Optimisation of orthopaedic implant design using statistical shape space analysis based on level sets

Nina Kozic [a,*], Stefan Weber [a], Philippe Büchler [a], Christian Lutz [b], Nils Reimers [b], Miguel Á. González Ballester [c], Mauricio Reyes [a]

[a] *Institute for Surgical Technology and Biomechanics, Bern, Switzerland*
[b] *Stryker Trauma GmbH, Kiel, Germany*
[c] *Alma IT Systems, Barcelona, Spain*

Abstract

Statistical shape analysis techniques have shown to be efficient tools to build population specific models of anatomical variability. Their use is commonplace as prior models for segmentation, in which case the instance from the shape model that best fits the image data is sought. In certain cases, however, it is not just the most likely instance that must be searched, but rather the whole set of shape instances that meet certain criterion. In this paper we develop a method for the assessment of specific anatomical/morphological criteria across the shape variability found in a population. The method is based on a level set segmentation approach, and used on the parametric space of the statistical shape model of the target population, solved via a multi-level narrow-band approach for computational efficiency. Based on this technique, we develop a framework for evidence-based orthopaedic implant design. To date, implants are commonly designed and validated by evaluating implant bone fitting on a limited set of cadaver bones, which not necessarily span the whole variability in the population. Based on our framework, we can virtually fit a proposed implant design to samples drawn from the statistical model, and assess which range of the population is suitable for the implant. The method highlights which patterns of bone variability are more important for implant fitting, allowing and easing implant design improvements, as to fit a maximum of the target population. Results are presented for the optimisation of implant design of proximal human tibia, used for internal fracture fixation.

Key words: Statistical shape models, image registration, principal component analysis, level sets, orthopaedics, implant design

1 Introduction

Statistical shape analysis techniques have shown to be efficient tools to build population specific models of anatomical variability. Their flagship, the Active Shape Model (ASM), proposed by Cootes et al. (1995) provides a method to study the variability encountered across a population in a compact representation based on a decomposition via principal components analysis (PCA) (Bishop, 1995). Statistical shape models representing the variation of shape and gray-level appearance, namely Active Appearance Models (AAM) (Cootes et al., 2004; Cootes and Taylor, 2004), have been extensively used in image segmentation to locate structures of interest and to solve many medical image interpretation problems. For instance, they have been used to locate vertebrae in DXA images of the spine (Cootes and Taylor, 2004; Roberts et al., 2006; Smyth et al., 1996), structures in MR images of the brain (van Ginneken et al., 2002; Hill et al., 1994), the femoral head in MR images (Cootes and Taylor, 2004), the prostate in MR images (Haslam et al., 1994), and the outlines of ventricles of the heart in echocardiograms (Hill et al., 1994; Mitchell et al., 2000), amongst others. A comprehensive review of statistical shape models for 3D medical image segmentation is given by Heimann and Meinzer (2009). More recently, statistical shape models have been used for shape estimation in image-free computer assisted surgery (Rajamani et al., 2007).

In all these applications, the approach is to find the instance in the statistical shape model that best approximates the input data, subject to some regularisation constraints (Davies et al., 2002; Rajamani et al., 2007). Optimisation in shape space of more complex criteria based on clinically meaningful shape measures related to anatomical locations has not been fully explored. Sierra et al. (2006) formulate a minimisation process based on Lagrange multipliers to incorporate such additional constraints, and then optimise this criterion based on a gradient descent algorithm starting from the mean of the shape distribution. This is used in their application to generate virtual anatomical models for surgery simulation, instantiated by specifying clinical parameters, such as fundus/cervix length/width, that depend non-linearly on the shape coefficients. However, it is not guaranteed that their optimisation algorithm will produce the instance of the shape space that best meets the constraints. Further, in common to other existing works, the aim is to find a single instance from the statistical shape model as the solution to their problem. In certain cases, it may be interesting to find *all* instances of the shape model that meet

* Corresponding author. Address: Institute for Surgical Technology and Biomechanics, Stauffacherstrasse 78, 3014 Bern, Switzerland. Tel: +41 31 631 59 50; fax: +41 31 631 59 60

Email addresses: kozic.nina@gmail.com (Nina Kozic), mauricio.reyes@istb.unibe.ch (Mauricio Reyes).

a certain criterion. That is, one may be interested in estimating which range of the population falls within a given anatomical criterion, thus establishing a partition of the shape space into "valid" and "invalid" shapes.

In this work our aim is to develop a framework to evaluate a given anatomical/morphological criterion across the full PCA shape space, in order to find the group of shape instances that satisfy the criterion. The method is based on level sets on the parametric domain of the shape coefficients. Level set methods define a powerful optimisation framework that, in combination with statistical shape priors, has been used to recover objects of interest by the propagation of curves or surfaces (Bresson et al., 2006; Chen et al., 2002; Cremers, 2006; Leventon et al., 2000; Rousson et al., 2004). However, these previous works are of a very different nature to ours, as they deal with the extraction of structures of interest in medical images, employing level sets as their choice of shape representation. The shape prior is thus defined as a PCA of levels set representations, and the segmentation method finds the most likely shape. In our case, we do not employ level sets as a shape modelling tool, but as an optimisation framework to assess complex criteria in PCA space. The level set is therefore defined in the parametric shape coefficient space, not in image space. The high dimensionality of level sets allows for the segmentation of the space of any dimension, determined by the number of principal components retained. Moreover, the ability to represent complex topologies can be used to identify disconnected subsets of the shape space that meet the criterion.

The ultimate goal of an orthopaedic implant is to stabilise the fractured bone, to enable fast healing of the injured bone, and to return early mobility and full function of the injured extremity. These aspects are related to the shape of the implant, its material and the mechanical response it produces to decrease the stress at the fracture site. Although these three aspects should be considered when designing an orthopaedic implant, in this work we focus on the shape of the implant, and its ability to fit to the bone surface. Mechanical and material aspects are out of the scope of the presented study, although some comments about mechanical considerations are included in the conclusions section.

Current practice in orthopaedic research involves the evaluation of implants for fracture fixation by manual fitting and fixation procedures, applied on a small set of cadaver bones in a trial-and-error process to find the optimal implant shape and position (Goyal et al., 2007). More recently, a noninvasive semi-automatic method for quantifying implant fitting was developed (Schmutz et al., 2008). Although the authors discussed recommendations for optimising fitting, there are no real results on how these modifications would improve the fitting. Moreover, the method was tested on a small set of 21 CT data sets. Using limited amount of CT data or cadaver specimens does not necessarily describe the diversity in a population, such as age, gender or ethnic origin. This diversity can be studied using statistical shape analysis techniques. In this

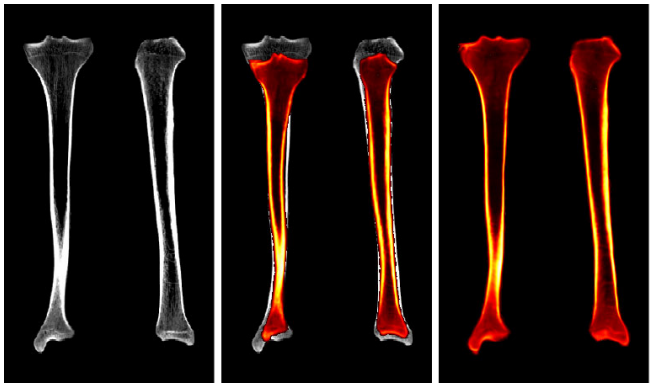

Fig. 1. Image registration. CT slice of the left human tibia, chosen as a reference bone; overlay of the rigid registration; overlay of the non-rigid registration matching the reference.

work, we show how our framework can be used as an evidence-based design methodology, assessing implant fitting on samples drawn from a statistical shape-and-intensity model by means of an automatic fitting procedure. We thus evaluate which proportion of the whole population is correctly fit by the proposed design. Then, by correlating segmented instances to the PCA manifold we are able to propose modifications to implant shape design as to fit a maximum of the target population.

Section 2 will briefly introduce the basic concepts behind statistical shape models based on PCA. In section 3, the key idea will be presented, that is the use of level set segmentation for PCA shape space optimisation. In section 4 we describe our framework for orthopaedic implant fitting assessment, and show results on the optimisation of the design of human tibial plates. Finally, discussion and conclusions are provided in section 5.

2 Statistical shape model

2.1 Image registration

The first step in generating a statistical model from a training set of images or shapes is to establish correspondences across the samples in the training set. Numerous approaches have been proposed in the literature, but since our aim in this paper is to construct shape-and-intensity models we will focus on non-rigid image registration techniques, and will illustrate the approach on CT images of human tibiae.

First, an image from the training set is selected as the reference, using an average box size as a reference, to which all other images will be registered.

In order to compensate for the different positioning during CT acquisition, we spatially align the remaining images of the training data set with the selected reference, via rigid registration. This allows to overcome the pose disparity and to maintain the size variation of the tibia (Figure 1). The next step in our model construction consists in warping the instances in the training set to the reference image. To capture the entire anatomical variability, we apply an intensity-based non-rigid registration algorithm (Rueckert et al., 2001, 2003). This algorithm defines the deformation as a B-spline mapping, defined by a uniformly-spaced grid of control points and the corresponding B-spline coefficients.

For the registration of CT data sets in our particular application, we employ sum of square distances (SSD) as the similarity metric, and gradient descent as the optimisation function. Based on the deformation fields obtained from the registration process, we build vectors of corresponding positions and image intensities. The reference image can be described as in Generalized Image Models (González et al., 2004):

$$v_R = (x_1, y_1, z_1, I_1, ..., x_n, y_n, z_n, I_n), \tag{1}$$

where n is the number of voxels in the region of interest and I_i is the intensity at voxel (x_i, y_i, z_i). Similarly, each of the other images can be described as a vector of the same length:

$$\begin{aligned} v_j = (&x_1 + \Delta x_1^j, y_1 + \Delta y_1^j, z_1 + \Delta z_1^j, I_1^j, ..., \\ &x_n + \Delta x_n^j, y_n + \Delta y_n^j, z_n + \Delta z_n^j, I_n^j), \end{aligned} \tag{2}$$

where $(\Delta x_i^j, \Delta y_i^j, \Delta z_i^j,)$ is the displacement vector at position (x_i, y_i, z_i), and I_i^j is the intensity of the voxel $(x_i + \Delta x_i^j, y_i + \Delta y_i^j, z_i + \Delta z_i^j)$ in image j.

2.2 Principal Component Analysis

The resulting image vectors described in Eq. (2) are high dimensional data, because we consider every point coordinate in the region of interest. To reduce the dimensionality of the data and obtain a compact parametric description, we apply principal component analysis. PCA is a multivariate factor analysis technique aiming at finding a low-dimensional manifold in the space of the data, such that the distance between the data and its projection on the manifold is small (Bishop, 1995). PCA is the best, in the mean-square error sense, linear dimension reduction technique.

Given a set of training data $\{\vec{t}_1, \vec{t}_2, ..., \vec{t}_N\}$, with $\vec{t}_i = (\vec{x}_i, \vec{y}_i, \vec{z}_i)$ and N equal to number of training instances, PCA finds a new orthonormal basis $\{\vec{u}_1, ..., \vec{u}_D\}$

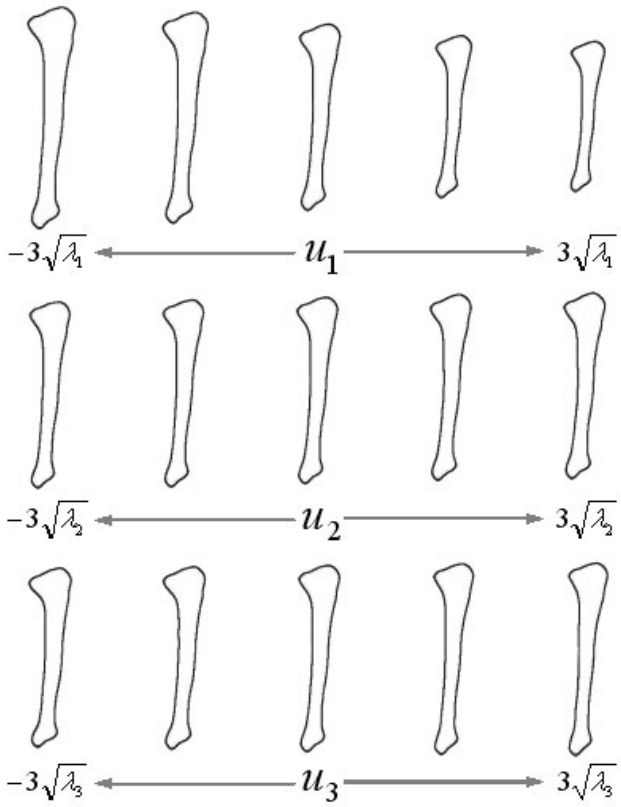

Fig. 2. The three first modes of variation for left human tibia are visualized individually. The first mode affects the change in the tibia length; the second mode influences the changes of the lateral condyle and a slight torsion of the lateral surface of the tibia; the third mode affects abduction of the medial condyle, changes of medial malleolus and in medial surface of the tibia.

with its axes ordered. This new basis is rotated such that the first axis is oriented along the direction in which the data has its highest variance. The second axis is oriented along the direction of maximal variance in the data, orthogonal to the first axis. Similarly, subsequent axes are oriented so as to account for as much as possible of the variance in the data, subject to the constraint that they must be orthogonal to the preceding axes. Consequently, these axes have associated decreasing "index" λ_d, $d = 1, ..., D$, corresponding to the variance of the data set when projected on the axes. The principal components are the set of new ordered basis vectors.

The principal components are found by computing the sample covariance matrix of the data set, $\vec{S}$, and then finding its eigenstructure

$$\vec{S}\vec{U} = \vec{U}\vec{\Lambda}.$$

$\vec{U}$ is a $D \times D$ matrix which has the unit length eigenvectors $\vec{u}_1, ..., \vec{u}_D$ as its columns, and $\vec{\Lambda}$ is a diagonal matrix with the corresponding eigenvalues $\lambda_1, ..., \lambda_D$. The eigenvectors are the principal components and the eigenvalues

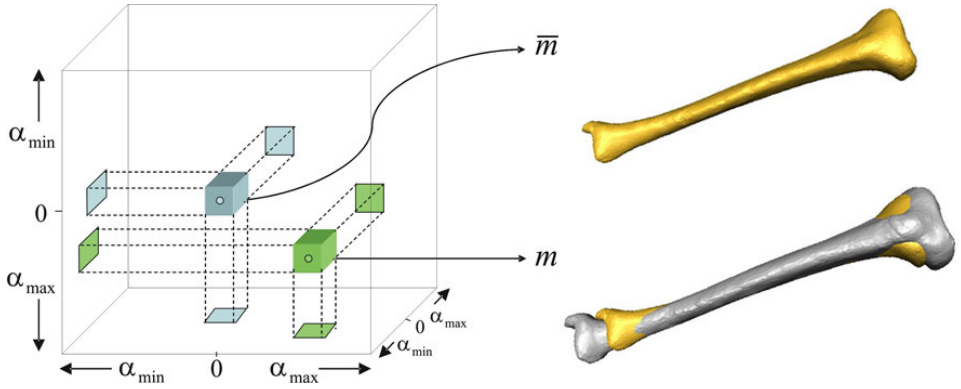

Fig. 3. Shape space defined by the three first principal components. The center element (labeled in the figure $\bar{m}$) corresponds to the mean of the population. Each element in this shape space is formed by a linear combination of the principal components, in this case $m = \bar{m} + \alpha_1\sqrt{\lambda_1}\vec{u}_1 + \alpha_2\sqrt{\lambda_2}\vec{u}_2 + \alpha_3\sqrt{\lambda_3}\vec{u}_3$).

their corresponding projected variances (Figure 2).

3 Optimisation in PCA space using level sets

3.1 PCA shape space mapping

Let us consider the shape space defined by the weighted linear combination of the first $L \leq D$ eigenvectors $\vec{u}_1, ..., \vec{u}_L$ of the PCA decomposition of a set of training shapes in $\mathcal{R}^D$. Each element $m \in \mathcal{R}^D$ in this shape space is defined by a set of coefficients $\alpha_1, ..., \alpha_L$ (Figure 3):

$$m = \bar{m} + \sum_{i=1}^{L} \alpha_i \sqrt{\lambda_i}\vec{u}_i, \tag{3}$$

where $\lambda_1, ..., \lambda_L$ are the eigenvalues corresponding to each principal component, and $\bar{m}$ is the arithmetic mean of the training sets. Now let us consider a scalar mapping $\mathcal{M} : A = [\alpha_{min}, \alpha_{max}]^L \rightarrow \mathcal{R}$. This mapping can represent a clinically meaningful anatomical criterion derived from the shapes in the PCA space (e.g. femoral inclination angle (Kozic et al., 2008)). We now would like to find all instances in the shape space that meet a certain criterion dependent on the scalar measure. This problem is approached as a segmentation in the PCA shape space defined by the mapping $\mathcal{M}$ defined above, and solved using the level sets framework described in the following section.

3.2 Level set segmentation

Segmentation techniques based on active contours, or deformable models, have been widely used in image processing for different medical applications (Kass et al., 1987; McInerney and Terzopoulos, 1996). The idea behind active contours is to extract the boundaries of homogeneous regions within the image, while keeping the model smooth during deformation. In such models, the initial contour, specified by the user, is evolved to the boundaries of the object by balancing two energy forces. The first force, computed from image data, represents external energy that attracts the curve toward image features, while the second force, defined within the curve, represents the internal energy and affects the smoothness of the curve. A particular instantiation of this paradigm is that of active contours based on level sets (Chan and Vese, 2001; Chen and Guan, 2004; Mumford and Shah, 1989; Tsai et al., 2001).

Let us consider a parameterized closed surface $C(s) : S = [0,1]^{L-1} \rightarrow \mathcal{R}^L$ defined in a bounded region $\Omega \in \mathcal{R}^L$. In order to segment the observed image $\mu : \Omega \rightarrow \mathcal{R}$ we propose to minimize the following energy functional:

$$E(C) = a \int_{\omega} (\mu - \epsilon)\, \partial\Omega + b \int_{S} |C'|\, ds, \tag{4}$$

where $\omega \subset \Omega$ and $C = \partial\omega$ is the closed surface. The first term represents the boundary force that attracts the evolving surface toward a predefined segmentation constraint $\epsilon = const$, while the second term regulates the smoothness of the surface. Here, a and b are positive scalar weights.

The energy functional proposed in Eq. (4) is not easy to solve because of the unknown set of complex surfaces C and unidentified image topologies. The segmentation algorithm developed in this work is based on the implicit representation of deformable models implemented within the framework of level sets. This implicit representation for evolving curves, introduced by Osher and Sethian (1988), allows automatic change of topologies without reparametrization. Using the level set formulation, the boundary surface $C = \partial\omega$ can be modeled as a zero level set of a Lipschitz function ϕ, defined on the entire image domain Ω as (Figure 4):

$$\begin{aligned} &C = \partial\omega = \{x \in \Omega : \phi(x) = 0\}, \\ &inside(C) = \omega = \{x \in \Omega : \phi(x) > 0\}, \\ &outside(C) = \Omega \setminus \omega = \{x \in \Omega : \phi(x) < 0\}. \end{aligned}$$

Having the Heaviside function $H(\phi)$ defined on the whole image domain as $\int_{\omega} \partial\Omega = \int_{\Omega} H(\phi) dx$, for $\omega \subset \Omega$, and its corresponding Dirac function $\delta(\phi) = \frac{d}{d\phi} H(\phi)$, we can replace the unknown variable C by the level set function $\phi(x)$

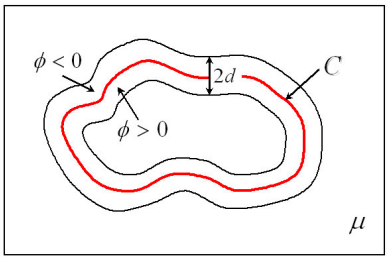

Fig. 4. Narrow band level set approach allows us to compute mapping values only for the points in a narrow band around the zero level set (red line).

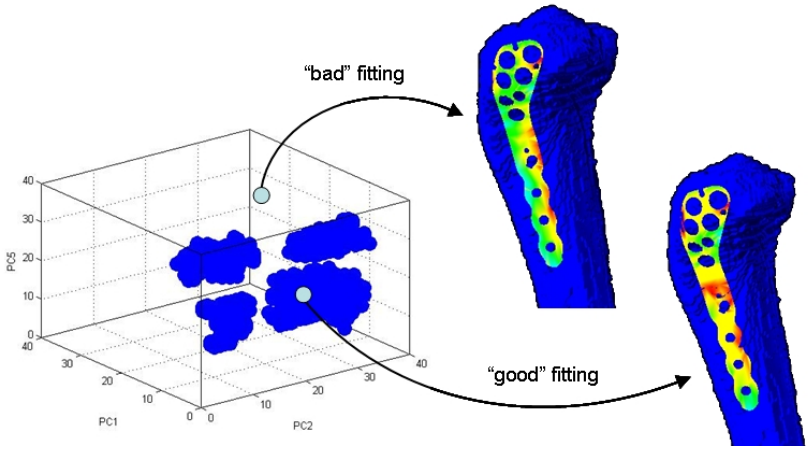

Fig. 5. Level set segmentation of the PCA shape space. Clinical criteria is chosen to be the implant fitting distance to the proximal human tibia. Segmented regions (in blue) satisfy the given segmentation criterion for the 'good' fitting (average fitting distance less than 1mm). Results of the fitting for the instances from 'good' and 'bad' fitting area in the PCA shape space are visualised.

as:

$$E(\phi) = a \int_{\Omega} (\mu - \epsilon)\, H(\phi)\, dx + b \int_{\Omega} \delta(\phi)\, |\nabla(\phi)|\, dx, \tag{5}$$

where the surface value $|C(\phi = 0)| = \int_{\Omega} \delta(\phi)\, |\nabla(\phi)|\, dx$ is estimated directly from the level set function (Evans and Gariepy, 1992). By minimizing the energy functional with respect to ϕ we get a model associated Euler-Lagrange equation for boundary flow:

$$\frac{\partial \phi}{\partial t} = a\,(\mu - \epsilon)\,\delta(\phi) + b\, div(\frac{\nabla \phi}{|\nabla \phi|})\,\delta(\phi), \tag{6}$$

where t is an artificial time $t \geq 0$ for boundary flow and $\int_{\Omega} |\nabla(\phi)| dx = div \frac{\nabla(\phi)}{|\nabla(\phi)|}$.

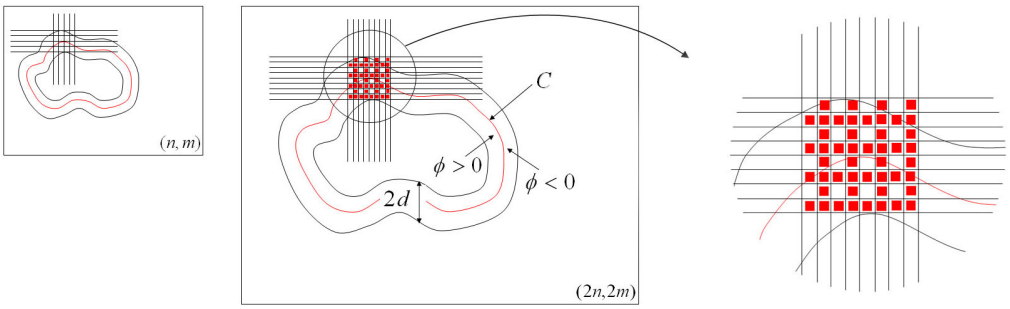

(a) (b) (c)

Fig. 6. Hierarchical approach to narrow band zero-level set evolution. (a) Initial low resolution 2D image space map with a stable zero level set in red colour and a narrow band around it. (b) Higher-resolution map with the augmented narrow band and zero level set, adopted from the low-resolution map. (c) The values of white pixels in the grid map come from the low resolution map, while the values of the red pixels that come from the augmented map still need to be calculated.

3.3 Level set optimisation in PCA shape space

In the framework of our application to the evaluation of anatomical criteria in PCA shape space, shape space will be the L-dimensional "image" μ to be segmented, defined in the domain of shape coefficients $\Omega = A$. Thus, level sets are used to find the region in the shape space defined by the weights applied to the principal components, in which the criterion is met (Figure 5). The flexibility of level sets allows to identify disconnected regions of the shape space. Further, the generality of the method allows to define any criterion, including complex functions that depend non-linearly on the shapes defined by the principal components.

It must be noted that in this work level sets are not used as a shape representation method, as is the case in all previous works that combine level sets with statistical shape models (employed as prior in the segmentation process). Rather, we do the analysis in the statistical shape space directly, not in image space, and we deal with the identification of a population, rather than a particular image.

3.4 Hierarchical approach to zero level set evolution

In order to decrease the computational complexity of the standard level set method we extend a narrow band level set approach, which uses only the points close to the evolving front at every time step (Adalsteinsson and Sethian, 1995) to hierarchical narrow band level set (HNBLS) approach. First we initialize our level set function using automatic seed initialisation on a low resolution

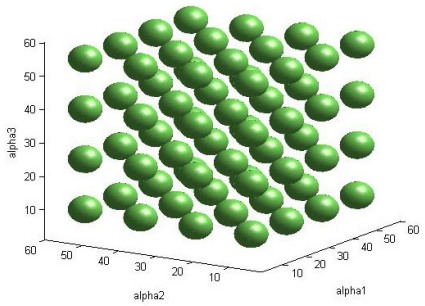

Fig. 7. Seed initialisation of the level set function.

image map. The seed initialisation consists of partitioning the data image u_0 into N windows W_n, $n = 1..N$ (Figure 7). Windows are of predefined size and do not overlap. The size is selected empirically to be $dim(\mu)/15$ in order to detect all the small "irregularities" in the image space and to decrease computational time. Level set function is computed only in these seed points. Then, minimisation of the energy functional (Eq. (5)) is performed to evolve the surface towards the segmented region.

We define a thin band around the zero-level set, that contains the neighboring points with distance to the zero-level less than d_{max} and we update the level set only on these points (Eq.(6)), instead of re-calculating it for each grid point (Figure 6a). As the zero-level set corresponding to the front evolves, we must ensure that it stays within the band. We re-initialise the band after 10 iterations, when the front is close to the edge of the domain, using the current zero-level set as the initial surface. Once the stable boundaries of the low resolution map are reached we increase the resolution of the image space and continue zero-level set surface evolution in the augmented low-resolution narrow band (Figure 6b).

The hierarchical narrow band level set algorithm is as follows (Figure 8):
Step 1. Initialise the zero level set function ϕ^0, as a corresponding circular signed distance on each window W_n. Construct a thin band around zero-level set $\beta^0 = N(\phi^0)$.
Step 2. Update ϕ^{k+1} for all pixels on β^k (Eg.(6)). If $k(mod 10) = 0$ then go to Step 4, else if k is equal to a maximum number of iterations, then stop.
Step 3. Update narrow band β^k and assign values to new pixels on narrow band. Outside the domain the value is defined as: $\phi^{k+1} = +d_{max}$ if the point is inside of the curve and $\phi^{k+1} = -d_{max}$ if the point is outside of the curve. Go back to Step 2.
Step 4. Increase the resolution of the image space and compute the values of the missing pixels in the augmented low-resolution narrow band. Go back to Step 2 (Figure 6).

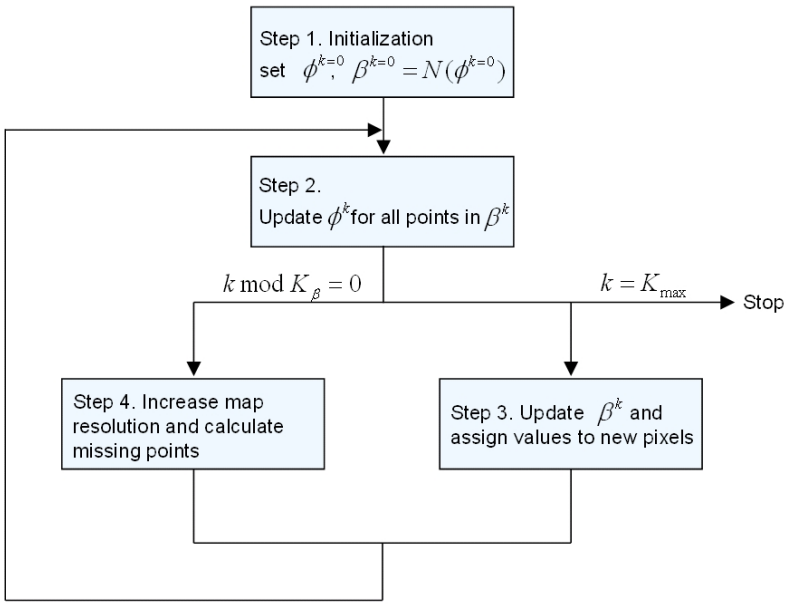

Fig. 8. Hierarchical level set segmentation algorithm.

4 Optimisation of orthopaedic implant designs

4.1 Clinical context

Since the late 1950s, open reduction and internal fracture fixation has been used to restore bone anatomy and enable early mobilization. Internal fixation alters the biology of fracture healing and reduces strain at the fracture site. Plate contouring is an important step in osteosynthesis. Plates are precontoured before or during a surgery to match either patient-specific or an average bone anatomy. The safety and ease of this procedure depends on certain material properties of the plate, such as the yield point and fatigue endurance (Frankel and Burstein, 1970). In addition to this, contouring is affected by the complexity of the bone shape to which the plate has to fit.

Nowadays, with an annual incidence of over a half million fractures of the tibia and fibula in the US (Russell and Levine, 1996), manufactures are moving from costly patient-specific implant design to the average implant shape that can fit to a given population.

Currently in orthopaedic research, the evaluation of implants for fracture fixation is done by manual fitting and fixation procedures, applied on a small set of cadaver bones in a trial-and-error process to determine the optimal implant shape and position (Figure 9). Goyal et al. (2007) investigated the accuracy of the periarticular tibial plate fit using 101 cadaver specimens of human tibia, on whom the implants were manually fixed by visually finding the best implant

Fig. 9. Internal fixation of the proximal tibia implant.

position. More recently, a noninvasive semi-automatic method for quantifying implant fitting was developed (Schmutz et al., 2008). In this study the surface of the plate was fitted to 21 computer tomography (CT) based 3D models of human tibia. Although the recommendations for implant modifications were discussed, there are no conclusive results on how these modifications would improve fitting.

4.2 Automatic implant fitting algorithm

A modified Iterative Closest Point (ICP) technique (Besl and McKay, 1992), developed in our group (Reyes et al., 2008), was used for the specific task of bone implant fitting. The method initialises the position of the implant close to the bone surface and optimises its position as to fit the bone as closely as possible, subject to specified positioning constraints. Based on this result, it computes the distance map from each point in the implant to the closest point in the bone surface.

In this work, the method is refined by a modified collision constraint to ensure that no points in the implant mesh model fall inside the bone model. Collision detection is performed by tracking the change of direction between the vector pointing from the inspected point to its closest point to the mesh and its normal vector. In addition, fitting guidelines provided by the implant manufacturer were included as fitting constraints, this in order to find plausible implant fittings. These further specific constraints favor fittings of the implant that are collinear with the bone main axis, and do not take place above the bone plateau (Figure 10).

The constrained ICP algorithm is based on the optimization of the following functional:

$$\operatorname{argmin} \sum_i W_i \|e_i\|, \tag{7}$$

where W_i and e_i are the corresponding weight and distance error for point i in the implant mesh model, respectively. The weights W_i are computed as

a linear combination of constraint-specific weights for collision W_i^C, implant-bone collinearity $W^{||}$, and tibia plateau W_i^p:

$$W_i = W_i^C + W_i^{||} + W_i^p. \tag{8}$$

The collision weight W_i^C is computed as follows:

$$W_i^C = \begin{cases} 1 & p_i \notin V_{in} \\ k_i^c \|e_i\| & p_i \in V_{in} \end{cases}, \tag{9}$$

where V_{in} is the 3D space inside the bone model. To detect if a point p_i is inside or outside the bone model, the sign of the dot product between the normal vector on the bone surface closest to p_i and the vector formed by p_i and its closest point on the bone surface is computed.

In order to avoid biases due to the number of points inside and outside the volume, the variable k_i^c in Eq. (9) is proposed by the following inequality:

$$k_i^c \geq (N_{tot} - N_{in}) / \sum_{i \in V_{in}} \|e_i\|, \tag{10}$$

with N_{tot} the number of points of the implant mesh, and N_{in} the number of points falling inside the bone model. We have found that adjusting the weight k_i^C we avoid biases due to the number of points inside and outside as the iterations proceed.

Similarly as for the collision constraint,weights $W_i^{||}$, and W_i^p are computed as follows:

$$W_i^{||} = \begin{cases} 1 & \alpha \leq \alpha_{th} \\ k^{||} \|\alpha_{th} - \alpha\| & \alpha > \alpha_{th} \end{cases}, \tag{11}$$

$$W_i^p = \begin{cases} 1 & p_i \in \Gamma \\ k_i^p \|p_i - \Upsilon\| & p_i \notin \Gamma \end{cases}, \tag{12}$$

where α is the angle between the implant main axis and the bone main axis, α_{th} is a threshold angle chosen by the user, $k^{||}$ is a scalar chosen empirically and used to weigh the global effect of the parallelism constraint, Υ is the z-coordinate of the plateau region interface, and Γ is the 3D space above the bone plateau (Figure 10).

For the computation of α the main axis of the implant model and the bone are required. This is performed through a Oriented-Bounding-Box (OBB) decomposition of both shapes (Figure 11b). Furthermore, for the implant model, only the lower region is used in order to improve the alignment between the bone shaft and the implant.

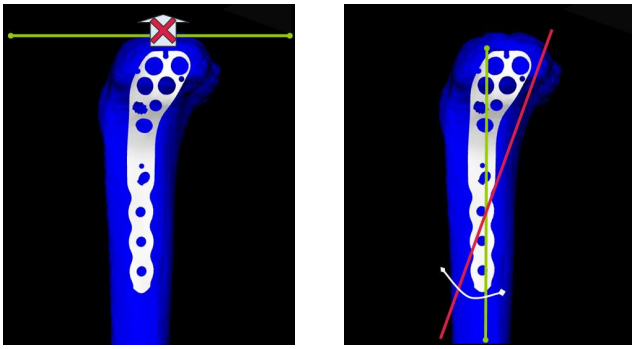

(a) (b)

Fig. 10. The constraints proposed by the implant manufacturer: (a) plateau constraint and (b) parallelism constraint.

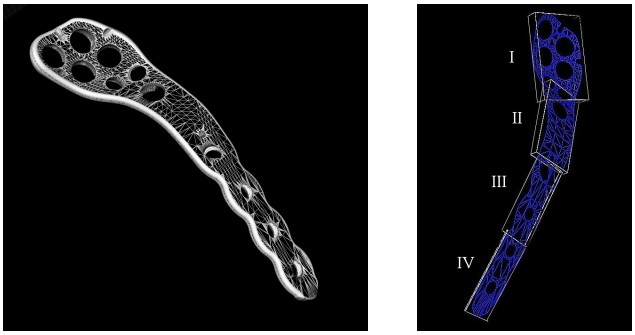

(a) (b)

Fig. 11. (a) The original implant model and the extracted inner surface. (b) The Oriented-Bounding-Boxes of the implant.

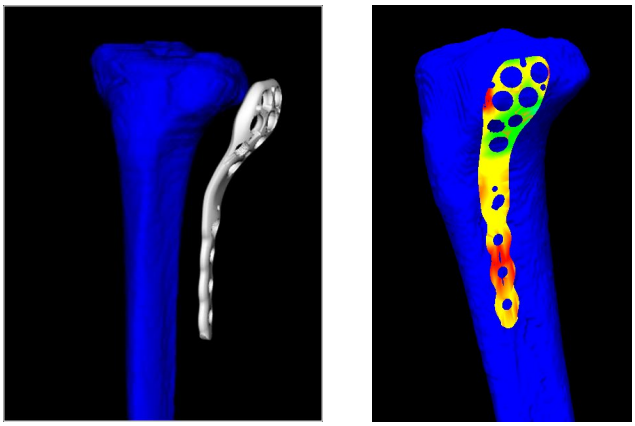

(a) (b)

Fig. 12. (a) The initialisation of the implant fitting. (b) The final result of the implant fitting shows the distance map of fitting error, where the red colour represents the perfect fitting of the implant to the bone and the green colour represents the distance of 3mm to the bone surface.

Figure 12 shows the initialisation step of the automatized implant fitting procedure and the final result of the fitting where the colour map of the implant represents the distance map of the fitting error.

4.3 Population-based evaluation of implant fitting

We apply our method to evaluate the performance of orthopaedic implants, used for internal fracture fixation of the proximal tibia, within the context of the PCA space level set evaluation framework described in Section 3, to optimise the implant shape as to fit a majority of the target population. Figure 13 illustrates the complete procedure.

We present results obtained from a training set of tibia surface models extracted from CT data. The training set consists of 92 left human tibiae from which Asian, Caucasian, male and female are equally present. Statistical shape modeling was then performed, as explained in Section 2. We retain the first five principal components, which account for 92% of shape variability in the population. Using more than five modes to explain the statistical model would give us more subtle changes which, however, do not bring modifications in the area of implant placement (Figure 14). We define the mapping transformation $\mathcal{M}$ as the mean error distance from 844 points sampled on the implant surface to their corresponding best fitted points on the bone surface. The PCA shape space is then built by sampling the space of shape coefficients, generating the corresponding shape, and then computing the mapping $\mathcal{M}$ to obtain the measure of interest. We use the range $-3 \leq \alpha_i \leq 3$ for every shape coefficient. This accounts for 99.7% of the shape variability encompassed in each principal component.

We start with a low resolution sampling of the PCA space, namely a sampling step of $\Delta\alpha_i = 0.1$ for each principal component (which would result in a map of 60x60 instances if two principal components were retained). We initialise the zero level set by applying seed initialisation on the PCA shape space, and then proceed with the hierarchical narrow-band zero level set evolution, as explained in Section 3.4. We do not need to explicitly generate all instances and compute mean error fitting for every point in the shape space, but only in the narrow band around the evolving zero level set. We continue with a hierarchical narrow band approach by reducing successively the sampling step to $\Delta\alpha_i = 0.05$ and $\Delta\alpha_i = 0.025$, respectively. As sampling resolution increases, the narrow band level set approach becomes mandatory to decrease high computation times and to reduce the search space of shape parameters.

For the given implant fitting problem, which includes space optimisation, instance creation and fitting without manual initialisation, we need less than

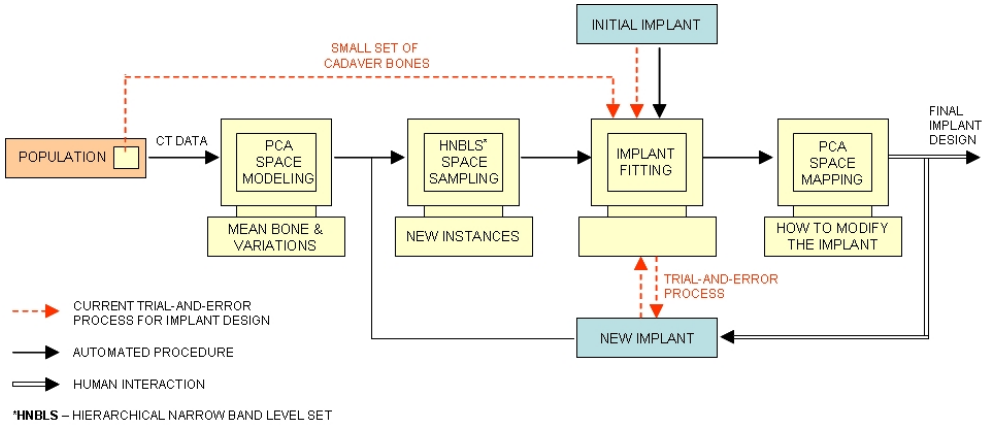

Fig. 13. The complete procedure of the implant design process. First a statistical model of a given population is computed and new instances are created (PCA shape space modeling and HNBLS 'sampling'). The implant is automatically fitted to these virtual bones and a mean error of the fitting is computed (PCA shape space 'mapping'). Using hierarchical narrow band level set segmentation 'good' fitted bones are selected. From observing the selected instances and the fitting results on them we could propose modifications to the implant shape. Finally, we repeat the process of automatized fitting for new implant to verify its performance.

1 minute per bone (Dual CPU @2.2 GHz, RAM 2GB). In combination with hierarchical narrow band and a given segmentation criterion the fitting process was performed on 1'504, 1'168 and 4'904 instances, respectively for the 3 resolution levels mentioned above. This results in reducing the computation to only 13.5% of the whole shape space (i.e. 57'600 instances), which drastically reduces the computation time.

The segmented areas in Figure 15a represent the range of parametric values that generate tibia shapes satisfying the segmentation criterion that was provided as a requirement from the implant designer, i.e. mean fitting error of less than 1mm. The 2D shape space map is built using 2 principal components, u_1 and u_2, in order to illustrate the strong effect of the first principal component for the implant shape design. Figure 15b shows an example of a construction of a 3D PCA shape space (i.e. using 3 principal components to generate the shape instances) and the result of the level set optimisation for the fitting error less then 1mm. It can be visualised that the first and fifth PCs have higher influence on the implant shape design, whereas the second PC does not interfere much as it covers the whole space $-3\sqrt{\lambda_2}, +3\sqrt{\lambda_2}$. We decided to exclude principal components u_3 and u_4 since their variations do not affect the bone in the area of the implant placing (Figure 14).

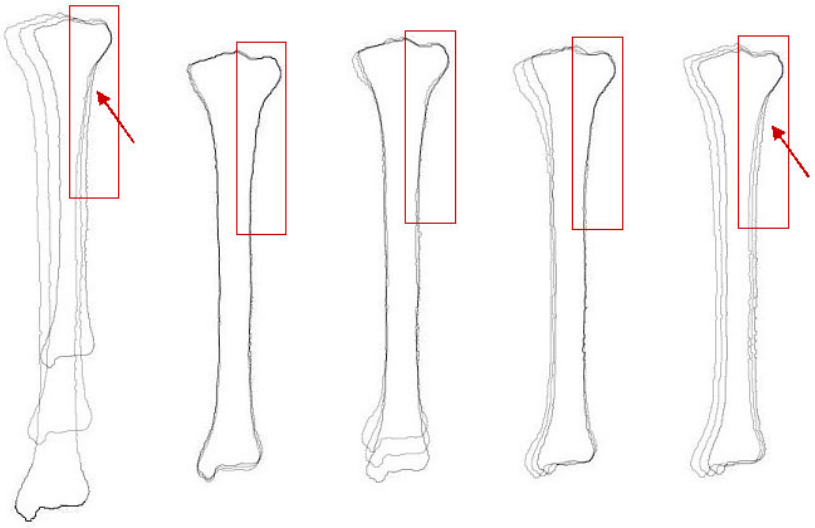

Fig. 14. The first five modes of variation for the left human tibia (anterior view) are visualised. For each principal component, we show $\bar{m} - 3\sqrt{\lambda_i u_i}$, $\bar{m}$, and $\bar{m} + 3\sqrt{\lambda_i u_i}$, where $\bar{m}$ represents the mean bone. The arrows point to the area of implant placement, which is most affected by the first and fifth principal component. The first five modes account for 71, 11, 6, 3 and 1% of shape variability in the population, respectively.

4.4 Implant design modifications and analysis

Implant design modifications followed analysis of the segmented spectrum of shapes. It can be seen in Figure 15a that the result of the fitting depends mostly on the first principal component, as the segmented area falls in the negative values of u_1. Since the negative values of the first principal component favor 'good' fitting, this leads to the conclusion that the implant works better for longer bones. Our aim is to optimise the fitting as to cover the whole shape space, i.e. the majority of the population. Having the measure of variation between positive and negative values for the first principal component (Figure 14), it can be concluded that changes of the length of the tibia affect the result of the fitting, since these changes affect as well changes of the oblique line of the tibia and a slight torsion of the lateral surface of the tibia. In other words, the analysis allows us to conclude that the angles and curvatures in the first and second OBB of the implant geometry (Figure 11b) as well as the curvature of the third and fourth OBB of the implant are responsible for the fitting.

In agreement with the previous conclusions we proceed with the optimisation of implant shape design by applying the following modifications. First, we decrease the angles and flatten the curvatures in the first and second bounding box of the implant (Figure 16a), to follow the oblique line of the mean tibia

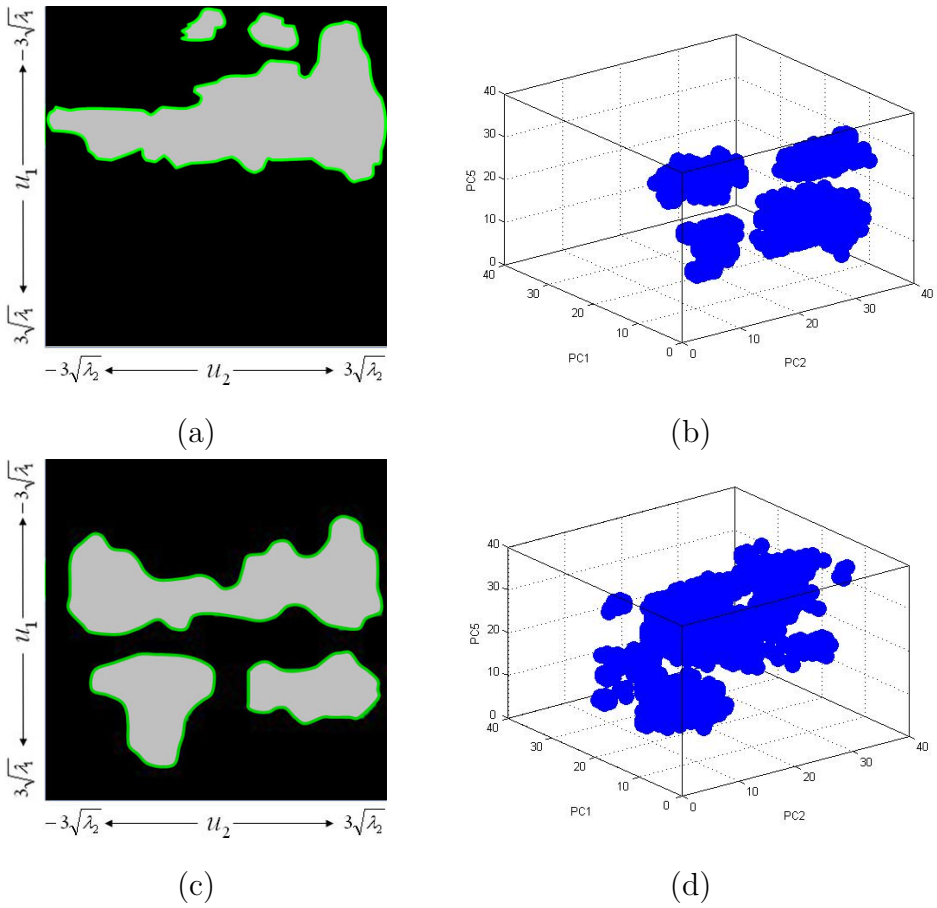

Fig. 15. (a) Automatic hierarchical 2D level set segmentation gives the spectrum of shapes that have fitting error less than 1mm for the implant given by the manufacturer. (b) 3D level set segmentation gives the spectrum of shapes that have the fitting error less than 1mm for the implant given by the manufacturer. (c) Spectrum of shapes that have fitting error less than 1mm for the modified implant design. (d) 3D level set segmentation for the modified implant design.

bone (Figure 14). In addition, we follow the distance dimensions between bone head and implant from the implant fitting distance map (Figure 12b). We apply further modifications to the implant shape by increasing the torsion of the distal part of the plate. We rotate the third and fourth bounding box along the center of the plate to bring the left anterior edge of the implant closer to the lateral surface of tibia (Figure 16b).

To evaluate the new design we perform a re-fitting in the PCA shape space using the modified implant shape. The results of the segmented space are shown in Figures 15c and 15d. It can be seen that the modified implant expands the space of segmented bones by covering different shape variability and therefore fits better to the majority of the population. With the new implant design we found that there is an increase of 40% on the number of instances that satisfy the given fitting criterion.

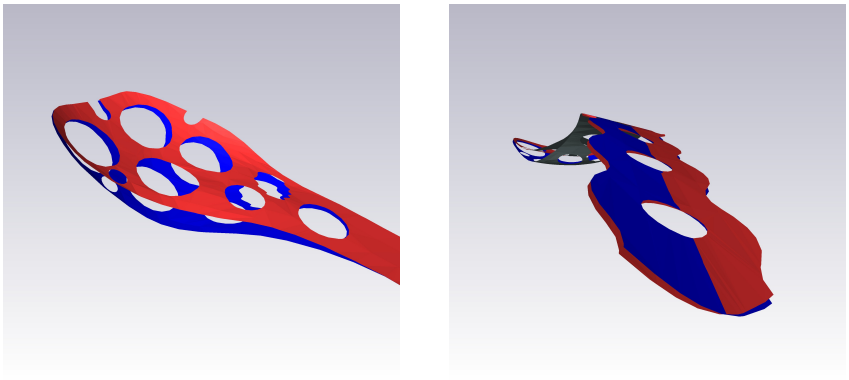

(a) (b)

Fig. 16. New implant design (in red). (a) Curvatures in the first and second bounding box of implant are flattened. (b) Implant surface in the third and fourth Oriented-Bounding-Box is twisted inside.

5 Conclusions

In this paper we have presented a methodology for the evaluation of a functional criterion (that could represent an anatomical/physiological measure) across a target population. Our framework is based on building a statistical model via PCA and finding the region of the parametric space defined by the principal component weights that matches the criterion. The mechanism to search for this partition is based on a level set evolution in parametric space, optimised via a multi-level narrow-band approach for computational efficiency. To our knowledge, this is the first work that tackles the issue of finding a partition of PCA space based on a criterion, and the first time that level sets are used within this context. Existing previous works combining PCA and level sets used the later as a shape representation, and evolve the level set in image space. This is fundamentally different to our work.

Current evaluation and optimisation of orthopaedic implants is done by manual fitting and fixation procedures, applied on a small set of cadaver bones in a trial-and-error process. The method that we propose allows to virtually test the implants on a representative set of bones generated by sampling the statistical model. Using level sets a spectrum of shapes is segmented in the PCA shape space, based on a given fitting criterion. By correlating the principal components of the selected instances to the given implant geometry the modifications to the implant design/geometry can be assessed directly from the segmented map. The proposed method highlights which patterns of bone variability are more important for implant fitting, allowing and easing implant design improvements, as to fit a maximum of the target population. A hierarchical narrow band approach is used to avoid exhaustive search of the instances in the high resolution space, and to search for the instances only in

the neighborhood of the zero level set and not in the whole shape space.

To our knowledge this is the first research into the problem of estimating how a given implant fits to the wide population and how the morphological features in implant design can be improved. The practical use of the proposed concept are of great importance for the implant manufacturer, due to the huge potential benefits in terms of patient satisfaction and financial gains in this high-volume market. Further validation of the method is ongoing work.

Future work will include automatic correlation of the principal components to the given implant geometry, so that the modifications to the implant design/geometry could be assessed directly from the segmented map and automatically proposed. A parametric model for the implant design could be established, including design parameters such as diameters, lengths, positions of the holes, etc. Such parameters could be automatically optimised by maximising the fitted volume in the PCA space.

Furthermore, we intend to include the application of the proposed method to bone implant fitting assessment taking into account shape and biomechanical properties. A combined shape and intensity statistical bone model will be built, and the intensity values, which are linked to bone density, will be used to do a finite element analysis of the performance of the implant (Belenquer et al., 2006), which will be used as the criterion to be evaluated in the level set evolution.

6 Acknowledgments

This research has been supported by the Swiss National Science Foundation through its National Center of Competence in Research (NCCR) on Computer Aided and Image Guided Medical Interventions (Co-Me) and the Swiss Innovation Promotion Agency, CTI (7961.2 LSPP-LS).

References

Adalsteinsson, D., Sethian, J. A., 1995. A fast level set method for propagating iterfaces. J. Comput. Phys. 118, 269–277.

Belenquer, L., Büchler, P., Rückert, D., Nolte, L. P., Ballester, M. A. G., 2006. Statistical finite element model for bone shape and biomechanical properties. In: Proc. MICCAI. pp. 405–411.

Besl, P., McKay, N., 1992. A method for registration of 3D shapes. IEEE Trans. Pattern. Anal. Mach. Intell. 14 (2), 239–256.

509 Bishop, C., 1995. Neural Networks for Pattern Recognition. Oxford University
510 Press, USA.
511 Bresson, X., Vandergheynst, P., Thiran, J. P., 2006. Variational model for
512 object segmentation using boundary information and shape prior driven by
513 the Mumford-Shah functional. Int. J. Comput. Vis. 68 (2), 145–162.
514 Chan, T. F., Vese, L. A., 2001. Active contours without edges. IEEE Trans.
515 Image Process. 10 (2), 266–277.
516 Chen, X. F., Guan, Z. C., 2004. Image segmentation based on Mumford-Shah
517 functional. J. Zhejiang University SCIENCE 5, 123–128.
518 Chen, Y., Tagare, H. D., Thiruvenkadam, S., Huang, F., Wilson, D., Gopinath,
519 K. S., Briggs, R. W., Geiser, E. A., 2002. Using prior shapes in geometric
520 active contours in a variational framework. Int. J. Comput. Vis. 50 (3),
521 315–328.
522 Cootes, T. F., Marsland, S., Twining, C. J., Smith, K., Taylor, C. J., 2004.
523 Groupwise diffeomorphic non-rigid registration for automatic model build-
524 ing. In: Proc. 8th European Conference on Computer Vision. Vol. 4. pp.
525 316–327.
526 Cootes, T. F., Taylor, C. J., 2004. Anatomical statistical models and their role
527 in feature extraction. Br. J. Radiol. 77, S133–S139.
528 Cootes, T. F., Taylor, C. J., Cooper, D. H., Graham, J., 1995. Active shape
529 models - their training and applications. Comput. Vis. Image Underst.
530 61 (1), 38–59.
531 Cremers, D., 2006. Dynamical statistical shape priors for level set based track-
532 ing. IEEE Trans. Pattern. Anal. Mach. Intell. 28 (8), 1262–1273.
533 Davies, R. H., Twining, C. J., Cootes, T. F., Waterton, J. C., Taylor, C. J.,
534 2002. A minimum description length approach to statistical shape modeling.
535 IEEE. Trans. Med. Imaging. 21 (5), 525–537.
536 Evans, L. C., Gariepy, R. F., 1992. Measure Theory and Fine Properties of
537 Functions. Boca Raton, FL: CRC.
538 Frankel, V. H., Burstein, A. H., 1970. Orthopaedic Biomechanics. Lea and
539 Febiger, Philadelphia.
540 González, M. A., Pennec, X., Linguraru, M., Ayache, N., 2004. Generalised
541 image models and their application as statistical models of images. Med.
542 Image Anal. 8 (3), 361–369.
543 Goyal, K. S., Skalak, A. S., Marcus, R. E., Vallier, H. A., Cooperman, D. R.,
544 2007. Analysis of anatomic periarticular tibial plate fit on normal adults.
545 Clin. Orthop. Relat. Res. 461, 245–457.
546 Haslam, J., Taylor, C. J., Cootes, T. F., 1994. A probabalistic fitness mea-
547 sure for deformable template models. In: Proc. 5th British Machine Vision
548 Conference. pp. 33–42.
549 Heimann, T., Meinzer, H. P., 2009. Statistical shape models for 3d medical
550 image segmentation: A review. Med. Image Anal. 13, 543–563.
551 Hill, A., Cootes, T. F., Taylor, C. J., Lindley, K., 1994. Medical image inter-
552 pretation: A generic approach using deformable templates. Inform. Health.

Soc. Care. 19 (1), 47–59.
Kass, M., Witkin, A., Terzopoulos, D., 1987. Snakes: Active contour models. Int. J. Comput. Vis. 1 (4), 321–331.
Kozic, N., González, M. A., Tannast, M., Nolte, L. P., Reyes, M., 2008. Statistical shape space analysis based on level sets. In: Proc. Medical Imaging and Augmented Reality. pp. 160–167.
Leventon, M. E., Faugeras, O., Grimson, W. E. L., Wells, W. M., 2000. Level set based segmentation with intensity and curvature priors. In: Proc. IEEE Workshop on Mathematical Methods in Biomedical Image Analysis. pp. 4–11.
McInerney, T., Terzopoulos, D., 1996. Deformable models in medical images analysis: A survey. Med. Image Anal. 1 (2), 91–108.
Mitchell, S. C., Lelieveldt, B. P., van der Geest, R., Schaap, J., Reiber, J. H., Sonka, M., 2000. Segmentation of cardiac MR images: An active appearance model approach. In: Proc. SPIE Medical Imaging. Vol. 3979. pp. 224–234.
Mumford, D., Shah, J., 1989. Optimal approximation by piecewise smooth functions and associated variational problems. Commun. Pure and Applied Mathematics 42, 577–685.
Osher, S., Sethian, J. A., 1988. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulation. J. Comput. Phys. 79, 12–49.
Rajamani, K. T., Styner, M. A., Talib, H., Zheng, G., Nolte, L. P., González, M. A., 2007. Statistical deformable bone models for robust 3D surface extrapolation from sparse data. Med. Image Anal. 11 (2), 99–109.
Reyes, M., Büchler, P., Nolte, L. P., Reimers, N., Lutz, C., Ballester, M. A. G., 2008. Evidence-based implant design using a statistical bone model and automated implant fitting. In: Proc. CAOS International. pp. 379–381.
Roberts, M. G., Cootes, T. F., Adams, J. E., 2006. Automatic segmentation of lumbar vertebrae on digitised radiographs using linked active appearance models. In: Proc. MIUA. pp. 120–124.
Rousson, M., Paragios, N., Deriche, R., 2004. Implicit active shape models for 3D segmentation in MRI imaging. In: Proc. MICCAI. pp. 209–216.
Rueckert, D., Frangi, A. F., Schnabel, J. A., 2001. Automatic construction of 3D statistical deformation models using non-rigid registration. In: Proc. MICCAI. pp. 77–84.
Rueckert, D., Frangi, A. F., Schnabel, J. A., 2003. Automatic construction of 3D statistical deformation models of the brain using non-rigid registration. IEEE. Trans. Med. Imaging. 22 (8), 1014–1025.
Russell, T. A., Levine, A. M., 1996. Fractures of the tibial diaphysis. Orthopedic knowledge update trauma, 171–179.
Schmutz, B., Wullschleger, M. E., Kim, H., Noser, H., Schütz, M. A., 2008. Fit assessment of anatomic plate for the distal medial plate. Clin. Orthop. Relat. Res. 22, 258–263.
Sierra, R., Zsemlye, G., Szekely, G., Bajka, M., 2006. Generation of variable anatomical models for surgical training simulators. Med. Image Anal. 10 (2),

598 275–285.
599 Smyth, P. P., Taylor, C. J., Adams, J. E., 1996. Automatic measurement of
600 vertebral shape using active shape models. In: Proc. 7th British Machine
601 Vision Conference. pp. 705–714.
602 Tsai, A., Yezzi, A., Willsky, A. S., 2001. Curve evolution implementation of the
603 Mumford-Shah functional for image segmentation - denoising, interpolation,
604 and magnification. IEEE Trans. Image Process. 10 (8), 1169–1186.
605 van Ginneken, B., Frangi, A. F., Stall, J. J., ter Haar Romeny, B. M.,
606 Viergever, M. A., 2002. Active shape model segmentation with optimal fea-
607 tures. IEEE Trans. Med. Imaging. 21 (8), 924–933.

Figure1

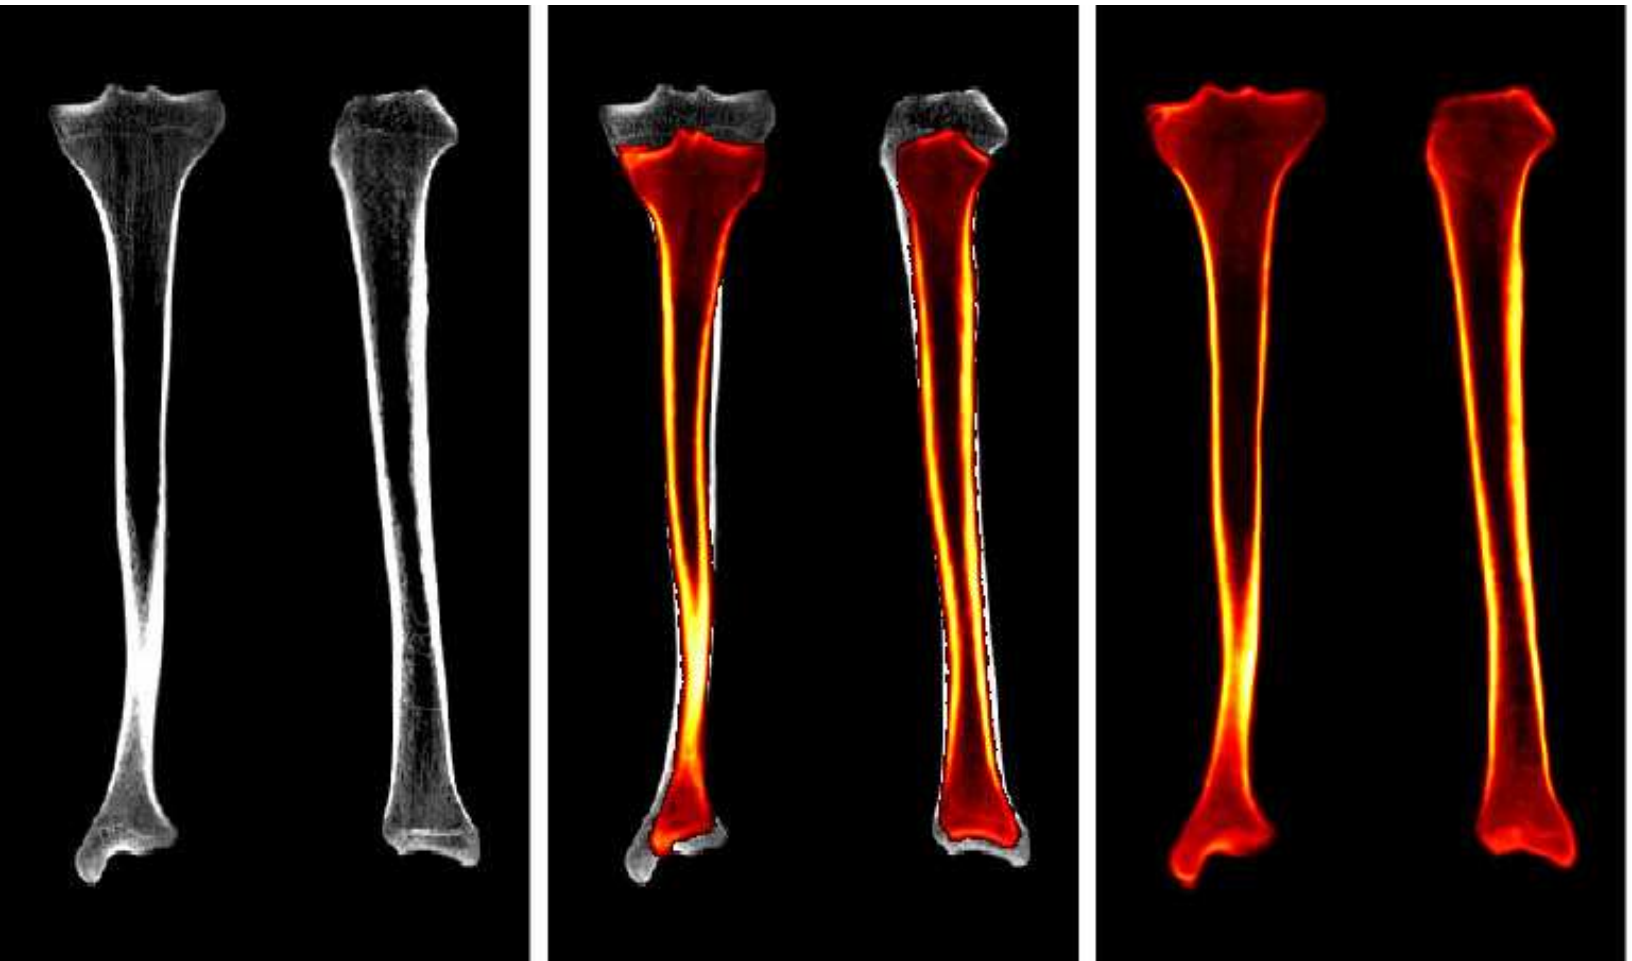

Figure2

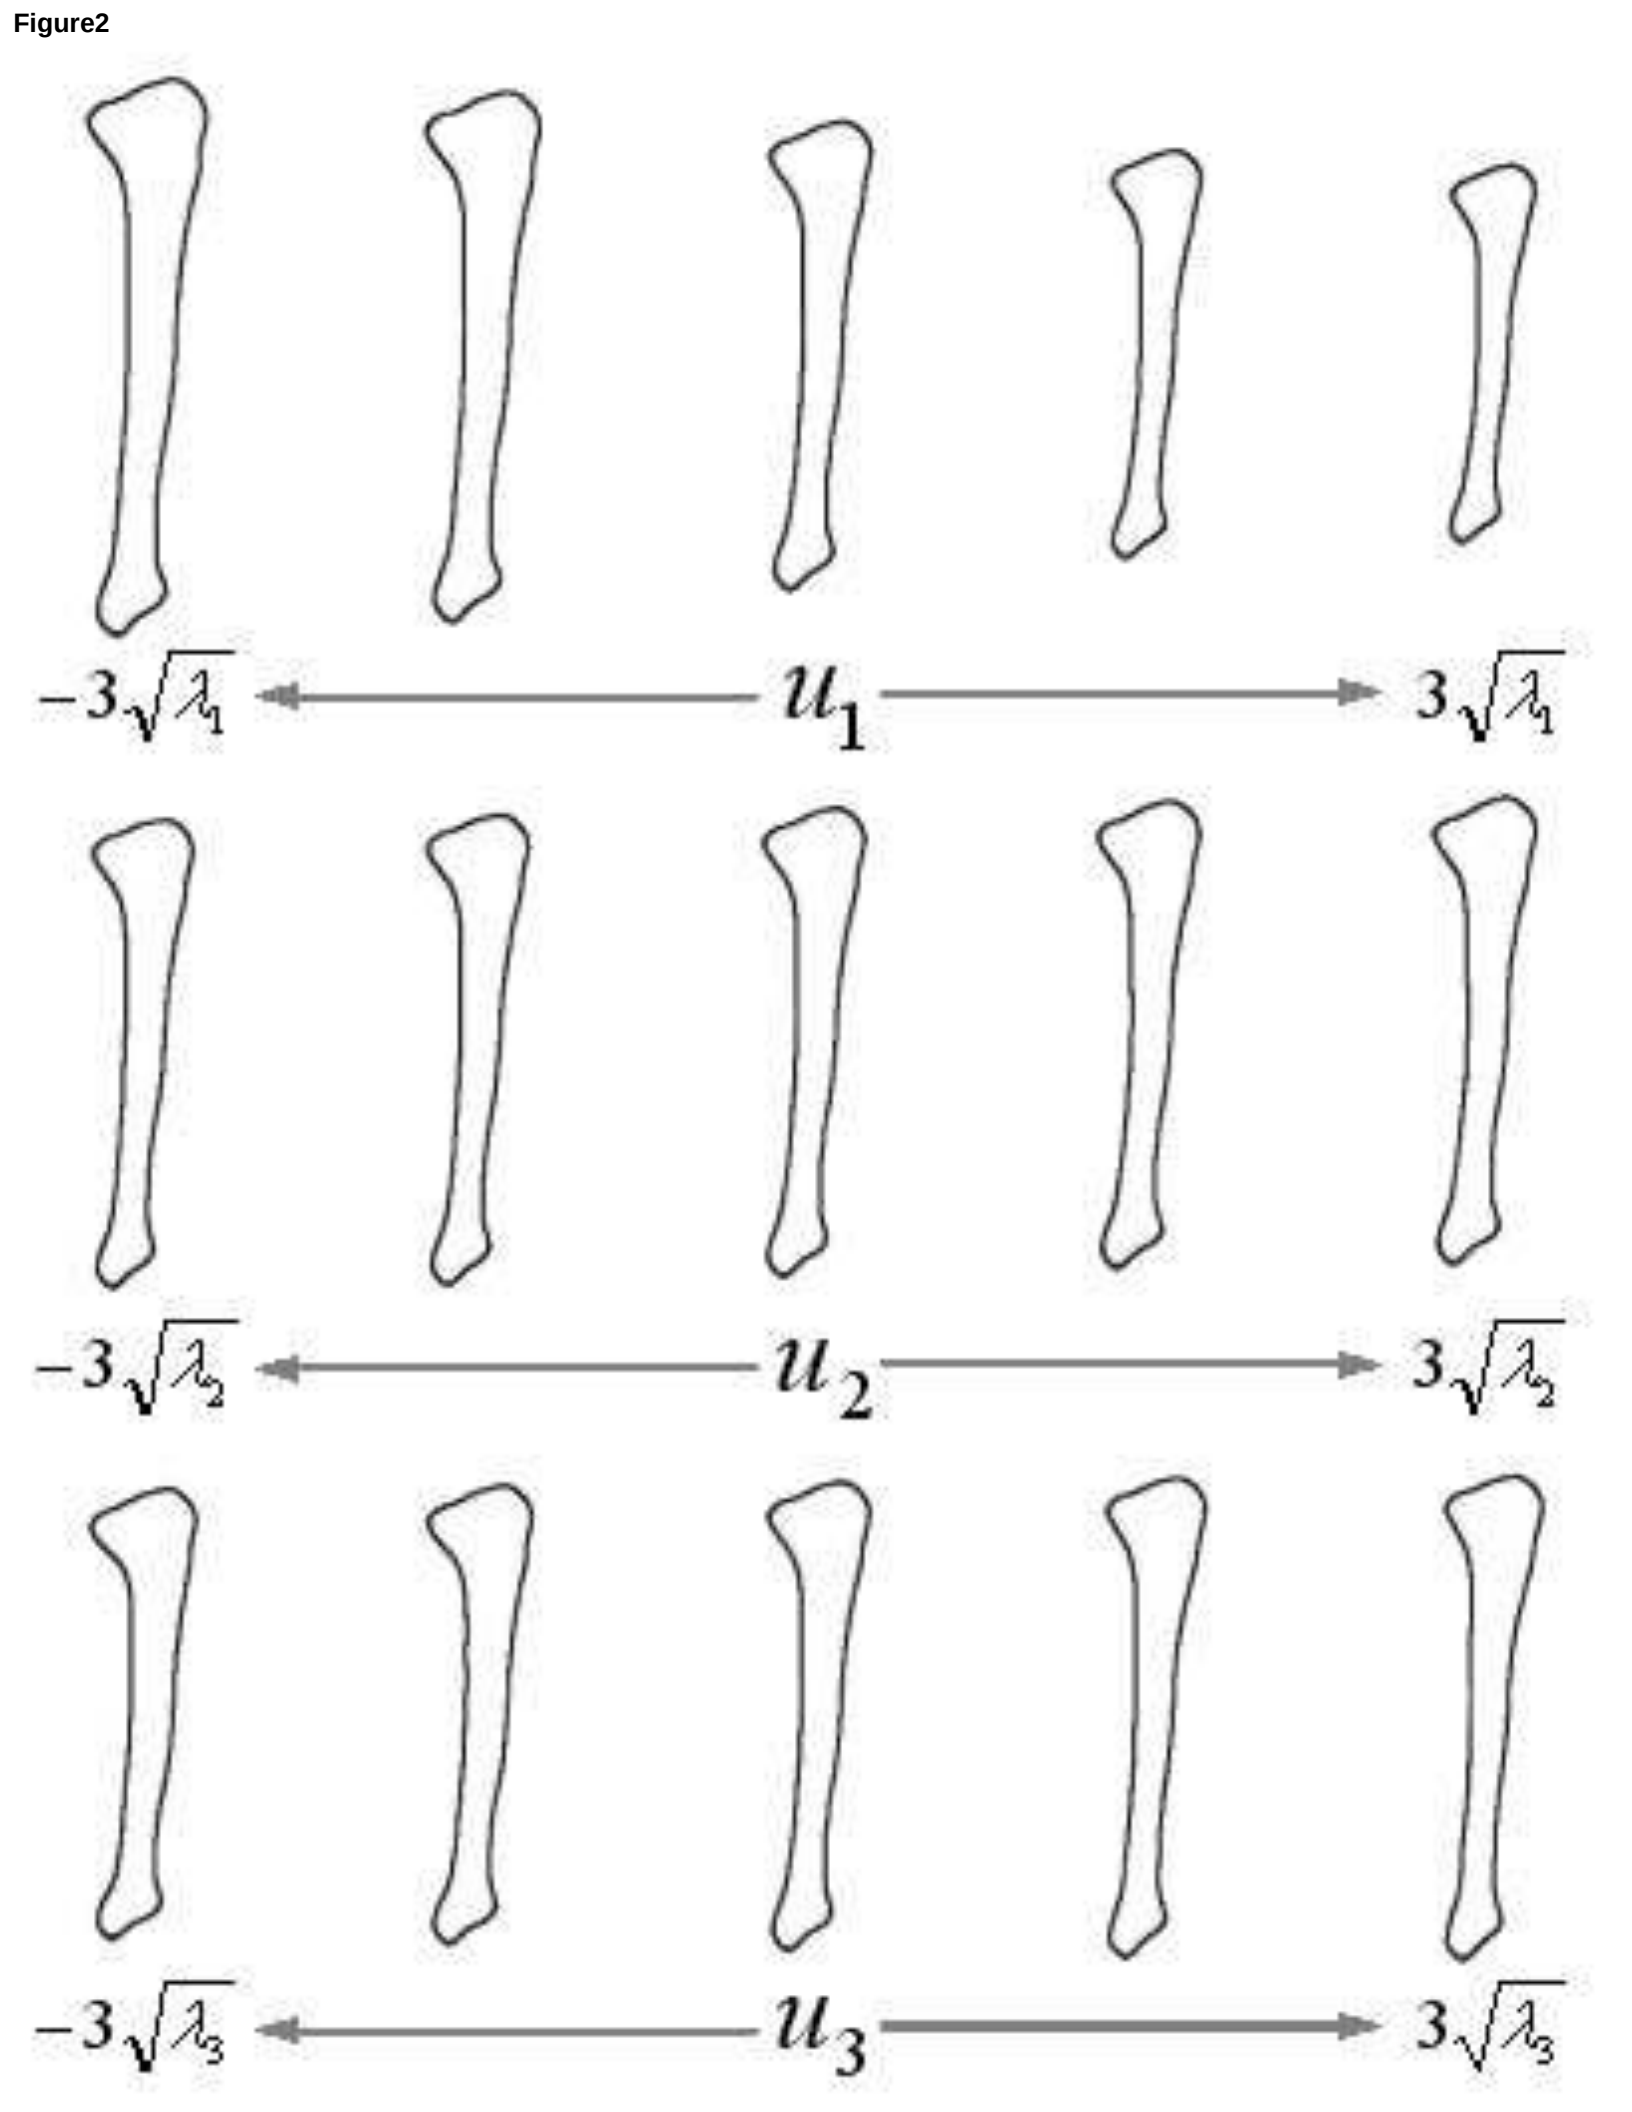
$-3\sqrt{\lambda_1}$
u_1
$3\sqrt{\lambda_1}$
$-3\sqrt{\lambda_2}$
u_2
$3\sqrt{\lambda_2}$
$-3\sqrt{\lambda_3}$
u_3
$3\sqrt{\lambda_3}$

Figure3

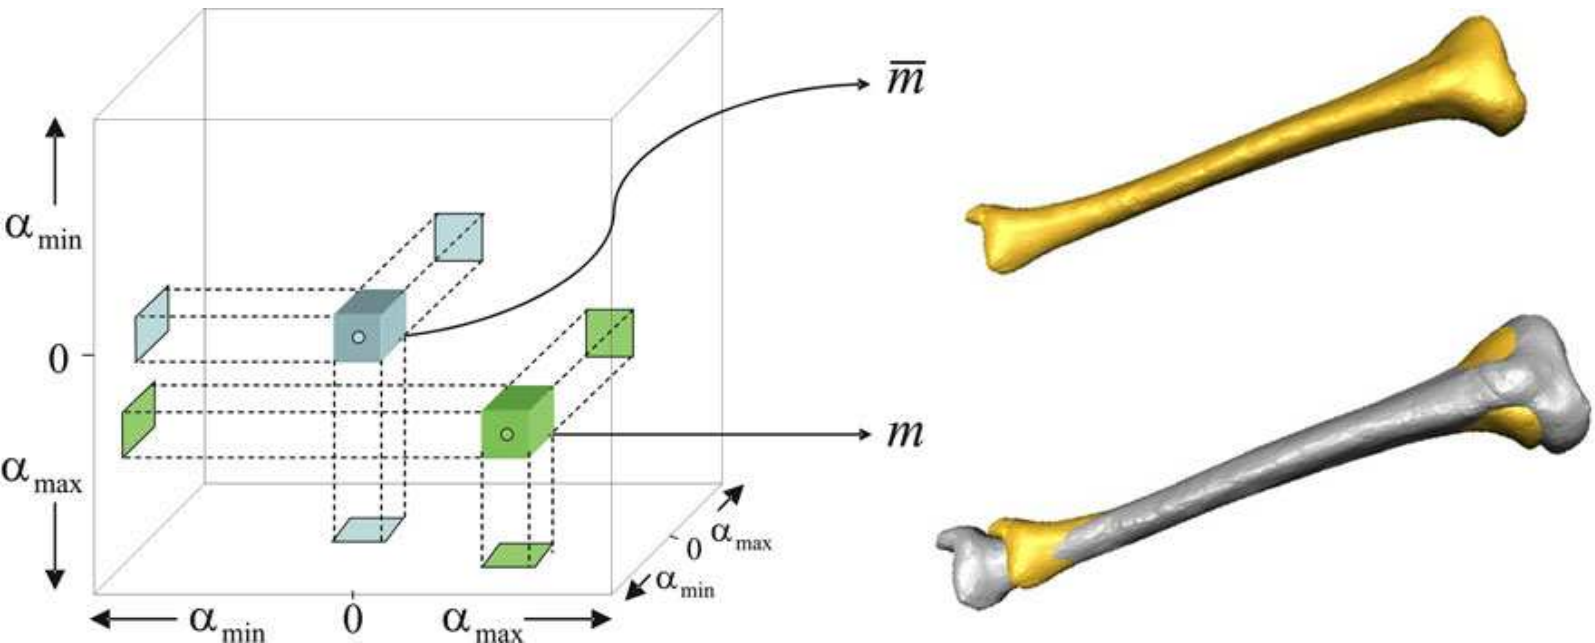
$\overline{m}$
m
α_{min}
0
α_{max}
α_{min}
0
α_{max}
α_{min}
0
α_{max}

Figure4

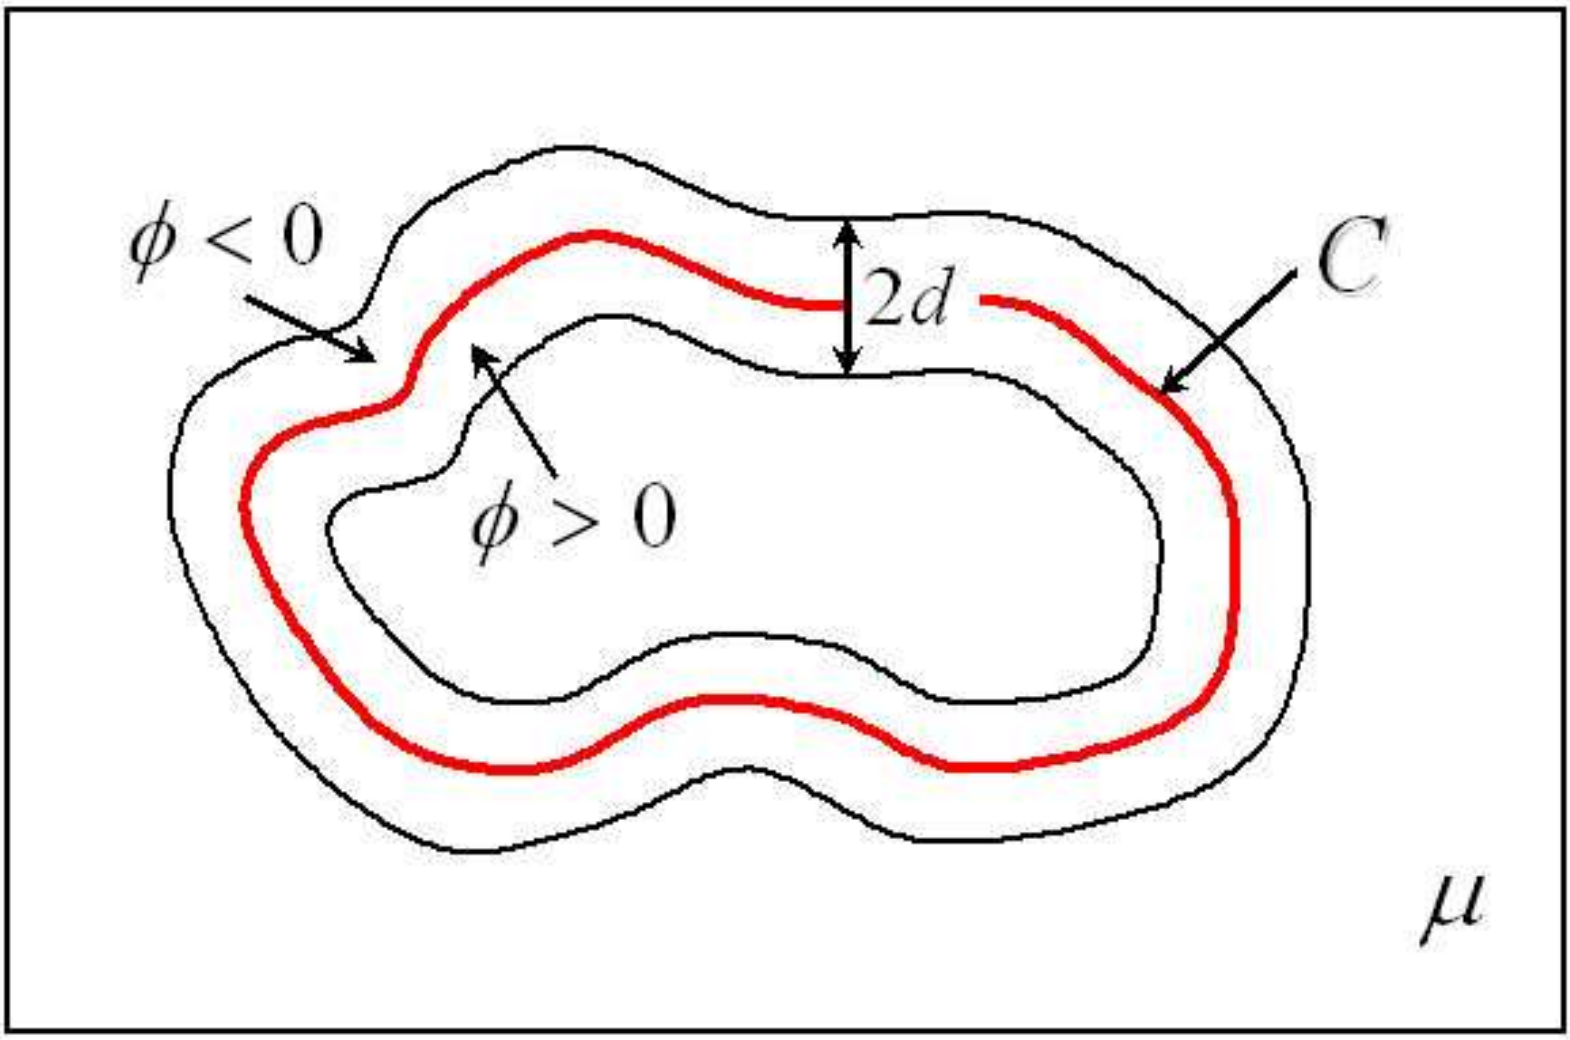

Figure5

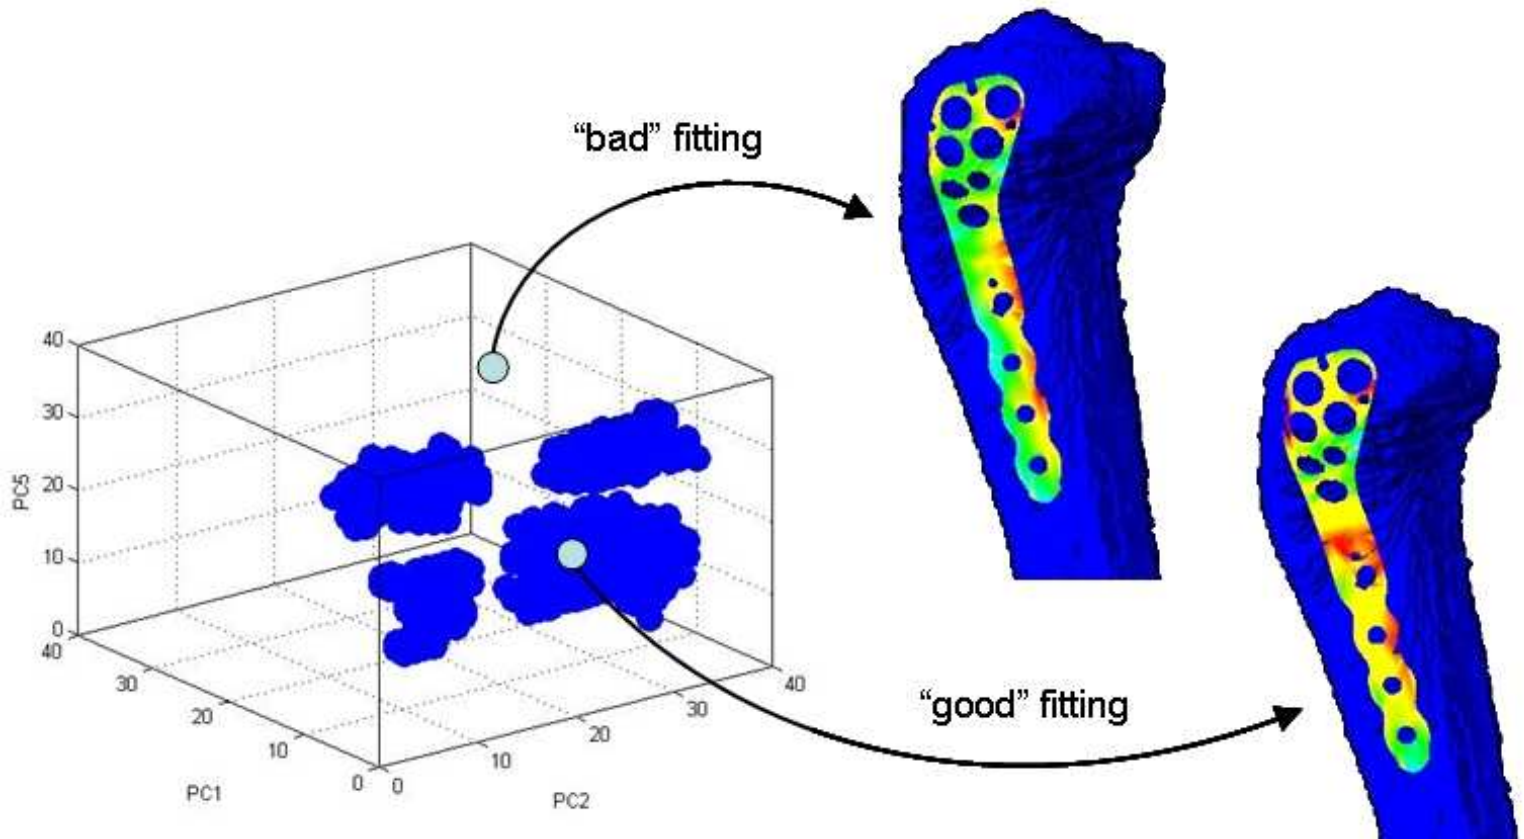

Figure6a

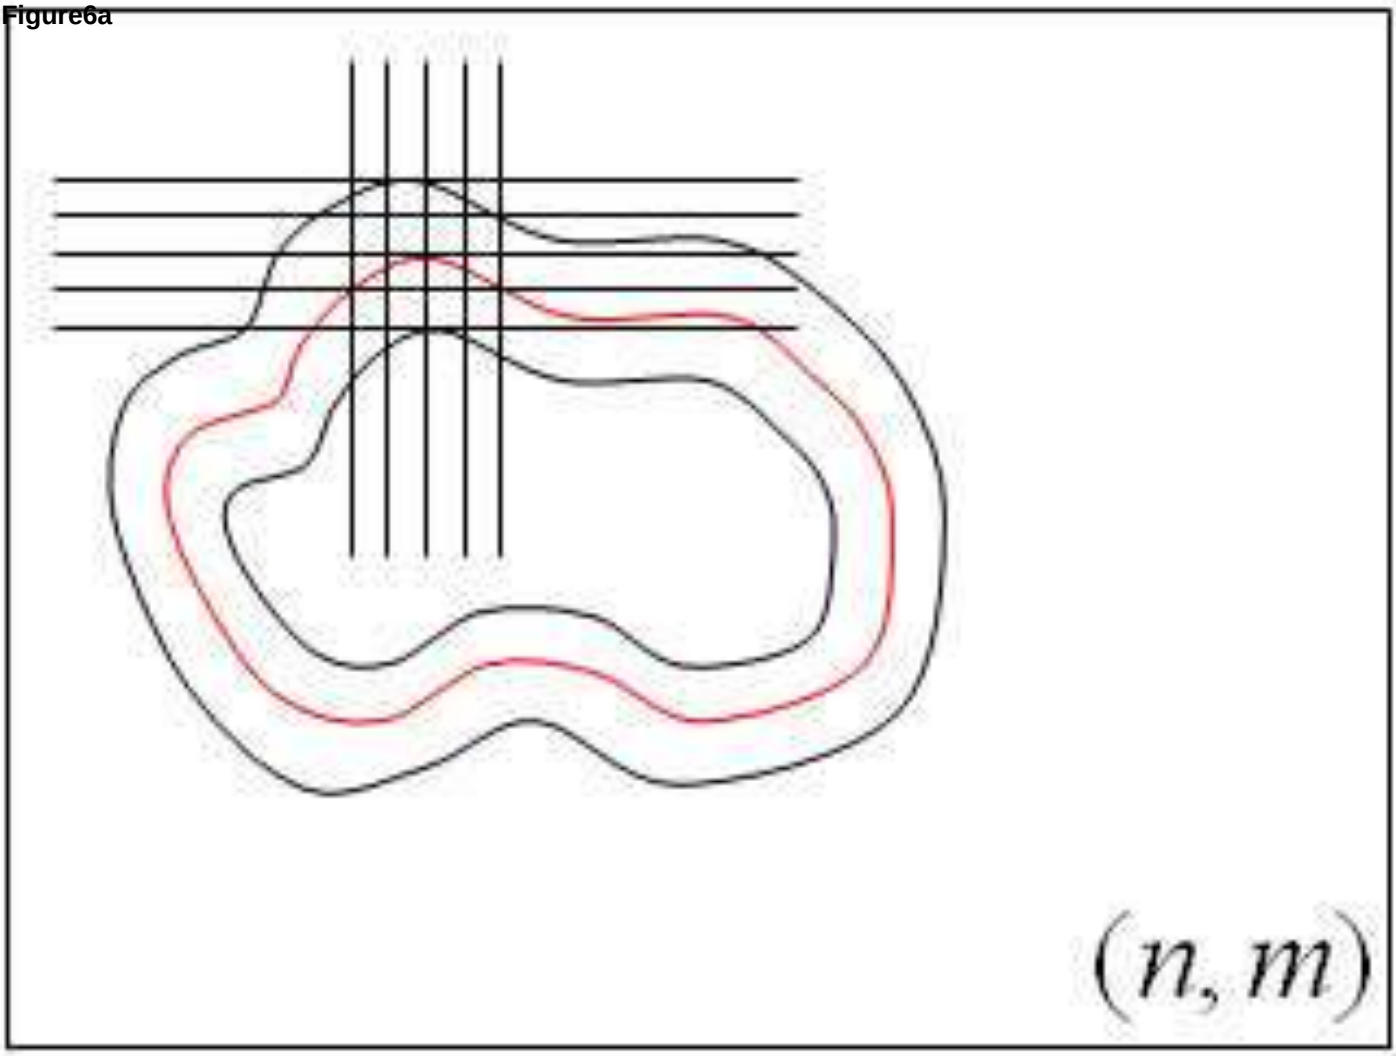

Figure6b

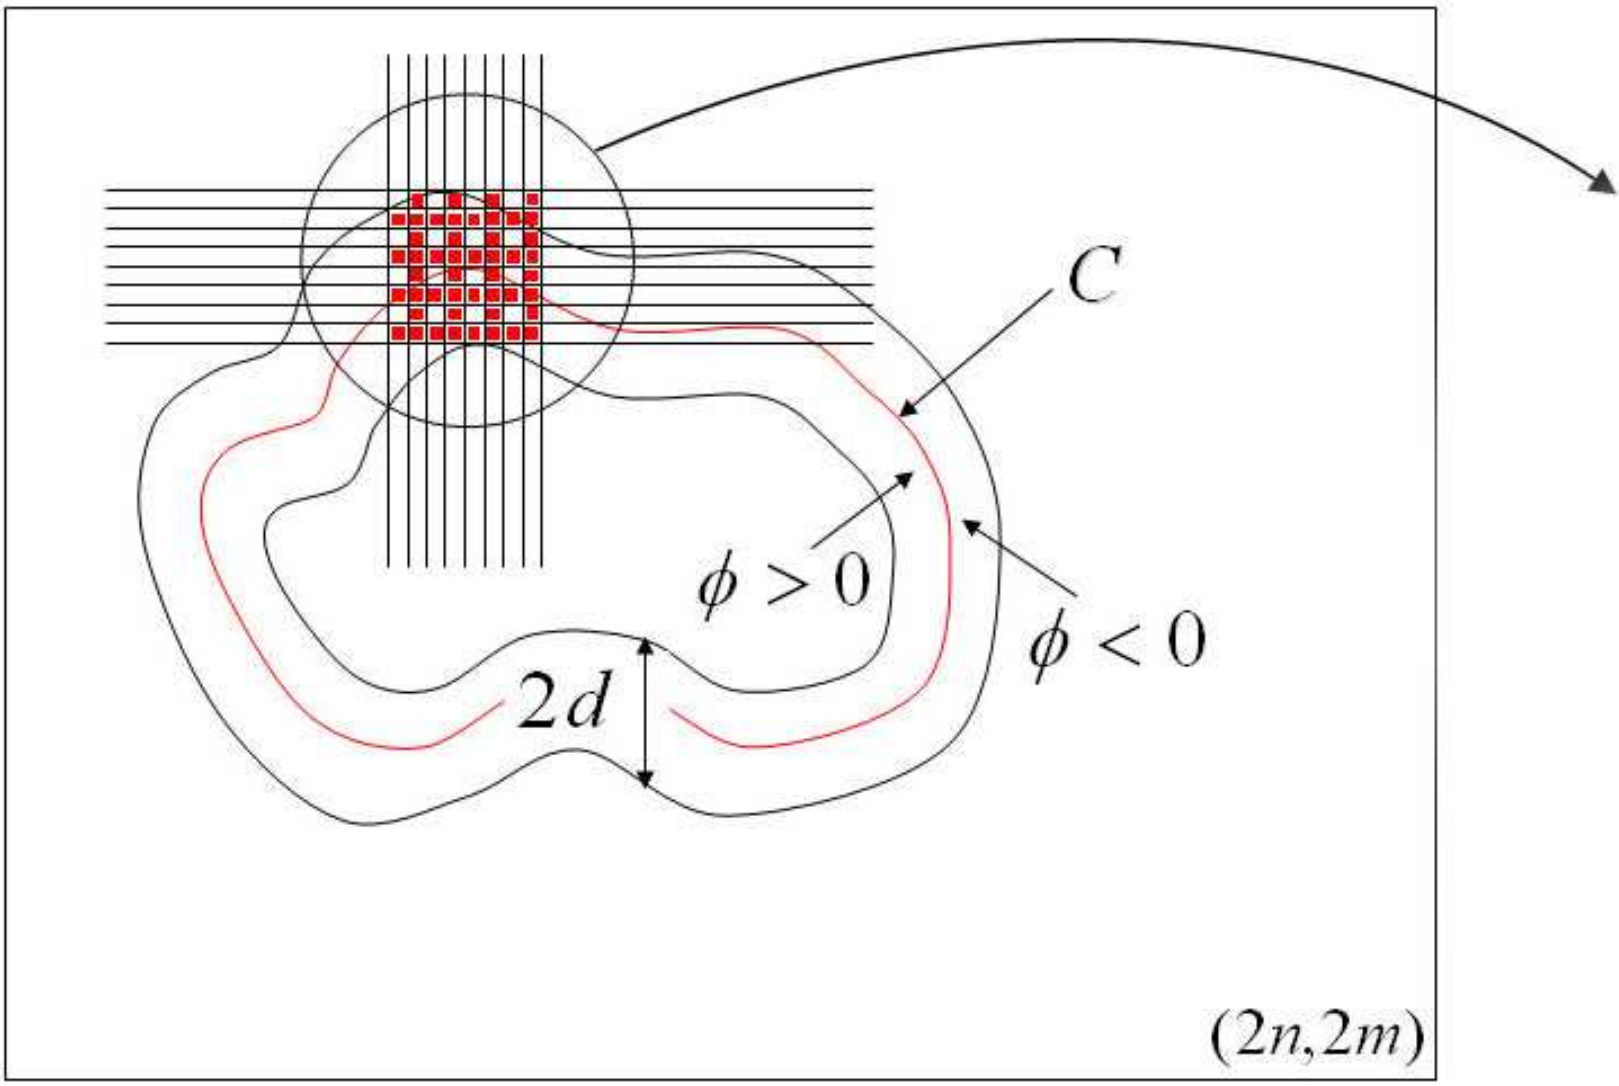

Figure6c

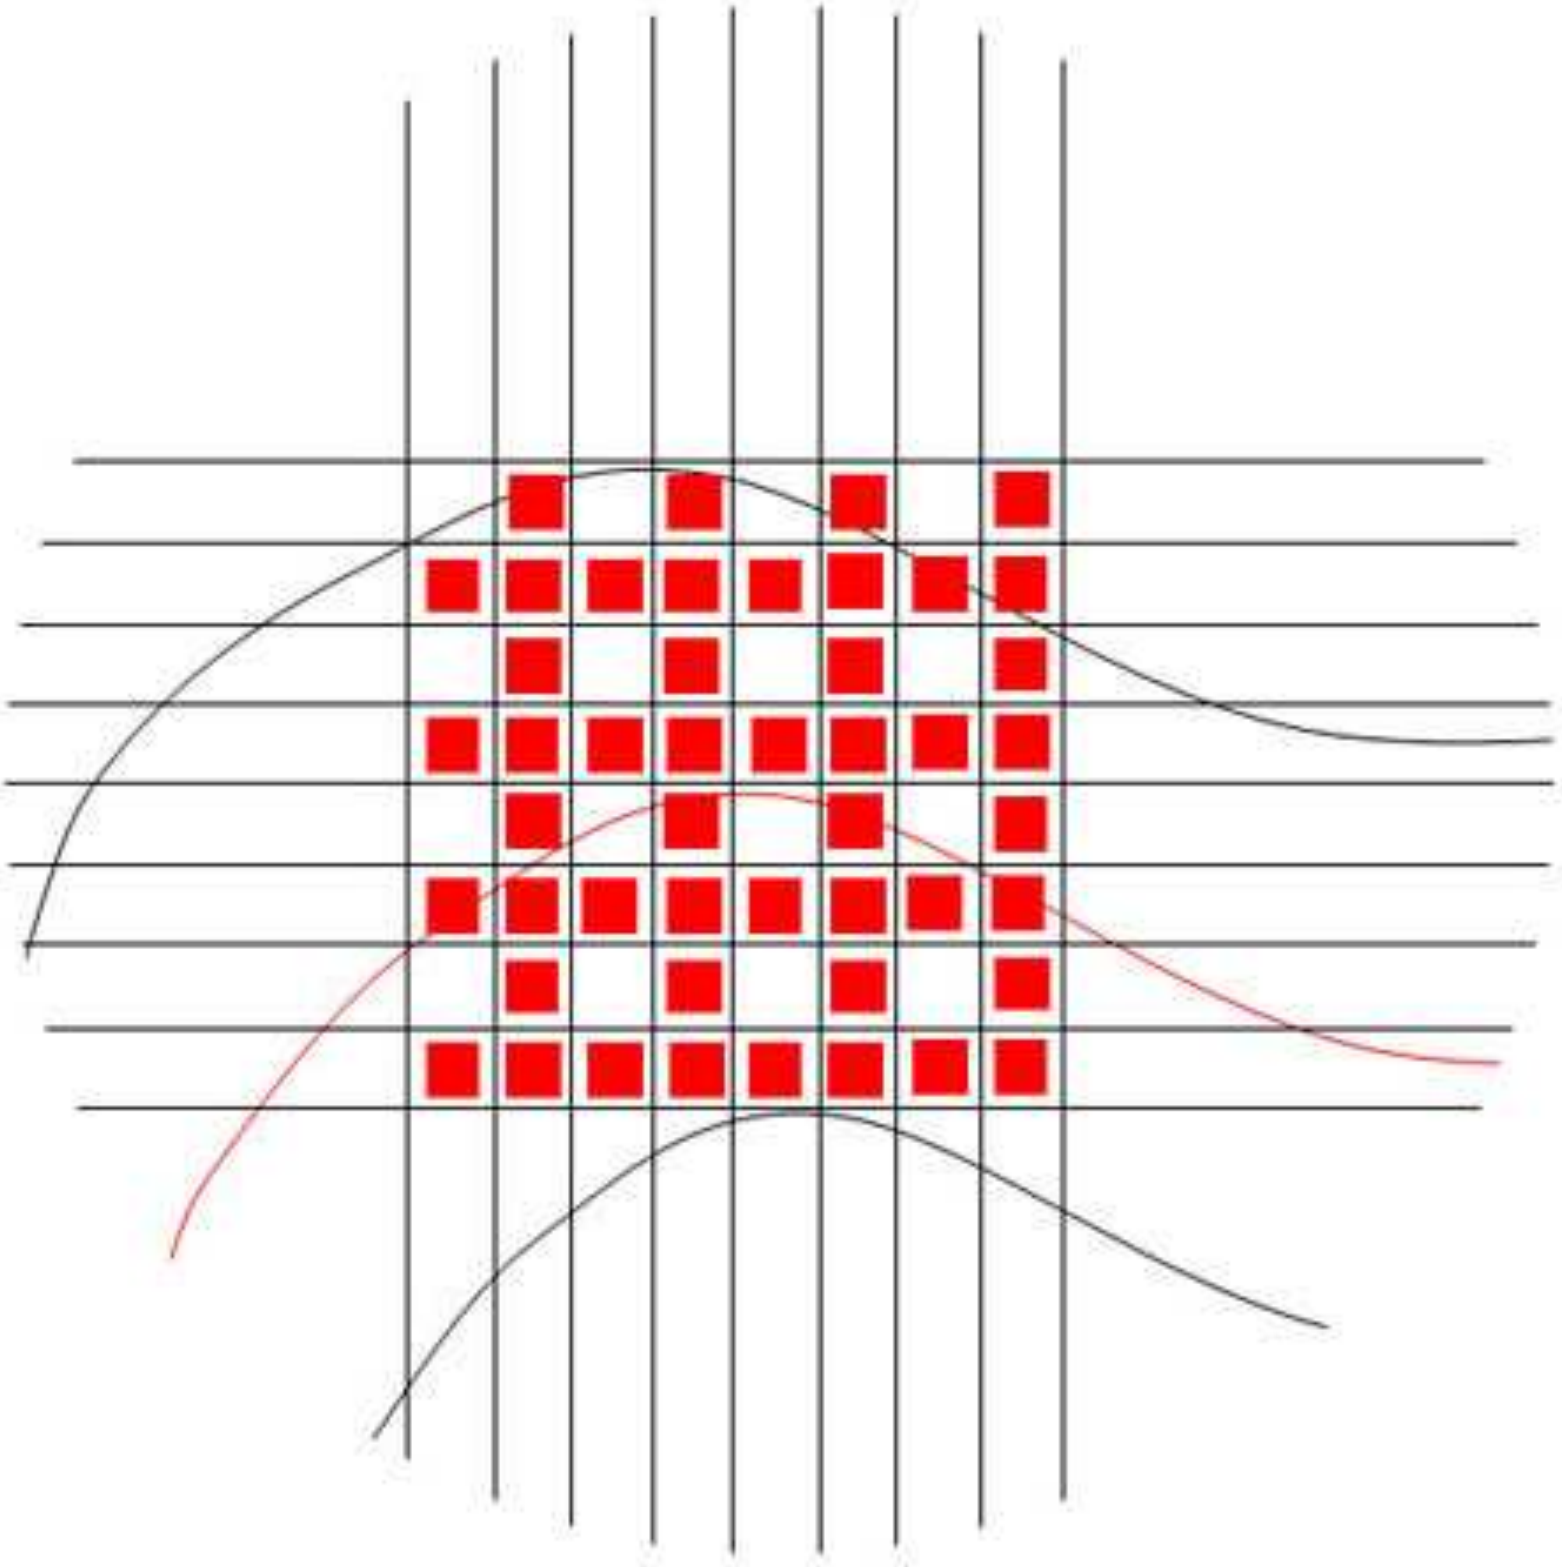

Figure7

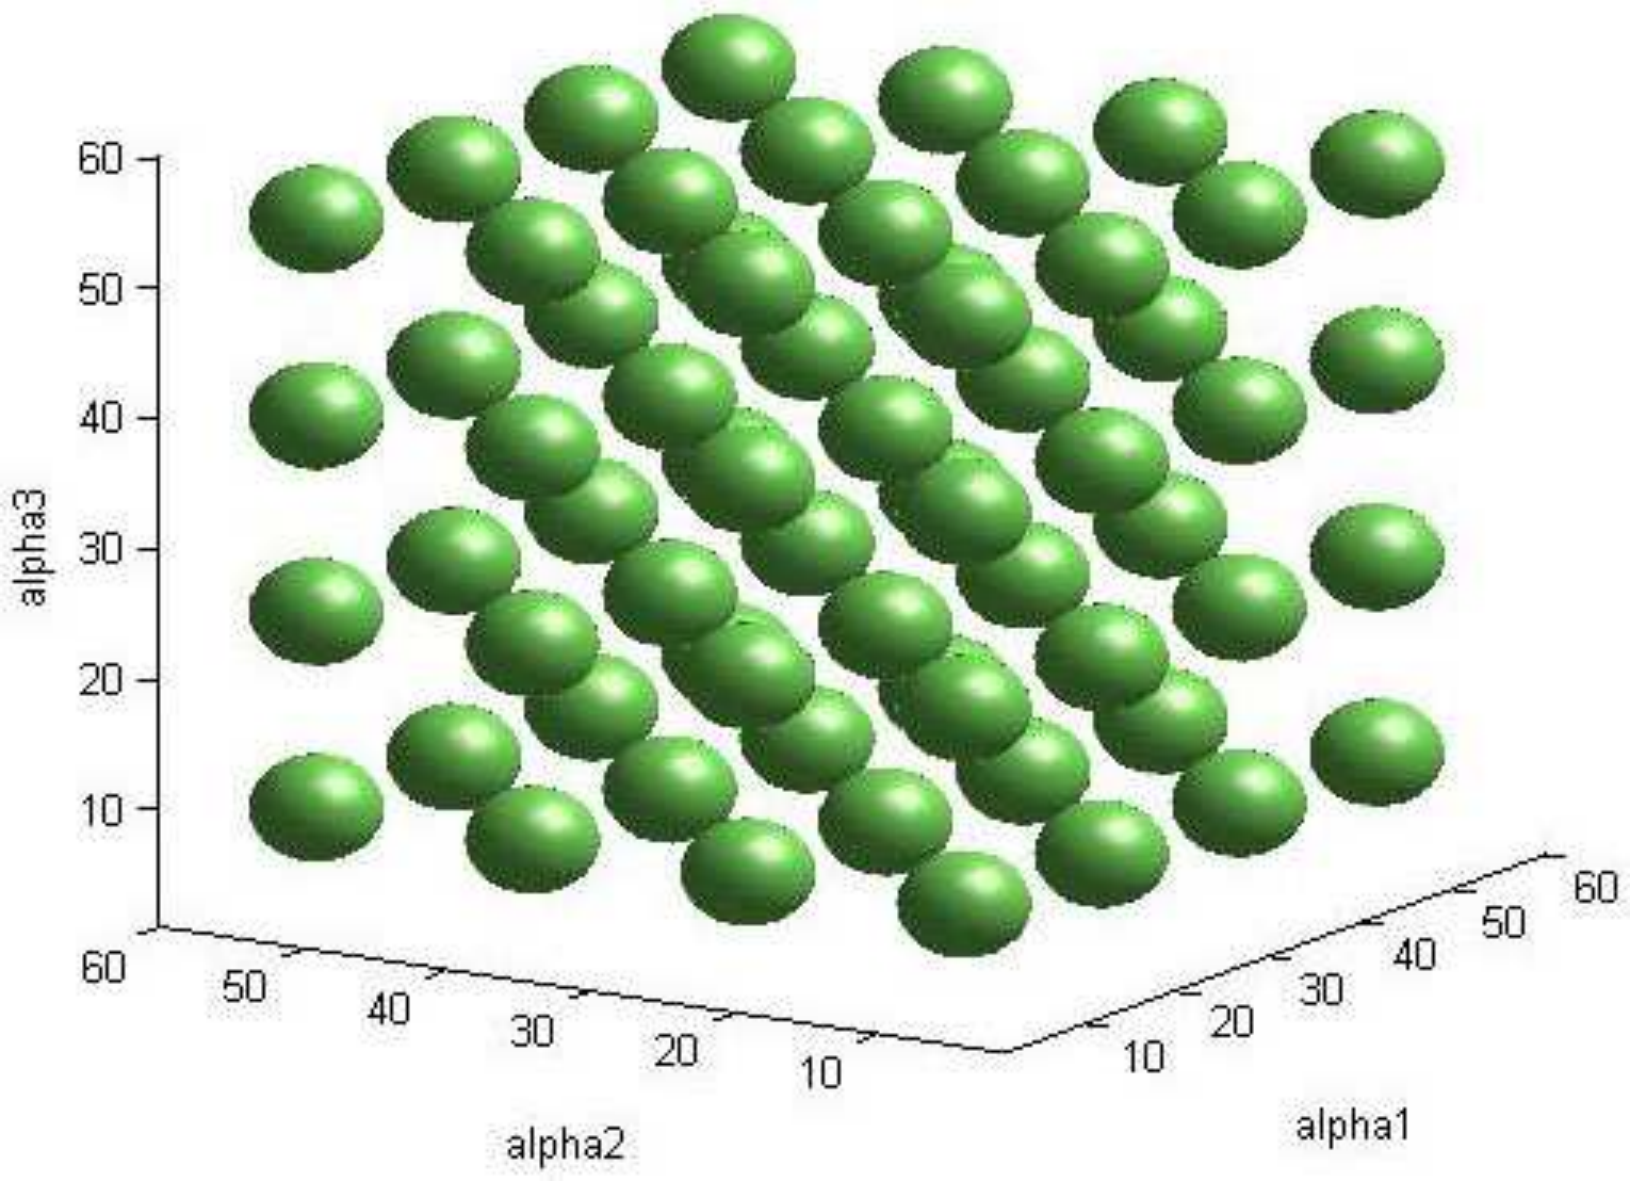
alpha3
60
50
40
30
20
10
60
50
40
30
20
10
alpha2
10
20
30
40
50
60
alpha1

Figure8

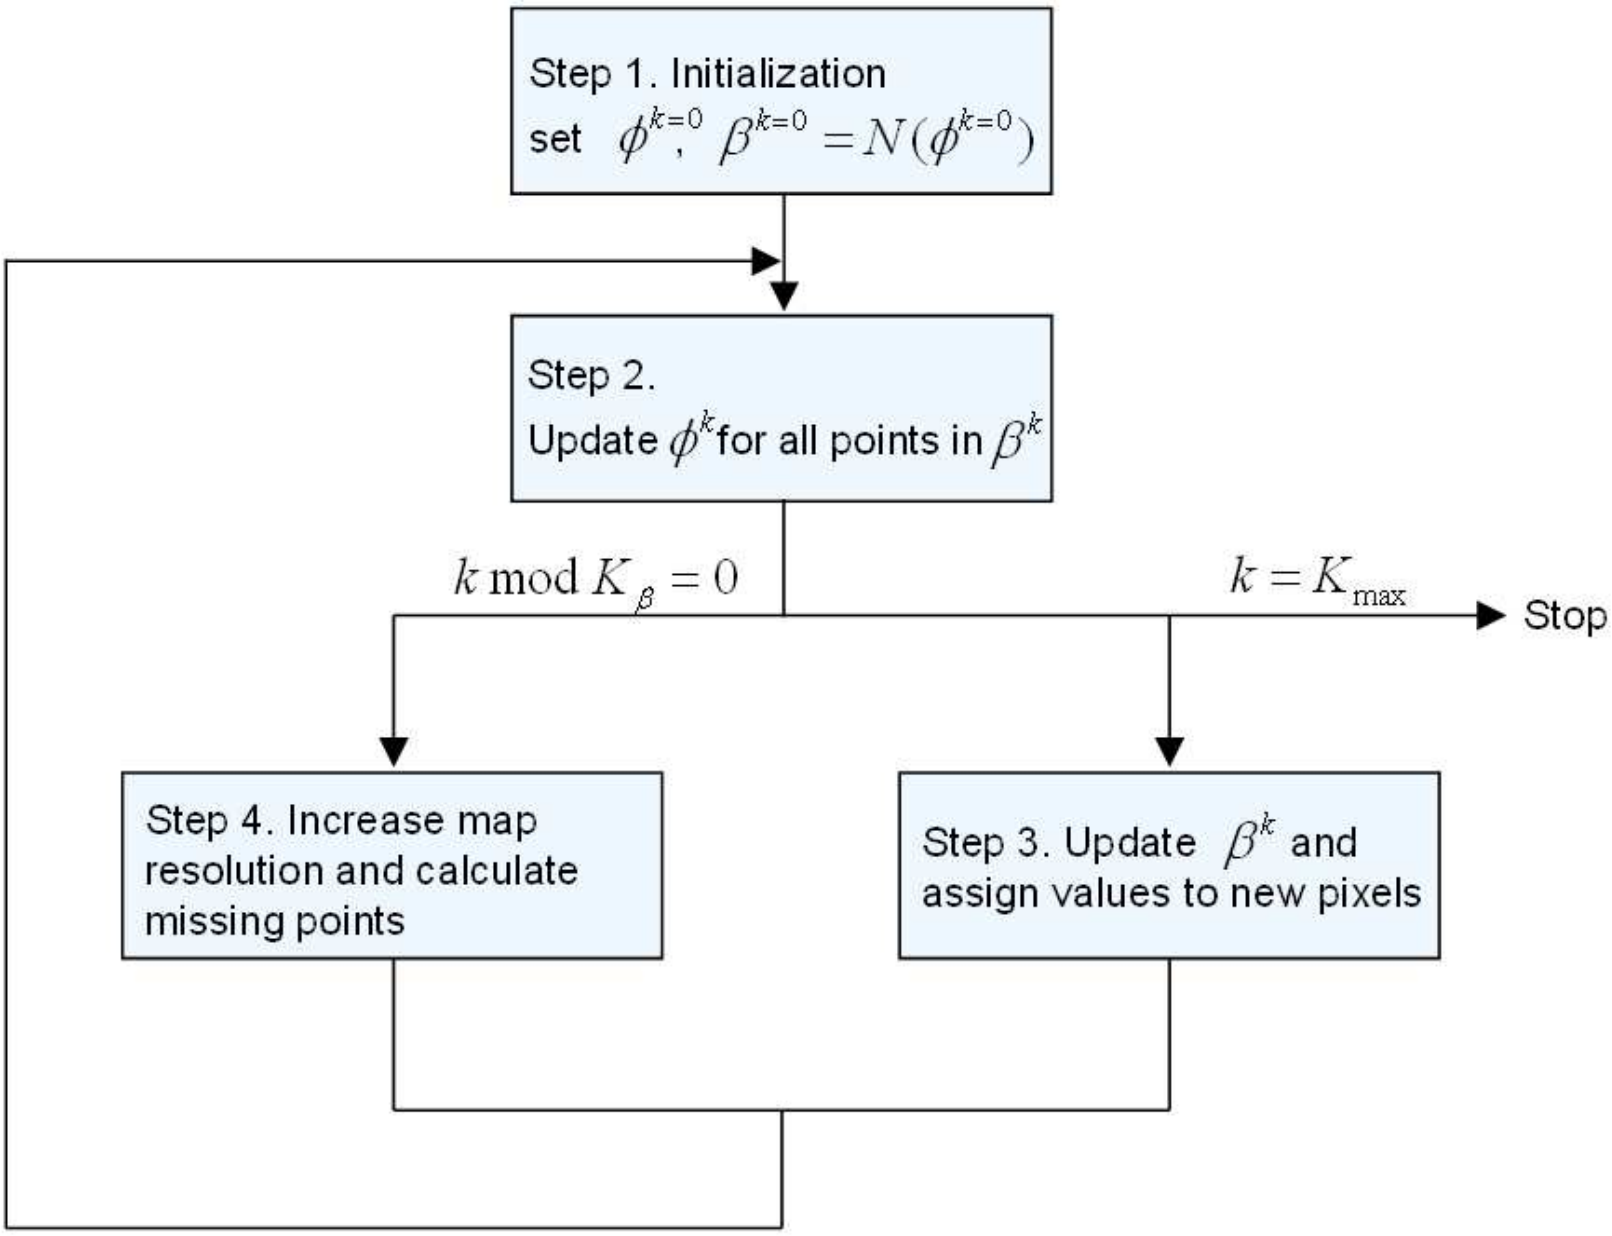
Step 1. Initialization
set $\phi^{k=0}$, $\beta^{k=0} = N(\phi^{k=0})$
Step 2.
Update ϕ^k for all points in β^k
$k \bmod K_\beta = 0$
$k = K_{\max}$
Stop
Step 4. Increase map resolution and calculate missing points
Step 3. Update β^k and assign values to new pixels

Figure9

Figure10a

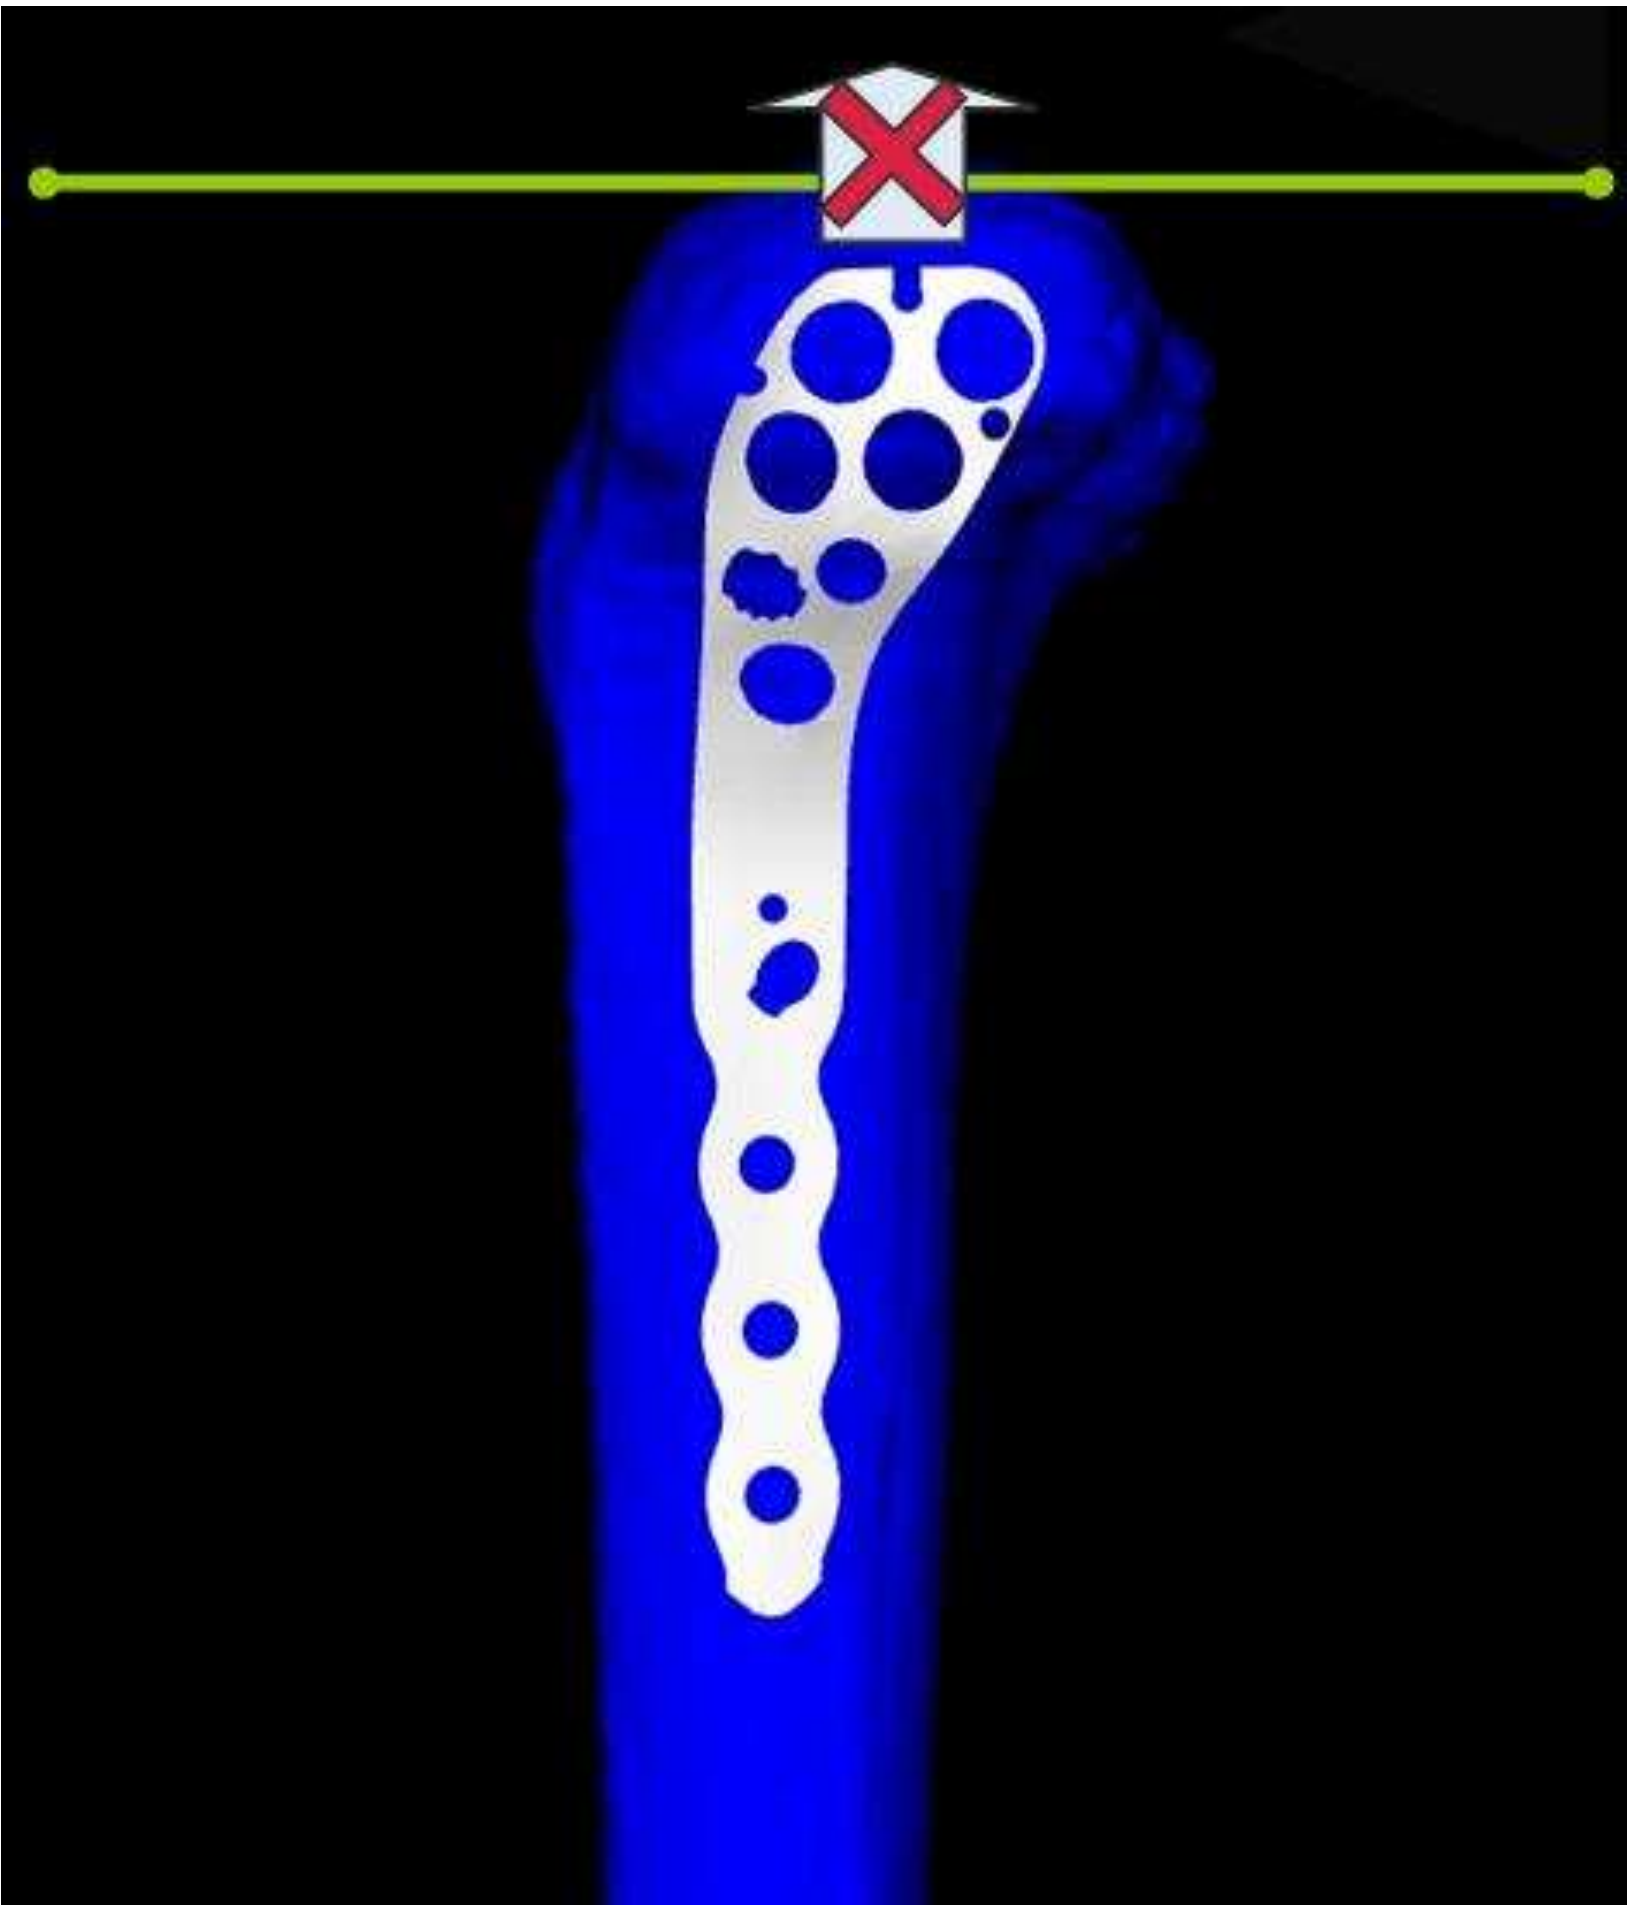

Figure10b

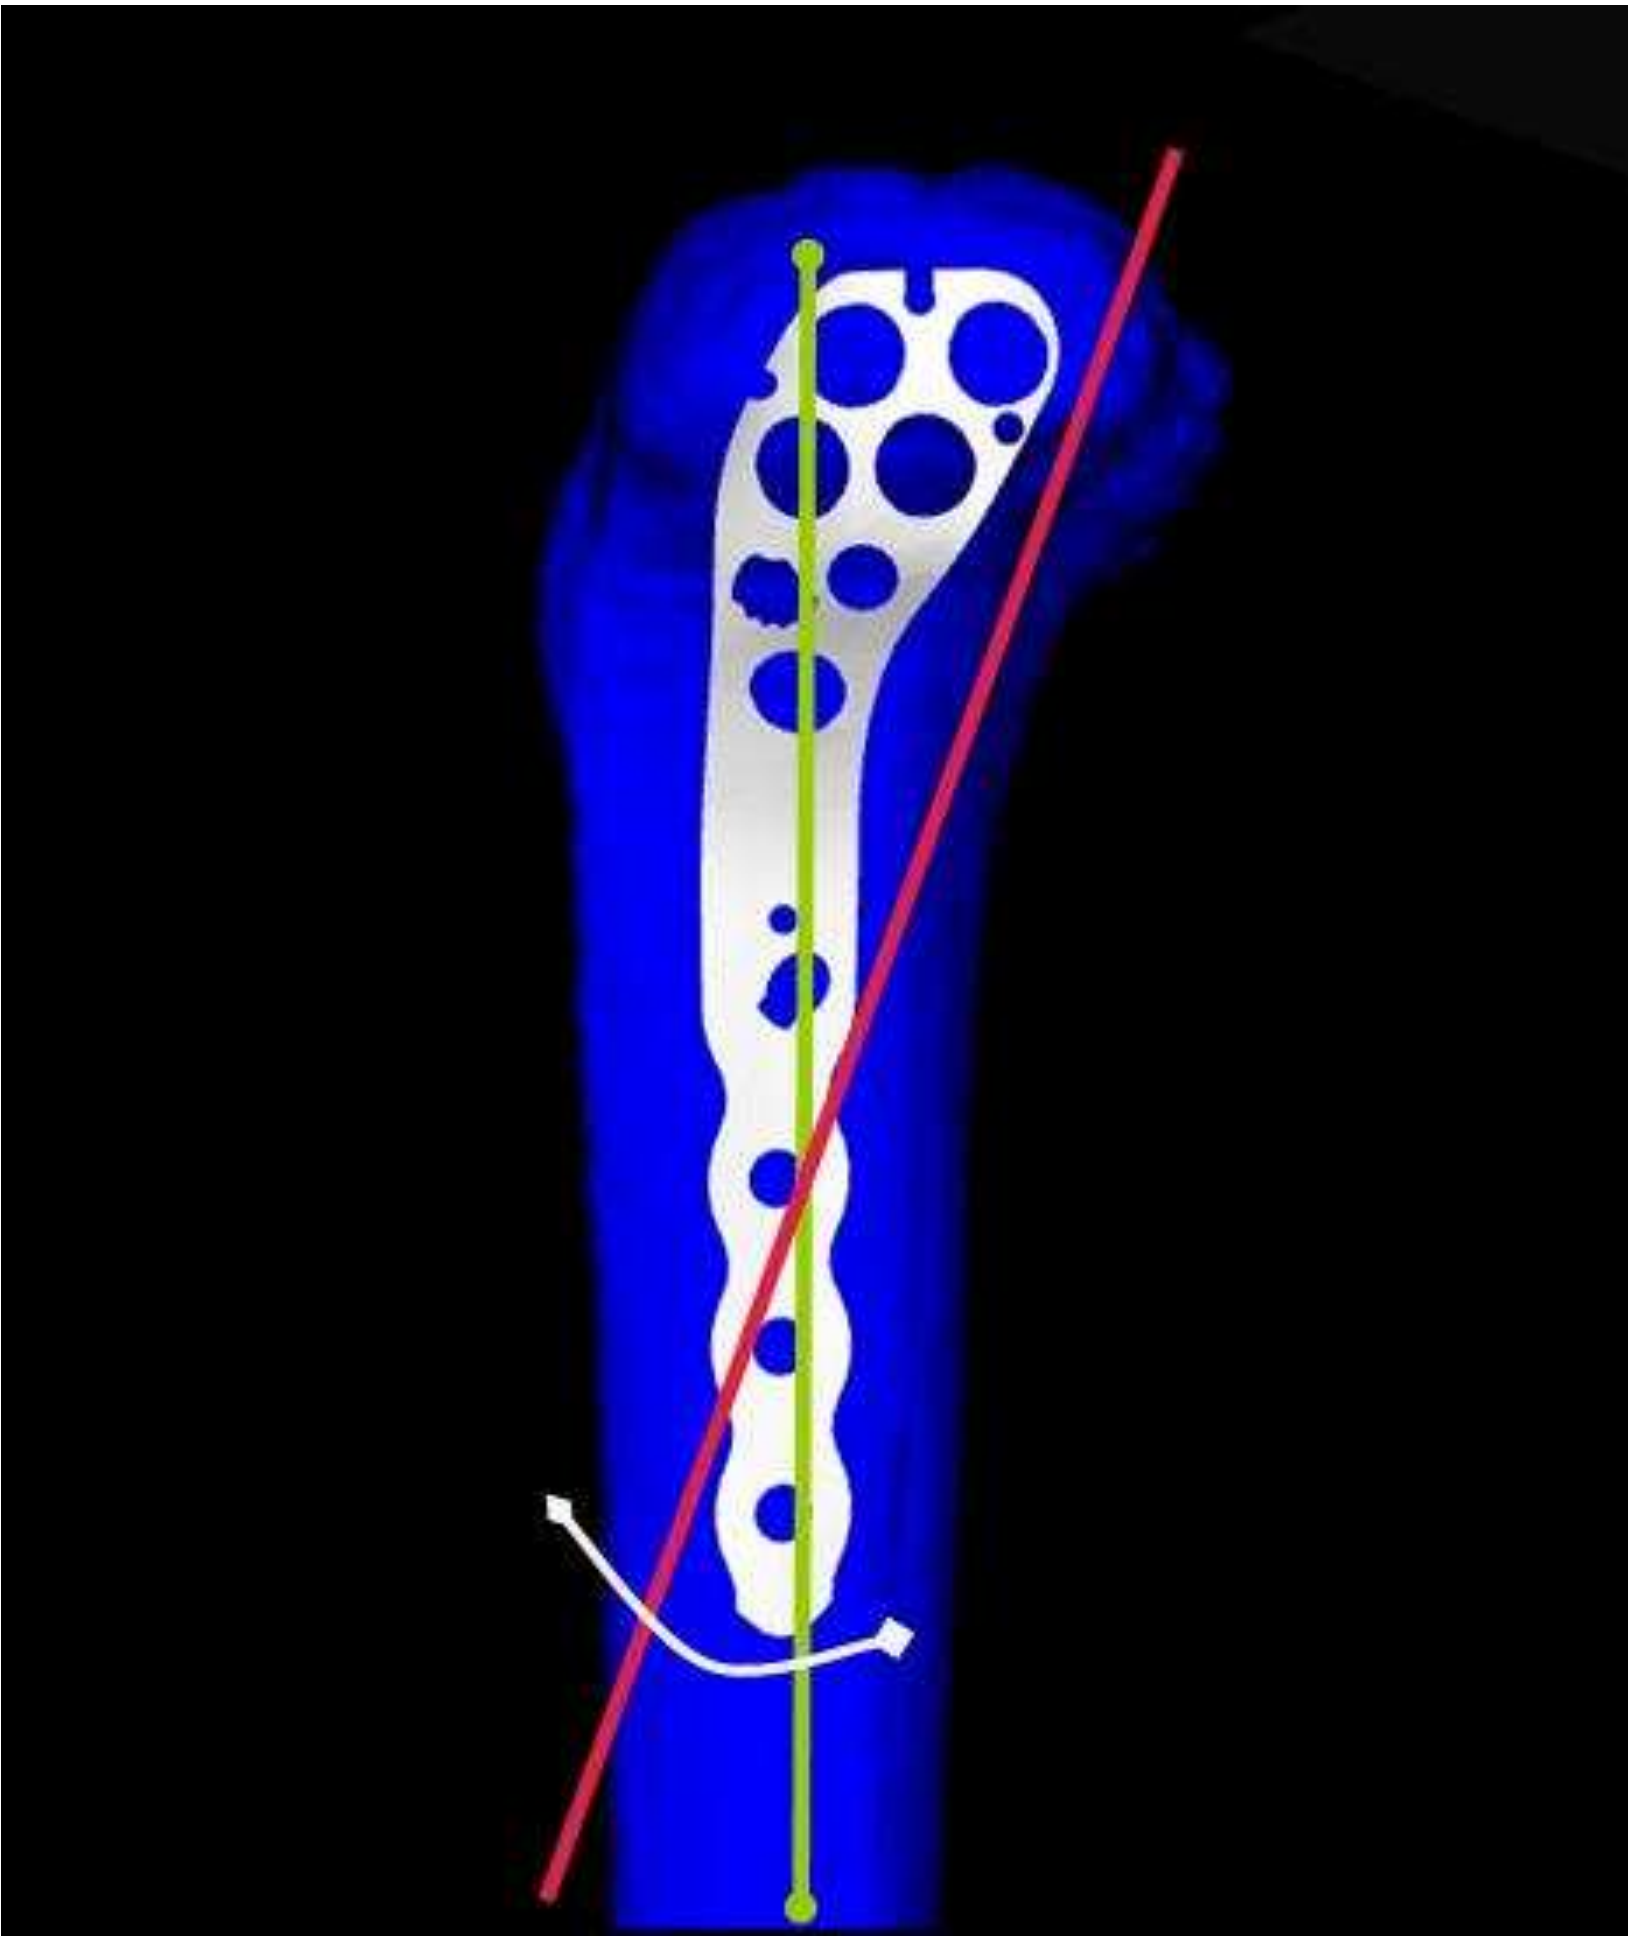

Figure11a

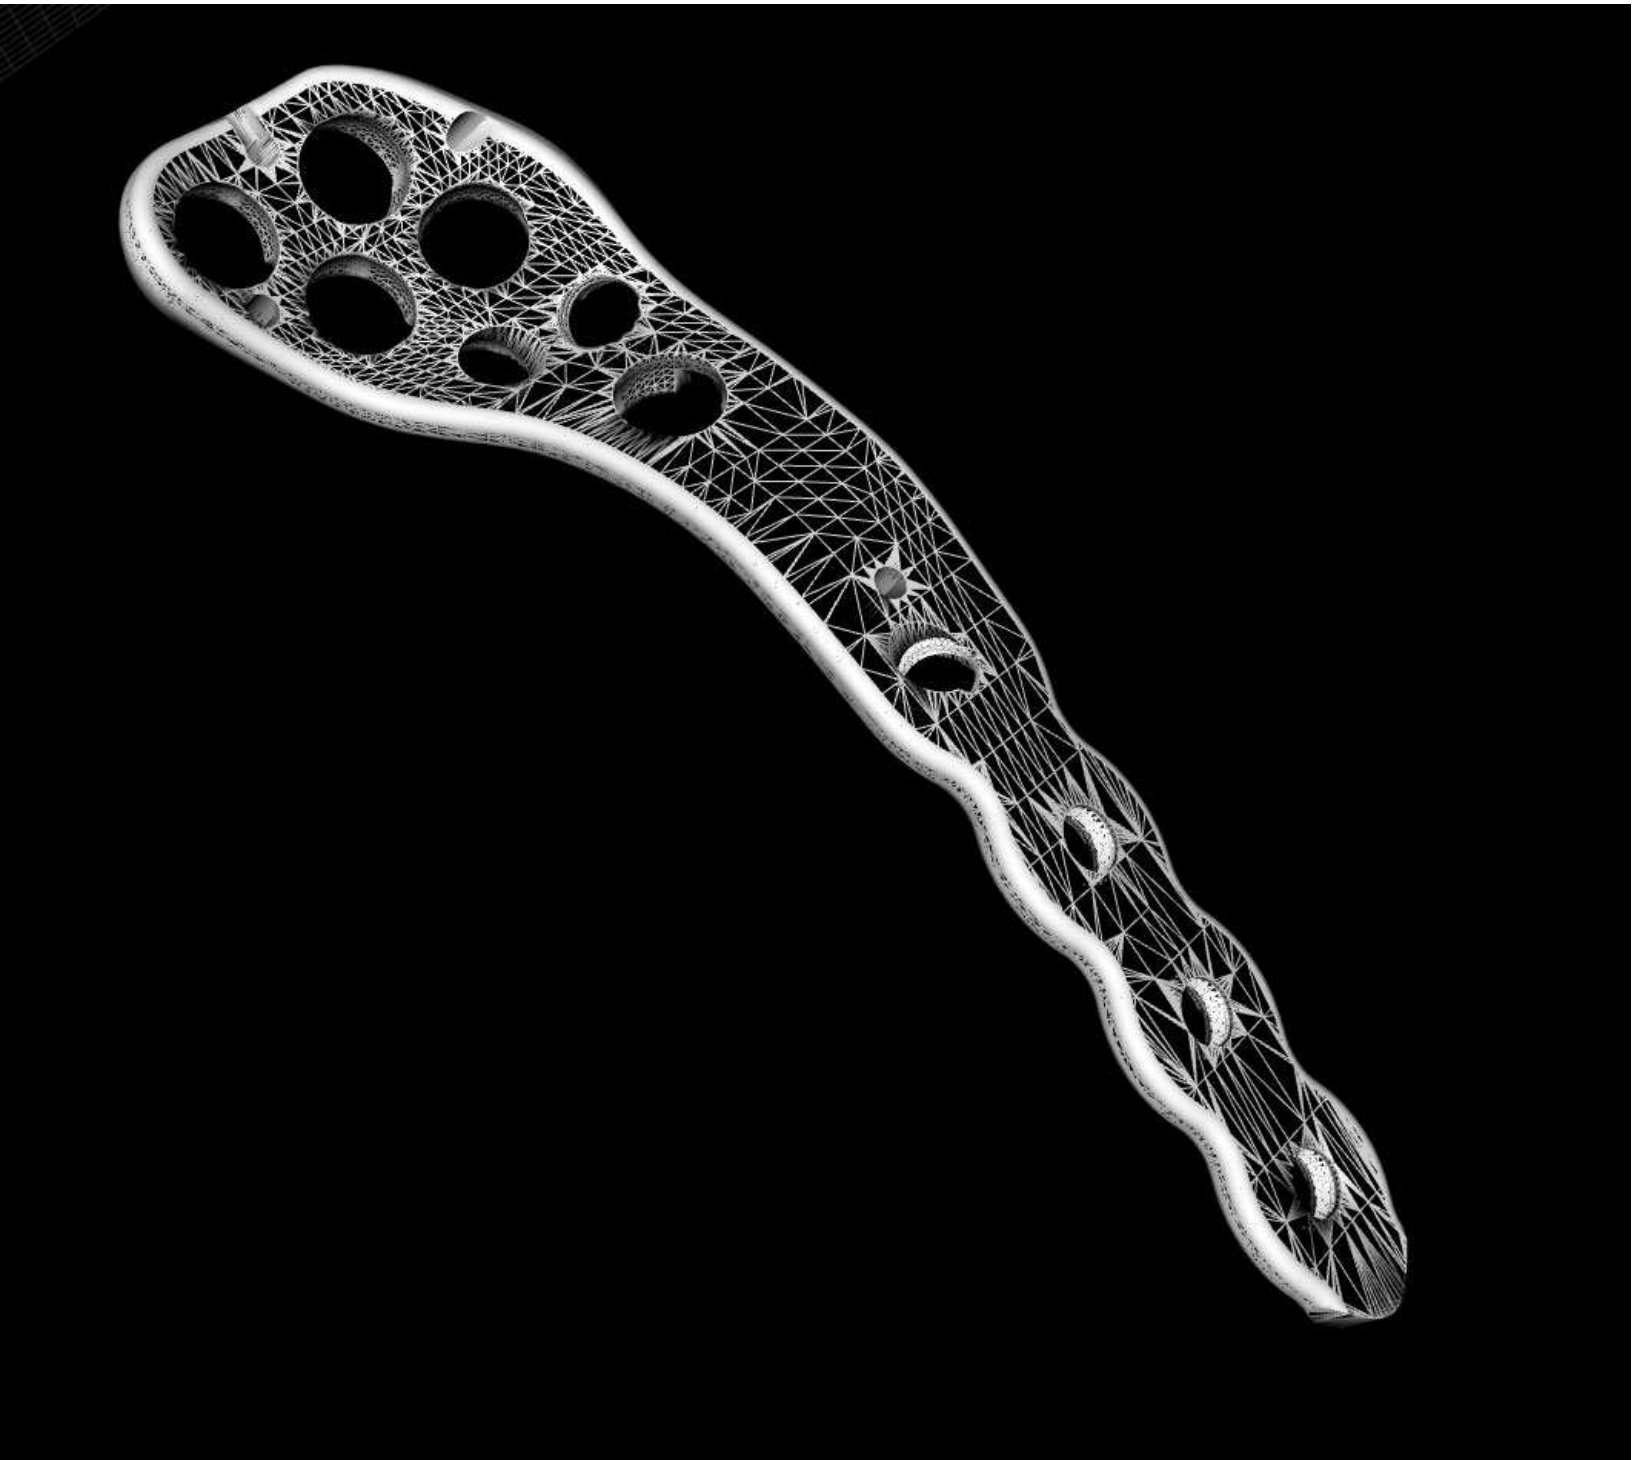

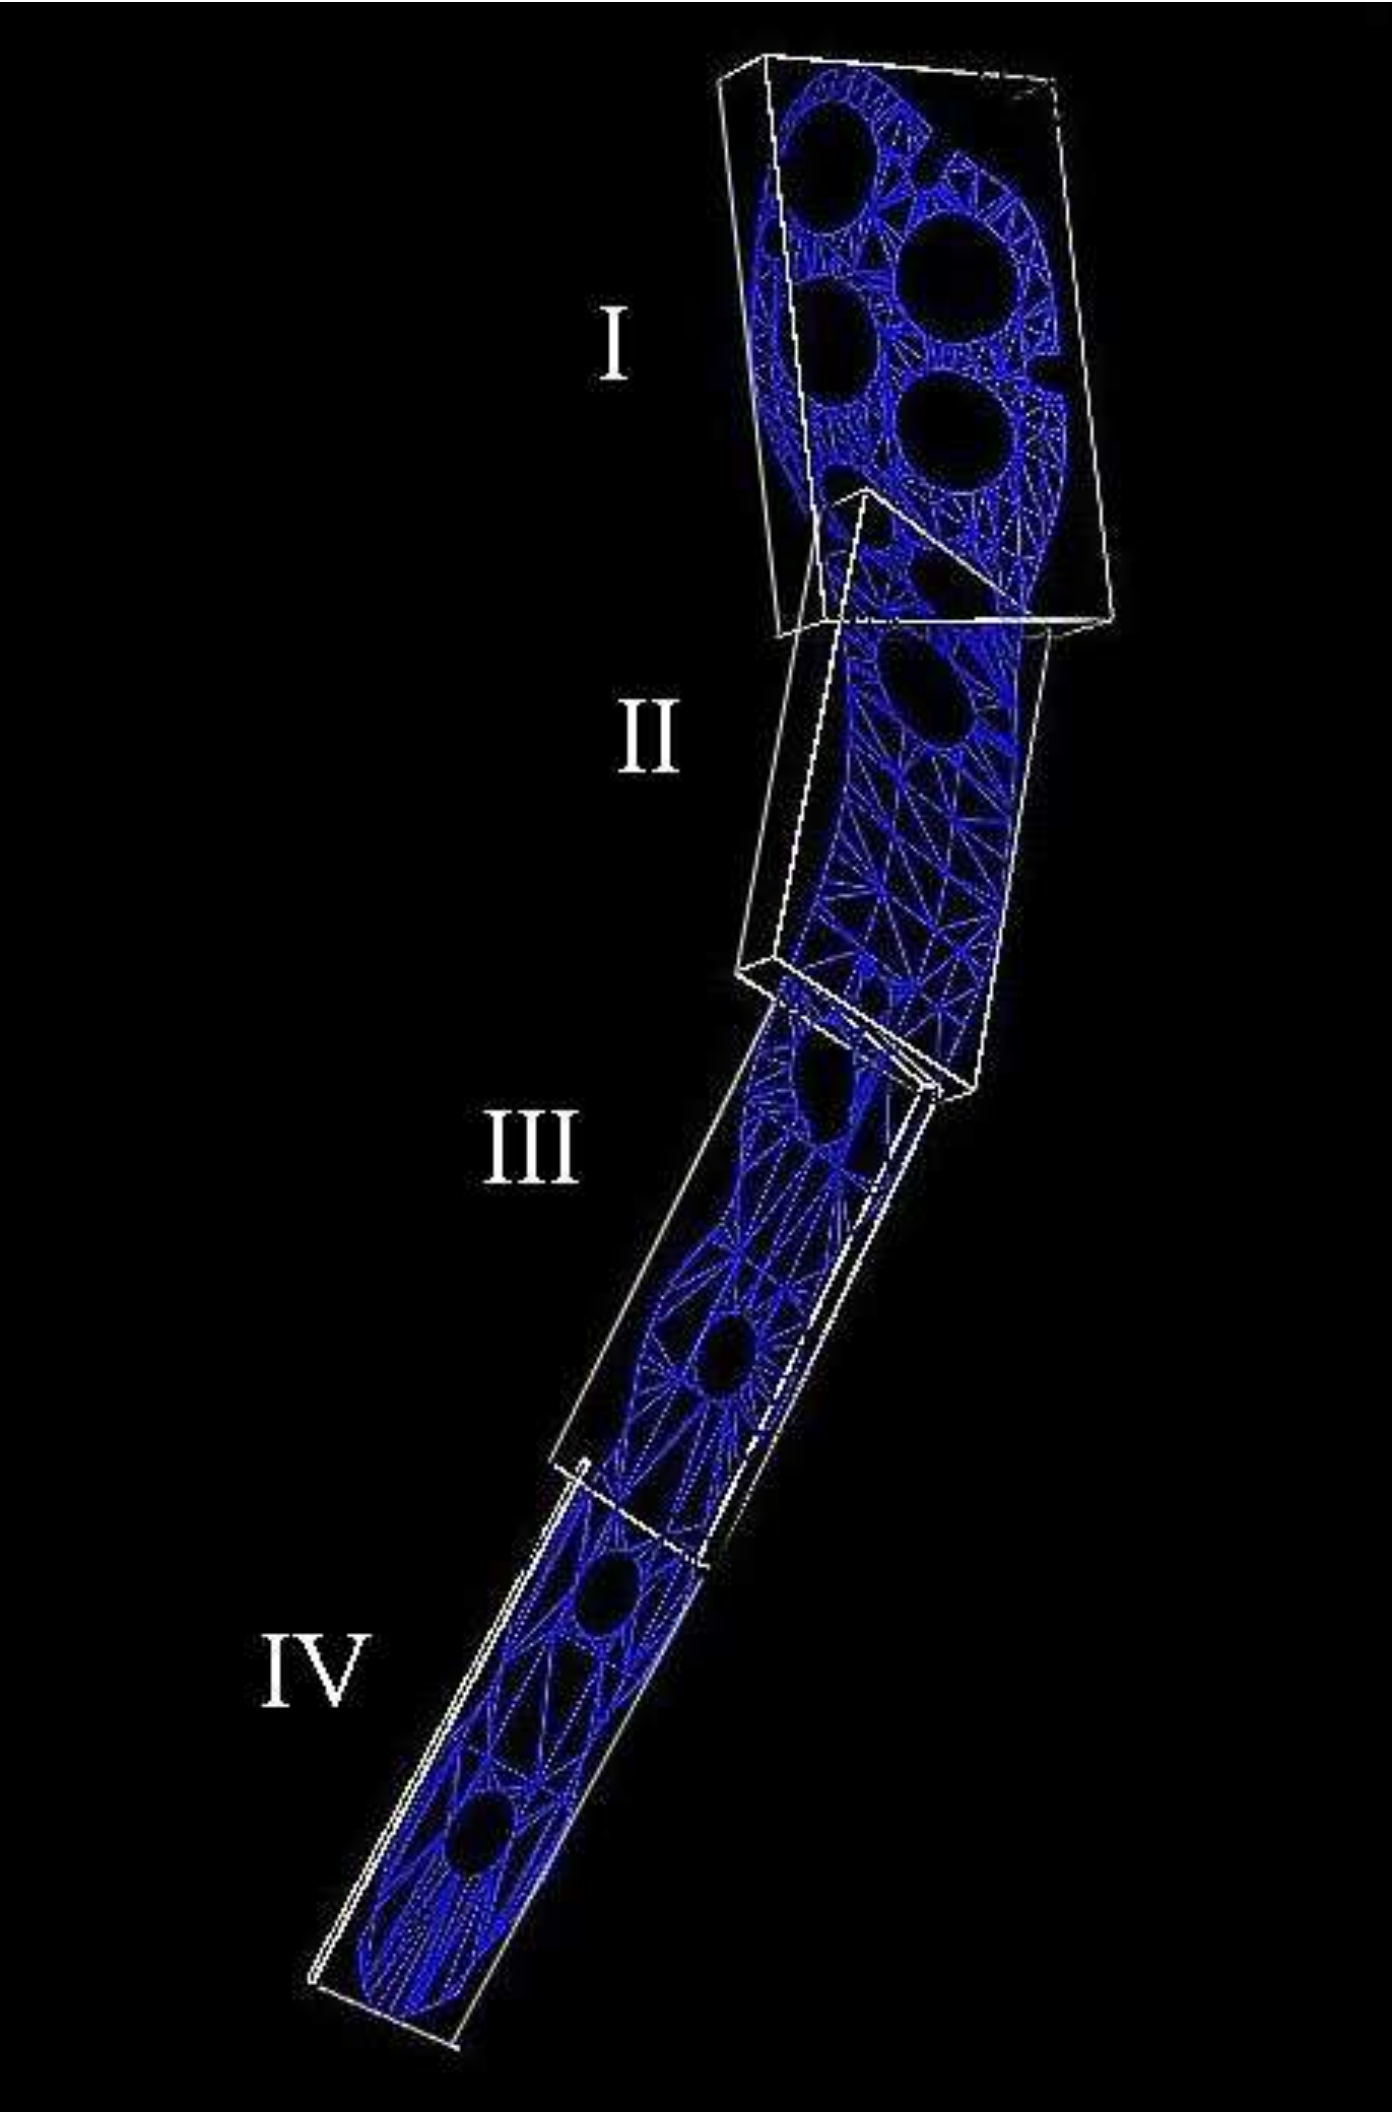
I
II
III
IV

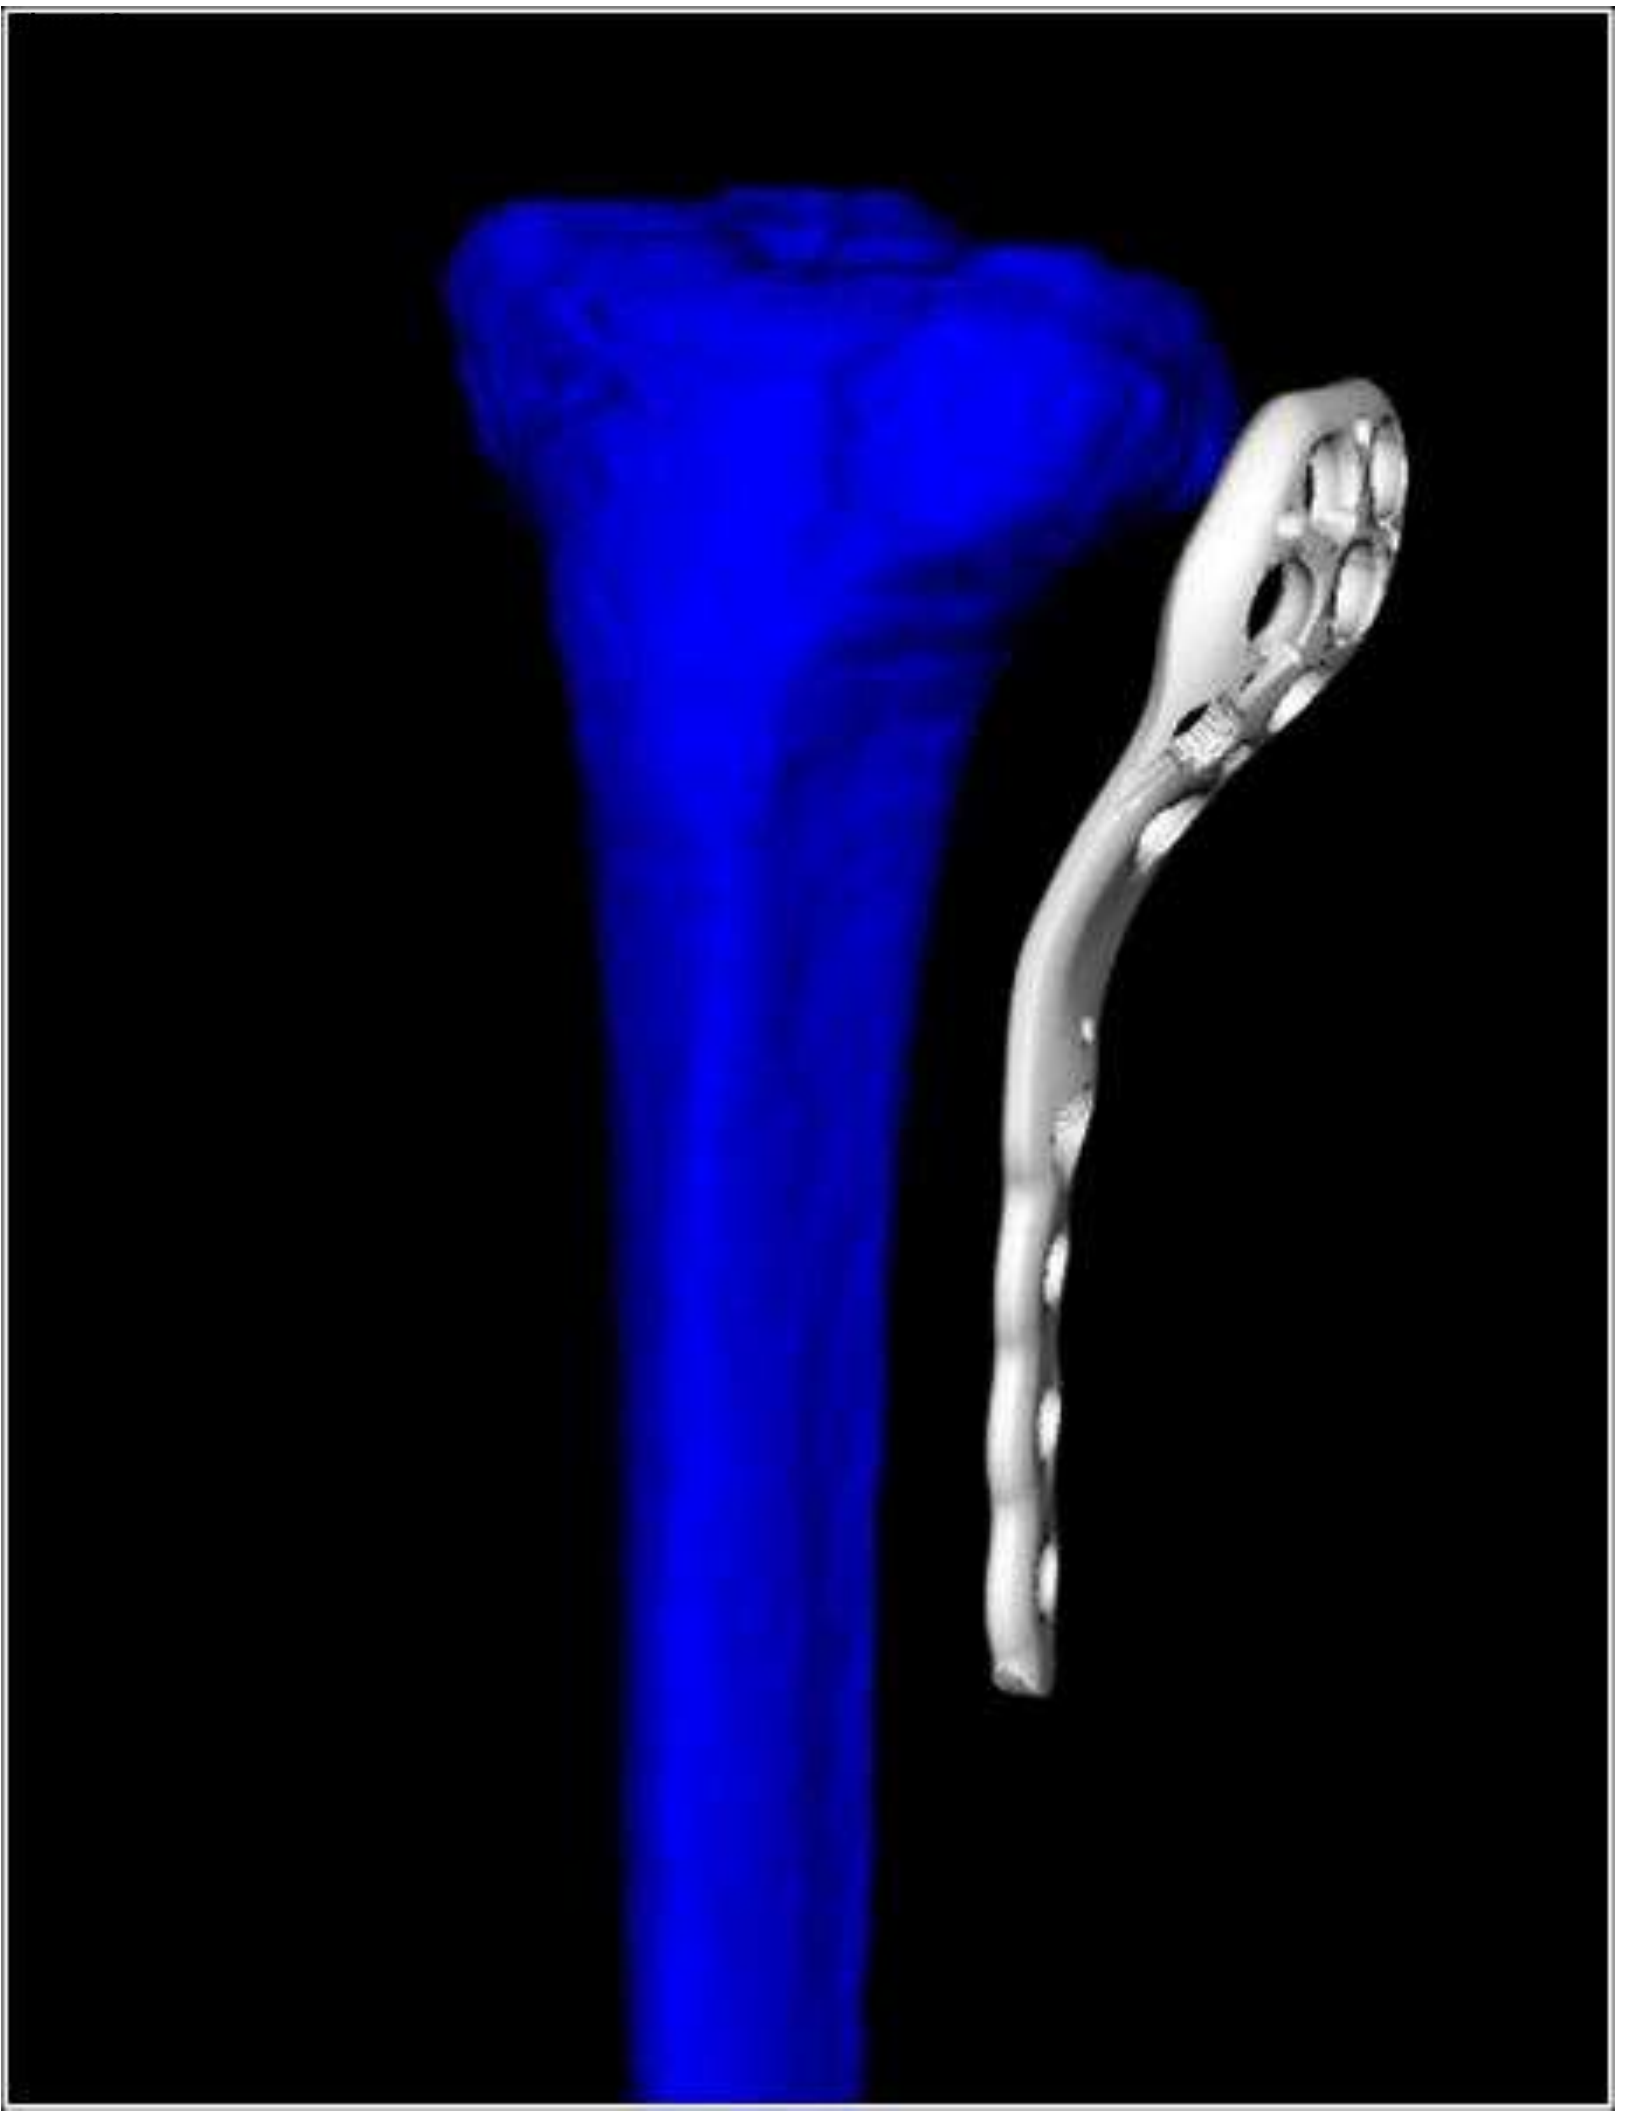

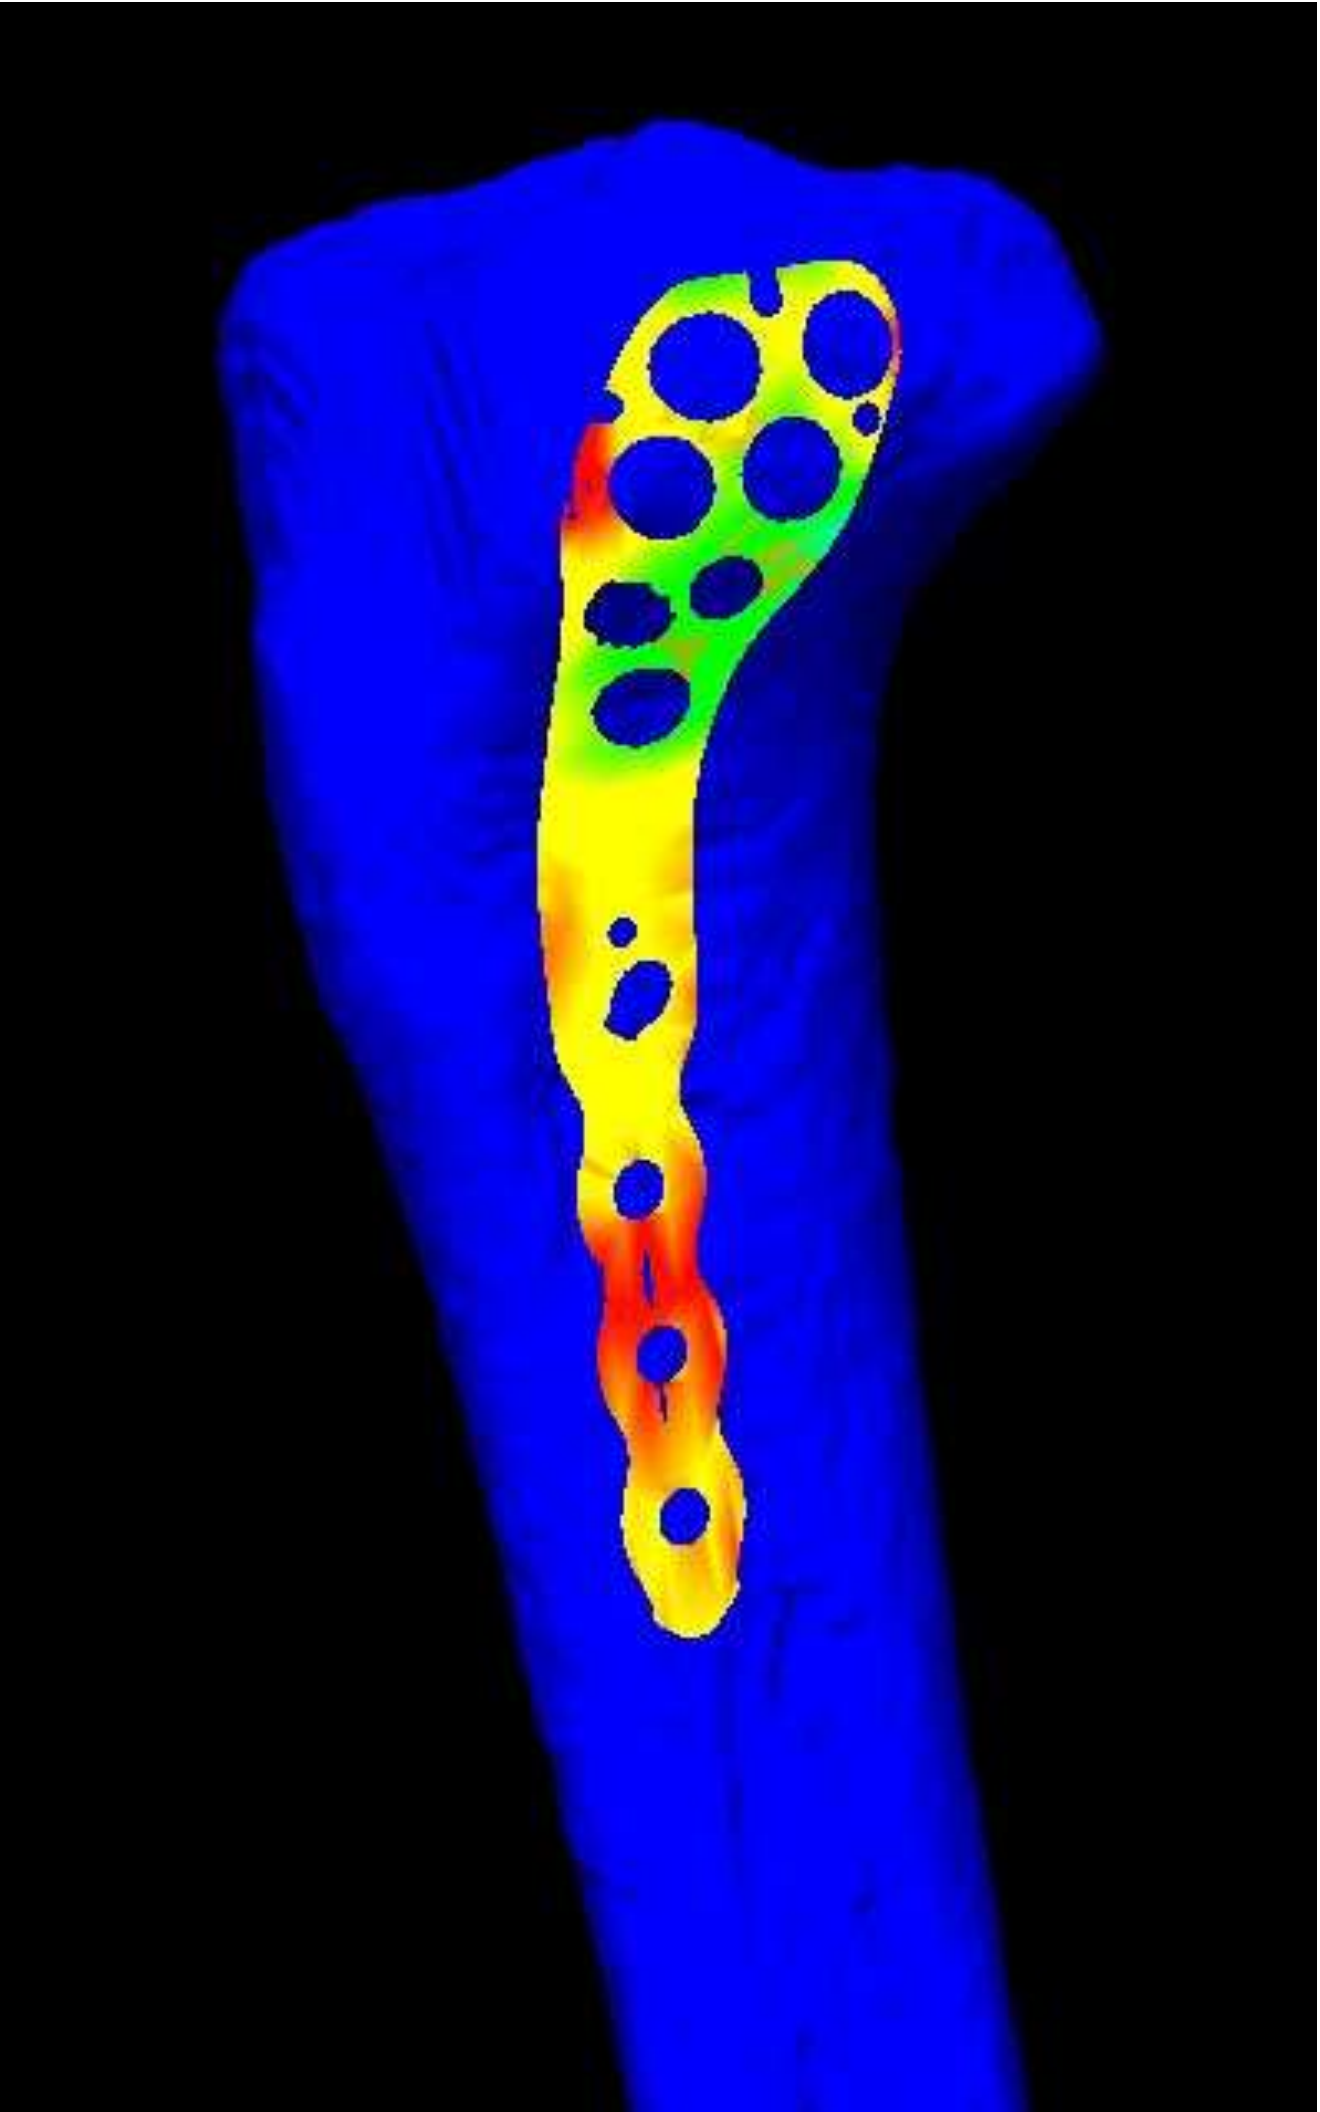

Figure13

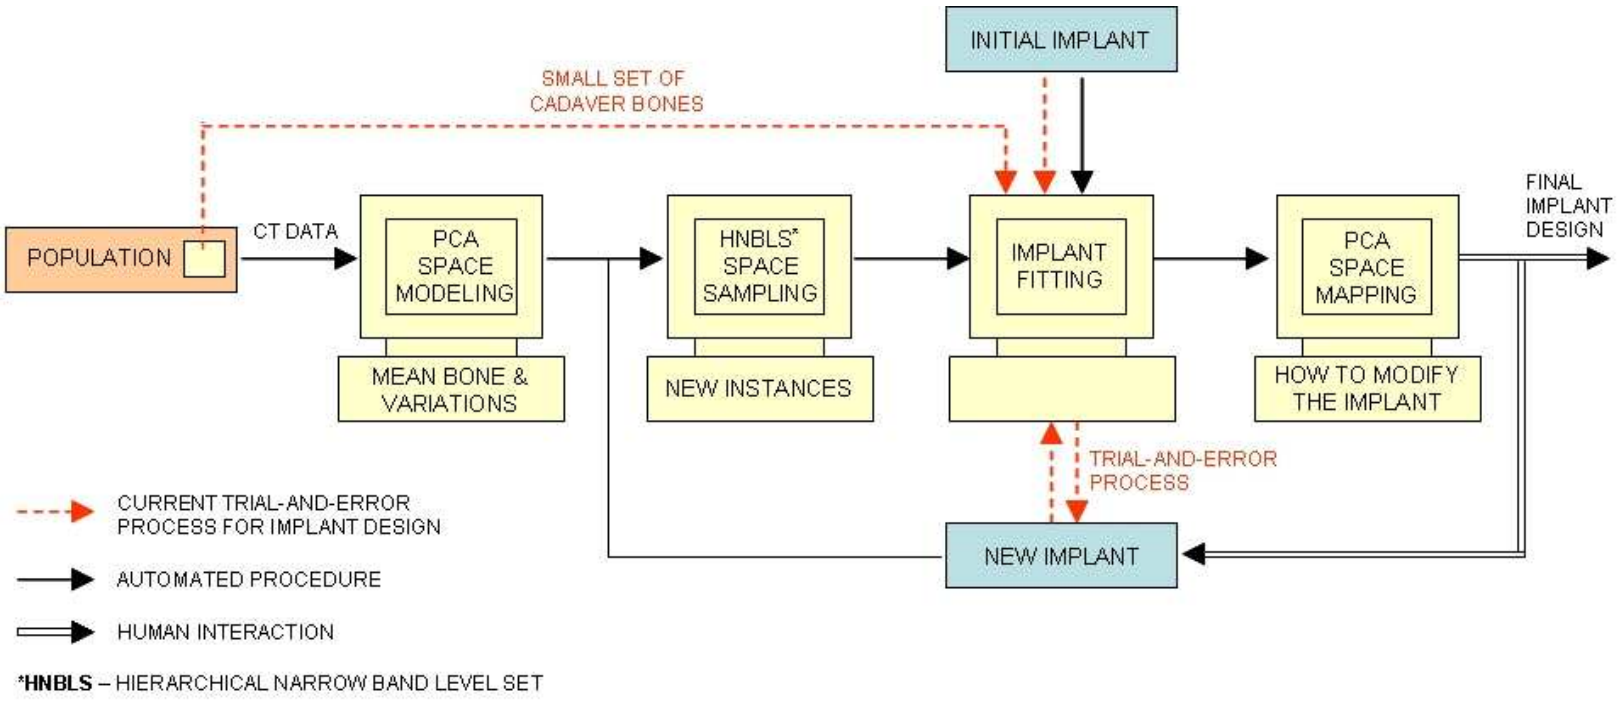
INITIAL IMPLANT
SMALL SET OF CADAVER BONES
POPULATION
CT DATA
PCA SPACE MODELING
HNBLS* SPACE SAMPLING
IMPLANT FITTING
PCA SPACE MAPPING
FINAL IMPLANT DESIGN
MEAN BONE & VARIATIONS
NEW INSTANCES
HOW TO MODIFY THE IMPLANT
TRIAL-AND-ERROR PROCESS
NEW IMPLANT
CURRENT TRIAL-AND-ERROR PROCESS FOR IMPLANT DESIGN
AUTOMATED PROCEDURE
HUMAN INTERACTION
*HNBLS – HIERARCHICAL NARROW BAND LEVEL SET

Figure14

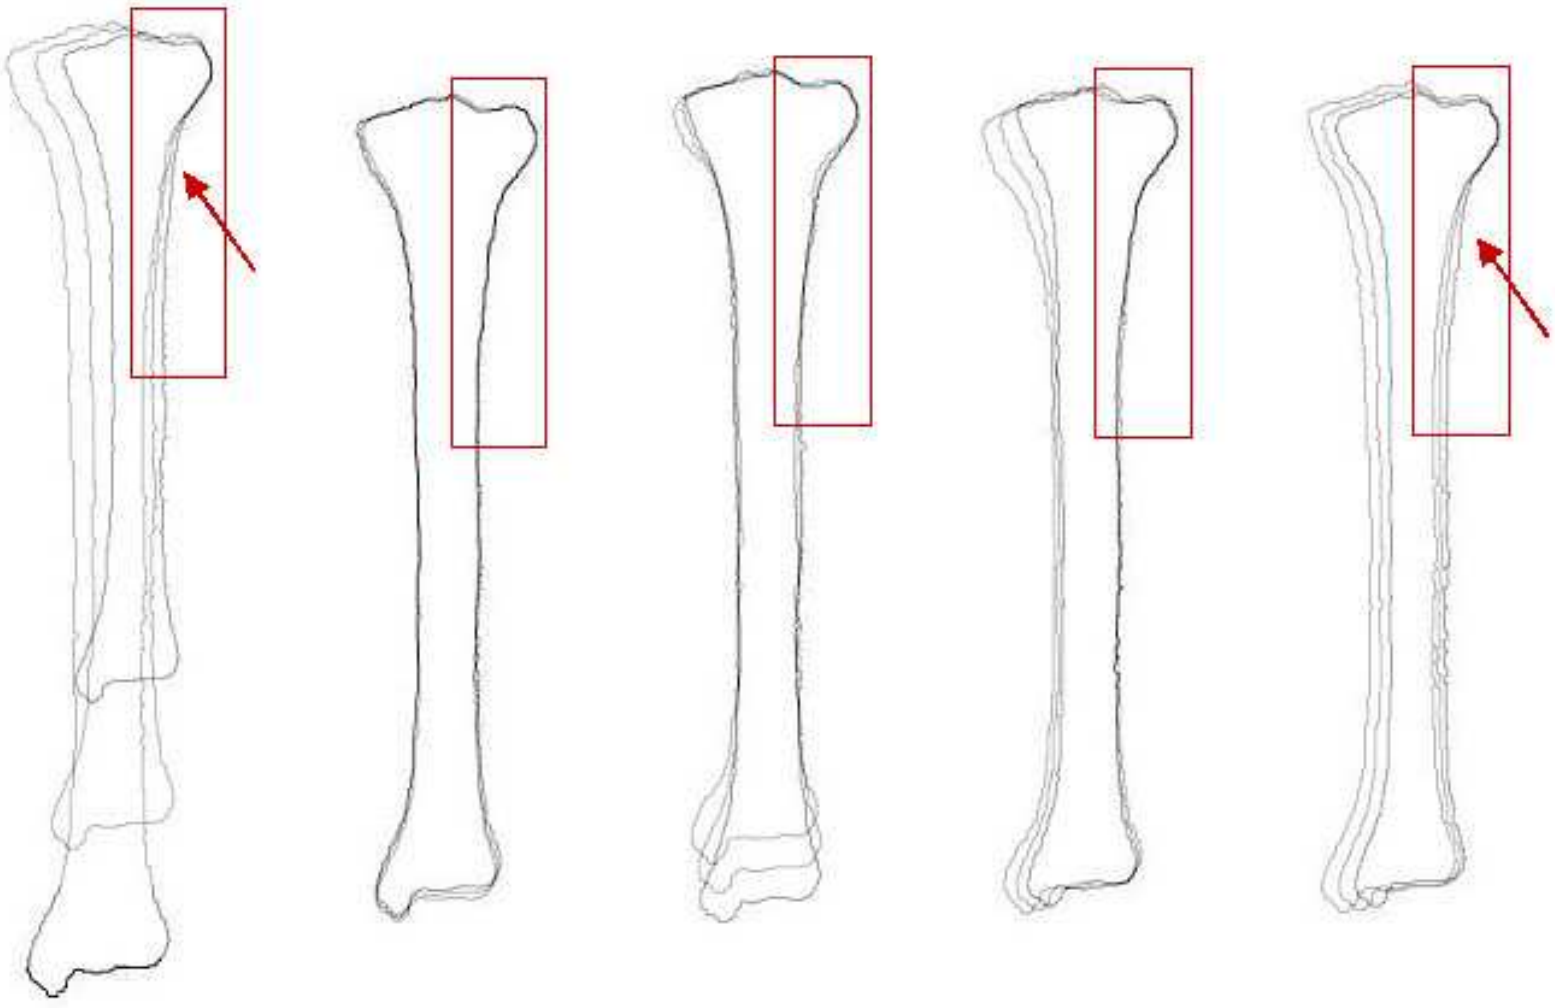

Figure15a

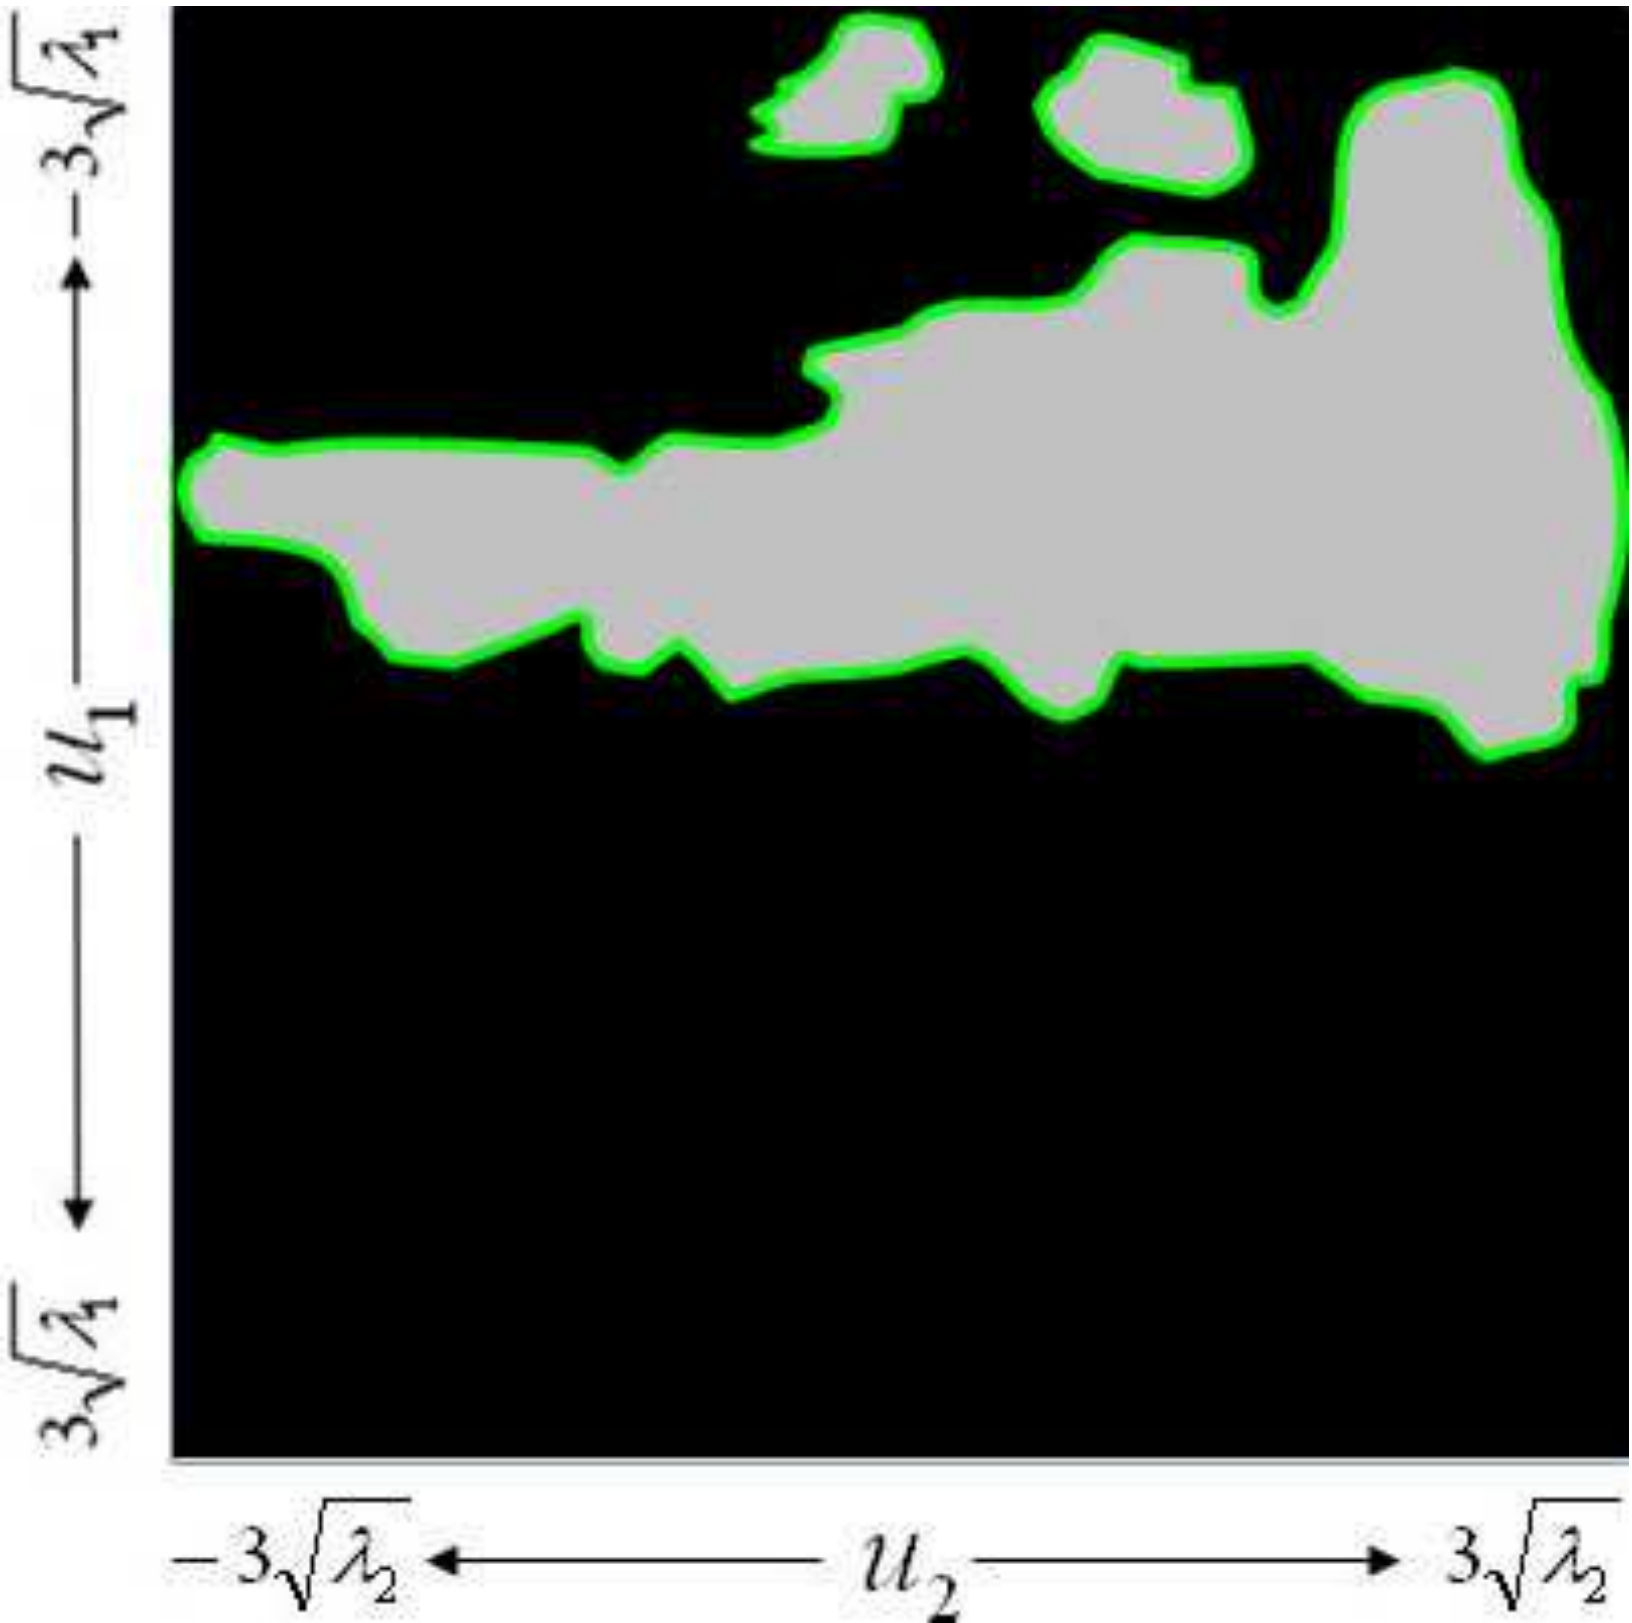
$-3\sqrt{\lambda_1}$
u_1
$3\sqrt{\lambda_1}$
$-3\sqrt{\lambda_2}$
u_2
$3\sqrt{\lambda_2}$

Figure15b

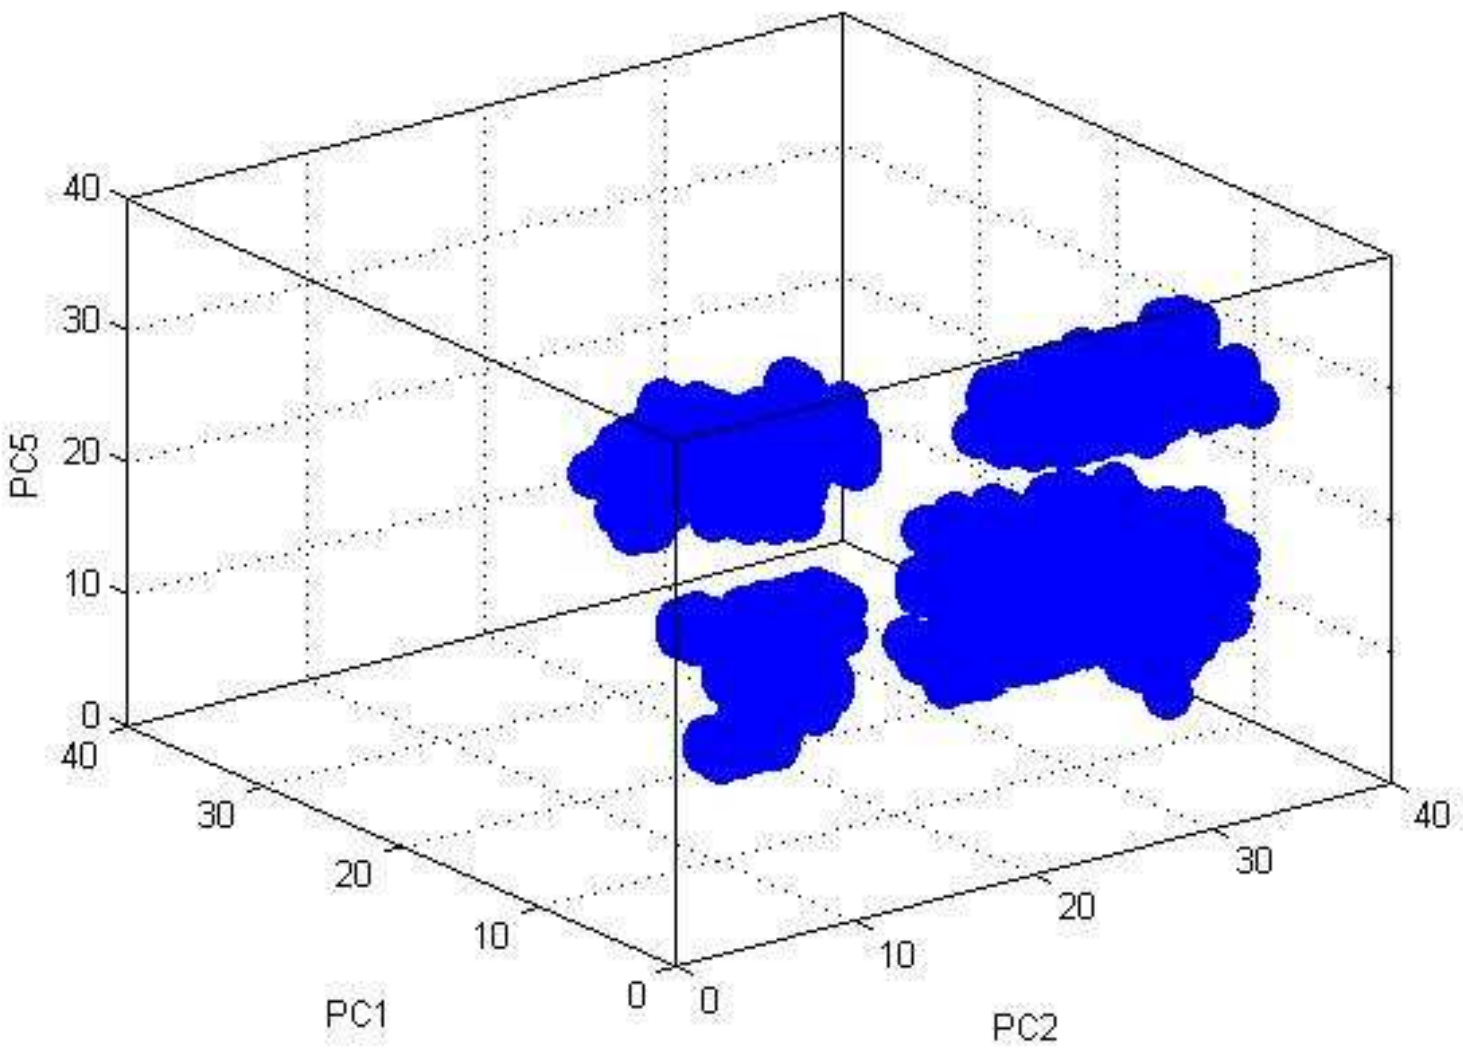

Figure15c

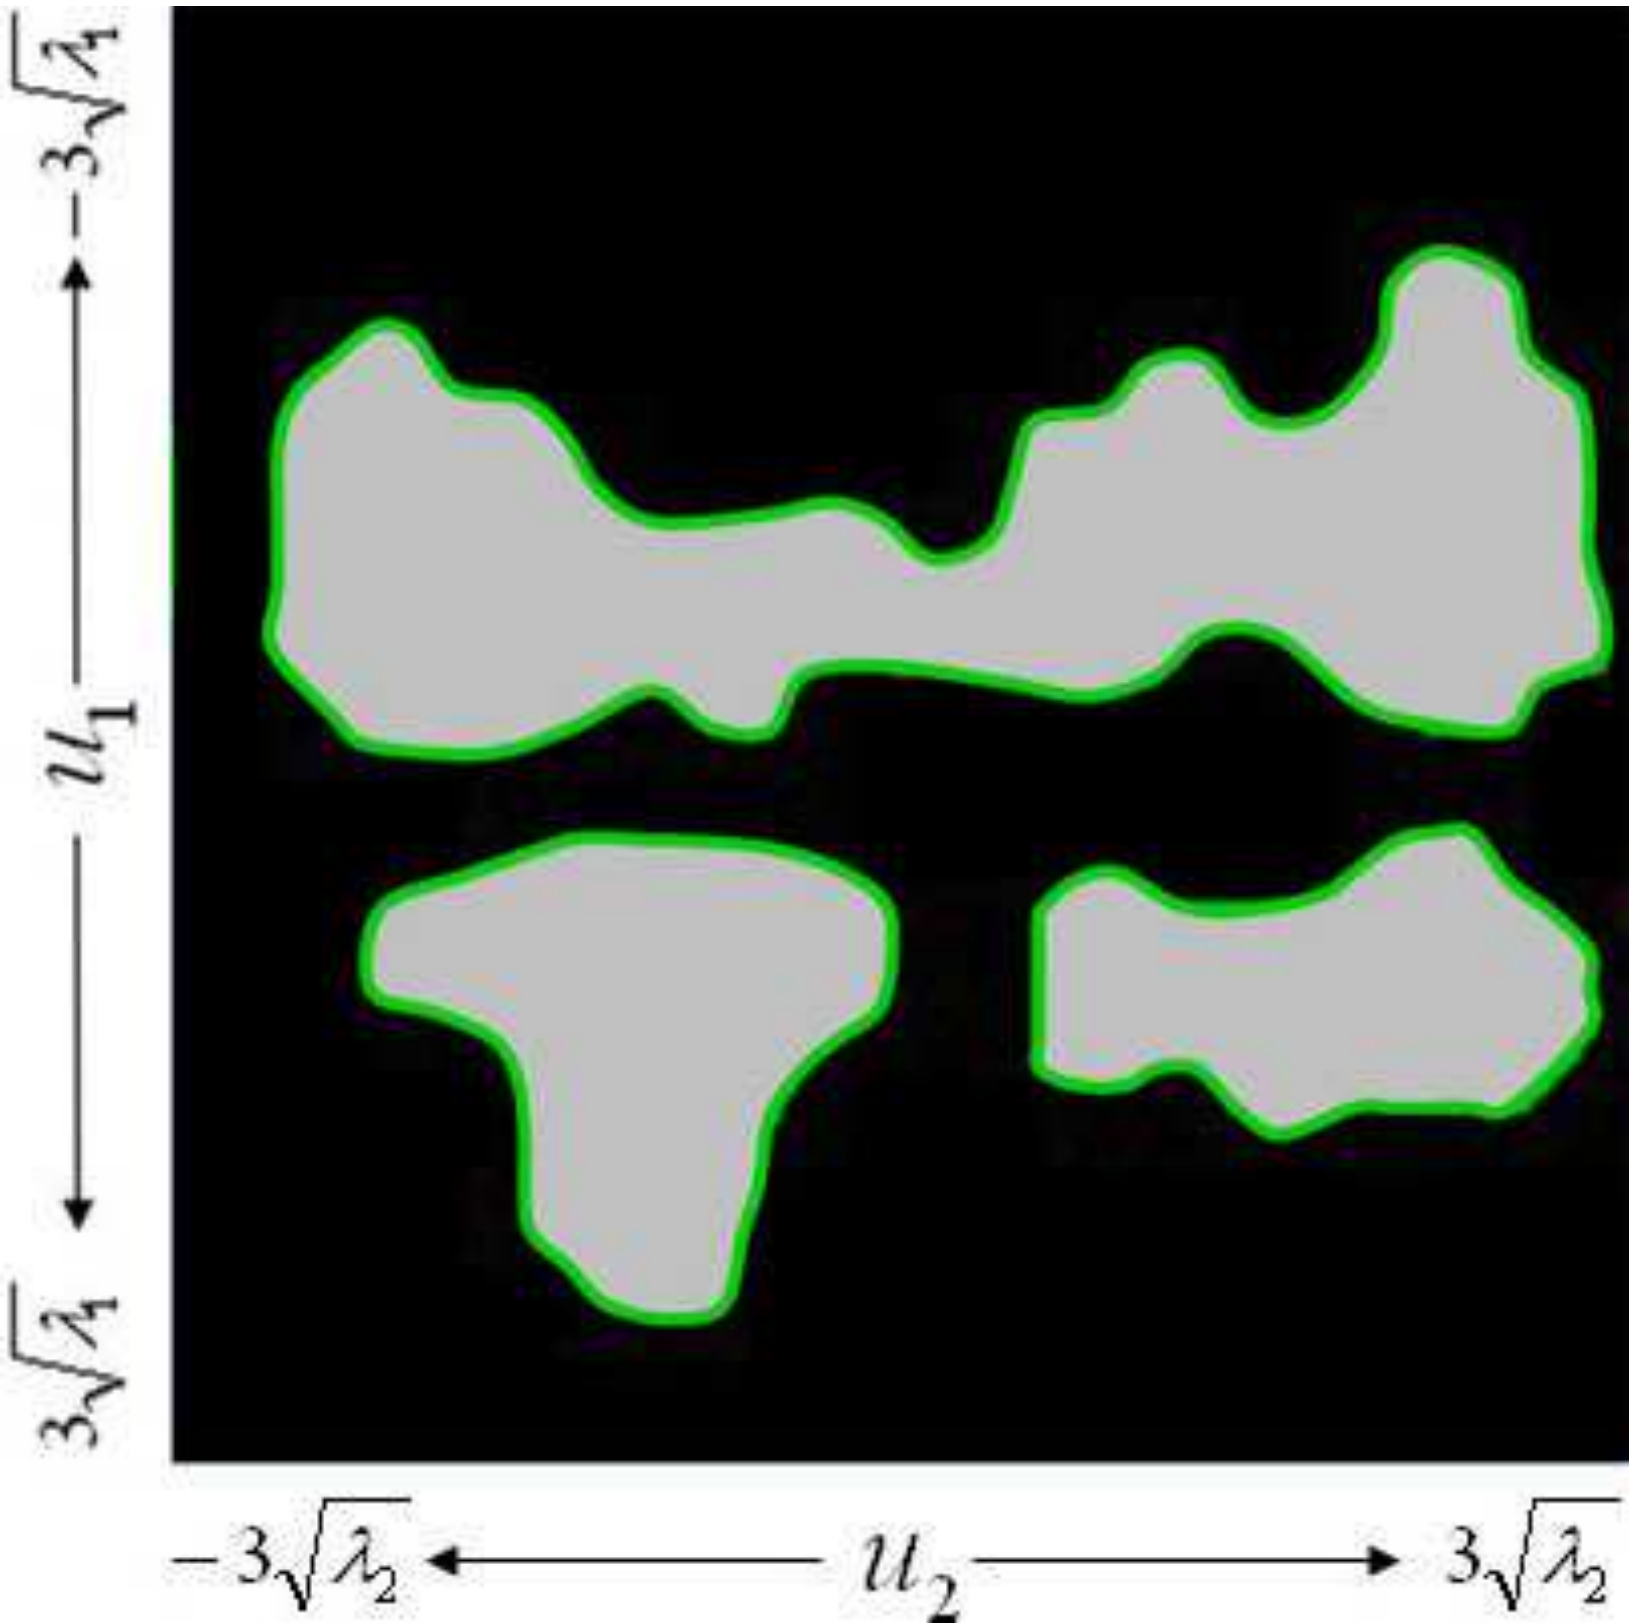
$-3\sqrt{\lambda_1}$
u_1
$3\sqrt{\lambda_1}$
$-3\sqrt{\lambda_2}$
u_2
$3\sqrt{\lambda_2}$

Figure15d

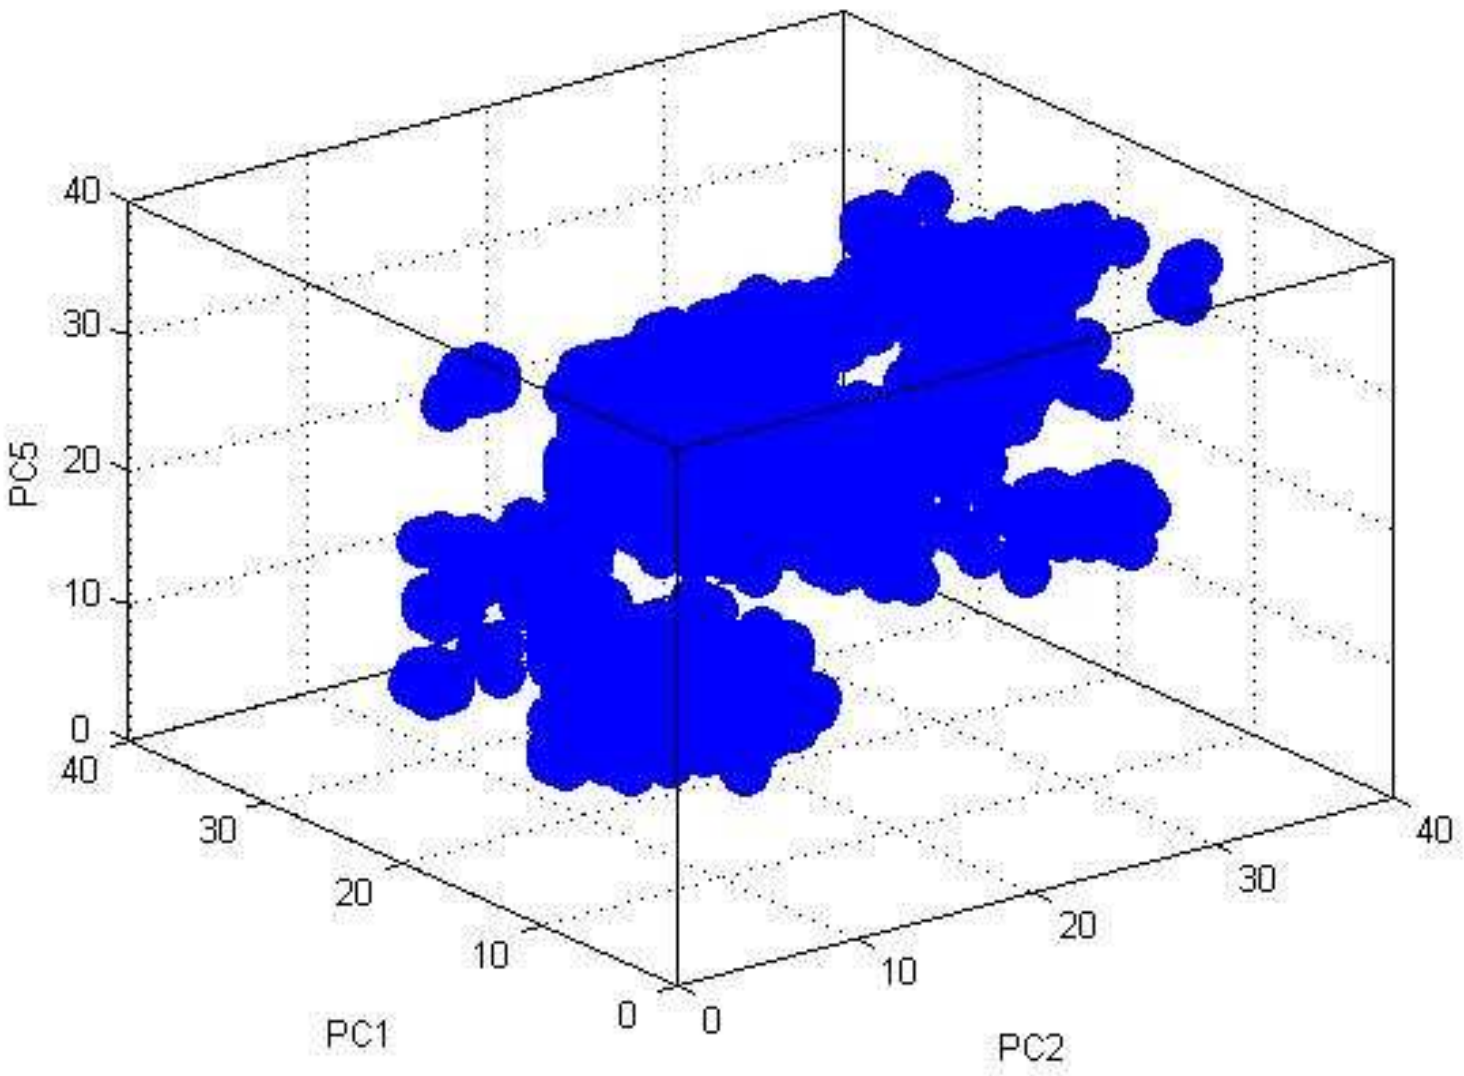

Figure16a

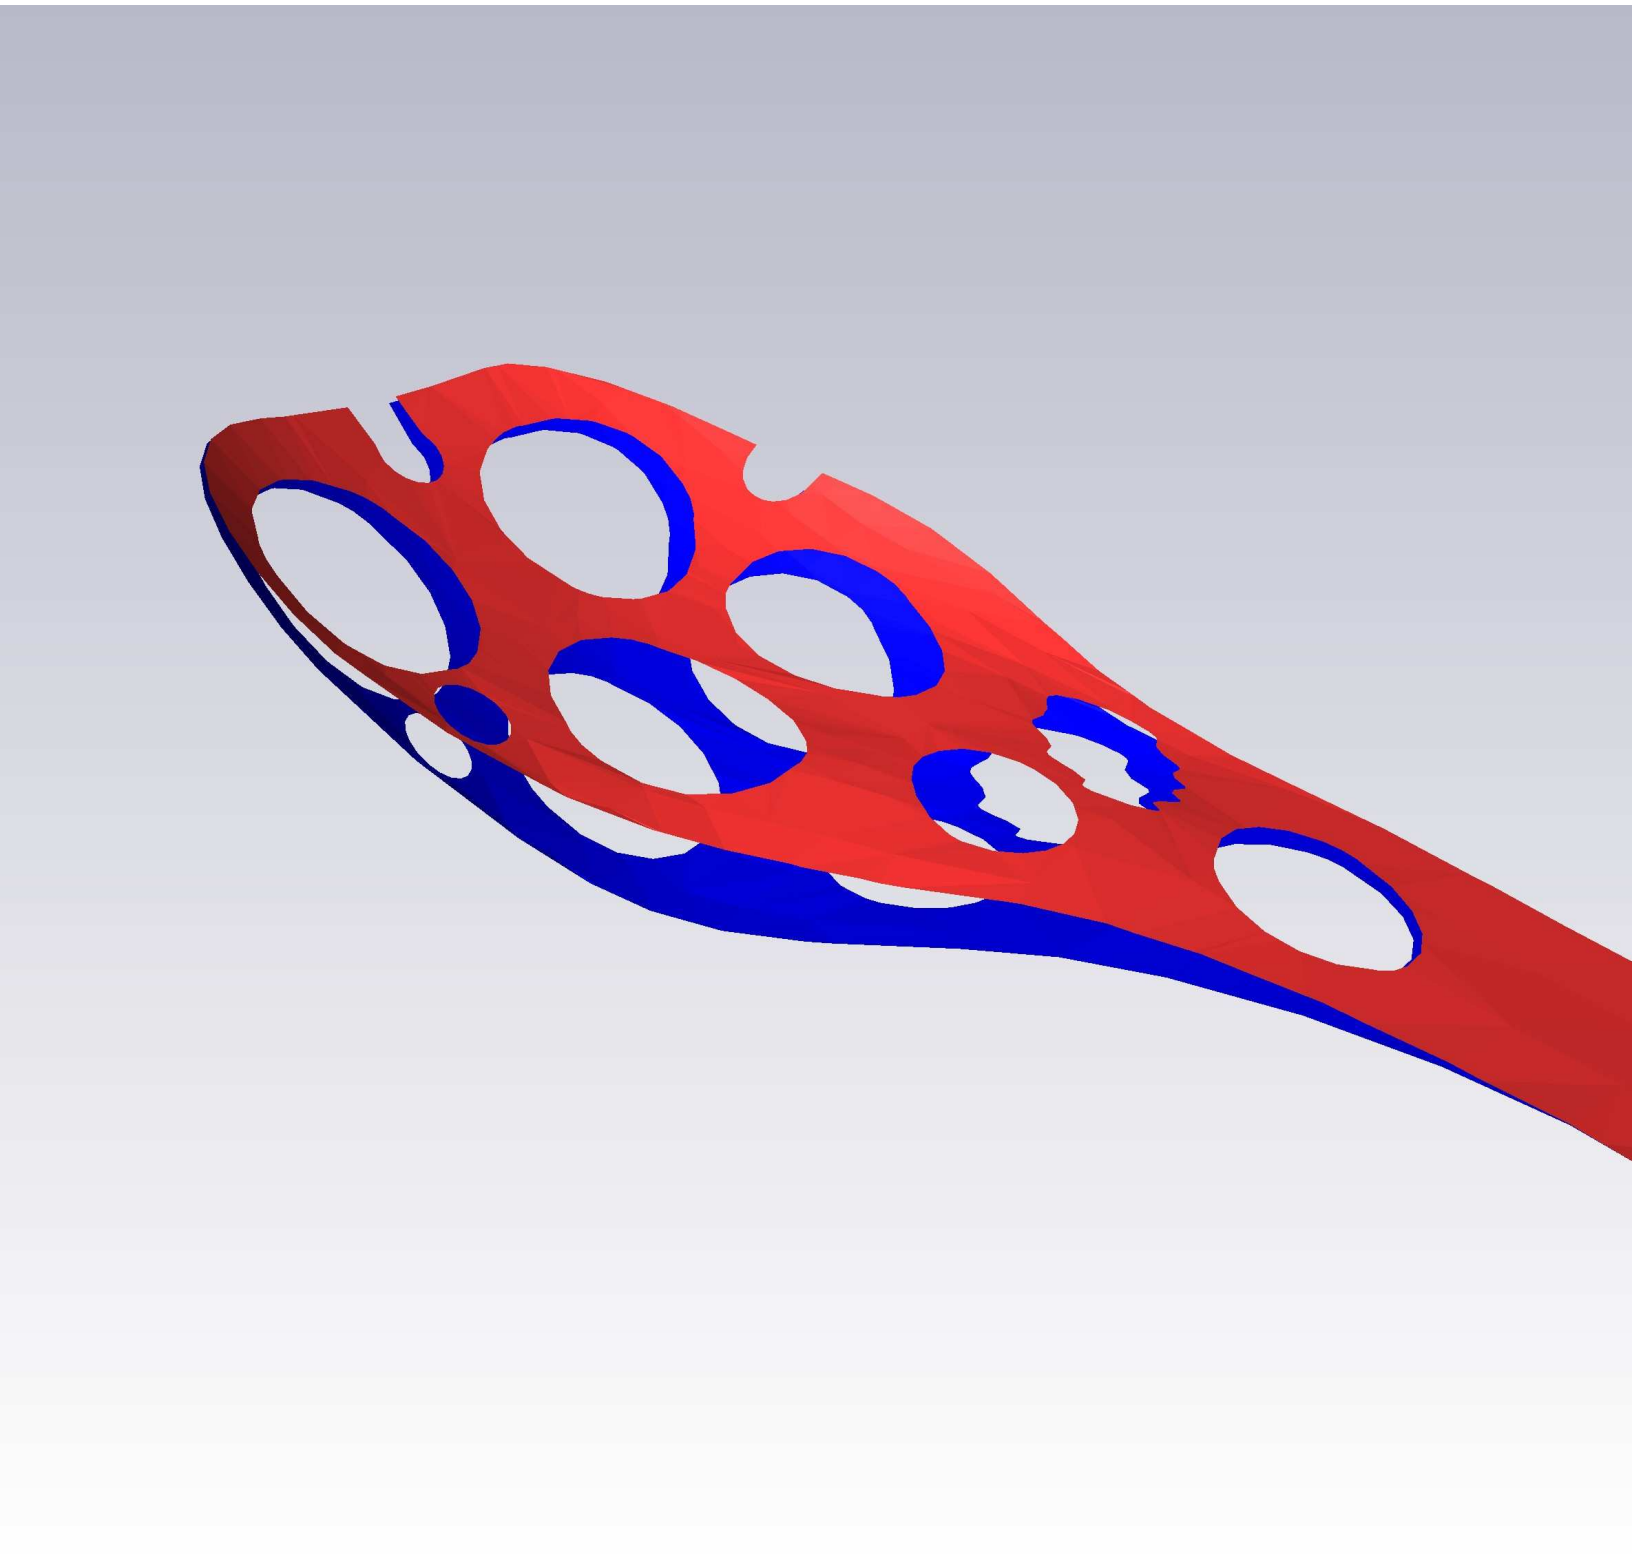

Figure16b

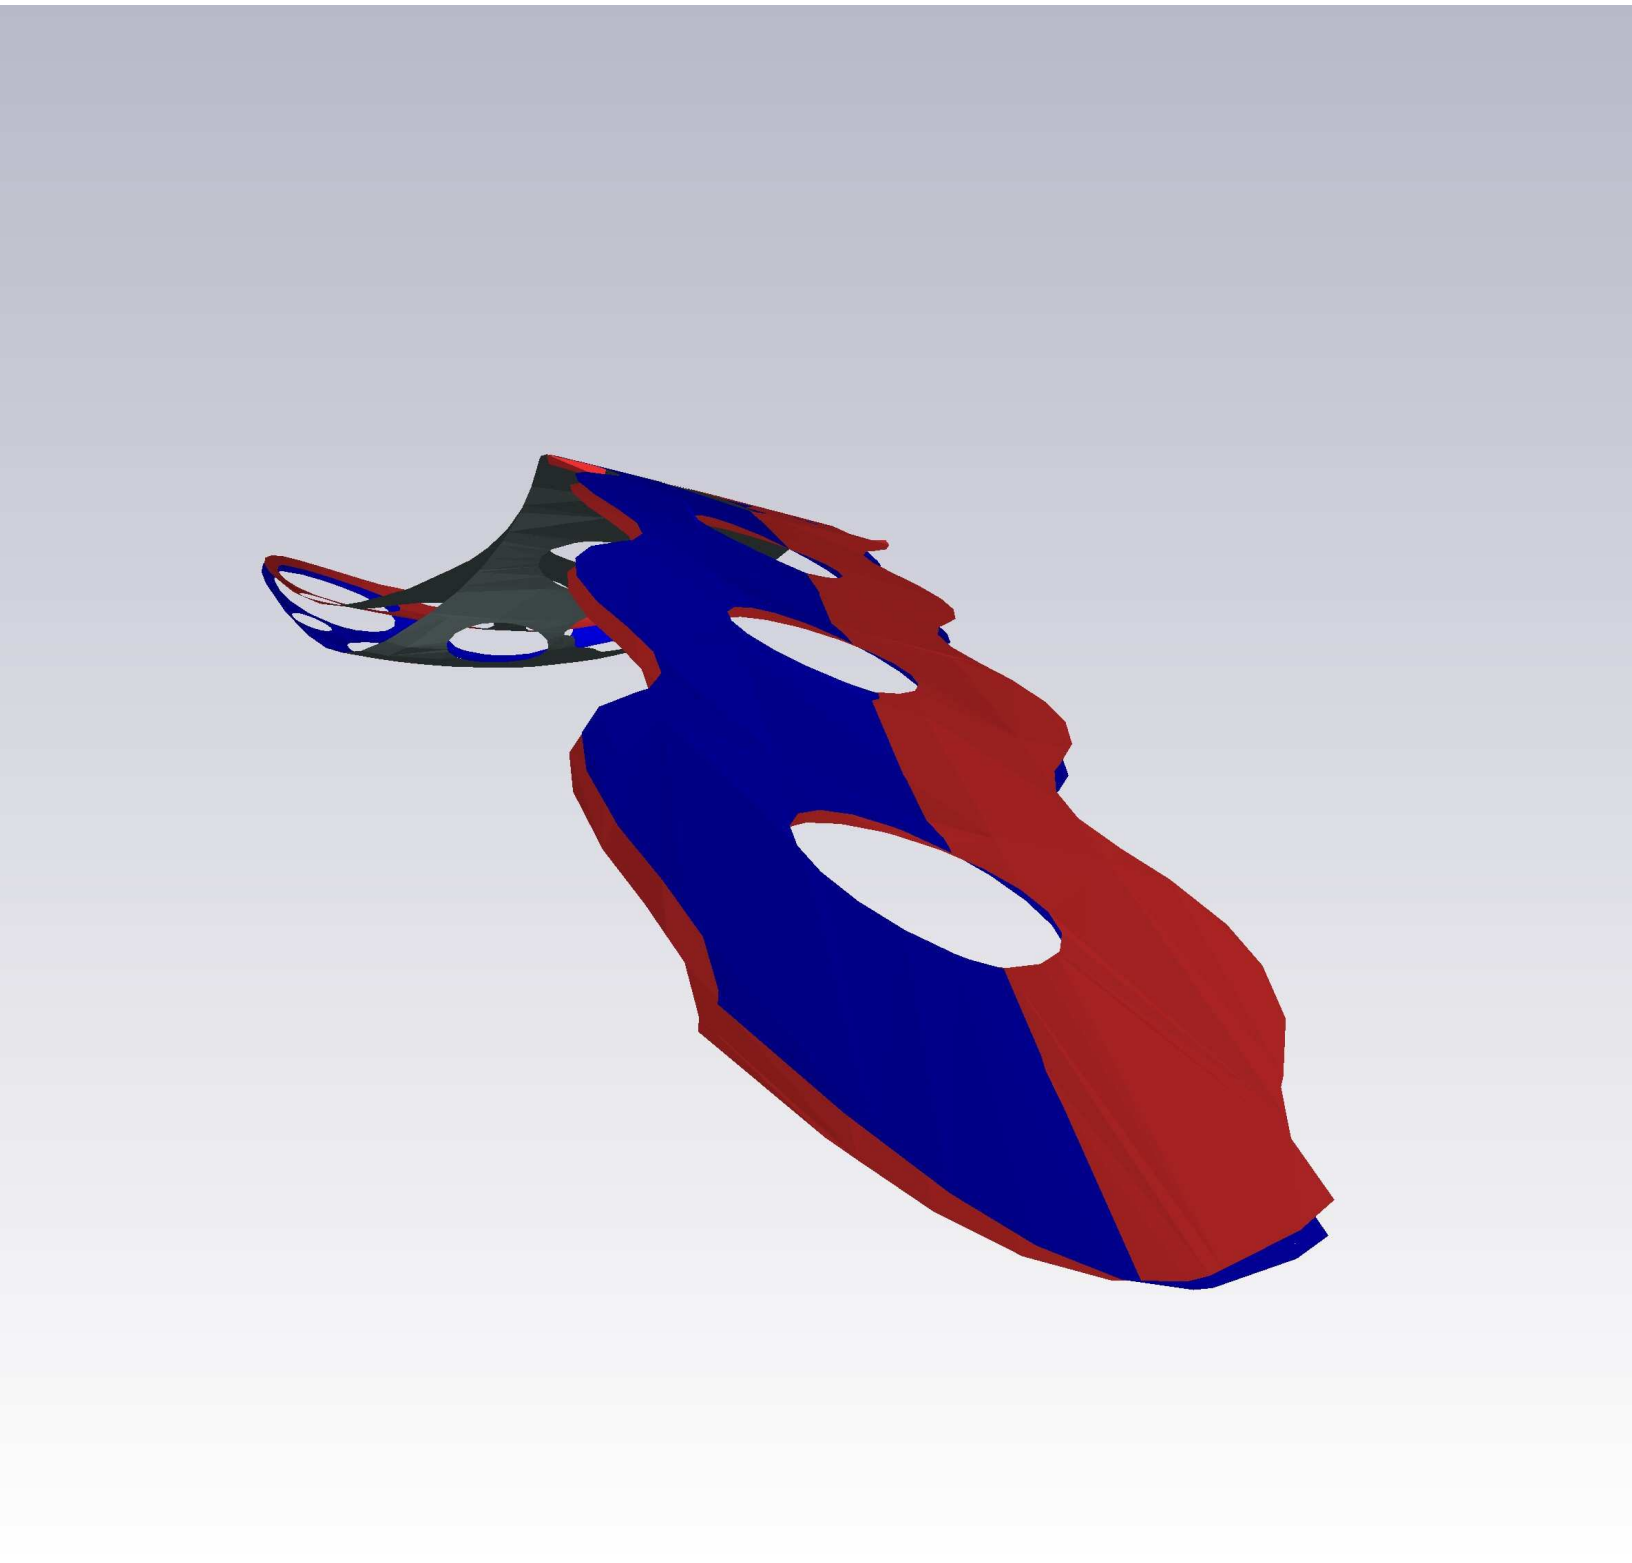

Figure captions

Click here to download Supplemental data for online publication only: Figure_Captions.pdf

Figure software

Click here to download Supplemental data for online publication only: Figure_Software.pdf