



A Variational Approach to Vesicle Membrane Reconstruction from Fluorescence Imaging

Kalin Kolev¹, Norbert Kirchgeßner^{2,5}, Sebastian Houben², Agnes Csiszár²,
Wolfgang Rubner², Christoph Palm⁴, Björn Eiben³, Rudolf Merkel², Daniel Cremers¹

Abstract

Biological applications like vesicle membrane analysis involve the precise segmentation of 3D structures in noisy volumetric data, obtained by techniques like magnetic resonance imaging (MRI) or laser scanning microscopy (LSM). Dealing with such data is a challenging task and requires robust and accurate segmentation methods. In this article, we propose a novel energy model for 3D segmentation fusing various cues like regional intensity subdivision, edge alignment and orientation information. The uniqueness of the approach consists in the definition of a new anisotropic regularizer, which accounts for the unbalanced slicing of the measured volume data, and the generalization of an efficient numerical scheme for solving the arising minimization problem, based on linearization and fixed-point iteration. We show how the proposed energy model can be optimized globally by making use of recent continuous convex relaxation techniques. The accuracy and robustness of the presented approach are demonstrated by evaluating it on multiple real data sets and comparing it to alternative segmentation methods based on level sets. Although the proposed model is designed with focus on the particular application at hand, it is general enough to be applied to a variety of different segmentation tasks.

Keywords:

3D segmentation, convex optimization, vesicle membrane analysis, fluorescence imaging.

1. Introduction

1.1. Model system for cell membrane deformation processes

Cellular membranes permanently sense their environment. They generate, stabilize or destabilize local regions of membrane curvature depending on changes of environmental conditions like temperature, ionic strength, and osmotic or mechanical stress [1]. To cope with environmental osmotic stress, prokaryotes have a wide repertoire of mechanisms like ion pumping or accumulation of mechanosensitive proteins on the cell surface. For example, s-layers formed by two dimensional crystals of certain proteins on the surface of the membrane serve as structural osmoprotectant proteins [2]. They increase the membrane rigidity under hyperosmotic conditions. Under these conditions the

¹Department of Computer Science, TU München, Boltzmannstraße 3, 85748 Garching, Germany

²Institute of Bio- and Nanosystems, IBN-4, Biomechanics, Forschungszentrum Jülich GmbH, 52425 Jülich, Germany

³Institute of Neuroscience and Medicine (INM-1), Forschungszentrum Jülich GmbH, Germany

⁴Department of Computer Science, Regensburg University of Applied Sciences, Regensburg, Germany

⁵present address: Institute of Agricultural Sciences, ETH Zurich, LFW A4, Universitätsstrasse 2, CH-8092 Zurich

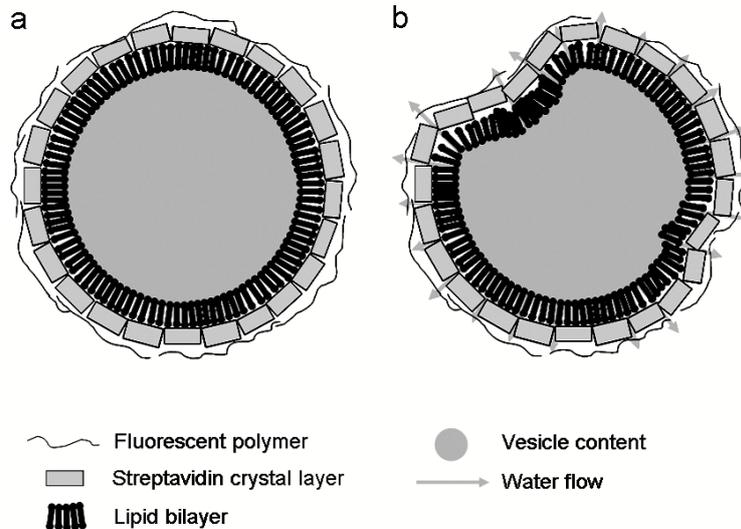


Figure 1: a) Streptavidin coated giant unilamellar vesicle under isoosmotic conditions and b) under hyperosmotic conditions. In the latter case, the solute concentration outside of the vesicle was higher than in the vesicle lumen. To reduce this concentration difference, water flowed out of the vesicle, i. e. its volume decreased while its surface remained constant. Due to this excess surface after osmotic shrinkage, the vesicle underwent shape deformation and folding processes.

rigid crystal layer resists higher mechanical stress and the cell membrane undergoes reversible weak surface wrinkling inhibiting an irreversible membrane budding or fission. However, the importance of this mechanical protection against osmotic stress compared to molecular signaling processes is not yet elucidated. The pressure range which can be compensated just by mechanical osmoprotection is also unknown.

Since the separation of the molecular signalling and the mechanical properties is difficult to realize in living cells, it is more promising to investigate model membrane systems. To mimic a rigid cell membrane under hyperosmotic pressures, *Ratanabanangkoon et al.* [3] and recently *Csiszár et al.* [4] used a model system consisting of phospholipid vesicles and a two-dimensional crystalline protein layer of streptavidin on the vesicle surfaces. Like cell membranes, the phospholipid molecules built spherically closed bilayers, while streptavidin simulated the role of surface proteins like s-layers [2], spam [5] or COPI [6] (see Fig. 1). Under hyperosmotic conditions, the solute concentration outside of the vesicle membrane was higher than in the vesicle lumen. To reduce this concentration difference, water streamed out of the vesicle and its volume decreased while its surface remained constant. Due to this excess of surface after osmotic shrinkage, the vesicle underwent shape deformation and folding processes. The amount of deformation depended on the material parameters of the vesicle shell. From the physical point of view, condensation of bending and stretching energy in the folds of a crumpled surface structure depending on material parameters has been of high interest [7]. Such deformation energies can be calculated from the 3-dimensional shape of the coated vesicle when exposed to osmotic pressure [8].

Among the major computational challenges regarding the understanding of these vesicle membrane deformation processes is therefore the precise and robust estimation of the vesicle membrane over time. To this end, one can rely on the fluorescence properties of materials as observed in volumetric image data, in our case collected by confocal laser scanning microscopy [9]. Since typical measurements are quite noisy (see Fig. 2) the reliable segmentation is a major challenge.

1.2. Variational Estimation of the Membrane Shape

Inspired by recent advances in shape optimization [10, 11], we propose a novel variational approach for 3D segmentation which is aimed at exploiting the properties of fluorescence images obtained by confocal laser scanning microscopy of vesicle membranes. The model combines regional intensity subdivision, edge alignment and orientation information and can be *globally* minimized by means of convex relaxation techniques. Starting with a specified user input in the form of roughly marked interior and exterior regions in one of the volume slices, corresponding intensity

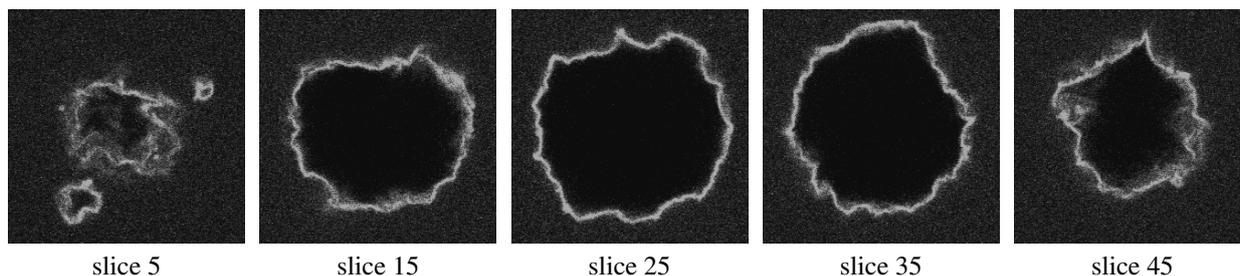


Figure 2: Several slices of a typical volumetric data set, obtained by microscopic imaging of a fluorescently labeled vesicle membrane. While interior and exterior regions exhibit similar intensity characteristics, the membrane surface is identified by bright areas due to its fluorescence properties.

histograms are estimated and used to derive regional statistics. Secondly, the model encourages the extracted surface to pass through locations of high intensity with correspondingly high probability of the presence of the fluorescent vesicle membrane. And, thirdly, the proposed anisotropic formulation allows to impose local surface orientations to align with the local image gradient. Due to the variety of considered image features and the robust regularization scheme, the proposed approach enjoys remarkable resilience to noise and deviations from the primary assumptions which are frequently encountered in real environments. The contributions of the current work can be summarized as follows:

- We unify various established image cues like regional statistics, edge attraction and orientation information into a continuous convex energy model.
- We propose an anisotropic generalization of the classical isotropic minimal surface model, which accounts for the sparse slicing of the measured volume data in the z -direction, e. g. inherent in confocal microscopy. We show that the generalized model retains convexity and thus all globality guarantees of its isotropic counterpart.
- We adopt an efficient numerical scheme for solving the arising optimization problem, based on linearization and fixed-point iteration.

Although the proposed approach is designed with a focus on the particular application scenario of vesicle membrane estimation in fluorescence imaging, the proposed concepts are easily adapted to other segmentation tasks. Finally, the curvature of the extracted surface is computed in order to facilitate further analysis of the membrane deformations.

In the next section, we present and discuss related work. In Section 3 we introduce a variational model integrating regional intensity statistics, edge attraction and orientation fitting. Section 4 is devoted to the optimization of the proposed model. Section 5 briefly sketches an algorithm for estimating the curvature of the segmented surface, used for further analysis. In Section 6, we show experimental results on challenging real data sets, demonstrating the effect of each individual component of the energy on the quality of computed solutions and showing a clear superiority over level set based solutions. We conclude with a brief summary.

2. Related Work

Image segmentation is one of the fundamental problems in Computer Vision and has been extensively studied for decades. Among the myriad of existing techniques variational methods prevail due to their accuracy, robustness and mathematical elegance. Since the pioneering work of *Kass et al.* [12] considerable efforts have been made on exploring different energy models relying on various image cues. Three paradigms have proven to be particularly useful: regional statistics, edge terms and orientation information. In [13] the segmentation problem is posed in terms of a piecewise smooth or piecewise constant image approximation with minimal boundary length. This methodology is generalized in [14] and [15], where regional intensity of color statistics are estimated to achieve an optimal region assignment. The segmentation is obtained by alternately minimizing the underlying energy functional with respect to the contour and the parameters of the probability models. In [16] and [17] the region parameters are estimated interactively, which allows for globally optimal segmentations according to the specified user input. Although the

reliance on regional image criteria entails considerable robustness, it often results in imprecise or oversmoothed segmentations. In order to account for this drawback and increase the accuracy of the segmentation, researchers suggested *edge terms* attracting the curve (in 2D) or the surface (in 3D) towards the locations of sudden brightness changes that commonly arise at object boundaries. *Geodesic active contours* [18, 19] pose the segmentation problem as a weighted minimal surface model, where the weights reflect a local Riemannian metric based on the magnitude of the image gradient and therefore account for the presence of edges. While the incorporation of edge terms can result in substantial improvements of object boundary alignment, it is still unable to handle thin or small-scale geometric details. To this end, *Vasilevskiy and Siddiqi* [20] demonstrated the potential of using *orientation information* to segment narrow elongated structures like blood vessels. The curve (or surface) orientation is provided in the form of a dense vector field defined on the image domain. A typical choice for the normal field is the normalized or unnormalized image gradient. In this work, we propose a unified framework integrating regions, edges and orientation with the goal of combining the advantages of all utilized paradigms to achieve robust and accurate 3D segmentations regardless of the available geometric structure. Although the presented model is built upon established concepts, it differs from previous methods in their concrete realization. In particular, it relies on a novel anisotropic generalization of the classical weighted minimal surface model and a new edge attraction term.

It should be noted that there are some efforts to combine local and global shape retrieval criteria, e. g. [21] in the context of image segmentation and [22] in the context of video interpolation. In contrast, the proposed model completely relies on global image features. While the integration of local characteristics is straightforward, we didn't find it to be necessary for the particular application under consideration.

Along with the underlying energy model, a crucial issue regarding the design of a segmentation approach is the choice of contour/surface representation and the choice of optimization method. A simple and straightforward way to represent the subdividing interface is as a set of line segments (in 2D) or triangles (in 3D), where the derived evolution equations are applied to its vertices [12]. A major limitation of such *explicit* schemes is that numerical instabilities can occur during the evolution process, especially in case of topological changes like merging or splitting. In order to address this weakness, level set methods have been proposed [23]. The key idea is to switch to an implicit representation of the contour/surface by assigning each point within the domain a value that encodes its membership to the interior or exterior region. Hence, the evolution of the subdividing interface is replaced by the evolution of an implicit function. Various energy models have been adopted to the level set framework including regional statistics [24], edge terms [18] and orientation integration via flux maximization [20]. Unfortunately, level sets are susceptible to local minima, which could result in unexpected and undesired segmentations. Since respective energies are generally non-convex in the space of level set functions, solutions rely on an appropriate initialization and do not come with any optimality guarantee. In contrast, combinatorial methods like graph cuts [25, 26, 27, 28] do guarantee global optimality of the reconstruction. The key idea is to pose the segmentation problem as a mincut/maxflow problem of a graph reflecting the specific image content. Subsequently, the globally optimal solution (corresponding to the minimal cut) can be computed in polynomial time by determining the maximal flow in the corresponding graph. Unfortunately, the global minimum is obtained in a *spatially discrete* setting. As a consequence, respective solutions typically exhibit metrication errors. In particular, it is not a trivial task to represent arbitrary surface orientations on a fixed grid and for that reason relying on a discrete framework often leads to the introduction of an orientation bias in the shape retrieval process. Recent investigations led to the development of convex optimization techniques [10, 29], which combine the advantages of spatially continuous formulations like the level set methods with the guaranteed optimality of spatially discrete methods like graph cuts. A discussion of these differences with an experimental evaluation can be found in [30]. Yet, more research is needed to precisely specify the range of energy models amenable to continuous convex optimization. The current work exhibits a step in this direction. In particular, we show that established concepts like regional probability estimation, edge attraction and orientation integration can be unified in a single cost functional that can be optimized globally in a spatially continuous setting by means of convex relaxation techniques.

3. Anisotropic Energy Model Integrating Regions, Edges and Orientation

This paragraph introduces the underlying energy functional accounting for various complementary information sources encoded in the data. In particular, the proposed model imposes regional intensity subdivision, edge alignment and orientation integration.

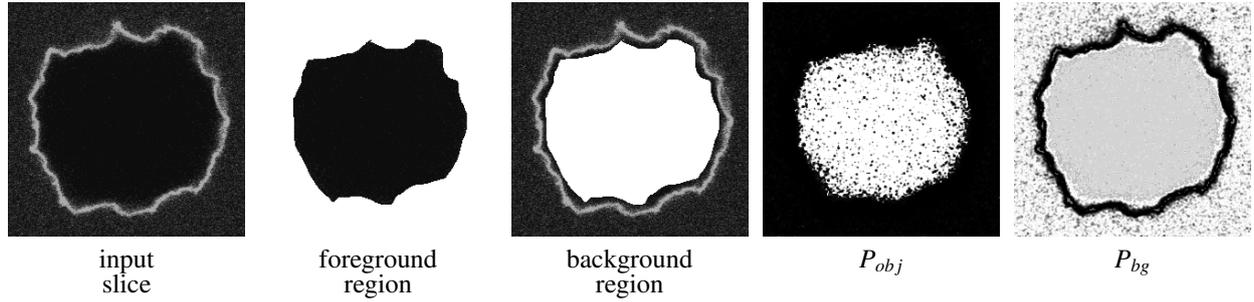


Figure 3: Regional subdivision via interactive user input. The user is required to roughly segment one of the slices of the volume data. Based on the provided segmentation, foreground/background histograms are built, which are in turn used to assign intensity likelihoods to each voxel conditioned on being inside or outside the membrane surface. An example is illustrated for one of the slices in Fig. 2. Note that due to the noisy input data and the similarity in the intensity characteristics of both regions, these likelihoods alone do not allow a reliable separation of object and background.

We start with some notations. Let $I : V \subset \mathbb{R}^3 \rightarrow [0, 1]$ denote the measured volumetric data (see Fig. 2 for an example). The goal is to partition the volume into two subsets by identifying the vesicle membrane surface $S \subset V$. We will indicate the interior region according to a certain surface estimate S as R_{obj}^S and the exterior one as R_{bg}^S , respectively.

As previously mentioned, the user is required to specify a rough identification of both regions by marking them in one of the input volume slices (Fig. 3). This user input is used to build initial foreground/background histograms which are adaptively refined during the evolution process. To this end, precise segmentation is not needed. Alternatively, strokes or scribbles specifying both regions could be provided. After smoothing and normalization derived histograms are used to define probabilities $P_{obj}, P_{bg} : V \rightarrow [0, 1]$ for observing a certain intensity, given that the respective voxel is inside or outside the membrane surface. However, due to the noisy input data and the similarity in the intensity characteristics of both regions these intensity likelihoods do not permit a reliable separation of object and background. Based on the probabilistic volume subdivision a surface estimate is desired which separates voxels with high values for P_{obj} and low values for P_{bg} from voxels with opposite property. Simultaneously, favorable shape characteristics like smoothness should be imposed. These observations naturally lead to the following energy model

$$E(S) = -\mu \left(\int_{R_{obj}^S} \log P_{obj}(x) dx + \int_{R_{bg}^S} \log P_{bg}(x) dx \right) + \int_S ds, \quad (1)$$

where $\mu \in \mathbb{R}_{\geq 0}$ is a parameter controlling the weighting of data term and regularization and ds denotes an Euclidean area element. Here, for the sake of simplicity we assume a given parametrization of the surface $S : \Omega \rightarrow \mathbb{R}^3$. In the next paragraph we will get rid of this assumption and switch over to an implicit representation. For a good surface estimate P_{obj} is expected to be close to 1 within the interior region R_{obj}^S and P_{bg} close to 1 within the exterior region R_{bg}^S . Hence, the overall energy costs will be low. Analogously, for inaccurate estimates the corresponding probability values will generally vary, which after applying the logarithm would result in high data fidelity costs. In effect, minimization of the functional in (1) specifies a shape with the desired properties. Recall that both foreground/background histograms are updated permanently by reinitializing them with the voxel values within the current regions R_{obj}^S and R_{bg}^S . This adaptive procedure allows to adjust the histograms to the input data rather than relying on the specified user input.

Unfortunately, regional intensity statistics is not sufficient to achieve the required precision of the extracted surface due to the noise in the estimated probability maps (see Fig. 3). Yet, a closer inspection of the volume data reveals that the boundary of the interior region can be identified as areas, where the intensity changes from dark to bright. This requirement can be imposed by encouraging the shape normal at each point to align with the local intensity gradient. Taking this into account we come up with the following energy model

$$E(S) = -\mu_1 \left(\int_{R_{obj}^S} \log P_{obj}(x) dx + \int_{R_{bg}^S} \log P_{bg}(x) dx \right) - \mu_2 \int_S \langle N_S(s), \nabla I(s) \rangle ds + \int_S ds, \quad (2)$$

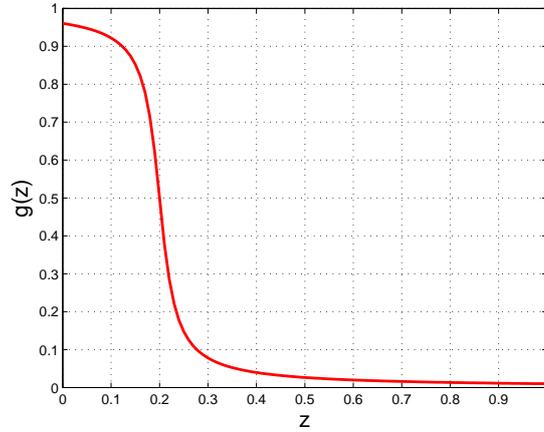


Figure 4: Edge attraction function. Plotted is the proposed function $g(z) = \frac{1}{2} - \frac{\text{atan}(\frac{z-\beta_1}{\beta_2})}{\pi}$ with parameters $\beta_1 = 0.2$ and $\beta_2 = 0.025$.

where N_S denotes the (normalized) outward surface normal at s . Again, $\mu_1, \mu_2 \in \mathbb{R}_{\geq 0}$ are constant weighting parameters. The second term in (2), referred in the literature as flux [20], favors the alignment of the surface orientation with the vector field defined by ∇I . In the ideal case when shape normals exactly coincide with the intensity gradient, the corresponding costs will be negative and maximal in magnitude. On the other hand, surface normals pointing in the opposite direction will cause positive energy costs scaled by the respective deviation. It is important to note that abrupt intensity changes in the input data, i. e. areas of high gradient, can be observed not only at the boundary between vesicle membrane and interior region but also at the boundary between membrane and exterior region. However, the estimated region statistics are additionally used to remove this ambiguity. In fact, regional information and flux are exploited as complementary cues. The balancing between them via the weighting parameters μ_1 and μ_2 should be established based on the confidence in both information sources in the particular situation.

The energy model in (2) allows further extensions by exploiting the fact that the vesicle membrane can be identified in the input volume data by locations of relatively high intensity due to its fluorescence properties. A simple but very effective way to take this observation into account is to replace the Euclidean metric used by the regularization term in (2) by a more general Riemannian measure based on the local data content. This yields the following functional

$$E(S) = -\mu_1 \left(\int_{R_{obj}^S} \log P_{obj}(x) dx + \int_{R_{bg}^S} \log P_{bg}(x) dx \right) - \mu_2 \int_S \langle N_S(s), \nabla I(s) \rangle ds + \int_S g(I(s)) ds, \quad (3)$$

where $g : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ is a non-negative, monotonically decreasing function. The new weighted minimal surface term alleviates the smoothing and encourages the shape to pass through locations of low costs according to g . To this end, the function should be monotonically decreasing in order to assign low costs to bright locations which vote for the presence of membrane edges and to penalize dark regions specifying interior and exterior. Yet, this is not the only requirement that should be followed when designing the function g . When conducting different measurements, the particular intensity distribution of the volume data, including the brightness characteristics of the vesicle membrane, may vary. Hence, from a practical point of view it is preferable to provide the user with the flexibility to specify the particular complexion of the weighting function g . We propose the following definition

$$g(z) := \frac{1}{2} - \frac{\text{atan}(\frac{z-\beta_1}{\beta_2})}{\pi}, \quad (4)$$

where $\beta_1 \in [0, 1], \beta_2 \in (0, 1]$ are constant parameters. The parameter β_1 determines the location of the inflection point of the function, while β_2 controls the smoothness of the falling edge (see Fig. 4). In effect, the value of β_1 indicates the lower limit of the range in which a pixel is categorized as “bright”. Large values will restrict the range and classify only high intensities as vesicle membrane. The parameter β_2 can be interpreted as a kind of tolerance of the estimated

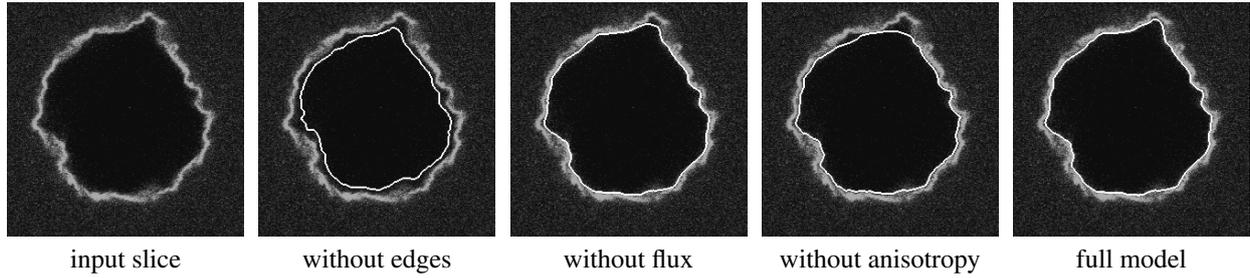


Figure 5: Contribution of different data features. Results of various energy models obtained by combining different cues applied to the 3D data set in Fig. 2. The segmentation is visualized for slice #38 and superimposed on the input data. Note that only the complete model (5) – shown on the right – yields tight and accurate segmentation results.

intensity splitting. Large values will give a smooth function and will smear the boundary, while small values will sharpen the splitting.

Although the energy functional in (3) incorporates a broad variety of data features, it still does not address one relevant point. The applied scanning technique produces distorted volume data in a sense that the spacing in different directions is not equal. In fact, the resolution and, as a consequence the sampling in the z -direction is significantly coarser than along the other two spatial dimensions. As a consequence, the isotropic weighted minimal surface term $\int_S g(I(s)) ds$, which imposes the shape regularity, treats the slicing direction with the same priority as the other two, which leads to a different amount of smoothing due to the unequal spacing. Hence, different scaling is needed along this particular direction, which is achieved by the following anisotropic generalization

$$E(S) = -\mu_1 \left(\int_{R_{obj}^s} \log P_{obj}(x) dx + \int_{R_{bg}^s} \log P_{bg}(x) dx \right) - \mu_2 \int_S \langle N_S(s), \nabla I(s) \rangle ds + \int_S g(I(s)) \sqrt{N_S^T(s) D(s) N_S(s)} ds \quad (5)$$

where $\{D(x) \in \mathbb{R}^{3 \times 3} \mid x \in V\}$ denotes a family of positive semidefinite symmetric tensors tolerating certain directional selectivity. In addition to encouraging the surface to pass through preselected locations according to the edge map g , the anisotropic energy model favors certain shape orientations while suppressing others. This can be achieved with the appropriate design of the tensors $D(x)$. Note that D is defined pointwise. However, in the remainder of this paragraph we will omit the argument for the sake of simplicity. The term $\int_S g(I(s)) \sqrt{N_S^T(s) D(s) N_S(s)} ds$ can still be interpreted as a minimal surface formulation according to the metric induced by D (i. e. $\|v\|_D = \sqrt{v^T D v}$). The isotropic metric used in (3) appears as a special case with $D = I$, where $I \in \mathbb{R}^{3 \times 3}$ denotes the identity matrix. Our intention is to define D in a way that allows for different scaling in the z -direction. This can be achieved by setting

$$D := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \gamma \end{pmatrix}, \quad (6)$$

where $\gamma \in [0, 1]$ is a constant parameter. Ideally, γ should be set to the squared ratio between the sampling in the z -direction and the x/y -direction, respectively. Note again that for $\gamma = 1$ we obtain the original isotropic model.

Now, the final model defined in (5) contains all constructs necessary to achieve our primary goal. The relevance of the various features used in the model depends on the particular input data. In cases, where one of the cues prevails, the corresponding energy term will play a dominant role, which should be reflected by an adequate parameter setting. However, in a typical practical situation all utilized paradigms will have a certain contribution to the final segmentation result. An example demonstrating the contribution of the different data features is illustrated in Fig. 5. Superimposed segmentation results for one of the slices with different models, obtained by omitting one of the cues, are visualized. The regional terms alone, which are based on the specified user input, are not sufficient to produce accurate segmentations, even in combination with a flux term. This is due to the blurred membrane measurements

resulting in high intensity characteristics in the vicinity of the surface (see Fig. 3). Note that a more accurate user input will not drastically improve the segmentations, since more bright pixels will enter the foreground histogram and diminish the discriminative power of the regional terms. Substantial increase of the accuracy of segmentation can be observed by adding the edge attraction term to the model. Yet, the absence of a flux term is still noticeable, especially at small-scale contour structures. Finally, the effect of the anisotropic generalization is investigated by comparing it to the original isotropic model (i. e. for $\gamma = 1$). As expected, the isotropic regularizer produces significant oversmoothing at areas of abrupt slice-to-slice changes. This can be suppressed by reducing the overall regularization but only at the cost of degrading the contour alignment along the slice-plane. In summary, regional terms and edge term seem to be decisive for the overall quality of the segmentation results but the flux term and the anisotropic scaling also contribute significantly to the precision of the estimated boundary.

4. Variational Minimization via Partial Differential Equations

This paragraph deals with the minimization of the energy functional proposed in (5). Here, we focus on continuous optimization methods due to their numerical accuracy, memory efficiency and parallelizability. We refer to [30] for a detailed comparison to discrete approaches.

First, we observe that the flux term in (5) can be rewritten as a regional term over the surface interior via the *divergence theorem*. Now, we end up with the following energy functional, equivalent to (5)

$$\begin{aligned} E(S) = & - \int_{R_{obj}^S} (\mu_1 \log P_{obj}(x) + \mu_2 \Delta I(x)) dx - \mu_1 \int_{R_{bg}^S} \log P_{bg}(x) dx \\ & + \int_S g(I(s)) \sqrt{N_S^T(s) D(s) N_S(s)} ds. \end{aligned} \quad (7)$$

Note that $\text{div}(\nabla I) = \Delta I$. Basically, the above formulation states that a flux term can be incorporated by adequately adopting the interior region statistics. This observation will considerably facilitate further optimization.

4.1. Solution I: Minimization by Level Sets

The level set method [31, 23] is an established technique for minimizing energy functionals for shape estimation like (7). The key idea is to represent the surface implicitly by a function $\phi : V \rightarrow \mathbb{R}$, whose values reflect the signed distance from it, where negative values indicate the interior region and non-negative values - the exterior, respectively. The formulation proposed in this section is based on the framework of [24]. Note, however, that the model in (7) surpasses the one in [24] in its completeness. In particular, the anisotropic regularization scheme poses the main challenge when developing a valid generalization.

The construction of an implicit surface representation implies that both interior/exterior regions can easily be accessed by means of the Heaviside function

$$H(z) = \begin{cases} 1, & \text{if } z \geq 0 \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

In the level set framework a surface deformation equation is converted to an evolution equation with respect to the implicit function ϕ . Hence, topological changes are handled automatically without additional efforts. By means of the Heaviside function, we can rewrite the functional in (7) as a functional over the set of implicit functions, which yields

$$\begin{aligned} E(\phi) = & - \int_V (\mu_1 \log P_{obj}(x) + \mu_2 \Delta I(x)) (1 - (H \circ \phi)(x)) dx - \mu_1 \int_V \log P_{bg}(x) (H \circ \phi)(x) dx \\ & + \int_V g(I(x)) \sqrt{\nabla(H \circ \phi)(x)^T D(x) \nabla(H \circ \phi)(x)} dx. \end{aligned} \quad (9)$$

The terms $1 - (H \circ \phi)(x)$ and $(H \circ \phi)(x)$ in the above formulation act as filtering functions by specifying interior and exterior voxels, respectively. Note that the composite function $H \circ \phi$ is a binary labeling function identifying both regions. More precisely $H \circ \phi \equiv \mathbf{1}_{R_{bg}^S}$, where $\mathbf{1}_{R_{bg}^S}$ denotes the indicator function of the exterior region. Note also the relations $N_S(x) = \frac{\nabla(H \circ \phi)(x)}{\|\nabla(H \circ \phi)(x)\|}$ and $ds = \|\nabla(H \circ \phi)(x)\| dx$ which give rise to the last term in (9).

The first step towards minimizing the functional in (9) is to compute its Euler-Lagrange equations. We obtain (the argument x is omitted for simplicity)

$$\begin{aligned}
\frac{\partial E}{\partial \phi} &= \delta(\phi) (\mu_1 \log P_{obj} + \mu_2 \Delta I - \mu_1 \log P_{bg}) + \delta'(\phi) g(I) \sqrt{\nabla \phi^T D \nabla \phi} \\
&- \operatorname{div} \left(\delta(\phi) g(I) \frac{D \nabla \phi}{\sqrt{\nabla \phi^T D \nabla \phi}} \right) \\
&= \delta(\phi) (\mu_1 \log P_{obj} + \mu_2 \Delta I - \mu_1 \log P_{bg}) + \delta'(\phi) g(I) \sqrt{\nabla \phi^T D \nabla \phi} \\
&- \delta'(\phi) g(I) \frac{\nabla \phi^T D \nabla \phi}{\sqrt{\nabla \phi^T D \nabla \phi}} - \delta(\phi) \operatorname{div} \left(g(I) \frac{D \nabla \phi}{\sqrt{\nabla \phi^T D \nabla \phi}} \right) \\
&= \delta(\phi) (\mu_1 \log P_{obj} + \mu_2 \Delta I - \mu_1 \log P_{bg}) - \delta(\phi) \operatorname{div} \left(g(I) \frac{D \nabla \phi}{\sqrt{\nabla \phi^T D \nabla \phi}} \right).
\end{aligned} \tag{10}$$

Thereby, $\delta(\cdot)$, which denotes the Dirac delta function, appears as the derivative of the Heaviside function, i. e. $\delta(z) = H'(z)$. It is not surprising that all terms in the above expression are scaled by $\delta(\phi)$, since only the 0-level of ϕ representing the surface is of interest. In practice, $\delta(\phi)$ is replaced by a smoothed version $\delta_\epsilon(\phi)$ for numerical reasons.

The functional in (9) can be minimized via gradient descent by solving the following partial differential equation (PDE)

$$\phi_t = -\frac{\partial E}{\partial \phi}. \tag{11}$$

The above equation describes an evolution process with respect to an artificial time variable t . Updating the implicit function ϕ accordingly gradually decreases the energy and leads to a (local) minimum of the functional, identified as a steady state of (11). Numerically, the above equation can be solved by applying classical discretization schemes like the forward Euler method. It should be noted that after each update, ϕ will no longer reflect the exact distance to the surface. For that reason, the function should permanently be reinitialized in order to circumvent numerical difficulties.

Although the level set method leads to a tractable numerical scheme and avoids the need for explicit handling of topological changes, it suffers from one important limitation. Due to the binary nature of the representation, whereas the localization of the surface is implicitly encoded by the sign of a distance map ϕ , this technique always involves optimization over a set of binary labeling configurations. As a consequence, due to the minimization over a non-convex domain any iterative procedure can get stuck in an undesired local minimum that may be far from the expected solution. In effect, the final result strongly depends on the initialization. This issue will be discussed in more detail in the next section.

4.2. Solution II: Minimization by Convex Relaxation

In [10] it has been proven that a simple generalization of the level set method based on convex relaxation allows for global optimization in certain cases. We build upon this work and show that the class of energy functionals amenable to convex relaxation techniques can be extended to include our primary model given in (5).

First, we observe that the implicit function ϕ in (9) always appears in concatenation with the binary Heaviside function H . Hence, an equivalent formulation can be obtained by replacing $u := H \circ \phi$. This leads to the following minimization problem

$$\begin{aligned}
E(u) &= - \int_V (\mu_1 \log P_{obj}(x) + \mu_2 \Delta I(x)) (1 - u(x)) dx - \mu_1 \int_V \log P_{bg}(x) u(x) dx \\
&+ \int_V g(I(x)) \sqrt{\nabla u(x)^T D(x) \nabla u(x)} dx \rightarrow \min, \\
\text{s. t. } &u : V \rightarrow \{0, 1\}.
\end{aligned} \tag{12}$$

Now, one can notice that the above functional is convex. The convexity of the first two terms follows trivially from their linearity. The last anisotropic regularization term is also convex, since it represents the weighted total variation

of u in the norm induced by g and D . However, the minimization is carried out over the set of all binary labeling functions which is not convex. As a result, we have a constrained non-convex optimization problem and obtaining an exact solution is not straightforward. A simple and effective remedy to this problem is to relax the domain of feasible functions and to allow also intermediate values in the interval $[0, 1]$. Thus, we obtain a more tractable constrained convex optimization problem

$$\begin{aligned} E(u) &= - \int_V (\mu_1 \log P_{obj}(x) + \mu_2 \Delta I(x)) (1 - u(x)) dx - \mu_1 \int_V \log P_{bg}(x) u(x) dx \\ &+ \int_V g(I(x)) \sqrt{\nabla u(x)^T D(x) \nabla u(x)} dx \rightarrow \min, \\ \text{s. t. } &u : V \rightarrow [0, 1]. \end{aligned} \quad (13)$$

Now, we have a convex functional, optimized over a convex domain. Hence, a global minimum can be obtained by any iterative local optimization procedure like gradient descent. However, recall that we are interested in solving the original “binary” problem (12). Similar to [10], we can infer that thresholding the solution of (13) will give us a global minimum of (12) due to favorable properties of the particular variational model at hand. In [10] this has been shown only for a certain class of functionals including regional subdivision terms and weighted isotropic regularization. In a recent work [32] it has been proven that the thresholding property is still fulfilled when the minimal surface formulation is generalized to anisotropic Finsler metrics. Building upon these works we state the following theorem:

Theorem 1. *If $u^* : V \rightarrow [0, 1]$ is a minimizer of the functional in (13), then for almost any threshold $\nu \in (0, 1)$ the binary function $\mathbf{1}_{\Sigma_{\nu, u^*}}(x) : V \rightarrow \{0, 1\}$ with $\Sigma_{\nu, u^*} := \{x \mid u(x) > \nu\}$ is also a (global) minimizer of (13).*

Proof. The proof is given in the Appendix. □

Any thresholded (global) minimizer of (13) is binary and fulfills the constraints in (12). Trivially, it is also a global solution of the *non-convex* optimization problem in (12), since the set of binary labeling functions is a subset of the domain of relaxed functions.

In summary, we obtain the following approach for globally optimizing (12):

1. Find a minimizer u_{rel} of (13).
2. Threshold the result $u_{bin}(x) = \begin{cases} 1, & \text{if } u_{rel}(x) > \nu \\ 0, & \text{otherwise} \end{cases}$ to obtain a minimizer of (12).

In our experiments we chose $\nu = 0.5$ but we obtained virtually the same results with $\nu \in [0.1, 0.9]$.

We have seen that solving the non-convex problem (12) globally reduces to solving a more tractable convex problem (13). As mentioned previously, due to the convexity the applied optimization procedure will not affect the final result but only the speed of convergence. Here, we propose a generalization of a fixed-point iteration scheme based on Successive Overrelaxation (SOR). Similar techniques have been successfully applied for solving related problems like image denoising [33] and multiview 3D reconstruction [11]. The first step is to set up the Euler-Lagrange equation which states a necessary condition for a minimum of (13) (the argument x is omitted for simplicity):

$$0 = \mu_1 (\log P_{obj} - \log P_{bg}) + \mu_2 \Delta I - \operatorname{div} \left(g(I) \frac{D \nabla u}{\sqrt{\nabla u^T D \nabla u}} \right). \quad (14)$$

The key idea is to solve the above equation directly instead of deriving an evolution PDE like (11). Yet, (14) exhibits a non-linear differential equation which is not trivial to be solved directly. The source of non-linearity is given by the diffusivity $\rho := \frac{g(I)}{\sqrt{\nabla u^T D \nabla u}}$. Thus, a straightforward way to address this difficulty is to apply a linearization scheme. Starting with an initialization $u^0 = 0.5$, we can compute ρ and keep it constant. For constant ρ , (14) is linear and spatial discretization yields a sparse linear system of equations which can be solved efficiently with iterative approaches like SOR. This means, we iteratively compute the k -th update of u at voxel i by

$$u_i^{l,k+1} = (1 - \omega) u_i^{l,k} + \omega \frac{\sum_{j \in N(i), j < i} \rho_{i \sim j}^l u_j^{l,k+1} + \sum_{j \in N(i), j > i} \rho_{i \sim j}^l u_j^{l,k} - d_i}{\sum_{j \in N(i)} \rho_{i \sim j}^l}, \quad (15)$$

where $\omega \in (0, 2)$ is a fixed relaxation parameter. Thereby, the index l counts the updates in the diffusivity terms, $\mathcal{N}(i)$ denotes the 6-neighborhood of i and $d_i := \mu_1 (\log P_{obj,i} - \log P_{bg,i}) + \mu_2 \Delta I_i$ summarizes the constant part of (14) that does not depend on u , i. e. the righthand side of the linear system. Finally, $\rho_{i \sim j}^l$ denotes the diffusivity between voxel i and its neighbor j while distinguishing between different spatial directions according to the anisotropic tensor D . It is defined as

$$\rho_{i \sim j}^l = \begin{cases} \gamma \frac{\rho_i^l + \rho_j^l}{2}, & \text{if voxels } i \text{ and } j \text{ are neighbors in the } z\text{-direction} \\ \frac{\rho_i^l + \rho_j^l}{2}, & \text{otherwise.} \end{cases} \quad (16)$$

Thus, diffusivities between neighboring voxels are estimated by averaging the corresponding values, which are in turn computed as

$$\rho_i^l := \frac{g(I)_i}{\sqrt{\nabla u_i^{lT} D \nabla u_i^l + \epsilon^2}}, \quad (17)$$

where $\epsilon := 0.001$ is a small constant that prevents the diffusivity to become infinite when $\nabla u_i^{lT} D \nabla u_i^l = 0$ and ∇u_i^l is approximated by standard central differences. The over-relaxation parameter ω has to be chosen in the interval $(0, 2)$ for the method to converge. The optimal value depends on the linear system to be solved. Empirically, for the system at hand we obtained the fastest convergence rate for $\omega = 1.85$. Recall that for $\omega = 1$ the SOR method is equivalent to the Gauss-Seidel method. After the linear solver yields a sufficiently good approximation (we iterated for $k = 1, \dots, 10$), one can update the diffusivities ρ_i^l and solve the next linear system. Iterations are stopped as soon as the energy decay in one iteration is in the range of number precision.

5. Curvature Estimation

For further processing steps we extracted the surface of the segmented volume data to a Delaunay-like triangulation using the CGAL library [34]. This enabled us to easily estimate the volume and the surface area of the vesicle and additionally to exclude segmented regions that were not connected to the main vesicle.

To estimate the curvature on the vesicle surface we used a method proposed by *Rusinkiewicz* [35] which proved both stable and accurate. This procedure works on triangulations of arbitrary vertex distributions and topology following a finite-differences approach to approximate the second fundamental tensor of every facet, a 2×2 matrix in our case. Recall that the eigenvalues of the second fundamental tensor can be used to derive the mean and Gaussian curvature. To get a vertex-wise curvature measure the fundamental tensors of all aligning facets are averaged by a weight considering the Voronoi area of each vertex, i. e. the area of each facet that lies closest to the vertex.

To receive the normal vectors for the finite-differences approach we used the gradient of the segmented volume achieving stability by both the smooth triangulation and the continuity of the volume gradients.

6. Experiments

This paragraph presents an experimental validation of the proposed approach. Since the segmentation method was developed with focus on a particular application, we give some details on the data acquisition process before demonstrating its performance on generated real-world data sets.

6.1. Experimental Setup

6.1.1. Materials

To mimic mechanical properties of cell membranes, 1,2-dioleoyl-sn-glycero-3-phosphocholine (DOPC) was chosen as main lipid component. The addition of 1,2-dipalmitoyl-sn-glycero-3-phosphoethanolamine-N-(cap Biotinyl) (capBioDPPE) gave rise to a drastically increased rigidity of the coupling of streptavidin to the lipid membrane. Lipid components were dissolved in chloroform and mixed to 1 mg/ml total lipid concentration in a molar ratio of DOPC/capBioDPPE = 10/1. All lipids were purchased from Avanti Polar Lipids Inc. (AL, USA) in a purity of at least 99.9%.

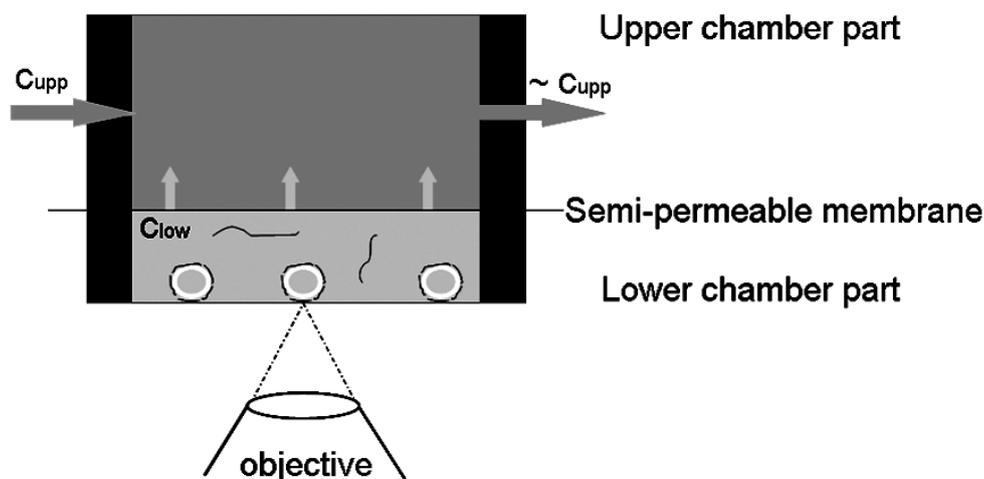


Figure 6: Experimental setup. Dialysis through flow chamber for fluorescence microscopy observation of vesicle folding under continuously increasing hyperosmotic pressures was used. Vesicles were deposited in the lower chamber in glucose solution with a concentration of c_{low} . This chamber part was separated from the upper one by a semi-permeable membrane, which allowed a water flow and simultaneously excluded any other transport between the two chamber parts. The upper chamber had a constantly higher glucose concentration (c_{upp}) than the lower chamber. The concentration difference was reduced by water flow from the lower chamber (see small arrows), including vesicle volumes, into the upper chamber. The changed osmolarity induced vesicle volume loss at constant surface area thereby giving rise to vesicle folding. The process was observed by a laser scanning microscope.

6.1.2. Vesicle preparation

Vesicles were prepared by the electroswelling technique [36]. 10 μl of the lipid mixture were carefully deposited on indium tin oxide (ITO) coated glass slides (Prazisionsglas & Optik GmbH, Iserlohn, Germany). The lipid films were then dried under vacuum for at least 1 h. For the electroswelling step, the plates were placed in a chamber containing 2 ml of 130 mM sucrose solution (130 mosm/L) in pure water and separated by a 1 mm teflon spacer. Vesicles were swollen by applying an altering electric field of 1.5 V and 10 Hz for 2 h. After electroswelling, vesicles were incubated in 70 $\mu\text{g/ml}$ protein solution for at least 2 h at pH 5.5 under isoosmotic conditions.

6.1.3. Experimental environment

Increasing osmotic pressure was generated in a special microscopic chamber designed for experiments using gradual osmotic pressure change (see Fig. 6). It consisted of two parts. The lower chamber part was filled with 2 ml of 130 mM glucose solution (130 mosm/L) and 50 μl of isoosmotic vesicles suspension. To visualize the model membranes in fluorescence microscopy, fluorescently labeled polymer (dextran-TRITC from Sigma Aldrich Inc., St. Louis, MO, USA) was added to the glucose solution in a final concentration of 1 $\mu\text{g/ml}$ while the vesicle lumen remained unlabeled. The fluorescently labeled polymer was uniformly dispersed in the glucose solution emitting low background intensity, which was bleached during the experiments. Streptavidin coated vesicle membranes are slightly negatively charged at pH 5.5. The electrostatic interaction between protein coating and fluorescently labeled polymer resulted in enrichment of polymer molecules on the protein surface increasing the fluorescence intensity on the vesicle membranes. This process occurred until the accumulation of polymers on the vesicle surfaces was saturated. As lipid bilayers were impermeable to polymers, vesicle lumens remained completely without fluorescent and appeared dark in the fluorescent channel, correspondingly.

The upper chamber part was a flow chamber. Isoosmolar glucose solution streamed through it with a flow rate of approximately 20 mL/min. The two parts were separated and simultaneously contacted with each other by a semi-permeable membrane (Reichelt Chemietechnik GmbH Co, Heidelberg, Germany). This membrane enabled the diffusion of water molecules between the two chamber parts simultaneously but retained the other dissolved components. Under isoosmotic conditions ($c_{upp} = c_{low}$), the two sides were equilibrated. Subsequently, high concentrated glucose solution (400 mosm/L) flowed through the upper chamber part ($c_{upp} \gg c_{low}$) disturbing the equilibrium state.

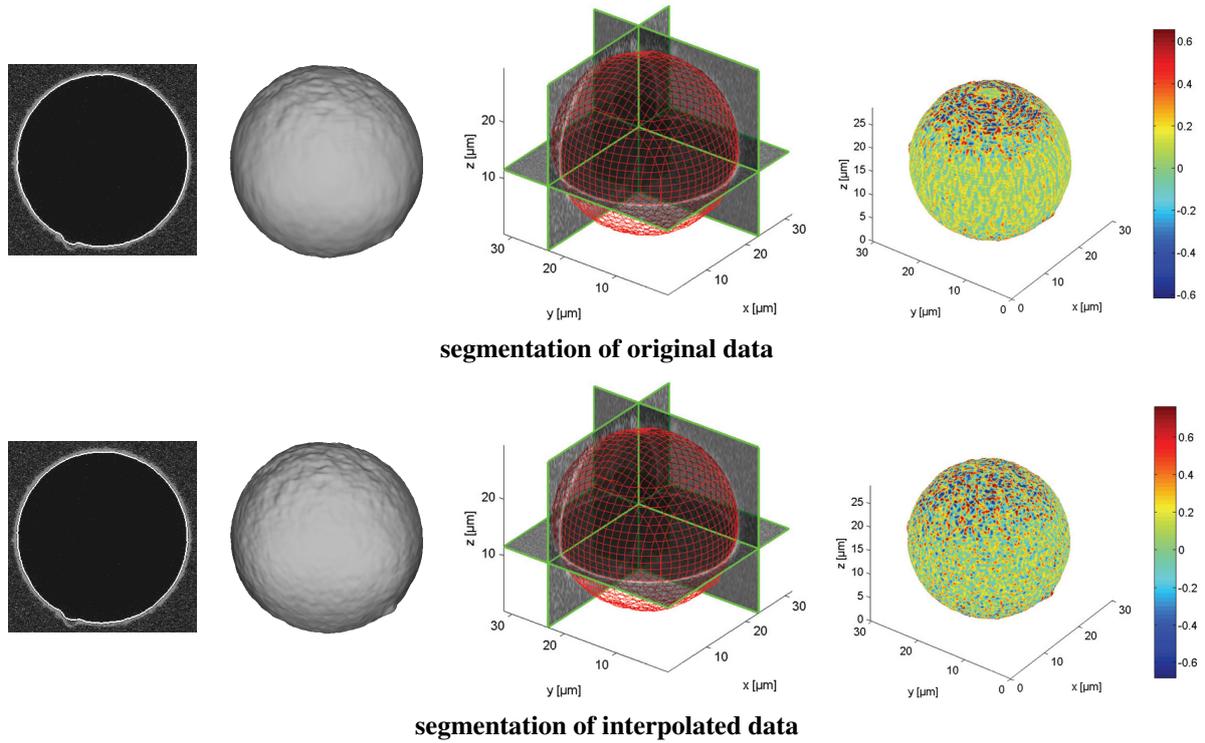


Figure 7: Data set #1, vesicle under isoosmotic conditions. *First row:* Segmentation of raw microscopic data. *Second row:* Segmentation of interpolated data. *From left to right:* Superimposed segmentation for one of the volume slices, rendered view of the extracted shape, cross-sections through the volume with superimposed wireframe segmentation result and color-coded curvature of the surface [μm^{-1}]. The figure is best viewed in color. This experiment demonstrates that upsampling the resolution in the z -direction does not improve the segmentation result, but rather increases the computation time by about a factor of 4.

To compensate the new concentration difference between the two sides, water flowed from the lower chamber into the upper flow-chamber until the equilibrium state was reached again. Consequently, the osmolarity in the lower chamber was continuously changed by 0 to 140 mosm/L in 1.5 h. The osmolarities of glucose and vesicle solutions were determined using a freezing point depression osmometer (Osmomat 030, Gonotec, Berlin, Germany).

6.1.4. Microscopic Imaging

Three dimensional membrane curvatures of streptavidin coated giant vesicles were imaged with a laser scanning microscope (LSM 510, Carl Zeiss MicroImaging GmbH, Jena, Germany) equipped with a helium-neon laser (543 nm). To detect the fluorescent signal of dextran-TRITC, a long pass filter LP600 nm was used. The microscope was focused on the vesicles using a C-Apochromate $40\times/1.20$ water immersion objective (Carl Zeiss). Images were recorded with a typical resolution of 512×512 pixels in x - y -plane (pixel size 154 nm) and 60 slices in the z -direction (pixel size 390 nm).

6.2. Evaluation on Real Data Sets

We demonstrate the performance of the proposed approach on four challenging real data sets generated with the described laser scanning methodology. Additionally, we present a comparison to alternative segmentation techniques based on level sets.

6.2.1. Performance

The data sets used in our experiments and the corresponding segmentation results are illustrated in Fig. 7, 8, 9 and 10. The data sets, acquired under different osmotic pressures, capture various deformations of the imaged vesicle. The

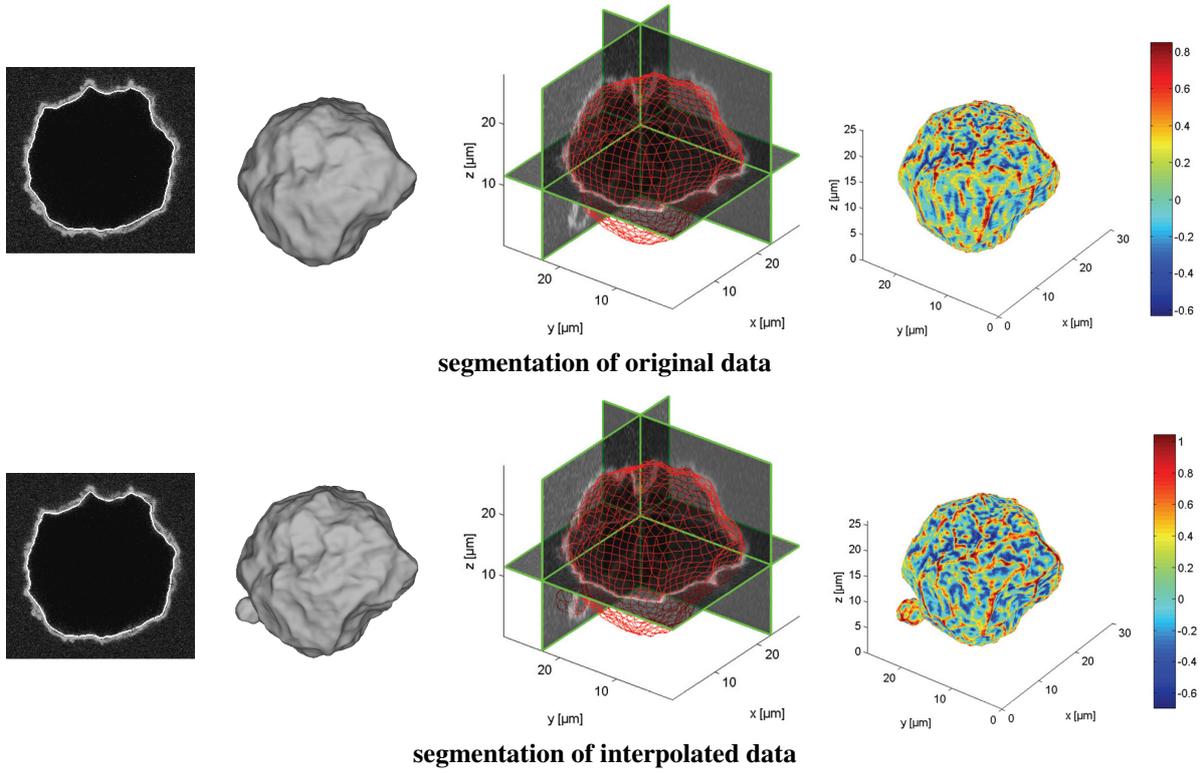


Figure 8: Data set #2. See text in Fig. 7 for reference. The slight differences in the reconstruction are most likely due to ghosting artefacts introduced in the interpolation process.

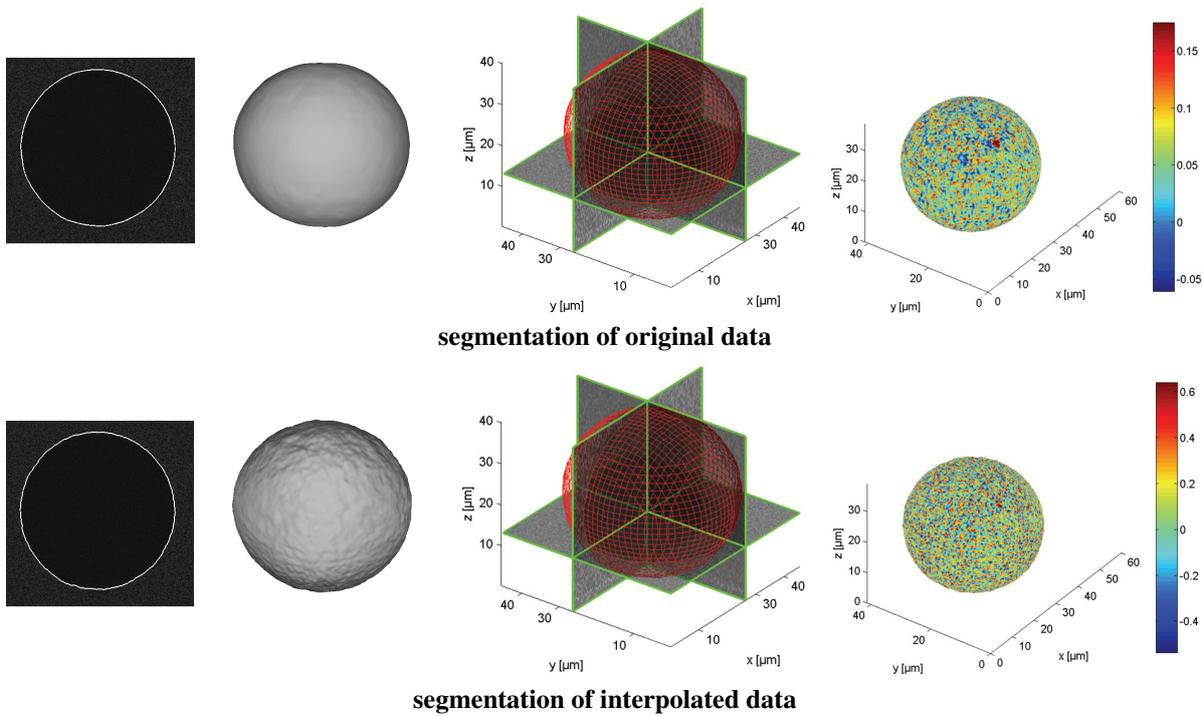


Figure 9: Data set #3. See text in Fig. 7 for reference. Again the interpolation gives rise to ghosting artefacts.

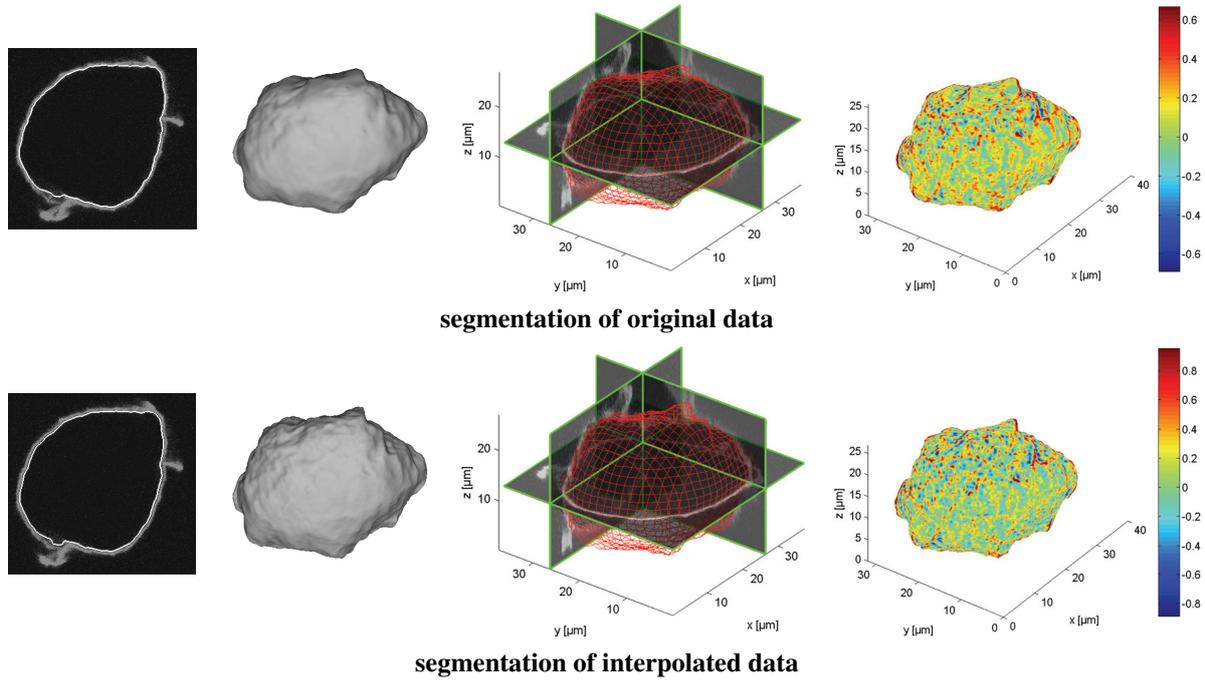


Figure 10: Data set #4. See text in Fig. 7 for reference.

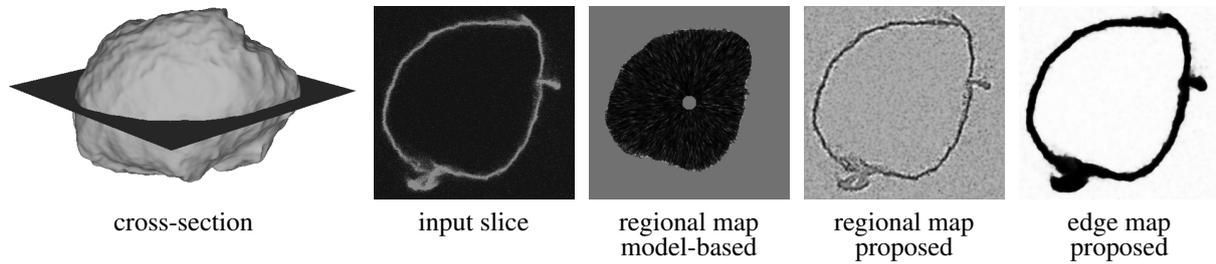


Figure 11: Comparison of the data terms utilized by the model-based approach in [37] and the presented method for slice 104 of the interpolated version of data set #4 (see Fig. 10). Although the proposed regional term exhibits lower separability of object and background than the model-based one, it offers higher precision due to the underlying global scheme and the integration of shape orientation. Note that the model-based approach doesn't use explicit edge information in contrast to the proposed one. See Fig. 12 for comparison of the segmentation results.

data set	resolution	μ_1	μ_2	β_1	β_2	γ
#1 orig.	$267 \times 274 \times 59$	10^{-4}	10^{-4}	0.2	0.025	0.05
#1 interp.	$267 \times 274 \times 256$	10^{-4}	10^{-4}	0.2	0.025	1.0
#2 orig.	$258 \times 257 \times 56$	10^{-4}	10^{-4}	0.2	0.025	0.05
#2 interp.	$258 \times 257 \times 243$	10^{-4}	10^{-4}	0.2	0.025	1.0
#3 orig.	$345 \times 342 \times 80$	10^{-4}	10^{-4}	1.0	10^{-6}	0.05
#3 interp.	$345 \times 342 \times 305$	10^{-4}	10^{-4}	1.0	10^{-6}	1.0
#4 orig.	$283 \times 275 \times 60$	10^{-4}	$3 \cdot 10^{-4}$	0.2	0.025	0.05
#4 interp.	$283 \times 275 \times 205$	10^{-4}	$3 \cdot 10^{-4}$	0.2	0.025	1.0

Table 1: Parameter setting of the proposed approach for all demonstrated experiments.

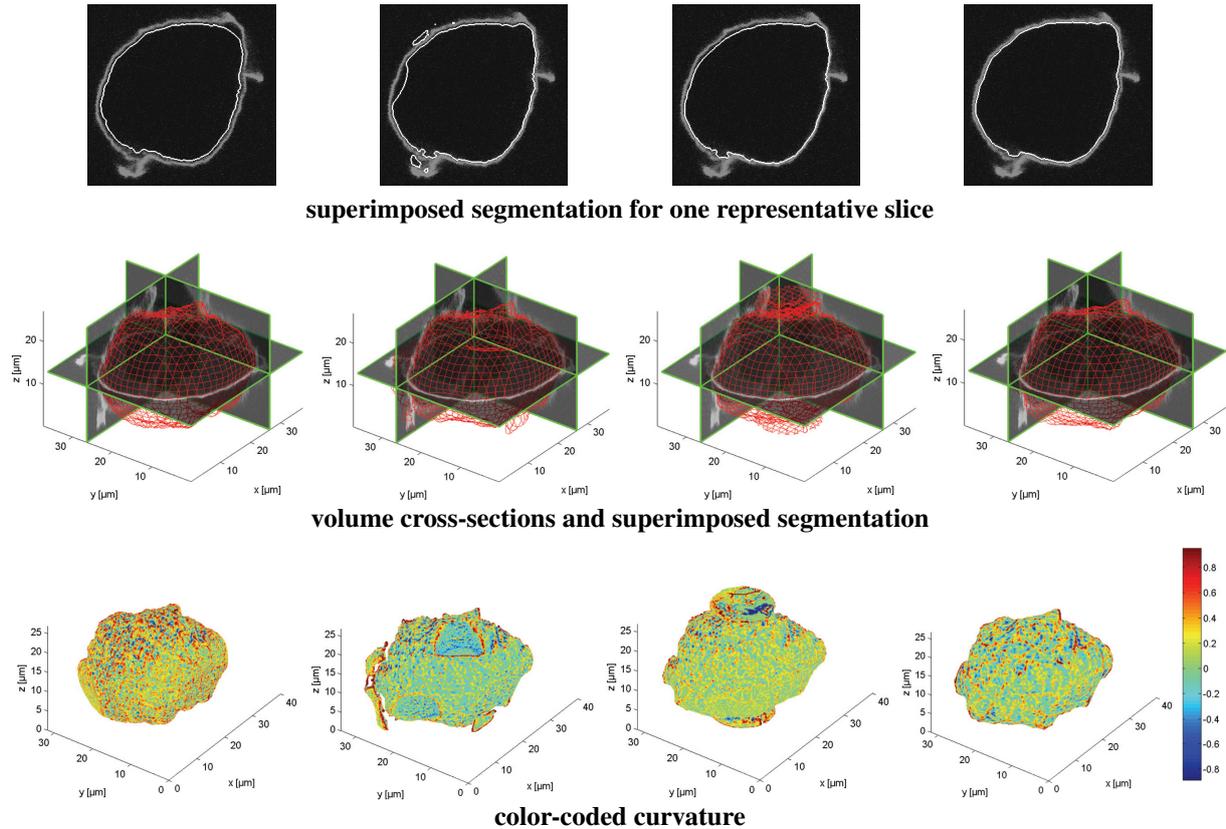


Figure 12: Comparison between the proposed approach and alternative segmentation techniques based on level sets. *First row:* Segmentation for one of the volume slices. *Second row:* Cross-sections through the volume with superimposed wireframe segmentation result. *Third row:* Color-coded curvature of the surface. *From left to right:* A variant of the model-based approach in [37], a level set implementation of the proposed energy model, as described in Section 4.1, by using a box-shaped and spherical initialization, respectively, and convex relaxation, as described in Section 4.2. Note that while the formulation in [37] involves only local data term estimation and requires a specific initialization, the presented method is general and based on global computations.

figures depict a superimposed segmentation for one of the volume slices, rendered view of the extracted 3D shape, cross-sections through the volume with superimposed wireframe segmentation result and color-coded curvature of the surface. In some cases intensity variations within the membrane regions caused high gradients, which lead to undesired structures within the membrane. We could overcome this problem by limiting the gray values to an adequate range. For all test cases we show segmentation results for both the original raw volumetric data as obtained with the described laser scanning microscopy technique and an interpolated version of the data on an approximate cubic grid. The parameter setting used in our experiments is specified in Table 1. As can be expected, the anisotropy parameter γ has been adapted appropriately so as to compensate for the sparse volume slicing in the z -direction of the original raw data. Yet, none of the other parameters has been changed. This allows to assess the effect of the interpolation process on the quality of the segmentation results. As can be observed, in both cases the segmentations are quite accurate (see the selected slices). However, they possess different degrees of smoothness, which is confirmed by the curvature measurements. This phenomenon can be explained with the appearance of small-scale ghost structures in the interpolation process. As for the runtimes, on a PC with 2.83 GHz and 8 GB of main memory we measured computational times in the range 10 – 15 minutes for the original raw data sets and in the range 40 – 60 minutes for the interpolated data sets.

The demonstrated experiments cover a wide range of test scenarios. While some of the data sets (#1 and #3) exhibit little deformation of the imaged vesicle membrane, others (#2 and #4) capture substantial bending and stretching. Furthermore, data set #3 is characterized by completely missing edge information, whereas data set #4 features lacking

foreground/background intensity discrimination. The remarkable accuracy of all segmentation results emphasizes the high practical value of the proposed approach.

6.2.2. Comparison to Alternative Approaches

In order to demonstrate the accuracy and robustness of the proposed approach, we present a twofold comparative evaluation (see Fig. 12). Firstly, we explore the viability of the energy model by comparing it to a state-of-the-art method. In particular, we implemented a variant of [37] tailored to the application at hand. Secondly, we compare an implementation of the proposed energy model with level sets, as described in Section 4.1, and convex relaxation, as described in Section 4.2.

Next, we will briefly overview the approach in [37]. Similar to the proposed formulation, it is based on a variational framework. Concretely, the following energy functional is being minimized:

$$E(u) = \int_V S(u) u(x) dx + \int_V |\nabla u(x)| dx \quad (18)$$

s. t. $u : V \rightarrow \{0, 1\}$,

where the speed function $S : V \rightarrow \mathbb{R}$ is defined as follows

$$S(x) := \frac{I_{max}}{I_{min}} - T. \quad (19)$$

Thereby, the function is evaluated only in a tight vicinity around the evolving interface by considering lines of certain length along its local normals. In practice, S is updated only on the surface and at voxels in direct contact with it. I_{max} and I_{min} denote the maximum and minimum intensity along the current line. $T \in \mathbb{R}_{\geq 0}$ is a parameter specified by the user, that controls the desired intensity contrast and is used to stop the evolution process. Note that due to its greedy nature the described procedure requires a specific initialization. It relies on the assumption that the surface can be identified by locations of high edge contrast, while the propagation scheme is initialized within a homogeneous region. Usually, the initial estimate is placed within the object interior and evolved outwards until the speed function reaches values close to zero. In our implementation, the initialization was fixed to a sphere centered in the middle of the volume with radius 10 voxels. Although the above method is able to handle data of weak foreground/background discrimination effortlessly, its precision may suffer. The utilized local update procedure could lead to premature termination and oversmoothing effects (see Fig. 12).

In the following, we summarize the main conceptual differences between the approach in [37] and the proposed one. They are revealed by a direct comparison between both variational models, given in (12) and (18).

- While the method in [37] uses a regional term based on the local intensity contrast, the proposed formulation relies on a more sophisticated computation involving foreground/background distribution modeling and surface orientation alignment. As a result, the estimated regional map is imprecise, susceptible to noise and only locally defined in the first case and offers a global range and a high degree of accuracy in the latter case, see Fig. 11. It is important to note that while the model-based approach in [37] requires a specific initialization due to the local nature of its evolution scheme, the proposed model is global and does not depend on initialization.
- The method in [37] relies on a classical Euclidean regularizer to impose smoothness of the surface. In contrast, the proposed model uses explicit edge information to avoid oversmoothing effects and increase the accuracy of the segmentation results. See Fig. 11 for an example edge map.
- We even take a further step and develop an anisotropic generalization which allows to adequately guide the smoothing along particular directions, taking into account the density of the slicing produced by the utilized imaging technology. As a consequence, the proposed approach can be applied to the measured raw volume data without any modifications, while the model in [37] is restricted to cubic voxels and requires interpolation as a preprocessing step.
- While [37] is based on the level set framework to solve the underlying minimization problem, we propose the use of convex relaxation which allows for global optimization. However, this is not a crucial issue for the method in [37], since only local data term updates are performed.

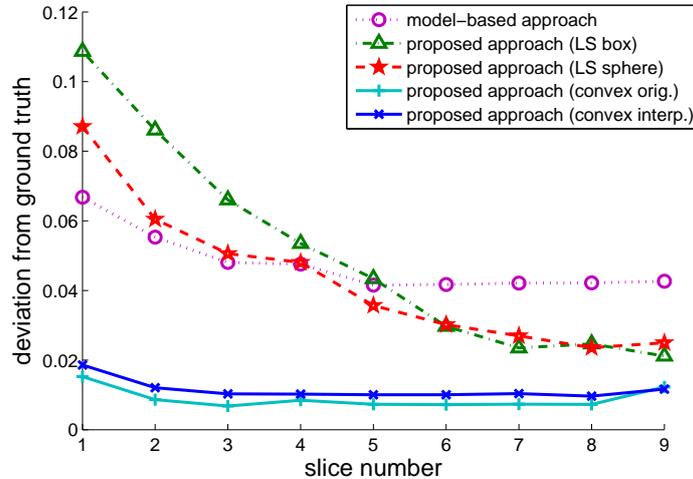


Figure 13: Quantitative evaluation of the segmentation results in Fig. 12. The accuracy of the computed segmentations is measured in terms of their deviation from a manually obtained ground truth for a couple of representative slices. Note that the proposed model optimized with level sets exhibits a quite jumpy behavior due to the local minimization scheme in combination with the globally estimated data terms. Yet, the proposed approach based on convex optimization proves superior to all other methods.

Fig. 12 compares the segmentations obtained with both methods on data set #4. We used this data, since it seems to exhibit the most challenging scenario due to the lack of foreground/background intensity discrimination. It can be observed that the method in [37] gives a relatively good segmentation result for this challenging data set, although it is a portent of a clear shrinking bias limiting the precise localization of the vesicle membrane. In contrast, the proposed approach demonstrates clear superiority and provides a quite accurate segmentation.

In order to emphasize the importance of a robust optimization scheme, we compare an implementation of the proposed energy model with level sets, as described in Section 4.1, and convex relaxation, as described in Section 4.2. Note that the level set method leads to a local minimum, which makes it highly susceptible to initialization. To this end, we provide segmentation results with two different initializations - a box centered in the middle of the volume with side length equal to half of the resolution along the corresponding dimension and a sphere centered in the middle of the volume with radius 100 voxels. In contrast, the convex relaxation technique always converges to a global minimum of the underlying energy functional and is completely independent of initialization. Expectedly, both initialization procedures for the level set implementation lead to fairly different segmentations. Note that the optimization scheme is coupled with the process of updating the foreground/background histograms. The box-shaped initialization produces a quite poor segmentation result. In particular, due to the weak foreground/background intensity discrimination and the dominance of the flux term, the generated surface consists of multiple components situated around the vesicle boundary. The spherical initialization provides a closer capturing of the imaged shape and leads to a more accurate segmentation result. The elongated structures at the top and the bottom of the surface may seem surprising at first but not after a more extensive inspection of the input data set (see the visualization of the cross-sections). And still, some oversmoothing effects of the resulting boundary can be observed. Hence, the convex relaxation technique proves essential for the precision of the final segmentation result. Recall that the proposed approach relies on a global data term computation.

The discussed conclusions are additionally confirmed by a quantitative evaluation of the above segmentation results, shown in Fig. 13. To this end, a few equidistantly spaced slices out of the volume interior have been manually segmented by an expert and used as a ground truth. Note that the slices close to the boundary are quite unclear and a ground truth segmentation is hard to obtain. If $u_{gt} : V \rightarrow \{0, 1\}$ denotes an implicit labeling representing this ground truth surface (being 0 within the interior region and 1 within the exterior) and $u : V \rightarrow \{0, 1\}$ the obtained 3D

segmentation, we measure the misalignment between them as

$$\epsilon = \frac{\int_V |u_{gr}(x) - u(x)| dx}{\int_V u_{gr}(x) dx + \int_V u(x) dx}. \quad (20)$$

In particular, we have $\epsilon \in [0, 1]$ with $\epsilon = 0$ if and only if both segmentations are identical and $\epsilon = 1$ if u is the empty set or if both labelings are complementary to each other. As evident from the plot, the proposed model optimized with level sets exhibits a jumpy behavior for both a box-shaped and a spherical initialization due to the local minimization scheme in combination with the globally estimated data terms. In contrast, the model-based method in [37] as well as the proposed approach based on convex relaxation demonstrate a stable behavior, whereas the latter offers substantially increased accuracy. This clear superiority over all other methods validates both the derived energy model and the utilized optimization technique. Note that applying the method to the original volume data without interpolation, which is not possible with the formulation in [37], entails additional accuracy gains.

In addition to the improved accuracy, the proposed approach is substantially faster than the method in [37]. For the interpolated data set, the measured computational time was about an hour for the proposed algorithm as opposed to more than 6 hours for the model-based one. The relatively high computational costs of [37] are mainly due to the time-consuming data term updates within an entire vicinity around the current surface estimate as well as the process of reinitialization of the underlying implicit function, inherent to the level set framework. The level set implementation of the proposed model took about 1.5 hours, also due to the permanently performed reinitialization.

7. Conclusion

We proposed a new robust and accurate approach for volumetric 3D segmentation with applications to vesicle membrane analysis. The method is designed to operate adequately with noisy data sets with a specific focus on biological applications like vesicle and cell segmentation. Despite this emphasis on the particular application area the approach is general and can be easily adapted to other segmentation tasks. The key idea is to tackle the segmentation problem as the minimization of an appropriate energy functional integrating various features that can be extracted from the input volume data like regional statistics, edge identification and surface orientation. Moreover, the proposed model replaces the traditionally used isotropic regularizer by an anisotropic one in order to address the non-uniform resolution and therefore also sampling of the measurements along different spatial directions. We showed that the resulting energy functional can be optimized globally in a continuous setting by means of convex relaxation. In particular, our work extends the class of functionals which can be globally minimized by convex relaxation and thresholding. Experiments on challenging fluorescence imaging data demonstrate that the proposed approach allows to accurately reconstruct the vesicle membranes despite the limited discriminative characteristics of intensity distributions and despite blurred and noisy surface boundaries. Furthermore, the reconstruction results are systematically improved by incorporating the different energy terms. Experimental results also confirm the superiority of the proposed convex relaxation approach over state-of-the-art level set methods, since it gives rise to robust and highly accurate membrane reconstructions, independent of initialization.

Appendix A.

Here, the technical details on the proof of Theorem 1 in Section 4.2 are given.

Theorem 1. *If $u^* : V \rightarrow [0, 1]$ is a minimizer of the functional in (13), then for almost any threshold $v \in (0, 1)$ the binary function $\mathbf{1}_{\Sigma_{v,u^*}}(x) : V \rightarrow \{0, 1\}$ with $\Sigma_{v,u} := \{x \mid u(x) > v\}$ is also a (global) minimizer of (13).*

Proof. We observe that the first term of $E(u^*)$ according to (13) is of the form $\int_V b(x)(1 - u^*(x))$, where $b(x) := -(\mu_1 \log P_{obj}(x) + \mu_2 \Delta I(x))$. Using the layer cake representation of the function $u^* : V \rightarrow [0, 1]$ (see [10])

$$u^*(x) = \int_0^1 \mathbf{1}_{\Sigma_{v,u^*}}(x) dv$$

we obtain

$$\begin{aligned}
& \int_V b(x)(1 - u^*(x)) dx \\
&= \int_V b(x) \left(1 - \int_0^1 \mathbf{1}_{\Sigma_{v,u^*}}(x) dv\right) dx \\
&= \int_V b(x) \left(\int_0^1 dv - \int_0^1 \mathbf{1}_{\Sigma_{v,u^*}}(x) dv\right) dx \\
&= \int_0^1 \int_V b(x) (1 - \mathbf{1}_{\Sigma_{v,u^*}}(x)) dx dv.
\end{aligned}$$

Analogously, we can rewrite the second term of $E(u^*)$ which is of the form $\int_V c(x)u^*(x) dx$ with $c(x) := -\mu_1 \log P_{bg}(x)$ as

$$\int_V c(x)u^*(x) dx = \int_0^1 \int_V c(x)\mathbf{1}_{\Sigma_{v,u^*}}(x) dx dv.$$

Furthermore, the anisotropic coarea formula [38] can be used to express the last total variation term of $E(u^*)$ as an integral over the length of all level lines of u^* , measured in the norm induced by both the weighting map g and the anisotropic tensor D , i. e.

$$\begin{aligned}
& \int_V g(I(x)) \sqrt{\nabla u^*(x)^T D(x) \nabla u^*(x)} dx = \int_0^1 |\partial \Sigma_{v,u^*}|_{g,D} dv \\
&= \int_0^1 \int_V g(I(x)) \sqrt{\nabla \mathbf{1}_{\Sigma_{v,u^*}}(x)^T D(x) \nabla \mathbf{1}_{\Sigma_{v,u^*}}(x)} dx dv,
\end{aligned}$$

where $\partial \Sigma_{v,u^*}$ denotes the boundary of Σ_{v,u^*} . Finally, plugging all together yields

$$E(u^*) = \int_0^1 E(\mathbf{1}_{\Sigma_{v,u^*}}) dv.$$

Clearly, the functional is now merely an integral over the binary characteristic functions of the upper level sets of u^* . Now, we can deduce

$$E(u^*) = \int_0^1 E(\mathbf{1}_{\Sigma_{v,u^*}}) dv \geq \min_{v \in (0,1)} E(\mathbf{1}_{\Sigma_{v,u^*}}).$$

Since u^* is assumed to be a global minimizer of E , this implies that the equality holds in the above expression and $\mathbf{1}_{\Sigma_{v_0,u^*}}$ is also a global minimizer where

$$v_0 := \operatorname{argmin}_{v \in (0,1)} E(\mathbf{1}_{\Sigma_{v,u^*}}).$$

Now, we can conclude that up to a null set the function $E(\mathbf{1}_{\Sigma_{v,u^*}})$ is constant with respect to v , i. e. $\mathbf{1}_{\Sigma_{v_0,u^*}}$ is a global minimizer of (13) for almost any $v \in (0, 1)$. \square

References

- [1] H. T. McMahon, J. L. Gallop, Membrane curvature and mechanisms of dynamic cell membrane remodelling, *Nature* 438 (2005) 590–596.
- [2] H. Engelhardt, Are s-layers exoskeletons? the basic function of protein surface layers revisited., *Journal of structural biology* 160 (2007) 115–124.
- [3] P. Ratanabangkoon, M. Gropper, R. Merkel, E. Sackmann, A. P. Gast, Two-dimensional streptavidin crystals on giant lipid bilayer vesicles, *Langmuir* 18 (2002) 4270–4276.
- [4] A. Csiszár, B. Hoffmann, R. Merkel, Double-shell giant vesicles mimicking gram-negative cell wall behavior during dehydration., *Langmuir* 25 (2009) 5753–61.
- [5] W. B. M. B. Cook, R. W. Hardy, C. S. Zuker, Preserving cell shape under environmental stress, *Nature* 452 (2008) 361–364.
- [6] J. B. Manneville, J. F. Casella, E. Ambroggio, P. Gounon, J. Bertherat, P. Bassereau, J. Cartaud, B. Antony, B. Goud, COPI coat assembly occurs on liquid-disordered domains and the associated membrane deformations are limited by membrane tension, *PNAS* 105 (2008) 16946–16951.
- [7] G. A. Vliegenthart, G. Gompper, Forced crumpling of self-avoiding elastic sheets, *Nature Materials* 5 (2006) 216–221.

- [8] T. Kosawada, K. Inoue, G. W. Schmid-Schönbein, Mechanics of curved plasma membrane vesicles: Resting shapes, membrane curvature, and in-plane shear elasticity, *Journal of Biomechanical Engineering* 127 (2005) 229–236.
- [9] J. B. Pawley, *Handbook of biological confocal microscopy*, 2nd Edition, Plenum Press, New York, 1995.
- [10] T. Chan, S. Esedoğlu, M. Nikolova, Algorithms for finding global minimizers of image segmentation and denoising models, *SIAM Journal on Applied Mathematics* 66 (5) (2006) 1632–1648.
- [11] K. Kolev, M. Klodt, T. Brox, D. Cremers, Continuous global optimization in multiview 3D reconstruction, *International Journal of Computer Vision* 84 (1) (2009) 80–96.
- [12] M. Kass, A. Witkin, D. Terzopoulos, Snakes: Active contour models, *International Journal of Computer Vision* 1 (1988) 321–331.
- [13] D. Mumford, J. Shah, Optimal approximations by piecewise smooth functions and associated variational problems, *Communications on Pure and Applied Mathematics* 42 (1989) 577–685.
- [14] S.-C. Zhu, A. Yuille, Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18 (9) (1996) 884–900.
- [15] J. Kim, J. Fisher, A. Yezzi, M. Cetin, A. Willsky, Nonparametric methods for image segmentation using information theory and curve evolution, in: *IEEE International Conference on Image Processing*, Vol. 3, Rochester, NY, 2002, pp. 797–800.
- [16] C. Rother, V. Kolmogorov, A. Blake, GrabCut: interactive foreground extraction using iterated graph cuts, *ACM Transactions on Graphics* 23 (3) (2004) 309–314.
- [17] M. Unger, T. Pock, W. Trobin, D. Cremers, H. Bischof, TVSeg - interactive total variation based image segmentation, in: *British Machine Vision Conference 2008*, Leeds, UK, 2008.
- [18] V. Caselles, R. Kimmel, G. Sapiro, Geodesic active contours, in: *Proc. Fifth International Conference on Computer Vision*, IEEE Computer Society Press, Cambridge, MA, 1995, pp. 694–699.
- [19] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, A. Yezzi, Gradient flows and geometric active contour models, in: *Proc. Fifth International Conference on Computer Vision*, IEEE Computer Society Press, Cambridge, MA, 1995, pp. 810–815.
- [20] A. Vasilevskiy, K. Siddiqi, Flux maximizing geometric flows, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (2002) 1565–1578.
- [21] X. Ren, J. Malik, Learning a classification model for segmentation, in: *Proc. International Conference on Computer Vision and Pattern Recognition*, 2003, pp. 10–17.
- [22] A. Karargyris, N. Bourbakis, An elastic video interpolation methodology for wireless capsule endoscopy videos, in: *Proceedings of the 2010 IEEE International Conference on Bioinformatics and Bioengineering*, BIBE '10, IEEE Computer Society, Washington, DC, USA, 2010, pp. 38–43.
- [23] S. Osher, J. A. Sethian, Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton–Jacobi formulations, *Journal of Computational Physics* 79 (1988) 12–49.
- [24] T. Chan, L. Vese, Active contours without edges, *IEEE Transactions on Image Processing* 10 (2) (2001) 266–277.
- [25] D. Greig, B. Porteous, A. Seheult, Exact maximum a posteriori estimation for binary images, *Journal of the Royal Statistical Society B* 51 (2) (1989) 271–279.
- [26] Y. Boykov, O. Veksler, R. Zabih, Fast approximate energy minimization via graph cuts, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (11) (2001) 1222–1239.
- [27] V. Kolmogorov, R. Zabih, What energy functions can be minimized via graph cuts?, in: *Proc. European Conference on Computer Vision*, London, UK, 2002, pp. 65–81.
- [28] Y. Boykov, V. Kolmogorov, Computing geodesics and minimal surfaces via graph cuts, in: *Proc. International Conference on Computer Vision*, IEEE Computer Society, Washington, DC, USA, 2003, p. 26.
- [29] B. Appleton, H. Talbot, Globally optimal geodesic active contours, *J. Math. Imaging Vis.* 23 (1) (2005) 67–86.
- [30] M. Klodt, T. Schoenemann, K. Kolev, M. Schikora, D. Cremers, An experimental comparison of discrete and continuous shape optimization methods, in: *European Conference on Computer Vision (ECCV)*, Marseille, France, 2008.
- [31] A. Dervieux, F. Thomasset, A finite element method for the simulation of Rayleigh–Taylor instability, in: R. Rautman (Ed.), *Approximation Methods for Navier–Stokes Problems*, Vol. 771 of *Lecture Notes in Mathematics*, Springer, Berlin, 1979, pp. 145–158.
- [32] C. Zach, M. Niethammer, J. M. Frahm, Continuous maximal flows and Wulff shapes: Application to MRFs, in: *Proc. International Conference on Computer Vision and Pattern Recognition*, Miami, FL, 2009.
- [33] C. R. Vogel, M. E. Oman, Fast, robust total variation-based reconstruction of noisy, blurred images, *IEEE Transactions on Image Processing* 7 (1998) 813–824.
- [34] Computational geometry algorithms library, <http://www.cgal.org>.
- [35] S. Rusinkiewicz, Estimating curvatures and their derivatives on triangle meshes, in: *Symposium on 3D Data Processing, Visualization, and Transmission*, 2004.
- [36] M. I. Angelova, S. Soléau, P. Méléard, J. F. Faucon, P. Bothorel, Preparation of giant vesicles by external AC electric field. kinetics and applications., in: *Progress in Colloid & Polymer Science*, Vol. 89, 1992, pp. 127–131.
- [37] A. H. Zhuang, D. J. Valentino, A. W. Toga, Skull-stripping magnetic resonance brain images using a model-based level set., *Neuroimage* 32 (1) (2006) 79–92.
- [38] C. Olsson, M. Byröd, N. C. Overgaard, F. Kahl, Extending continuous cuts: Anisotropic metrics and expansion moves, in: *International Conference on Computer Vision*, 2009.