, are all accompanied by images post-registration. These registered images were produced by contributors from previous work on this dataset, and they were obtained through the application of SIFT on segmented images. The experiments in this paper primarily involve analyses conducted on these registered segmentation images.




Figure 5. Image 21 and its corresponding annotation from set 04. The coronary angiogram reveals contrast fluid clouds leaking into the body, and these contrast fluid clouds are marked as blood vessels in the annotation.

To investigate the distribution of images in each sequence and the quality of segmentation, we analyzed the trend of the proportion of the image area segmented as vessels with respect to the image sequence number in Figure 6. Since changes in the area of the vessel region are caused by the diffusion of the contrast agent, the alterations in the curves in the images can roughly reflect the projected blood flow velocity. The results show that all sequences exhibit a similar trend in blood flow, with a slower flow in the initial 10 frames. In the subsequent 10 frames, the contrast agent rapidly diffuses with the blood flow, resulting in a steep increase in the curve. In the final frames, as the contrast agent fills the vessels in the area captured by the machine, the proportion of the segmented vessel area relative to the whole no longer shows a significant upward trend, and the curve tends to flatten out.

In general, the area segmented as vessels in the coronary angiography images should increase over time. Considering the possibility of introducing various noises during the image acquisition




Figure 6. The trend of the proportion of the image area segmented as vessels with respect to the image number in each sequence.

process, some fluctuations in the vessel area proportion curve are acceptable. However, if these fluctuations become too pronounced at a specific frame, typically showing a significant drop followed by a rise, it indicates potential quality issues with the segmentation of the image at that point. For example, in the 40th frame of set 06, there is a noticeable fluctuation in the vessel area proportion, corresponding to a clear deficiency in the segmented image data at that location (see in Figure 7). However, there are no evident weaknesses in the original image that could significantly contribute to segmentation issues, suggesting the inadequacy of the segmentation algorithm itself as the possible cause of this issue.

Overall, the changes in the proportion of the image area segmented as vessels generally conform to a sigmoid function, and the gradient changes of the curve can serve as a rough estimate of blood flow velocity. To illustrate this, we took set03 as an example and conducted a sigmoid fitting, the results of which are shown in Figure 8.

4 Methods

In this section, we develop a comprehensive framework for estimating projected velocity and 3D velocity. Specifically, given any two registered segmented cardiac angiographic images within the same sequence, our algorithms will initially calculate the projected velocity of blood flow between these two frames. Subsequently, utilizing this estimated projected velocity and referencing the actual cardiac surface model, the algorithm will transform it into 3D velocity.

Due to the high contrast, contrast agent are more prominent in the images. Consequently, segmented images essentially entail the segmentation of the contrast agent rather than the segmentation of vessels or blood. However, this does not imply that the diffusion speed of the contrast agent in the angiographic image sequence can fully reflect the blood flow velocity. In reality, the diffusion speed is influenced by various factors, including blood flow velocity, vascular permeability, tissue uptake, and clearance of the contrast agent. Building a comprehensive mapping model is too complex and beyond the scope of this study. To simplify the model, we make the following assumption: the diffusion speed of the contrast agent is equal to the true blood flow velocity. Thus, we transform the prediction of blood flow velocity into detecting the diffusion speed of the contrast agent in the vessels. To avoid confusion, unless specified otherwise, both terms in the following text will refer to blood flow velocity.




Figure 7. The original images from frames 38 to 42 in Set 06 and their corresponding segmented images. Despite the absence of any apparent quality issues in the original images, the segmented image at the 40th frame exhibits significant deficiencies.




Figure 8. The curve depicting the variation of the proportion of the segmented vessel area with respect to the image sequence number in set 03, along with the sigmoid function fitted based on it.

In addition, because of intrinsic motion and noise, the change in the angiographic image sequence is not solely caused by the diffusion of the contrast agent. We assume that the quality of registration is high enough so that intrinsic bodily movements of patients, such as heartbeat and respiration, do not introduce errors in the observed vascular angiographic images. This implies that, within an image sequence, the coordinates of all vessels that appeared in previous frames remain unchanged in subsequent frames, and all newly appearing vessel pixels are solely due to the diffusion of the contrast agent.

4.1 Projected velocity estimation

Angiographic images are projections of the heart onto the imaging plane; therefore, the diffusion speed exhibited by the contrast agent on the projection plane is also the velocity after the projection transformation. Given a series of vascular angiographic images, the smallest observable change is the interval between adjacent frames. According to the definition,

$$\text{Velocity} = \frac{\text{Displacement}}{\text{Time}}$$

It is important to note that when the term "velocity" is used, it typically refers to a vector, where the direction of the vector represents the direction of object movement. Considering that in real-world scenarios, the flow direction of blood in vessels is complex and variable, we make a rough assumption that its direction is along the vessel. In the remainder of this paper, the term "velocity" specifically denotes the magnitude (i.e., speed) of the corresponding vector, and the direction will not be explicitly specified.

Assuming we want to predict the blood flow velocity of branch n between frames i and j , then the time taken is the difference in image indices between the two frames, i.e., $|i - j|$ (units in frames). And the displacement refers to the diffusion distance of the contrast agent along the blood vessel in branch n between the two frames i and j , which is reflected in the geometric differences of branch n shown in the angiographic images. This geometric difference can be translated into the difference in the locations of the "endpoints" of branch n in each frame (See in Figure 9). We can quantify the geometric difference by measuring the difference in the coordinate positions of endpoints. That is, the blood flow velocity of branch n between frame i and j can be estimated as


$$v_{i,j}^n = \frac{D(\text{Pos}_i^n, \text{Pos}_j^n)}{|i - j|}$$

where Pos_k^m represents the endpoint of branch m in frame k , and the distance function $D(\text{Pos}_1, \text{Pos}_2)$ denotes the path length along the vessel from point Pos_1 to point Pos_2 . It is noteworthy that, in this paper, we consider pixels as the smallest units, and employ the Manhattan distance as the distance calculation method. In this scenario, all potential displacements are integer values. The advantage of using the Manhattan distance is its ease of computation in the proposed algorithm later in the paper, but it also implies a restriction on the blood flow direction to four directions. When the distance between two points is diagonally large, this may introduce some errors.

Considering the ideal scenario where the segmented vessel region in the image is a complete and connected area, the displacement between points Pos_1 and Pos_2 can be *approximated* as the difference in displacement between points Pos_1 and Pos^{root} (the beginning point of the vessels), and Pos_2 and Pos^{root} , that is,

$$v_{i,j}^n = \frac{|D(\text{Pos}^{\text{root}}, \text{Pos}_j^n) - D(\text{Pos}^{\text{root}}, \text{Pos}_i^n)|}{|i - j|} \quad (1)$$

The above formula calculates the units in pixels/frame. Up to this point, we have transformed the prediction of blood flow in each branch into a comparison of the variation in distances between branch endpoints and the root point of the vessel. Thus, we need to design a suitable distance function D to measure the displacement between two points. A straightforward solution involves treating the segmented images as maps and employing a search algorithm (such as breadth-first



(a) Segmented image 18 from set 07



(b) Segmented image 19 from set 07

Figure 9. Segmented images 18 and 19 from set 07. The white annotations represent the "end-points" of each branch in this frame, indicating the farthest extent reached by the contrast agent within that branch at this moment. The differences in endpoints between different frames can be used to measure displacement.

search (BFS) or Dijkstra's algorithm) to traverse the entire image and seek the minimum distance between two points. However, as the size of the map increases, the runtime of both BFS and Dijkstra's algorithm will significantly decrease. To enhance algorithm efficiency, we have designed an algorithm based on the A* algorithm to accomplish the calculation of distances at various positions within the vessel.

Due to the diffusion of the contrast agent along the vessels, we can consider the segmented vessel part as the "passable region" for the contrast agent, while other tissue regions are the "impassable region." Therefore, a reasonable distance function $D(\text{Pos}_1, \text{Pos}_2)$ can be designed as follows: the shortest (optimal) path length from Pos_1 to Pos_2 in the passable region. If we consider each segmented vessel pixel as the smallest node, the ideal properties of the segmented vessel region are as follows:

1. The problem domain is finite.
2. For any starting point and endpoint, there exists a shortest (optimal) path from the starting point to the endpoint.
3. The cost from any node to its child nodes (adjacent nodes) is greater than 0.

Consider the following heuristic function design:

$$f(n) = g(n) + h(n)$$

where $f(n)$ is the evaluation of each potential probing point, consisting of two parts: the first part is $g(n)$, representing the cost from the initial search point to the current point; the other part, the heuristic function $h(n)$, represents the estimate from the current node to the goal node. Searching for the optimal heuristic function is not the focus of this paper, so we simply design $h(n)$ as the Manhattan distance. It is clear that in a two-dimensional space, the following holds:

$$h(n) \leq h^*(n)$$

where $h^*(n)$ is the actual cost value of the problem, which is unknown during the search process. This means that a heuristic algorithm with the strategy $f(n) = g(n) + h(n)$ can become the A* algorithm and will definitely find the optimal solution. It is worth noting that the smaller the




Figure 10. Schematic diagram of the heuristic function. The green dots represents the start point of path, the red dot represents the goal point of path, the blue dot represents the furthest point n of the currently traversed optimal path, and the black square represents impassable areas. The length of the blue solid line is the cost from the initial search point to the current point $g(n)$, the length of the red dashed line is the estimate from the current node to the goal node $h(x)$, and the length of the green dashed line is the possible optimal path length from the current node to the goal node $h^*(x)$.

difference between $h(n)$ and $h^*(n)$, the stronger the discriminative ability and heuristic ability of $h(n)$, and the higher the algorithm efficiency.

Based on the aforementioned theory, we designed the distance indexing algorithm, as shown in Algorithm 1 in Appendix A. By traversing vessel pixels, we generated a distance matrix with the same dimensions as the original image. In this matrix, each coordinate corresponding to a vessel pixel stores the distance from that point to the root point, calculated through the A*-based algorithm. Non-vessel pixels are assigned a distance of -1, indicating non-passability. This distance matrix serves as the return value for the function D . Subsequently, by retrieving the distance matrix, we can calculate blood flow velocity between different frames using (1).

However, in most cases, a single angiographic image will exhibit multiple vascular branches, and the endpoints of these branches will have different coordinates across different frames. Therefore, it is necessary to design an algorithm to detect the coordinates of endpoints for all branches in one angiographic image. This information can then be used as input for the distance function D . Considering that the endpoints of branches have the property that their values in the distance matrix are always greater than the values of the surrounding elements, an intuitive solution is to use distances as the measuring object and search for local maxima in the distance matrix. By setting an appropriate threshold for the size of the local region, we can filter out the coordinates of suitable local maxima, which correspond to the positions of endpoints. The specific implementation can be found in Algorithm 2 in Appendix B.

Thus, the complete process of our designed projected velocity estimation is illustrated in Figure 11.




Figure 11. The process of our designed projected velocity estimation.

4.2 3D velocity estimation

We continue to estimate 3D velocity based on the calculated projected velocity from the previous subsection. At the same time, we maintain the assumption that the applied registration has completely eliminated any intrinsic movement of the patient. This allows us to consider the angiographic image sequence as originating from a motionless patient. Under this assumption, the positions of the heart and the projection plane are fixed during the image sequence acquisition process. In other words, within the same sequence group, the distance from the same part of the coronary arteries to the projection plane remains constant. It ensures a fixed projection angle and enables us to proceed with the conversion from the projected speed to the 3D real speed.

For the fixed heart and projection plane, we use distance lines as an auxiliary tool for estimating 3D speed. Unlike contour lines in geography, which have equal height decrease between two adjacent lines, we use distance lines (see in Figure 12) where the perpendicular distance along the heart's surface between two adjacent lines is constant. This property grants the distance lines the following characteristics:

1. The shortest distance line is a single point, and this point is the one on the surface of the heart closest to the projected plane. We refer to this point as the "peak point".
2. On each line, the distance along heart surface from every point to the peak point is constant, and the difference in this distance between any two adjacent distance lines is constant.
3. Any two distance lines do not intersect.
4. All distance lines are closed curves, although their projections on the image plane may not be. This might be due to the image not containing the complete outline of the heart or some distance lines being on the back of the heart from the projection plane's perspective.
5. The larger the angle between the heart surface and the projected plane, the denser the projection of distance lines on the projected plane. This implies that the projected velocity at these locations is more influenced by the projection.

Assuming the true three-dimensional velocity of blood within the vessels on the heart's surface is v , with its direction tangential to the local surface of the heart. Let the component of velocity along the distance line be denoted as v_h , and the component along the perpendicular direction to the distance line as v_p . Due to their orthogonal directions, mathematically, we have

$$v^2 = v_h^2 + v_p^2$$

When multiplied by the square of time, we obtain

$$D^2 = d_h^2 + d_p^2$$




Figure 12. One possible set of distance lines. The black lines represent the contour of the heart, with solid lines indicating the visible portion on the projected plane and dashed lines representing the invisible portion. The orange lines denote distance lines, where adjacent distance lines have equal distances along the surface of the heart. The projected distance on the projected plane for each of these lines is denoted as l_i .

where D represents the distance traveled by the blood during this time period, d_h is the component of this distance along the direction of the distance line, and d_p is the component along the perpendicular direction to the distance line.

Assuming that, from the perspective of the projection surface (i.e., the image), the blood flow paths during this time period span n distance lines, and when the distance lines are sufficiently dense, we can *approximate* that

$$d_p = nd \quad \text{and} \quad d_h = \sqrt{D_p^2 - \left(\sum_{k=1}^n l_k\right)^2}$$

Here, d is the spacing along the heart surface between adjacent distance lines, D_p is the projected distance estimated in the previous section, l_k is the projection spacing between the k -th distance line and the $(k+1)$ -th distance line on the projected plane (see in Figure 12), and $\sum_{k=1}^n l_k$ is the sum of projection spacing of the distance lines spanned by the blood flow path.

Therefore, based on the estimated projected distance from the previous section, we can calculate the 3D velocity between frame i and frame j using the following formula (units in pixels/frame):

$$v_{i,j} = \frac{\sqrt{D_p^2 + (nd)^2 - \left(\sum_{k=1}^n l_k\right)^2}}{|i-j|} \quad (2)$$

We can observe that when the heart surface forms a small angle with the projected plane, i.e., when $d \approx l_k$, the influence from projection changes is very small and the term $(nd)^2 - \left(\sum_{k=1}^n l_k\right)^2$ in the square root of the formula is approximately zero. This implies that the projected velocity is close to the 3D velocity. Conversely, when the angle is large, l_k will be significantly smaller than d , indicating a substantial impact from projection changes. In this case, the value of the square




Figure 13. Segmented images after registration. The first row corresponds to the first five images from set 07, while the second row depicts the first five images from Set 08. It is evident that the image quality in set 07 is noticeably superior to that in set 08.

root term in the formula is notably different from D_p , suggesting a strong correction is needed to convert the projected velocity to the 3D velocity.

5 Experiments

The experiments in this paper are grounded on two assumptions outlined in the methods section:

1. The diffusion velocity of the contrast agent in the blood equals the blood flow velocity.
2. Registration has effectively eliminated potential noise introduced by intrinsic movements and other factors.

Consequently, with registered segmentation images, set 07 and set 08 have emerged as viable sub-datasets. However, between the two, the image sequence in set 07 exhibits noticeably superior quality upon visual inspection compared to set 08 (see in Figure 13). Therefore, set 07 is chosen as the primary dataset for the experiments.

5.1 Projected velocity estimation

To estimate projected velocity, we implemented the process outlined in Figure 11 on set 07. Specifically, for each pair of consecutive frames, we sequentially computed the distance matrix based on them. Subsequently, we detected the coordinates of all potential branch endpoints in each frame and utilized them for the estimation of projected velocity.

For the sake of research, we visualized the original segmented image after computing the distances matrix (see in Figure 14). Specifically, for the regions segmented as the body, we kept them black, while for the vessel areas, we assigned blue to red labels based on their corresponding distances in the distances matrix. We normalized the distance matrix before visualization, ensuring that the color always transitions from (0,0,255) to (255,0,0) from the root pixel to the pixel with the maximum distance in each image. The visualization provides an intuitive reflection of how our algorithm operates. It is observable that the depths of different branches vary.



Figure 14. Distance visualization of the 18th segmented image in set 07. The color gradient from blue to red represents the distance from the root pixel, indicating proximity to distance.

In the data section, we mentioned that the variation in the proportion of the vessel area in the sequence might reflect the blood flow speed to some extent. We aimed to explore this possibility in detail. Therefore, we fitted a sigmoidal curve (see in Figure 15) to the changing proportion of segmented vessel areas in the registered set07 and obtained the following results:

$$y = -\frac{0.08}{1 + e^{0.25x - 4.75}} + 0.09$$

where y represents the proportion of the vessel area, and x represents the image number.

Generally, for the same vessel, the blood flow velocity should remain relatively stable over a period. This implies that the proportion change curve of the segmented vessel region in Figure 15 should exhibit a consistently increasing trend over time. However, the actual situation portrayed by the dataset deviates from this expectation. In fact, the curve corresponding to set 07 shows a trend of initially gradual increase, followed by a steep rise, and ultimately returning to a gentle or even stagnant trend. The initial and subsequent transitional segments of the curve are contradictory to the theoretical expectations and it suggests that if velocity estimation is based on the entire image sequence, the results obtained may incur significant errors. To account for the observed shape of the actual curve, we offer the following reasonable conjectures:

In the initial segment of the curve, corresponding to frames 1 to 15 of the image sequence, the contrast agent has just been introduced into the vessels. During this period, the contrast agent gradually accelerates from a stationary state until its velocity aligns with that of the blood. This may be due to molecular interactions with the blood or artificial interventions (such as the injection process by doctors). This phase is reflected in the curve as a gradual transition from a gentle slope to a steep incline. Subsequently, in frames 16 to 25 of the image sequence, the contrast agent rapidly disperses through the vessels at the blood flow velocity. This stage is evident in the curve as a sharp upward surge in the middle portion. Finally, after frame 25 of the image sequence, the contrast agent has almost completely filled the vessels within the imaging range of the device. At this point, although blood continues to flow, the proportion of the contrast agent in the image no longer changes. Consequently, the curve exhibits a final plateau or even a cessation of increase.

Based on the analysis above, the segments corresponding to the initial and final parts of the curve, i.e., the image frames before frame 16 and after frame 25, are not suitable for velocity estimation. This is because they contradict our assumption 1, which states that the diffusion velocity of the contrast agent in the blood equals the blood flow velocity. In contrast, frames 15 to




Figure 15. The curve depicting the variation of the proportion of the segmented vessel area with respect to the image sequence number in the registered set 07, along with the sigmoid function fitted based on it.

25 appear to more accurately reflect the blood flow velocity. Consequently, subsequent experiments were conducted on the selected frames to obtain more reliable velocity estimations.

For frames 15 to 25, we conducted branch endpoint detection separately. There are two main branches, and their pixel coordinates in each image, along with the estimated projected velocity obtained based on the distance matrix indices, are presented in the Table 2:

Frame number	Branch 1	Velocity 1	Branch 2	Velocity 2
16	(147, 467)	\	(397, 508)	\
17	(219, 523)	114	(401, 505)	33
18	(188, 533)	93	(416, 498)	76
19	(197, 530)	72	(543, 561)	262
20	(413, 700)	400	(651, 621)	202
21	(517, 703)	163	(720, 644)	134
22	-	\	(722, 645)	67
23	(448, 719)	\	(715, 644)	57
24	(431, 718)	20	(679, 642)	10
25	(442, 716)	65	(660, 644)	53

Table 2. The coordinates of the endpoints of the two main branches in frames 16 to 25 and their estimated projected velocity (in pixels/frame) between adjacent frames. The symbol '-' indicates that the endpoint could not be detected due to image quality issues, while '\' signifies that the velocity retrieval from the distance matrix was not possible due to missing endpoint coordinates. Additional bias has been introduced to prevent the occurrence of negative values.

From the results in Table 2, we can observe that the trend of blood flow velocity changes in the two branches aligns with our predictions in the previous text. It is noteworthy that the estimated projected velocity of branch 1 between frames 19 and 20 is significantly higher than its velocities before and after this frame. For this outlier, we examined the original images and segmented images corresponding to it. In the original angiographic image, the contrast agent's diffusion in branch 1 is complex. Specifically, a small portion of the contrast agent rapidly diffuses into most




Figure 16. The derivative of the fitted sigmoid function and the variation of the estimated projected blood flow velocities for the two main branches with respect to frame number. To align their vertical scales, a calibration has been applied. Specifically, we multiplied the derivative of the sigmoid function by the image area (in pixels) and the blood flow velocity of each branch by the vessel width (in pixels). The vertical axis of the figure is in units of $\text{pixel}^2/\text{frame}$.

of the area of branch 1, but it appears faint in the original image. Therefore, it is not labeled as a vessel region in the segmentation until more contrast agent fills this area. Thus, in reality, the projected velocity of branch 1 between frames 19 and 20 is slightly lower than the estimated value in the table. Considering this, we regard this kind of error as acceptable.

Afterward, we compared the derivative of the sigmoid function and the projected speeds of the two main branches (Figure 16). To ensure comparability on the y-axis, we adjusted them as follows: we multiplied the derivative by the area of the image, and we multiplied the projected speed by the width of the vessels (in pixels).

From Figure 16, it is evident that the changing trends of the estimated projected blood flow velocities for the two branches align with the variations in the overall vessel region proportion in the images. Specifically, all three exhibit a peak around frame 18, with rising and falling trends on the left and right sides, respectively. Considering that the global vessel region increases with the diffusion of the contrast agent in the vessels, this consistency is in line with expectations.

It is worth noting that the dataset used in this experiment lacks ground truth information regarding the blood flow velocities in each branch. Therefore, assessing the algorithm’s performance by comparing estimated velocities with actual values is not feasible. However, the algorithm we designed does not directly rely on any information about the vessel region area. In fact, it focuses solely on the optimal path variations between two points in the image map. This means it utilizes less information than the vessel region proportion. Despite this, the algorithm’s output still maintains a certain consistency with the latter, which can to some extent indicate the accuracy and reliability of the algorithm.

5.2 3D velocity estimation

Our 3D velocity estimation is based on projected velocity. Therefore, this part of the experiment is also conducted on set 07. Since we have assumed the effectiveness of registration, a fixed set of distance lines can be applied to the entire image sequence. Specifically, we will use (2) to calculate the 3D velocity, and the prerequisite for this equation is that the distance lines are dense enough. This requires that, within a certain time period, the length of the blood flow path should not be




Figure 17. Our plotted set of distance lines based on the real heart model. The left image shows them overlaid on the angiographic image from frame 20 of set 07, while the right image displays the overlay on its corresponding aligned segmented image. In the right image, the blue regions represent the paths of contrast agent diffusion in branch 1 between adjacent frames, and the red regions represent the paths of contrast agent diffusion in branch 2 between adjacent frames.

significantly smaller than the distance d between adjacent distance lines and the blood flow path needs to span a certain number of distance lines. Therefore, based on the results in Table 2, we selected frame 20 as an example for the experiment.

To conduct the experiment, we hand-drew a set of feasible distance lines based on the surface model of the actual heart. These distance lines were then overlaid onto the angiographic image of frame 20, corresponding to the aligned segmented image. Based on the earlier experimental results, the blood flow projected velocity in branch 1 at frame 20 is 400 pixels/frame (blue region in Figure 17b), while the blood flow projected velocity in branch 2 is 202 pixels/frame (red region in Figure 17b).

Since we drew the distance lines based on a reference to the real heart model, this implies that the measurement parameters of these distance lines are in real distances (specifically, in millimeters). However, the projected velocity we calculated is based on image pixels, so it is necessary to establish a conversion relationship between millimeters and pixels. In general, the diameter of the coronary arteries is typically around 4 millimeters [33] and the width of the human heart is approximately 80 millimeters [34]. Based on the given data, we designed the spacing between our distance lines to be 10 millimeters. Assuming the scale ratio between pixels and actual distances can be determined by the vessel width and the actual vessel width, then 1 pixel in the image corresponds to approximately 0.2 millimeters. Applying (2), for branch 1, we calculate its 3D velocity as

$$v_{19,20}^1 = \frac{\sqrt{(400 * 0.2)^2 + (3 * 10)^2 - ((37 + 37 + 18) * 0.2)^2}}{20 - 19} = 83.4 \text{ mm/frame}$$

and for branch 2, we calculate its 3D velocity as

$$v_{19,20}^2 = \frac{\sqrt{(202 * 0.2)^2 + (3 * 10)^2 - ((60 + 55 + 25) * 0.2)^2}}{20 - 19} = 41.8 \text{ mm/frame}$$

At the same time, we can calculate the projected blood flow velocities for branch 1 and branch 2 between frames 19 and 20, which are $400 * 0.2 = 80 \text{ mm/frame}$ and $202 * 0.2 = 40.4 \text{ mm/frame}$,

respectively. From the results, it appears that the 3D blood flow velocities in both branches are slightly greater than their projected velocities. This implies that, through the process of projection transformation, there is some degree of loss in the apparent path lengths on the projection plane for both branches. However, this impact is relatively minor. Specifically, for vessel branch 1, its direction is mainly parallel to the distance lines, suggesting that it has not descended significantly in the vertical direction. For vessel branch 2, although it spans many distance lines, the projection spacing between the distance lines on the projection surface is considerable. This indicates that the surface of the heart in this region is relatively flat, meaning it has not descended significantly in the vertical direction either.

Generally, the closer to the edge of the heart, the more densely packed the distance lines are. If a branch continuously spans multiple closely spaced distance lines at the edge, it indicates that it has traveled a greater distance in the vertical direction. However, from the perspective of the projection plane, this distance may appear short. Conversely, if a branch does not cross any distance lines or the projected spacing of the distance lines it traverses is large (i.e., $d \approx l$), it indicates that the branch has descended or ascended only a small distance in the vertical direction. In such cases, the projected velocity is approximately equal to the 3D velocity.

6 Conclusion

In this paper, we proposed a proof-of-concept framework encompassing multiple algorithms to estimate blood flow velocity. Specifically, a plane distance calculation algorithm based on A* was employed to estimate projected velocity, and an algorithm based on distance lines was used to convert projected velocity into 3D velocity. Furthermore, we applied the framework to cardiac angiography image sequences from the Department of Cardiology at the University Hospital of Stavanger, demonstrating its feasibility.

Several aspects of future work could potentially enhance the performance of the framework:

More Detailed Datasets The existing datasets providing angiographic images are limited in their scope. Utilizing datasets with more detailed and diverse scenarios, such as smaller imaging intervals, higher resolutions, additional patient information, and variations in scanner parameters, could offer a more comprehensive evaluation of the framework’s performance.

Improved Registration and Segmentation Methods The framework in this paper relies on the assumption that registration is effective enough, and segmentation is accurate. The quality of both processes directly influences the accuracy of velocity estimation. Enhancements in the methods for image registration and segmentation, such as incorporating deep learning techniques, can contribute to more accurate and reliable results.

Refined Distance Lines Distance lines serve as a direct reference for estimating 3D velocity and play a crucial role in evaluating the velocity loss in the vertical projection process. Improving the precision of distance lines, possibly through advanced modeling or selection methods, may lead to more accurate velocity estimations.

Enhanced Distance Metric The choice of distance metric determines the standard for velocity calculation. Exploring and adopting more sophisticated distance metrics, such as Euclidean distance or Minkowski distance, could refine the accuracy of distance calculations and potentially mitigate errors introduced by the limitations of Manhattan distance.

Appendix A

Algorithm 1 A*-based distance matrix generation

```
1: function DISTANCE_MATRIX_GENERATION(image, root_point)
2:   rows, cols  $\leftarrow$  image.size
3:   distances  $\leftarrow$   $[[\infty] * \text{cols for } i \text{ in range}(\text{rows})]$ 
4:   heap  $\leftarrow$  [(0, start_point)]
5:   distances[start_point[0]][start_point[1]]  $\leftarrow$  0
6:   while heap do
7:     distance, (row, col)  $\leftarrow$  heap.pop()
8:     for dr, dc in [(-1, 0), (1, 0), (0, -1), (0, 1)] do
9:       new_row, new_col  $\leftarrow$  row + dr, col + dc
10:      if 0  $\leq$  new_row  $<$  rows and 0  $\leq$  new_col  $<$  cols and
        map_array[new_row][new_col] == 1 then
11:        new_distance  $\leftarrow$  distances[row][col] + 1
12:        if new_distance  $<$  distances[new_row][new_col] then
13:          distances[new_row][new_col]  $\leftarrow$  new_distance
14:          priority  $\leftarrow$  new_distance + heuristic((new_row, new_col), start_point)
15:          heap.pop(priority, (new_row, new_col))
16:        end if
17:      end if
18:    end for
19:  end while
20:  distances[distances ==  $\infty$ ]  $\leftarrow$  -1
21:  return distances
22: end function
```

Appendix B

Algorithm 2 Branch endpoint detection

```
1: function FIND_BRANCH_ENDPOINTS(image_array, root_point, threshold)
2:   distances  $\leftarrow$  Distance_Matrix_Generation(image_array, root_point)
3:   endpoints  $\leftarrow$  []
4:   rows, cols  $\leftarrow$  distances.size
5:   for  $i \leftarrow 0$  to rows do
6:     for  $j \leftarrow 0$  to cols do
7:       if distances[ $i$ ][ $j$ ]  $\geq 0$  then
8:         surrounding  $\leftarrow$  distances[max(0,  $i - \text{threshold}$ ) : min(rows,  $i + \text{threshold} + 1$ ),
9:                               max(0,  $j - \text{threshold}$ ) : min(cols,  $j + \text{threshold} + 1$ )]
10:        if max(surrounding) == distances[ $i$ ][ $j$ ] then
11:          endpoints.append(( $i, j$ ))
12:        end if
13:      end if
14:    end for
15:  end for
16:  return endpoints
17: end function
```

References

- [1] M. Litviňuková, C. Talavera-López, H. Maatz, D. Reichart, C. L. Worth, E. L. Lindberg, M. Kanda, K. Polanski, M. Heinig, M. Lee *et al.*, “Cells of the adult human heart,” *Nature*, vol. 588, no. 7838, pp. 466–472, 2020.
- [2] A. M. Katz, *Physiology of the Heart*. Lippincott Williams & Wilkins, 2010.
- [3] S. D. Joshi, S. S. Joshi, and S. A. Athavale, “Origins of the coronary arteries and their significance,” *Clinics*, vol. 65, pp. 79–84, 2010.
- [4] P. Libby and P. Theroux, “Pathophysiology of coronary artery disease,” *Circulation*, vol. 111, no. 25, pp. 3481–3488, 2005.
- [5] M. Miyazaki and V. S. Lee, “Nonenhanced mr angiography,” *Radiology*, vol. 248, no. 1, pp. 20–43, 2008.
- [6] N. Navab, S.-M. Heining, and J. Traub, “Camera augmented mobile c-arm (came): calibration, accuracy study, and clinical applications,” *IEEE transactions on medical imaging*, vol. 29, no. 7, pp. 1412–1423, 2009.
- [7] H. R. Fazlali, N. Karimi, S. M. R. Soroushmehr, S. Sinha, S. Samavi, B. Nallamotheu, and K. Najarian, “Vessel region detection in coronary x-ray angiograms,” in *2015 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2015, pp. 1493–1497.
- [8] R. Collins, J. Burch, G. Cranny, R. Aguiar-Ibáñez, D. Craig, K. Wright, E. Berry, M. Gough, J. Kleijnen, and M. Westwood, “Duplex ultrasonography, magnetic resonance angiography, and computed tomography angiography for diagnosis and assessment of symptomatic, lower limb peripheral arterial disease: systematic review,” *Bmj*, vol. 334, no. 7606, p. 1257, 2007.
- [9] T. T. Hormel, T. S. Hwang, S. T. Bailey, D. J. Wilson, D. Huang, and Y. Jia, “Artificial intelligence in oct angiography,” *Progress in Retinal and Eye Research*, vol. 85, p. 100965, 2021.
- [10] L. G. Brown, “A survey of image registration techniques,” *ACM computing surveys (CSUR)*, vol. 24, no. 4, pp. 325–376, 1992.
- [11] J.-W. Hsieh, H.-Y. M. Liao, K.-C. Fan, M.-T. Ko, and Y.-P. Hung, “Image registration using a new edge-based approach,” *Computer vision and image understanding*, vol. 67, no. 2, pp. 112–130, 1997.
- [12] X. Dai and S. Khorram, “A feature-based image registration algorithm using improved chain-code representation combined with invariant moments,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 5, pp. 2351–2362, 1999.
- [13] Y. Pei, H. Wu, J. Yu, and G. Cai, “Effective image registration based on improved harris corner detection,” in *2010 International Conference on Information, Networking and Automation (ICINA)*, vol. 1. IEEE, 2010, pp. V1–93.
- [14] G. Shrivakshan and C. Chandrasekar, “A comparison of various edge detection techniques used in image processing,” *International Journal of Computer Science Issues (IJCSI)*, vol. 9, no. 5, p. 269, 2012.
- [15] Y. Wu, W. Ma, M. Gong, L. Su, and L. Jiao, “A novel point-matching algorithm based on fast sample consensus for image registration,” *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 1, pp. 43–47, 2014.
- [16] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International journal of computer vision*, vol. 60, pp. 91–110, 2004.

- [17] H. Bay, T. Tuytelaars, and L. Van Gool, “Surf: Speeded up robust features,” in *Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006. Proceedings, Part I* 9. Springer, 2006, pp. 404–417.
- [18] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “Orb: An efficient alternative to sift or surf,” in *2011 International conference on computer vision*. Ieee, 2011, pp. 2564–2571.
- [19] S. Miao, Z. J. Wang, and R. Liao, “A convolutional neural network approach for 2d/3d medical image registration,” *CoRR abs/1507.07505*, 2015.
- [20] B. D. De Vos, F. F. Berendsen, M. A. Viergever, M. Staring, and I. Išgum, “End-to-end unsupervised deformable image registration with a convolutional neural network,” in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings* 3. Springer, 2017, pp. 204–212.
- [21] M. C. Lee, O. Oktay, A. Schuh, M. Schaap, and B. Glocker, “Image-and-spatial transformer networks for structure-guided image registration,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II* 22. Springer, 2019, pp. 337–345.
- [22] J. Rogowska, “Overview and fundamentals of medical image segmentation,” *Handbook of medical imaging, processing and analysis*, pp. 69–85, 2000.
- [23] L. Lu, Y. Zheng, G. Carneiro, and L. Yang, “Deep learning and convolutional neural networks for medical image computing,” *Advances in computer vision and pattern recognition*, vol. 10, pp. 978–3, 2017.
- [24] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. Springer, 2015, pp. 234–241.
- [25] V. Badrinarayanan, A. Kendall, and R. Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [26] Z. Guo, X. Li, H. Huang, N. Guo, and Q. Li, “Deep learning-based image segmentation on multimodal medical imaging,” *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 3, no. 2, pp. 162–169, 2019.
- [27] C. Li, H. Sun, Z. Liu, M. Wang, H. Zheng, and S. Wang, “Learning cross-modal deep representations for multi-modal mr image segmentation,” in *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II* 22. Springer, 2019, pp. 57–65.
- [28] N. Shahcheraghi, H. Dwyer, A. Cheer, A. Barakat, and T. Rutaganira, “Unsteady and three-dimensional simulation of blood flow in the human aortic arch,” *J. Biomech. Eng.*, vol. 124, no. 4, pp. 378–387, 2002.
- [29] J. Labadin and A. Ahmadi, “Mathematical modeling of the arterial blood flow,” in *Proceedings of the 2nd IMT-GT Regional Conference on Mathematics, Statistics and Applications, Universiti Sains Malaysia, Penang*, 2006.
- [30] J. Lantz, L. Henriksson, A. Persson, M. Karlsson, and T. Ebberts, “Patient-specific simulation of cardiac blood flow from high-resolution computed tomography,” *Journal of biomechanical engineering*, vol. 138, no. 12, p. 121004, 2016.
- [31] S. Sankaran, D. Lesage, R. Tombropoulos, N. Xiao, H. J. Kim, D. Spain, M. Schaap, and C. A. Taylor, “Physics driven real-time blood flow simulations,” *Computer Methods in Applied Mechanics and Engineering*, vol. 364, p. 112963, 2020.

- [32] M. Behbahani, M. Behr, M. Hormes, U. Steinseifer, D. Arora, O. Coronado, and M. Pasquali, "A review of computational fluid dynamics analysis of blood pumps," *European Journal of Applied Mathematics*, vol. 20, no. 4, pp. 363–397, 2009.
- [33] J. T. Dodge Jr, B. G. Brown, E. L. Bolson, and H. T. Dodge, "Lumen diameter of normal human coronary arteries. influence of age, sex, anatomic variation, and left ventricular hypertrophy or dilation." *Circulation*, vol. 86, no. 1, pp. 232–246, 1992.
- [34] J. G. Betts, P. Desaix, E. Johnson, J. Johnson, O. Korol, D. Kruse, B. Poe, J. Wise, M. Womble, and K. Young, "Anatomy and physiology. openstax," *Rice University: Houston, TX, USA*, 2013.