Using the Stochastic Collocation Method for the Uncertainty Quantification of Drug Concentration Due to Depot Shape Variability

J. Samuel Preston, Tolga Tasdizen, *Member, IEEE*, Christi M. Terry, Alfred K. Cheung, and Robert M. Kirby*, *Member, IEEE*

Abstract—Numerical simulations entail modeling assumptions that impact outcomes. Therefore, characterizing, in a probabilistic sense, the relationship between the variability of model selection and the variability of outcomes is important. Under certain assumptions, the stochastic collocation method offers a computationally feasible alternative to traditional Monte Carlo approaches for assessing the impact of model and parameter variability. We propose a framework that combines component shape parameterization with the stochastic collocation method to study the effect of drug depot shape variability on the outcome of drug diffusion simulations in a porcine model. We use realistic geometries segmented from MR images and employ level-set techniques to create two alternative univariate shape parameterizations. We demonstrate that once the underlying stochastic process is characterized, quantification of the introduced variability is quite straightforward and provides an important step in the validation and verification process.

Index Terms—Drug diffusion, finite-element modeling, level set, porcine model, segmentation, shape model, stochastic collocation, uncertainty quantification.

I. INTRODUCTION

IN ANY numerical simulation, assumptions are made that do not precisely reflect the physical reality of the problem. These simplifications yield tractable problems and also introduce error. As part of the traditional simulation validation and verification (V&V) process, it is important to understand the impact of a given assumption on the final result. Assuming that a given modeling decision represents a specific choice from a set of possible decisions, it is possible to quantify the impact of that choice on the simulation. Knowing the variability introduced by the range of possible choices allows confidence in the final result to be adjusted accordingly. Probabilistic characterization of the effect of modeling decisions on simulation results is a powerful tool for validation, but is generally prohibitively costly due to the number of simulations that must be run to obtain a reasonable confidence level for the statistics. To quantify the impact of a given modeling decision using a standard Monte Carlo approach, individual samples (each representing a specific modeling choice) are randomly taken from the population of possible choices according to the underlying probability distribution. However, this method converges quite slowly (as $1/\sqrt{N}$ [1]), often requiring thousands of samples (and therefore thousands of simulations), which generally proves a prohibitive requirement. If instead we can model this population as arising from an underlying stochastic process, it is possible to apply the generalized polynomial chaos–stochastic collocation (gPC-SC) method [2], [3] to efficiently sample the probability space, yielding accurate statistics from a nearly minimal number of samples. The stochastic collocation method can be seen as a "sampling" extension to generalized polynomial chaos, which represents the stochastic process as a linear combination of orthogonal polynomials of random variables. This gPC representation has been used to solve stochastic differential equations in a number of fields [4]–[7]. Stochastic collocation has been successfully used to quantify the effects of heart position in ECG forward simulation [8] and in modeling uncertainty in diffusion simulation due to microstructure variability [9]. Since stochastic collocation builds statistics based on deterministic solutions for sampled stochastic parameter values, it only requires a standard discrete solver for the problem of interest. This makes for easy implementation and allows its use on problems with complicated governing equations for which a nonsampling gPC formulation would be difficult or impossible.

The specific experiment we are simulating in this paper attempts to test the efficacy of the antiproliferative drug rapamycin delivered via an injected triblock copolymer gel in limiting neointimal hyperplasia growth at the venous anastomosis of an implanted arteriovenous expanded polytetrafluoroethylene (ePTFE) graft. This experiment is part of a study analyzing the efficacy of such a procedure for possible future use in hemodialysis patients with implanted ePTFE grafts. The current study

J. S. Preston is with the Scientific Computing and Imaging Institute and the School of Computing, University of Utah, Salt Lake City, UT 84112 USA (e-mail: jsam@sci.utah.edu).

T. Tasdizen is with the Scientific Computing and Imaging Institute and the Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, UT 84112 USA (e-mail: tolga@sci.utah.edu).

C. M. Terry is with the Department of Medicine, University of Utah, Salt Lake City, UT 84112 USA (e-mail: christi.terry@hsc.utah.edu).

A. K. Cheung is with the Department of Medicine, University of Utah, Salt Lake City, UT 84112 USA, and also with the Medical Service, Veterans Affairs Salt Lake City Healthcare System, Salt Lake City, UT 84112 USA (e-mail: alfred.cheung@hsc.utah.edu).

*R. M. Kirby is with the Scientific Computing and Imaging Institute and the Department of Mathematics and Bioengineering, University of Utah, Salt Lake City, UT 84112 USA (e-mail: kirby@cs.utah.edu).

TABLE I
PHYSICAL REGION ATTRIBUTES

	Vein	Artery	Graft	Gel	NVT[1]
Diffusion Coefficient (mm^2/s)	4.09×10^{-7}	7.80×10^{-8}	2.91×10^{-6}	3.89×10^{-6}	1.0×10^{-5}
Structure Thickness (mm)	0.3	0.7	0.6	n/a	n/a
Init. Drug Concentration (mol/mm^3)	0.0	0.0	0.0	2.73×10^{-9}	0.0

uses a porcine model with grafts implanted between the common carotid artery and ipsilateral external jugular vein [10].

We are specifically interested in drug transport via diffusion from the injected gel deposit containing drug (the "drug depot") to the venous anastomosis (the "target site"). The geometry of the vein, artery, graft, and depot are modeled from MR images and diffusion coefficients of each are experimentally determined [11] (see Table I). All other structures in the test region are combined and labeled as homogeneous "nonvascular tissue." Additionally, the vein, artery, and graft wall thicknesses cannot be measured reliably from MR images. Therefore, we use measurements obtained from histological studies.

One of the modeling decisions involved in this simulation is the geometric shape of the drug depot. While the volume of gel and amount of drug are known and constant from experiment to experiment, each injection is a nonrepeatable process that yields a different depot shape based on the properties of the local tissue and variability in the injection technique. For an individual simulation, a single depot shape must be chosen from the range of possible shapes. This choice introduces variability to the experimental outcome; hence, it is important to study. In this paper, we propose a method for studying the effect of depot shape on the outcome of the experiment. In order to determine the variability in drug concentrations at some simulation time due to the choice of depot shape, we develop a shape parameterization to represent the underlying population of shapes. We then use the gPC-SC method to sample this population at well-chosen parameter values and derive statistics of the effect of shape on drug concentration at our chosen simulation time. While the specific shape models developed in this paper are characterized by a single parameter and are too simple to capture the entire variability observed from all injections, the contribution of this paper is the development of a method for studying the effect of shape variability on simulation outcomes. More complex shape models with multiple parameters can be easily inserted into the proposed framework in future studies.

The remainder of the paper is organized as follows. Section II introduces the notion of probabilistic characterization of shape variability and the use of the stochastic collocation method that take advantage of this characterization to effectively compute simulation outcome statistics. Section III outlines the image segmentation pipeline for obtaining realistic geometries (Section III-A), the specific transformations used in characterizing shape variability (Section III-B), the discretization of segmentation results (Section III-C), the simulation model (Section III-D), and its numerical approximation (Section III-E). Results of experiments with two different shape variability characterizations are presented in Section IV, followed by a concluding discussion in Section V.

II. MODELING UNCERTAINTY IN DEPOT SHAPE

To begin our study, we must first understand and clearly articulate the mathematical problem we are attempting to solve. We accomplish this by first describing our "experiment" (in the sense that a probabilist would), then by articulating a probability space that captures the salient features of the experiment about which we want to reason, and finally, by describing how the stochastic collocation method can be employed to allow computation of statistics of interest. Our "experiment" of interest is as follows. We assume that we have been given a test volume containing the static structures of interest—in our case, the vein, artery, and graft. These objects are modeled as remaining fixed in both position and material properties (see Table I) throughout the course of all experiments. We assume that we have also been given a parameterization of drug depot shapes such that a specific parameter value corresponds to a specific geometric depot model. Having both the static test volume and a specific depot model, diffusion simulations can be run yielding drug concentrations for positions within the test volume at any future time, respecting assumptions regarding our simulation method.

For our probabilistic experiment, we limit the set of admissible depot shapes to those represented by a univariate parameterization. In Section III-B, we will formulate two specific parameterizations: one that generate shapes with varying regularity in surface curvature and another that generates shapes by morphing between two realistic depot shapes as segmented from a magnetic resonance image. The initial concentration of drug in the depot is assumed to be uniform, and the volume of all depot geometries generated by this parameterization is constant. The volume, position, and gross shape of the depot were chosen to represent a "standard" injection result as seen by postinjection MRI. Variation in depot shape can now be expressed by assuming that any specific shape is drawn from a set consisting of shapes resulting from our parameterization. We assume that there is an equal probability of obtaining any specific depot shape from our parameterization, i.e., the probability distribution is uniform. Using this parameterization, we aim to compute statistics of drug concentration at some time t_f induced by the variations in depot shape expressed by this simple (but quantifiable) probabilistic experiment. To formalize this process, let $(\chi, \mathcal{A}, \mu)$ be a complete continuous probability space that expresses variation in the depot shape where χ is the event space consisting outcomes corresponding to depot shapes, $\mathcal{A} \subset 2^\chi$ the σ-algebra used to define measurable events, and μ the probability measure expressing the distribution from which outcomes are drawn; in our case, we choose a uniform distribution. We can now express the depot shape as a function of the uniform random variable ξ where the length of the distribution interval represents the range of admissible depot shapes.

We note that although a uniform distribution is used throughout this study, the stochastic collocation approach can be used for any compactly supported second-order random process. The random field of interest (and, in particular, its statistical characterization) in this study is the final drug concentration at a time t_f. Given that the depot shape in our experiment can be completely expressed in terms of ξ (given a single-parameter diffeomorphism of admissible shapes), and given that the drug concentration follows as a direct consequence of a depot shape, the drug concentration can be expressed as a function of ξ.

Let the distribution of drug concentration due to ξ be denoted by $f(\boldsymbol{x}, t_f, \xi)$ where $\boldsymbol{x}$ denotes points within our test volume and t_f denotes the final time of interest. Mathematically, the quantities we are interested in computing are given by the following expressions:

$$\text{mean}(f)(\boldsymbol{x}, t_f) = \mathbb{E}[f(\boldsymbol{x}, t_f, \xi)]$$

$$= \int_\Gamma f(\boldsymbol{x}, t_f, \xi) d\mu \qquad (1)$$

$$\text{var}(f)(\boldsymbol{x}, t_f) = \mathbb{E}[(f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^2]$$

$$= \int_\Gamma (f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^2 d\mu \qquad (2)$$

or more generally for the pth moment

$$\mathbb{E}[(f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^p]$$

$$= \int_\Gamma (f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^p d\mu \qquad (3)$$

where $\text{mean}(f)$ and $\text{var}(f)$ (and other high-order moments) are themselves fields defined over our test volume giving the mean and variance of drug concentration at any $\boldsymbol{x}$ due to depot shape variability, and Γ is the domain over which the random variable varies. Also note that $\text{mean}(f)$ and $\text{var}(f)$ are functions of t_f. Based upon these quantities, we can make steps toward quantifying the impact of depot shape variability on drug transport in our test volume.

We note that the process of moving from the uniform shape distribution to the distribution of $f(\boldsymbol{x}, t_f, \xi)$ (the drug concentration distribution) is highly nonlinear; the resulting distribution is unlikely to be uniform. We present means and variances (i.e., the first two moments) as the lowest order statistical information one would examine, but higher order moments may (and, in general, will) be necessary to fully characterize the statistics of the resulting process. An example analysis showing higher order moments is presented in Section IV.

The gPC-SC approach takes advantage of quadrature rules to integrate the stochastic process of interest over the stochastic domain, thus allowing the efficient computation of quantities such as means and variance. Under assumptions of smoothness, which, in this case, equate to the assumption that drug concentration varies smoothly with changes to our parameterized depot shape model, we can gain exponential convergence in the accuracy of our computed statistics as a function of the number of diffusion simulations we are required to run. In the discussion of our experiment, we will refer to the simplified diagram in Fig. 1. In the stochastic collocation approach, a collection
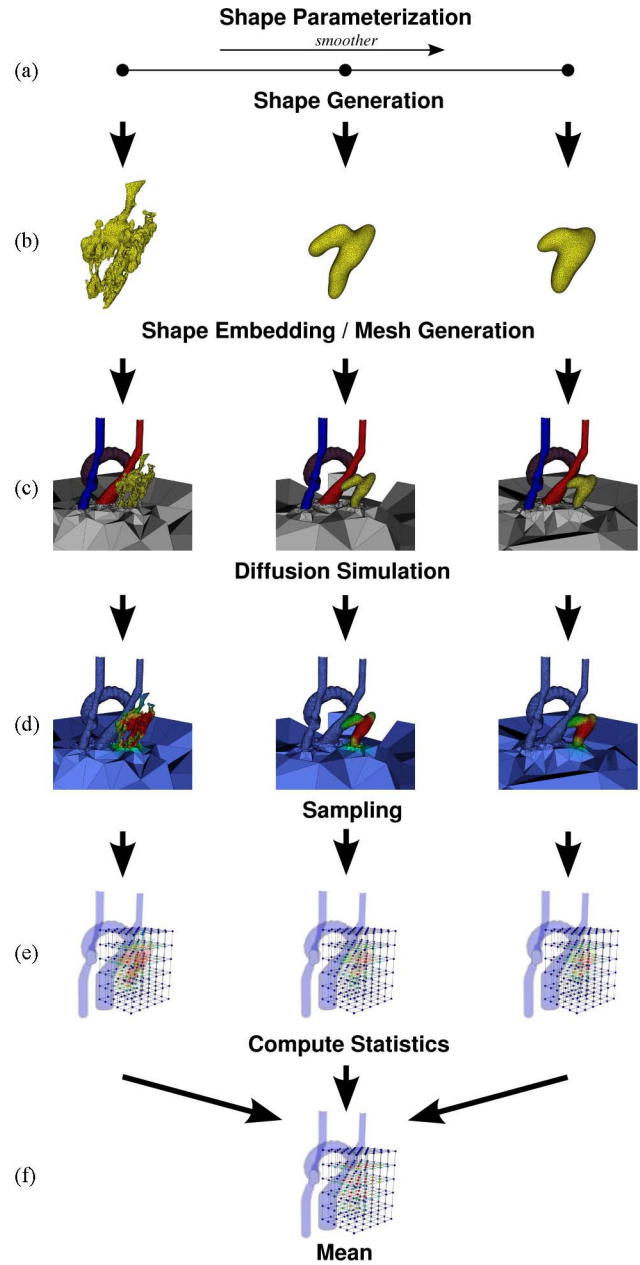

Fig. 1. Calculation of the mean drug concentration for a shape parameterization using first-order Smolyak points for collocation. (a) Univariate shape parameterization and three collocation points. (b) Three gel shapes corresponding to the collocation points. (c) Cutaway of tetrahedral mesh showing embedding of gel shapes. (d) Cutaway of tetrahedral mesh colored by concentration solution at day 14 of simulation. (e) Sampling of concentration values on a regular grid spacing over a region of interest. (f) Mean concentration at each sample point is calculated as a weighted sum of concentration values from (e).

of points at which the random field is to be sampled is specified, and a set of corresponding weights is computed, which account for the probability density function (pdf) characteristics underlying the set from which the points (or outcomes) are drawn. In the context of our study, each collocation point ξ represents a particular gel shape selected from our outcome set. Fig. 1(a) shows an example of a univariate smoothing parameterization sampled at three collocation points, and Fig. 1(b) shows the resulting depot shapes. Since the embedding process

is deterministic given a depot shape, the entire volume [shown in Fig. 1(c)] is a consequence of the chosen collocation point. At each collocation point, we can then compute the drug concentration $f(\boldsymbol{x}, t_f, \xi)$ through traditional finite-element diffusion simulation techniques, as shown in Fig. 1(d). Now, unlike traditional Monte Carlo in which large numbers of collocation points are used in an attempt to form accurate statistics, only a limited number of samples are used. The gPC-SC approach exploits smoothness assumptions on the random process to orchestrate the selection of points and weights such that accurate results can be obtained. We use third-order Smolyak points that require $q = 9$ points for a single random dimension, as this level of integration is sufficiently accurate for the needs of this study (note that Fig. 1 shows only the three first-order Smolyak sample points). In contrast, the number of realizations required for equivalent solution accuracy via the Monte Carlo method is significantly larger. While in the 1-D case the gPC-SC approach reduces to Clenshaw–Curtis quadrature [12], we will continue to formulate our method as stochastic collocation in order to motivate the extension to multidimensional parameterizations. Once diffusion computations have been done for each depot shape dictated by the collocation sampling, the statistics of drug concentration as a consequence of depot shape is given by the following expressions:

$$\text{mean}(f)(\boldsymbol{x}, t_f) = \mathbb{E}[f(\boldsymbol{x}, t_f, \xi)]$$

$$\approx \sum_{j=1}^{q} w_j f(\boldsymbol{x}, t_f, \xi_j) \tag{4}$$

$$\text{var}(f)(\boldsymbol{x}, t_f) = \mathbb{E}[(f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^2]$$

$$\approx \sum_{j=1}^{q} w_j (f(\boldsymbol{x}, t_f, \xi_j) - \text{mean}(f))^2 \tag{5}$$

where q is the number of points used and w_j denotes the weights. As discussed further in Section IV-A, calculation of these statistics over the distinct diffusion solutions requires sampling the solutions at a set of points uniform between simulation solutions, as shown in Fig. 1(e). The means and variances can then be computed at these points [see Fig. 1(f)].

Similarly, higher order centralized moments can be approximated by

$$\mathbb{E}[(f(\boldsymbol{x}, t_f, \xi) - \text{mean}(f))^p]$$

$$\approx \sum_{j=1}^{q} w_j (f(\boldsymbol{x}, t_f, \xi_j) - \text{mean}(f))^p. \tag{6}$$

Discussion of the formulation and computation of further statistics based upon the stochastic collocation method can be found in [3]. Of note is that increasing the order of the moments under examination may lead to an increase in the number of collocation points required.

III. METHODS

In order to run our diffusion simulation, we must first have a discretized representation of the tissue in which we are inter-ested. For our statistical experiment, the model will be a combination of "static" components not changed from simulation to simulation and a model of the injected drug depot generated from our parameterization that changes based on the parameter value chosen. The depot shape is assumed to be unchanged within a single diffusion simulation. While this ignores degradation of the gel over time, changes to the shape or diffusive properties of the depot would be difficult to experimentally determine and greatly increase the complexity of the simulation. Both the models of the static components and the depot shape parameterization rely on segmentations of MRI volumes. Once specific structures are segmented, geometric models of these objects are generated. These models are used to create a finite-element diffusion simulation for each parameter value. This process is described in greater detail in the following sections.

A. Image Segmentation

Lumen and depot geometries were constructed from MR images taken at weekly intervals following surgery. Gel injection was performed at these times, with pre- and postinjection scans acquired. Lumen geometry was obtained via a "black blood" 3-D turbo spin echo pulse sequence. This sequence uses a nonselective inversion pulse followed by a selective reinversion pulse in preparation for imaging in order to suppress signal from blood flowing into the imaging plane. A single lumen geometry was selected as the "standard" and used for all simulations; this is the static part of the model. Fig. 2 depicts the processing pipeline for image segmentation and mesh generation. A rough segmentation of the lumen was performed using user-defined seed points within the lumen as inputs to a connected component/thresholding algorithm [see Fig. 2(b)]. Due to the extreme narrowing of the vessels near the anastomosis and the common presence of flow and pulsatile motion artifacts in this region, a minimum cost path (MCP) algorithm is necessary to complete the lumen connection through this area [see Fig. 2(c)]. This technique is similar to the stenosis-robust segmentation technique described in [13]. Starting from the initial segmentation, this algorithm defines a graph in which voxels are considered nodes and each voxel has an edge to each immediate neighbor (six-connected). Since lumen should be the darkest region in this black blood scan, the cost associated with each edge is based on the intensity of the neighboring pixel. A search is then performed to find the MCP connecting the disconnected segmentation regions. An active-contour-based region growing step (implemented in the Symbolic Numeric Algebra for Polynomials (SNAP) [14] software package) was then used to refine and "fill in" the segmentation [Fig. 2(d)].

Gel images were taken immediately following injection using a specially developed pulse sequence for depot imaging [15]. Due to the improved contrast introduced by this sequence, a simple thresholding algorithm was sufficient for segmentation. The segmented volume agreed well with the known injected volume—taking a maximum and minimum "reasonable" threshold level changed the segmented volume to approximately $\pm 20\%$ of the known volume. The threshold level used was
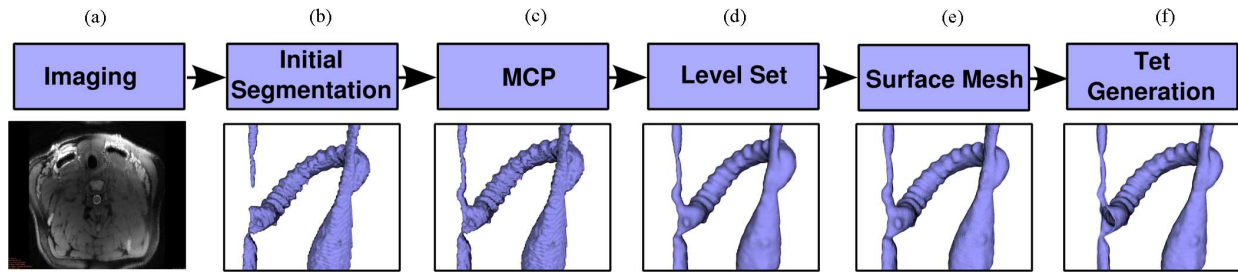
Fig. 2. Segmentation and discretization pipeline for lumen geometry. (a) Black blood MR image. (b) Initial segmentation from user-defined seed points. (c) Disconnected segments of lumen (see artery in the left of previous image) are connected using an MCP algorithm. (d) Segmentation is refined using a level-set-based region growing step. (e) Surface triangulation is generated from refined segmentation. (f) Tetrahedral mesh of vessel walls is generated using an offsetting technique.

chosen such that the segmented volume agreed with the known injected volume.

B. Characterization of Object Variability

To apply the stochastic collocation method, we require a parameterization of depot shape—the variable component of the geometric model. To physically model the injection process would require a complex, high-dimensional parameterization based on properties of surgically damaged tissue, gel properties, and injection techniques, and would result in a more complex simulation than the diffusion problem we are attempting to solve. Instead, our parameterization is based on the results we have—namely segmentations of the injected depot from postinjection MRI. In this way, we can develop a simple, low-dimensional parameterization that still respects the available data.

1) Level-Set Shape Representation: We explore two parameterizations both based on the evolution of existing depot segmentations using level-set methods [16]. Since level-set methods operate directly on a volume of data, they are an obvious choice for working with images acquired from MRI and have the advantages of easily allowing deformations based on differential properties of a surface and automatically handling changes in surface topology. For both parameterizations, we require that the volume of the depot be constant for any parameter value in order to keep the amount of drug uniform for all simulations while still using the known value for drug concentration in the injected gel.

Our parameterization will be formed by deforming an initial depot surface over time. Note that "time" in this section does not refer to the diffusion simulation time, but an arbitrary unit describing the amount of evolution of our level-set surface has undergone. In order to set up our level-set methods, assume that we start with some initial surface $\mathcal{S}_0$ at time $t = 0$. We will find a scalar function θ such that $\mathcal{S}_{0,k}$ is an isosurface of θ with value k, i.e.

$$\mathcal{S}_{0,k} = \{\boldsymbol{x} | \theta(\boldsymbol{x}, t_0) = k\}. \tag{7}$$

We can then modify this surface by changing the function θ due to properties of $\mathcal{S}_t$. More generally, any point $\boldsymbol{x}$ in a volume dataset at time t can be thought of as a point on the isosurface $\mathcal{S}_{t,\theta(\boldsymbol{x})}$. The function θ can be modified at all points based on properties of the isosurface passing through that point (the

"local isosurface"). The most obvious choice for our embedding θ is a distance transform to the isosurface $\mathcal{S}_t$. In this case, the surface we are interested in is $\mathcal{S}_{t,0}$, which we will now refer to as simply $\mathcal{S}_t$. For surface evolution, we move the isosurface $\mathcal{S}_t$ in the surface normal direction $\boldsymbol{n}$

$$\boldsymbol{n} = \frac{\nabla\theta}{\|\nabla\theta\|} \tag{8}$$

by modifying the values of θ over a series of time steps

$$\mathcal{S}_{t+1}(\boldsymbol{x}) = \mathcal{S}_t(\boldsymbol{x}) + V(\boldsymbol{x})\boldsymbol{n} \tag{9}$$

where $V(\boldsymbol{x})$ is the "speed function," a function of local and global properties of the local isosurface and independent forcing functions. As described in Section II, this parameterization will be sampled at specified collocation points ("times" in our current formulation). In order to ensure stability, we use a variable time step for our level-set evolution based on the fastest moving point on our surface. We do not attempt to sample at the exact collocation point specified, but use a "nearest neighbor" approach, ensuring that our surface is sampled within one time step of the prescribed sample time. Since the purpose of the variable time step is to limit the movement of any point on the surface to one voxel length in a single time step, the maximum error introduced by this approximation is one grid spacing, and for most of the surface, the actual error is much less.

Since we are only interested in a single isosurface of θ, it is unnecessary to update the entire volume at each time step. Only voxels near $\mathcal{S}_t$ at time step t are needed to calculate $\mathcal{S}_{t+1}$. Therefore, we use a sparse-field method [17] for calculating updates. This greatly reduces the computational burden while giving subvoxel accuracy for surface position.

2) Parameterization 1—Mean Curvature Flow: The first parameterization takes an existing gel shape chosen as an example of a particularly nonregular surface and smooths it using a volume-preserving mean curvature flow (MCF) transformation. MCF smooths a surface at each time step by moving each point on the surface in the surface normal direction by an amount proportional to the mean curvature. In its original, nonvolume-preserving formulation, MCF is the gradient descent for the surface area. The mean curvature H of the local isosurface is proportional to the divergence of the local isosurface normal

$$H = \frac{1}{2}(\nabla\boldsymbol{n}). \tag{10}$$
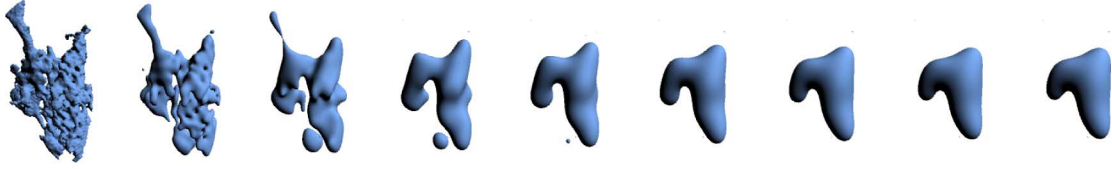
Fig. 3.   Gel shape smoothing using volume preserving MCF.
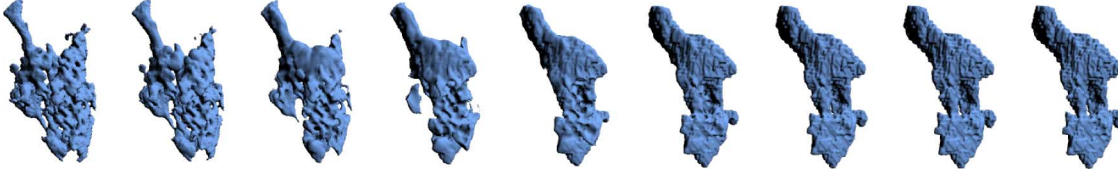

Fig. 4.   Volume-preserving gel shape metamorphosis.

This technique, however, reduces the interior volume of the surface, generally resulting in a singularity at t_∞. In order to enforce volume preservation, the movement of the surface is adjusted by the average mean curvature so that the total offset of the surface in the normal direction is zero

$$V(\boldsymbol{x}) = (H(\boldsymbol{x}) - \bar{H}) \tag{11}$$

where

$$\bar{H} = \frac{\int_S H(\boldsymbol{x})dx}{\int_S dx}. \tag{12}$$

The effect of volume preserving MCF is to create progressively smoother shapes that approach a sphere, as shown in Fig. 3. For a more rigorous treatment, see [18]. Using this parameterization in conjunction with stochastic collocation, we can study the effect of depot shape smoothness on experimental outcomes of drug diffusion. Equivalently, for a fixed depot volume, this can be seen as studying the effect of surface area.

3) Parameterization 2—Metamorphosis: The second parameterization performs a volume-preserving metamorphosis between two depot segmentations chosen to represent the extremes of surface regularity for the existing segmentations. The less-regular shape was taken as the "source" surface and the more-regular shape as the "target" surface for the morph. In order to allow volume preservation, the target surface was scaled and resampled to have the same interior volume as the source surface.

The morphing technique used is based on Breen and Whitaker's work on level-set shape metamorphosis [19], where transitioning between the two surface models occurs via a level-set formulation with the speed function given by

$$V(\boldsymbol{x}) = \gamma(\boldsymbol{x}) \tag{13}$$

where $\gamma(\boldsymbol{x})$ is the distance transform of the target image. In this way, points on the source surface will move toward the target surface whether they are inside or outside the target surface. This shape metamorphosis, however, is not volume preserving. While we know that the interior volumes at the beginning and end of the transformation are equal, the volume of intermediate surfaces is not constrained. As Breen and Whitaker noted, the

volume inside an intermediate surface Ω_t can be divided into two regions—Ω_t^i, a region of voxels inside the target surface that will grow to match the target, and Ω_t^o, a region of voxels outside the target surface that will shrink to zero volume. For any $\boldsymbol{x} \in \Omega_t$, the two regions can be distinguished by the sign of $\gamma(\boldsymbol{x})$. In order to ensure constant volume throughout the transformation, we introduce two scaling factors, ρ_i and ρ_o, which scale the speed function of voxels in Ω_i and Ω_o, respectively. The values of the scaling factors rely on the total movement of the surface of the two regions $S_t^i = \{\boldsymbol{x} \in S_t | \gamma(\boldsymbol{x}) > 0\}$ and $S_t^o = \{\boldsymbol{x} \in S_t | \gamma(\boldsymbol{x}) < 0\}$. Let $v_i = \int_{S_t^i} V(\boldsymbol{x})dx$ and $v_o = \int_{S_t^o} V(\boldsymbol{x})dx$. We then constrain the change in the two regions to be equal at each time step by updating ρ_i and ρ_o

$$\rho_i = \begin{cases} \dfrac{v_o}{v_i}, & \text{if } v_i > v_o \\ 1.0, & \text{otherwise} \end{cases} \qquad \rho_o = \begin{cases} 1.0, & \text{if } v_i > v_o \\ \dfrac{v_i}{v_o}, & \text{otherwise} \end{cases}$$

where the level-set formulation now becomes

$$\mathcal{S}_{t+1}(\boldsymbol{x}) = \begin{cases} \mathcal{S}_t(\boldsymbol{x}) + \rho_i V(\boldsymbol{x})\boldsymbol{n}, & \boldsymbol{x} \in \Omega_i \\ \mathcal{S}_t(\boldsymbol{x}) + \rho_o V(\boldsymbol{x})\boldsymbol{n}, & \boldsymbol{x} \in \Omega_o. \end{cases}$$

Fig. 4 shows the intermediate shapes obtained by the metamorphosis method between two segmented gel shapes. One of the gel shapes is less regular than the other, providing a second way of studying the effect of surface area for fixed-volume depot shapes.

C. Geometric Discretization

Once segmented, geometric surface models of the structures of interest were created and embedded in a control volume. This multisurface structure is used to define material regions and voids in a surface-conforming tetrahedralization of the control volume.

Since the imaged region is smaller than the desired control volume, the vessels are extended to the boundaries of the control volume. The vessels travel in roughly the z-direction within the imaged region, so this extension is accomplished by simply repeating the boundary z slices of the segmentation to the extents of the control volume.

A surface mesh of the lumen segmentation was created by slightly blurring the segmentation with a Gaussian kernel, then using this gray-level image as input to an isosurface meshing algorithm. The algorithm used is part of the Computational Geometry Algorithms Library (CGAL) [20] and uses iterative point insertion on the isosurface to create a triangle surface mesh with constraints on the maximum size and minimum angle of any member triangle.

Each of the vascular structures (vein wall, artery wall, and graft) has a different thickness and different diffusive properties. The portions of the lumen contained in the different structures cannot be clearly distinguished in the available MR images, so the generated surface mesh of the lumen was hand-annotated with this information based on knowledge of the anatomical structures. The resolution of the MRI sequence used to image the lumen was not high enough to directly image the vessel walls or graft. In order to create the geometry of the graft and vessel walls, thickness measurements from histology were used. Using the labeled surface dataset, a tetrahedral model of the vascular structures was generated by duplicating the surface mesh and offsetting each duplicate mesh point in the surface normal direction by a distance equal to the thickness of the vascular structure (see Table I for thickness values). Connecting each vertex to its offset duplicate produces a mesh of triangular prisms. The maximum triangle size of the surface mesh was chosen with respect to the thickness of the vascular structure in order to produce well-shaped prisms. This mesh is then subdivided such that each prism becomes three tetrahedra and all tetrahedra have coherent faces. Coherent prism splits are computed using the "rippling" algorithm described in [21]. The tetrahedra generated by this process then have maximum edge lengths roughly equal to the thickness of the structure they represent, 0.3–0.7 mm. While the simple offsetting technique used has the possibility of producing intersecting or incorrectly oriented tetrahedra, experimental results show that the curvature of our surface is low enough at all points relative to the offset distance to avoid producing any invalid elements. The surface mesh of a specific drug depot shape is generated from the output of our parameterized depot shape generation algorithm by the same method as used for the lumen surface.

The depot surface and the outer (offset) lumen surface are embedded in a cube representing the extent of our test volume. These surfaces are combined and used as the input piecewise linear complex (PLC) to TetGen [22]. The cube and the offset lumen surface define the boundaries of the PLC (i.e., the offset lumen surface represents a "hole" in the volume), and the elements of the depot surface mesh are used as internal constraining faces of the generated constrained Delaunay tetrahedralization (CDT) defining the drug depot region. The area outside of the depot in this CDT is considered nonvascular tissue. Elements generated are constrained to have a maximum radius–edge ratio of 2.0, allowing element size to grow away from the structures of interest. Finally, the previously computed tetrahedral mesh of the vascular structures is merged with this tetrahedral volume to create our final test domain containing five regions—the graft, arterial wall, venous wall, drug depot, and nonvascular tissue. The tetrahedralization step is constrained to preserve boundary

faces along the offset lumen surface, so this merging process requires only these duplicate points be merged. This process yields a discretization containing 210 000–270 000 elements depending on the specific depot shape.

D. Model Equations

For the purposes of our simulation, we assume that drug transport is only due to diffusion. Furthermore, we assume that diffusion rates are isotropic and constant for each material; let D be a function that is piecewise discontinuous and consists of diffusion coefficients D_m associated with each material m. The drug concentration c at position $\boldsymbol{x}$ and time t is then governed by the equation

$$\frac{\partial c(\boldsymbol{x}, t)}{\partial t} = D \boldsymbol{\nabla}^2 c(\boldsymbol{x}, t), \qquad \boldsymbol{x} \in \Omega \qquad (14)$$

where appropriate boundary conditions are applied. We assume that any drug diffusing into the interior of the vascular structures would be immediately removed by blood flow, leading to a known (zero) concentration at the lumen boundary (zero Dirichlet conditions). We also require that no diffusion occur at the boundary of the test volume.

E. Numerical Approximation

Solving (14) for any but the simplest geometries requires a numerical approximation, typically based on methods such as boundary elements or finite elements. For our simulation, we use the standard finite-element formulation with piecewise linear basis functions.

For the domain Ω representing our control volume, let $T(\Omega)$ be the tetrahedralization of the volume, as described in Section III-C. Also, let the set $\mathcal{N}$ contain the indexes of the vertices of $T(\Omega)$, which are the nodes of our finite-element mesh. We then decomposed the set $\mathcal{N}$ into two nonintersecting sets, $\mathcal{B}$ and $\mathcal{I}$, representing nodal indexes that lie on the boundary and nodal indexes of the interior, respectively. Let us furthermore decompose $\mathcal{B}$ into $\mathcal{B}_q$, the outer boundary of the test volume, and $\mathcal{B}_c$, the lumen boundary. To satisfy our requirements from Section III-D, we impose zero Neumann boundary conditions at $\mathcal{B}_q$ and zero Dirichlet boundary conditions at $\mathcal{B}_c$.

Let $\phi_i(\boldsymbol{x})$ denote the global finite-element interpolating basis functions, which have the property that $\phi_i(\boldsymbol{x}_j) = \delta_{ij}$ where $\boldsymbol{x}_j$ denotes a node of the mesh for $j \in \mathcal{N}$. Solutions are then of the form

$$u(\boldsymbol{x}) = \sum_{k \in \mathcal{N}} \hat{u}_k \phi_k(\boldsymbol{x}) \qquad (15)$$

$$= \sum_{k \in \mathcal{U}} \hat{u}_k \phi_k(\boldsymbol{x}) + \sum_{k \in \mathcal{B}_c} \hat{u}_k \phi_k(\boldsymbol{x}) \qquad (16)$$

where $\mathcal{U} = \mathcal{I} \cup \mathcal{B}_q$, which represents the unknown values of the problem. The first term of (16) then denotes the sum over the degrees of freedom (the values at the unknown vertices weighted by the basis functions) and the second term denotes the same sum for the (known) Dirichlet boundary conditions of the solution. Substituting the expansion (16) into the differential equation (14), multiplying by a function from the test space

$\{\phi_j(\boldsymbol{x}); j \in \mathcal{U}\}$, taking inner products, integrating by parts, and applying boundary conditions yield a linear system of the form

$$\mathbf{M}\frac{\partial \vec{u}}{\partial t} + \mathbf{K}\vec{u} = \vec{f} \tag{17}$$

where $\vec{u}$ denotes a vector containing the solution of the system (i.e., concentration values at the nodal positions in $\mathcal{U}$) and $\mathbf{M}$, $\mathbf{K}$, and $\vec{f}$ denote the mass matrix, stiffness matrix, and right-hand side function, respectively, given by the following expressions:

$$\mathbf{M} = M_{jk} = (\phi_j, \phi_k) \tag{18}$$

$$\mathbf{K} = K_{jk} = (\boldsymbol{\nabla}\phi_j, D \cdot \boldsymbol{\nabla}\phi_k). \tag{19}$$

In the previous expressions, $j, k \in \mathcal{U}$, and $(\cdot, \cdot)$ denotes the inner product taken over the entire spatial domain Ω. In our case, all $\hat{u}_i$ for $i \in B_c$ are zero and the flux on the exterior boundary was taken to be zero, so $f_j = 0 \forall j$.

The time-dependent equation (17) can be discretized using the Crank–Nicolson scheme. Arranging the unknown (future time step) values on the left-hand side yields the time stepping equation

$$\left(\mathbf{M} - \tfrac{\Delta t}{2}\mathbf{K}\right)\boldsymbol{u}^{n+1} = \left(\mathbf{M} + \tfrac{\Delta t}{2}\mathbf{K}\right)\boldsymbol{u}^n \tag{20}$$

where Δt is our time step and $\boldsymbol{u}^n$ represents our solution at time step n. To solve using this scheme, we must also be given an initial condition $c(\boldsymbol{x}, t = 0)$, which in our case represents the time of injection, with a known, constant drug concentration within the depot and zero concentration elsewhere. We retain the solutions to the diffusion simulation at a set of time points $\mathbf{T}$ at which we will calculate our statistics. Our simulation is carried out over a period of 80 days with a time step of 2 h, which has shown sufficiently accurate results for our purposes.

IV. RESULTS

The data presented in this section are computed from simulation results for input data representing the discretization of the simulation domain for depot shapes obtained by sampling our shape parameterizations at the nine Smolyak collocation points specified for third-order convergence.

Note that while these results give accurate characterization of the variability based on our modeling and simulation assumptions, validation based on *in vivo* experimental studies has not yet been performed.

We will describe and present data for three methods of obtaining statistics based on these simulation results. Section IV-A deals with obtaining statistics over an area for which the location and number of nodes of the solution $u(\boldsymbol{x})$ differ for different shape parameters. Section IV-B shows the mean and standard deviation of total drug delivery to a region of interest over time, as well as an example analysis of higher order statistics on this data. Section IV-C shows the simpler case of computing statistics over a static region of the discretized domain.

A. Domain Resampling

Since a different finite-element mesh is generated for each depot shape (as described in Section III-C), the solution vectors $\boldsymbol{u}$ from different discretizations do not represent weights for the
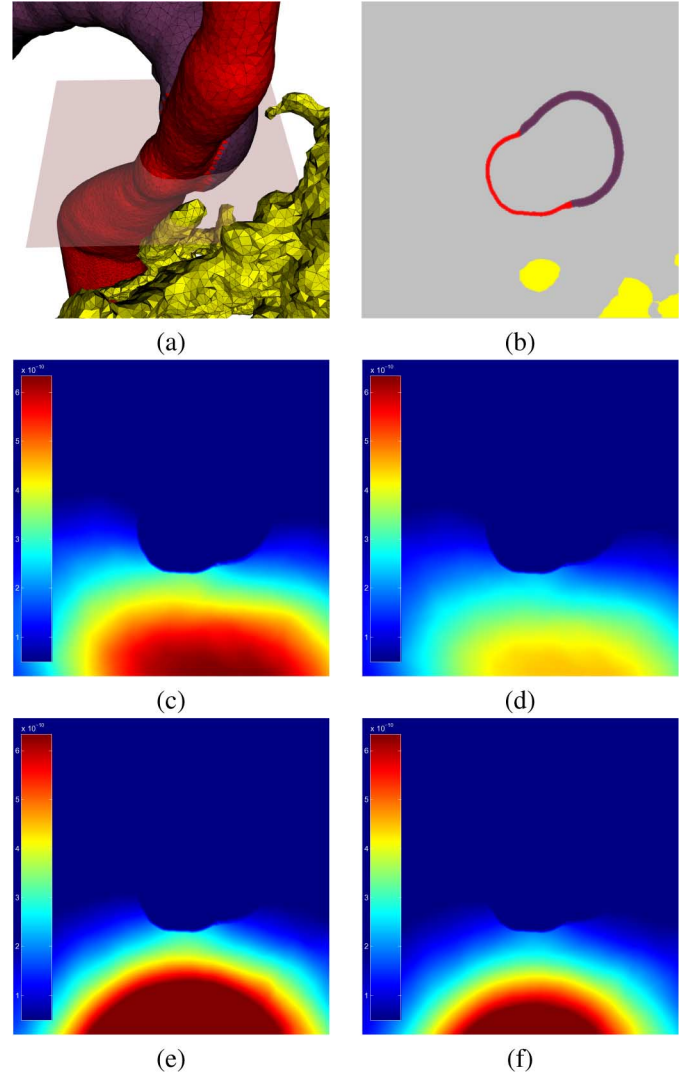

Fig. 5. Slice in the Z-plane of sampled collocation data at day 14. (a) Location of collocation slice. (b) Material locations in slice plane, colored to match the regions as shown in (a). (c) MCF diffusion mean values. (d) MCF diffusion standard deviation values. (e) Morph diffusion mean values. (f) Morph diffusion standard deviation values.

same basis functions, or even the same number of basis functions. However, the solutions $u(\boldsymbol{x})$ are defined everywhere over the same domain Ω. We can therefore take samples at a constant set of points in Ω, and compute statistics at these points. Since our control volume is specifically chosen to be larger than the area where we expect consequential levels of drug to diffuse during our simulation period, there is no reason to sample the entire control volume. For simplicity, we set a minimum level of drug concentration in which we are interested and choose an axis-aligned rectangular region encompassing all areas with at least this drug concentration in any of the simulation solutions. This region is then sampled on a regular grid, and the mean and standard deviation calculated at these points. This process is shown in Fig. 1(e) and (f). Fig. 5 shows an 11×11 mm 2-D region of the venous anastomosis for the solution at simulation day (t_f) 14 sampled at a resolution of 512×512. The location of the slice is shown in Fig. 5(a), with the structures

containing each sample point shown in Fig. 5(b). Note that the gel region changes for each choice of parameter ξ, and the specific depot region shown in yellow in Fig. 5(a) and (b) is for reference only. The mean and standard deviation at each sample point is shown in Fig. 5(c) and (d), respectively, for the shape parameterization based on MCF. The mean and standard deviation are similarly shown in Fig. 5(e) and (f) for the metamorphosis-based shape parameterization.

B. Total Drug Delivery

Since we are interested in the amount of drug delivered to the venous anastomosis over time, we also compute mean and standard deviation of total drug in the vein wall at each time point $t \in \mathbf{T}$. Each element of our tetrahedralization $\varepsilon \in T(\Omega)$ represents one of the materials in our domain. In particular, we define $\mathcal{E}_v$, the set of elements representing the tissue of the vein wall (see Fig. 6). From this subset of elements, the total amount of drug in the vein wall $V_j(t) = \int_{\mathcal{E}_v} u(\boldsymbol{x}) d\boldsymbol{x}$ is calculated for each solution time point $t \in \mathbf{T}$ in each simulation arising from a parameter value ξ_j. The mean and standard deviation at each time point is then given by

$$\text{mean}(V(t)) = \sum_{j=1}^{q} w_j V_j(t) \tag{21}$$

$$\text{var}(V(t)) = \sum_{j=1}^{q} w_j \big(V_j(t) - \text{mean}(V(t))\big)^2 \tag{22}$$

while as previously noted, we can compute higher order moments in a similar manner to (22).

Fig. 6(b) and (c) shows the amount of drug over time for individual shape parameter simulations as well as the mean and standard deviation as calculated earlier for our two shape parameterizations.

To look more specifically at the amount of drug reaching our target site, we define a spherical region of interest surrounding the venous anastomosis [see Fig. 7(a)]. We then define the set of elements $\mathcal{E}_s$ as those falling completely within this sphere [see Fig. 7(b)]. Of interest to us are the elements $\mathcal{E}_{sv} = \mathcal{E}_s \cap \mathcal{E}_v$. By the same equations [(21) and (22)], we can compute similar statistics on this reduced region. The results are shown in Fig. 7(c) and (d).

As noted in Section II, the resulting drug concentration distribution due to our shape parameterization is not completely specified by a mean and variance. Further statistics can be calculated, as shown in Fig. 8. This shows the skewness and kurtosis of the distributions of total drug delivered to the venous tissue at our sampled time points under the MCF shape parameterization. The skewness and kurtosis are calculated from the normalized third and fourth moments of the distribution as given in [23]. The pdf (whose domain is normalized to $[0, 1]$) of the distribution at day 10 is also shown. The means and variances for this simulation are shown in Fig. 6. Comparing these figures, we see that the skewness and kurtosis correspond to the increased "lopsidedness" of the distributions from the mean, as we would expect. The pdf is also as we would expect from the given time
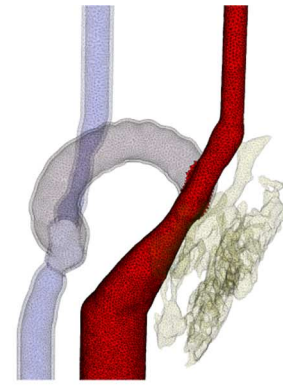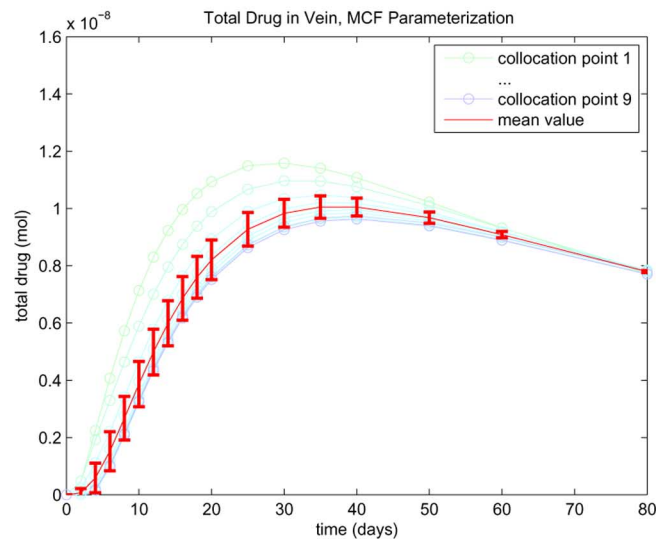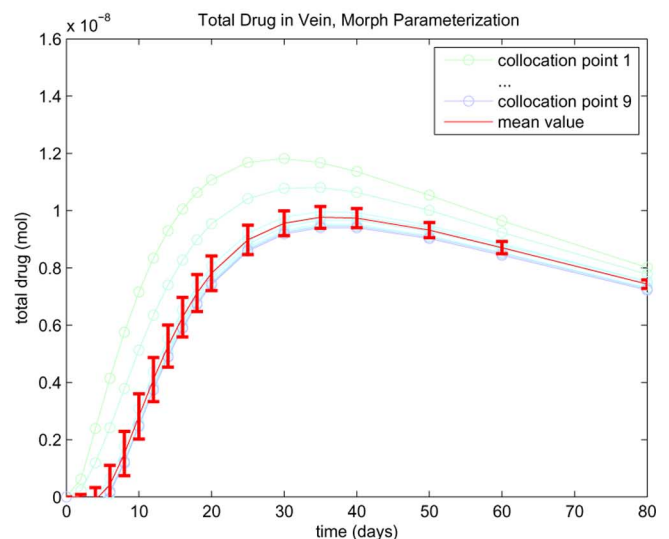
Fig. 6. Total drug deposited in vein wall [opaque geometry in (a)] over time. (b) MCF parameterization. (c) Morph parameterization. Results for the nine individual collocation points (depot shapes) are shown in blue to green, with the mean and variance shown in red.
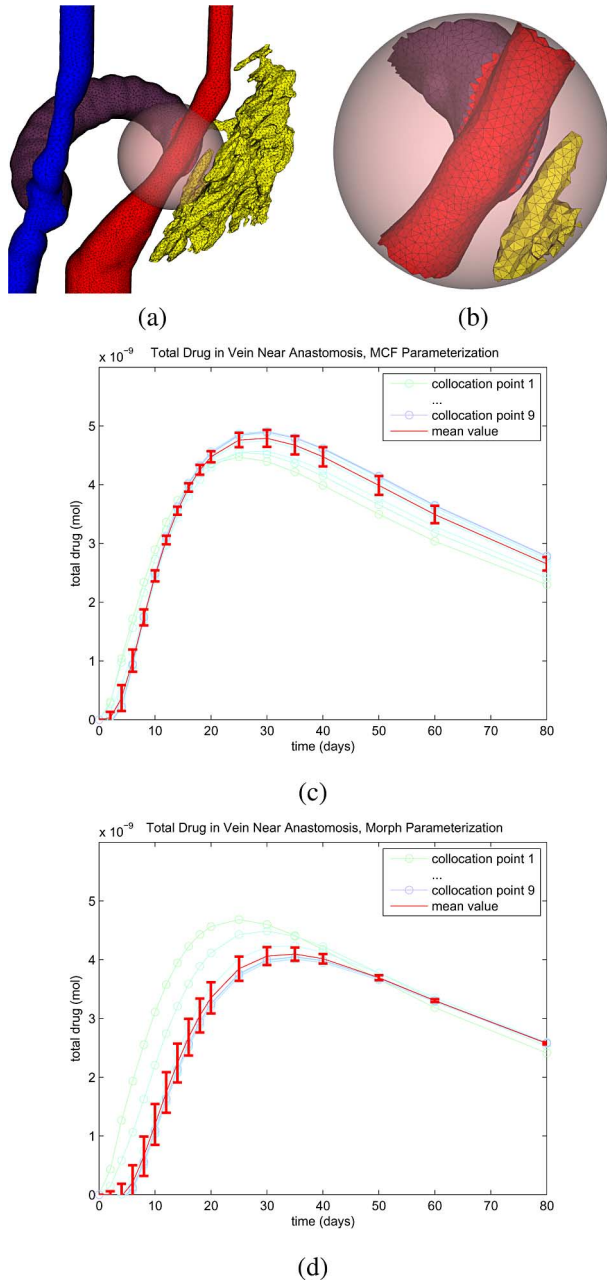
Fig. 7. (a) and (b) Region of interest around venous anastomosis. Total drug deposited in the vein wall over time. (c) MCF parameterization. (d) Morph parameterization. Results for the nine individual collocation points (depot shapes) are shown in blue to green, with the mean and variance shown in red.
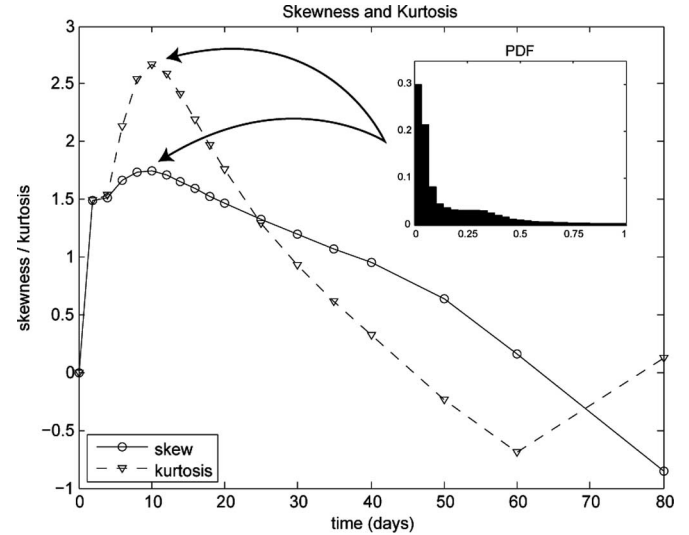

Fig. 8. For sampled time points, skewness and kurtosis are given for concentration distributions of total drug in venous tissue due to shape variability. The pdf of the distribution at day 10 is shown inset.
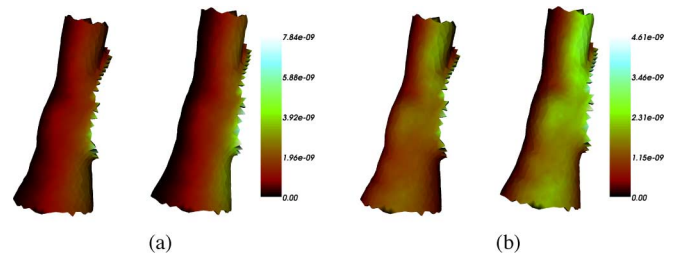

Fig. 9. Drug concentration mean and standard deviation at day 14 given by nodal statistics. (a) MCF parameterization. (b) Morph parameterization.

point's distribution, which has several samples grouped at the lower end of the concentration range—note that the mean is near the lowest concentration at the time point.

C. Nodal Statistics

While, in general, the discretization $T(\Omega)$ is different for different parameter values of ξ, using the discretization method described in Section III-C, the elements comprising the static geometry are always the same. In particular, the elements $\mathcal{E}_v$ comprising the tissue we are interested in are the same for any

value of ξ. Since the concentration field within a tetrahedral linear finite element is given by a simple linear weighting of the values at the element's nodes, the mean and standard deviation fields within $\mathcal{E}_v$ can be completely described by computing the means and standard deviations at the points $\mathcal{E}_v$ and using these as the weights of the linear basis functions. Since sampling, as described in Section IV-A, can be computationally expensive at high sampling resolutions or fine discretizations of the test volume due to the search for containing cell of each sample point, this method can improve performance when applicable. Fig. 9 shows the mean and standard deviation on the surface of $\mathcal{E}_{sv}$ as computed by this method.

V. DISCUSSION

The results of Section IV explore both the variability introduced by changes in depot shape and the effect of a specific parameterization of gel shape on this variability. From the results of Section IV-B, we see that the total amount of drug delivered to the venous tissue is quite similar for both shape parameterizations. We also see that the variance due to shape variation within each parameterization is fairly small, and that the amount of drug delivered decreases for "smoother" depot shapes
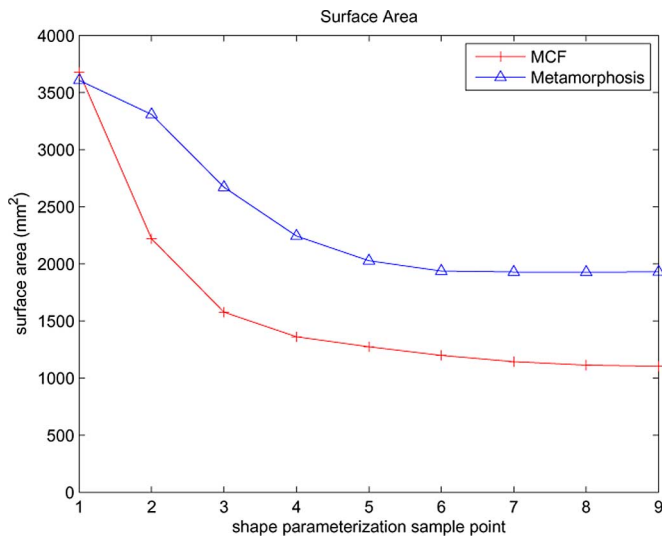
Fig. 10.   Surface area of shapes from shape parameterization.

(see Fig. 6). This agrees with our notion that increased smoothness corresponds to decreased surface area, which should lead to lower total diffusion rates. Fig. 10 shows how depot surface area changes over our two parameterizations. While the surface area decreases for both parameterizations, the MCF parameterization has a smaller final surface area. From this, we would expect a lower diffusion rate; however, as we can see in Fig. 6, the curves for total delivery over time are nearly identical for the later collocation points of the two parameterizations. This may be due to changes in the initial location and distribution of drug relative to the tissue in which we are interested. This effect is more pronounced if we restrict our interest to a more specific target site. Fig. 7 shows that if we look at only drug delivered to the venous tissue near the anastomosis, the two parameterizations are not equivalent. The MCF parameterization now has higher peak delivery as well as higher variance compared to the morph parameterization. A side effect of the "smoothing" caused by MCF is the reduction of the convex hull of the depot, concentrating more drug near the center of mass. We believe that this explains the increased delivery seen here, as more drug is located near the target site. As can be seen in Fig. 4, the depot retains a larger positional "spread" during our metamorphosis parameterization, and therefore, this effect is not apparent there. The results from Section IV-A show that the variance at any given point is closely related to the local mean concentration. This also agrees with expectations about a smooth process such as diffusion varying by a smooth process, e.g., such as our shape parameterization. These results are also quite helpful in determining how drug concentration varies across the target site and how the "front" of diffusing drug moves through the test volume.

While the parameterization of properties such as shape can be cumbersome and still requires judgment in determining that the appropriate range of variation is expressed, we have demonstrated that through methodologies such as the stochastic collocation method, the problem of studying the impact of shape

variability on a particular forward simulation is tractable. The basic requirements are that:

1) the variation can be expressed by a parameterized model;
2) the resulting stochastic field varies smoothly with changes to model parameters.

Once the underlying stochastic process is characterized, quantification of the introduced variability is quite straightforward and provides an important step in the V&V process, without which modeling (geometric, parameter, and mathematical) and simulation error bounds on the final solution must be viewed with some skepticism.

While our univariate shape parameterization is simple, in many cases, such a simple model will suffice. As noted earlier, development of this model will require judgment on the part of the modeler and affect the accuracy of the computed statistics. At the very least, however, a formal statement of confidence contingent on experimental realities respecting assumptions of the parameterization can be made.

It is worth noting that we are not limited to a univariate shape parameterization. If such a simple parameterization is not adequate, a multidimensional shape representation can be created. In this case, we will have a multidimensional random variable $\vec{\xi}$ and will use multidimensional sample points as described in [3]. The primary restriction on this method is that the entries of $\vec{\xi}$ must be independent. If this cannot be shown directly, the covariance matrix of $\vec{\xi}$ can be generated in a similar manner to the statistics previously described, and independence generated through a Gramm–Schmidt-like procedure. Although requiring more sample points than the univariate case, this extension scales well to a limited number of dimensions and is quite straightforward as the simulation process is identical regardless of the dimensionality of the underlying parameterization.

REFERENCES

[1] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*.   New York: Springer-Verlag, 2004.
[2] D. Xiu and J. S. Hesthaven, "High-order collocation methods for differential equations with random inputs," *SIAM J. Sci. Comput.*, vol. 27, no. 3, pp. 1118–1139, 2005.
[3] D. Xiu, "Efficient collocational approach for parametric uncertainty analysis," *Commun. Comput. Phys.*, vol. 2, no. 2, pp. 293–309, 2007.
[4] D. Xiu and G. E. Karniadakis, "Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos," *Comput. Methods Appl. Mech. Eng.*, vol. 191, no. 43, pp. 4927–4948, 2002.
[5] D. Xiu and G. E. Karniadakis, "Modeling uncertainty in flow simulations via generalized polynomial chaos," *J. Comput. Phys.*, vol. 187, no. 1, pp. 137–167, 2003.
[6] V. A. B. Narayanan and N. Zabaras, "Stochastic inverse heat conduction using a spectral approach," *Int. J. Numer. Methods Eng.*, vol. 60, no. 9, pp. 1569–1593, 2004.

[7] S. E. Geneser, R. M. Kirby, and R. S. MacLeod, "Application of stochastic finite element methods to study the sensitivity of ECG forward modeling to organ conductivity," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 1, pp. 31–40, Jan. 2008.

[8] S. E. Geneser, J. G. Stinstra, R. M. Kirby, and R. S. MacLeod, "Cardiac position sensitivity study in the electrocardiographic forward problem using stochastic collocation and BEM," to be submitted for publication.

[9] B. Ganapathysubramanian and N. Zabaras, "Modeling diffusion in random heterogeneous media: Data-driven models, stochastic collocation and the variational multiscale method," *J. Comput. Phys.*, vol. 226, no. 1, pp. 326–353, 2007.

[10] T. Kuji, T. Masaki, K. Goteti, L. Li, S. Zhuplatov, C. M. Terry, W. Zhu, J. K. Leypoldt, R. Rathi, D. K. Blumenthal, S. E. Kern, and A. K. Cheung, "Efficacy of local dipyridamole therapy in a porcine model of arteriovenous graft stenosis," *Kidney Int.*, vol. 69, no. 12, pp. 2179–2185, 2006.

[11] R. J. Christopherson, C. M. Terry, R. M. Kirby, A. K. Cheung, and Y.-T. Shiu, "Finite element modeling of the continuum pharmacokinetics of perivascular drug delivery at the anastomoses of an arterio-venous graft," submitted for publication.

[12] L. N. Trefethen, *Spectral Methods in MATLAB.*. Philadelphia, PA: SIAM, 2000.

[13] O. Wink, A. F. Frangi, B. Verdonck, M. A. Viergever, and W. J. Niessen, "3D MRA coronary axis determination using a minimum cost path approach," *Magn. Reson. Med.*, vol. 47, no. 6, pp. 1169–1175, 2002.

[14] P. A. Yushkevich, J. Piven, H. C. Hazlett, R. G. Smith, S. Ho, J. C. Gee, and G. Gerig, "User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability," *Neuroimage*, vol. 31, no. 3, pp. 1116–1128, 2006.

[15] C. M. Terry, H. Li, E. Kholmovski, S. Kim, and I. Zhuplatov, "Serial magnetic resonance imaging (MRI) to monitor the precise location of a biodegradable thermal gelling copolymer for drug delivery after ultrasound-guided injection in a porcine model of arteriovenous (AV) hemodialysis graft stenosis," submitted for publication.

[16] J. A. Sethian, "Level set methods and fast marching methods," in *Level Set Methods and Fast Marching Methods*, J. A. Sethian, Ed. Cambridge, U.K.: Cambridge Univ. Press, Jun. 1999, p. 400.

[17] R. T. Whitaker, "A level-set approach to 3D reconstruction from range data," *Int. J. Comput. Vis.*, vol. 29, no. 3, pp. 203–231, 1998.

[18] J. Escher and G. Simonett. (1998). The volume preserving mean curvature flow near spheres. *Proc. Amer. Math. Soc.* [Online]. *126*(*9*), pp. 2789–2796. Available: http://www.jstor.org/stable/119205.

[19] D. E. Breen and R. T. Whitaker, "A level-set approach for the metamorphosis of solid models," *IEEE Trans. Vis. Comput. Graphics*, vol. 7, no. 2, pp. 173–192, Apr. 2001.

[20] J.-D. Boissonnat and S. Oudot, "Provably good sampling and meshing of surfaces," *Graph. Models*, vol. 67, no. 5, pp. 405–451, Sep. 2005.

[21] S. D. Porumbescu, B. Budge, L. Feng, and K. I. Joy, "Shell maps," in *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*. New York, NY, USA: ACM, 2005, pp. 626–633.

[22] H. Si and K. Gartner, "Meshing piecewise linear complexes by constrained Delaunay tetrahedralizations," in *Proc. 14th Int. Meshing Roundtable*. Berlin, Germany: Springer-Verlag, 2005, pp. 147–164.

[23] G. W. Snedecor and W. G. Cochran, *Statistical Methods*. Oxford, U.K.: Blackwell, 1989.

Tolga Tasdizen (S'98–M'00) received the B.S. degree in electrical engineering from Bogazici University, Istanbul, in 1995. He received the M.S. and Ph.D. degrees in engineering from Brown University, Providence, RI, in 1997 and 2001, respectively.

From 2001 to 2004, he was a Postdoctoral Research Scientist in the Scientific Computing and Imaging Institute, University of Utah. From 2004 to 2008, he was a Research Assistant Professor in the School of Computing at the University of Utah. Since 2008, he has been an Assistant Professor in the Electrical and Computer Engineering Department and the Scientific Computing and Imaging Institute at the University of Utah, Salt Lake City. He also holds adjunct faculty positions in the School of Computing and the Department of Neurology.

Prof. Tasdizen is a member of the IEEE Computer Society.

Christi M. Terry received a B.S. degree in biology from Mesa State College, Grand Junction, CO, in 1988, and received the Ph.D. degree in pharmacology and toxicology from the University of Utah, Salt Lake City, in 1997.

She was with the biotechnology industry for approximately three years. She did a one-year Postdoctoral Fellowship on iron homeostasis under the University of Utah Human Molecular Biology and Genetics Program. She completed a three-year fellowship as a Nati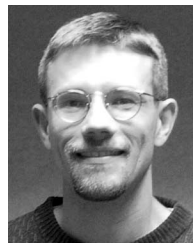onal Research Service Award Fellow in 2004 at the University of Utah, where she was engaged in research on the effect of polymorphisms on the expression and activity of the clotting protein, tissue factor. She is currently a Research Assistant Professor in the School of Medicine, Division of Nephrology and Hypertension, and an Adjunct Assistant Professor in the Department of Pharmaceutics and Pharmaceutical Chemistry, both at the University of Utah. Her current research interests include *in vivo* imaging of hemodialysis arteriovenous graft hyperplasia, the role inflammatory cells play in stenosis development, and characterization of drug pharmacokinetics and dynamics in an in vivo model of arteriovenous graft stenosis.

Alfred K. Cheung received the Graduate degree from Albany Medical College, Albany, NY, in 1977, and completed the Nephrology Fellowship from the University of California at San Diego, San Diego, in 1983.

He is a Board-Certified Clinical Nephrologist and a Professor of Medicine at the University of Utah, Salt Lake City. He holds the Dialysis Research Foundation Endowed Chair for Excellence in Clinical and Translational Research in kidney dialysis. His current research interests include laboratory and clinical research efforts focusing on the medical problems of patients with chronic kidney disease and the technical aspects of hemodialysis treatments. Other interests include the development of local pharmacological strategies for the prevention of vascular access stenosis.

J. Samuel Preston received the B.S. degree in computer engineering from the University of Florida, Gainesville, in 2003. He is currently working toward the M.S. degree in computer science at the University of Utah, Salt Lake City.

From 2004 to 2007, he was a Research Associate at the Center for Advanced Engineering Environments, Old Dominion University. His current research interests include image analysis and scientific computing.

Robert M. Kirby (M'04) received the M.S. degree in applied mathematics, the M.S. degree in computer science, and the Ph.D. degree in applied mathematics from Brown University, Providence, RI, in 1999, 2001, and 2002, respectively.

He is currently an Associate Professor of Computer Science at the School of Computing, an Adjunct Associate Professor in the Departments of Bioengineering and of Mathematics, and a member of the Scientific Computing and Imaging Institute, all at the University of Utah, Salt Lake City. His current research interests include scientific computing and visualization.