

Two methods for modified Doo–Sabin modeling of nonsmooth surfaces—applied to right ventricle modeling

Håkon Strand Bølviken^{Ⓞ, a,*} Jørn Bersvendsen,^b Fredrik Orderud,^b
Sten Roar Snare,^b Pål Brekke^{Ⓞ, c} and Eigil Samset^{a,b}

^aUniversity of Oslo, Gaustadalléen, Oslo, Norway

^bGE Healthcare Cardiovascular Ultrasound, Gaustadalléen, Oslo, Norway

^cOslo University Hospital, Department of Cardiology, Rikshospitalet, Sognsvannsveien, Oslo, Norway

Abstract

Purpose: In recent years, there has been increased clinical interest in the right ventricle (RV) of the heart. RV dysfunction is an important prognostic marker for several cardiac diseases. Accurate modeling of the RV shape is important for estimating the performance. We have created computationally effective models that allow for accurate estimation of the RV shape.

Approach: Previous approaches to cardiac shape modeling, including modeling the RV geometry, has used Doo–Sabin surfaces. Doo–Sabin surfaces allow effective computation and adapt to smooth, organic surfaces. However, they struggle with modeling sharp corners or ridges without many control nodes. We modified the Doo–Sabin surface to allow for sharpness using weighting of vertices and edges instead. This was done in two different ways. For validation, we compared the standard Doo–Sabin versus the sharp Doo–Sabin models in modeling the RV shape of 16 cardiac ultrasound images, against a ground truth manually drawn by a cardiologist. A Kalman filter fitted the models to the ultrasound images, and the difference between the volume of the model and the ground truth was measured.

Results: The two modified Doo–Sabin models both outperformed the standard Doo–Sabin model in modeling the RV. On average, the regular Doo–Sabin had an 8-ml error in volume, whereas the sharp models had 7- and 6-ml error, respectively.

Conclusions: Compared with the standard Doo–Sabin, the modified Doo–Sabin models can adapt to a larger variety of surfaces while still being compact models. They were more accurate on modeling the RV shape and could have uses elsewhere.

© 2020 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.7.6.067001](https://doi.org/10.1117/1.JMI.7.6.067001)]

Keywords: medical segmentation; right ventricle modeling; Doo–Sabin; computational topology; shape modeling applications; medical imaging.

Paper 20125R received May 14, 2020; accepted for publication Dec. 1, 2020; published online Dec. 23, 2020.

1 Introduction

The clinical importance of the right ventricle (RV) of the heart has become increasingly understood in recent years.¹ While most attention has traditionally been given to the left ventricle (LV),² RV performance has been shown to have prognostic and therapeutic consequences in a variety of heart diseases, from arrhythmogenic cardiomyopathy to pulmonary hypertension and left ventricular failure.^{3,4} With this increasing clinical interest, there has also been an increased interest in accurate modeling of the RV shape. Accurate assessment of the shape is needed for accurate estimation of the RV volume and RV ejection fraction, the latter being a predictor of moderate heart failure.^{5,6} RV ejection fraction is a predictor of survival of the

*Address all correspondence to Håkon Strand Bølviken, haaksb@ifi.uio.no

patients with idiopathic dilated cardiomyopathy⁷ and in unoperated patients with chronic severe mitral regurgitation. Subnormal RV ejection fraction at rest is correlated with decreased exercise tolerance, complex arrhythmias, and mortality.⁴

The RV is more challenging to model than the LV.¹ Viewed from a right anterolateral aspect, the RV has a triangular shape, with an inlet leading blood from the tricuspid valve and an outlet leading to the pulmonary valve. The RV is wrapped around the anterolateral section of the oblong spheroid-shaped LV, which makes the RV volume concave along the RV/LV interface and convex along the free wall. Due to the lower intracavitary pressures in the RV, the muscular walls are thinner than in the LV. In a two-dimensional (2D) short axis view, the RV is crescent-shaped with sharp corners. An example of this with approximate walls of the RV drawn in are shown in Fig. 1. In a long axis view, the RV can appear wedge-shaped.

Echocardiography is the first line modality for imaging the heart due to its low cost, real-time visualization, ease of use, and its safety compared to CT or x-ray due to the radiation-factor.⁸ Because of its complex shape, the RV is challenging to visualize with 2D imaging modalities. However, three-dimensional (3D) echocardiography shows promise in improving RV imaging.

Two challenges of 3D ultrasound imaging of the RV are image quality and limited field of view.⁹ Given the difficulties, several previous attempts on RV shape models rely on foundational models that include relatively strong presuppositions regarding the final model’s geometry. One such model was made by Bersvendsen et al.¹⁰ who used a Doo–Sabin model to provide a compact geometrical representation of the RV shape.

The Doo–Sabin model gives round, organic-looking surfaces, which appear well suited for medical shape modeling, as demonstrated by Orderud and Rabben¹¹ as well as Dikici.¹² In addition, it is easy to compute and its partial derivatives are well understood and easily calculated. However, the Doo–Sabin model is not ideal for modeling sharp corners without having a large number of control vertices around the sharpened area, increasing model complexity and computational cost. The RV has this type of sharp geometry, particularly in the apical region³ and where the free wall meets the interventricular septum.¹ This makes it a challenge for the traditional Doo–Sabin model.

In this work, we propose a modification that allows more flexibility in terms of modeling smooth and sharp surfaces while keeping the key properties of the Doo–Sabin model. This paper is an extension of the conference paper “Modified Doo-Sabin Modeling of the Right Heart Ventricle.”¹³ It extends the previous implementation by introducing a second approach to the construction of sharp surfaces, which gives different end results to what the surface looks like.

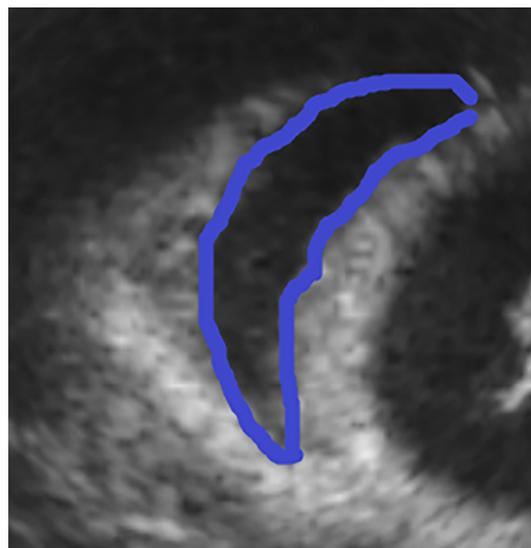


Fig. 1 A 2D slice of an ultrasound image of the RV, with approximate walls of the ventricle marked in blue.

This paper will also compare the standard Doo–Sabin model to the two modified, sharp Doo–Sabin models on modeling the RV shape based on 3D ultrasound images. As the model is meant to help with blood pool segmentation, the models should follow the endocardium. The Doo–Sabin models on the RV will be based on results from Bersvendsen et al.¹⁰ They used statistical shape modeling (SSM)¹⁴ to find an average RV shape by tuning the model on 14 MRI images. This model is already used in commercial applications, the EchoPAC 3D AutoRVQ (GE Vingmed Ultrasound, Horten, Norway), so an improvement could also be of commercial interest.

The inputs of the fitting process are the ultrasound image, six points manually placed by a trained cardiologist, and the relevant Doo–Sabin model. The first two inputs are kept the same between models, and the models are not changed depending on the image. A Kalman filter¹⁵ will be used to fit the models to ultrasound images of the RVs of patients. Kalman filters have been used before for segmentation of cardiac ultrasound images, for instance, by Orderud and Rabben,¹¹ Smistad and Lindseth,¹⁶ and Bersvendsen et al.¹⁰

2 Methods

2.1 Original Doo–Sabin Model

The Doo–Sabin model, originally described by Doo and Sabin,¹⁷ is a generalization of the quadratic B-splines. It has similarities to the Catmull–Clark model¹⁸ but has some advantages in terms of computation.

Let V_0 be a grid of vertices forming faces (simple 2D polygons in 3D space). A series of new, finer grids are created by constructing new vertices V_n and faces based on the previous V_{n-1} . There are three steps to this process:

- any face will be replaced by a smaller face with the same number of edges. This new face is called an F-face.
- any edge will be replaced by a four-sided face, in a sense fattening the edge into a face. This new face is called an E-face.
- any vertex will be replaced by a face with the same number of edges as the valence of the original vertex. For instance, in a cube, each vertex is connected to three others. Thus, each vertex is replaced by a triangle-shaped face. This new face is called a V-face. An example of the first two iterations of a Doo–Sabin process is shown in Fig. 2.

Each iteration of this process creates a finer grid of vertices, arriving at a limit surface when iterated an infinite number of times. A useful property of the Doo–Sabin model is that the only information needed to determine the surface is the coordinates of the vertices and the topology between them. The exact placement of each new vertex is then determined by a subdivision matrix, turning the placement of new vertices into a linear process. The subdivision matrix S is determined locally depending on topology and is multiplied by the array of local vertices V_{n-1} and the output gives the local vertices V_n of the new grid.

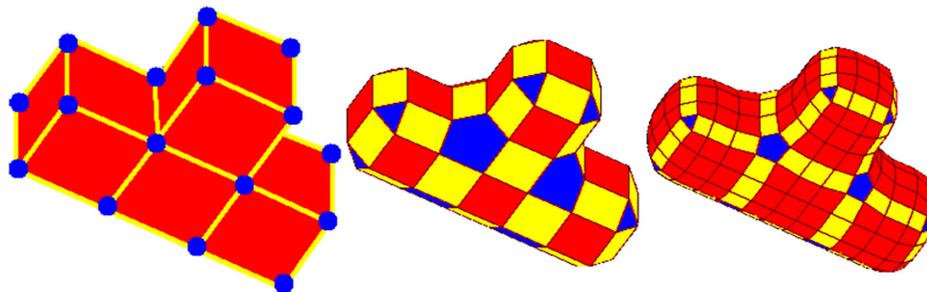


Fig. 2 The first two iterations of the Doo–Sabin algorithm. Note that the blue vertices in the first figure are replaced by polygons with as many corners as the vertex was connected to, and that the yellow edges are replaced by yellow four-sided faces. The red faces are replaced by red faces of the same shape, but smaller. Reproduced with permission, courtesy of Fredrik Orderud.

However, in practical applications, the limit surface can be approximated by doing one step of the process and then using an analytic method rather than going through repeated steps. This process, as well as more information on the subdivision matrix, has been detailed by Orderud and Rabben.¹¹ Briefly, for any point on the surface, the surface is subdivided until the point can be evaluated using the basis functions for B-splines in 2D. The process of subdivision is accomplished by repeated multiplication of the subdivision matrix.

In this work, the Doo–Sabin method was modified in two different ways to allow weighting of vertices and edges, where a higher weight will create a sharper vertex or edge. This was in both cases done by modifying the first grid created from the original starting grid, moving vertices close to the original vertice or edge.

While all Doo–Sabin models used in this paper are quadratic in degree, the mathematical methods for constructing sharp Doo–Sabin models can be applied to Doo–Sabin models of all degrees.

2.2 Sharp Doo–Sabin Models

We will present two modified Doo–Sabin models and will refer to the first as the connected sharp Doo–Sabin and the latter as the simple sharp Doo–Sabin. The modifications were in both cases done by changing the first subdivision matrix S , which means that unless the topology of the vertices resulting from the subdivision matrix is changed, the Doo–Sabin model retains its locality and derivatives properties. There are two different types of sharpness considered here:

- A sharp vertex, meaning the surface around that vertex can be sharpened, the sharpness centering in one point like a cone. Figures 3(a) and 3(b) show the effect of a sharp vertex in a simple case.
- A sharp edge, meaning a sharp ridge between two vertices. Figures 3(c) and 3(d) show the effect of a sharp edge in a simple case.

The following sections will describe how sharpness was achieved for both the connected and simple modified models.

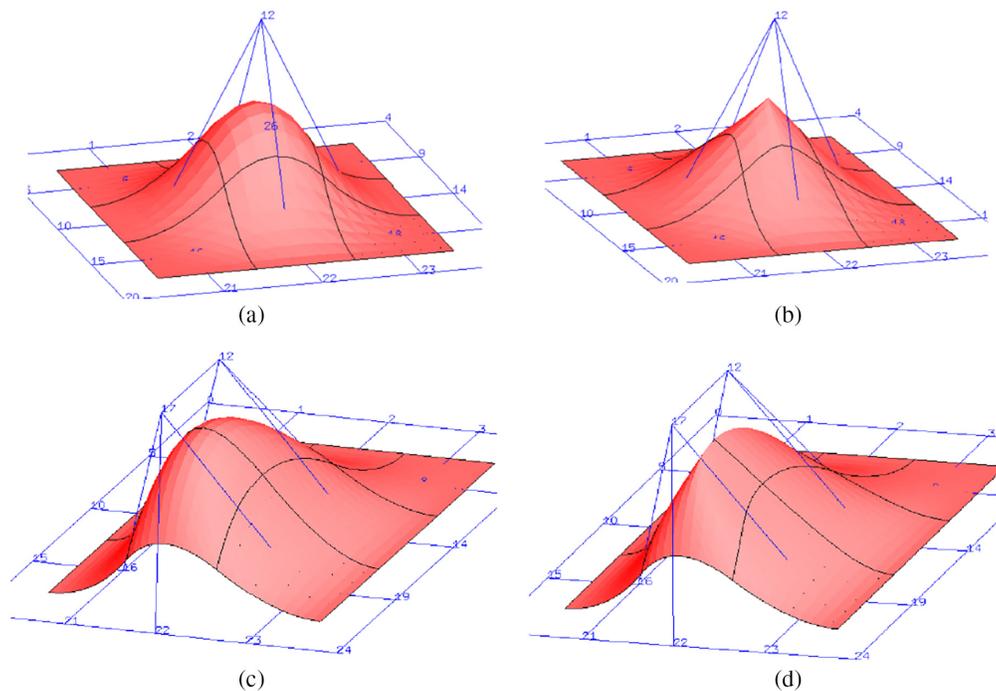


Fig. 3 A comparison of regular and sharp Doo–Sabin models for one implementation. (a) and (b) The same surface but with a varying degree of sharpness at vertex 12, (a) the regular Doo–Sabin and (b) a sharp vertex. (c) and (d) The effect of adding a sharp edge between vertices 12 and 17, and (c) a regular Doo–Sabin left and (d) a sharp edge.

2.2.1 Connected sharpness

The first part of this section will describe how to construct a connected sharp vertex, followed by a description of the connected sharp edges.

Let N be the number of vertices that a certain vertex is connected to. Then, for each iteration of the Doo–Sabin subdivision, the vertex of valence N will be replaced by a face consisting of N new vertices q_i . For the connected model, a sharp vertex is achieved by moving those vertices toward their common center, in effect shrinking the size of the face. This means that the curved limit surface of the face is also made smaller, and the surface getting reparametrized as the face shrinks. This causes the surface to have larger surface derivatives and take on a sharper appearance. A 2D version of the idea can be seen in Fig. 4.

To start, define w , a weight defined to determine how sharp the surface should be around any vertex. The process is defined such that when $w = 0$ the local surface would be the same as traditional Doo–Sabin, but as w increases the topology would be constant while the V-face based on the former vertex would shrink, each vertex being moved linearly toward the center of the V-face. By Nasri,¹⁹ this center is on the surface of the original Doo–Sabin model. Moving each vertex toward the center linearly means that the center will still be on the modified surface.

The basis B-splines that are used to evaluate a regular face gives values depending on a parametrization of the face. Thus, when the face shrinks, the result is the same curve but reparametrized to fit the face, giving a sharper appearance, but keeping all the Doo–Sabin properties including partial derivatives. This holds as long as $w \in [0,1)$. When $w = 1$, the face would be a point, creating a higher knot multiplicity and thus partial derivatives can not be defined.

The average of the vertices in Q is calculated in terms of the subdivision matrix S . Each vertex in Q is constructed by multiplying a vector V with a row of S . A row of S will be denoted s_i , and V is the set of vertices of the original control nodes that are local, that is, relevant to the calculation of the new vertices:

$$q_i = s_i V. \tag{1}$$

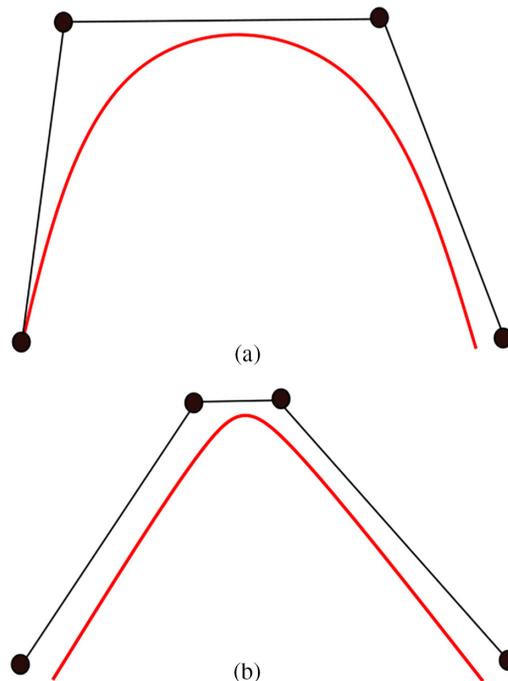


Fig. 4 A 2D illustration of the idea behind modifying the Doo–Sabin model. The red line shows the surface and the black dots the control nodes. As the top two control nodes are moved closer together, the surface gets a sharper curve. How close the two top nodes are is determined by a weight. (a) A low weight and (b) a larger one.

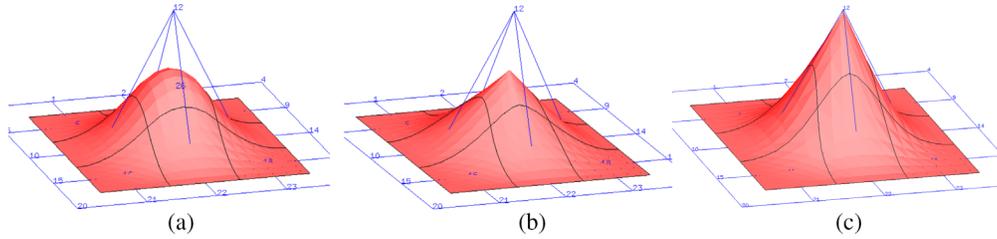


Fig. 5 A comparison of regular and sharp Doo–Sabin models. (a) The regular Doo–Sabin, (b) the connected sharp Doo–Sabin, and (c) the simple sharp Doo–Sabin. The sharp Doo–Sabins have a weight close to 1 at vertex 12 and 0 elsewhere.

Taking the rows s_i of S that fit the above equation and averaging them gives the center c expressed in terms of the original vertices:

$$c = \frac{1}{n} \sum_{j=1}^n q_j = \frac{1}{n} \left(\sum_{j=1}^n s_j \right) V, \quad (2)$$

where n is the number of vertices in Q . Next, change S by replacing s_i with

$$s_i^1 = w \frac{1}{n} \left(\sum_{j=1}^n s_j \right) + (1-w)s_i, \quad (3)$$

for any row s_i corresponding to a vertex $q_i \in Q$. From this, we get

$$s_i^1 V = \left[w \frac{1}{n} \left(\sum_{j=1}^n s_j \right) + (1-w)s_i \right] V = wc + (1-w)q_i. \quad (4)$$

An example of regular versus the sharp Doo–Sabins can be found in Fig. 5.

The method for a connected sharp edge is similar to the method for connected sharp vertices. An edge is turned into a four-sided E-face in one step of the Doo–Sabin algorithm, and we denote the four vertices in this E-face by Q .

As in Sec. 2.2.1, the sharpness is related to the size of the face. Though in general, the edges of the E-face are not parallel, simplifying it slightly the E-face can be seen as a rectangle, the longer side running parallel to the original edge in what can be called the X direction and the shorter going in the Y direction. To create a smaller face, the E-face will be shortened by moving the nodes of the E-face in the Y direction, closer to the original edge. Figure 6 shows this concept.

For the connected sharp Doo–Sabin edge, the vertices of the E-face are paired up, and the vertices of each pair are moved toward their common center. The pairs are based on what side of the E-face they are on. Vertices having the same coordinate in X direction, or equivalently those that are closest to the same vertex on the original edge, are paired together. A weight $w \in [0,1)$ is used to determine the degree of sharpness, as at $w = 1$ partial derivatives are undefined at the center of the face. For the two vertices in each pair, q_i and q_j , the center c is defined as

$$c = \frac{q_i + q_j}{2} = \frac{s_i + s_j}{2} V. \quad (5)$$

Then, s_i and s_j is respectively replaced with

$$s_{i,j}^1 = w \frac{s_i + s_j}{2} + (1-w)s_{i,j}, \quad (6)$$

$$s_{i,j}^1 V = wc + (1-w)q_{i,j}. \quad (7)$$

w acts as a weight determining how close to the original edge the new vertices are. As in the calculation of connected sharp vertices, the center of the surface of the E-face is by Nasri¹⁹ kept

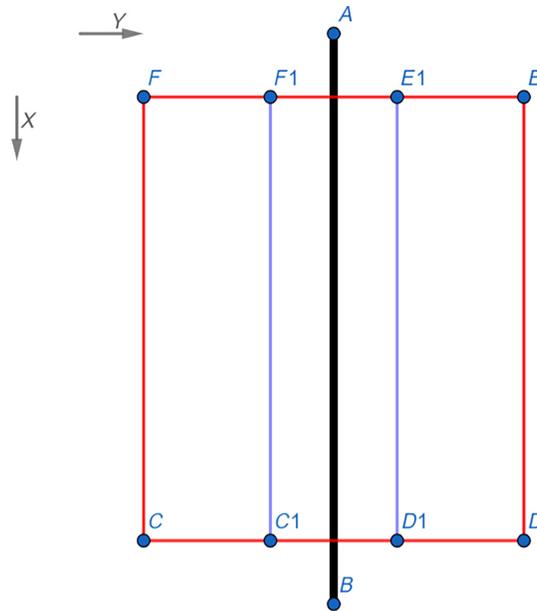


Fig. 6 A 2D illustration of the idea behind sharp edges. A weight is placed on the edge between A and B, which goes in the X direction. The E-face formed by points C, D, E, and F is modified by moving the points in the Y direction, making a new E-face formed by C1, D1, E1, and F1. Edges of the original E-face are shown in red, and the new edges in X direction are shown in blue.

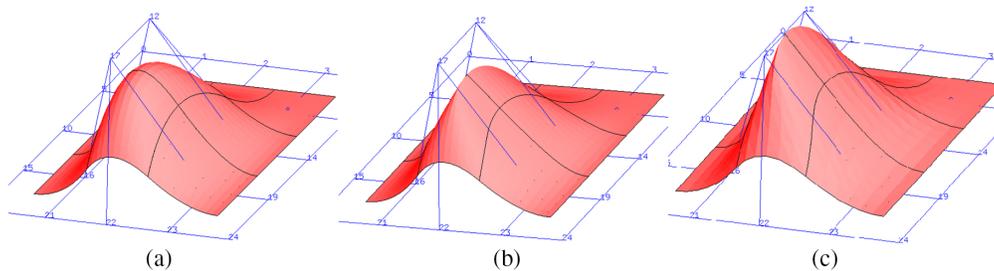


Fig. 7 A comparison of the regular and the sharp Doo–Sabin models. (a) The regular Doo–Sabin, (b) the connected sharp Doo–Sabin, and (c) the simple sharp Doo–Sabin. The sharp Doo–Sabin has weights close to 1 between the vertices 12 and 17, and also between 12 and 7 (further behind 12), and 0 elsewhere.

constant. This is because Nasri states that the center of the E-face will be on the final surface and because q_i and q_j are moved toward their common center in equal amount, so the center is unchanged. This means that while the shape changes, it is not a dramatic change to the surface in terms of the distance a point is moved. An example of a sharp edge can be seen in Fig. 7.

2.2.2 Simple sharpness

Simple sharpness uses many of the same terms and ideas defined in Sec. 2.2.1 and will not be redefined or explained in detail here.

This version, like the previous, manipulates S to move certain nodes together to simulate sharpness. In this case, for a sharp vertex V_w with weight w , all vertices Q that are part of V_w 's V-face are moved linearly toward V_w . As before when $w = 0$ the surface is locally a standard Doo–Sabin surface, and at $w = 1$ all nodes in Q are at V_w , which gives a completely linear surface but an increased knot multiplicity.

Note that $V_w \in V$ and can be written as a weighted sum of the elements of V . This means that it can be written as a vector that can be added to a row s_i of S . To get the desired effect, the relevant rows s_i are replaced with

$$s_i^1 = wV_w + (1 - w)s_i. \quad (8)$$

An example can be seen in Fig. 5.

As for simple sharp edges, Fig. 6 still gives an understanding of the general concept. Let w be a weight set between two vertices V_w^1 and V_w^2 in V_0 . Vertices in the E-face Q resulting from the edge between V_w^1 and V_w^2 are moved toward a point V_w between V_w^1 and V_w^2 :

$$V_w = W_i V_w^1 + (1 - W_i) V_w^2. \quad (9)$$

V_w^1 , V_w^2 , and V_w can all be written as weighted averages of the elements of V_0 . This allows them to be added to s_i . We will refer to the row of weights used in these averages as v_w^1 , v_w^2 , and v_w , respectively.

V_w and W_i are different for each $q_i \in Q$ and is based on the row s_i of S corresponding to q_i . W_i will be the ratio between the values of s_i corresponding to V_w^1 and V_w^2 . In other words, if s_1 and s_2 are the multipliers of V_w^1 and V_w^2 in the calculation of q_i , then

$$W_i = \frac{s_1}{s_1 + s_2}. \quad (10)$$

This results in

$$s_i^1 = wv_w + (1 - w)s_i = w[W_i v_w^1 + (1 - W_i)v_w^2] + (1 - w)s_i. \quad (11)$$

Note that the ordering of V_w^1 and V_w^2 does not affect the final outcome, as $(1 - W_i) = \frac{s_2}{s_1 + s_2}$, so changing the order would lead to the same weights. An example of a simple sharp edge can be seen in Fig. 7.

The connected Doo–Sabin of Sec. 2.2.1 keeps the center of the modified V-face constant, which by Nasri¹⁹ means that the center of the local surface is kept constant. This leads to comparatively small changes to the surface compared to the simple Doo–Sabin, which moves the surface closer to the control vertices.

On the other hand, the simple sharp Doo–Sabin calculates the modified rows of S without needing to know the values of surrounding nodes. For the connected sharp Doo–Sabin, it is needed to know which of the surrounding vertices are part of the same V- or E-face. For E-faces even more knowledge of the local topology is needed. The modifications made to the connected version also lead to a larger area of support, as each node q_i now also depends on the control vertices of the other elements of Q .

2.3 Evaluating Modified Doo–Sabin on Modeling the Right Ventricle

We evaluated the behavior of the sharp Doo–Sabin methods to model the RV geometry and compared it with the original Doo–Sabin method. This was done using a Doo–Sabin model based on a general RV shape. There were three models based on the general shape, one used the original Doo–Sabin model, one used the connected sharp Doo–Sabin, and one used the simple sharp Doo–Sabin. The two latter models included added weights at the apex and along the boundary between the free wall and the septum, but in all other aspects, such as topology, number of vertices and their positions in 3D space, and the three models were identical. The weights were the same in the two modified models.

Using the EchoPAC software with a custom version of 3D AutoRVQ (GE Vingmed Ultrasound, Horten, Norway), the three models were fitted to 16 3D ultrasound images. AutoRVQ uses a Kalman filter¹⁵ approach to fit the images, and the initial state of models was the same.

AutoRVQ uses six points placed manually for each file. These points are used to determine the scale and location of the model, and the models are attracted to the landmarks to make sure that the model fits with the users choices. One landmark is placed at the apex, one at the free-wall, one at each border of the septum and freewall on posterior and anterior side, and two are placed around the tricuspid valve. These were placed by a senior cardiologist and were the same for each of the three models.

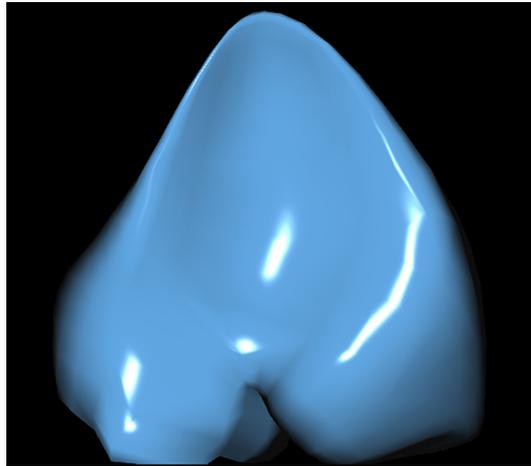


Fig. 8 An example of the ground truth, picture taken at end diastole.

The workflow for placing landmarks in the custom version of 3D AutoRVQ was not altered. The user aligns the view and places landmarks on the image, then AutoRVQ calculates the output surface, a process that takes a couple of seconds. This process has good stability, meaning that minor changes in the input image or in the landmarks placements should only lead to small changes in the final surface.

The three models were compared with ground truths, created by the cardiologist using EchoPAC. EchoPAC has the ability to manually fit the models after initial fitting using the Kalman filter. This was done by manually changing the initial surface where it does not match the image. The cardiologist would use EchoPAC to get an initial result surface and then manually edit it wherever it seemed necessary. This manual editing was only done on the ground truth not on the testing models. Measurement and analysis were performed in two specific frames, one at end systole and another at end diastole. An example of the ground truth at end diastole can be seen in Fig. 8.

2.3.1 Doo–Sabin models of the right ventricle

Bersvendsen et al.¹⁰ used the Doo–Sabin models to model the shape of individual right heart ventricles. The Doo–Sabin model in this paper was based on the Bersvendsen model but simplified by removing 6 out of 32 vertices and making minor changes to vertice positions. This was done as we wanted to show that a model could accurately model sharpness using the weighted methods discussed in Secs. 2.2.1 and 2.2.2 rather than additional vertices.

The edges along the interface of the interventricular septum and the RV free wall were made sharp by setting nonzero weights as explained in Secs. 2.2.1 and 2.2.2. These weights were all set to 0.9. The vertex weight in apex was set to 0. The sharp edges already create a sharp geometry in the apical region. The weights were manually tuned using training data separate from the data used for the comparison.

2.3.2 Measuring the error of the models

For each comparison, four meshes were used: the manually segmented ground truth, a regular Doo–Sabin model, and the two sharp Doo–Sabin model differing from the regular Doo–Sabin only in the added weights and differing from each others only in how the sharpness is implemented as discussed in Sec. 2.2. The distribution of weights was discussed in Sec. 2.3.1.

Several metrics were used to measure the error of the models compared to the ground truth. All focused on the models at peak diastole and peak systole.

The first measure used was the Hausdorff distance. Two frames from each 3D echocardiographic cine loop were used: one at end diastole and another at end systole. Each Doo–Sabin model was compared with the ground truth in each frame, and the errors for each frame were

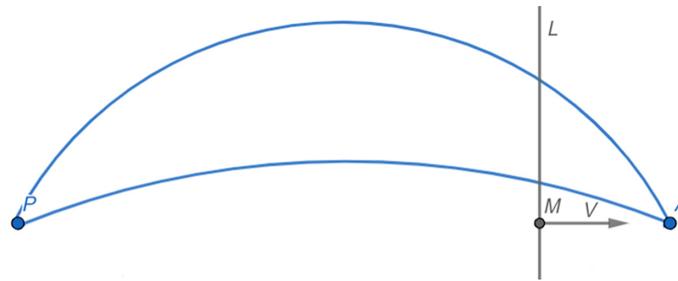


Fig. 9 A simple figure showing how points were removed on the regions close to the anterior of the RV. The blue lines form the RV, and A and P were set by the cardiologist. M is a weighted average of A and P , $M = 0.2P + 0.8A$. A vector V is determined by $V = A - P$. In the image, it is set to start in M . Together M and V determine a plane L . Any point on the RV to the right of L was removed.

added together, giving a total error for each model. This was done across all 16 3D ultrasound images, and the errors across files were averaged.

The regions close to the anterior were not taken into account due to low image quality.⁹ This decision was taken by the cardiologist, as the confidence of the manual segmentation was also low. The exact part to remove was determined based on the data from the cardiologist's ground truth: the cardiologist had established two points as the posterior P and the anterior A of the RV. A plane cropping away the most anterior part of the model was determined using a point and a normal vector. The point was $0.2P + 0.8A$ and the normal vector was $A - P$. Any point on the anterior side of the plane was removed. The 4:1 ratio of the point was chosen based on a manual assessment of where the image quality usually deteriorated. A figure illustrating how the removal was done is seen in Fig. 9.

The second measure used for the evaluation was the mean surface distance (MSD). The parts closest to the anterior were not taken into consideration. The error on both peak systole and diastole was averaged, and an improvement in percent was calculated.

The third measure was the volume of the models at peak diastole and systole. This was calculated using a surface integral based on the Divergence Theorem. The error in milliliters was calculated, as was stroke volume and ejection fraction.

3 Results

3.1 Hausdorff Measure

Figure 10 shows the error as described in Sec. 2.3.2 in millimeters of both models for each 3D ultrasound image compared to the ground truth. Note that these errors are averaged over both measured frames in each ultrasound image.

For each file, we calculated the percentwise improvement of connected and simple sharp models over the regular Doo–Sabin, calculated by $100 \frac{R-S}{R}$ where R is the error of the regular Doo–Sabin and S is the error of the sharp Doo–Sabin. These values were then averaged. On average, the connected sharp model improved on the regular model by 0.6% and the simple sharp model by 0.1%.

An example of the three models fitted to the same ultrasound image is seen in Fig. 11. Sharpness is clearly seen at the top of the two right-most images.

3.2 Mean Surface Distance

Table 1 shows the error in terms of MSD. The connected sharp model has an average improvement of 5.3% and a median improvement of 2.6%. The simple sharp model has an average improvement of 2.7% and a median improvement of 6.2%.

An example of the MSD of each surface type can be found in Fig. 12. Highest error can be seen at the border between septum and freewall and at the atrioventricular plane. For the

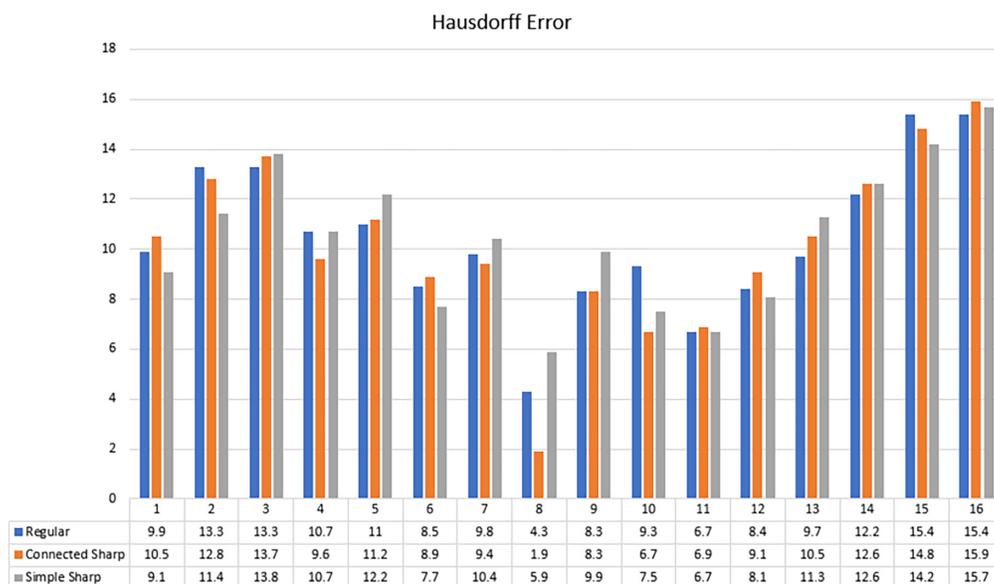


Fig. 10 Hausdorff error of each file for each of the three models. Error is in millimeters.

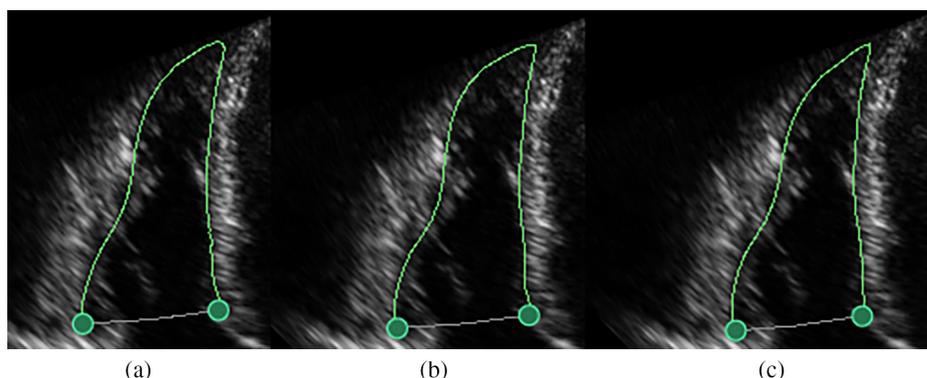


Fig. 11 The regular Doo–Sabin and the two sharp models fitted to a 3D ultrasound image. (a) The smooth model, (b) the connected sharp, and (c) the simple sharp. Sharpness can clearly be seen at the top in the two right-most models, marking the border of the septum and free-wall.

sharp models, the error is significantly lower compared with the regular model at the posterior side.

3.3 Volume Measures

Improvement in terms of volume was calculated by taking the absolute value of the difference between the models and ground truth at both peak diastole and systole and averaging them. On average, the smooth Doo–Sabin had an error of 10 ml, the connected sharp had an error of 8 ml, and the simple sharp had an error of 6 ml. This means the connected sharp had an average improvement of 20% and simple sharp 40%.

Bland–Altman plots²⁰ of the volumes at ED are shown in Fig. 13.

Stroke volumes were calculated based on volumes, and for the three test-models, the absolute difference in stroke volume was determined. The results are shown in Fig. 14. On average, the regular Doo–Sabin had an error of 6 ml, whereas both of the sharp models had an average error of 5 ml, meaning an improvement of 16%.

Taking the absolute value of the difference between the ground truth and each of the models, on average the error of ejection fraction is 4.9 for the regular Doo–Sabin, 4.9 for the connected sharp Doo–Sabin, and 6.0 for the simple sharp Doo–Sabin.

Table 1 Results of comparison of normal Doo–Sabin models and sharp Doo–Sabin models against a ground truth traced by a cardiologist. Error between ground truth and the other models were determined using the MSD. Calculation of improvements are detailed in Sec. 3.1.

File number	Error connected sharp Doo–Sabin (mm)	Error simple sharp Doo–Sabin (mm)	Error regular Doo–Sabin (mm)	Improvement connected sharp Doo–Sabin (%)	Improvement simple sharp Doo–Sabin (%)
1	2.31	2.36	2.16	−6.79	−9.24
2	2.6	2.14	2.78	6.41	22.96
3	2.32	2.25	2.38	2.81	5.51
4	2.06	1.93	2.15	4.56	10.34
5	2.55	2.28	2.47	−3.1	7.74
6	1.85	1.62	1.91	2.95	15.18
7	1.99	1.92	2.06	3.23	6.81
8	0.33	1.19	0.71	53.87	−67.23
9	1.87	1.89	1.85	−1.42	−2.45
10	1.33	1.67	1.71	22.18	2.74
11	1.58	1.54	1.63	2.9	5.26
12	2.24	1.79	2.2	−2.04	18.55
13	1.99	1.82	1.91	−4.4	4.92
14	2.13	2.02	2.11	−0.76	4.21
15	3.35	3.13	3.4	1.34	8.00
16	3.17	2.92	3.25	2.4	10.17
Average	2.11	2.03	2.17	5.26	2.72

4 Discussion

The aim of this work was to describe a compact data model that can adapt to a large variety of shapes for real life applications and to illustrate its usefulness in medical modeling by testing the model on the complex RV geometry.¹ This work builds on the previous studies by Bersvendsen et al.,¹⁰ Orderud and Rabben,¹¹ and others.

We have shown that it is possible to construct versions of the Doo–Sabin model that allows for sharper vertices and edges, similar to the way that deRose et al.²¹ did for the Catmull–Clark method. These models are easy to compute as they are local processes, and the calculation of B-splines and subdivision matrices is well-known and easy. The simple Doo–Sabin is the easiest to calculate of the two. The sharp Doo–Sabins gives a more flexible model compared with the classic Doo–Sabin while keeping useful properties such as locality and the easy calculation of partial derivatives as long as $w \neq 1$.

This work focuses on the modeling of 3D images. The Doo–Sabin process forms a 2D surface that models the walls in 3D space. A generalization to 2D images could be possible by replacing the Doo–Sabin surface with one-dimensional (1D) B-splines, a close 1D equivalent. Weighting of B-splines to achieve similar effects as presented here could be possible. Several measures were used to determine the error of the models compared with the ground truth, as detailed in Sec. 2.3.2. Not much improvement was found in terms of Hausdorff measure. This is to be expected, as the Hausdorff measure solely determines the location with the highest error. As detailed in Sec. 2.3, the models are attracted to the landmarks set by the user. This in particular

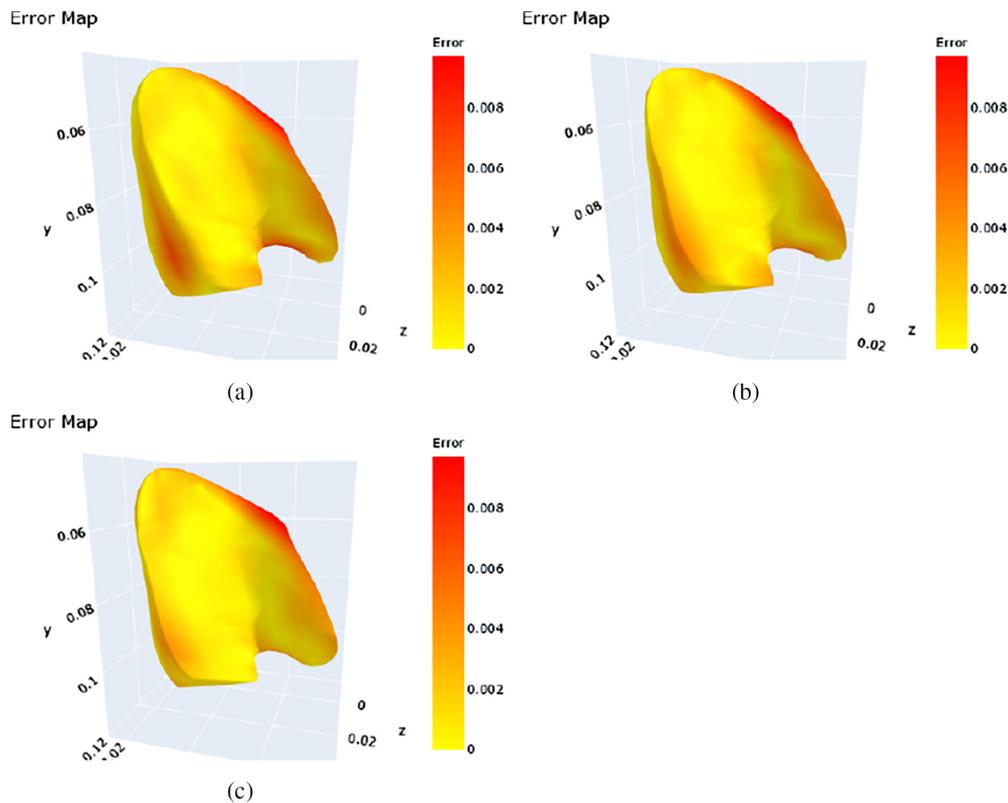


Fig. 12 A comparison of the error of the smooth, connected sharp, and simple sharp models in one example, compared to the ground truth. Each point on each surface is colored depending on its distance to the ground truth. Yellow indicated low error and red high error, all models use the same scale. (a) The regular Doo–Sabin, (b) the connected sharp Doo–Sabin, and (c) the simple sharp Doo–Sabin.

holds for the borders of the septum and freewalls. These are the regions where one might expect a smooth model struggle with filling in the RV shape, as there is often a sharp edge here. The moving of the model to the landmarks ensures that does not happen, but the regular Doo–Sabin might still struggle with fitting the general sharp shape and instead go into the septum or freewall. Both the connected and simple sharp model do better in terms of MSD, and Fig. 12 indicates that the regular Doo–Sabin struggles at modeling the sharp corners of the RV, as it has larger error at the border of the septum and freewall.

Both of the sharp models show good improvement in terms of volume accuracy and stroke volume as shown in Sec. 3.3. Despite this, no improvements were found in terms of ejection fraction. A likely explanation of this is that EchoPAC uses a Kalman filter¹⁵ process to estimate the walls of the RV. That means that the estimate at each frame is based on the previous frame, and so an error in model shape can carry over from one frame to the next if the image data do not have good enough quality to correct it in the failing region. It is possible that if a model is too large at peak diastole it will also be too big at peak systole in roughly the same proportion, making the ejection fraction about the same.

The connected sharp Doo–Sabin of Sec. 2.2.1 was developed so that modifying already existing Doo–Sabin models would be easy. By Nasri,¹⁹ the method described in this article ensures that the center of each face is unchanged by changed weights and that the center of the local surface remains the same. This makes it easy to modify models that already exist. These methods take advantage of the fact that the Doo–Sabin model is local, so modifying the subdivision matrix in some places will not cause changes outside of a small region. It has an average improvement in terms of MSD of 5.3% and an improvement in volume of 14.4%.

The connected sharp Doo–Sabin did better on MSD on average, but the simple sharp Doo–Sabin did better in median. A partial cause of this is file nr 8, as seen in Table 1, which is unusual

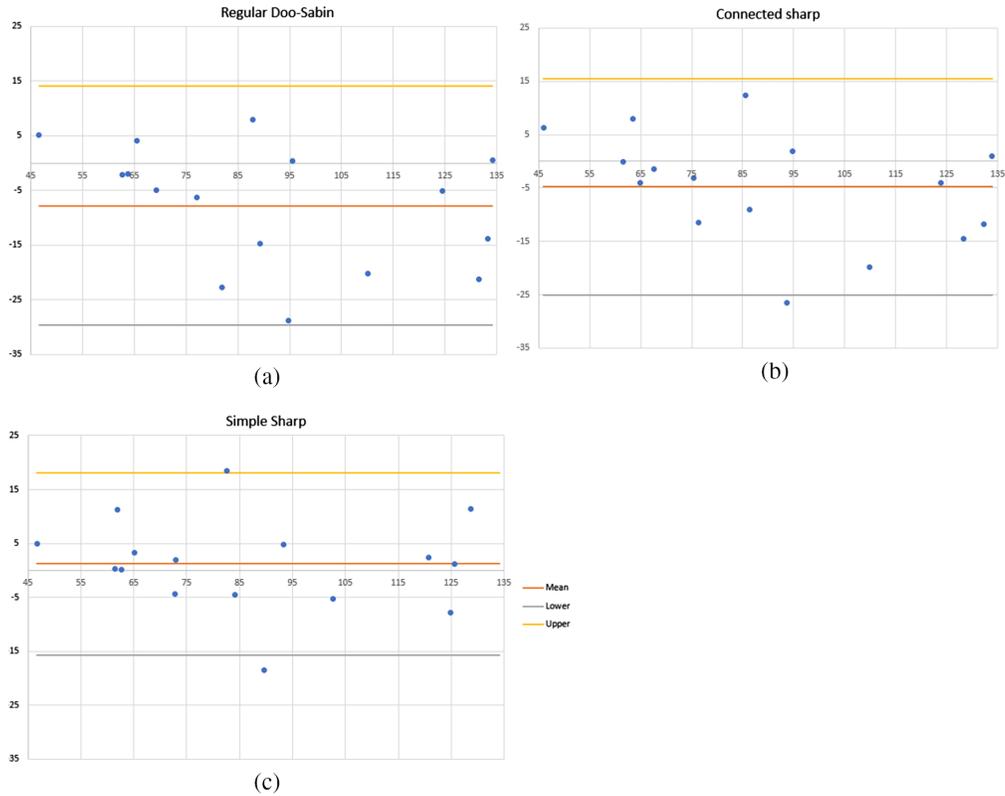


Fig. 13 Three Bland–Altman graphs comparing the regular Doo–Sabin and the two sharp models to the ground truth. The measure is error in volume estimation at peak diastole. (a) The regular Doo–Sabin and has average -7.8 and standard deviation 11.2 . (b) The connected sharp model and has average -4.8 and standard deviation 10.4 . (c) The simple sharp model and has average 1.2 and standard deviation 8.6 .

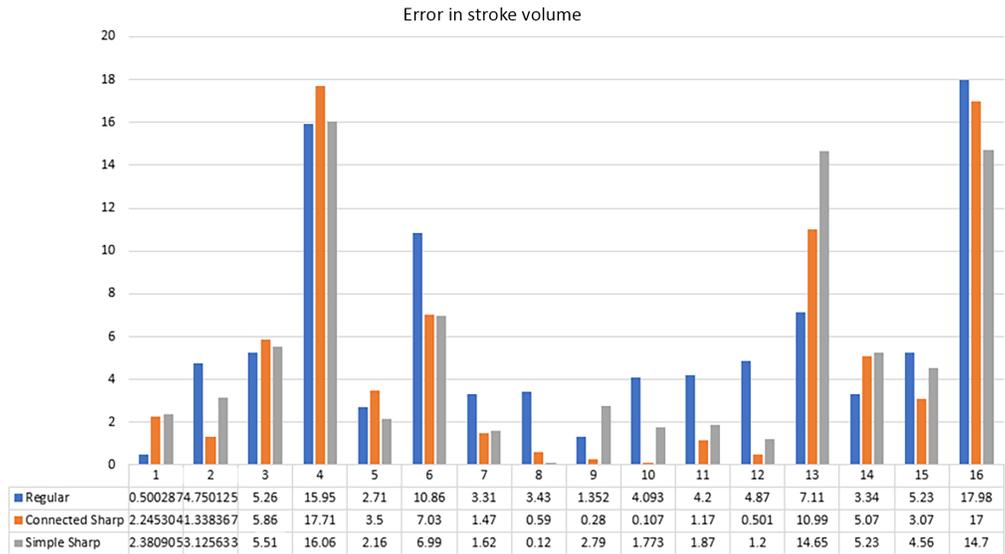


Fig. 14 The error in stroke volume for each model for each file. Error is in absolute difference in stroke volume and is measured in millimeters.

for its low error for all three models, meaning that any small deviation leads to a big increase in percent-wise error.

The simple sharp Doo–Sabin has the advantage of being easier to compute. It is a more local process and needs less information about the local topology. The substantial improvement of the

simple sharp Doo–Sabin in terms of volume makes it a promising tool, and it is possible that the support for larger deformations are an advantage in RV shape segmentation.

The sharp Doo–Sabin models could prove useful in a many applications, and in this article we have shown improvements in capturing the shape of the RV. While there was no improvement in terms of Hausdorff measure, the sharp Doo–Sabins did better than the original when compared by MSD and volume measures. This shows the advantages of the sharp Doo–Sabin when dealing with anatomy that has sharper and edgier structure than what the regular Doo–Sabin excels at. Here, the comparison was done on ultrasound images, but similar improvements would be expected in other imaging modalities or even nonmedical use cases.

The models have good stability, small differences in user placed landmarks or in ultrasound image should only lead to minor changes in output surface. The steps of the model fitting process are the same as for regular EchoPAC, and there is not a big increase in time. After landmarks are placed, the calculation of the model only takes a couple of seconds.

The study has some limitations. Looking at the results of Table 1, we see that in 6 out of 16 3D ultrasound images the regular Doo–Sabin did better compared with the connected sharp Doo–Sabin model on MSD. The same is true in 3 out of 16 images in the case of simple sharp Doo–Sabin. Out of the 32 frames volume was evaluated at, the regular Doo–Sabin was more accurate in 12 compared to both models.

Segmentation of an ultrasound image will always be subject to human interpretation, and it is therefore not possible to establish an indisputable ground truth. Also, the image quality of some of the ultrasound images was poor, leading to uncertainty in both manual segmentation and the software’s detection of RV walls. The dataset was limited, with little variation in heart size and pathology. The weights in the sharp Doo–Sabin model are added manually when constructing the model and determined using training data. This obviously creates some subjectivity in the model.

For further research, applying the modified Doo–Sabin methods to other aspects of the heart could lead to improvements. The model could also be used on nonmedical modeling. Sub-division surfaces are popular in computer animation, and these changes would allow Doo–Sabin to model a larger variety of surfaces. An improvement to the current sharp Doo–Sabins would be for the weights to be automatically adjusted, instead of manually.

5 Conclusion

The sharp Doo–Sabin models both outperformed the normal Doo–Sabin model when it comes to modeling the RV shape in terms of MSD and volume. While both the connected sharp Doo–Sabin model and the simple sharp Doo–Sabin model on average came closer to the ground truth than the regular Doo–Sabin, the simple sharp model did better in terms of the volume accuracy. It also did better in median MSD, whereas the connected sharp did better on the average MSD. The connected sharp Doo–Sabin model also did slightly better on Hausdorff measures and was better on stroke volume. On ejection fraction, the connected sharp did as well as the regular, whereas the simple did worse.

Needing fewer control nodes to model sharp contours, the sharp models can achieve a more effective calculation and a more robust model. This indicates its improved usefulness which might be interesting to explore in further research.

Disclosures

Several of the authors work at GE Vingmed Ultrasound.

Acknowledgments

Funding for the first author’s PhD comes from the University of Oslo, at the INIUS project. Software and ultrasound images were provided by GE Vingmed Ultrasound.

References

1. F. F. Faletta, J. Narula, and S. Y. Ho, *Atlas of Non-Invasive Imaging in Cardiac Anatomy*, pp. 91–94, Springer Nature, Cham, Switzerland (2020).
2. J. A. Noble and D. Boukerroui, “Ultrasound image segmentation: a survey,” *IEEE Trans. Med. Imaging* **25**, 987–1010 (2006).
3. L. Mertens and M. Friedberg, “Imaging the right ventricle: current state of the art,” *Nat. Rev. Cardiol.* **7**, 551–563 (2010).
4. F. Hadda et al., “Right ventricular function in cardiovascular disease, part II: pathophysiology, clinical importance, and management of right ventricular failure,” *Circulation* **117**, 1717–1731 (2008).
5. P. de Groot et al., “Right ventricular ejection fraction is an independent predictor of survival in patients with moderate heart failure,” *J. Am. Coll. Cardiol.* **32**, 948–954 (1998).
6. S. Apostolakis and S. Konstantinides, “The right ventricle in health and disease: insights into physiology, pathophysiology and diagnostic management,” *Cardiology* **121**, 263–273 (2012).
7. N. F. Voelkel et al., “Right ventricular function and failure,” *Circulation* **114**, 1883–1891 (2006).
8. T. L. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*, 2nd ed., pp. 23–35, Academic Press, Cambridge, Massachusetts (2013).
9. E. Ostfeld et al., “Manual correction of semi-automatic three dimensional echocardiography is needed for right ventricular assessment in adults; validation with cardiac magnetic resonance,” *Cardiovascular Ultrasound* **10**, 1 (2012).
10. J. Bersvendsen et al., “Automated segmentation of the right ventricle in 3D echocardiography: a Kalman filter state estimation approach,” *IEEE Trans. Med. Imaging* **35**, 42–51 (2016).
11. F. Orderud and S. Inge Rabben, “Real-time 3D segmentation of the left ventricle using deformable subdivision surfaces,” in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, pp. 1–8 (2008).
12. E. Dikici, “Ultrasound cardiac modeling, segmentation and tracking,” PhD Thesis, Norwegian University of Science and Technology (2013).
13. H. S. Bølviken et al., “Modified Doo–Sabin modeling of the right ventricle,” *Proc. SPIE* **11319**, 1131909 (2020).
14. T. Heimann and H.-P. Meinzer, “Statistical shape models for 3d medical image segmentation: a review,” *Med. Image Anal.* **13**, 543–563 (2009).
15. R. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.* **82**, 35–45 (1960).
16. E. Smistad and F. Lindseth, “Real-time tracking of the left ventricle in 3D ultrasound using Kalman filter and mean value coordinates,” in *Proc. MICCAI Challenge Echocardiogr. Three-Dimensional Ultrasound Segmentation*, pp. 65–72 (2014).
17. D. W. H. Doo, “A subdivision algorithm for smoothing down irregular shaped polyhedrons,” in *Proc. Interactive Tech. Comput. Aided Des.*, pp. 157–165 (1978).
18. E. Catmull and J. Clark, “Recursively generated B-spline surfaces on arbitrary topological meshes,” *Comput.-Aided Des.* **10**, 350–355 (1978).
19. A. Nasri, “Polyhedral subdivision methods for free-form surfaces,” *ACM Trans. Graphics* **6**, 29–73 (1987).
20. J. M. Bland and D. G. Altman, “Statistical methods for assessing agreement between two methods of clinical measurement,” *Lancet* **327**(8476), 307–310 (1986).
21. T. Derose, M. Kass, and T. Truong, “Subdivision surfaces in character animation,” in *Proc. ACM SIGGRAPH Conf. Comput. Graphics*, Vol. 32, pp. 85–94 (1998).

Håkon Strand Bølviken received his master’s degree in theoretical mathematics from the Norwegian University of Science and Technology (NTNU) in Trondheim. Currently, he is pursuing his PhD in informatics at the University of Oslo, specializing in analysis and segmentation of cardiac images.

Sten Roar Snare received his MSc degree in engineering cybernetics and his PhD in medical technology from NTNU. He worked at the Department of Circulation and Medical Imaging at NTNU until fall 2011, when he was hired at GE. His research interests are medical ultrasound technology, and in particular image analysis for automated quantification of physiological parameters from ultrasound images.

Biographies of the other authors are not available.