Detection and Visualization of Surface-Pockets to Enable Phenotyping Studies

Kishore Mosaliganti*, Firdaus Janoos, Richard Sharp, Randall Ridgway, Raghu Machiraju, Kun Huang, Pamela Wenzel, Alain deBruin, Gustavo Leone, and Joel Saltz

Abstract—In this paper, we propose a technique for detecting pockets on a surface-of-interest. A sequence of propagating fronts converging to the target surface is used as the basis for inspection. We compute a correspondence function between the initial and the target surface. This leads to a natural definition of the local feature size measured as the evolution distance between mapped points. Surface pockets are then extracted as salient clusters embedded in the feature space. The level-set initialization also determines the scale-space of the extracted pockets. Results are presented on a case-study in which the focus is to chronicle the phenotyping differences in genetically modified mouse placenta. Our results are validated based on manually verified ground-truth.

Index Terms—Computational phenotyping, level-sets, surface-pockets.

I. INTRODUCTION

ADVANCEMENTS in medical image analysis, computer graphics, and pattern recognition research over the last decade have permitted the segmentation, reconstruction and visualization of 3-D structures such as vasculature, anatomical surfaces, and tissue-tissue interfaces. With the aid of computerized morphometry algorithms, it is possible to quantify and visualize morphological features such as surface areas, curvature, thickness, and volume. In this paper, detection of pockets that are present on surfaces is of specific interest to us. Pockets refer to the localized regions on the surface that are similar to finger-like infiltrations on either side of the surface. The presence of pockets has a direct correlation with surface morphological parameters such as interface surface-area, convolutedness, and the extent of tissue infiltration. We are interested in quantifying the 3-D finger-like infiltration that occurs on the labyrinth-spongiotrophoblast tissue interface of the mouse placenta [Fig. 1(a) and (b)].

Several research problems benefit from surface pocket analysis. Protein docking for structure-based drug design is an important application wherein the high-resolution structural configurations of a protein and a drug are separately analyzed for fitting the two molecules together in 3-D space [4]. The optimal fit is characterized by a concave pocket on the surface of the protein into which the drug may dock effectively. Tracking the temporal growth of dendrites and spines on the surfaces of neuronal cells is another recent problem [20].

The detection of surface pockets is complicated owing to several reasons. There is no formal definition in medical or computer vision literature and it is not easily represented by a mathematical model or distribution. Further, determining the correct scale for pocket exploration is not trivial. Depending on the scale, pockets could be classified as high-frequency noise. Also, erroneous segmentation/reconstruction creates artifacts that are very similar to pockets. From a visualization standpoint, occlusion is a major obstacle, especially for concave intrusions into the surface.

It may be pointed out that one can alleviate the problem of visualizing the surface pockets to some extent through a choice of viable shading techniques, light sources, and opacity transfer functions (Fig. 1(c), bottom). While this perspective is valid, our focus is on the objective quantification and validation of the observed surface pockets and, therefore, explicit extraction was deemed necessary (see Fig. 6). In our approach, we propose evolving a front in close vicinity of the target surface. (Please refer to Fig. 1(c), top.) The front initially represents a global shape of the surface without pockets. As the front progresses towards the target surface, it acquires the features on the surface and finally converges to it. This leads to a natural definition of feature size as the distance traveled by a point from the initial front to the target surface. Surface pockets have larger feature sizes compared to the flat regions owing to the larger distances traversed. Hence, they are suitably extracted.

II. RELATED WORK

The idea of employing curvature in the specification of TFs has been presented by Hladůvka *et al.* [6] and Kindlmann *et al.* [10]. Such a TF helps a user to enhance or suppress shapes of a specific curvature class. However, pockets are not necessarily characterized by specific curvature values. They are *deep* infiltrations on the surface. Kanda *et al.* [9] and Zhou *et al.* [23] used distances as one of the dimensions in the specification of TFs. Seed point locations were specified by the user or implied by the camera position, and distance fields were computed based
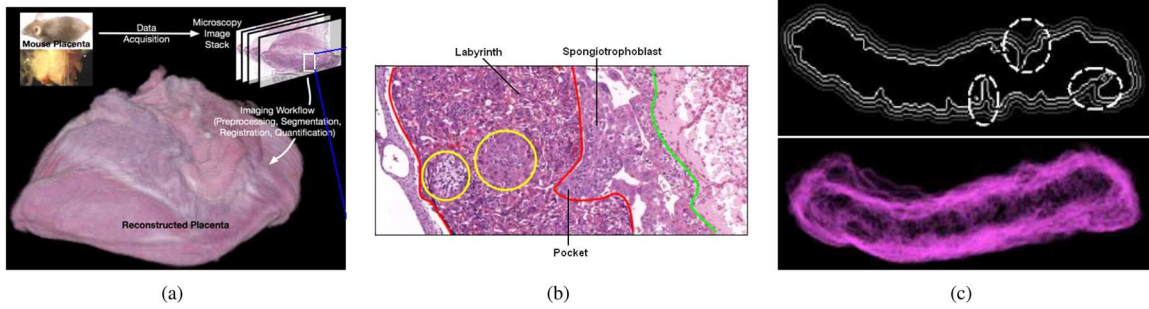
Fig. 1. (a) Reconstructed mouse placenta with the imaging workflow. (b) A slice of the placenta tissue showing the labyrinth-spongiotrophoblast interface (red). The interface presents pocket-like infiltrations and extrusions. The yellow circles show glycogen tissue embedded within the labyrinth. The green boundary is the interface between the spongiotrophoblast and the remaining maternal layers. (c) Top: An initialized front (outermost) that evolves into the target surface (innermost contour). Notice that infiltrations (within circles) are initially absent and gradually develop on the front. Below: 3-D rendering of many consecutive labyrinth tissue contours with a constant opacity.

on these seed points. Our work is related to the spatial TFs in terms of our overall goal to focus on localized regions in the 3-D image space. We, however, require a different approach as the pockets locations are not explicitly available.

Skeletonization approaches [11], [12] in 3-D may be considered for detecting pockets. The pockets are extrusions from the surface and will reflect as extended branches on the medial axis backbone. We have implemented a similar approach for 2-D neuron images (with dendritic spines) in the recent past [13]. A skeletonization approach presents several conceptual problems. 1) There is no direct method of computing the true medial axis for a stack of 2-D edge contours—only approximations on polyhedral models are realizable. 2) The medial axis is very sensitive to surface variation. Minor variations on the surface cause significant changes to the medial axis resulting in high false-positive rates of detection. Given the large data sizes, such an approach may even prove intractable.

A good introduction to level set techniques and applications can be found in [17]. In [1], Caselles *et al.* proposed the *geodesic active contours* model with a level sets based implementation. It extended previous work on active contour models with a new component in the speed function derived from the image properties. The new component creates an attraction force that caused the contour to *halt* at the object boundaries. We adopt this formulation for propagating the fronts.

Edelsbrunner *et al.* [4] provided a mathematical definition of pockets in proteins and other macromolecules. Pockets were constructed by an algorithm based on alpha complexes, and the algorithm was applied to proteins with known three-dimensional conformations. This algorithm uses information specific to proteins and is not immediately suitable for general pocket identification.

III. SURFACE POCKET MODELING

We attempt to capture the intuitive notion of a surface pocket as a narrow and deep infiltration of a surface. A surface pocket is similar to the high-frequency noise in that both are infiltrations or extrusions present on the surface. However, certain differences exist between noise and a surface-pocket that allow for easier mutual discrimination. Fig. 2 shows different kinds of infiltrations. Clearly, the only differentiating factor among them
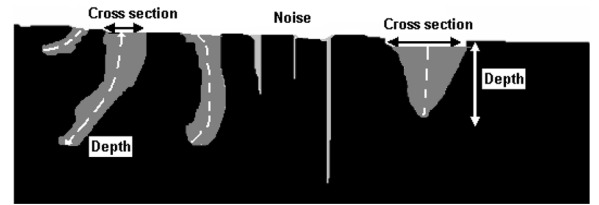

Fig. 2. Different types of infiltrations characterized by their height and depth. Note that the pockets are characterized with higher aspect ratios (cross-section area to depth) as compared to noise.

is their depth of infiltration and the cross-sectional diameter at the neck. This leads us to make the following observations.

1) The depth of the infiltration needs to be computed relative to a similar surface without the infiltration.
2) A notion of scale needs to be incorporated to discriminate pockets with narrow cross section.
3) The ratio of depth to cross-sectional area is important to locate spurious pockets.

In Section III-A, we define a measure of pocket feature size and in Section III-B, we describe a way of measuring the convolutedness of a surface. Section III-C describes the model used for identifying a pocket, after we have measured the feature sizes of all the points on the surface. Please note that a pocket is composed of points that have a certain feature size.

A. Feature Size

Let Γ be a surface front in $\mathbb{R}^3$ propagating with a well-defined speed F such that $\Gamma(t)$ gives the position of the front at time t. $\Gamma(t)$ is the zero level-set of the signed distance function $\psi(t)$: $\mathbb{R}^3 \to \mathbb{R}$.

Let $d_{\mathrm{ev}}(\Gamma(t), \mathbf{r})$ denote the evolution distance of the point $\mathbf{r} \in \Gamma(t)$ from a corresponding point in $\Gamma(0)$. We define the *local feature size* of a point $\mathbf{r}$ on the final surface $\Gamma(T)$ (i.e., the target surface S) as the evolution distance $d_{\mathrm{ev}}(\Gamma(T), \mathbf{r})$. We now describe methods for computing the local feature size.

1) Evolution Distance Metric: To accurately measure the depth of curved pockets, the computation of the feature size needs to follow the evolution of the surface front. In general, the propagation paths may be curved [5]. Fig. 3 illustrates the need of measuring the distance traveled by a point on the final surface along the successive fronts. The point on the pocket $\mathbf{r_1}$
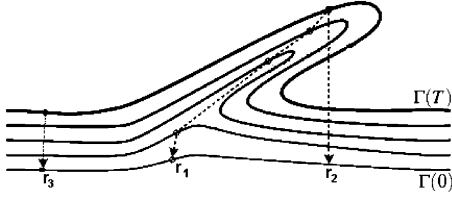
Fig. 3. The image shows three points that are being tracked. Points $\mathbf{r_1}$ and $\mathbf{r_2}$ lie on the pocket. While $\mathbf{r_1}$ evolves along a curved path, $\mathbf{r_2}$ follows a Euclidean path between the initial and final contours. Point $\mathbf{r_3}$ lies on the flat portions of the surface and evolves along a Euclidean path as well.

travels along a curved path. However, simply measuring the distance to the nearest point on the starting front, $\mathbf{r_2}$ is not accurate. Therefore, we need to track the evolution of each point on the initial front (surface without pockets) to the final front (surface with pockets).

A good approximation for the evolution distance is obtained by determining the distance of a point on the front at a given time point to the nearest point on the front at an earlier time-point. The values may then be accumulated across successive fronts to yield the final traversed distances.

Formally, we define the mapping function for a point $\mathbf{r} \in \Gamma(t)$ to the corresponding point $\mathbf{p} \in \Gamma(t-1)$ as

$$\begin{aligned} f_t &: \Gamma(t) \to \Gamma(t-1) \\ f_t(\mathbf{r}s) &= \arg \min_{\forall \mathbf{p} \in \Gamma(t-1)} (\|\mathbf{p} - \mathbf{r}\|) \end{aligned} \tag{1}$$

where $\| \cdot \|$ represents the Euclidean distance metric. The evolution distance is recursively defined as

$$d_{\text{ev}}(\Gamma(t), \mathbf{r}) = d_{\text{ev}}(\Gamma(t-1), f_t(\mathbf{r})) + \|\mathbf{r} - f_t(\mathbf{r})\|. \tag{2}$$

2) Euclidean Distance Metric: Usually, pockets are not oriented perpendicular to the geometry. Hence, the distance needs to be computed along the evolution path. However, pockets occupy only a small fraction of the surface area and the remaining area evolves in a near normal direction. For example, consider the point $\mathbf{r_3}$ on the flat regions of the surface in Fig. 3 The Euclidean path coincides with the evolution and the distance traversed may be computed directly rather than recursing through the different fronts. Hence, using a direct mapping f from the final to the initial front is sufficiently accurate

$$\begin{aligned} f &: \Gamma(T) \to \Gamma(0) \\ f(\mathbf{r}) &= \arg \min_{\forall \mathbf{p} \in \Gamma(0)} (\|\mathbf{p} - \mathbf{r}\|) \end{aligned} \tag{3}$$

where T is the total time expended.

3) Combined Distance Metric: The method in Section III-A1 computes the evolution distance accurately. This method requires the determination of matching pairs of points across successive contours resulting in high computational costs. The method in Section III-A2 computes the direct Euclidean evolution distance efficiently since correspondence needs to be established only for the initial and the final fronts. However,

the method is accurate only in the nonpocketed parts of the contour. We therefore propose the *combined distance metric* $d_C(\Gamma(t), \mathbf{r})$ to compute the feature size that incorporates the advantages of both the methods.

We start by computing the Euclidean metric between the final and initial surface fronts. Pocketed regions exhibit high distances although their exact values may not be correct. A threshold (τ) is applied to extract such regions, which are then re-evaluated using the evolution distance metric. The threshold is selected based on the nature of the surface to be identified. A lower threshold will evaluate shallower variations in the surface using the evolution distance. A higher threshold will make the distance computation more efficient, at the risk of missing out pockets that are shallow or nonorthogonal to the surface.

The initial front $\Gamma(0)$ is now selected based on a *convolutedness criterion* that is described next.

B. Convolutedness Criterion

We wish to select an initial front $\Gamma(0)$ as a reference for making measurements of the depth of the pockets on the surface S. This initial front must have the property that while it approximates the overall shape of the surface S, it should not contain pockets. The distance traversed by a point is measured relative to this front.

As the initial front becomes smoother (aggregating over the high-frequency variations in S), the computed feature size function changes. Initially, the feature size function at the pockets has a high rate of change while it is fairly stable for nonpocketed regions. This rate of change stabilizes throughout the surface as the initial front becomes even smoother and devoid of surface pockets. This initial smoothed surface is computed by extracting an iso-surface in the distance map (D) of the surface S (see Section IV–B for more details). Larger iso-values in the distance map of the surface result in smoother iso-surfaces.

We employ a *surface convolutedness metric* in selecting such a suitable front [17]. Let $\Gamma(u, v)$ be a surface front parameterized by the variables, and whose convolutedness is determined by the function $\text{conv}(\Gamma)$. Further, let $\rho(u, v)$ represent the local mean curvature of the surface at position (u, v). The following metric determines the extent of "wrinkling" or oscillations present in the surface:

$$\text{conv}(\Gamma) = \frac{\oint_\Gamma |\rho(u, v)| g(u, v) du\, dv}{\oint_\Gamma g(u, v) du\, dv} \tag{4}$$

where g is a "stretch" function given by

$$g(u, v) = \sqrt{x_{uv}^2 + y_{uv}^2 + z_{uv}^2}. \tag{5}$$

We compute surface principal curvatures using the algorithm given by Chen and Schmitt [2].

Fig. 4 shows a graph of the ratio of the convolutedness of the smoothed initial front $(\Gamma(0))$ to the final surface (S). For an iso-value of 0 in the distance map D, the initial surface is the actual surface and hence the ratio is 1. The ratio is initially high as expected and then begins to drop monotonically. There is a distinctive knee-point in the convolutedness ratio as the initial
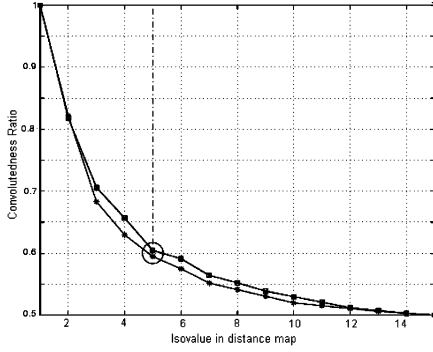
Fig. 4. The graph in this figure demonstrates how the ratio of the convolutedness metric starts to drop for initial fronts that are progressively smoothed versions of the surface. Larger iso-values in the distance map of the surface result in smoother initial fronts. The wild-type and mutant placenta data are marked with a plus and a box, respectively.

front is further smoothed after which point the surface does not get any less convoluted. This point is marked with a circle in Fig. 4. We have found that selecting a front at the knee-point yields good results (see Section VI). The convolutedness metric rapidly drops and saturates at about same iso-value in both the cases since the pockets constitute a small percentage of the surface. Our framework was implemented on four placenta datasets and, therefore, manually selecting the knee-point was considered appropriate. A fully automatic implementation was not the aim although it is possible to realize it. On the other hand, some user inspection always helps. The knee-point is intuitive since the disappearance of a large number of pockets is expected to be correlated with higher value of the convolutedness ratio.

C. Pocket Identification

After having computed the feature size at all the points on the target surface S, the next step involves using this information to identify the individual pockets. For this purpose, we select points that have feature size (d_C) above a certain threshold Θ_d, as candidates for belonging to a possible pocket.

To enforce the definition of a pocket as a narrow and deep infiltration or extrusion on the surface, we apply the following constraints on the pocket:

1) *The maximum allowable cross-sectional diameter at the base of the pocket*. To determine the cross section at the base, we take all the points in the pocket with the lowest feature size. If the cross-sectional diameter is above a lower limit (Π_{cs}), then the pocket is considered invalid.
2) *The minimum allowable depth of a pocket*. The depth of the pocket is defined by the largest feature size of the points in that pocket. Typically, the point at the tip of the pocket has the largest feature size. If the depth of the pocket is below a lower limit (Π_{depth}), then the pocket is considered invalid.

IV. LEVEL SET FORMULATION

We use the level set method [1], [17] as a powerful numerical technique for implementing the front propagation scheme. The two key steps for the level set method are as follows.

1) *Embedding the surface:* We represent the surface $\Gamma(t)$ as the zero level-set of a signed distance function $\psi(\mathbf{x}, t)$, i.e., $\Gamma(t) = \{\mathbf{x} : \psi(\mathbf{x}, t) = 0\}$.
2) *Embedding the motion*: We need to derive the time evolution PDE such that the motion of the zero level set has the desired properties as described below. For this purpose, we use the active contour formulation developed by Caselles *et al.* [1]

$$\psi_t + g(P)(\beta F_A + \gamma\rho)|\nabla\psi| + \alpha\nabla g \cdot \nabla\psi = 0. \quad (6)$$

Note that ψ_t refers to the temporal derivative of the signed distance function ψ.

The attraction term $(\nabla g \cdot \nabla\psi)$ causes the zero-level set to be attracted to the object boundaries. The advection force term F_A is independent of the geometry and causes the front to expand or contract depending on the local gradient magnitude. The diffusion term F_G is proportional to the local curvature ρ. The parameters α, β, and γ are user-defined settings for the relative scaling of the three speeds. The speed function is scaled with a quantity $g(P) = (1)/(1 + |\nabla G_\sigma * P|)$ that has values proportional to the image gradient.

A. Segmentation Probability Maps

We employ material characterization techniques from the material science field to characterize microscopic structure (Torquato [19]). Recently, Ridgway *et al.* [8], [16] investigated the use of the N-Point correlation functions for medical image segmentation to generate probability maps of the different tissue types. Our region of interest is characterized by a high value (close to 1) in the probability map P while the remaining region has a lower value.

B. Initial Surface Determination

We require an initialization that captures the shape of the labyrinth layer without the presence of surface pockets. Our distance-based methods will only be able to provide accurate measurements of feature size under such conditions. Therefore, we adopt the following procedure to obtain such an initialization. Fig. 5 illustrates the method on a single slice in 2-D.

A threshold (T_S) is applied to the smoothed version of the probability map P. Smoothing reduces the number of floating components in the image. The largest component is chosen as representative of the labyrinth layer. The obtained object-of-interest captures the overall shape of the labyrinth layer. However, the surface is not smooth and shows interdigitation (pockets). To obtain a smooth version of the surface, we first construct a Euclidean distance map from the surface [3]. An iso-value in the distance map provides a smoother version of the surface that captures the overall shape of the labyrinth structure without the pockets. The choice of an appropriate iso-value (ω) is guided by the knee-point transition in the convolutedness ratio. Please refer to Fig. 4 in Section III-B. The graph shows that for large iso-values, the surface gets less convoluted. At the knee-point, the surface is sufficiently smooth with no or few pockets. The
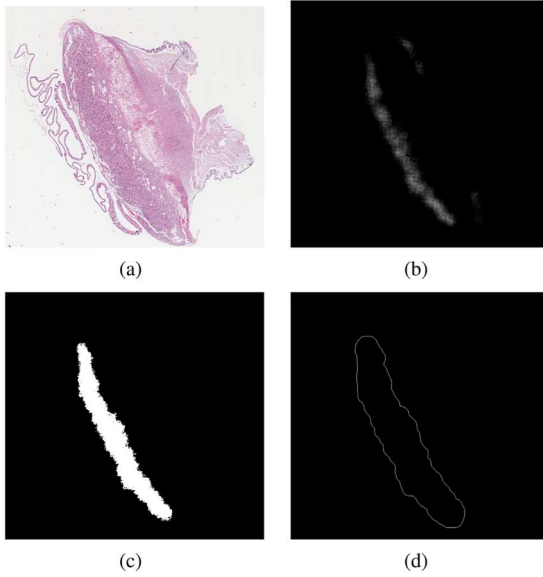
Fig. 5. (a) A slice of the 04-2069 placenta. (b) Segmentation probability map of the labyrinth layer. (c) Thresholding the probability map and choosing the largest component provides a contour marked with interdigitation. (d) A smoother version is obtained as an iso-contour in the Euclidean distance map of the surface.

iso-value parameter ω is therefore initialized to the knee-point and this provides us with the initial front.

V. IMPLEMENTATION

Algorithm 1 Pipeline for pockets

// Get probability map

$P \leftarrow \text{NPointCorrelation}(I)$

//Use convolutedness metric to identify initial front

$\Gamma(0) \leftarrow \text{InitialSurfaceExtraction}(P, T_S, \omega)$

//Extract the set of all evolving fronts

$\Gamma(\cdot) \leftarrow \text{rm LevelSetMethod}(P, \Gamma(0), \sigma, \alpha, \beta, \gamma)$

// Compute the feature size at each point

$D \leftarrow \text{CombinedDistance}(\Gamma(\cdot), T, \mathbf{r}(T), \tau)$

// Identify pockets using the constraints in Section III-C

// $F \leftarrow \text{PocketIdentification}(D, \Theta_d, \Pi_{cs}, \Pi_{\text{depth}})$

The overall pipeline is described in Algorithm 1. We begin by obtaining a probability map for the labyrinth layer in the volume. A starting front $(\Gamma(0))$ is constructed from the distance map of the labyrinth layer, to initialize the level-set propagation (Section IV-B). We apply the convolutedness criterion (Section III-B) to select this initial front. The probability map is also used to obtain motion laws for propagating the front (Section IV). The front halts when it encompasses the labyrinth layer $(\Gamma(T))$. The distance metric (Section III-A3) is computed between the initial front $\Gamma(0)$ and final front $\Gamma(T)$ to determine

TABLE I
VALUES USED IN THE ITK IMPLEMENTATION

T_S	Probability map threshold	0.4
ω	Iso-value parameter for initial front	7
τ	Threshold for combined distance metric	$\mu_1 + \sigma_1$
Θ_d	Threshold for feature size	$\mu_2 + \sigma_2$
Π_{cs}	Maximum allowable cross-sectional diameter	8
Π_{depth}	Minimum allowable depth	12
σ	Gaussian kernel width	1.0
α	Attraction parameter	1
β	Advection term parameter	1
γ	Diffusion parameter	2.0

the feature size. Finally, pockets are extracted using the method described in Section III-C. Table I lists the typical parameter settings that were used in our experiments.

Note: The parameter τ is estimated after using the Euclidean measure to determine contour points that are recomputed using the evolution metric. Note that μ_1, σ_1 are the mean and standard deviation respectively of the resultant feature values after using the Euclidean measure. Later on, after applying the evolution measure, the parameter Θ_d is initialized to separate the pocketed and nonpocketed regions. Note that the μ_2, σ_2 are the final mean and standard deviation of the feature values.

We implemented our framework using the National Library of Medicine's (NIH/NLM) Insight Segmentation and Registration Toolkit (ITK) [7] and the Visualization Toolkit (VTK) from Kitware Inc. We used the ITK implementation of the level-set formulation by Caselles *et al.* [1]. The feature vectors at each voxel are loaded into Kitware's VolView volume visualization software to render the surface appropriately. All our experiments were conducted on a 2.5-GHz Pentium machines running Linux with 1-GB main memory.

VI. CASE STUDY: GENETIC PHENOTYPING OF THE MOUSE PLACENTA

In this paper, we are interested in quantifying the finger-like infiltration that occurs on the labyrinth-spongiotrophoblast tissue interface of the mouse placenta. Please refer to Fig. 1. The infiltration affects the oxygen and nutrient exchange taking place between the maternal and fetal tissues. The normal tissue interface presents well-developed pockets that are also well distributed over the interface. A recent study has revealed that with the knockout of the *Rb* gene, the interdigitation (pockets) wildly increases and clumps together with poor infiltration [22]. As a results, fetal death occurs. Our biomedical collaborators are interested in quantifying this phenotype difference [21].

Two pairs of wild-type and mutant (Rb^-) placenta were harvested at 13.5 days of gestation and prepared using a standard histological protocol. They were fixed in formalin, embedded in paraffin and sectioned at 5 μm thickness using a microtome. Serial sections were mounted on glass slides and scanned at 200× magnification using a light microscope. The image dimensions on average were (15 K × 15 K × 3). Each placenta data-set produced approximately 600–800 color images, ranging in size from 300 to 500 GB. The images are extremely noisy and require preprocessing, down-sampling, image registration, and tissue segmentation prior to the visualization stage. Considerable efforts were undertaken to prove the medical hypothesis that has led to new algorithmic
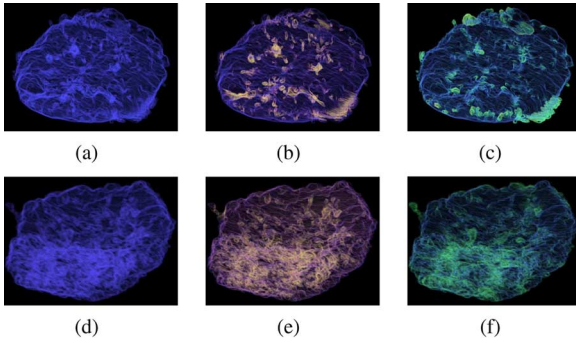
Fig. 6. The top row contains images of normal mouse placenta labyrinth while the bottom row shows mutant (Rb^-) placentas. The left column (a,d) show images of the overall mouse labyrinth layers; no specific infiltrations or extrusions in the tissue can be seen. The middle column (b,e) highlights the infiltrations in increasing intensity from blue to yellow. The right column (c,f) show the result of highlighting extrusions in increasing intensity from blue to green.
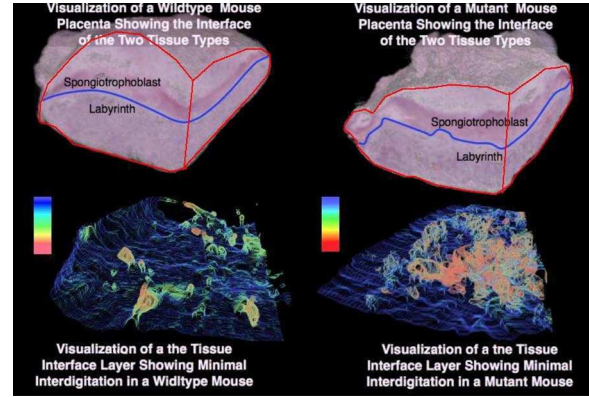

Fig. 7. The top row shows the RGB visualization of the 05-829 wild-type and the 05-1903 mutant placentae. The 3-D volume has been sectioned using two vertical planes to reveal the two different tissue layers at a macroscopic resolution. The interface was then visualized using our framework in the bottom row. Note the well-defined infiltrations in the wild-type and the randomly clumped pockets in the mutant placenta.

developments. Mosaliganti *et al.* [14] developed a two-level optimizer for the image stack registration. Pan *et al.* [15] and Ridgway *et al.* [8], [16] independently developed probabilistic segmentation methods that helped in extracting the placenta tissue layers. Registration and segmentation methods were used by Sharp *et al.* [18] for visualizing the placenta in 3-D. In this work, we realize the final goal of visualizing and quantifying the infiltration in the labyrinth-spongiotrophoblast interface by detecting surface pockets.

A. Experiments and Results

This section reports on a number of experiments that were conducted to evaluate the utility of the framework towards pocket detection. Our goals are twofold. 1) We show that superior visualization is obtained by incorporating the feature size information into a volume rendering application. 2) Validation studies are performed in 2-D/3-D with manually marked up pockets on the labyrinth contour.

The image shown in Fig. 1(c), bottom shows the 3-D rendering after a large portion of the placenta is cropped away. The infiltrations and extrusions are occluded by the rest of the labyrinth surface. Fig. 6 shows the results of highlighting infiltrations and extrusions on the labyrinth surface in wild-type (05-977) and mutant (Rb^-) (05-1904) mouse placenta. Specifically, Fig. 6(a) and (d) shows the entire original surface of the labyrinth in blue. No infiltrations or extrusions are obvious in the image. The image in Fig. 6(b) and (e) highlights infiltrations in opaque yellow using an outward marching level-set to obtain feature size information. Similarly, the image in Fig. 6(c) and (f) uses an inward marching level set to highlight extrusions on the data. Fig. 7 provides a more distinct phenotyping difference in the interface topology. The following observations can be readily made. 1) The wild-type placenta seems to have well-developed infiltrations/extrusions that run deep into the other tissue. The mutant placenta is however marked by their absence. Instead, there is random interdigitation that is clumped together at different sites. 2) It also can be seen that the wild-type labyrinth has more uniformly distributed features than the mutant placenta.

We validate our results using manually segmented surface pockets that are present in a contiguous section of the 3-D stack

of images. The contiguous section is selected approximately from the center of the stack. The central section is usually associated with better consistency in staining, section thickness and reduced deformations of the tissue. This usually results in better segmentation and higher registration accuracy. Surface pockets were then identified in sections that were drawn from each placenta dataset. Each detected pocket in this stack can be placed in four categories: 1) a counts the number of pockets that are marked in the ground-truth and fit our model (true positives); 2) b counts the pockets that are not marked in ground-truth but fit our model (false positives); 3) c counts the pockets that are marked-up but do not fit our model (false negatives); and 4) d counts the pixels that are neither marked nor fit our model (true negatives). Note that a high value of specificity $((d)/(b+d))$ indicates that it is easy for the algorithm to rule out a spurious pocket with a high probability. On the other hand, a value of high sensitivity $((a)/(a+c))$ can identify a bona fide pocket with a high probability of success. The error may be defined as $(b+c)/(a+b+c+d)$.

Table II tabulates the sensitivity, specificity, and error rates observed for different settings of the minimum allowable pocket depth (Π_{depth}) and the maximum allowable cross-sectional base diameter (Π_{cs}). The following observations are made.

1) Within every placenta category, as the range of allowable pocket sizes widens ($\Pi_{\text{depth}} = 14$) or ($\Pi_{\text{cs}} = 10$), the sensitivity increases and the specificity decreases as compared to the highlighted case when ($\Pi_{\text{depth}} = 12$ and ($\Pi_{\text{cs}} = 8$). This is an expected trend since nonpocket features now fit our constraints thereby lowering specificity. Simultaneously, all true pockets also readily fit the model, thereby increasing the sensitivity.

2) Across the placentas, we always obtain a good sensitivity with moderate specificities. Error rates are maintained below 0.10. The wild-type placentas (05-829 and 05-1903) have well-developed fingers and therefore have better detection rates as compared to the mutants (05-977 and 05-1904).

3) Fig. 8(a) plots the family of ROC iso-contour curves of the Π_{cs} parameter in the 05-829 dataset for different settings

TABLE II
SENSITIVITY AND SPECIFICITY VALUES FOR POCKET DETECTION

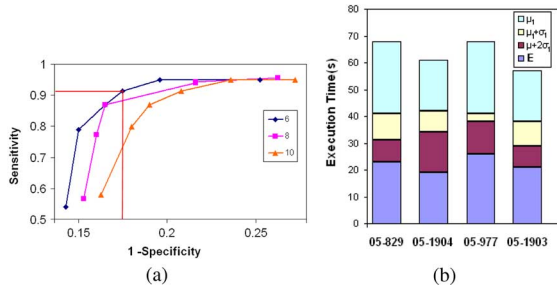| Dataset	Depth	Diameter	Sens.	Spec.	Error
05-829	≥ 14	≤ 8	0.9223	0.8253	0.0324
	≥ 12	≤ 8	**0.9134**	**0.8527**	**0.0215**
	≥ 12	≤ 10	0.9690	0.8041	0.0785
05-1904	≥ 14	≤ 8	0.9021	0.8475	0.0542
	≥ 12	≤ 8	**0.8907**	**0.8822**	**0.0433**
	≥ 12	≤ 10	1.000	0.8173	0.0851
05-977	≥ 14	≤ 8	0.9853	0.8238	0.0662
	≥ 12	≤ 8	**0.8594**	**0.8735**	**0.0544**
	≥ 12	≤ 10	0.9637	0.7835	0.0951
05-1903	≥ 14	≤ 8	0.9522	0.8325	0.0483
	≥ 12	≤ 8	**0.9162**	**0.8756**	**0.0357**
	≥ 12	≤ 10	1.000	0.5281	0.0642

Fig. 8. (a) ROC family of curves for different Π_{depth} settings with $\Pi_{cs} = 6, 8,$ and 10 respectively. An optimal setting is given by $\Pi_{cs} = 8$ and $\Pi_{\text{depth}} = 12$ with sensitivity and specificity values at (0.9134, 0.8527) (red markers). (b) Effect of different settings of the τ parameter on the time-performance while using the combined distance metric accumulated over the Euclidean metric (denoted by E when $\tau = \infty$).

TABLE III
EXECUTION TIME IN SECONDS FOR THE THREE METRICS

Dataset	Euclidean	Evolution	Combined
05-829	23	221	**41**
05-1904	19	215	**42**
05-977	26	274	**41**
05-1903	21	225	**38**

TABLE IV
THRESHOLDING OF CONTOUR PIXELS

Dataset	finger	non-finger	sens.	spec.	error
05-829-1	288	150	0.8866	0.9028	0.1027
05-829-2	435	206	0.767	0.8506	0.0963
05-1904-1	325	188	0.8564	0.9292	0.0974
05-1904-2	396	252	0.8492	0.9419	0.941

of the Π_{depth} parameter. A step-size of 0.5 was considered while varying the Π_{depth} parameter in the ranges (12–14) while Π_{cs} was varied as 6, 8, and 10. The ROC curve indicate stability in the observed results and validates the choice of their settings.

In Table III, we compare the time performance of using the three distance metrics as observed in the four placenta datasets. It is easy to see that the combined distance metric (Section III-A3) is very close to the Euclidean distance metric (Section III-A2) which is in turn, about 7–10 times faster than the evolution distance metric (Section III-A1). The trend is expected since the evolution mapping operates only in pocketed regions that are a fraction of the overall surface.

In Fig. 8(b), different settings for the threshold parameter τ are explored based on the observed mean (μ_1) and standard deviation (σ_1) of the feature size values. We observe comparable execution performance to the Euclidean metric for three different settings of τ $(\mu_1, \mu_1 + \sigma_1, \mu_1 + 2\sigma_1)$, and yet the time is still a fraction of that expended for the evolution metric (refer to Table III). Hence, this justifies using the combined scheme as a speed-up mechanism. Note that τ was initialized to equal $\mu_1 + \sigma_1$ in the combined metric.

Recall that the threshold Θ_d derived from the final feature values was used to extract the pocket regions on the surface. We validate its choice based on the performance in Table IV. Manual ground-truth for finger and non-finger presence on a voxel basis is generated on two images drawn randomly from both the wild-type (05-829) and mutant placenta (05-1904) stacks. In each image, all pixels on the contour that correspond to a finger are identified manually and the remaining contour pixels are classified as non-finger. These are shown in columns 2 and 3. Using a threshold of $\Theta_d = \mu_2 + \sigma_2$ on the feature size values, we classify the surface pixels again into four categories as described earlier.

On average, we obtain good sensitivity and specificity of 0.84 and 0.90, respectively. The error rate is observed to be around 0.09.

VII. CONCLUSION

In this paper, we describe a method to detect surface pockets. Using a sequence of fronts generated by the level-set technique, evolution distances to the final segmentation are calculated from an initial front. The initial front is selected so that it captures the overall shape of the surface without the presence of pockets. We present a definition of feature size, such that pockets have higher feature size than nonpocketed regions. We employ our framework for mouse placenta phenotyping studies where validation is performed on 3-D pockets found in a normal (wild-type) placenta. The sensitivity of our detection algorithm was measured at $> 85\%$ with a specificity of $> 85\%$. We also use the feature size information to enhance the expressivity of visualization of the surface. In future, we shall explore alternative methods of computing the evolution distance metrics along the fronts. These methods will utilize flow-fields that arise from computing the gradients of the level-set function. Furthermore, we shall consider using alternative generic level set formulations that provide intuitive parameter settings.

REFERENCES

[1] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours.," *Int. J. Comput. Vis.*, vol. 22, no. 1, pp. 61–97, 1997.
[2] X. Chen and F. Schmitt, "Intrinsic surface properties from surface triangulation," in *Proc. Eur. Conf. Computer Vision*, May 1992, pp. 739–743.
[3] P. Danielsson, "Euclidean distance mapping," *Comput. Graph. Image Process.*, vol. 14, pp. 227–248, 1980.
[4] H. Edelsbrunner, M. Facello, and J. Liang, "On the definition and the construction of pockets in macromolecules," *Discrete Appl. Math. Combin. Oper. Res. Comput. Sci.*, vol. 88, pp. 83–102, 1998.
[5] J. Gomes and O. Faugeras, "Reconciling distance functions and level sets," *In Scale-Space Theories Comput. Vis.*, pp. 70–81, 1999.
[6] J. Hladůvka, A. König, and E. Gröller, "Curvature-based transfer functions for direct volume rendering," in *Proc. Spring Conf. Computer Graphics*, 2000, vol. 16, pp. 58–65.

[7] L. Ibáñez and W. Schroeder, The ITK Software Guide. Kitware, Inc., 2003.

[8] F. Janoos, M. O. Irfanoglu, K. Mosaliganti, R. Machiraju, K. Huang, P. Wenzel, A. deBruin, and G. Leone, "Multi-resolution image segmentation using the 2-point correlation functions," in *Proc. IEEE Int. Symp. Biomed. Imag.*, 2007, pp. 56–63.

[9] K. Kanda, S. Mizuta, and T. Matsuda, "Volume visualization using relative distance among voxels," in *Proc. SPIE Medical Imaging 2002: Visualization, Image-Guided Procedures and Display*, 2002, pp. 641–648.

[10] G. Kindlmann, R. Whitaker, T. Tasdizen, and T. Möller, "Curvature-based transfer functions for direct volume rendering," in *Proc. IEEE Conf. Visualization*, 2003, pp. 513–520.

[11] T. Lee, R. Kashyap, and C. Chu, "Building skeleton models via 3-d medial surface/axis thinning algorithms," *Graph. Models Image Process.*, vol. 56, no. 6, pp. 462–478, 1994.

[12] C. Ma and M. Sonka, "A fully parallel 3d thinning algorithm and its applications," *Comput. Vis. Image Understand.*, vol. 64, no. 3, pp. 420–433, 1996.

[13] K. Mosaliganti, F. Janoos, X. Xu, R. Machiraju, K. Huang, and S. Wong, "Temporal matching of dentritic spines in confocal microscopy images of neuronal tissue sections," in *Proc. Workshop Microscopic Image Analysis with Applications in Biology*, 2006, pp. 106–113.

[14] K. Mosaliganti, T. Pan, R. Sharp, R. Ridgway, A. Gulacy, S. lyengar, P. Wenzel, A. deBruin, R. Machiraju, K. Huang, G. Leone, and J. Saltz, "Registration and 3D visualization of large microscopy images," in *SPIE Conf. Medical Imaging*, 2006, pp. 923–934.

[15] T. Pan and K. Huang, "Virtual mouse placenta: Tissue layer segmentation," in *Proc. 27th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, 2005, pp. 3112–3116.

[16] R. Ridgway, O. Irfanoglu, R. Machiraju, and K. Huang, "Image segmentation with tensor-based classification of N-point correlation functions," in *Proc. MICCAI Workshop Medical Image Analysis with Applications in Biology*, 2006, pp. 106–113.

[17] J. Sethian, *Level Set Methods and Fast Marching Methods*. Cambridge, U.K.: Cambridge University Press, 1999.

[18] R. Sharp, R. Ridgway, K. Mosaliganti, P. Wenzel, T. Pan, A. Bruin, R. Machiraju, K. Huang, G. Leone, and J. Saltz, "Volume rendering phenotype differences in mouse placenta microscopy data," *Comput. Sci. Eng. (Special Issue on Anatomic Rendering and Visualization)*, vol. 9, no. 1, pp. 38–47, 2007.

[19] S. Torquato, *Random Heterogenous Material*. New York: Springer-Verlag, 2004.

[20] C. Weaver, P. Hof, S. Wearne, and W. Lindquist, "Automated algorithms for multiscale morphometry of neuronal dendrites," *Neural Comput.*, vol. 16, pp. 1353–1383, 2004.

[21] P. Wenzel, L. Wu, A. deBruin, W. Chen, G. Dureska, E. Sites, T. Pan, A. Sharma, K. Huang, R. Ridgway, K. Mosaliganti, R. Sharp, R. Machiraju, J. Saltz, H. Yamamoto, J. Cross, M. Robinson, and G. Leone, "Rb is critical in a mammalian tissue stem cell population," *Genes Develop.*, vol. 21, no. 1, pp. 85–97, 2007.

[22] L. Wu, A. de Bruin, H. I. Saavedra, M. Starovic, A. Trimboli, Y. Yang, J. Opavska, P. Wilson, J. Thompson, M. Ostrowski, T. Rosol, L. Woollett, M. Weinstein, J. Cross, M. Robinson, and G. Leone, "Extra-embryonic function of Rb is essential for embryonic development and viability," *Nature*, vol. 421, pp. 942–947, 2003.

[23] J. Zhou, A. Doring, and K. D. Tönnnies, "Distance transfer function based rendering," Inst. Simulation and Graphics, Univ. Magdeburg, Magdeburg, Germany, Tech. Rep., 2004.