

Visualization of Myocardial Perfusion Derived from Coronary Anatomy

Maurice Termeer, Javier Oliván Bescós, Marcel Breeuwer, Anna Vilanova, *Member, IEEE*,
Frans Gerritsen, M. Eduard Gröller, *Member, IEEE* and Eike Nagel

Abstract—Visually assessing the effect of the coronary artery anatomy on the perfusion of the heart muscle in patients with coronary artery disease remains a challenging task. We explore the feasibility of visualizing this effect on perfusion using a numerical approach. We perform a computational simulation of the way blood is perfused throughout the myocardium purely based on information from a three-dimensional anatomical tomographic scan. The results are subsequently visualized using both three-dimensional visualizations and bull's eye plots, partially inspired by approaches currently common in medical practice. Our approach results in a comprehensive visualization of the coronary anatomy that compares well to visualizations commonly used for other scanning technologies. We demonstrate techniques giving detailed insight in blood supply, coronary territories and feeding coronary arteries of a selected region. We demonstrate the advantages of our approach through visualizations that show information which commonly cannot be directly observed in scanning data, such as a separate visualization of the supply from each coronary artery. We thus show that the results of a computational simulation can be effectively visualized and facilitate visually correlating these results to for example perfusion data.

Index Terms—Cardiac visualization, coronary artery territories, myocardial perfusion.

1 Introduction

Coronary Artery Disease (CAD) is a condition where one or more of the coronary arteries have developed a stenosis, a narrowing of the vessel that may hamper or prevent blood from flowing through. The coronary arteries are small vessels supplying the heart muscle tissue, the myocardium, with blood. Figure 4 gives an overview of the various parts of the heart. In current clinical practice, X-ray angiography is the reference standard used for the assessment of CAD. Besides two-dimensional X-ray angiography, three-dimensional tomographic techniques including Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) can contribute to the diagnosis of CAD. Information on the hemodynamic significance of a stenosis can be obtained from a stress test, e.g. stress echocardiography or dobutamine stress MRI.

In the field of cardiac MRI, various scanning protocols have been developed to image different aspects of the heart. Among common scanning protocols, a whole heart scan provides anatomical information. A cine scan provides functional information by showing the heart at different time intervals of the cardiac cycle. A late enhancement scan provides information on tissue viability or scar by imaging areas where a previously injected contrast agent has accumulated. A perfusion scan provides information on how blood flows throughout the myocardium. During the analysis the information provided by several of these scanning protocols can be combined for a more comprehensive diagnosis. In practice only those scanning protocols that are relevant to the condition of the patient are used. Good correlations between MRI perfusion imaging, X-ray angiography and SPECT have been reported [21, 15].

Due to its higher resolution, cardiac CT can image the coronary arteries with more detail than cardiac MRI. Although perfusion and functional scans are possible with CT, they are rarely used in practice due to a range of issues, including heart beat limitations, a lower scanning speed and radiation.

CAD may change the way blood perfuses throughout the myocardium, which in turn may cause decreased motion of the heart muscle. In severe cases even scar, the death of part of the myocardium, may occur. A whole heart MRI or CT scan images the anatomical defect directly, while scanning protocols such as cine, perfusion and late enhancement image the effect of this defect on the functioning of the heart. The reason these latter scanning protocols are used is that it is very difficult to estimate the relevance of a stenosis based on the coronary artery anatomy.

We present an indirect approach to visualize the coronary anatomy through the combination of a computational simulation and a set of comprehensive visualization methods. From a whole heart scan we segment the myocardial surface of both the left and the right ventricle. We also segment the main coronary arteries including diameter information. These segmentations form the input for our computational simulation, where we compute the perfusion of the myocardium based on the patient-specific coronary anatomy. Finally we apply our comprehensive visualization methods on the results of this simulation. This provides detailed information on the effects of defects in the coronary anatomy on the distribution of blood throughout the myocardium.

Our main contribution is that we demonstrate that information from such a computational simulation can be effectively visualized. We show that it can provide a more comprehensive visualization of the information contained within a whole heart scan as compared to common methods that render the anatomical data directly. The visualization methods we provide are visually similar to those often used for the analysis of data from different scanning protocols. This simplifies the task of finding correlations between various sources of data.

The structure of the paper is as follows. In section 2, we discuss previous work used by or related to our approach. In section 3, we provide details on our computational simulation we use to compute the flow of blood throughout the myocardium. In section 4, we present our methods for visualizing the results from the computations. In section 5, we demonstrate the effectiveness of the approach by comparing a healthy case to one with a stenosis. In section 6, we discuss feedback from an expert clinician and the limitations of our approach. Finally, in section 7, we present the conclusions and future work.

-
- Maurice Termeer and M. Eduard Gröller are with the Vienna University of Technology, E-mail {maurice|groeller}@cg.tuwien.ac.at.
 - Javier Oliván Bescós, Marcel Breeuwer and Frans Gerritsen are with Philips Healthcare, E-mail: {javier.olivan.bescos|marcel.breeuwer|frans.gerritsen}@philips.com.
 - Anna Vilanova is with the Eindhoven University of Technology, E-mail: a.vilanova@tue.nl.
 - Eike Nagel is with the King's College London, E-mail: eike.nagel@kcl.ac.uk.

Manuscript received 31 March 2008; accepted 1 August 2008; posted online 19 October 2008; mailed on 13 October 2008.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

2 Related Work

The anatomy of the coronary arteries varies greatly among individuals. A set of recommendations published by the American Heart Association (AHA) on the segmentation of the myocardium included a standardized mapping of myocardial segments to supplying coronary arteries [4]. It was presented with the remark that there is indeed tremendous variability in the coronary artery blood supply. The variation among patients was later verified in a clinical study that evaluated the correspondence between the 17 segments of the recommended model and the coronary arteries [19]. This great variability calls for using a patient-specific approach towards assessing the coronary arteries and their relation to the myocardium. Beliveau *et al.* [1, 2] proposed a method of creating a patient-specific mapping between myocardial segments and coronary arteries using tomographic scans to determine the closest coronary arteries for each segment.

There has been extensive research in the area of simulating blood flow in the coronary arteries, or the arterial system in general. Several studies have been performed on the correlation between flow simulations and quantitative measurements of the flow in vessel trees. Good correlations between the two have been reported [3, 12, 7]. The majority of these studies focus on the numerical results and the correlation to flow measurements. They do not provide effective and intuitive visualization methods.

In the area of cardiac medicine, a commonly recurring visualization primitive applied to various types of data is the bull’s eye plot. Although several variations of the same idea are in use, their common purpose is to provide a two-dimensional overview of the left ventricle. The recommendations of the AHA also included a proposal how to arrange the 17 segments in a bull’s eye plot [4]. These types of bull’s eye plots are often applied when the resolution of the data is fairly low, for example in the visualizations of Hennemuth *et al.* [9] of perfusion data. Occasionally the right ventricle is included as well [5]. When more precision is available, a bull’s eye plot with a higher resolution can be constructed. This is demonstrated in the work of Noble *et al.* on combining late enhancement and cine data [17]. In their work the bull’s eye plot visualization is combined with a mapping of the results on the three-dimensional left ventricular surface. Termeer *et al.* [20] introduced the volumetric bull’s eye plot which uses a continuous transformation of the volume data to obtain a diagram similar to the bull’s eye plot, yet preserving continuity and wall thickness.

The most common approach to examine the anatomy of the coronary arteries is to directly inspect the three-dimensional tomographic data, although methods exist to simplify this process. A typical approach is demonstrated in the work of Kuehnel *et al.* [11], where a segmentation of the coronary artery centerlines is used to provide cross-sectional views of the volume data. Curved planar reformations [10] are popular for examining coronary arteries. Since these methods only visualize the coronary anatomy, they do not facilitate a correlation to data from other sources.

When additional data is available, these can be simultaneously visualized to allow a correlation between these data and the coronary anatomy. Oeltze *et al.* [18] visualize perfusion data and use an anatomical model of the coronary arteries to relate the two. Here the great variance of coronary anatomy could cause problems, as is noted by the authors. Nakaura *et al.* [16] apply fused volume rendering to visualize the coronary anatomy from a CT scan and relate it to SPECT data.

We present an approach to visualize the relative distribution of blood throughout the myocardium using a computational simulation of the perfusion of the myocardium. Our simulation uses patient-specific segmentations of the coronary and myocardial anatomy. Figure 1 gives a schematic overview of our approach. We apply an automatic whole heart segmentation algorithm [6] to obtain segmentations of the left and right ventricles, both represented as unstructured polygonal meshes. We use a semi-automatic vessel tracking algorithm [14] to obtain segmentations of the main coronary arteries, represented by their centerlines and a set of diameter measurements along those centerlines. This approach involves manually indicating seed points at the beginning and end of each coronary artery branch. All informa-

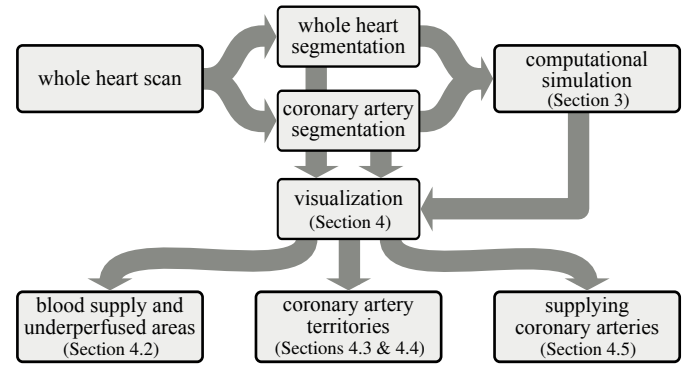


Fig. 1: Overview of our approach

tion is extracted from a single whole heart scan. These patient-specific segmentations are the input for a computational simulation, which is discussed in Section 3. The results of these computations are subsequently visualized, which is elaborated on in Section 4. The visualizations provide insight into the overall distribution of blood supply, the areas each coronary artery is supplying and a method to query the relative supply from each coronary artery to a region of interest.

3 Computation of Coronary Flow

Since our prime task is a visualization feasibility study, we have developed a simulation which is sufficient for this purpose, but may prove to be too simplistic for clinical applications. Our approach does not take the dynamic nature of cardiac flow into account. We also perform all our computations on the epicardium, the outer surface of the myocardium, rather than on a true three-dimensional representation such as a tetrahedral mesh. The additional complexity of a more advanced computational model would make it harder to integrate such an approach. We therefore chose to use a simpler method at first.

3.1 Perfusion through the Coronary Arteries

The amount of blood that can flow through a coronary artery is limited by the radius of that artery. The radius of each coronary artery varies along its path. Coronary arteries typically are thinner towards the end, near the apex, than they are at their origin near the base of the left ventricle. Also a stenosis can be a source of a varying diameter. The amount of blood flowing through a point in an artery also determines the amount that can flow out of the artery at that point and out of the remainder of the artery.




Fig. 2: The two building blocks used to construct a network of resistors representing the coronary artery tree. The block in (a) represents a linear segment of an artery while the block in (b) represents a bifurcation.

We have modeled the influence of the diameter of the vessel on the amount of blood that can flow through it by representing the coronary artery tree as a network of resistors. The entire coronary artery tree is constructed using two simple building blocks depicted in Figure 2. Using a discrete representation of the centerline of a coronary artery, the artery is modeled by a connected chain of building blocks as depicted in Figure 2a. The resistor labeled R_{through} controls the flow through the artery while the resistor labeled R_{out} controls the flow out of the artery at that point. We base the actual values of R_{through} on the diameter of the

artery at that point using Poiseuille’s law, which states that the resistance is inversely proportional to the radius of the artery to the fourth power [8].

For the value of R_{out} we use a predefined constant. This does not mean that an equal amount of blood flows out of an artery at each point. The blood volume present at a point is dependent on the radius at that point and the previous part of the artery. This constant is typically several orders of magnitude larger than the average value of R_{through} to prevent all blood from flowing out at the beginning of an artery. If it is set too high however, the effect of the variance of R_{through} due to a varying artery diameter is reduced.

Once all arteries are represented as chains, the final tree can be constructed by inserting building blocks as depicted in Figure 2b at the bifurcations. Note that no additional resistors are necessary. The split of the flow is automatically taken care of by the parallel circuit created and the resistors following a bifurcation, which represent the diameters of the two branches.

Once the tree is complete, we compute the relative amount of flow out of the artery at each point along the artery. By applying the laws for series (Equation 1) and parallel (Equation 2) circuits we can compute the total equivalent resistance of the remainder of the tree at any point. Using Ohm’s law we find that, as shown in Equation 3, the ratio between the flow through and the flow out of the artery is inversely proportional to the ratio between the reciprocals of R_{through} and R_{out} .

$$R_{\text{total}} = \sum_n R_n \quad (1)$$

$$\frac{1}{R_{\text{total}}} = \sum_n \frac{1}{R_n} \quad (2)$$

$$\frac{I_{\text{through}}}{I_{\text{out}}} = \frac{R_{\text{out}}}{R_{\text{through}}} \quad (3)$$

We treat the end of each artery as closed, meaning that all incoming flow has to leave the artery through the set of R_{out} resistors. This is a valid assumption, as all the blood has to leave the coronary arteries in order to be lead back into the coronary veins.

3.2 Modeling of a Stenosis

With the methods described above, we can easily model a stenosis by modifying the radius of an artery at a location of choice. In Figure 3 the dashed blue line shows the relative out-flow, the flow of blood out of the artery, at each point along an artery. Note that there is less flow out of the artery towards the end, as the artery narrows down along its path.




Fig. 3: Plot showing the relative flow out of a healthy artery and the same artery with a simulated stenosis.

Around the middle of the artery we have subsequently simulated a stenosis by narrowing the artery by 70%. The relative out-flow is shown by the solid green line. The relative out-flow after the stenosis severely decreases, as can be expected. Since we are computing relative flows, the out-flow before the stenosis increases. This does not necessarily mean that the absolute flow out of the artery has also increased in that segment, since the overall resistance of the artery

has also increased. A stenosis in one artery can thus lead to a global change in the distribution of blood throughout the entire coronary artery tree.

3.3 Diffusion of Flow

Once the flow out of the coronary artery tree is determined, it needs to be diffused throughout the myocardium. We approximate the myocardium by the surface of the epicardium, represented as an unstructured triangle mesh. This mesh is extracted from a whole heart scan using an automatic segmentation algorithm based on an approach by Ecabert *et al.* [6].

The diffusion process consists of three parts. First, blood from the coronary artery tree is diffused onto the myocardial surface. Second, blood on the myocardial surface itself is diffused. Finally, blood on the myocardial surface is absorbed. The last step simulates the flow of oxygen-depleted blood out of the myocardium through a network of veins.

Both the diffusion from the coronary arteries and the diffusion on the myocardial surface are implemented as iterative convolutions with a Gaussian kernel. All computations are performed on the unstructured triangle mesh of the epicardium. To obtain the desired accuracy during the diffusion process, we apply a subdivision scheme to the epicardial mesh to ensure the area of each triangle is below a predefined threshold. This subdivision scheme guarantees that the kernels used for the convolution process are approximately isotropic. We employ an approach loosely based on Loop’s subdivision scheme [13]. Instead of subdividing all triangles in each iteration, only those triangles with an area exceeding the threshold are subdivided. Neighboring triangles are split if necessary to avoid T-junctions. We have found that the smoothing step proposed by Loop may introduce undesired modifications of the mesh near sharp edges. Therefore we choose not to smooth our new vertices. Also the smoothness of the surface is of less importance for our purpose. The subdivision is applied iteratively until all triangles have an area below the specified threshold. The dynamic subdivision approach, compared to Loop’s original subdivision scheme, significantly reduces the number of triangles in the resulting mesh, which gives better performance during the diffusion process.

Given the subdivided triangle mesh, we compute a diffusion neighborhood for each node of the coronary artery tree. This neighborhood forms the kernel for the convolution process. It consists of a set of vertices of the epicardial mesh and a diffusion weight for each vertex. The diffusion weight is the Gaussian function applied to the Euclidean distance of the node to that particular vertex. This is shown in Equation 4, where $w(d)$ denotes the diffusion weight, d denotes the Euclidean distance between the node and the vertex, σ is a parameter of the diffusion process and ϵ denotes the minimum diffusion weight of vertices to include in the diffusion neighborhood.

$$w(d) = \begin{cases} e^{-\frac{d^2}{2\sigma}} & \text{if } d \leq -2\sigma \ln(\epsilon) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

A similar neighborhood is computed for each vertex of the epicardial mesh, except that here between vertices geodesic distances along the epicardial surface are used to compute the diffusion weights. The average size of the diffusion neighborhoods is significantly decreased by removing vertices that have a very small diffusion weight, controlled by ϵ in Equation 4. The weights of each neighborhood are normalized to ensure that the total amount of blood remains equal.

We also compute an absorption ratio for each vertex of the mesh. The demand for oxygenated blood at a certain point is related to the amount of myocardial tissue near that point. We incorporated this in our simulation by relating the absorption coefficient at a vertex to the area of the myocardium on a slice perpendicular to the long axis of the heart and intersecting the corresponding vertex. While some degree of absorption is necessary to prevent the continuous accumulation of blood, our experience is that the distribution of the absorption coefficients is not a significant factor in the outcome of the computational simulation.

After the above values are precomputed, we set a constant flow of blood into the root node of the coronary artery tree. We then iteratively execute the three phases of the diffusion process, i.e., the diffusion from the coronary artery tree to the myocardial surface, the diffusion within the myocardial surface and the absorption. The absorption phase consists of decreasing the amount of supply at each vertex of the myocardial surface by the absorption coefficient of that vertex. After a sufficient number of iterations, this process will converge to a stable state. We then stop the iterative process. The amount of blood at each point in the myocardium in this stable state represents the amount of supply to that point. To retrieve additional information on which coronary artery is supplying which area, we perform the entire diffusion process separately for each of the three primary coronary arteries. This information can easily be combined to obtain global supply information. In the following part we demonstrate how these quantities can be effectively visualized.

4 Visualization of Coronary Flow

The previously discussed model is simple, but sufficiently effective for our purposes. In this section we demonstrate that the results of our computational simulation can be effectively visualized. The term effective here refers to being able to assess the effect of the coronary anatomy on the perfusion of the heart. This should be possible without detailed technical knowledge or understanding of the computational simulation. The latter requirement is essential for the viability of our techniques in clinical practice.

We present visualization methods that directly use the three-dimensional segmentation data to portray the results of the computational simulation as well as two-dimensional projections of these segmentations. We first discuss how these projections are constructed. Second we demonstrate the direct visualization of supply and underperfused regions using color, isocontours and textures. Third we demonstrate how the area each coronary artery is supplying can be visualized by using borders of the coronary artery territories. We also visualize the supply area of each coronary artery separately. Finally we show the inverse of the latter, i.e., finding the supplying coronary arteries of a region of interest.

4.1 Bull's Eye Plot Representation

The bull's eye plot is a visualization primitive that is commonly used in cardiac medicine. Its main goal is to provide a two-dimensional overview of the left ventricle and optionally the right ventricle. The classical approach to generate a bull's eye plot is to segment the myocardium in a stack of short-axis slices and map these segments to a set of concentric rings.

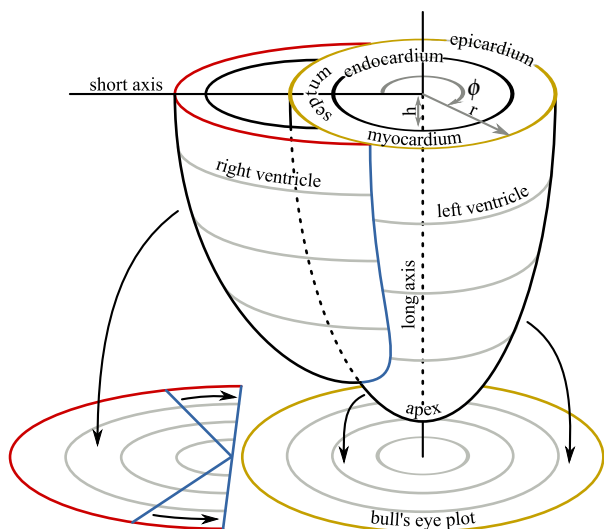


Fig. 4: The parameterization of the left and right ventricles and the projection onto a bull's eye plot.

We however construct a bull's eye plot as a continuous unfolding of the left ventricle along the long axis. The segmentation of the heart we use is an unstructured grid of vertices along the epicardium. Each of these vertices can be represented by three parameters: the angle ϕ with the short axis on a plane perpendicular to the long axis, the distance h to the apex along the long axis and the distance r to the long axis. This is illustrated in Figure 4. Any point can then be projected onto a cylinder by interpreting this triplet as cylindrical coordinates. Since the surface will be projected onto a plane, we do not need to maintain the information provided by r . We therefore only use ϕ and h as polar coordinates to form a circle.

Our segmentation covers the entire left and right ventricles, including the top areas, near the bigger arteries and valves. These areas do not correspond to contracting myocardium and are not commonly part of a bull's eye plot. The little variation in h in these areas would also cause a cluttered projection. We therefore do not include these top areas in the bull's eye plot. Using our parameterization, we only include areas that have a value of h that is less than the maximum value of h in the septum, the myocardial wall shared by the left and right ventricle.

Since the right ventricle is also part of our simulation, we also want to depict it in the bull's eye plot. While several approaches to visualize the right ventricle in a two-dimensional manner exist, the most common approach is to represent it by a half circle. This approach is based on the idea that the right ventricle is essentially half of a left ventricle in shape. This approach is common in clinical practice [5].

To realize a half-circle unfolding we first apply the same parameterization to the right ventricle as we do to the left ventricle. In the bull's eye plot we translate the right ventricle along the short axis to prevent overlap with the projection of the left ventricle. Next we normalize ϕ to the range $[-\frac{1}{2}\pi; \frac{1}{2}\pi]$ to form a half circle. This step also eliminates any inter-patient shape variations of the right ventricle.

The orientation of the bull's eye plot that is used in clinical practice would correspond to a viewpoint below the bull's eye plot looking upwards. We use this orientation whenever we show a separate bull's eye plot. This orientation appears to be flipped with respect to the projection below a three-dimensional view of the heart.

4.2 Blood Supply Area and Underperfused Regions

The primary outcome of the computational simulation is the blood supply present at each point in the myocardium. We directly visualize this using color encoding, as can be seen in Figure 5. To avoid visual clutter due to too many colors, we use a light gray to dark red discrete color gradient. For additional insight in the structure of the supplying regions, we add isocontours as white dashed lines. These contours delineate borders where the supply of blood is equal. Finally we apply a striped pattern to the underperfused regions, i.e., regions where the supply is below a user-specified threshold.

Due to the simplicity of our computational simulation, the absolute values we obtain may not correspond close to reality. Currently we therefore cannot define intervals that correspond to a healthy supply. We solve this by showing relative supply instead, which is based on the maximum supply in the simulation under analysis. In order to perceive both large and small differences in the distribution of blood throughout the myocardium, we apply a logarithmic scaling prior to applying color encoding and computing the isocontours. This eases the interpretation of the shape of regions with near-equal supply. The striped pattern provides a way of marking underperfused areas, i.e., areas where the supply of blood is below a predefined threshold. The borders of the pattern are made fuzzy, to indicate an area of uncertainty.

The striped pattern is implemented using h , the distance to the apex along the long axis. All parts where the fractional part of $\delta \cdot h$ is larger than $\frac{1}{2}$ are made opaque, the others transparent. The factor δ controls the thickness of the stripes. This gives stripes perpendicular to the long axis in the three-dimensional view, and circles or arcs in the two-dimensional projection. An advantage of this pattern is that it is rotation invariant. The orientation of the pattern is perceived equally throughout the entire mesh and the two-dimensional projection.




Fig. 5: Visualization of coronary artery blood supply showing the supply of blood using color encoding, isocontours delineating regions of equal supply as white dashed lines, lines indicating equal supply from multiple coronary arteries as black dotted lines and underperfused regions using a striped pattern. Below the mesh, a two-dimensional projection shows the same data.

We display the coronary arteries in the three-dimensional view as well as on the bull's eye plot. In the three-dimensional view the coronary arteries are displayed as tubes with radii corresponding to the actual artery radii. In order to prevent intersections of the coronary arteries with the myocardium, we modify parts of the coronary arteries. These modifications are done only for visualization purposes. The computational simulation uses the original, unmodified coronary arteries. The coronary arteries are first projected onto the myocardium and subsequently translated a constant distance along the surface normal. This causes each coronary artery to have a constant distance to the myocardium. We apply black (high) to blue (low) color encoding to show the relative outflow at each point in the coronary artery tree.

The coronary arteries are also projected onto the bull's eye plot using the same projection technique as described in Section 4.1. Whether a point along an artery should be mapped on the projection of the left or right ventricle is based on which ventricle it is closest to. We apply the same color encoding showing relative outflow on the projected coronary arteries.

4.3 Coronary Artery Territories

To gain an insight into which region each coronary artery is supplying, we visualize the coronary artery territories. We divide the coronary arteries into three groups: left anterior descending (LAD), left circumflex (LCX) and right coronary artery (RCA). This division is common in clinical practice and is also recommended by the AHA [4]. In the direct visualization of blood supply, we visualize the coronary territories by drawing black dotted lines along which the supply from two or more coronary arteries is equal according to our simulation. This is demonstrated in Figure 5.

While visualizing the equi-perfusion lines already gives a good basic insight into the area each coronary artery group is supplying, these areas are not discrete. Although there is generally little overlap in the coronary territories in practice, our simulation suggests that areas near the borders are supplied by multiple coronary artery groups. This is due to our diffusion approach. The exact details of how the coronary arteries supply the myocardium are not yet fully understood. We can thus also interpret the overlap generated by our simulation as an uncer-




Fig. 6: Individual coronary territories visualized using separate colors. A striped pattern with fuzzy borders is used in regions of uncertainty where the exact territory is unknown. Underperfused regions are not colored.

tainty of coronary territories. As is demonstrated in Figure 6, we visualize this by showing the full territory each coronary artery group is supplying in separate colors. Where these regions overlap, we apply a striped two-color pattern to indicate an uncertainty on which coronary artery groups is supplying this region. This striped pattern is implemented using the same approach as for the striped pattern to visualize underperfused regions. We again make the border of the pattern fuzzy, to further indicate an area of uncertainty.

The boundaries of each coronary territory are controlled by a user-specified threshold. If the supply from a coronary artery group in a certain point is above this threshold, that point is considered to be part of the respective coronary territory. Figure 6 shows that using this approach, the regions each coronary artery group is supplying as well as areas of uncertainty can be clearly identified.

4.4 Separate Coronary Artery Groups

The computational simulation provides separate information on the supply for each coronary artery group. This allows for separate visualization of the blood supply by each group. It gives a focused view on the area a particular group is supplying.

Figure 7 shows an example of the separate visualization of coronary artery groups by showing four bull's eye plots. Figure 7a shows the supply of all three coronary artery groups combined, while Figures 7b, 7c and 7d show the supply of the LAD, LCX and RCA, respectively. Due to the separation, more detailed visualization techniques can be applied than when visualizing the coronary territories in a combined fashion, as was proposed in Section 4.3. A combined visualization of the coronary territories allows for assessing the relation between them, while separate visualization allows for a more comprehensive analysis of a particular group.

Showing supply from coronary arteries separately benefits from the two-dimensional projection, as this allows for a quick comparison of the regions each coronary artery group is supplying. The ease of comparison argument holds for the bull's eye plot in general. Comparisons cannot be as easily performed using different viewpoints of the three-dimensional view. In the bull's eye plot the relation to the three-dimensional anatomy is however partially lost.

4.5 Querying Supplying Coronary Arteries

The previously discussed techniques visualize the region each coronary artery group is supplying. We can also do the inverse, i.e., to visualize the coronary artery group that a region of interest is supplied by. This approach is demonstrated in Figure 8. The user indicates a region of interest by specifying a point on either the three-dimensional mesh or the bull's eye plot and a radius to determine the size of the region. We then compute the relative supply from each of the three coronary artery groups. Again taking into consideration that there is minimal overlap between coronary territories in practice, we can in-




Fig. 7: Separate visualization of the supply by each coronary territory group in two-dimensional bull's eye plots. Figure (a) shows the supply of the three groups combined, (b) the supply of the LAD, (c) the supply of the LCX and (d) the supply of the RCA. Only the relevant ventricles are depicted.

interpret this as the *likelihood* that the region is supplied by one or more arteries. We then draw arrows from the closest point of each coronary artery group to the specified point of interest. We project the start and end point of each arrow on a cylinder with the long axis as cylinder axis. We then project the shortest path on the cylinder between these two points back onto the myocardium. On the bull's eye plot we draw a straight path between the two projected points. The relative supply by each group is expressed in the width of each arrow.

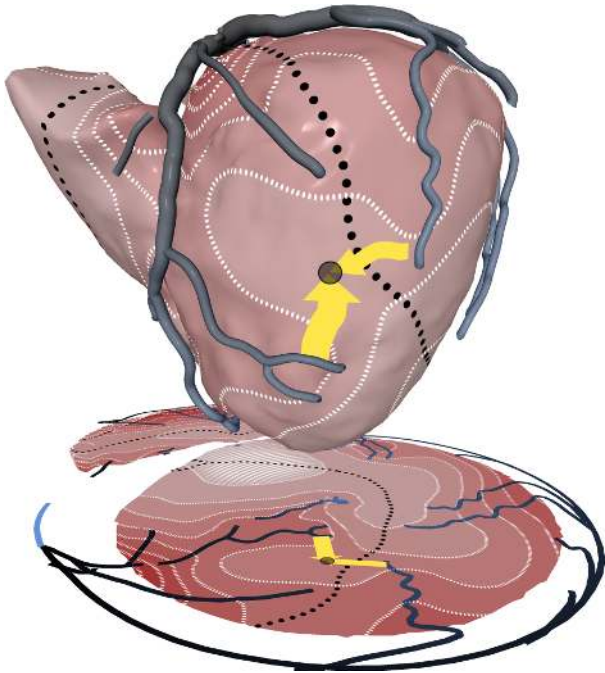


Fig. 8: For a selected region arrows indicate the relative supply by, or likelihood of being supplied by, each of the three coronary artery groups. The relative supply is encoded in the width of the arrows.

This approach allows for a quick identification of the supplying coronary arteries in a visualization where this may not be directly apparent. An example is to query the supplying coronary arteries of an underperfused region in visualizations as discussed in section 4.2.

5 Simulation and Visualization of a Stenosis

In order to evaluate the effectiveness of the visualization techniques previously described, we perform an experiment. Starting out with a

scan of a healthy subject, we visualize the coronary anatomy using our computational simulation and the proposed visualization techniques. Next we artificially induce a stenosis in one or more of the coronary arteries and observe the change in the visualization of the coronary perfusion.

Our approach requires a whole heart scan of sufficient accuracy to allow a detailed segmentation of the coronary arteries with diameter measurements. CT currently provides better resolution than MRI, so we chose to use CT scans for our experiment. When MRI has advanced sufficiently to allow accurate segmentation of the coronary arteries, using MRI may be preferable over CT. MRI allows for the acquisition of additional data that show the functioning of the heart, which can lead to a more comprehensive diagnosis.

5.1 Supply and Underperfused Regions

A stenosis often causes a perfusion defect in the area normally supplied by the coronary artery that contains the stenosis. When simulating a stenosis, we would expect to observe this perfusion defect in the resulting visualizations. Figure 9 shows a comparison of a healthy case (Figure 9a) and the same case with an artificially induced stenosis blocking part of the upper LAD by approximately 70% (Figure 9b). This means that the cross-sectional area of the vessel has decreased by 70% in that area. Conforming to expectations, Figure 9b shows an increased underperfused region near the lower segment of the LAD. The shape of the isocontours also expresses the change in supply in the affected area. Note that Figure 9a shows an underperfused region in the basal septal wall, which is commonly not regarded as an indicator of CAD.

5.2 Coronary Artery Territories

A stenosis is also expected to change the shape of the territory of the respective coronary artery. Figure 10 compares the coronary artery territories of a normal case and the same case with an artificially induced stenosis in the RCA. In Figure 10b the severe reduction in size of the RCA territory can clearly be observed, especially the posterior and septal parts of the left ventricle are affected. Note that the isocontours, representing the overall supply, have also somewhat changed in the other territories.

5.3 Querying Supplying Coronary Arteries

Finally, the effect of a stenosis can also be observed in the visualization of the supplying coronary arteries of a region of interest. Figure 11 compares a region of interest in a healthy case and a case with an artificially induced stenosis in the LAD. The region is primarily supplied by the LAD in the healthy case. It is primarily supplied by the LCX and lies at the border of an underperfused region in the stenosed case. Note that because the arrows encode *relative* supply, the thicker arrow from the LCX does not necessarily indicate an increase in supply.




Fig. 9: Comparison of supply between (a) a healthy case and (b) the same case with an artificially induced stenosis in the upper segments of the LAD indicated by the arrow.

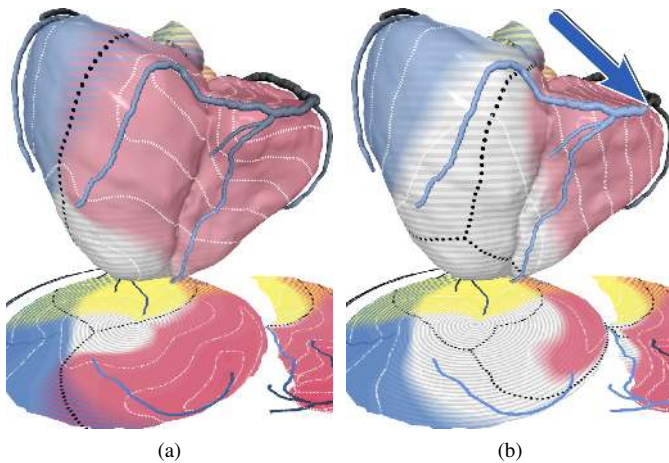


Fig. 10: Comparison of the coronary territories between (a) a healthy case and (b) the same case with an artificially induced stenosis in the RCA indicated by the arrow.

6 Discussion

In current clinical practice the 17-segment model from the AHA is used as a reference for the relation between the coronary arteries and the myocardium. This model is not based on patient-specific information. We have extended this model by including such information. Applying a computational simulation, we obtain a more detailed relation between the coronary arteries and the myocardium. While our simulation is primitive, it gives more detailed information than is currently available in medical practice. This approach can be seen as a top-down approach. There are more advanced approaches to compute myocardial perfusion based on computational fluid dynamics which can be seen as a bottom-up approach.

In clinical practice a more advanced perfusion simulation is probably preferable. Our simulation only uses a two-dimensional surface of the myocardium, does not take the dynamic nature of the heart cycle into account and the perfusion method may not correspond close to reality. This may hamper correlating results to for example perfusion imaging, although the basic functionality is captured by our approach. We used CT data since our approach requires a detailed coronary artery tree segmentation. When MRI advances enough to allow for more detailed imaging of the coronary arteries, additional scanning protocols can be used to provide a more comprehensive visualization

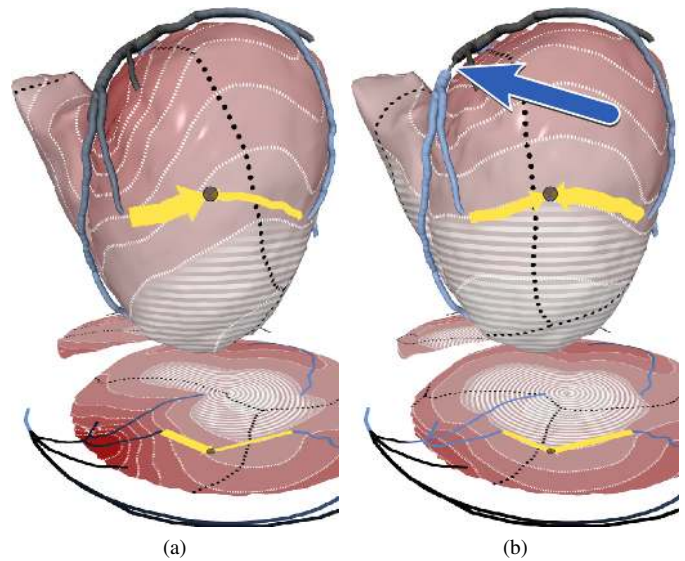


Fig. 11: Comparison of supplying coronary arteries between (a) a healthy case (b) and the same case with an artificially induced stenosis in the LAD indicated by the arrow. The relative supply is encoded in the thickness of the arrows.

of the patient.

6.1 Feedback from Expert Clinicians

In a discussion with two expert radiologists we informally evaluated the clinical relevance of the various aspects of our visualizations. We received very positive feedback on the visualization of supply and underperfused regions. It displays the relevant data in a clear way. The black dotted lines give a good indication of the coronary territories, but could perhaps be annotated with uncertainty information and colored in case the relations to arteries are not obvious.

The detailed visualization of the coronary territories may be good for assessing the area at risk in more detail, but gives a less comprehensive view than the previous visualization. One expert believes that for many cases this level of detail is not necessary. It may thus be worthwhile to integrate more of the coronary territory information in our visualization of supply such that a separate visualization thereof may no longer be necessary. The separate visualization of a single coronary territory provides a more focussed view, but similarly it may not be necessary for all patients. In Figure 7d it was not immediately obvious that a bull's eye plot of the right ventricle was shown instead of part of a left ventricular plot.

One expert was critical towards the approach of querying the supplying coronary arteries. He questions whether the cost of having to interact with a three-dimensional visualization would outweigh the limited additional information it provides. In summary, although both experts commented that the simulation may not be sufficiently realistic, we received positive feedback on the presented visualizations. One expert noted that they may also have potential for communication between doctor and patient.

6.2 Implementation Details

Both the left and right coronary arteries branch off from the aorta. The aorta is not part of our segmentation of the coronary artery tree, which is therefore not connected. In order to construct a tree with a single root node, we connect the two subtrees of the left and right coronary arteries with a building block from Figure 2b. No additional resistors are required in this area, as no significant flow out of the arteries takes place in the corresponding segments. Since typically also no flow out of the left main artery occurs, we also model this artery only using building blocks from Figure 2b at the bifurcations.

In the computational simulation only the two types of resistors R_{through} and R_{out} determine the distribution of blood throughout the coronary artery tree. We base the value of R_{through} on the local artery radius, but we use a manually specified constant value of R_{out} instead of basing it on physical properties. This causes the supply in a region to be partially related to the amount of close-by vessel surface area.

We have noticed perfusion defects near the apex in some cases. This is probably caused by an incomplete segmentation of the coronary arteries, mainly near the apex. The diameter of the coronary arteries decreases towards the apex, making segmentation in that area more difficult. Our segmentations thus may contain shorter coronary arteries near the apex as compared to reality, which manifests itself as a perfusion defect in the simulations. Perfusion defects also tend to show up between two vessels, where the true perfusion situation, both in reality and in our simulation, is not fully understood.

In our experiments we subdivide the triangle mesh of the myocardial surface such that all triangles have an area of at most 1mm^2 . The resulting surface consists of about 35,000 triangles. The diffusion neighborhoods for the coronary artery tree and the myocardial surface are computed using a σ of 5 and 7, respectively. We use a kd-tree to compute the neighborhoods efficiently. For both diffusion processes an ε of 10^{-2} is used. Elements with a diffusion weight below this threshold are not included in the diffusion neighborhood. We use a value of R_{out} of 10^3 . The radius-based values of R_{through} typically vary between 0.1 and 4.5. With these parameters the diffusion process stabilizes after approximately 500 iterations.

We implemented all visualizations using OpenGL 2.0. All of the images were generated on a desktop PC (Intel Xeon 3Ghz, 2GB RAM, NVIDIA GeForce 8800GTX). The precomputation phase of the computational simulation takes approximately two seconds. Performing sufficient iterations until a stable state is obtained typically takes 15 to 20 seconds. All visualizations can be computed at interactive rates.

7 Conclusion

We have presented a computational simulation for the perfusion of blood throughout the myocardium and a set of visualization methods for the visual assessment of the results of these computations. Our computational simulation computes the flow of blood through the coronary arteries using a network of resistors, while the successive distribution of blood is implemented using iterative convolution with a Gaussian kernel. This computational simulation advances on current clinical practice by including patient-specific information. It establishes a more advanced mapping between the myocardium and the coronary arteries based on perfusion. In contrast current approaches assume that each area is supplied by its closest coronary artery. Our simulation produces sufficiently accurate results to demonstrate the effectiveness of our visualizations.

We have demonstrated comprehensive visualizations of the computed myocardial perfusion. We presented a direct visualization of blood supply. Furthermore we visualize the coronary artery territories and the supplying coronary arteries of a region of interest. The computational simulation gives separate information on the supply from each coronary artery group. This demonstrates the advantages of a simulation approach, as this information is not commonly available from scanning technology. We have given an example how our techniques visualize a stenosis by artificially inducing one.

Our techniques extend visualization primitives currently common in clinical practice, including the bull's eye plot. This approach simplifies the process of visually correlating the anatomical data to information from other imaging modalities and should ease clinical adoption. We have discussed our work with radiologists and overall received positive feedback. The combination of two- and three-dimensional approaches provides a good blend between a clear overview and a detailed relation to the three-dimensional anatomy.

Acknowledgements

This work was performed in the scope of the COMRADE project funded by Philips Healthcare, Best, The Netherlands.

References

- [1] P. Beliveau, R. Setser, F. Cheriet, and T. O'Donnell. Patient-specific coronary territory maps. *Proc. SPIE*, 6511:65111J, 2007.
- [2] P. Beliveau, R. Setser, F. Cheriet, R. White, and T. O'Donnell. Computation of coronary perfusion territories from CT angiography. *Computers in Cardiology*, 34:753–756, 2007.
- [3] D. Bessems. *On the propagation of pressure and flow waves through the patient-specific arterial system*. PhD thesis, Eindhoven University of Technology, 2007.
- [4] M. D. Cerqueira, N. J. Weissman, V. Dilsizian, A. K. Jacobs, S. Kaul, W. K. Laskey, D. J. Pennell, J. A. Rumberger, T. Ryan, and M. S. Verani. Standardized myocardial segmentation and nomenclature for tomographic imaging of the heart. *Circulation*, 105:539–542, 2002.
- [5] E. G. DePuey, M. E. Jones, and E. V. Garcia. Evaluation of right ventricular regional perfusion with technetium-99m-sestamibi SPECT. *Journal of Nuclear Medicine*, 32:1199–1205, 1991.
- [6] O. Ecabert, J. Peters, and J. Weese. Modeling shape variability for full heart segmentation in cardiac computed-tomography images. In *Proc. SPIE*, volume 6144, pages 1199–1210, 2006.
- [7] T. Frauenfelder, E. Boutsianis, T. Schertler, L. Husmann, S. Leschka, D. Poulikakos, B. Marincek, and H. Alkadhi. In-vivo flow simulation in coronary arteries based on computed tomography datasets: feasibility and initial results. *European Radiology*, 17:1291–1300, 2007.
- [8] Y. Fung. *Biomechanics: Circulation*. Springer, 1996.
- [9] A. Hennemuth, S. Behrens, C. Kuehnel, S. Oeltze, O. Konrad, and H. O. Peitgen. Novel methods for parameter based analysis of myocardial tissue in MR-images. In *Proc. SPIE*, volume 6511, page 65111N, 2007.
- [10] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and E. Gröller. CPR - Curved Planar Reformation. In *IEEE Visualization*, pages 37–44, 2002.
- [11] C. Kuehnel, A. Hennemuth, S. Bock, S. Oeltze, T. Boskamp, S. Krass, B. Preim, and H. O. Peitgen. New software assistants for cardiovascular diagnosis. In *GI-Workshop "Softwareassistenten - Computerunterstützung für die medizinische Diagnose und Therapieplanung" within the GI annual conference*, pages 491–498, 2006.
- [12] C. D. Lazzari, M. Trivella, M. Micalizzi, F. Bernini, M. Varanini, G. Ferrari, A. Macerata, D. Neglia, and A. L'Abbate. Coronary blood flow: Comparison between in vivo and numerical simulation data. *Computers in Cardiology*, 33:881–884, 2006.
- [13] C. Loop. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, 1987.
- [14] C. Lorenz, S. Renisch, T. Schlathoelter, and T. Buelow. Simultaneous segmentation and tree reconstruction of the coronary arteries in MSCT images. In *Proc. SPIE*, volume 5031, pages 167–177, 2003.
- [15] E. Nagel, C. Klein, I. Paetsch, S. Hettwer, B. Schnackenburg, K. Wegscheider, and E. Fleck. Magnetic resonance perfusion measurements for the noninvasive detection of coronary artery disease. *Circulation*, 108:432, 2003.
- [16] T. Nakaura, D. Utsunomiya, S. Shiraishi, S. Tomiguchi, T. Honda, H. Ogawa, K. Awai, and Y. Yamashita. Three-dimensional cardiac image fusion using new CT angiography and SPECT methods. *American Journal of Roentgenology*, 185:1554–1557, 2005.
- [17] N. M. Noble, D. L. Hill, M. Breeuwer, and R. Razavi. The automatic identification of hibernating myocardium. *MICCAI*, 0:890–898, 2004.
- [18] S. Oeltze, A. Kuß, F. Grothues, A. Hennemuth, and B. Preim. Integrated visualization of morphologic and perfusion data for the analysis of coronary artery disease. In *Proc. EuroVis*, pages 131–138, 2006.
- [19] O. Pereztol-Valdes, J. Candell-Riera, C. Santana-Boado, J. Angel, S. Aguade-Bruix, J. Castell-Conesa, E. V. Garcia, and J. Soler-Soler. Correspondence between left ventricular 17 myocardial segments and coronary arteries. *European Heart Journal*, 26:2637–2643, 2005.
- [20] M. Termeer, J. Oliván Bescós, M. Breeuwer, A. Vilanova, F. Gerritsen, and M. E. Gröller. CoViCAD: Comprehensive visualization of coronary artery disease. *TVCG*, 13(6):1632–1639, 2007.
- [21] H. Thiele, S. Plein, M. Breeuwer, J. P. Ridgway, D. Higgins, P. J. Thorley, G. Schuler, and M. U. Sivananthan. Color-encoded semiautomatic analysis of multi-slice first-pass magnetic resonance perfusion: Comparison to tetrofosmin single photon emission computed tomography perfusion and X-ray angiography. *International Journal of Cardiovascular Imaging*, 20:371–84, 2004.