

Harmonic active contours

Virginia Estellers*, Dominique Zosso*, Xavier Bresson^o, Jean-Philippe Thiran*

Abstract—We propose a segmentation method based on the geometric representation of images as two-dimensional manifolds embedded in a higher dimensional space. The segmentation is formulated as a minimization problem, where the contours are described by a level set function and the objective functional corresponds to the surface of the image manifold. In this geometric framework, both data-fidelity and regularity terms of the segmentation are represented by a single functional that intrinsically aligns the gradients of the level set function with the gradients of the image and exploits this directional information to overcome image inhomogeneities and fragmented contours. The proposed formulation combines this robust alignment of gradients with attractive properties of previous methods developed in the same geometric framework: the natural coupling of image channels proposed for anisotropic diffusion in [1] and the ability of subjective surfaces [2] to detect weak edges and close fragmented boundaries. The potential of such a geometric approach lies in the general definition of Riemannian manifolds, which naturally generalizes existing segmentation methods (the geodesic active contours of Caselles et al. [3], the active contours without edges of Chan and Vese [4] and the robust edge integrator of Kimmel and Bruckstein [5]) to higher dimensional spaces, non-flat images and feature spaces. Our experiments show that the proposed technique improves the segmentation of multi-channel images, images subject to inhomogeneities, and images characterized by geometric structures like ridges or valleys.

Index Terms—Image segmentation, edge detection, active contours, Beltrami.

I. INTRODUCTION

IMAGE segmentation is a first fundamental step in many applications of computer vision and machine learning because it simplifies the understanding of an image from thousands of pixels to a few regions. The goal of image segmentation is thus to partition the image domain Ω into homogeneous regions corresponding to individual objects or, equivalently, to find the contours \mathcal{C} that define the boundaries of these objects. To that purpose, the segmentation of an image is formulated as a minimization problem, where the objective functional specifies the segmentation criterion and the unknown variables describe the contours of the different regions. The level set method [6], for instance, adopts an implicit parametrization of the contours as the zero-level set of a function ϕ , which becomes then the minimization variable, and writes the objective functional in terms of this level set function.

Two representative segmentation methods within this context are the geodesic active contours (GAC) of Caselles et al. [3] and the active contours without edges (ACWE) of Chan and Vese [4]. The GAC model proposes an edge-based

segmentation criterion and defines an objective functional that weights the length of the contour with an inverse edge detector. On the other hand, ACWE proposes a region-based segmentation criterion that considers the characteristics of the different regions, and defines an objective functional that measures the variance of the grey-level values within each region in the segmentation. Both segmentation criteria can be combined and result in a segmentation method simple and extensively used in imaging, for which fast convex implementations [7]–[11] have been proposed recently. The resulting methods are extremely fast and reliable for the segmentation of cartoon-like images, but fail in the segmentation of complex images with low contrast, inhomogeneities or ridge structures. To tackle these cases, more complex segmentation criteria are necessary, and objective functionals based on second order derivatives of the image along the contours have been proposed by several authors [5], [12]–[14]. The key idea of Kimmel and Bruckstein’s robust edge integrator [5], [12], [13] is based on the observation that the direction of the image gradient is a good estimator of the orientation of an edge contour, and the segmentation criterion can be improved by introducing a geometric term in the objective functional that aligns the normal vectors to the contours with the gradients of the image. With a level set parametrization of the contours, the normal vectors to the contours are described by $\mathbf{n} = \frac{\nabla\phi}{|\nabla\phi|}$, the robust edge integrator term of Kimmel and Bruckstein becomes $-\int_{\mathcal{C}} |\nabla I \cdot \frac{\nabla\phi}{|\nabla\phi|}|$ and tries to align the normalized gradient of the level set function $\frac{\nabla\phi}{|\nabla\phi|}$ with the gradient of the image ∇I along the contour \mathcal{C} . The inclusion of this alignment term in the objective functionals of GAC or ACWE leads to more accurate segmentations, but the resulting objective functionals are not convex and suffer therefore from local minima and slow numerical minimizations. Figure 1(a) illustrates the main idea behind this alignment term and Figures 1(b)–1(c) the improvement it brings to GAC and ACWE segmentation methods.

The method that we propose goes a step further and aligns the gradients of the image ∇I with the gradients of the level set function $\nabla\phi$ for all the level sets of the image, not only for the evolving active contour \mathcal{C} . This results in a functional that is able to also exploit the alignment of the neighbouring level sets to pull the contours to the right position. Compared to the robust edge integrator, we are able to detect new contours from the alignment of the neighbouring level sets, which makes us more robust to local minima and less sensitive to initialization. A clear example is shown in Figure 1(d).

Our method also answers the question “how can we align gradients and normal vectors in images with k channels?” Generalizing the alignment term proposed by Kimmel and Bruckstein to multi-channel images is not trivial because

V. Estellers and J.P. Thiran are with the Signal Processing Laboratory (LTS5), Ecole Polytechnique Fédérale de Lausanne (EPFL).

D. Zosso is with the Department of Applied Mathematics, University of California, Los Angeles (UCLA).

X. Bresson is with the Department of Computer Science, City University of Hong Kong.

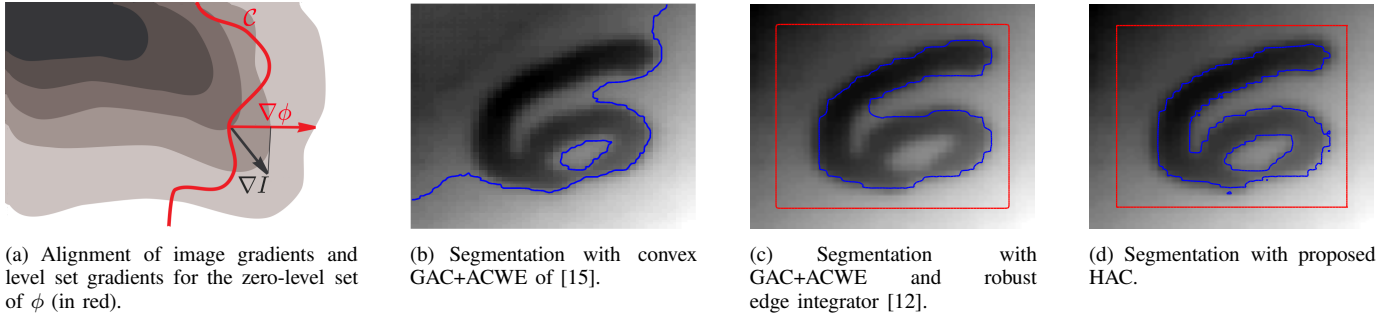


Fig. 1: Segmentation of gray-scale image subject to inhomogeneity with different methods. Initial level set set in red, final segmentation in blue. Figure 1(a) illustrates the mechanism of the proposed alignment term for the zero-level set of ϕ (in red).

considering images as vector fields in the image domain Ω , the coupling of the different channels must be defined heuristically. The natural way to treat multi-channel images is to interpret them as two-dimensional manifolds or surfaces embedded in \mathbb{R}^{k+2} and make use of differential geometry to define equivalent alignment terms. To that purpose, we develop our method in the Beltrami framework, originally proposed for color image denoising by Sochen et al. [1] and later extended to image segmentation with the subjective surfaces of Sarti and Sethian [2]. In that sense, our method introduces an alignment term in the subjective surface model, in the same way that Kimmel and Bruckstein introduced an alignment term in the GAC model.

The proposed harmonic active contours (HAC) also include region and edge-based segmentation criteria. Our objective functional, however, is not a sum of alignment terms with region and edge-based functionals, but is defined through the embedding of different image features in the Beltrami framework. In this framework, the objective functional of GAC has been generalized to two-dimensional manifolds in the subjective surfaces of Sarti and Sethian [2] by embedding the gray-level values of the image and defining the objective functional as the weighted surface of the embedded manifold. Our method presents then a generalization of GAC, ACWE and the robust edge integrator to higher dimensional spaces by embedding image channels, region features and level set function and defining the objective functional as the surface of the embedded manifold. This functional offers three main advantages. First, the alignment term between the normal vectors to the contours and the gradients of the image is naturally extended to all the level sets of ϕ , and exploits the alignment of neighbouring level sets to overcome local minima. Secondly, the definition of images as manifolds can naturally include and handle multi-channel images and take into account the coupling of the different channels. Thirdly, the proposed formulation is easily generalized to non-flat parametric manifolds [16] and feature spaces [17], where the usual segmentation methods cannot be directly applied. Compared to our conference paper [18], we have further developed the theoretical connections to existing methods, we propose a new and faster algorithm for the numerical minimization and we have extended the experimental results.

The rest of the paper is organized as follows. Section II

reviews basic concepts of differential geometry and introduces the Beltrami framework. In Section III we define our segmentation criteria as a minimal surface manifold, specifying the proposed embedding, features and metric. Section IV, describes the proposed algorithm for a fast numerical implementation of the minimization problem and Section V shows experimental results with different types of images. Finally, conclusions are drawn in Section VI.

II. DIFFERENTIAL GEOMETRY IN IMAGE PROCESSING

In order to understand the Beltrami framework, it is necessary to formalize two basic concepts of differential geometry: smooth manifolds and the metric tensor. In this Section we review basic definitions of differential geometry to provide the reader not only with the set of resulting equations, but with a more intuitive understanding of the proposed geometric framework. We introduce the Beltrami framework by particularizing the definitions to images, which will allow us later to draw the connections between our method and previous literature. We refer the reader to [19] for a more detailed description.

The simplest manifold that we can consider is a curve in the plane, which corresponds to a one-dimensional manifold embedded in \mathbb{R}^2 . Intuitively, a curve in the plane is a flexible cord that can be straighten to locally look like the real line \mathbb{R} . Formally, we describe it by means of a smooth parametrization between \mathbb{R} and the embedding space \mathbb{R}^2 , ensuring that there are no critical self-intersections in the curve by constraining the parametrization to be locally invertible. In order to measure the length of the curve or the angle between its different directions, we consider its trajectory in \mathbb{R}^2 and measure length and angles in terms of the usual coordinates and scalar product of \mathbb{R}^2 . In this Section, we generalize these concepts to n -dimensional manifolds embedded in \mathbb{R}^m , $m > n$. The definitions, however, just formalize the previous intuitive explanation of a curve in \mathbb{R}^2 (the manifold) and how we measure its length (the metric).

Definition II.1. $\mathcal{M} \subset \mathbb{R}^m$ is an n -dimensional smooth manifold in \mathbb{R}^m if for every point p in \mathcal{M} there is a local chart (σ, U) satisfying

- 1) $U \subset \mathbb{R}^n$ and $V \subset \mathbb{R}^m$ are open sets with $p \in V$

2) σ is a parametrization of \mathcal{M} around p such that

$$\begin{aligned} \sigma: U \subset \mathbb{R}^n &\longrightarrow V \subset \mathbb{R}^m \\ (u_1, \dots, u_n) &\mapsto \sigma(u_1, \dots, u_n). \end{aligned} \quad (1)$$

- σ is a differential function with $\text{rank } J\sigma = n$, where $J\sigma$ is the Jacobian of σ .
- $\sigma: U \longrightarrow V \cap \mathcal{M}$ is an homeomorphism i.e it has a continuous inverse function σ^{-1} .

The collection of local charts $\{\sigma, U\}$ of \mathcal{M} is called an atlas.

The inherent idea of the definition is the same than for the one-dimensional curve in \mathbb{R}^2 . The first point is only a technical condition to ensure differentiability. The second condition formalizes the idea that an n -dimensional smooth manifold is something that locally looks like \mathbb{R}^n and can be parametrized with a set of differential functions (the atlas) that are locally invertible (i.e., the manifold has no critical points associated with self-intersections).

For instance, the plane \mathbb{R}^2 is a 2-dimensional manifold that can be parametrized with local coordinates (x, y) and a single atlas $\{\text{id}, \mathbb{R}^2\}$. A more interesting case is the sphere $S^2 \subset \mathbb{R}^3$, a 2-dimensional manifold in \mathbb{R}^3 which cannot be covered by one single atlas and has curvilinear coordinates. In this case we can construct an atlas with the stereographic projections from the north and south poles, which cover all the sphere except the projecting point. The stereographic projection from the north pole, for instance, reads $\sigma(x, y) = \left(\frac{2x}{1+x^2+y^2}, \frac{2y}{1+x^2+y^2}, \frac{x^2+y^2-1}{1+x^2+y^2} \right)$. This will be useful later.

In the Beltrami framework introduced in [1], images are considered as 2-dimensional Riemannian manifolds embedded in the so-called space-feature manifold. For a two-dimensional grey-scale image defined over a rectangle $\Omega \subset \mathbb{R}^2$, for instance, the manifold is parametrized with the following single chart:

$$\begin{aligned} \sigma: \Omega \subset \mathbb{R}^2 &\longrightarrow \Omega \times \mathbb{R} \subset \mathbb{R}^3 \\ (x, y) &\mapsto (x, y, f(x, y)), \end{aligned} \quad (2)$$

where $f(x, y)$ is the grey level associated to point (x, y) . Assuming f to be differentiable, it is trivial to prove that σ verifies the properties of a local chart and the image can be considered a two-dimensional manifold \mathcal{M} embedded in a higher dimensional space $\mathcal{M} \subset \Omega \times \mathbb{R}$. The potential of this geometric framework lies in the general definition of the space-feature manifold and the choice of its metric. The features are not restricted to scalar values but include vector features encountered in color, texture or multi-spectral image analysis [20]. Similarly, the embedding is not limited to 2-dimensional space Ω and generalizes naturally to n -dimensional manifolds associated to volumetric data or time varying images. Moreover, the choice of the metric enables the study of complex geometries inherent to scale-space methods [21] and non-flat images generated by catadioptric or omnidirectional cameras [16].

To measure distances, angles and areas in the manifold \mathcal{M} we require the concept of a metric, that is, a scalar product in vectorial space.

Definition II.2. Given an n -dimensional smooth manifold $\mathcal{M} \subset \mathbb{R}^m$ and $p \in \mathcal{M}$, we say that $\mathbf{v} \in \mathbb{R}^m$ is a **tangent vector** to \mathcal{M} in p if there is $\epsilon > 0$ and a parametric curve $\gamma: (-\epsilon, \epsilon) \longrightarrow \mathcal{M}$ such that $\gamma(0) = p$ and $\gamma'(0) = \mathbf{v}$.

The collection of all tangent vectors to \mathcal{M} at p is called the **tangent space** $T_p\mathcal{M}$. It is easy to prove that $T_p\mathcal{M}$ is a vectorial space of dimension n , with basis $\left\{ \frac{\partial \sigma}{\partial u_1}, \dots, \frac{\partial \sigma}{\partial u_n} \right\} \Big|_{u=\sigma^{-1}(p)}$.

As in the case of the curve in \mathbb{R}^2 , distances and angles in the manifold \mathcal{M} are inherited from the standard scalar product in the embedding space \mathbb{R}^m . This operation is known as pull-back of the metric of \mathbb{R}^m to the manifold and denoted σ^* , for obvious reasons.

Definition II.3. For every point $p \in \mathcal{M}$ we consider the vectorial space $T_p\mathcal{M} \subset \mathbb{R}^m$ with the inner product \cdot induced by \mathbb{R}^m . The pull-back operation σ^* gives us a linear symmetric definite positive form \mathcal{I}_p for every $p \in \mathcal{M}$, that we call the **metric or first fundamental form** of the manifold \mathcal{M}

$$\begin{aligned} \mathcal{I}_p: T_p\mathcal{M} \times T_p\mathcal{M} &\longrightarrow \mathbb{R} \\ (\mathbf{u}, \mathbf{v}) &\mapsto \mathbf{u} \cdot \mathbf{v}. \end{aligned} \quad (3)$$

Given a local chart (σ, U) and the basis $\left\{ \frac{\partial \sigma}{\partial u_1}, \dots, \frac{\partial \sigma}{\partial u_n} \right\} = \{\sigma_1, \dots, \sigma_n\}$ of $T_p\mathcal{M}$, we can compute the expression of \mathcal{I}_p in that basis. The resulting $n \times n$ matrix G^* has entries $g_{ij}^* = \sigma_i \cdot \sigma_j$. As a bilinear form \mathcal{I}_p does not depend on the local coordinates, but its matrix expression G^* depends on the basis used for $T_p\mathcal{M}$ and, consequently, on the local chart.

In our previous examples, as the plane can be described by one single atlas $\{\text{id}, \mathbb{R}^2\}$ we have $T_p\mathbb{R}^2 = \mathbb{R}^2$ with metric given by the identity matrix I_2 . The sphere S^2 , on the contrary, requires multiple atlases and the induced metric with the stereographic projection reads $G^* = \frac{4}{(1+x^2+y^2)^2} I_2$.

In the case of images with the parametrization defined in Equation (2), we have $\{\sigma_1 = (1, 0, f_x), \sigma_2 = (0, 1, f_y)\}$ and the induced metric is given by

$$G^* = \begin{pmatrix} 1 + f_x^2 & f_x f_y \\ f_x f_y & 1 + f_y^2 \end{pmatrix}. \quad (4)$$

Naturally, the distance between two points $p_1 = (x, y)$ and $p_2 = (x + dx, y + dy)$ in the manifold is measured by the length of vector $p_1 \vec{p}_2 = \mathbf{v}$ in metric G^* . Therefore, the squared distance between these two points is given by $\mathbf{v}^T G^* \mathbf{v} = dx^2 + dy^2 + (D_v f)^2$, where $D_v f = \nabla f \cdot \mathbf{v}$ is the directional derivative of f in the direction \mathbf{v} . As a consequence, two points in the manifold are close if they are physically close in the coordinate space and their grey level values are similar. In other words, the notion of distance between two points in the image refers not only to the spatial distance, but also to the information available in the feature space. In this case, a scaling factor α is usually introduced in order to bring feature and space variables to the same scale.

To measure areas in the manifold, we recall first that, given \mathbf{u}, \mathbf{v} two non-parallel vectors in \mathbb{R}^3 the area of the parallelogram they form is given by

$$\sqrt{|\mathbf{u} \wedge \mathbf{v}|} = \sqrt{\|\mathbf{u}\|^2 \|\mathbf{v}\|^2 - (\mathbf{u} \cdot \mathbf{v})^2} = \sqrt{\det(\mathbf{u}, \mathbf{v})^T (\mathbf{u}, \mathbf{v})},$$

where (\mathbf{u}, \mathbf{v}) is the matrix of columns \mathbf{u} and \mathbf{v} . Consequently, for a 2-dimensional manifold the area of a bounded region $R \subset \mathcal{M}$ is measured by

$$\int_{\sigma^{-1}(R)} \sqrt{\|\sigma_1\|^2 \|\sigma_2\|^2 - (\sigma_1 \cdot \sigma_2)^2} = \int_{\sigma^{-1}(R)} \sqrt{\det G^*}.$$

In general, for higher dimensional manifolds $\int_{\Omega} \sqrt{\det G^*}$ measures the hyper-surface of the manifold in its metric. As a result, $g^* = \det G^*$ is defined as the squared hyper-surface element on the manifold.

In the case of two dimensional grey-scale images, the surface of the manifold is given by $\int_{\Omega} \sqrt{1 + \alpha^2 |\nabla f|^2}$ and reduces to a regularization term on the grey level values of the pixels. The role of the scale parameter α allows us to consider different norms. If $\alpha \rightarrow \infty$, the 1 inside the square root becomes negligible, and the energy approaches the total variation semi-norm commonly used in image denoising. On the other hand, if $\alpha \rightarrow 0$ the minimizing flow approaches the isotropic heat diffusion. In other words, the ‘‘size’’ or hyper-surface of the manifold associated to an image involves in a single functional both smoothness in spatial and feature coordinates, and the weight given to each term is controlled by the metric parameter α .

For images with k feature channels f^1, \dots, f^k , we make use of Einstein’s summation convention¹ to write the surface element as $g^* = 1 + \alpha_i^2 |\nabla f^i|^2 + \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2$, which takes into account the coupling of the different channels in the surface measure. Indeed, the terms $[\nabla f^i, \nabla f^j] = f_x^i f_y^j - f_y^i f_x^j$ correspond to the magnitude of the cross product of the vectors ∇f^i and ∇f^j and measure the coupling of feature channels f^i and f^j in terms of their gradients. As a consequence, Kimmel in [20] noted that the determinant of the induced metric works as a generalized edge indicator. When the determinant of the metric has a value larger than unity, it indicates the presence of a strong gradient on the manifold. A value which is close to unity indicates a region where the manifold is almost flat. In the case of vectorial images it does not only exploit the gradient ∇f^i in terms of its norm, but also in terms of its direction, that is, it preserves the information on orientation to take into account the coupling of the different channels in the definition of edges. In fact, this is the main advantage of this geometric framework for image segmentation: it offers a general tool for the evaluation of gradients in terms of an image manifold regardless of the features used.

Finally, the notion of metric and distances in \mathcal{M} , allows us to define differential operators on functions defined on the manifold. The differential operators are then defined taking into account the metric of the manifold $G = (g_{ij})$. For instance, the gradient is defined imposing that the directional derivative of function f in direction \mathbf{v} , denoted as $D_{\mathbf{v}} f$, verifies $D_{\mathbf{v}} f = \nabla f \cdot \mathbf{v}$ with the scalar product defined with the metric of the manifold. The gradient operator corresponds then to $\nabla_{\mathcal{M}} f = g^{ij} f_j du^i$ and the squared norm of the gradient of a function is given by $\|\nabla_{\mathcal{M}} f\|_{\mathcal{M}}^2 = g^{ij} f_i f_j$, where (g^{ij}) is the inverse of the metric matrix G and $f_i = \frac{\partial f}{\partial u^i}$.

¹Summation is assumed for variables with the same sub- and super-indices

III. HARMONIC ACTIVE CONTOURS

We formulate our segmentation method as a minimization problem. Our minimization variable is the level set function ϕ used to describe the contours associated to each region of the segmentation and the objective functional is the surface associated to the image manifold, which is derived from a particular space-feature embedding. Consequently the minimizing function is an harmonic map, that can be interpreted as a generalization of a geodesic curve to higher dimensional manifolds, and identifies our technique as an harmonic active contour.

We start this Section describing our geometric interpretation of images, introducing the objective functional and justifying its use in image segmentation. In our experiments and numerical implementation we consider two-dimensional images defined in the plane, but the proposed method is easily extended to higher dimensional spaces and more complex space geometries, e.g., omnidirectional images or measurements in the sphere. For this reason, we keep the initial formulation general and particularize it later for two dimensions.

A. Proposed manifold metric for robust image segmentation

In our formulation, the segmentation is defined by the zero-level set of a function $\phi(\mathbf{u})$ and the segmentation criterion is given by the contour of the features f^1, \dots, f^k , which might depend on ϕ . Both the features f^1, \dots, f^k , and level set function are defined in the same space manifold Σ than the images, take values in \mathbb{R} and are considered as differentiable functions in Σ . The space manifold Σ has coordinates $\mathbf{u} (u^1, \dots, u^n)$, with metric $G = [g_{\mu\nu}]$ on these coordinates². For instance, for a flat two-dimensional image, we have $\mathbf{u} = (x, y)$, $\Sigma = \Omega$ and $G = I_2$, while for omnidirectional images in the sphere parametrized with the stereographic projection we have $G = \frac{4}{(1+x^2+y^2)^2} I_2$.

We now make use of the Beltrami framework and consider an n -dimensional manifold \mathcal{M} associated to each possible segmentation of the image. This manifold \mathcal{M} is defined by the following embedding into the space-feature $\Sigma \times \mathbb{R}^{k+1}$

$$\begin{aligned} \sigma: \Sigma & \longrightarrow \Sigma \times \mathbb{R}^{k+1} \\ \mathbf{u} & \longmapsto (\mathbf{u}, f^1, \dots, f^k, \phi). \end{aligned} \quad (5)$$

Compared to previous approaches, our space-feature manifold includes both image features and level set function in the embedding. We keep the name space-feature manifold for simplicity, but we point out that **the inclusion of both level set function and features is the key point that allows the definition of an alignment term between the gradients of the level sets and the gradients of the image**. We now pull-back the metric defined in \mathcal{M} by space and features together.

²To simplify notation, we use greek indexes μ, ν to refer to space coordinates in Σ and latin ones i, j for the features. To that purpose also, the usual gradient and norms are assumed in \mathbb{R}^n and we specify with a sub-index ∇_{Σ} , $\|\cdot\|_{\Sigma}$ any other case.

In $\Sigma \times \mathbb{R}^{k+1}$ we consider the metric

$$\begin{pmatrix} G & 0 & \dots & 0 & 0 \\ 0 & \alpha_1 & \dots & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & \dots & \alpha_k & 0 \\ 0 & 0 & \dots & 0 & \beta \end{pmatrix}, \quad (6)$$

which offers the following interpretation. The distances on the coordinates \mathbf{u} associated to the space manifold are measured in the corresponding metric G of Σ , which is considered orthogonal to the features space. Similarly, the different feature channels f^i and the level set function ϕ live in the Euclidean space \mathbb{R}^{k+1} and are considered orthogonal to each other. We use the pull-back operator to induce the space-feature metric $(g_{\mu\nu}^*)$ on the image manifold \mathcal{M} and obtain

$$g_{\mu\nu}^* = g_{\mu\nu} + \alpha_i f_\mu^i f_\nu^i + \beta \phi_\mu \phi_\nu. \quad (7)$$

With the proposed embedding both the image features and segmenting functions are included in the induced metric, leading to the terms weighted by α_i and β . Previous geometric schemes for segmentation [2], [21] included only the level set function in the embedding and had to weight the hyper-surface element by an inverse edge detector in order to drive the contours of the level set function to the edges of the image, which resulted in a generalization of GAC to two-dimensional manifolds and did not include region-based criteria in the segmentation. We avoid this issue by introducing additional dimensions on the embedding. The extra dimensions in the feature space, however, do not affect the intrinsic dimension of the image manifold \mathcal{M} , which does not depend on the embedding but on the space Σ .

In this work we limit ourselves to 2-dimensional images, i.e. $\mathbf{u} = (x, y)$ and the determinant of the induced metric tensor is then computed as $g^* = g_{11}^* g_{22}^* - g_{12}^* g_{21}^*$, with the squared hyper-surface element given by

$$g^* = g + \alpha_i g g^{\mu\nu} f_\mu^i f_\nu^i + \beta g g^{\mu\nu} \phi_\mu \phi_\nu + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (8)$$

In Equation (8), ∇f is the gradient computed in coordinates \mathbf{u} , g is the surface element associated to the original metric tensor $(g_{\mu\nu})$ on Σ and $(g^{\mu\nu})$ its inverse. If we take into account the definition of differential operators in the manifold Σ , Equation (8) simplifies to

$$g^* = g + \alpha_i g \left\| \nabla_\Sigma f^i \right\|_\Sigma^2 + \beta g \|\nabla_\Sigma \phi\|_\Sigma^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (9)$$

This allows us to interpret the terms $\|\nabla_\Sigma f^i\|_\Sigma^2$ and $\|\nabla_\Sigma \phi\|_\Sigma^2$ as a simple measure of smoothness of the features and level set functions in the original metric of the space Σ . On the other hand, the terms $[\nabla f^i, \nabla f^j] = f_x^i f_y^j - f_y^i f_x^j$ correspond to the magnitude of the cross product of the vectors ∇f^i and ∇f^j and measure the coupling of the different features f^i and f^j . An equivalent term for the coupling of the level set function and the different feature channels is present in $[\nabla f^i, \nabla \phi]$. As

these features are orthogonal to each other in the space-feature manifold, their associated terms do not involve the metric of the space manifold $(g_{\mu\nu})$.

For images on manifolds conformally equivalent to the Euclidean metric, i.e. whose metric matrix is diagonal $G = \sqrt{g} I_2$, the squared surface element simplifies to

$$g^* = g + \alpha_i \sqrt{g} \|\nabla f^i\|^2 + \beta \sqrt{g} \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2.$$

This type of manifolds are commonly encountered in image processing, from the usual images defined on the plane ($\sqrt{g} = 1$) to non-flat images on the sphere ($\sqrt{g} = \frac{2}{1+x^2+y^2}$ with the stereographic projection). For the usual images on the plane, the previous expression simplifies to

$$g^* = 1 + \alpha_i \|\nabla f^i\|^2 + \beta \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (10)$$

B. Segmentation as a minimization problem

We formulate our segmentation technique as the following minimization problem

$$\min_{\phi} \int_{\Omega} \sqrt{g^*}. \quad (11)$$

We claim that the optimal segmentation function ϕ corresponds to the level set function that minimizes the surface of the image manifold. Indeed, as can be seen in Equation (10), minimization of the surface element aligns the gradients of the level set function ϕ with the gradients of the embedded features and, in fact, leads to piece-wise constant level set functions ϕ . Let us analyse it in more detail. In regions of the image where edge information exists, the gradients of the features f^1, \dots, f^k are large and minimization of the cross-terms attracts the level sets of ϕ to the edges of the image, the spatial gradient of ϕ increases and the level set function develops a discontinuity. On the other hand, inside homogeneous regions the minimization is driven by the smoothness penalty $\|\nabla \phi\|$, the level set function tends to be flat because $\sqrt{g} \approx \sqrt{1 + \beta \|\nabla \phi\|^2}$ and the objective functional approximates total variation or isotropic diffusion depending on the value of parameter β . It is interesting to analyse regions in the image corresponding to weak edges and fragmented contours. As discussed above, in regions with well defined edge information, the level curves of ϕ accumulate and the spatial gradient $\nabla \phi$ increases. Due to the regularity constraint on the level set function and the coupling terms $[\nabla f^i, \nabla \phi]$, the information about size and orientation of the fragmented edge is smoothly extended from the neighbouring pixels and the level set function can complete the missing boundaries.

The trade-off between gradient fidelity and level set regularity is controlled by the metric parameters α_i, β and the segmentation criterion is defined by the choice of features. Next, we particularize these expressions for different features f^i to obtain both contour- and region-based segmentation.

C. Segmentation criteria and feature definition

Contour-based segmentation is obtained by choosing the features to be local image descriptors. In the easiest case, we simply embed the pixel's grey level or color intensities, but more elaborate features such as semi-local texture descriptors, Wavelet or Gabor coefficients could be equally used. For simplicity we define $f^i(x, y) = I^i(x, y)$ for each channel I^i in the image

At the same time, in order to detect objects that are not defined by local gradients, but by homogeneous regions, we introduce region-based features. We adopt the piece-wise constant instance of the image segmentation model of Mumford-Shah [22] and approximate the image by a two-valued piece-wise constant function. We define two image regions $R_+ = \{(x, y) | \phi(x, y) > 0\}$ and $R_- = \{(x, y) | \phi(x, y) < 0\}$ and characterize them by the mean values of the local features inside the region, μ_+^i and μ_-^i . The resulting region features define the segmentation in two regions in terms of the zero-level set of the function ϕ , which results in an implicit parametrization of our harmonic active contours.

For each channel, the associated region features are inspired by ACWE and introduce the following region descriptor associated to R_+

$$f^i = \left(I^i - \mu_+^i \right)^2 H(\phi) = s^i H(\phi), \quad (12)$$

where H is the Heaviside function and f^i a measure of the square error associated to the piecewise approximation of R_+ . The region descriptor for R_- can be included in a similar manner, but we omit it here to shorten the equations. Its treatment and minimization are analogous to the region descriptor R_+ .

The proposed segmentation method combines contour and region-based criteria by including both contour and region-based features in the embedding. The resulting squared surface element reads

$$\begin{aligned} g = & 1 + \alpha_i \left\| \nabla I^i \right\|^2 + \gamma_i \left\| \nabla f^i \right\|^2 + \frac{1}{2} \alpha_i \alpha_j \left[\nabla I^i, \nabla I^j \right]^2 \\ & + \alpha_i \gamma_j \left[\nabla I^i, \nabla f^j \right]^2 + \frac{1}{2} \gamma_i \gamma_j \left[\nabla f^i, \nabla f^j \right]^2 + \beta \left\| \nabla \phi \right\|^2 \\ & + \alpha_i \beta \left[\nabla I^i, \nabla \phi \right]^2 + \gamma_i \beta \left[\nabla f^i, \nabla \phi \right]^2, \end{aligned} \quad (13)$$

where $\left\| \nabla \phi \right\|^2$ controls the smoothness of the level set function, $\left\| \nabla f^i \right\|^2$ the smoothness of the region decomposition of the image, $\left[\nabla I^i, \nabla f^j \right]^2$ measure the coherence between region and contour criteria and the terms $\left[\nabla I^i, \nabla \phi \right]^2$ and $\left[\nabla f^i, \nabla \phi \right]^2$ the alignment between the contours of level set function and the edges of the original image and its region decomposition. The parameters of the metric associated to each feature are α_i for the image channel I^i and γ_i for its associated region term f^i . We can further simplify the metric to $\alpha_i = \frac{\alpha}{m_\alpha^i}$ for the contour features and $\gamma_i = \frac{\gamma}{m_\gamma^i}$ for the region ones, where m_α^i and m_γ^i normalize the feature range of each channel I^i and f^i . We have now three parameters controlling the segmentation: α controls the contour criterion associated to the edge detector of the image features, γ controls the region criterion corresponding

to a piece-wise constant decomposition of the image and β regularizes the segmentation function ϕ .

D. Relation to subjective surfaces and anisotropic diffusion

We analyse now the relation of the proposed technique with existing geometric methods proposed in the Beltrami framework [1], [2], [20], [21], [23] in order to understand the advantages of the proposed technique. For simplicity we limit the analysis to two-dimensional images in the plane.

We start by reviewing the embedding and resulting metric proposed for anisotropic diffusion and texture denoising in [1], [20], [23]. Given a multi-valued image in the plane with features f^1, \dots, f^k , the natural embedding for denoising is given by

$$\begin{aligned} \sigma^f: \Omega & \longrightarrow \Omega \times \mathbb{R}^k \\ (x, y) & \mapsto (x, y, f^1, \dots, f^k). \end{aligned} \quad (14)$$

The metric induced in the image manifold is then

$$G^f = \begin{pmatrix} 1 + \alpha_i f_x^i{}^2 & \alpha_i f_x^i f_y^i \\ \alpha_i f_x^i f_y^i & 1 + \alpha_i f_y^i{}^2 \end{pmatrix}. \quad (15)$$

The resulting surface element is given then by $g^f = 1 + \alpha_i^2 |\nabla f^i|^2 + \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2$ and works as a robust edge detector. It measures the gradients of each feature channel and defines the edges of the image considering both of the norms of the gradients (terms $|\nabla f^i|^2$ in g^f) and the coupling of the different channels (terms $[\nabla f^i, \nabla f^j]$).

A similar geometric framework has been applied to segment images in [2], where the level set function is interpreted as a subjective surface. Sarti and Sethian in [2] substitute the previous image features f^1, \dots, f^k by the level set function ϕ to define the embedding, that is,

$$\begin{aligned} \sigma^\phi: \Omega & \longrightarrow \Omega \times \mathbb{R}^k \\ (x, y) & \mapsto (x, y, \phi). \end{aligned} \quad (16)$$

With the usual scalar product in the space-feature $\Omega \times \mathbb{R} \subset \mathbb{R}^3$, the induced surface element is given by $g^\phi = 1 + \beta \left\| \nabla \phi \right\|^2$ and measures the area of the level set function seen as a surface. In order to detect contours, the authors weight this surface element with an inverse edge detector $w(x, y)$ and minimize the weighted surface of the manifold $\int_\Omega w \sqrt{1 + \beta \left\| \nabla \phi \right\|^2}$. Consequently, this method is a generalization of GAC to two-dimensional manifolds, as shown in [21]. An equivalent formulation of this functional is obtained weighting directly the elements of the metric of $\Omega \times \mathbb{R}$ by the edge detector, that is, considering the scalar product in $\Omega \times \mathbb{R}$ given by the following symmetric positive semi-definite matrix

$$\begin{pmatrix} w & 0 & 0 \\ 0 & w & 0 \\ 0 & 0 & \beta w \end{pmatrix}. \quad (17)$$

The edge detector w is usually a simple function of the image gradient, like $w = \frac{1}{1 + a |\nabla I|^2}$, and exhibits none of the properties of the robust edge detector g^f used in diffusion. However, embedding of the level set function ϕ in the definition of

the manifold results in a segmentation technique able to cope with weak edges or fragmented boundaries [2]. The subjective surface technique, however, does not include a region-based criterion in the objective functional of the segmentation, nor does it explicitly define a contour (the level sets of ϕ that describes the contour is not specified in the objective functional, i.e., in the segmentation criterion). The inclusion of our region features of Equation (12) overcomes these two limitations, i.e., the ability to segment images based on the homogeneity of the regions and the explicit parametrization of the contour as the zero-level set of function ϕ .

In fact, it is easy to prove that our embedding corresponds to a generalization of the subjective surfaces with the following metric in the space-feature $\Omega \times \mathbb{R}$

$$\begin{pmatrix} G^f & 0 \\ 0 & \beta \end{pmatrix} = \begin{pmatrix} 1 + \alpha_i f_x^{i2} & \alpha_i f_x^i f_y^i & 0 \\ \alpha_i f_x^i f_y^i & 1 + \alpha_i f_y^{i2} & 0 \\ 0 & 0 & \beta \end{pmatrix}. \quad (18)$$

Compared to subjective surfaces, our metric in the space-feature $\Omega \times \mathbb{R}$ substitutes the edge detector w with the induced metric given by the embedding used in diffusion [1], [20], [23]. Interpreting the metric G^f as a robust edge detector, we can easily see that the proposed technique combines the advantages of both geometric approaches: diffusion methods for the robust detection of edges and subjective surfaces for its ability to close fragmented or weak edges.

E. Relation to Kimmel's robust alignment term

The proposed HAC shares a strong connection with the segmentation method proposed by Kimmel and Bruckstein in [5], [12], where a geometric alignment term between the normals to the contour and the gradient of the image is combined with the objective functionals of GAC and ACWE to improve image segmentation. In the following, we analyse the differences between the proposed HAC and Kimmel's alignment term in more detail.

In [5] Kimmel and Bruckstein observe that the direction associated to the gradient vector of an image, ∇I , is a good estimator of the orientation of the edge contour, and introduce a robust edge integrator in the objective functional that encourages the alignment between the normal vector to the contour and the gradient of the image along the contour. With an implicit parametrization of the contours $\mathcal{C} = \{\phi(x, y) = 0\}$, this alignment term or robust edge integrator reads

$$- \int_{\mathcal{C}} |\nabla I(x, y) \cdot \frac{\nabla \phi}{|\nabla \phi|}| dx dy. \quad (19)$$

From this equation it is clear that this alignment term is only active in \mathcal{C} for the zero-level set of ϕ . On the other hand, the HAC aligns the gradients of the level set function (not the normals) with the gradients of the image for all the level sets of ϕ by integrating the cross-terms $[\nabla f^i, \nabla \phi]^2$ throughout the whole image domain Ω . This gives us two theoretic differences with respect to the robust edge integrator for grey-scale images: 1) the weight given to the alignment term in HAC is proportional to the norm of the gradient of the level set function, which gives more weight to the alignment

of gradients when the level set function has larger gradients, i.e., stronger edges. 2) HAC aligns the gradients of all the level sets of the level set function, not only along the active contour, and is therefore less likely to get trapped in local minima, see Figures 1(c) and 1(d) and the experiments of Section V.

Moreover, HAC naturally generalizes to multi-channel images where both the coupling of the different image channels and the coupling of their alignment with the level set function are taken into consideration by surface of manifold of Equation (13). Finally, the proposed HAC can also be applied in non-flat parametric space and feature spaces.

In terms of implementation, Kimmel in [13] proposes an efficient algorithm to minimize the resulting objective functional with a PDE descent flow in the level set function. The algorithm combines additive operator-splitting (AOS) and a narrow-band implementation, but requires periodic redistancing of ϕ as a signed distance function. As a result, and despite the level set formulation, the alignment term is only active in a narrow band close to the zero-level set of ϕ . On the other hand, HAC uses modern optimization techniques to minimize the objective functional and is therefore faster. Our objective functional is not compatible with a narrow-band implementation or periodic redistancing because it exploits the alignment of all the level sets of ϕ with the image features.

IV. NUMERICAL MINIMIZATION

To segment an image, ultimately, we need to solve the following optimization problem

$$\min_{\phi} \int_{\Omega} \sqrt{g} dx dy \quad \text{with } g \text{ given in Equation (13)}. \quad (20)$$

In our conference paper [18], we adopted an iterative procedure based on gradient descent. In particular, we used Euler-Lagrange to obtain the optimality conditions for the minimization problem, in terms of partial differential equations, and evolved ϕ with the following flow

$$\phi_t = - \frac{1}{2\sqrt{g^*}} \left(\frac{\partial g^*}{\partial \phi} + \frac{1}{2g^*} \frac{\partial g^*}{\partial \sigma^\mu} \frac{\partial g^*}{\partial \phi_\mu} - \frac{\partial}{\partial \sigma^\mu} \frac{\partial g^*}{\partial \phi_\mu} \right) \quad (21)$$

until a fixed-point was encountered and the optimality conditions were met. To guarantee stability of the iterative scheme, the time step of the flow was limited by the CFL conditions [24], which resulted in a slow minimization technique. To avoid this limitation, we make use of recent advances in convex minimization [15], [25] and develop an efficient and easy-to-code algorithm.

The key idea of the proposed decomposition algorithm is to split the original problem into sub-optimization problems which are easy to solve and combine them together. Different options based on variable-splitting and equality constrained optimization are possible to that purpose: quadratic-penalties [25], Bregman iterations [15], [26] or the equivalent augmented Lagrangian method [27]. In our algorithm, we rewrite problem (20) as a constraint minimization and use augmented Lagrangians to solve it. The resulting Lagrangian is minimized with respect to each variable independently and the multipliers are then updated in a cyclic way. Since all the minimizations

can be analytically solved or are decoupled for each pixel, the resulting algorithm is fast, stable and easy to implement.

Let us consider the following constrained minimization problem

$$\min_{\phi, \mathbf{p}, \varphi, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k} \int_{\Omega} \sqrt{g_c} \quad \text{subject to} \quad \begin{cases} \mathbf{p} = \nabla \phi \\ \mathbf{q}^i = \nabla f^i \quad 1 \leq i \leq k \\ \varphi = \phi \\ f^i = s^i H(\varphi) \quad 1 \leq i \leq k \end{cases} \quad (22)$$

with

$$g_c = 1 + \alpha_i \left\| \nabla I^i \right\|^2 + \gamma_i \left\| \mathbf{q}^i \right\|^2 + \beta \left\| \mathbf{p} \right\|^2 + \frac{1}{2} \alpha_i \alpha_j \left[\nabla I^i, \nabla I^j \right]^2 + \alpha_i \gamma_j \left[\nabla I^i, \mathbf{q}^j \right]^2 + \frac{1}{2} \gamma_i \gamma_j \left[\mathbf{q}^i, \mathbf{q}^j \right]^2 + \alpha_i \beta \left[\nabla I^i, \mathbf{p} \right]^2 + \gamma_i \beta \left[\mathbf{q}^i, \mathbf{p} \right]^2.$$

Problem (22) is equivalent to the original problem (20).

Next, we reformulate the constrained minimization of problem (22) as an unconstrained optimization task. This can be done with an augmented Lagrangian approach [28], which translates the constraints into pairs of Lagrangian multiplier and penalty terms. Let us define the augmented Lagrangian energy $\mathcal{L}(\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi, \Lambda)$ associated to the previous problem (22):

$$\begin{aligned} \mathcal{L} = & \int_{\Omega} \sqrt{g_c} + \lambda_1 \cdot (\mathbf{p} - \nabla \phi) + \frac{r_1}{2} \left\| \mathbf{p} - \nabla \phi \right\|^2 + \lambda_{2,i} \cdot (\mathbf{q}^i - \nabla f^i) \\ & + \frac{r_{2,i}}{2} \left\| \mathbf{q}^i - \nabla f^i \right\|^2 + \lambda_3 (\varphi - \phi) + \frac{r_3}{2} (\varphi - \phi)^2 \\ & + \lambda_{4,i} \cdot (f^i - s^i H(\varphi)) + \frac{r_{4,i}}{2} (f^i - s^i H(\varphi))^2, \quad (23) \end{aligned}$$

where the Lagrange multipliers $\Lambda = (\lambda_1, \lambda_{2,1}, \dots, \lambda_{2,k}, \lambda_3, \lambda_{4,1}, \dots, \lambda_{4,k})$ are functions in Ω and $r_1, \dots, r_{4,k}$ positive constants. The constraint minimization problem (22) reduces then to finding the saddle-point of the augmented Lagrangian energy \mathcal{L} . The solution to this saddle point problem can be approximated by the following iterative algorithm. Given an initial ϕ^0 , we set the Lagrange multipliers to zero and initialize the split variables. At each iteration, an alternating minimization method is used to find an approximate minimizer of \mathcal{L} with respect to the variables $\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k$ and then the Lagrange multipliers are updated with the residuals associated to each constraint. The process is repeated until convergence and leads to Algorithm 1. The next step is to determine the solutions of the sub-minimization problems (24)-(28), which can actually be computed efficiently. In the following, we simplify notation by omitting the super-index in the different sub-minimizations.

A. Notation

We discretize the image domain $\Omega \subset \mathbb{R}^2$ with a regular grid of size $n = n_x \times n_y$. We use forward differences to compute the discrete gradients and backward differences for the divergence in order to preserve the ad-joint relationship $\text{div} = -\nabla^*$ in the discrete setting. In matrix-vector notation, we can efficiently compute the spatial derivatives multiplying the discrete functions arranged as a column vector with the sparse finite difference matrices $\nabla_x u = D_x u$, $\nabla_y u = D_y u$. Similarly, the discretization of the L_2 inner product in Ω corresponds to the usual dot product of vectors.

Algorithm 1 Lagrangian method for minimization of HAC method.

- 1: Initialize $\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi, \Lambda$
- 2: At iteration l , find a minimizer of \mathcal{L} with respect to variables $(\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi)$ with fixed Lagrange multipliers Λ^{l-1} :

$$\phi^l = \arg \min_{\phi} \mathcal{L}(\phi, \mathbf{p}^{l-1}, f^{1l-1}, \dots, f^{kl-1}, \mathbf{q}^{1l-1}, \dots, \mathbf{q}^{kl-1}, \varphi^{l-1}, \Lambda^{l-1}) \quad (24)$$

$$\mathbf{p}^l = \arg \min_{\mathbf{p}} \mathcal{L}(\phi^l, \mathbf{p}, f^{1l-1}, \dots, f^{kl-1}, \mathbf{q}^{1l-1}, \dots, \mathbf{q}^{kl-1}, \varphi^{l-1}, \Lambda^{l-1}) \quad (25)$$

$$f^{i^l} = \arg \min_{f^i} \mathcal{L}(\phi^l, \mathbf{p}^l, f^1, \dots, f^k, \mathbf{q}^{1l-1}, \dots, \mathbf{q}^{kl-1}, \varphi^{l-1}, \Lambda^{l-1}) \quad (26)$$

$$\mathbf{q}^{i^l} = \arg \min_{\mathbf{q}^i} \mathcal{L}(\phi^l, \mathbf{p}^l, f^{1^l}, \dots, f^{k^l}, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi^{l-1}, \Lambda^{l-1}) \quad (27)$$

$$\varphi^l = \arg \min_{\varphi} \mathcal{L}(\phi^l, \mathbf{p}^l, f^{1^l}, \dots, f^{k^l}, \mathbf{q}^{1^l}, \dots, \mathbf{q}^{k^l}, \varphi, \Lambda^{l-1}) \quad (28)$$
- 3: Update the Lagrange multipliers and region terms

$$\lambda_1^l = \lambda_1^{l-1} + r_1 (\mathbf{p}^l - \nabla \phi^l)$$

$$\lambda_{2,i}^l = \lambda_{2,i}^{l-1} + r_{2,i} (\mathbf{q}^{i^l} - \nabla f^{i^k}) \quad i = 1, \dots, k$$

$$\lambda_3^l = \lambda_3^{l-1} + r_3 (\varphi^l - \phi^l)$$

$$s^i = (I^i - \mu_+^i)^2 \quad \text{with} \quad \mu_+^i = \frac{\int_{\Omega} I^i H(\varphi)}{\int_{\Omega} H(\varphi)} \quad i = 1, \dots, k$$

$$\lambda_{4,i}^l = \lambda_{4,i}^{l-1} + r_{4,i} (f^{i^l} - s^i H(\varphi^l)) \quad i = 1, \dots, k$$
- 4: Stop the iterative process when $\left\| \phi^l - \phi^{l-1} \right\| < \epsilon$.

B. Minimization associated to ϕ and f^i

The sub-minimization problem (24) can be written as follows:

$$\min_{\phi} \int_{\Omega} \frac{r_1}{2} \left\| \mathbf{v} - \nabla \phi \right\|^2 + \frac{r_3}{2} (z - \phi)^2,$$

where $z = \varphi + \frac{\lambda_3}{r_3}$ and $\mathbf{v} = \mathbf{p} - \frac{\lambda_1}{r_1}$. The corresponding Euler-Lagrange equation is:

$$r_3 \phi - r_1 \Delta \phi = r_3 z - r_1 \text{div} \mathbf{z}. \quad (29)$$

With the notation explained above, Equation (29) is discretized as a linear system of equations $A_{\phi} \phi = b_{\phi}$, with matrix A_{ϕ} and vector b_{ϕ} given by

$$\begin{aligned} A_{\phi} &= r_3 I_n + r_1 D_x^T D_x + r_1 D_y^T D_y \\ b_{\phi} &= r_3 z - r_1 D_x^T \mathbf{v}_x - r_1 D_y^T \mathbf{v}_y. \end{aligned}$$

Matrix A_{ϕ} is symmetric, definite positive and block-circulant and we can use the Fourier transform \mathcal{F} to decompose it as $A_{\phi} = \mathcal{F}^T D_{\phi} \mathcal{F}$, with D_{ϕ} a diagonal matrix. Consequently, the system $A_{\phi} \phi = b_{\phi}$ can easily be solved in the Fourier domain. In practice we use the FFT transform instead of doing the matrix multiplications with \mathcal{F} and \mathcal{F}^T , which gives us the following solution $\phi = \mathcal{F}^T (D_{\phi}^{-1} \mathcal{F} b_{\phi})$ of complexity $\mathcal{O}(n \log n)$

The minimization problem associated to the features f^1, \dots, f^k have all the same form and can also be solved in the frequency domain. Without loss of generality, we present here the minimization associated to f^1 , which reads

$$\min_{f^1} \int_{\Omega} \frac{r_{2,1}}{2} \|\nabla f^1 - \mathbf{v}\|^2 + \frac{r_{4,1}}{2} (f^1 - z)^2,$$

where we have defined the auxiliary variables $\mathbf{v} = \mathbf{q}^1 + \frac{\lambda_{2,1}}{r_{2,1}}$ and $z = s^1 H - \frac{\lambda_{4,1}}{r_{4,1}}$. The corresponding Euler-Lagrange equation is

$$r_{4,1} f^1 - 2r_{2,1} \Delta f^1 = r_{4,1} z - r_{2,1} \operatorname{div} \mathbf{v},$$

which is discretized also as a linear system of equations $A_f f^1 = b_f$. Matrix A_f and vector b_f are given by

$$\begin{aligned} A_f &= r_{4,1} I_n + r_{2,1} D_x^T D_x + r_{2,1} D_y^T D_y \\ b_f &= r_{4,1} z - r_{2,1} D_x^T \mathbf{v}_x - r_{2,1} D_y^T \mathbf{v}_y, \end{aligned}$$

where A_f is also symmetric, positive definite and block-circulant and the system is solved again in the Fourier domain.

C. Minimization associated to vector fields \mathbf{p} and $\mathbf{q}^1, \dots, \mathbf{q}^k$

We first note that the minimization problem associated to the vector fields \mathbf{p} and each $\mathbf{q}^1, \dots, \mathbf{q}^k$ are decoupled for each pixel and can be solved by point-wise minimization of the functions being integrated. As the problems associated to the vector fields $\mathbf{q}^1, \dots, \mathbf{q}^k$ have all the same form, we develop only the expressions for \mathbf{q}^1 and \mathbf{p} , which read

$$\min_{\mathbf{p}} \int_{\Omega} \sqrt{g_c} + \lambda_1 (\mathbf{p} - \nabla \phi) + \frac{r_1}{2} \|\mathbf{p} - \nabla \phi\|^2 \quad (30)$$

$$\min_{\mathbf{q}^1} \int_{\Omega} \sqrt{g_c} + \lambda_{2,1} (\mathbf{q}^1 - \nabla f^1) + \frac{r_{2,1}}{2} \|\mathbf{q}^1 - \nabla f^1\|^2. \quad (31)$$

The main difficulty involved in (30) and (31) is the square root affecting the hyper-surface element $\sqrt{g_c}$, which does not allow for a closed-form solution. To overcome this issue, we use the *iterative re-weighted least squares* (IRLS) technique, which has been successfully used in the context of Beltrami in [29].

IRLS iteratively minimizes the square root term $\sqrt{g_c}$ by the following process. At iteration l the square root is approximated by the weighted hyper surface element $\sqrt{g_c^l} \approx \frac{g_c}{\sqrt{g_c^{l-1}}} = \frac{g_c}{w^l}$. Each iteration the weight is fixed to the value of the square root term from the previous iteration $w^l = \sqrt{g_c^{l-1}}$. The problem reduces to a series of a quadratic minimizations on \mathbf{p} (\mathbf{q}^1), for which a closed-form solution is available. For instance, applied to (30), at each iteration the minimization problem is equivalent to

$$\min_{\mathbf{p}} \int_{\Omega} \frac{\beta \|\mathbf{p}\|^2 + \alpha_i \beta [\nabla I^i, \mathbf{p}]^2}{w} + \frac{r_1}{2} \|\mathbf{p} - \mathbf{v}\|^2,$$

where $\mathbf{v} = \nabla \phi - \frac{\lambda_1}{r_1}$. By simple differentiation with respect to each component of $\mathbf{p} = (\mathbf{p}_x, \mathbf{p}_y)$ we obtain a 2×2 linear system of equations $A_p \mathbf{p} = r_1 w \mathbf{v}$ for each pixel. Matrix

$A_p = 2\beta \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$ is symmetric positive definite and has entries

$$\begin{aligned} A_{11} &= 1 + \alpha_i [I_y^i]^2 + \gamma_i [\mathbf{q}_y^i]^2 + \frac{r_1 w}{2\beta} \\ A_{12} &= A_{21} = \alpha_i I_x^i I_y^i + \gamma_i \mathbf{q}_x^i \mathbf{q}_y^i \\ A_{22} &= 1 + \alpha_i [I_x^i]^2 + \gamma_i [\mathbf{q}_x^i]^2 + \frac{r_1 w}{2\beta}, \end{aligned}$$

with determinant

$$|A| = \left(\frac{2\beta}{w} \right)^2 g^f + r_1 \frac{2\beta}{w} \left(2\alpha_i \|\nabla I^i\|^2 + \gamma_i \|\mathbf{q}^i\|^2 \right),$$

where g^f is the induced metric associated to the embedding $(x, y, I^1, \dots, I^k, f^1, \dots, f^k)$. With a 2×2 linear system to solve at each pixel, we have an equivalent to a closed-form solution for each IRLS update. These formulas can be further simplified by ignoring the coupling of the x, y components of \mathbf{p} and updating each one of them iteratively.

The same procedure is used to solve (31). In that case we have the following linear system for each pixel $B^1 \mathbf{q}^1 = \frac{r_{2,1}}{2\gamma_1} w \nabla f^i - \lambda_{2,1}$, with the symmetric positive matrix B^1 given by³

$$B_{11} = 1 + \sum_{i=1}^k \alpha_i [I_y^i]^2 + \sum_{i=2}^k \gamma_i [\mathbf{q}_y^i]^2 + \beta [\mathbf{p}_y]^2 + \frac{w r_{2,1}}{2\gamma_1}$$

$$B_{12} = B_{21} = \sum_{i=1}^k \alpha_i I_x^i I_y^i + \sum_{i=2}^k \gamma_i \mathbf{q}_x^i \mathbf{q}_y^i + \beta \mathbf{p}_x \mathbf{p}_y$$

$$B_{22} = 1 + \sum_{i=1}^k \alpha_i [I_x^i]^2 + \sum_{i=2}^k \gamma_i [\mathbf{q}_x^i]^2 + \beta [\mathbf{p}_x]^2 + \frac{w r_{2,1}}{2\gamma_1},$$

with determinant

$$|B^1| = \frac{1}{w^2} g^{f_1} + r_{2,1} \frac{1}{2\gamma_1 w} \left(2 \sum_{i=1}^k \alpha_i \|\nabla I^i\|^2 + \sum_{i=2}^k \gamma_i \|\mathbf{q}^i\|^2 \right),$$

where g^{f_1} is the induced metric associated to the embedding $(x, y, I^1, \dots, I^k, f^2, \dots, f^k)$. We have then also a closed-form for each IRLS update. In practice, we have found that 3 to 5 iterations of IRLS are enough for both minimization problems.

D. Minimization associated to φ

If we define the auxiliary variables $z = \phi - \frac{\lambda_3}{r_3}$, $r = r_{4,i} [s^i]^2$ and $w = \frac{1}{r} s^i (\lambda_{4,i} + r_{4,i} f^i)$, the minimization problem associated to φ is equivalent to the following minimization

$$\min_{\varphi} \int_{\Omega} \frac{r_3}{2} (\varphi - z)^2 + \frac{r}{2} (H(\varphi) - w)^2 \quad (32)$$

and can again be solved by pixel-wise minimization of the integrand $F(\varphi)$. Observe that for practical implementations, this minimization involves a smooth approximation H_ε of the Heaviside function. We propose two steps to find quickly a minimizer of (32).

³Einstein's convention is not used because the summations always exclude one index.

1) Find a solution φ^0 of (32) for $\epsilon = 0$ (i.e. for the distributional/non-smooth Heaviside function). A closed-form solution exists for this problem and can be computed as follows. The first term of $F(\varphi)$ is minimized for $\varphi^0 = z$. As the Heaviside function can take only values 0 or 1, the second term is minimized for $\varphi^0 < 0$ when $w < \frac{1}{2}$ and $\varphi^0 \geq 0$ when $w \geq \frac{1}{2}$. That means that both terms, and therefore the function, are minimized for $\varphi^0 = z$ if $w < \frac{1}{2}$ and $z < 0$ or $w \geq \frac{1}{2}$ and $z \geq 0$. Otherwise we must choose to minimize the greater of these terms and set $\varphi^0 = 0$ if $F(0) < F(z)$ and $\varphi^0 = z$ otherwise.

2) Find a solution φ of (32) for $\epsilon > 0$ using the standard Newton's method with φ^0 as initialization. The iterative Newton's method for finding the minimizer of (32) is as follows:

$$\varphi^{m+1} = \varphi^m - \frac{r_3 (\varphi^m - z) + r (H(\varphi^m) - w) \delta(\varphi^m)}{r_3 + r (H(\varphi^m) - w) \delta t(\varphi^m) + r \delta(\varphi^m)^2}.$$

Each iteration, Newton's method finds a second order polynomial approximation to the function around the current iterate and minimizes it. Initialized close to a minimum (as our first step assures), Newton's method converges fast because the second order approximation is accurate and requires only a few iterations (usually three to five in our case) to converge.

V. EXPERIMENTS

In this Section we apply the proposed method to segment different types of images to evaluate the different properties of our algorithm and compare it to related segmentation methods. All the code was implemented in Matlab and run in a laptop with an Intel Core i5 at 2.3GHz and the timings are given in seconds (s).

The first set of experiments is designed to characterize the properties of the proposed models. We start in Figures 2(a) and 2(b) by showing the ability of HAC to exploit the directional information of the image gradients to overcome the inhomogeneities present in the images and produce correct segmentations. We then show the ability of the proposed method to correctly segment ridge structures, first with a noisy image of a car-plate image in Figure 2(d) and then with the segmentation of blood vessels in a fundus image with low-contrast in Figure 2(c). Next, Figures 3(a)-3(d) show the ability of the proposed method to close fragmented contours and segment medical images with weak edges. Figures 4 present the results of the segmentation of color images, where the proposed method exploits the coupling of the different channels to detect meaningful edges. In this case, we provide also the segmentations obtained with the standard convex GAC+ACWE segmentation method, where neither the coupling of the different channels nor the alignment of image gradients with the contours are exploited to overcome the inhomogeneities of the background (Figure 4(c)) or textured areas (Figure 4(d)).

In a second set of experiments we compare the proposed HAC with the robust edge integrator of Kimmel and Bruckstein [12] combined with the GAC and ACWE models. In particular, we adopt the narrow-band implementation proposed by Kimmel in [13], with fast redistancing of the level set

function following [30]. Figures 2(e)-2(h) and 3(e)-3(h) show the results of the segmentations for Kimmel's and Bruckstein's method. For images not subject to inhomogeneities, Figures 3(e)-3(h), both methods perform similarly in terms of accuracy, but HAC is faster because we make use of variable-splitting and augmented Lagrangians to design an efficient minimization technique. The images of Figures 2(e)-2(g) are subject to inhomogeneities and cannot be segmented with a simple combination of GAC+ACWE models. In these cases, the alignment of image gradients and level set function is necessary to produce correct segmentations, and our experiments show that the proposed HAC produces slightly better segmentations in terms of accuracy (ridges of fundus image in Figures 2(c) and 2(g), or objects with pointed corners in Figures 2(a) and 2(e)), and speed. The key point, however, is the robustness of HAC to different initializations, as shown in Figure 5. In these synthetic images, Kimmel and Bruckstein's method cannot exploit the region term due to the image inhomogeneity and a balloon force is introduced to either shrink or inflate the original contour. With the wrong initialization, therefore, Kimmel's method is not able to detect new outer contours of the second of object of Figure 5(a) or the inner contours of the number of Figure 1(c) because they are too far from the zero-level set of ϕ . On the other hand, the proposed HAC aligns the gradients of all the level sets of ϕ with the gradients of the image features and is able to propagate the alignment of neighbouring level sets to detect new contours. As a result, the segmentations obtained from different initializations in Figures 2(a) and 5(c), and Figures 1(d) and 5(d) are not qualitatively different.

VI. CONCLUSIONS

In this paper we have developed a new segmentation method in a geometric framework where image segmentations are interpreted as two-dimensional manifolds embedded in a higher dimensional space, from which they inherit a metric. This metric defines distances between points in the image manifold that consider simultaneously the spatial distance between the points, the values of the features at these points and the labels assigned to them in the segmentation. Consequently, the resulting segmentation criterion incorporates region, edge and regularity terms in a single objective functional that measures the surface of the embedded manifold, naturally takes into account the coupling of the different image channels and the alignment of the level set gradients with the gradients of the image. Our formulation extends directly to higher dimensional spaces and non-flat images, where usual segmentation methods cannot be applied. For flat images in the plane, we have also developed an efficient numerical method to solve the resulting minimization problem.

Compared to the existing geodesic active contours [3], active contour without edges [4] and subjective surfaces [2], our technique exploits the directional information of the gradients of the image and level set function to align the contours of the segmentation with the edges of the image. As a result, our method is able to segment images with strong inhomogeneities. In comparison to Kimmel's and Bruckstein's

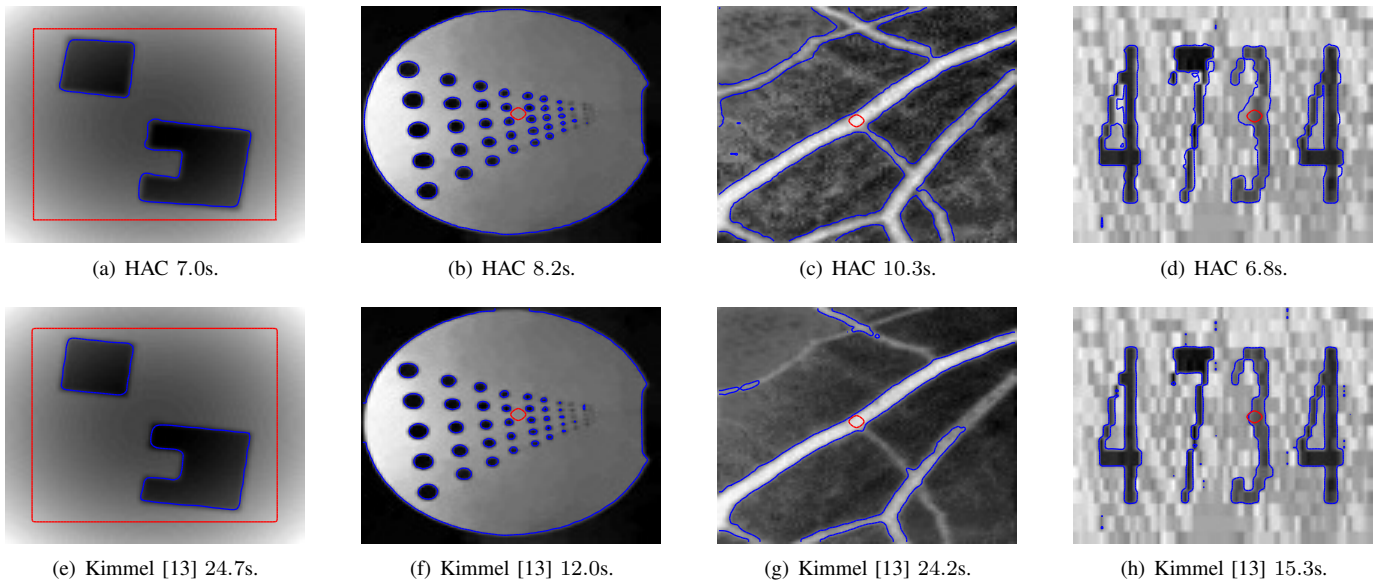


Fig. 2: Segmentation of images with ridge structures, inhomogeneities and noise with the proposed HAC (top row) and with Kimmel and Bruckstein's model [12] (bottom row). Initial level set set in red, final segmentation in blue.

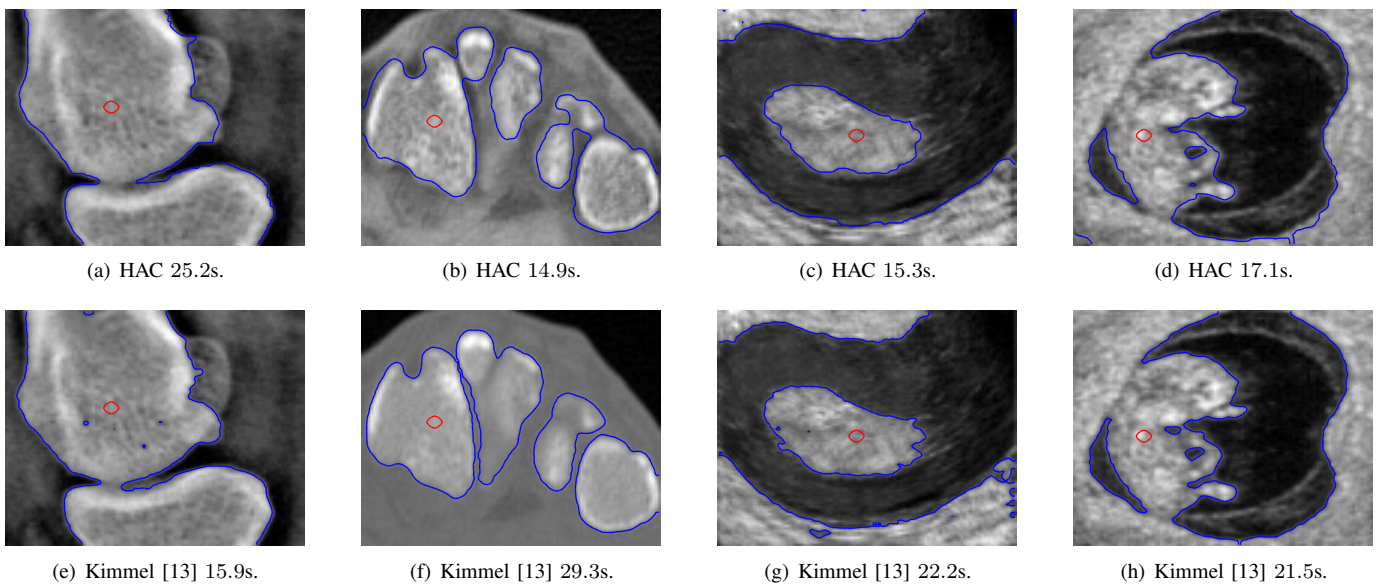


Fig. 3: Segmentation of medical images with the proposed HAC (top row) and with Kimmel and Bruckstein's model [12] (bottom row). Initial level set set in red, final segmentation in blue.

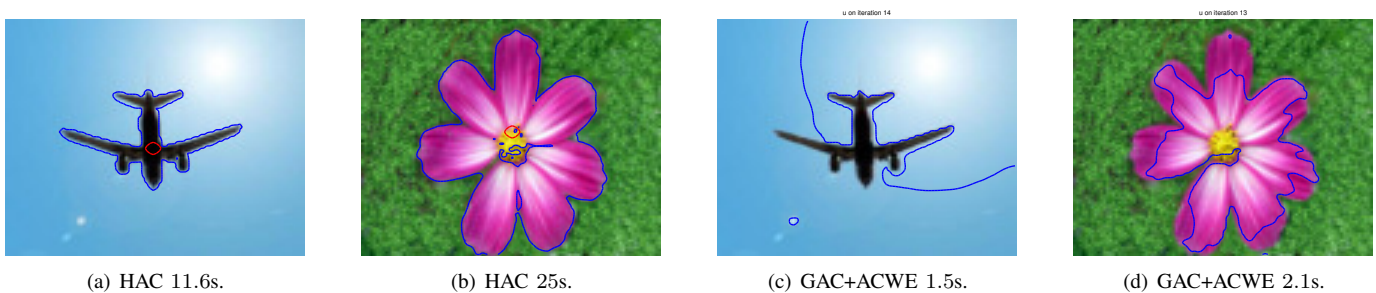


Fig. 4: Segmentation of color images with HAC model and a the convex formulation of GAC+ACWE models. Initial level set set in red, final segmentation in blue. Note that initialization does not affect the convex GAC+ACWE segmentation model.




Fig. 5: Comparison of image segmentation results with HAC and the inclusion of the robust edge integrator of Kimmel and Bruckstein [12], [13] into the GAC+ACWE model. Initial level set set in red, final segmentation in blue.

robust edge integrator [5], [12], the proposed technique aligns the gradients of the image with the gradients of the level set function for all the level sets, not only the zero-level one. As a consequence, the proposed harmonic active contour is able to exploit the alignment of the neighbouring level sets to pull the contours to the right position and discover new edges, making it less sensitive to initialization.

REFERENCES

- [1] N. Sochen, R. Kimmel, and R. Malladi, "A general framework for low level vision." *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 310–318, Jan. 1998.
- [2] A. Sarti, R. Malladi, and J. A. Sethian, "Subjective Surfaces : A Geometric Model for Boundary Completion," *International Journal of Computer Vision*, vol. 46, no. 3, pp. 201–221, 2002.
- [3] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours." *International Journal of Computer Vision*, vol. 22, no. 1, pp. 61–79, Jan. 1997.
- [4] T. F. Chan and L. Vese, "Active contours without edges." *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, Jan. 2001.
- [5] R. Kimmel and A. M. Bruckstein, "On edge detection, edge integration and geometric active contours," in *Proceedings of International Symposium on Mathematical Morphology, ISMM, 2002*, pp. 37–45.
- [6] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, Nov. 1988.
- [7] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1222–1239, 2001.
- [8] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124–1137, Sep. 2004.
- [9] T. F. Chan, S. Esedoglu, and M. Nikolova, "Algorithms for Finding Global Minimizers of Image Segmentation and Denoising Models," *SIAM Journal on Applied Mathematics*, vol. 66, no. 5, pp. 1632–1648, 2006.
- [10] X. Bresson, S. Esedoglu, P. Vandergheynst, J.-P. Thiran, and S. Osher, "Fast Global Minimization of the Active Contour/Snake Model," *Journal of Mathematical Imaging and Vision*, vol. 28, no. 2, pp. 151–167, Jul. 2007.
- [11] T. Goldstein, X. Bresson, and S. Osher, "Geometric Applications of the Split Bregman Method: Segmentation and Surface Reconstruction," *Journal of Scientific Computing*, vol. 45, no. 1-3, pp. 272–293, Nov. 2009.
- [12] R. Kimmel and A. M. Bruckstein, "Regularized Laplacian zero crossings as optimal edge integrators," *International Journal of Computer Vision*, vol. 53, no. 3, pp. 225–243, 2003.
- [13] R. Kimmel, "Fast Edge Integration," *Geometric Level Set Methods in Imaging Vision and Graphics*, pp. 59–77, 2003.
- [14] A. Vasilevskiy and K. Siddiqi, "Flux maximizing geometric flows," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 12, pp. 1565–1578, Dec. 2002.
- [15] T. Goldstein and S. Osher, "The Split Bregman Method for L1-Regularized Problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, 2009.
- [16] I. Bogdanova, X. Bresson, J.-P. Thiran, and P. Vandergheynst, "Scale space analysis and active contours for omnidirectional images." *IEEE Transactions on Image Processing*, vol. 16, no. 7, pp. 1888–901, Jul. 2007.
- [17] R. Kimmel and N. Sochen, "Orientation Diffusion or How to Comb a Porcupine," *Journal of Visual Communication and Image Representation*, vol. 13, no. 2000-02, pp. 238–248, Mar. 2002.
- [18] V. Estellers, D. Zosso, X. Bresson, and J.-P. Thiran, "Harmonic Active Contours for multichannel image processing," in *IEEE International Conference on Image Processing*, no. 200021, 2011, pp. 3141–3144.
- [19] M. P. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall, 1976.
- [20] R. Kimmel, R. Malladi, and N. Sochen, "Images as embedding maps and minimal surfaces: movies, color, and volumetric medical images," *IEEE International Conference of Computer Vision and Pattern Recognition*, vol. 39, no. 2, pp. 350–355, 1997.
- [21] X. Bresson, P. Vandergheynst, and J.-P. Thiran, "Multiscale Active Contours," *International Journal of Computer Vision*, vol. 70, no. 3, pp. 197–211, Dec. 2006.
- [22] D. Mumford and M. Shah, "Optimal approximations by piecewise smooth functions and associated variational problems," *Communications on pure and applied Mathematics*, vol. XLII, 1989.
- [23] C. Sagiv, N. Sochen, and Y. Y. Zeevi, "Integrated active contours for texture segmentation." *IEEE Transactions in Image Processing*, vol. 15, no. 6, pp. 1633–1646, Jun. 2006.
- [24] R. Courant, K. Friedrichs, and H. Lewy, "On the Partial Difference Equations of Mathematical Physics," *IBM Journal of Research*, no. March, pp. 215–234, 1967.
- [25] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM Journal on Imaging*, pp. 1–24, 2008.
- [26] W. Yin, S. Osher, D. Goldfarb, and J. Darbon, "Bregman Iterative Algorithms for l_1 -Minimization with Applications to Compressed Sensing," *SIAM Journal on Imaging Sciences*, vol. 1, no. 1, pp. 143–168, 2008.
- [27] C. Wu and X.-C. Tai, "Augmented Lagrangian method, dual methods, and split Bregman iteration for ROF, vectorial TV, and high order models," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 300–339, 2010.
- [28] R. Glowinski, J.-L. Lions, and R. Tremolieres, *Numerical Analysis of Variational Inequalities*. Elsevier, 1981.
- [29] G. Rosman, "Polyakov Action for Efficient Color Image Processing," *UCLA CAM Report 10-46*, pp. 1–21, 2010.
- [30] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proceedings of the National Academy of Sciences*, vol. 93, no. 4, pp. 1591–1595, Feb. 1996.