



Published in final edited form as:

*Med Image Anal.* 2018 May ; 46: 215–228. doi:10.1016/j.media.2018.03.009.

## SiSSR: Simultaneous Subdivision Surface Registration for the Quantification of Cardiac Function from Computed Tomography in Canines

Davis M. Vigneault<sup>a,b,c,\*</sup>, Amir Pourmorteza<sup>b</sup>, Marvin L. Thomas<sup>d</sup>, David A. Bluemke<sup>b</sup>, and J. Alison Noble<sup>a</sup>

<sup>a</sup>Institute of Biomedical Engineering, Department of Engineering, University of Oxford

<sup>b</sup>Department of Radiology and Imaging Sciences, Clinical Center, National Institutes of Health

<sup>c</sup>Tufts University School of Medicine, Sackler School of Graduate Biomedical Sciences

<sup>d</sup>Division of Veterinary Resources, National Institutes of Health


### Abstract

Recent improvements in cardiac computed tomography (CCT) allow for whole-heart functional studies to be acquired at low radiation dose ( $< 2\text{mSv}$ ) and high-temporal resolution ( $< 100\text{ms}$ ) in a single heart beat. Although the extraction of regional functional information from these images is of great clinical interest, there is a paucity of research into the quantification of regional function from CCT, contrasting with the large body of work in echocardiography and cardiac MR. Here we present the Simultaneous Subdivision Surface Registration (SiSSR) method: a fast, semi-automated image analysis pipeline for quantifying regional function from contrast-enhanced CCT. For each of thirteen adult male canines, we construct an anatomical reference mesh representing the left ventricular (LV) endocardium, obviating the need for a template mesh to be manually sculpted and initialized. We treat this generated mesh as a Loop subdivision surface, and adapt a technique previously described in the context of 3-D echocardiography to register these surfaces to the endocardium efficiently across all cardiac frames simultaneously. Although previous work performs the registration at a single resolution, we observe that subdivision surfaces naturally suggest a multiresolution approach, leading to faster convergence and avoiding local minima. We additionally make two notable changes to the cost function of the optimization, explicitly encouraging plausible biological motion and high mesh quality. Finally, we calculate an accepted functional metric for CCT from the registered surfaces, and compare our results to an alternate state-of-the-art CCT method.

### Graphical abstract

\*Corresponding author. davis.vigneault@gmail.com (Davis M. Vigneault).

**Publisher's Disclaimer:** This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



## Keywords

cardiac computed tomography; regional cardiac function; Loop subdivision surface; personalized cardiac mesh generation

## 1. Introduction

Regional cardiac dysfunction commonly precedes changes in global metrics such as ejection fraction (EF), and is useful in the characterization of subtle impairment and monitoring of treatment response (Tee et al., 2013). Techniques developed for cardiac magnetic resonance imaging (CMR), including harmonic phase (HARP) imaging (Osman et al., 1999), strain encoded (SENC) imaging (Osman et al., 2001), and displacement encoding with stimulated echos (DENSE) imaging (Aletras et al., 1999), have enabled the measurement of biomarkers for regional cardiac function, including strain, strain rate, and torsion. However, due to their complex and time-consuming acquisition and analysis, further complicated by contraindication of CMR in patients with pacemakers, these techniques have been confined to research studies in academic centers.

Echocardiographic techniques for measuring function are widely deployed in the clinic due to low cost and real-time, bedside acquisitions. However, in terms of cardiac function, echocardiography is subject to certain limitations. Tissue Doppler echocardiography (TDI) is used to measure cardiac motion, but is highly angle-dependent. Speckle-tracking, which has been developed both for 2-D and 3-D sequences, currently provides measurements which are less sensitive to probe angle than Doppler techniques (Tee et al., 2013). Transthoracic echocardiography (TTE) is non-invasive, but unfortunately is limited by the available acoustic window. Transesophageal echocardiography (TEE) avoids these issues, but is an uncomfortable, semi-invasive procedure. All echocardiographic techniques are limited by operator dependence, and 2-D techniques are additionally subject to through-plane motion errors.

Historically, cardiac computed tomography (CCT) has been clinically infeasible for functional studies due to high radiation dose and low temporal resolution (Suinesiaputra et al., 2016). However, recent advances in CCT have dramatically reduced the necessary radiation dose ( $< 2\text{mSv}$ ) and increased the temporal resolution ( $< 50\text{ms}$ ) for whole-heart functional studies (Tee et al., 2015), spurring interest in the use of CCT in assessing cardiac function. Moreover, CCT is an attractive modality for functional studies given its high spatial resolution (on the order of  $0.25\text{mm}$ ), inherently volumetric data, fast scan times

(single heart beat), ubiquity in clinic, and increasing inclusion in cardiac imaging protocols due to the depiction of coronary anatomy (Tee et al., 2013). Therefore, fast, automated, and robust algorithms for extracting regional functional information from CCT are of immediate interest.

Initial efforts to recover strain from CCT have largely focused on application of tools originally developed for 2-D ultrasound (Ogawa et al., 2006) to resliced CCT sequences (Helle-Valle et al., 2010; Tee et al., 2013). This approach is limited, as it relies on myocardial features which are largely absent from CCT and discards useful volumetric information. An alternative metric and approach has been proposed by Pourmorteza et al. (2012), who employed the coherent point drift (CPD) method (Myronenko et al., 2007; Myronenko and Song, 2010) to register a template mesh representing the left ventricular (LV) endocardium at end diastole to each subsequent frame. In CPD, the “moving” point set is represented as a Gaussian Mixture Model (GMM). Using Expectation Maximization (EM), the GMM is deformed so as to maximize the probability of the “fixed” point set having been sampled from it. “Coherent” motion is enforced by constraining the motion of the GMM centroids relative to one another. CPD can be used for non-rigid registration, and does not prescribe any specific transformation model. Having performed the registration, Pourmorteza et al. (2012) calculated for each face the square root of the ratio between the face’s area and the area of the corresponding face at end diastole. This metric, which they termed Stretch QUantifier for Endocardial Engraved Zones (SQUEEZ), may be interpreted as a surrogate for cardiac contraction ( $< 1$ ) or expansion ( $> 1$ ) in the circumferential-longitudinal plane. SQUEEZ is well suited to volumetric CCT data, and was subsequently shown to correlate with strain obtained by tagged CMR (Pourmorteza et al., 2015); to have clinical utility in the planning of cardiac resynchronization therapy (Behar et al., 2017); and to be robust to noise in low-dose CCT (Pourmorteza et al., 2018). After detailing the technical contributions of our proposed method, we calculate SQUEEZ and compare our results to those obtained from the CPD method.

The objective of our method is conceptually similar to Pourmorteza et al. (2012) in that we derive SQUEEZ from a series of registered meshes. However, rather than registering meshes using CPD, we adapt and extend a method which has been successfully used for cardiac modeling and segmentation in the context of 3-D echocardiography (Stebbing, 2014; Stebbing et al., 2015). In Stebbing (2014), a subdivision surface is registered to an arbitrary point set obtained using an off-the-shelf edge detector. A notable advantage of this method is that the subdivision surfaces are registered *across all frames simultaneously*, contrasting with the conceptually simpler “frame-to-frame” and “reference-frame” formulations. In the frame-to-frame formulation, consecutive pairs of frames are registered, leading to propagation of error throughout the cardiac cycle. In the reference-frame formulation, each frame is registered to the end diastolic frame, resulting in independent motion estimates. While this may be desirable in some scenarios, it may also lead to discontinuities in the displacements measured between successive frames (as the transformation estimated by a registration algorithm is not in general a smooth function of its inputs). By formulating the problem as a single, global optimization, the method constructed by Stebbing (2014) avoids both drawbacks.

In this paper, we present the Simultaneous Subdivision Surface Registration (SiSSR) method, which extends the state-of-the-art method for Loop subdivision surface registration developed by Stebbing (2014) in order to register an anatomical reference mesh representing the endocardium of the LV to a 3-D+time cardiac image sequence (Fig. 1). The most notable differences between SiSSR and its predecessor are:

- An anatomical reference mesh is constructed directly from the (segmented) image data, obviating the need for a bespoke template mesh representing the structure of interest, and accounting for the wide anatomical variability between canines.
- Two notable changes are made to the cost function. First, the velocity, rather than acceleration, of the (coarse) mesh vertices are regularized, more plausibly modeling real cardiac motion. Second, (coarse) mesh faces with high aspect ratios are penalized, thereby explicitly discouraging solutions with low mesh quality.
- Subdivision surfaces naturally lend themselves to a multiscale approach, which may decrease computation time while directing the optimization away from local minima. Therefore, we perform an initial registration, explicitly subdivide the result, and use the subdivided result to initialize a second pass.
- Stebbing (2014) evaluates the method primarily as a segmentation algorithm, and stops short of calculating functional parameters. We calculate SQUEEZ, and compare our results with the CPD approach.
- The subdivision method was originally developed in the context of 3-D echocardiography. Here, we instead investigate 3-D+time CCT, which has received relatively less attention and which we feel will become increasingly common clinically in the near future.

The structure of this paper is as follows. In Section 2 we present our methodology in detail, specifically highlighting points of difference between our and prior methods. In Section 3, we provide technical details pertaining to the implementation. In Section 4, we describe the canine infarct model, imaging parameters, and validation against the CPD method. In Section 5, we present qualitative and quantitative results. In Section 6 we offer a discussion of our findings, its limitations, and future directions. In Section 7, we provide a summary of the paper.

Unless otherwise noted or obvious from context, images and meshes are denoted by uppercase script letters (e.g.,  $\mathcal{Q}$ ); matrices by uppercase, boldface letters (e.g.,  $\mathbf{X}$ ); tuples and vectors by lowercase boldface letters (e.g.,  $\mathbf{t}$ ); real and integer scalars by lowercase, unweighted letters (e.g.,  $r$ ); and constants by uppercase, unweighted letters (e.g.,  $F$ ).

## 2. Methods

SiSSR takes as its input a sequence of contrast-enhanced CCT volumes,  $\mathcal{Q}$ , and produces as its output a sequence of registered mesh models,  $\mathcal{T}$ , from which functional parameters may be calculated. Broadly, the pipeline involves selecting boundary candidates,  $\mathcal{B}$ ,

corresponding to the endocardium of the LV (Section 2.1), generating a single anatomical reference mesh,  $\mathcal{S}$ , from the boundary candidate data at end diastole (Section 2.2), registering  $\mathcal{S}$  to the boundary candidates across all frames simultaneously (Section 2.3), and calculating functional parameters (Section 2.4).

For each of these quantities, a particular cardiac phase is indicated by the subscript  $f$ ; that is,  $\mathcal{Q}$  indicates the 3-D+time image sequence, whereas  $\mathcal{Q}_f$  indicates the particular volume at frame  $f$ .

## 2.1. Boundary candidate selection

The input volume sequence  $\mathcal{Q}$  is binarized by thresholding at 200 Hounsfield units (HU), resulting in a foreground segmentation corresponding to the intravenous contrast (Fig. 2); in one case where contrast was low, it was necessary to use a threshold of 125 HU. The blood pool of the LV is then separated from other structures using standard morphological techniques (binary opening/closing, connected component selection, and hole filling), and the surface of the binary image (the “boundary candidate mesh,”  $\mathcal{B}$ ) is extracted using marching cubes (Lorensen and Cline, 1987). Where necessary, the right heart is pruned from the LV before selecting the largest connected component. In order to aid comparison with the CPD method, the LV is segmented from the left atrium and LV outflow tract, though this is not a requirement of our pipeline. The result is a sequence of dense, triangular boundary candidate meshes,  $\mathcal{B}$ , representing the endocardium of the LV for each volume in the sequence.

## 2.2. Mesh model generation

$\mathcal{B}$  cannot be used directly to quantify cardiac function, because these meshes are not parametrically addressable. Therefore, it is necessary to generate a single anatomical reference mesh,  $\mathcal{S}$ , which can be registered to  $\mathcal{B}$  across all frames. We choose the boundary candidate mesh at end diastole,  $\mathcal{B}_{ED}$ , as the starting point for generating  $\mathcal{S}$  (Fig. 3).  $\mathcal{B}_{ED}$  is decimated to approximately 100 faces by iterative edge contraction, positioning new vertices so as to minimize quadric error relative to the prior iteration (Garland and Heckbert, 1997). Quadric error is preferred over a simpler metric, such as squared edge length, due to the tendency of the former to preserve regions of high curvature (in our application, we found that this was helpful in preserving, in particular, the cardiac apex and atrial appendage).

Our ability to generate  $\mathcal{S}$  directly from  $\mathcal{B}_{ED}$  is possible because the boundary candidates are stored as a mesh, as opposed to a point set as in Stebbing (2014). This advantage justifies our decision to use marching cubes in the previous stage (Section 2.1). The decimation step is necessary because it dramatically reduces the number of vertices in  $\mathcal{S}$  compared to  $\mathcal{B}_{ED}$  (a two order of magnitude reduction in our application), each of which contributes three parameters to the optimization (Section 2.3). Note that by generating  $\mathcal{S}$  directly from  $\mathcal{B}_{ED}$ , we avoid the need to provide a bespoke template mesh, as required elsewhere (Stebbing, 2014; Pourmorteza et al., 2012).

As a practical note, it is common for  $\mathcal{B}_{ED}$  to contain one or more regions which are perfectly flat, either due to manual pruning or to the region of interest extending beyond the bounds of the image. We found that when decimating this mesh directly, the face count could be

reduced by contracting edges in these flat regions with zero quadric error. The resulting mesh tended to score poorly in metrics of mesh quality, specifically the condition number of the triangle Jacobian matrix (a measure of the deviation of a triangle face from equilateral, where 1.0 indicates equilateral and greater magnitude indicates greater distortion).<sup>1</sup> We were able to largely circumvent these negative effects by perturbing the vertices of  $\mathcal{B}_{ED}$  with a small amount of Gaussian noise prior to decimation (Fig. 4). We refer the reader to Vigneault (2016) for a more thorough discussion.

## 2.3. Registration

Having generated the anatomical reference mesh,  $\mathcal{S}$ , and the boundary candidate meshes,  $\mathcal{B}$ , we now describe how the former may be registered to the latter across all frames simultaneously, yielding a sequence of registered meshes,  $\mathcal{T}$ . The parameters of the registration are the Cartesian coordinates of the vertices of  $\mathcal{S}$ . The primary residuals are the distances between points sampled *from the underlying surface* and the nearest boundary candidates. Additional regularizers encourage physically plausible motion and topology, and high mesh quality. Here we present background on Loop subdivision surfaces, the details for evaluating an arbitrary position on a Loop subdivision surface, formulation of the cost function for the optimization, and our multiresolution approach.

**2.3.1. Loop subdivision surfaces**—Loop subdivision is an approximating (as opposed to interpolating) procedure in which a coarse triangle mesh is refined (either explicitly or implicitly) to a smooth underlying surface. Loop subdivision generalizes quadratic triangle B-splines to triangle meshes with “extraordinary” ( $v \neq 6$ ) valency, where quadratic triangle B-splines are themselves one possible generalization of quadratic B-splines to a 2-D manifold in  $\mathbb{R}^3$ . Loop subdivision surfaces have properties which are highly desirable for cardiac mesh modeling.

From an anatomical standpoint:

- Iterative subdivision converges to a limit surface with  $C^1$  continuity, which is useful when representing an anatomical structure.
- Membrane and thin plate energies, which are physically-motivated measures of surface curvature, may be closely approximated for a given patch and used as regularizers to discourage excessive bending (Cashman and Fitzgibbon, 2013).
- Subdivision surfaces are “parametrically addressable;” that is, a particular point may be uniquely addressed by its patch index and parametric patch coordinates. This point exists and remains unique for all meshes in the sequence. This is in contrast to, e.g., level sets, where there is no uniform way of addressing corresponding points on two successive level set contours.

From a computational standpoint:

- Exact evaluation of a parametric position on the model surface requires at most one explicit subdivision (Stam, 1998).

<sup>1</sup>The implementation used here can be found at <https://www.vtk.org/doc/nightly/html/classvtkMeshQuality.html>



- Surface positions and measures of surface curvature depend only on the one-ring surrounding a patch,<sup>2</sup> meaning that the residuals of the optimization depend only on a small number of parameters. This sparsity can be exploited during the optimization in order to increase both computational and memory efficiency (Agarwal and Mierle, 2010).
- The models may be optimized jointly across all cardiac phases, circumventing the propagation of error inherent to “frame-to-frame” algorithms.
- Subdivision surfaces suggest a multiresolution approach, whereby the output of one registration pass is subdivided and used as the initialization to a subsequent registration pass.

**2.3.2. Subdivision surface evaluation**—In order to calculate the primary residual of the cost function, it is necessary to evaluate an arbitrary position on the Loop subdivision surface. The details of Loop subdivision surface evaluation are provided elsewhere (Stam, 1998; Stebbing, 2014); in order to aid in understanding the cost function formulation (Section 2.3.3), we briefly review the mathematics here.

We begin by defining our terminology. A “patch” is the subset of the limit subdivision surface which arises from a particular face in the initial, coarse mesh. A point  $\mathbf{u} \in \mathbb{R}^3$  on a subdivision surface is uniquely described by its patch index  $i$ , and its parametric coordinates  $\mathbf{t} = (r, s) \in \mathbb{R}^2 = \{(r, s) : r, s \in [0, 1], r + s \leq 1\}$ . A patch is termed “ordinary” if the three vertices defining it each have valency  $v = 6$ . A patch is termed “extraordinary” if exactly one vertex has valency  $v \neq 6$ . Patches with more than one extraordinary vertex are disallowed.<sup>3</sup> This distinction is useful because ordinary and extraordinary patches require distinct procedures for evaluating surface positions. The “one-ring” of a patch is the set of the three vertices defining the central triangle, plus those vertices connected by exactly one edge. Any surface position in a patch may be evaluated given its one-ring.

Let  $\mathbf{X} \in \mathbb{R}^{3 \times (C \times F)}$  be the matrix of Cartesian coordinates comprising the subdivision surface control points across all frames, where  $C$  is the number of control points in a single mesh and  $F$  is the number of frames.  $\mathbf{X}_f \in \mathbb{R}^{3 \times C}$  is the submatrix of  $\mathbf{X}$  containing the control points in frame  $f$ , and  $\mathbf{X}_{f,i} \in \mathbb{R}^{3 \times (v+6)}$  is the submatrix of  $\mathbf{X}_f$  containing the control points comprising the one-ring of patch  $i$ . We define a function  $\chi : \mathbf{t}, \mathbf{X}_{f,i} \rightarrow \mathbb{R}^3$  which gives the position of a surface point  $\mathbf{u}_{f,i,t}$  given the parametric coordinates and the one-ring matrix for a given frame and patch. For an ordinary patch,  $\chi$  can be evaluated directly as a barycentric combination of the 12 relevant control points, where  $\mathbf{b}(\mathbf{t}) \in \mathbb{R}^{12}$  is a vector of quartic polynomials (Stam, 1998).

$$\chi(\mathbf{t}, \mathbf{X}_{f,i}) = \mathbf{X}_{f,i} \cdot \mathbf{b}(\mathbf{t}) = \mathbf{u}_{f,i,t} \quad (1)$$

<sup>2</sup>In this context, we use the “one-ring” of a mesh face to mean to the collection of vertices which are at most one edge away from the vertices comprising the face itself.

<sup>3</sup>The requirement that each patch has at most one extraordinary vertex does not impose any practical limitation, as it can be satisfied by subdividing the input mesh exactly once prior to surface point evaluation.

In the case of a patch with a single extraordinary vertex, a simple set of barycentric weights does not exist. However, geometric subdivision of an extraordinary patch results in four “child” patches: three ordinary patches and one extraordinary patch with topology identical to the parent (Fig. 5). This process may be repeated iteratively. Therefore, excluding the origin  $\mathbf{t} = (0, 0)$  (which by convention is placed at the extraordinary vertex), the position at any local patch coordinate may be evaluated with (1) by subdividing a sufficient number of times, extracting the vertices of the relevant one-ring, and transforming the local patch coordinate into the subdivided patch. In Stam (1998), this concept is formalized in terms of matrix algebra. In particular, the “extended subdivision matrix”  $\mathbf{A} \in \mathbb{R}^{(v+6) \times (v+6)}$  is described, which subdivides the patch around the extraordinary vertex, yielding a patch of identical topology to the first; the “bigger subdivision matrix”  $\mathbf{B} \in \mathbb{R}^{(v+12) \times (v+6)}$ , which additionally generates the three ordinary child patches; and the “picking matrix”  $\mathbf{P}_k \in \mathbb{R}^{12 \times (v+12)}$ , which selects the one-ring around the  $k^{\text{th}}$  ordinary child patch ( $k \in \{0, 1, 2\}$ ). For the precise definitions of these matrices, we refer the reader to Stam (1998).

For a given  $\mathbf{t}$ ,  $n$  geometric subdivisions are required before the surface point of interest is contained within one of the three ordinary child patches:

$$n = \lfloor -\log_2(r + s) + 1 \rfloor.$$

The index of the ordinary child patch  $k$  is given by:

$$k = \begin{cases} 0: 2^{-n} < r \leq 2^{-n+1}, 0 \leq s \leq 2^{-n+1} - r \\ 1: 0 \leq r \leq 2^{-n}, 2^{-n} - r \leq s \leq 2^{-n} \\ 2: 0 \leq r \leq 2^{-n}, 2^{-n} < s \leq 2^{-n+1} - r, \end{cases}$$

and the patch coordinate  $\mathbf{t}$  may be transformed into the appropriate child patch given  $n$  and  $k$ :

$$\hat{\mathbf{t}} = \begin{cases} k = 0: (2^n r - 1, 2^n s) \\ k = 1: (1 - 2^n r, 1 - 2^n s) \\ k = 2: (2^n r, 2^n s - 1). \end{cases}$$

Given these definitions, the surface position for patch  $i$  and frame  $f$  at  $\mathbf{t} = \mathbf{0}$  is given by<sup>4</sup>:

$$\chi(\mathbf{t}, X_{f,i}) = X_{f,i} (\mathbf{P}_k \bar{\mathbf{A}} \mathbf{A}^{n-1})^T \mathbf{b}(\hat{\mathbf{t}}).$$

At the patch origin ( $\mathbf{t} = \mathbf{0}$ ), this Equation cannot be used as the power of  $\mathbf{A}$  approaches infinity. However, as shown in Stebbing (2014), the surface position at the origin can be

<sup>4</sup>Note that the powers of  $\mathbf{A}$  can be calculated efficiently through eigen-decomposition ( $\mathbf{A}^n = \mathbf{V}_\mathbf{A} \mathbf{\Lambda}_\mathbf{A}^n \mathbf{V}_\mathbf{A}^{-1}$ , where  $\mathbf{V}_\mathbf{A}$  is a matrix of eigenvectors and  $\mathbf{\Lambda}_\mathbf{A}$  is a diagonal matrix of eigenvalues).



evaluated by  $\chi(\mathbf{0}, \mathbf{X}) = \mathbf{X} \tilde{\mathbf{v}}_1$ , where  $\tilde{\mathbf{v}}_1$  is the (normalized) left subdominant eigenvector of  $\mathbf{A}$ .

**2.3.3. Registration metrics and optimization**—Levenberg-Marquardt least squares optimization (Marquardt, 1963) is used to register the anatomical reference mesh,  $\mathcal{S}$ , to the boundary candidates,  $\mathcal{B}$ . The parameters  $\mathbf{X} \in \mathbb{R}^{3 \times (C \times F)}$  of the optimization are Cartesian displacements to the vertices of all model surface meshes. A fixed number of surface points  $\mathbf{u}_{f,i,t}$  were sampled across the model surfaces (three per patch). At each step in the optimization, for each of these points, the nearest boundary candidate  $\phi(\mathbf{u}_{f,i,t})$  was calculated (that is, the nearest vertex in  $\mathcal{B}_f$ ), where  $\phi: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . This was computed efficiently by representing the vertices defining  $\mathcal{B}$  at each frame as a  $K_d$  tree. The Cartesian components of the distance between the surface point and nearest boundary candidate were the residuals of  $E_{cf}$ , the first term of the cost function,

$$E_{cf} = \sum_{f,i,t} \|\mathbf{u}_{f,i,t} - \phi(\mathbf{u}_{f,i,t})\|^2. \quad (2)$$

Additionally, three regularizers were included to enforce physical constraints of anatomical deformation and ensure high mesh quality: control point acceleration, subdivision patch thin-plate energy, and triangle aspect ratio. This differs from the cost function formulation by Stebbing (2014), where the regularizing terms were control point velocity and thin-plate energy.

In our cost function, control point acceleration is the only means by which information is shared between frames (that is, the only residual which is dependent upon parameters from more than one frame). At a minimum, regularizing against control point velocity as in Stebbing (2014) is necessary to maintain material consistency (the assumption that, for fixed  $i$  and  $t$ ,  $\mathbf{u}_{f,i,t}$  corresponds to the same material point  $\forall f$ ). By regularizing against acceleration rather than velocity, our method additionally encourages smooth, biologically plausible motion. The control point acceleration regularizer  $E_{ac}$  was defined as the Cartesian components of the second differences between corresponding vertices  $\mathbf{x}_{f,c}$  in three adjacent frames, where  $\mathbf{x}_{f,c}$  is the  $(f \times C) + c^{\text{th}}$  column of  $\mathbf{X}$ .

$$E_{ac} = \sum_{f,c} \left\| \begin{bmatrix} 1 & -2 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x}_{(f+2) \bmod F, c}^\top \\ \mathbf{x}_{(f+1) \bmod F, c}^\top \\ \mathbf{x}_{f, c}^\top \end{bmatrix} \right\|^2 \quad (3)$$

Thin-plate energy for patch  $i$  in frame  $f$  is the surface integral over the sum of the partial second derivatives with respect to  $\rho = r + \frac{1}{2}s$  and  $\sigma = \frac{\sqrt{3}}{2}s$  (the transformation yields an orthogonal coordinate system, which is a necessary condition for isotropic thin-plate energy) (Cashman and Fitzgibbon, 2013).

$$E_{tp} = \sum_{f,i} \int_{t \in \Delta} \left\| \frac{\partial^2 \chi}{\partial \rho^2} \right\|^2 + 2 \left\| \frac{\partial^2 \chi}{\partial \rho \partial \sigma} \right\|^2 + \left\| \frac{\partial^2 \chi}{\partial \sigma^2} \right\|^2 dt$$

Unfortunately, this integral is not straightforward to evaluate for an extraordinary patch. However, it is possible to define a matrix  $\mathbf{B} \in \mathbb{R}^{15 \times (v+6)}$  in terms of the position function  $\chi$  and its first derivatives, evaluated at a discrete number of points, which replaces of a Loop subdivision patch  $\mathbf{X}_{f,i}$  of any valence with an approximate quartic triangle Bézier patch  $\hat{\mathbf{X}}_{f,i}$  (Stebbing, 2014). For an ordinary Loop subdivision patch, the Bézier representation is equivalent, and for an extraordinary patch, the error is negligible in practice (the error may become relevant for extremely high valence patches, which were not encountered in this work).

$$\hat{\mathbf{X}}_{f,i} = \mathbf{B} \mathbf{X}_{f,i}$$

The surface position on the Bézier patch may then be evaluated as a barycentric combination of the 15 relevant control points, where  $\hat{\mathbf{b}}(t)$  is a vector of quartic polynomials (Stam, 1998).

$$\hat{\chi}(t, \hat{\mathbf{X}}_{f,i}) = \hat{\mathbf{X}}_{f,i} \times \hat{\mathbf{b}}(t)$$

The thin plate energy of a Bézier patch can be calculated exactly, regardless of the valence of the Loop subdivision patch which it approximates. In our notation, we indicate quantities pertaining to the Bézier approximation with a superposed chevron ( $\hat{\cdot}$ ).

$$E_{\hat{tp}} = \sum_{f,i} \int_{t \in \Delta} \left\| \frac{\partial^2 \hat{\chi}}{\partial \rho^2} \right\|^2 + 2 \left\| \frac{\partial^2 \hat{\chi}}{\partial \rho \partial \sigma} \right\|^2 + \left\| \frac{\partial^2 \hat{\chi}}{\partial \sigma^2} \right\|^2 dt \quad (4)$$

Though thin plate energy is useful for limiting excessive bending, it is only sensitive to deformations which affect the second derivative of the surface positions with respect to the parametric coordinates, and does not penalize affine transformations. Therefore, thin plate energy regularization allows long, thin patches to “bunch” along a ridge, resulting in biologically implausible compression. In order to counteract this, the mesh triangle aspect ratio was penalized.

$$E_{ar} = \sum_{f,i} \frac{\|\Delta_{\max}\|}{\|\Delta_{\min}\|} \quad (5)$$

Here  $\Delta_{\max}$  and  $\Delta_{\min}$  denote the longest and shortest edges of the innermost triangle in patch  $i$ , respectively.

The overall optimization problem may then be written in terms of Equations 2, 3, 4, and 5 and corresponding scaling factors.

$$E = \min (\alpha_{cf} E_{cf} + \alpha_{ac} E_{ac} + \alpha_{\hat{tp}} E_{\hat{tp}} + \alpha_{ar} E_{ar})$$

Recall that the sum for  $E_{cf}$  is taken over all frames, patches, and surface points;  $E_{ac}$  over all frames and control points;  $E_{\hat{tp}}$  over all frames and patches; and  $E_{ar}$  over all frames and patches.

Two points relating to computational efficiency are worth noting. First, the Jacobians of Equations 2, 3, 4, and 5 can all be calculated analytically. By providing explicit Jacobians, we avoid the need for numeric derivatives, which would slow computation precipitously. For a given set of correspondences between surface positions and boundary candidates, the Jacobians of Equations 2, 3, and 4 are linear with respect to the Cartesian displacements of the control points and therefore trivial to calculate. Equation 5 is the ratio of two Euclidean distances. Its Jacobian, therefore, can be calculated with the quotient rule and the knowledge that the partial derivative of Euclidean distance  $\|w\|$  with respect to a particular component  $w_d$  is  $w_d/\|w\|$ , while noting that whether a particular control point contributes to the numerator, denominator, or neither must be determined with each iteration. Second, each individual residual depends upon a very small number of parameters. Specifically, the individual terms comprising  $E_{cf}$  and  $E_{\hat{tp}}$  depend on  $3 \times (v + 6)$  parameters, and those comprising  $E_{ac}$  and  $E_{ar}$  depend on  $3 \times 3$  parameters. This sparsity is exploited during the optimization to limit the number of components of the Jacobian which must be evaluated, further reducing computational cost.

**2.3.4. Multiresolution approach**—Multiresolution techniques frequently have the effect of speeding convergence and avoiding local minima; we expect both of these benefits to apply in our application. Recalling Sect 2.2, the anatomical reference mesh  $\mathcal{S}$  was decimated to  $\sim 100$  faces. This target was chosen empirically as the minimum number of faces beyond which large discrepancies were observed between  $\mathcal{B}_{ED}$  and  $\mathcal{S}$ . This model is subdivided once prior to registration to ensure that no triangle has more than one extraordinary vertex; therefore, the input to the registration algorithm consists of  $\sim 100 \times 4$  faces. The output of the first pass ( $\mathcal{T}^1$ ) is subdivided and used to initialize a second pass, such that the final mesh model ( $\mathcal{T}^2$ ) contains  $\sim 100 \times 4^2$  faces. In principle, the user could perform as many passes as necessary. In practice, we find that two passes are sufficient to reflect the endocardial geometry with subpixel error; that is, for the mean Euclidean distance between all points sampled from the registered subdivision surface and their nearest boundary candidate is less than the resampled voxel size of 0.5mm.

## 2.4. Cardiac function

Following registration, the end diastolic mesh is subdivided until visually smooth (we find that 5 subdivisions are adequate) and the subdivided mesh is segmented according to the 16-segment American Heart Association (AHA) model based on standard clinical planes (Cerqueira et al., 2002). The patches in all other frames are assigned to an AHA segment based on the assignment of the corresponding patch in the end diastolic mesh. For each face in the limit subdivision surface of  $\mathcal{T}$ , SQUEEZ is calculated for a given frame  $f$  as a

function of the current and end diastolic surface areas  $a_f$  and  $a_{ED}$ , respectively (Pourmorteza et al., 2012).

$$\text{SQUEEZ}(f) = \sqrt{\frac{a_f}{a_{ED}}} \quad (6)$$

## 2.5. Statistics

Normally distributed variables are reported as mean  $\pm 1$  standard deviation. Non-normally distributed variables are reported as median [interquartile range, IQR]. Normally and non-normally distributed variables were compared using the paired Student's  $t$ -test and Wilcoxon signed-rank test, respectively. Linear relationships are assessed with Pearson's correlation, where  $0.7 < |R| \leq 1.0$  is interpreted as a strong correlation,  $0.5 < |R| \leq 0.7$  a moderate correlation,  $0.3 < |R| \leq 0.5$  a weak correlation, and  $|R| \leq 0.3$  no correlation.

## 3. Implementation details

Evaluation of Loop subdivision surfaces at arbitrary parameter values was implemented according to Stam (1998) in the C++ programming language. The core classes have been packaged as an external module for the Insight Toolkit (ITK) library (Johnson et al., 2016). Visualization and user interaction was implemented using the Visualization Toolkit (VTK) library (Schroeder et al., 2006). Levenberg Marquardt least squares optimization was performed using the Ceres Solver (Agarwal and Mierle, 2010), configured with OpenBLAS (Wang et al., 2013). Experiments were run on an Ubuntu 16.04 laptop computer with eight Intel Core i7-6700HQ 2.60GHz quad-core processors and 16GB RAM.

## 4. Experiments

We applied our algorithm to CCT images obtained from a canine infarct model, and compared our results to those obtained for the CPD algorithm. The details of the animal model, imaging parameters, and validation are as follows.

### 4.1. Canine infarct model

The study design was reviewed and approved by the Animal Care and Use Committee at the National Institutes of Health. Thirteen adult male mongrel canines underwent surgical ligation of the left anterior descending (LAD) coronary artery for 90 minutes to simulate myocardial infarction, followed by reperfusion. Canines were scanned between six and eight weeks following surgery.

### 4.2. Image acquisition

Images were acquired using the SOMATOM Force scanner (Siemens, Erlangen, Germany). For the cardiac functional analysis test scan, a 30mL bolus of 20% iopamidol 370 (I-370) contrast was used at 2.5mL/s, followed by a 10mL saline chaser. For the final scan, a 30mL bolus of 80% I-370 contrast was used at 2.5mL/s, followed by a 10mL saline chaser. The test scan was used to determine optimal timing; the SiSSR analysis was performed using the

final scan. Images were reconstructed over a single heart beat. The heart rate was controlled to 100bpm (normal for medium-sized dogs) with a  $\beta$ -blocking agent to the extent permitted by the protocol. The gantry rotation time was 250ms; the tube current was 1152 – 1688mA; the tube voltage was 90kV; and the collimation was 96 rows (full detector). The matrix size was  $512 \times 512$  (slice thickness: 0.5mm; pixel size:  $\sim 0.3\text{mm} \times \sim 0.3\text{mm}$ ). Twenty cardiac phases were reconstructed over the R-R interval using the Br36d convolution kernel. The images were resampled to be isotropic (0.5mm in all dimensions) prior to SiSSR analysis. The LV bloodpool, assessed in a  $5 \times 5$  region of interest in the center of the LV, was found to have a signal of  $426.7 \pm 16.7$  HU. The LV myocardium, assessed in a  $5 \times 5$  region of interest in the midwall of the septum, was found to have a signal of  $77.2 \pm 11.1$  HU. The contrast-to-noise ratio between these was  $20.9 = (426.7 - 77.2)/11.1$ .

### 4.3. Validation

The original SQUEEZ implementation (Pourmorteza et al., 2015) relies upon coherent point drift (CPD). The dataset was analyzed with this implementation and the results were compared with those from the subdivision algorithm proposed here. Hyperparameters used in the CPD algorithm were chosen to match those used previously in the literature, and are given in Table 1. The relationship between whole-heart and segmental peak SQUEEZ values measured by the two methods was assessed quantitatively by paired Student's t-test and Pearson's correlation, and visually by Bland-Altman plots and regression.

## 5. Results

### 5.1. Parameter selection

In order to determine optimal hyperparameters, the weight  $\alpha_{cf}$  associated with the primary cost function  $E_{cf}$  was initialized to unity, and a sequential parameter sweep was performed for the remaining residuals across three orders of magnitude (0.01, 0.1, 1.0, and 10.0). For each run, segmentation accuracy and mesh quality were assessed for all frames and canines (Fig. 6). Jaccard Index (the ratio between the intersection and union of the ground truth and predicted segmentations, where 0.0 indicates disjoint sets and 1.0 indicates perfect overlap) was used as a surrogate for segmentation accuracy, and the condition number of the triangle Jacobian matrix (a measure of the deviation of a triangle face from equilateral, where 1.0 indicates equilateral and greater magnitude indicates greater distortion) was used as a surrogate for mesh quality.

First, a range of values for  $\alpha_{tp}$  were attempted (Fig. 6a). As  $\alpha_{tp}$  increased, mesh quality improved considerably, both in terms of median condition number and interquartile range. However, segmentation performance deteriorated above  $\alpha_{tp} = 0.1$ ; for this reason,  $\alpha_{tp} = 0.1$  was selected.

Second, a range of values for  $\alpha_{ar}$  were attempted (Fig. 6b). As  $\alpha_{ar}$  increased, median condition number improved, and the interquartile range tightened considerably, without deterioration in segmentation accuracy;  $\alpha_{ar} = 10.0$  was selected.

Note that one may indirectly penalize aspect ratio by instead penalizing edge length. This may be understood the smallest perimeter is equilateral. Since edge length is mathematically simpler to define (and, therefore, possibly easier to optimize), edge length was investigated for comparison with aspect ratio (Fig. 6c). Indeed, median condition number does improve somewhat when edge length is weighted 1.0. However, whereas penalizing aspect ratio has little effect on segmentation accuracy, penalizing edge length brings about a precipitous decline, causing the mesh to collapse at 10.0. For this reason, it was concluded that the additional complexity of penalizing aspect ratio was justified.

Finally, a range of values for  $\alpha_c$  were attempted (Fig. 6d). As  $\alpha_{ac}$  increased, segmentation accuracy slightly decreased, but without appreciable deterioration in mesh quality. While this suggests that penalizing acceleration is not detrimental, it does not prove its benefit. To see the benefit, consider the effect of varying  $\alpha_{ac}$  on global SQUEEZ in an individual canines (Fig. 7) with a discontinuity in global SQUEEZ at frame four as measured by CPD. Interestingly, even without regularizing acceleration ( $\alpha_{ac} = 0.0$ ), the discontinuity is less pronounced in the SQUEEZ curves calculated by the SiSSR algorithm. As  $\alpha_{ac}$  increases, the discontinuity is mitigated, but likewise details of the SQUEEZ curve at end systole and late diastole are lost.  $\alpha_{ac} = 0.1$  was ultimately selected as a compromise between enforcing temporal coherence and preserving fine details of the SQUEEZ curves.

Note that, in enforcing temporal consistency, penalizing velocity and acceleration are both considerations. Broadly speaking, the cardiac cycle consists of contraction, relaxation, and diastasis. Acceleration is expected to be high only during the transitions between these periods, whereas velocity is expected to be high during all of contraction and relaxation. That is, the acceleration field is expected to be sparse compared with the velocity field, and is therefore a more appropriate term to penalize. However, we nonetheless investigated velocity (Fig. 6e), and found that the two performed similarly in terms of segmentation accuracy and mesh quality.

The final values of  $\alpha_{cf}$ ,  $\alpha_{ac}$ ,  $\alpha_{ip}$ , and  $\alpha_{ar}$  are reported in Table 2.

## 5.2. Optimization performance

The two optimization passes combined took 4.7 minutes per canine (283.7s, IQR = [246.4s, 298.8s]). The first pass contributed very little time to the total compared with the second pass (22.8s, IQR = [21.7s, 26.6s] vs 259.0s, IQR = [224.0s, 272.1s]).

Qualitatively, the registered mesh models  $\mathcal{T}^2$  (Fig. 8) were of high quality, closely tracked the endocardium, and reflected the observed myocardial contraction throughout the cardiac cycle (Fig. 9). Quantitatively, the meshes were of high segmentation accuracy (Jaccard Index: 0.966, IQR = [0.961, 0.970]) and high mesh quality (triangle Jacobian condition number: 1.03, IQR = [1.03, 1.04]).

In order to assess the benefit of the multiresolution technique, a separate, single-pass registration was performed where the anatomical reference mesh  $\mathcal{S}$  was generated by decimating directly to  $100 \times 4^2 = 1600$  mesh faces, followed by one registration. Compared with the multiresolution approach, this single “dense” pass took 60.3% longer (454.8s, IQR

= [418.8s, 548.3s],  $p < 0.001$ ), while producing similar segmentation accuracy (0.968, IQR = [0.964, 0.972]) and mesh quality (1.04, IQR = [1.03, 1.07]).

In order to assess the benefit of deriving the anatomical reference mesh,  $\mathcal{S}$ , from the boundary candidates at end diastole,  $\mathcal{B}_{ED}$ , compared with instead using a bespoke left ventricular model, the analysis was repeated, replacing  $\mathcal{S}$  with a prolate hemispheroid. The prolate hemispheroid was constructed in Blender 2.78,<sup>5</sup> and designed to have a number of triangles equal to  $\mathcal{S}$  after its final subdivision. The model was initialized by translating, rotating, and scaling to align with  $\mathcal{B}_{ED}$ . Compared with registering  $\mathcal{S}$ , the bespoke model produced lower segmentation accuracy (Jaccard Index: 0.966, IQR = [0.961, 0.970] vs 0.957, IQR = [0.952, 0.965],  $p < 0.0001$ ) and lower mesh quality (triangle Jacobian Condition Number: 1.03, IQR = [1.03, 1.04] vs 1.10, IQR = [1.08, 1.18],  $p < 0.0001$ ).

### 5.3. Sensitivity of SQUEEZ to optimization parameters

The ultimate goal of this pipeline is to calculate SQUEEZ as a surrogate of regional function; therefore, it is necessary to investigate the sensitivity of SQUEEZ to moderate perturbations of the optimized registration weights. As a baseline measurement, peak segmental SQUEEZ was calculated in all canines using the parameters in Table 2. In six additional registrations,  $\alpha_{ip}$ ,  $\alpha_{ar}$ , and  $\alpha_{ac}$  were each doubled, then halved relative to baseline, and peak segmental SQUEEZ was recalculated. Linear regression was then performed between baseline and each of the six runs (Table 3). Correlation was excellent ( $R > 0.99$ ,  $p < 0.0001$ ), with near-zero intercepts and near-unity slopes in all cases, suggesting that the SQUEEZ analysis is robust to perturbations of the registration hyperparameters over a moderate range centered on the optimized values.

### 5.4. Correlation with CPD

Compared with SiSSR, CPD produced lower segmentation accuracy (Jaccard Index: 0.966, IQR = [0.961, 0.970] vs 0.922, IQR = [0.909, 0.930],  $p < 0.0001$ ) and lower mesh quality (triangle Jacobian Condition Number: 1.03, IQR = [1.03, 1.04] vs 1.10, IQR = [1.08, 1.14],  $p < 0.0001$ ). Correlation between global left ventricular SQUEEZ as measured by SiSSR vs CPD was excellent (SiSSR = 0.74 CPD + 0.25,  $R = 0.92$ ,  $p < 0.0001$ ). A positive bias noted visually on Bland-Altman plot (Fig. 10) was statistically significant ( $0.866 \pm 0.025$  vs  $0.827 \pm 0.030$ ,  $p < 0.0001$ ). Inspection of the meshes suggests that error for the CPD algorithm is highest near the papillary muscles, perhaps due to implicit penalization of adjacent expanding and contracting regions due to the coherent motion constraint. SQUEEZ estimated by CPD is lower for the papillary muscles than that measured by SiSSR, which may account for the noted bias in peak global SQUEEZ.

Representative curves as calculated by SiSSR as well as CPD are shown in Fig. 11. We note that the timing of contraction is very consistent between the two methods, reaching end systole at frame 7 and recovering by frame 12 in both cases. Moreover, SQUEEZ curves calculated by the CPD method are slightly noisy, which could be problematic for the calculation of strain rate without smoothing. SQUEEZ curves calculated by SiSSR by

<sup>5</sup><https://blender.org>



contrast are smooth as a result of regularizing against control point acceleration, and could be used to calculate strain rate directly.

Infarcts were localized using contrast-enhanced CCT, and found to be small (median 7.1g myocardial mass), subendocardial, and localized to the anterior and septal apex, consistent with LAD occlusion. In order to assess whether regional SiSSR is correlated with CPD in both infarcted and remote myocardium, segmental SQUEEZ was compared in the anterior apex and anteroseptal base (AHA segments 13 and 2, respectively, Cerqueira et al. (2002)).

Correlation between SQUEEZ in the anteroseptal base as measured by SiSSR vs CPD was good ( $\text{SiSSR} = 1.04 \text{ CPD} + -0.08$ ,  $R = 0.75$ ,  $p < 0.01$ ), whereas correlation between SQUEEZ in the anterior apex as measured by SiSSR vs CPD was moderate ( $\text{SiSSR} = 0.52 \text{ CPD} + 0.41$ ,  $R = 0.53$ ,  $p = NS$ , Fig. 12).

## 6. Discussion

In this article, we have described an end-to-end pipeline for quantifying cardiac function from a 3-D+time CCT image sequence. Our method, which we term Simultaneous Subdivision Surface Registration (SiSSR), inherits a number of desirable features from the prior state-of-the-art method (Stebbing, 2014) that we have adapted and extended. Loop subdivision surfaces are useful structures for modeling cardiac anatomy. First, they allow smooth, detailed structures to be represented implicitly with a coarse mesh. Second, they impose minimal modeling requirements on the mesh (only that it be triangular, and avoid pathological topologies such as self-intersection). Third, each control point in the coarse mesh influences only a small portion of the underlying surface, leading to sparse dependence between parameters and residuals during the optimization and consequent computational gains. Besides making use of subdivision surfaces, the cost function formulation performs the optimization across all frames simultaneously, which has the effect of avoiding frame-to-frame errors which typically plague heart tracking algorithms.

Our contributions include a number of modifications which improve upon the method proposed by Stebbing (2014). By generating the anatomical reference mesh directly from the blood pool of the left ventricle, we account for broad variability in cardiac anatomy, provide an excellent initialization for mesh registration, and minimize the required human intervention. Additionally, we make two notable changes to the cost function. First, we regularize against acceleration (as opposed to velocity), encouraging smooth, biologically plausible motion. Second, we explicitly encourage high mesh quality by penalizing mesh faces with high aspect ratio, which we note has the added benefit of discouraging physically unreasonable SQUEEZ values. Our work also employs a multiresolution approach for mesh registration, which has the dual advantages of avoiding local minima and increasing computational efficiency. Finally, Stebbing (2014) stopped short of calculating metrics of regional function. We hypothesize that their template mesh, defined by 40 faces initially and  $40 \times 4$  faces after subdivision, would have been insufficiently detailed to adequately capture regional function. The final registered meshes calculated in the SiSSR method, by contrast, contained  $\sim 100 \times 4^2 = 1600$  faces; approximately 75% of these correspond to the LV endocardium (the remaining faces define the mitral valve plane), such that  $1600 \times 0.75/16 =$

75 faces contribute on average to each AHA segment. The spatial resolution of our method is therefore sufficient to characterize clinical motion abnormalities. Moreover, we have shown that SQUEEZ calculated by our pipeline correlates globally with SQUEEZ calculated with an independent methodology (Pourmorteza et al., 2012), giving us further confidence in our results beyond visual inspection, as well as segmentally in both infarcted and remote myocardium.

### 6.1. Alternate methods

Functional CCT studies have historically had limited applicability due to high radiation dose and low temporal resolution, and as such quantifying cardiac function from CCT has received relatively less attention compared with CMR and echocardiography. Recent advances in CCT acquisition and reconstruction have largely eliminated these concerns, sparking renewed interest in the literature.

The most straightforward approach sought to directly apply a block-matching technique originally developed for 2-D ultrasound to stacks of 2-D+time CCT images, resliced from 3-D+time volumes (Tee et al., 2015). This study used a porcine infarct model, with the ultimate goal of calculating circumferential and radial strain. Segmentally, correlation with HARP was moderate in the radial direction ( $R = 0.55$ ,  $p = 0.27$ ) and weak in the circumferential direction ( $r = 0.40$ ,  $p = 0.125$ ). Unfortunately, in treating the volume as a 2-D stack, through-plane motion error inherent to 2-D block-matching is unavoidable. The SiSSR method avoids through-plane errors by considering the volume as a whole, but is limited to quantifying function in the circumferential-longitudinal plane; the radial component of cardiac function cannot be quantified.

More recently, Wong et al. (2016) measured radial, circumferential, and principal strains from 3-D+time CCT images in a canine infarct model. This work used deformable image registration to track the left ventricular endocardium and epicardium, deformed a bespoke finite element model of the left ventricle with the resulting displacements, and attempted to detect the infarcted region using image intensity and functional metrics. As the primary focus of their work was on infarct detection, relatively less effort was devoted to motion tracking. They employed a B-spline transform registration method available in ITK (Johnson et al., 2016), using the mutual information similarity metric and L-BFGS optimizer. The authors allude to the fact that, because traditional registration implementations have been designed to warp a single “moving” image into a single “fixed” image, there is no single, optimal way to apply these techniques to an image sequence without foundational modification. The authors compared “frame-to-frame” and “reference-frame” formulations, and found similar accuracy in terms of infarct detection, but no direct comparison of tracking error was possible. In the present work, we avoid this ambiguity by formulating the registration globally across all cardiac frames.

A state-of-the-art technique for quantifying regional cardiac function from CCT was the SQUEEZ method proposed by Pourmorteza et al. (2012), which has been shown to correlate well with HARP Pourmorteza et al. (2015) and which we employed as an independent methodology for validation in the present work. SQUEEZ is defined as the square root of the ratio between current and end diastolic area in the circumferential-longitudinal plane, and

may be interpreted as a surrogate for cardiac contraction ( $< 1$ ) or expansion ( $> 1$ ). At each frame, the left ventricular blood pool was segmented and a boundary candidate mesh representing the endocardium was calculated. A deformable point-set registration algorithm known as Coherent Point Drift (CPD) was used to register a bespoke template mesh to the end diastolic boundary candidate mesh, and again to register the end diastolic template mesh to each of the remaining cardiac frames. This formulation is similar to the reference-frame registration method used by Wong et al. (2016), and therefore subject to the same errors. To address these concerns, Pourmorteza et al. (2012) utilize “shape index” (SI), a scale-independent shape measure (Koenderink and van Doorn, 1992), which they hypothesize will remain constant throughout the cardiac cycle for a given material point:

$$SI = \frac{2}{\pi} \text{atan} \frac{k_1 + k_2}{k_1 - k_2} : k_1 \geq k_2.$$

Here,  $k_1$  and  $k_2$  are the principal curvatures at a point of interest. In order to encourage each point in the end diastolic mesh to warp into the same anatomical position in the remaining frames, a penalty for warping a point in the template mesh into a point with dissimilar SI in the target mesh was incorporated into the CPD registration. We note that penalizing SI discrepancies may be interpreted as a means of *implicitly* sharing information between frames. This contrasts to the present work, where information is *explicitly* shared between frames via a global optimization framework where no restrictions are imposed on which parameters may be used in the calculation of a particular residual. This more flexible formulation allows, for example, the acceleration of control points to be regularized, which would not be possible using the CPD method using the present implementation.

## 6.2. Limitations and future directions

**Segmental validation**—Although there was excellent correspondence between SiSSR and CPD in terms of *global* SQUEEZ, segmental agreement was only moderate. We hypothesize that this may be attributed to the relative lack of trabeculation of the canine endocardium as compared with humans. Therefore, we hypothesize that human studies planned for the future will demonstrate higher segmental correlation. Additionally, although the CPD method has itself been validated against HARP, it would be beneficial in the future to validate our method directly against HARP, and to assess the relative accuracy of the two methods.

Subdivision surfaces have the desirable property of being *parametrically addressable*; i.e., that a unique point on the surface may be addressed with a given patch index and parametric position. In order to estimate functional parameters, this parametric addressability must further be assumed to be *materially consistent*, i.e., assumed to refer to the same tissue point. The degree to which this assumption holds must be investigated in future work.

**Optimization of the SiSSR hyperparameters**—Hyperparameters of the SiSSR registration were chosen so as to optimize mesh quality while minimizing segmentation error. While we recognize that it may have been preferable to optimize peak segmental SQUEEZ values, we regret that a true ground truth was unavailable, and were reluctant to

optimize against those values measured by CPD, in case differences between the SiSSR and CPD measurements represented an improvement of the SiSSR algorithm.

**Papillary muscles**—Papillary muscles are quite prominent in canine hearts compared with humans, such that the endocardium and papillary muscles are flush with one another for most or all of the cardiac cycle. The lack of contrast between papillary muscles and myocardium makes endocardial function difficult or impossible to assess in the lateral midwall (AHA segments 11 and 12). We regard this as a limitation of the underlying imaging data, rather than our method, and advise caution when interpreting these segments. Modification of the SiSSR method (particularly in the mesh generation step) may be required in order to extend this method to humans.

**Computational Efficiency**—In terms of computational efficiency, we focused in this study on exploiting sparsity and maximizing algorithmic efficiency, and spent relatively less time exploring parallelism beyond what was available by configuring Ceres with OpenBLAS and OpenMP. In the future, it may be possible to further improve the algorithm by adapting the implementation to run on a graphics processing unit (GPU).

**Protocol Optimization**—Although our image acquisition procedure was not optimized for right ventricular contrast, we note that in an appropriate dataset our procedure may be extended to other cardiac chambers with little or no modification.

**Alternative Applications**—We further expect the contributions made in this paper to translate to problems outside the heart, such as fetal neurosonography, where a closely related subdivision surface method has been successfully applied (Namburete et al., 2015).

**Boundary Candidate Selection**—Finally, although contrast between the myocardium and the left ventricular blood pool is high (contrast-to-noise ratio = 20.9), it was nonetheless necessary to manually prune the right ventricle in some cases. The authors recognize that the technique used for detecting the boundary candidate mesh is not state-of-the-art from the perspective of LV segmentation; significant effort has been devoted to this task in the literature. However, this approach was nonetheless preferred over other methods described in the literature for the following reasons:

- A significant advantage of computed tomography over other modalities is that it provides a quantitative, reproducible signal (measured in Hounsfield Units) which, when combined with a timed contrast bolus, gives a detailed separation between myocardium and bloodpool.
- Endocardial trabeculation and papillary muscles are generally excluded in LV CCT segmentation algorithms (Zreik et al., 2016),(Ecabert et al., 2008). While this is an advantage in certain contexts, it is detrimental in our application, which relies upon these structures as the primary trackable features. Our approach retains these features with high fidelity.
- Among the most successful techniques described to-date involve atlas registration (Zhuang et al., 2015) and convolutional neural networks (Zreik et al.,

2016). These algorithms require large training datasets, which were unavailable to us.

- Given the striking difference in canine and human anatomy (specifically, the relative lack of endocardial trabeculation in dogs), any LV segmentation algorithm developed in this dataset would likely need to be redesigned for human subjects.

## 7. Summary

In this study, we have presented the Simultaneous Subdivision Surface Registration (SiSSR) algorithm, a pipeline for the quantification of SQUEEZ (a regional metric of cardiac function) from 3-D+time contrast-enhanced cardiac computed tomography sequences. The primary technical contributions of our technique are in the extension and adaptation of a prior technique developed in the context of 3-D echocardiography to cardiac computed tomography. An anatomical reference mesh (stored as a Loop subdivision surface) was generated at end diastole, and fitted to the endocardium simultaneously across all frames using least squares minimization. The optimization was regularized against acceleration of the control points, thin plate energy, and triangle aspect ratio. We demonstrated that our algorithm could produce visually appealing results in a clinically viable timescale, and validated those results against a state-of-the-art described technique for quantifying cardiac function from computed tomography.

## Acknowledgments

This work was funded by the National Institutes of Health Intramural Research Program and the NIH-Oxford Scholars Program. The first author would additionally like to thank Richard Stebbing, DPhil, Christopher Bridge, DPhil, and Ana Namburete, DPhil for their invaluable input and encouragement throughout this work.

## References

- Agarwal, S., Mierle, K. Ceres Solver. 2010. URL: <http://ceres-solver.org>
- Aletras AH, Ding S, Balaban RS, Wen H. DENSE: displacement encoding with stimulated echoes in cardiac functional MRI. *Journal of Magnetic Resonance*. 1999; 137(1):247–52. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2887318&tool=pmcentrez&rendertype=abstract>. DOI: 10.1006/jmre.1998.1676 [PubMed: 10053155]
- Behar JM, Rajani R, Pourmorteza A, Preston R, Razeghi O, Niederer S, Adhya S, Claridge S, Jackson T, Sieniewicz B, Gould J, Carr-White G, Razavi R, McVeigh E, Rinaldi CA. Comprehensive use of cardiac computed tomography to guide left ventricular lead placement in cardiac resynchronization therapy. *Heart Rhythm*. 2017; 14(9):1364–72. URL: <http://dx.doi.org/10.1016/j.hrthm.2017.04.041>. DOI: 10.1016/j.hrthm.2017.04.041 [PubMed: 28479514]
- Cashman TJ, Fitzgibbon AW. What shape are dolphins? Building 3D morphable models from 2D images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2013; 35(1):232–44. DOI: 10.1109/TPAMI.2012.68 [PubMed: 22392707]
- Cerqueira MD, Weissman NJ, Dilsizian V, Jacobs AK, Kaul S, Laskey WK, Pennell JD, Rumberger JA, Ryan T, Verani MS. Standardized Myocardial Segmentation and Nomenclature for Tomographic Imaging of the Heart: A Statement for Healthcare Professionals From the Cardiac Imaging Committee of the Council on Clinical Cardiology of the American Heart Association. *Circulation*. 2002; 105(4):539–42. URL: <http://circ.ahajournals.org/cgi/doi/10.1161/hc0402.102975>. DOI: 10.1161/hc0402.102975 [PubMed: 11815441]

- Ecabert O, Peters J, Schramm H, Lorenz C, Von Berg J, Walker MJ, Vembar M, Olszewski ME, Subramanyan K, Lavi G, Weese J. Automatic model-based segmentation of the heart in CT images. *IEEE Transactions on Medical Imaging*. 2008; 27(9):1189–202. DOI: 10.1109/TMI.2008.918330 [PubMed: 18753041]
- Garland, M., Heckbert, PS. Surface simplification using quadric error metrics; SIGGRAPH. 1997. p. 209–16. URL: <http://portal.acm.org/citation.cfm?doid=258734.258849>
- Helle-Valle TM, Yu WC, Fernandes VR, Rosen BD, Lima JA. Usefulness of radial strain mapping by multidetector computer tomography to quantify regional myocardial function in patients with healed myocardial infarction. *The American Journal of Cardiology*. 2010; 106(4):483–91. URL: <http://www.ncbi.nlm.nih.gov/pubmed/20691305>. DOI: 10.1016/j.amjcard.2010.03.063 [PubMed: 20691305]
- Johnson HJ, Ibanez L, McCormick M. The ITK Software Guide: Book 1 (4). 2016 doi: 1930934-157.
- Koenderink JJ, van Doorn AJ. Surface shape and curvature scales. *Image and Vision Computing*. 1992; 10(8):557–64. DOI: 10.1016/0262-8856(92)90076-F
- Lorensen WE, Cline HE. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*. 1987; 21(4):163–9. URL: <http://portal.acm.org/citation.cfm?doid=37401.37422>. DOI: 10.1145/37402.37422
- Marquardt DW. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society of Industrial and Applied Mathematics*. 1963; 11(2):431–41. arXiv:arXiv:1011.1669v3. DOI: 10.1017/CBO9781107415324.004
- Myronenko A, Song X. Point Set Registration: Coherent Point Drift. *IEEE Trans Pattern on Pattern Analysis and Machine Intelligence*. 2010; 32(12):2262–75. URL: <http://www.ncbi.nlm.nih.gov/pubmed/20975122http://ieeexplore.ieee.org/ielx5/34/5611449/05432191.pdf?tp=&arnumber=5432191&isnumber=5611449>. DOI: 10.1109/TPAMI.2010.46
- Myronenko, A., Song, X., Carreira-Perpinan, MA. Non-rigid point set registration: Coherent Point Drift. *NIPS*; 2007.
- Namburete AI, Stebbing RV, Kemp B, Yaqub M, Papageorgiou AT, Noble JA. Learning-based prediction of gestational age from ultrasound images of the fetal brain. *Medical Image Analysis*. 2015; 21(1):72–86. URL: <http://dx.doi.org/10.1016/j.media.2014.12.006>. DOI: 10.1016/j.media.2014.12.006 [PubMed: 25624045]
- Ogawa K, Hozumi T, Sugioka K, Matsumura Y, Nishiura M, Kanda R, Abe Y, Takemoto Y, Yoshiyama M, Yoshikawa J. Usefulness of automated quantitation of regional left ventricular wall motion by a novel method of two-dimensional echocardiographic tracking. *The American Journal of Cardiology*. 2006; 98(11):1531–7. URL: <http://www.ncbi.nlm.nih.gov/pubmed/17126665>. DOI: 10.1016/j.amjcard.2006.06.060 [PubMed: 17126665]
- Osman NF, Kerwin WS, McVeigh ER, Prince JL. Cardiac Motion Tracking Using CINE Harmonic Phase (HARP) Magnetic Resonance Imaging. *Magnetic Resonance in Medicine*. 1999; 42(6): 1048–60. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2570035&tool=pmcentrez&rendertype=abstract>. [PubMed: 10571926]
- Osman NF, Sampath S, Atalar E, Prince JL. Imaging Longitudinal Cardiac Strain on Short-Axis Images Using Strain-Encoded MRI. *Magnetic Resonance in Medicine*. 2001; 46(2):324–34. URL: <http://www.ncbi.nlm.nih.gov/pubmed/11477637>. [PubMed: 11477637]
- Pourmorteza A, Chen MY, van der Pals J, Arai AE, McVeigh ER. Correlation of CT-based regional cardiac function (SQUEEZ) with myocardial strain calculated from tagged MRI: an experimental study. *International Journal of Cardiovascular Imaging*. 2015; 32(5):1–7. DOI: 10.1007/s10554-015-0831-7
- Pourmorteza, A., Keller, N., Chen, R., Lardo, AC., Halperin, H., Chen, MY., McVeigh, ER. Precision of Regional Wall Motion Estimates from ultra-low-dose cardiac CT using SQUEEZ; *International Journal of Cardiovascular Imaging*. 2018. p. 0in revisioURL: <http://dx.doi.org/10.1007/s10554-018-1332-2>
- Pourmorteza A, Schuleri KH, Herzka Da, Lardo AC, McVeigh ER. A new method for cardiac computed tomography regional function assessment: Stretch quantifier for endocardial engraved zones (SQUEEZ). *Circulation: Cardiovascular Imaging*. 2012; 5(2):243–50. DOI: 10.1161/CIRCIMAGING.111.970061 [PubMed: 22342945]




- Schroeder, W., Martin, K., Lorensen, B. The Visualization Toolkit. 4. Kitware; 2006.
- Stam, J. Evaluation of Loop subdivision surfaces; SIGGRAPH. 1998. p. 1-15. URL: <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Evaluation+of+Loop+Subdivision+Surfaces#6>
- Stebbing, R. DPhil. University of Oxford; 2014. Model-Based Segmentation Methods for Analysis of 2D and 3D Ultrasound Images and Sequences.
- Stebbing RV, Namburete AI, Upton R, Leeson P, Noble JA. Data-driven shape parameterization for segmentation of the right ventricle from 3D+t echocardiography. Medical Image Analysis. 2015; 21:29–39. URL: <http://linkinghub.elsevier.com/retrieve/pii/S1361841514001868>. DOI: 10.1016/j.media.2014.12.002 [PubMed: 25577559]
- Suinesiaputra A, McCulloch AD, Nash MP, Pontre B, Young AA. Cardiac Image Modelling: Breadth and Depth in Heart Disease. Medical Image Analysis. 2016; 33:38–43. URL: <http://dx.doi.org/10.1016/j.media.2016.06.027>. DOI: 10.1016/j.media.2016.06.027 [PubMed: 27349830]
- Tee M, Noble JA, Bluemke DA. Imaging techniques for cardiac strain and deformation: comparison of echocardiography, cardiac magnetic resonance and cardiac computed tomography. Expert Review of Cardiovascular Therapy. 2013; 11(2):221–31. URL: <http://www.ncbi.nlm.nih.gov/pubmed/23405842>. DOI: 10.1586/erc.12.182 [PubMed: 23405842]
- Tee MW, Won S, Raman FS, Yi C, Vigneault DM, Davies-Venn C, Liu S, Lardo AC, Lima JAC, Noble JA, Emter CA, Bluemke DA. Regional Strain Analysis with Multidetector CT in a Swine Cardiomyopathy Model: Relationship to Cardiac MR Tagging and Myocardial Fibrosis. Radiology. 2015; 277(1):88–94. URL: <http://pubs.rsna.org/doi/10.1148/radiol.2015142339>. DOI: 10.1148/radiol.2015142339 [PubMed: 25853636]
- Vigneault DM. Perturbing Mesh Vertices with Additive Gaussian Noise. Insight Journal. 2016; 80(1): 1–8. URL: <http://www.insight-journal.org/browse/publication/981>.
- Wang Q, Zhang X, Zhang Y, Yi Q. AUGEM: Automatically Generate High Performance Dense Linear Algebra Kernels on x86 CPUs. Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. 2013; 25:1. URL: <http://doi.acm.org/10.1145/2503210.2503219>. doi: 10.1145/2503210.2503219
- Wong KCL, Tee M, Chen M, Bluemke DA, Summers RM, Yao J. Regional infarction identification from cardiac CT images: a computer-aided biomechanical approach. International Journal of Computer Assisted Radiology and Surgery. 2016; 11(9):1573–83. URL: <http://link.springer.com/10.1007/s11548-016-1404-5>. DOI: 10.1007/s11548-016-1404-5 [PubMed: 27072840]
- Zhuang X, Bai W, Song J, Zhan S, Qian X, Shi W, Lian Y, Rueckert D. Multiatlas whole heart segmentation of CT data using conditional entropy for atlas ranking and selection. Medical Physics. 2015; 42(7):3822–33. URL: <http://scitation.aip.org/content/aapm/journal/medphys/42/7/10.1118/1.4921366>. DOI: 10.1118/1.4921366 [PubMed: 26133584]
- Zreik, M., Leiner, T., de Vos, BD., van Hamersvelt, RW., Viergever, MA., Isgum, I. Automatic Segmentation of the Left Ventricle in Cardiac CT Angiography Using Convolutional Neural Networks. ISBI; 2016. p. 40-3.







**Highlights**

- A pipeline for quantifying cardiac function from computed tomography is proposed.
- Patient-specific meshes modelling cardiac anatomy are automatically constructed.
- Cardiac mesh models are fitted to the endocardium in all frames simultaneously.
- Model fitting is formulated so as to encourage biologically plausible motion.
- Functional cardiac parameters are validated against the prior state-of-the-art.

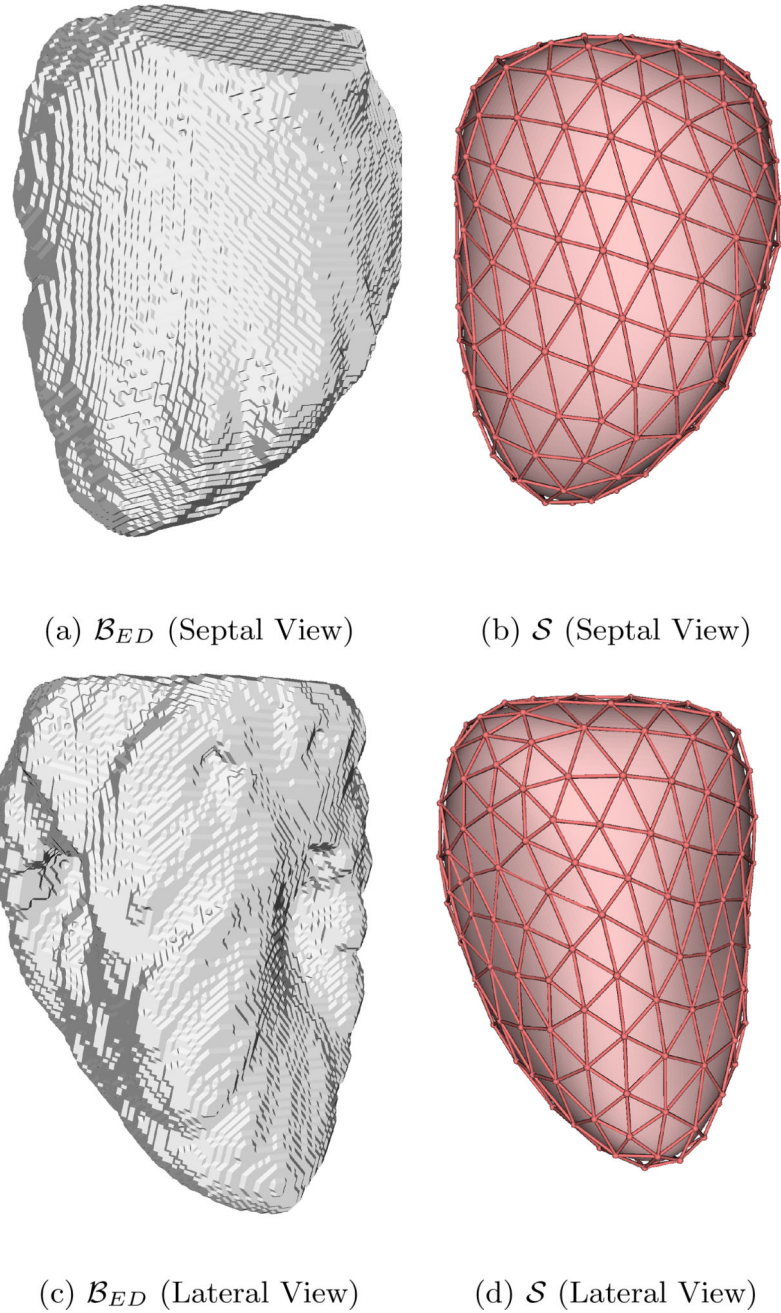


**Figure 1.**

Overview of the analysis pipeline. The volumetric CCT input images are used to generate the boundary candidate meshes. The boundary candidate mesh at end diastole is then decimated to generate the anatomical reference mesh. This mesh is then registered to the boundary candidate meshes across all frames in order to generate a sequence of registered models. The subdivided output of a pass may then serve as the initialization of a subsequent pass. The final registered models may then be used to compute a host of functional cardiac parameters.


(a)  $\mathcal{I}_{ED}$  (Septal)(b)  $\mathcal{I}_{ED}$  and  $\mathcal{B}_{ED}$  (Septal)(c)  $\mathcal{I}_{ED}$  (Lateral)(d)  $\mathcal{I}_{ED}$  and  $\mathcal{B}_{ED}$  (Lateral)**Figure 2.**

Boundary candidate selection. (a) A septal view of the left heart in a CCT volume,  $\mathcal{Q}_{ED}$ ; the bloodpool of the left heart is bright due to contrast enhancement. (b) A septal view of the boundary candidate mesh overlain with the underlying image data, showing coincidence of  $\mathcal{B}_{ED}$  with the endocardium. (c) A lateral view of  $\mathcal{Q}_{ED}$  (the mirror image of (a)). (d) A lateral view of  $\mathcal{Q}_{ED}$  overlain with  $\mathcal{B}_{ED}$ ; note the prominent papillary muscles visible in  $\mathcal{B}_{ED}$ .



**Figure 3.**


Generation of the anatomical reference mesh. (a) A septal view of a dense, triangular boundary candidate mesh,  $\mathcal{B}_{ED}$ , produced by the fast marching cubes algorithm for the end diastolic frame. (b) A septal view of the anatomical reference mesh,  $\mathcal{S}$ , generated through decimation of  $\mathcal{B}_{ED}$ . The dark red wireframe represents the control points and their connections; the smooth, pink surface represents the underlying subdivision surface defined by the wireframe. Note the tendency of the surface to “contract” somewhat during decimation. (c) A lateral view of  $\mathcal{B}_{ED}$ . (d) A lateral view of  $\mathcal{S}$ .



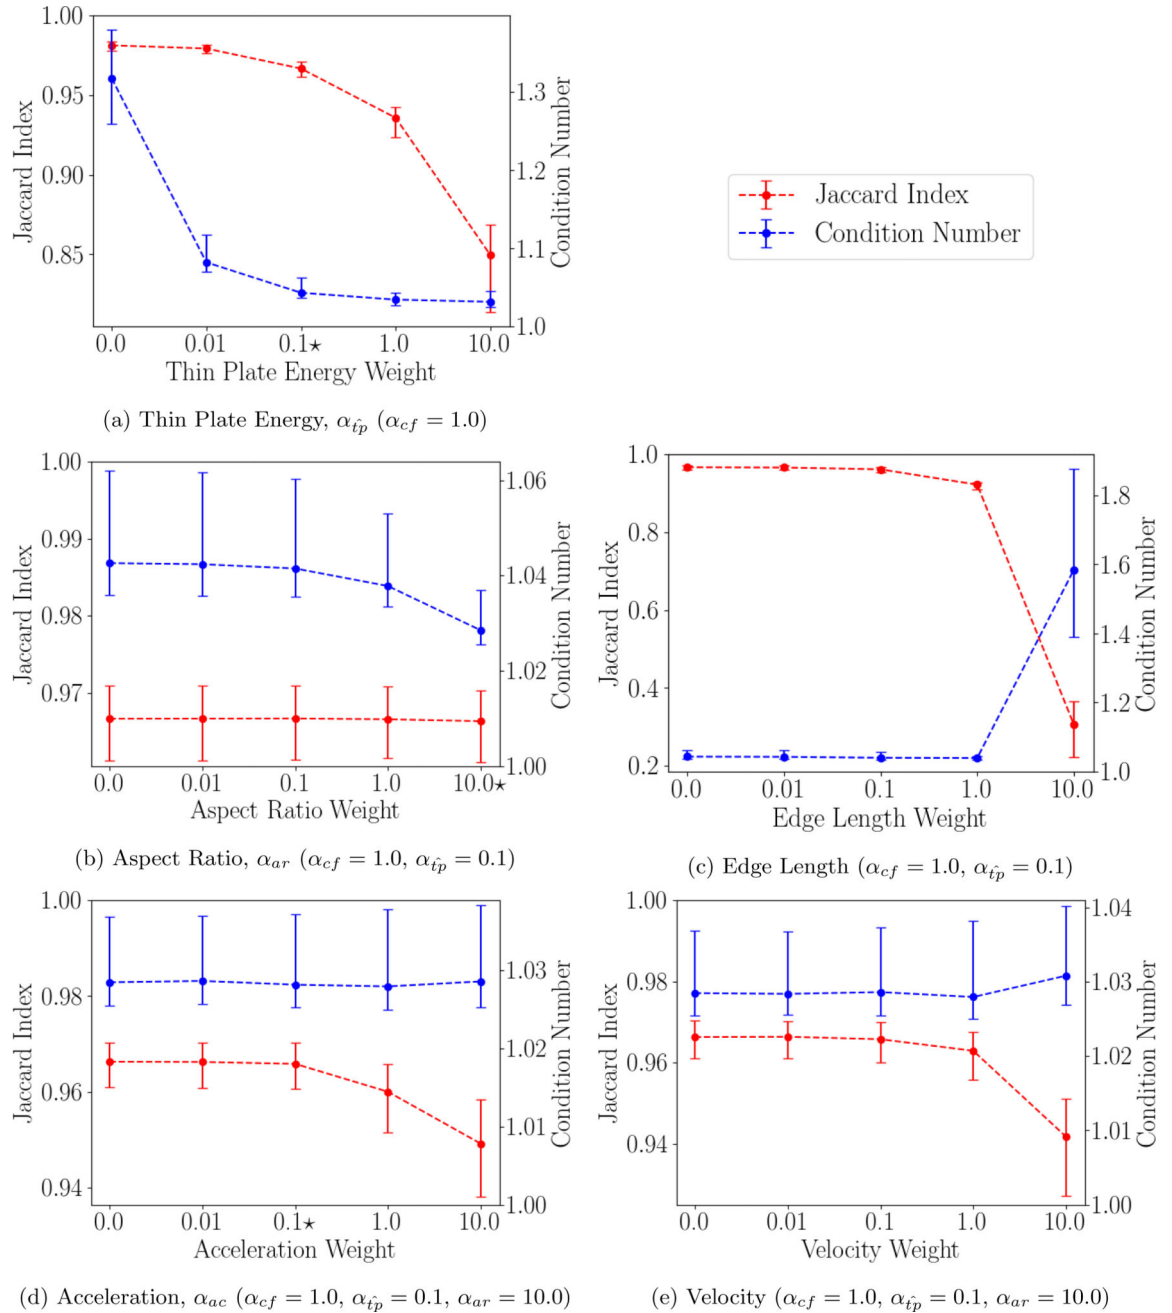
**Figure 4.**

Improvement of mesh quality with Gaussian noise. Marching cubes was used to generate a fine mesh from a binary half-sphere. This fine mesh was decimated both directly (4a) and after perturbing its vertices with a small amount of Gaussian noise (4b). A dramatic improvement in the quality of the decimated mesh is visually apparent. A version of this figure appears in Vigneault (2016).



**Figure 5.**


Geometric demonstration of loop subdivision. Here we have an extraordinary patch where  $v = 7$ , which we would like to evaluate at  $\mathbf{t} = (0.1, 0.2)$ . The number of required subdivisions  $n = 2$  and the index of the child patch  $k = 1$ . The original extraordinary patch is shown in white, and intermediate patches are shown in red. (a) We first subdivide using the “extended subdivision matrix,”  $\mathbf{A}$ . (b) We further subdivide using the “bigger subdivision matrix,”  $\bar{\mathbf{A}}$ . (c) From the result, we extract an ordinary patch containing the parametric point of interest, the surface position of which we may then evaluate directly.

**Figure 6.**

Stepwise selection of least squares optimization scaling factors. Each plot shows how the Jaccard Index (red, left vertical axis) and the condition number of the triangle Jacobian matrix (blue, right vertical axis) vary with the parameter of interest. Points and error bars represent median and interquartile range. Note that, for all but subplot (a), the leftmost datapoint is reproduced from a previous plot for convenience. The weight  $\alpha_{cf}$  associated with the primary residual  $E_{cf}$  was initialized to unity and the remaining residual weights  $\alpha_{\hat{tp}}$  (a),  $\alpha_{ar}$  (b), and  $\alpha_{ac}$  (d) were varied and tested sequentially. An edge length penalty (c) is shown for comparison to aspect ratio, and a velocity penalty (e) is shown for comparison to

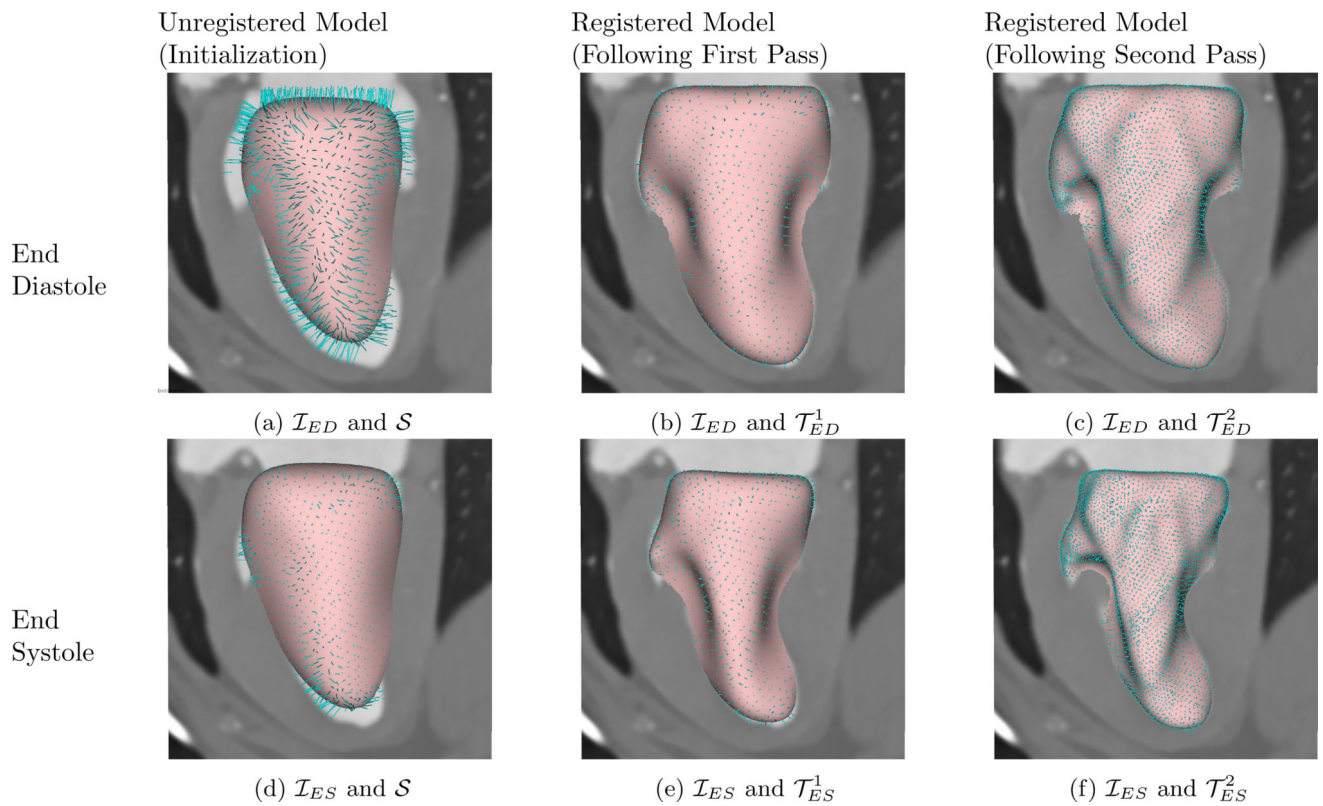


acceleration (these did not contribute to the final cost function). The selected parameter value is indicated on the horizontal axis with a star (\*). See text for further discussion.



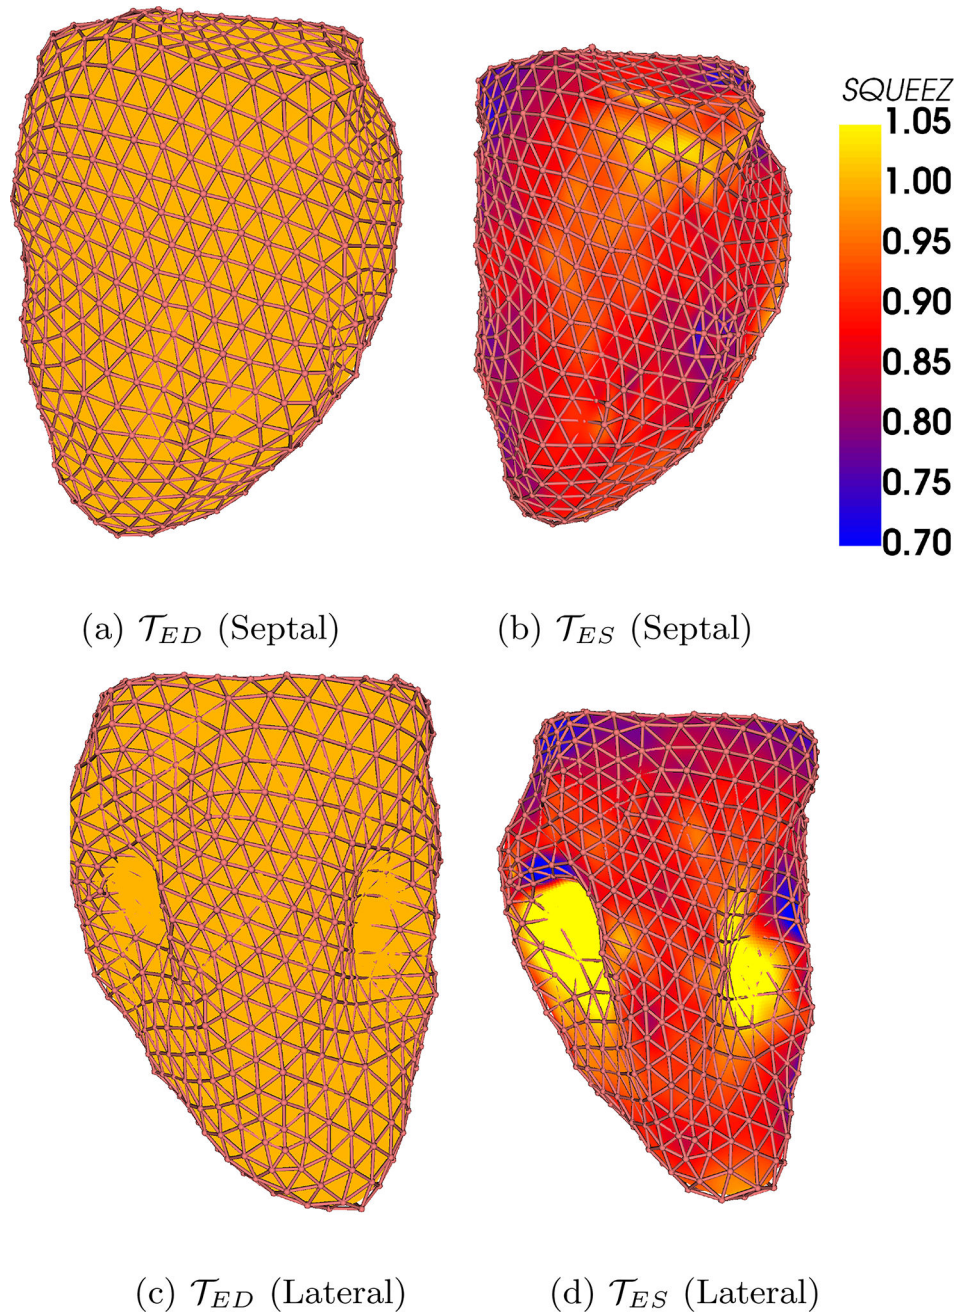
**Figure 7.**

Effect of varying  $\alpha_{ac}$  on global SQUEEZ. A discontinuity is noted in the CPD curve at frame 4. This discontinuity is visible in the  $\alpha_{ac} = 0.0$  SiSSR curve, but is less pronounced. Increasing  $\alpha_{ac}$  to 0.1 further diminishes this discontinuity without smoothing out potentially useful features at end systole and during diastasis. Values above 0.1 result in unwanted smoothing of these features.



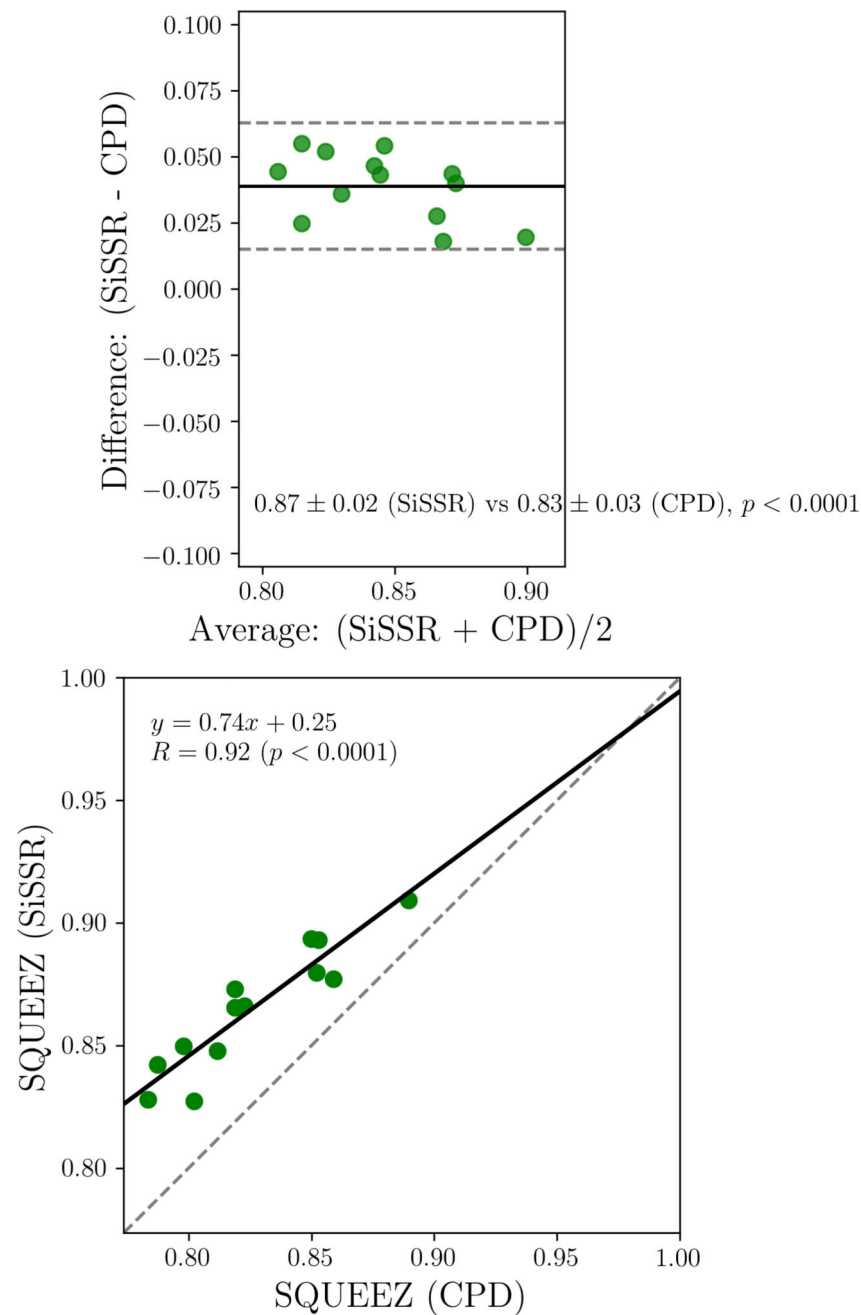
**Figure 8.**

Residual error at each registration stage. Residual error for the end diastolic (top) and end systolic (bottom) frames at various stages of the pipeline, overlaid on a long axis view of the heart. The pink meshes represent the subdivision surface of the model; the blue bars represent the primary residuals of the cost function, pointing from a point on the subdivision surface to the nearest boundary candidate. (a, d) The anatomical reference mesh,  $\mathcal{S}$ . (b, e) The registered model after the first pass,  $\mathcal{T}_f^1$ . (c, f) The registered model after the second pass,  $\mathcal{T}_f^2$ .




**Figure 9.**

Calculation of SQUEEZ from registered subdivision surfaces. (a) The septal wall at end diastole. (b) The septal wall at end systole. (c) The lateral wall at end diastole. (d) The lateral wall at end systole. Note that shading has been turned off for the subdivision surface to avoid any ambiguity in interpreting the colormap.



**Figure 10.**

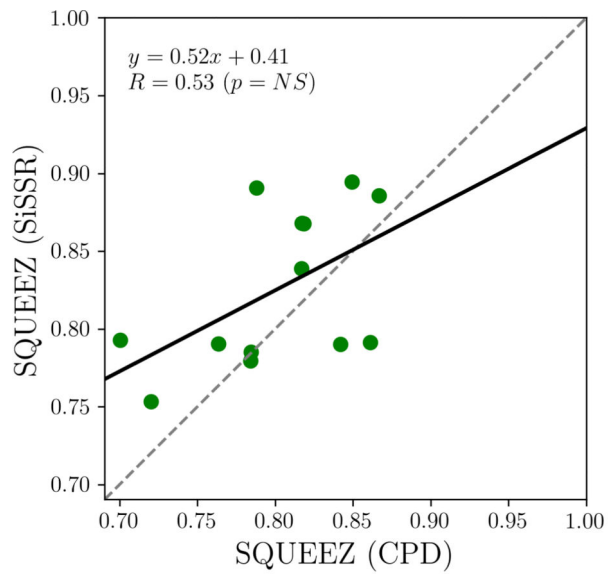
Peak global SQUEEZ measured by subdivision and CPD methods. Bland-Altman (top) and regression (bottom) comparing whole-heart, peak SQUEEZ as measured by SiSSR vs CPD. In the Bland-Altman, the solid line represents the mean difference, and the dashed lines represent  $\pm 1.96$  standard deviations. In the regression, the solid line represents the least-squares best fit, and the dashed line represents the ideal fit.




**Figure 11.**

Global SQUEEZ curves as measured by subdivision and CPD methods. Global SQUEEZ curves for all thirteen canines studied as measured by the subdivision regression method (above) and the coherent point drift method (below). Corresponding colors in the two plots indicate the same canine. SQUEEZ is by definition equal to one at the end diastolic frame, indicating no contraction.





(a) Anterior Apex (AHA Segment 13)



(b) Anteroseptal Base (AHA Segment 2)

**Figure 12.**

Peak segmental SQUEEZ as measured by the SiSSR and CPD methods. Segmental correlation is good in both infarcted (a) and remote (b) regions.



**Table 1**

Hyperparameters of the CPD registration.

Parameter	Value
Method	NONRIGID_LOW RANK
Spatial Normalization	true
Fast Gaussian Approximation	true
Fine Tuning	true
Function Tolerance	$10^{-5}$
Outliers	0.1
$\beta$ (Width of Gaussian)	2
$\lambda$ (Regularization Weight)	3

**Table 2**

Hyperparameters of the SiSSR registration.

Parameter	Value
Function Tolerance	$10^{-2}$
Parameter Tolerance	$10^{-2}$
Linear Solver Type	SPARSE_NORMAL_CHOLESKY
Minimizer Type	TRUST_REGION
$\alpha_{cf}$	1.0
$\alpha_{ac}$	0.1
$\alpha_{\hat{tp}}$	0.1
$\alpha_{gr}$	10.0

**Table 3**

Sensitivity of peak segmental SQUEEZ to regularization weights. Each row represents a single run, including all canines, where exactly one parameter is modified (either doubled or halved) relative to the baseline measurement in Table 2. The first column indicates which parameter was modified, and how. Linear regression was then performed between peak segmental SQUEEZ as calculated using the modified parameters versus baseline. The slope ( $m$ ), intercept ( $b$ ), coefficient ( $R$ ), and  $p$ -value of the linear regression are reported.

Parameter	$m$	$b$	$R$	$p$ -value
$\alpha_{tp} = 0.1 \times 2$	0.959	0.032	0.994	$p < 0.0001$
$\alpha_{tp} = 0.1 \div 2$	1.029	-0.022	0.993	$p < 0.0001$
$\alpha_{ar} = 10.0 \times 2$	0.975	0.021	0.996	$p < 0.0001$
$\alpha_{ar} = 10.0 \div 2$	1.020	-0.017	0.997	$p < 0.0001$
$\alpha_{ac} = 0.1 \times 2$	0.999	0.001	0.999	$p < 0.0001$
$\alpha_{ac} = 0.1 \div 2$	1.001	-0.001	1.000	$p < 0.0001$