

Variational Problems and Partial Differential Equations on Implicit Surfaces: Bye Bye Triangulated Surfaces?

M. Bertalmío, F. Méholi, L.T. Cheng, G. Sapiro, S. Osher

ABSTRACT A novel framework for solving variational problems and partial differential equations for scalar and vector-valued data defined on surfaces is described in this chapter. The key idea is to implicitly represent the surface as the level set of a higher dimensional function, and solve the surface equations in a fixed Cartesian coordinate system using this new embedding function. The equations are then both intrinsic to the surface and defined in the embedding space. This approach thereby eliminates the need for performing complicated and not-accurate computations on triangulated surfaces, as it is commonly done in the literature. We describe the framework and present examples in computer graphics and image processing applications, including texture synthesis, flow field visualization, as well as image and vector field intrinsic regularization for data defined on 3D surfaces.

1 Introduction

In a number of applications, variational problems and partial differential equations need to be intrinsically solved for data defined on arbitrary manifolds, three dimensional surfaces in particular. Examples of this exist in the areas of mathematical physics, fluid dynamics, image processing, medical imaging, computer graphics, and pattern formation. In computer graphics, examples of this include texture synthesis [60, 64], vector field visualization [61], and weathering [18]. In other numerous applications, data defined on surfaces often needs to be regularized, e.g., as part of a vector field computation or interpolation process [47, 62], for inverse problems [24], or for surface parameterization [19]. These last regularization examples can be addressed solving a variational problem defined on the surface, or its corresponding gradient-descent flow on the surface, using for example the well developed theory of harmonic maps [21, 22], which has recently been demonstrated to be of use for image processing and computer graphics applications as well, e.g., [11, 19, 45, 52, 54, 67]. All these equations are

generally solved on triangulated or polygonal surfaces. That is, the surface is given in polygonal (triangulated) form, and the data is discretely defined on it. Solving the problems then in this representation involves the non-trivial discretization of the equations in general polygonal grids, as well as the difficult numerical computation of other quantities like projections onto the discretized surface (when computing gradients and Laplacians for example). Although the use of triangulated surfaces is extremely popular in all areas dealing with 3D models, mainly in computer graphics, there is still *not* a widely accepted technique to compute simple differential characteristics such as tangents, normals, principal directions, and curvatures; see for example [17, 39, 57] for a few of the approaches in this direction. On the other hand, it is widely accepted that computing these objects for iso-surfaces (implicit representations) is straightforward and much more accurate and robust. This problem in triangulated surfaces becomes even bigger when we not only have to compute these first and second order differential characteristics of the surface, but also have to use them to solve variational problems and PDE's for data defined on the surface. Moreover, virtually no analysis exists on numerical PDE's on non-uniform grids in the generality needed for the wide range of applications mentioned above, making it difficult to understand the behavior of the numerical implementation and its proximity (or lack thereof) to the continuous model.

In this chapter we present the framework introduced in [4] to solve variational problems and PDE's for scalar and vector-valued data defined on surfaces. We use, instead of a triangulated/polygonal representation, an implicit representation: our surface will be the zero-level set of a higher dimensional *embedding* function (i.e., a 3D volume with real values, positive outside the surface and negative inside it). Implicit surfaces have been widely used in many areas including computational physics, e.g., [43], computer graphics, e.g., [7, 25, 63], and image processing [50], as an alternative efficient representation to triangulated surfaces. We smoothly extend the original (scalar or vector-valued) data lying on the surface to the 3D volume, adapt our PDE's accordingly, and then perform all the computations on the fixed Cartesian grid corresponding to the embedding function. These computations are nevertheless intrinsic to the surface. The advantages of using the Cartesian grid instead of a triangulated mesh are many: we can use well studied numerical techniques, with accurate error, stability, and robustness measures; the topology of the underlying surface is not an issue; and we can derive simple, accurate, robust and elegant implementations. If the original surface is not already in implicit form, and it is for example triangulated, we can use any of a number of implicitation algorithms that achieve this representation given a triangulated input, e.g., [20, 34, 56, 66]. For example, the public domain software [35] can be used. If the data is just defined on the surface, an extension of it to the whole volume is also easily achieved using a PDE, as we will later see. Therefore, the method here proposed works as well for non-implicit surfaces after the preprocess-

ing is performed. This preprocessing is quite simple and no complicated regridding needs to be done to go from one surface representation to another (see below). Finally, we will solve the variational problem or PDE only in a band surrounding the zero level set (a classical approach; see [44]). Therefore, although we will be increasing by one the dimension of the space, the computations remain of the same complexity, while the accuracy and simplicity are significantly improved. A final note before we proceed. As we have mentioned, the data defined on the surfaces will only be scalar or vector valued, i.e., if we see our variational problems and PDE's as involving maps from a domain manifold into a target manifold, then our domain manifolds in this chapter will be arbitrary implicit surfaces while our target manifolds will only be the Euclidean space \mathbb{R}^n or the unit ball. But the same formulation that is presented here to deal with the domain manifold as an implicit surface can be adapted to treat the target manifold as an implicit surface as well, always working on Cartesian grids, as shown in [36]. So in general there is no limitation as to the nature of the target manifold, and we can work not just with scalar and vector-valued data but with any sort of data defining a target manifold.

1.1 The background and our contribution

Representing *deforming* surfaces as level sets of higher dimensional functions was introduced in [43] as a very efficient technique for numerically studying the deformation (see [42] for a review of this technique and also [63] for studies on the deformation and manipulation of implicit surfaces for graphics applications). The idea is to represent the surface deformation via the embedding function deformation, which adds accuracy, robustness, and, as expected, topological liberty. When the velocity of the deformation is given by the minimization of an energy, the authors in [68] proposed a “variational level set” method, where they extended the energy (originally defined only on the surface) to the whole space. This allows for the implementation to be in the Cartesian grid. The key of this approach is to go from a “surface energy” to a “volume energy” by using a Dirac’s delta function that concentrates the penalization on the given surface.

We will follow this general direction with our fixed, *non deforming* surfaces. In our case, what is being “deformed” is the (scalar or vector-valued) data on the surface. If this deformation is given by an energy-minimization problem (as is the case in data smoothing applications), we will extend the definition of the energy to the whole 3D space, and its minimization will be achieved with a PDE, which, despite its being intrinsic to the underlying surface, is also defined in the whole space. Therefore, it is easily implementable. This is straightforward, as opposed to approaches where one maps the surface data onto the plane, performs the required operations there and then maps the results back onto the triangulated representation of the surface; or approaches that attempt to solve the problem directly on

a polygonal surface.

Very interestingly, the new framework proposed here also tells us how to translate into surface terms PDE's that we know that work on the plane but which do not necessarily minimize an energy (e.g., texture synthesis or flow visualization PDE's). Instead of running these PDE's on the plane and then mapping the results onto a triangulated representation of the surface, or running them directly on the triangulated domain, we obtain a 3D straightforward Cartesian grid realization that implements the equation intrinsically on the surface and whose accuracy depends only on the degree of spatial resolution.

Moreover, we consider that for computing differential characteristics and solving PDE's even for triangulated surfaces, it might be appropriate to run an implicitation algorithm as any of the ones used for the examples in this chapter and then work on the implicit representation. Current algorithms for doing this, some of them publicly available [35], are extremely accurate and efficient.

The contribution of this work is then a new technique to efficiently solve a common problem in many computational physics, computer graphics and engineering applications: the implementation of variational problems and PDE's on 3D surfaces. In particular, we show how to transform any intrinsic variational or PDE equation into a corresponding problem for implicit surfaces. In this chapter we are then proposing a new framework to better solve existent problems and to help in building up the solutions for new ones. To exemplify the technique and its generality, we implement and extend popular equations previously reported in the literature. Here we solve them with our framework, while these problems were solved in the literature with elaborate discretizations on triangulated representations.

Before proceeding, we should comment on a few of the basic characteristics of our framework. First, as stated before, although we solve the equations in the embedding space, the basic computational complexity of our technique is not increased, since all operations are performed on a narrow band surrounding the given surface. Secondly, since the work is now on a Cartesian grid, all classical numerical analysis results on issues like robustness and stability, apply here as well. Note that for triangulated representations, new theoretical results are needed to justify the common methods proposed in the literature, while with our framework, classical and well established numerical techniques can be used, as accurate, robust, and computationally efficient as dictated by the application.

2 The framework

2.1 Surface and data representation

As mentioned before, our approach requires us to have an implicit representation of the given fixed surface, and the data must be defined in a band surrounding it and not just on the surface. The implicit surfaces used in this chapter have been derived from public-domain triangulated surfaces via the computation of a (signed) distance function $\psi(x, y, z)$ to the surface \mathcal{S} . Arriving at an implicit representation from a triangulated one is currently not a significant issue, there are publicly available algorithms that achieve it in a very efficient fashion. To exemplify this, in our chapter we have used several of these techniques. For some surfaces the classical Hamilton-Jacobi equation $\|\nabla\psi\| = 1$ was solved on a pre-defined grid enclosing the given surface via the computationally optimal approach devised in [58]. Accurate implicit surfaces from triangulations of the order of one million triangles are obtained in less than two minutes of CPU-time with this technique. Alternatively we used the implementation of the Closest Point Transform available in [35]. The teapot and knot surfaces were obtained from unorganized data points using the technique devised in [41]. We therefore assume from now on that the three dimensional surface \mathcal{S} of interest is given in implicit form, as the zero level set of a given function $\psi : \mathbb{R}^3 \rightarrow \mathbb{R}$. This function is negative inside the closed bounded region defined by \mathcal{S} , positive outside, Lipschitz continuous a.e., with $\mathcal{S} \equiv \{x \in \mathbb{R}^3 : \psi(x) = 0\}$. To ensure that the data, which needs not to be defined outside of the surface originally, is now defined in the whole band, one simple possibility is to extend this data u defined on \mathcal{S} (i.e the zero level set of ψ) in such a form that it is constant normal to each level set of ψ . This means the extension satisfies $\nabla u \cdot \nabla\psi = 0$. (For simplicity, we assume now u to be a scalar function, although we will also address in this chapter problems where the data defined on \mathcal{S} is vector-valued. This is solved in an analogous fashion.) To solve this we numerically search for the steady state solution of the Cartesian PDE

$$\frac{\partial u}{\partial t} + \text{sign}(\psi)(\nabla u \cdot \nabla\psi) = 0.$$

This technique was first proposed and used in [12]. Note that this keeps the given data u on the zero level set of ψ (the given surface) unchanged.

Both the implicitation and data extension (if required at all by the given data), need to be done only once off line. Moreover, they will remain for all applications that need this type of data.

2.2 A simple example: Heat flow on implicit surfaces

We will exemplify our framework with the simplest case, the heat flow or Laplace equation for scalar data defined on a surface. For scalar data u

defined on the plane, that is, $u(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}$ it is well known that the heat flow

$$\frac{\partial u}{\partial t} = \Delta u \quad (1.1)$$

where $\Delta := \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}$ is the Laplacian, is the gradient descent flow of the Dirichlet integral

$$\frac{1}{2} \int_{\mathbb{R}^2} \|\nabla u\|^2 dx dy, \quad (1.2)$$

where ∇ is the gradient.

Eq. (1.1) performs smoothing of the scalar data u , and this smoothing process progressively decreases the energy defined in eq. (1.2). If we now want to smooth scalar data u defined on a surface \mathcal{S} , we must find the minimizer of the *harmonic* energy given by

$$\frac{1}{2} \int_{\mathcal{S}} \|\nabla_{\mathcal{S}} u\|^2 d\mathcal{S}, \quad (1.3)$$

The equation that minimizes this energy is its gradient descent flow:

$$\frac{\partial u}{\partial t} = \Delta_{\mathcal{S}} u. \quad (1.4)$$

Here $\nabla_{\mathcal{S}}$ is the intrinsic gradient and $\Delta_{\mathcal{S}}$ the intrinsic Laplacian or Laplace-Beltrami operator. These are classical concepts in differential geometry, and basically mean the natural extensions of the gradient and Laplacian respectively, considering all derivatives intrinsic to the surface. For instance, the intrinsic gradient is just the projection onto \mathcal{S} of the regular 3D gradient while the Laplace-Beltrami operator is the projected divergence of it [51].

Classically, eq. (1.4) would be implemented in a triangulated surface, giving place to sophisticated and elaborate algorithms even for such simple flows. We now show how to simplify this when considering implicit representations.

Recall that \mathcal{S} is given as the zero level set of a function $\psi : \mathbb{R}^3 \rightarrow \mathbb{R}$, ψ is negative inside the region bounded by \mathcal{S} , positive outside with $\mathcal{S} \equiv \{x \in \mathbb{R}^3 : \psi(x) = 0\}$. We proceed now to redefine the above energy and compute its corresponding gradient descent flow. Let \vec{v} be a generic three dimensional vector, and $P_{\vec{v}}$ the operator that projects a given three dimensional vector onto the plane orthogonal to \vec{v} :

$$P_{\vec{v}} := I - \frac{\vec{v} \otimes \vec{v}}{\|\vec{v}\|^2} \quad (1.5)$$

It is then easy to show that the harmonic energy (1.3) is equivalent to (see for example [51])

$$\frac{1}{2} \int_{\mathcal{S}} \|P_{\vec{N}} \nabla u\|^2 d\mathcal{S}, \quad (1.6)$$

where \vec{N} is the normal to the surface \mathcal{S} . In other words, $\nabla_{\mathcal{S}} u = P_{\vec{N}} \nabla u$. That is, the gradient intrinsic to the surface ($\nabla_{\mathcal{S}}$) is just the projection onto the surface of the 3D Cartesian (classical) gradient ∇ .¹ We now embed this in the function ψ :

$$\begin{aligned} \frac{1}{2} \int_{\mathcal{S}} \|\nabla_{\mathcal{S}} u\|^2 d\mathcal{S} &= \frac{1}{2} \int_{\mathcal{S}} \|P_{\vec{N}} \nabla u\|^2 d\mathcal{S} \\ &= \frac{1}{2} \int_{\Omega \in \mathbb{R}^3} \|P_{\nabla \psi} \nabla u\|^2 \delta(\psi) \|\nabla \psi\| dx, \end{aligned}$$

where $\delta(\cdot)$ stands for the Dirac delta function, and all the expressions above are considered in the sense of distributions. Note that first we got rid of intrinsic derivatives by replacing $\nabla_{\mathcal{S}}$ by $P_{\vec{N}} \nabla u$ (or $P_{\nabla \psi} \nabla u$) and then replaced the intrinsic integration ($\int_{\mathcal{S}} d\mathcal{S}$) by the explicit one ($\int_{\Omega \in \mathbb{R}^3} dx$) using the delta function. Intuitively, although the energy lives in the full space, the delta function forces the penalty to be effective only on the level set of interest. The last equality includes the embedding, and it is based on the following simple facts:

1. $\nabla \psi \parallel \vec{N}$.
2. $\int_{\Omega} \delta(\psi) \|\nabla \psi\| dx = \int_{\mathcal{S}} d\mathcal{S} = \text{surface area.}$

In [4] it is shown that the gradient descent of this energy is given by

$$\frac{\partial u}{\partial t} = \frac{1}{\|\nabla \psi\|} \nabla \cdot (P_{\nabla \psi} \nabla u \parallel \nabla \psi \parallel). \quad (1.7)$$

In other words, this equation corresponds to the intrinsic heat flow or Laplace-Beltrami for data on an implicit surface. But all the gradients in this PDE are defined in the three dimensional Cartesian space, not in the surface \mathcal{S} (this is why we need the data to be defined at least on a band around the surface). The numerical implementation is then straightforward. This is the beauty of the approach! Basically, for this equation we use a classical scheme of forward differences in time and a succession of forward and backward differences in space (see [4] for details). The other equations in this chapter are similarly implemented. This follows techniques as those in [48]. Once again, due to the implicit representation and embedding in a Cartesian grid, classic numerical techniques are used, avoiding elaborate projections onto discrete surfaces and discretization on general meshes, e.g., [17, 29]. Classical numerical approaches and theoretical findings on robustness, accuracy, and error bounds, apply then for our framework.

It is easy to show a number of important properties of this equation:

¹Note that using this fact, we have transformed the computation of the norm of the intrinsic 2D gradient into an equivalent 3D Euclidean computation, see below.

1. For any second embedding function $\phi = \phi(\psi)$, with $\phi' \neq 0$ and $\phi(0) = 0$, we obtain the same gradient descent flow. Since both ψ and ϕ have to share the zero level set, and we are only interested in the flow around this zero level set, this means that the flow is (locally) independent of the embedding function.
2. If ψ is the signed distance function, a very popular implicit representation of surfaces (obtained for example from the implicitization algorithms previously mentioned), the gradient descent simplifies to

$$\frac{\partial u}{\partial t} = \nabla \cdot (P_{\nabla \psi} \nabla u). \quad (1.8)$$

We have then obtained the basic approach for embedding intrinsic variational problems. We proceed now to embed general PDE's.

2.3 From variational problems to PDE's

We note that we could also have derived eq. (1.7), directly from the harmonic maps flow

$$\frac{\partial u}{\partial t} = \Delta_{\mathcal{S}} u,$$

via the simple geometry exercise of computing the Laplace-Beltrami $\Delta_{\mathcal{S}} u$ for \mathcal{S} in implicit form (this is simply done by means of the projected derivatives as explained above, e.g., [51]). That is, the same equation is obtained when embedding the energy and then computing the gradient descent and when first looking at the gradient descent followed by the embedding of all of its components. This property is of particular significance. It basically shows how to solve general PDE's, not necessarily gradient-descent flows, for data defined on implicit surfaces. All that we need to do is to recompute the components of the PDE for implicit representations of the surface. Note that in this way, conceptually, we can redefine classical planar PDE's on implicit surfaces, making them both intrinsic to the underlying surface and defined on the whole space.

2.4 Anisotropic diffusion on implicit surfaces

From this very simple example on the Laplace-Beltrami flow we have seen the key point of our approach. If the process that we want to implement comes from the minimization of an energy, we derive a PDE for the whole space by computing the gradient-descent of the whole-space-extension of that energy. Otherwise, given a planar PDE we recompute its components for an implicit representation of the surface. For instance, anisotropic diffusion can be performed on the plane by

$$\frac{\partial u}{\partial t} = \nabla \cdot \left(\frac{\nabla u}{\|\nabla u\|} \right), \quad (1.9)$$

which minimizes the TV energy $\int_{\mathbb{R}^2} \|\nabla u\| \, dx dy$ (see [48] and also [2, 6, 46] for related formulations).

If we now want to perform intrinsic anisotropic diffusion of scalar data on a surface \mathcal{S} , we can either recompute the gradient-descent flow for the intrinsic TV energy $\int_{\mathcal{S}} \|\nabla_{\mathcal{S}} u\| \, d\mathcal{S}$, which for \mathcal{S} in implicit form becomes $\int_{\Omega \in \mathbb{R}^3} \|P_{\nabla\psi} \nabla u\| \, \delta(\psi) \, \|\nabla\psi\| \, dx$, or just substitute in eq. (1.9) the corresponding expressions as explained in the previous section. Either way we obtain the same result, the following PDE, which is valid in the embedding Euclidean space:

$$\frac{\partial u}{\partial t} = \frac{1}{\|\nabla\psi\|} \nabla \cdot \left(\frac{P_{\nabla\psi} \nabla u}{\|P_{\nabla\psi} \nabla u\|} \|\nabla\psi\| \right). \quad (1.10)$$

Note that general *p-harmonic maps*, that is, maps for L_p , $p \neq 2$, norms of the intrinsic surface gradient have been studied in the literature as well, e.g., [11, 13, 26, 28, 54]. In the following section additional equations will be presented.

3 Experimental examples

We now exemplify the framework just described for a number of important cases. The numerical implementation used is quite simple, and requires a few lines of C++ code. The CPU time required for the diffusion examples is of a few seconds on a PC (512Mb RAM, 1GHz) under Linux. For the texture synthesis examples, the CPU time ranges from a few minutes to one hour, depending on the pattern and parameters chosen. All the volumes used contain roughly 128^3 voxels. Note once again that due to the use of only a narrow band surrounding the zero level set, the order of the algorithmic complexity remains the same. On the other hand, the use of straightforward Cartesian numerics reduces the overall algorithmic complexity, improving accuracy and simplifying the implementation.

3.1 Diffusion of scalar images on surfaces

The use of PDE's for image enhancement has become one of the most active research areas in image processing [?, 50]. In particular, diffusion equations are commonly used for image regularization, denoising, and multiscale representations (representing the image simultaneously at several scales or levels of resolution). This started with the works in [33, 65], where the authors suggested the use of the linear heat flow (1.1) for this task, where u represents the image gray values (the original image is used as initial condition). Note of course that this is the basic regularization needed for inverse problems defined on surfaces, e.g., [24]. By deriving the heat flow or Laplace-Beltrami equation on implicit surfaces we then derive the basic



FIGURE 1. Intrinsic isotropic diffusion. Left: original image. Middle: after 15 diffusion steps. Right: after 50 diffusion steps.

PDE used for image processing as well as the fundamental data regularization energy/flow. As we have seen, this flow is the gradient-descent of (1.2), and the generalizations of these equations for data on the surface are given by (1.4) and (1.3) respectively. In implicit form, the heat flow on surfaces is given by (1.7). Figure 1 shows a simple example of image diffusion on a surface. Please note that this is *not* equivalent to performing 3D smoothing of the data and then looking to see what happened on \mathcal{S} . Our flow, though using extended 3D data, performs smoothing *directly* on the surface, it is an intrinsic heat flow (Laplace-Beltrami on the surface and *not* Laplace on the 3D space). The complete details of the numerical implementation of this flow are given in [4] (they once again show how the implementation is significantly simplified with the framework here described). In Figure 2 we show an example for the anisotropic flow (1.10). In this case, we have a noisy image with known variance. We can then easily add this constraint to the flow and the corresponding variational formulation. The energy corresponding to this constraint is given by ($\lambda \in \mathbb{R}$ is a parameter and u_0 is the given noisy image)

$$\frac{\lambda}{2} \int_{\mathcal{S}} (u - u_0)^2 d\mathcal{S},$$

which after it is made intrinsic and implicit becomes

$$\frac{\lambda}{2} \int_{\mathbb{R}^3} (u - u_0)^2 \delta(\psi) \|\nabla \psi\| dx.$$

In order to incorporate the constraint on the noise variance into the diffusion/denoising process, we add to the flow (1.10) the corresponding Euler-Lagrange of this energy, given by

$$\lambda(u - u_0).$$

Note in the figure how the noise is removed while the image details are preserved, as expected from an anisotropic flow. The parameter λ is estimated




FIGURE 2. Intrinsic anisotropic diffusion with constraints (automatic stop term). Left: original noisy image. Middle: after 50 diffusion steps. Right: after 90 diffusion steps. Notice how the diffusion stops and information is not smeared.

in a way suggested in [48], see [4]. The same approach, that of anisotropic diffusion with a stopping term given by the constraint, may be used to perform intrinsic *deblurring*, see [15].

We should note before proceeding that [32] also showed how to regularize images defined on a surface. The author's approach is limited to graphs (not generic surfaces) and only applies to level set based motions. The approach is simply to project the deformation of the data on the surface onto a deformation on the plane.

3.2 Diffusion of directional data on surfaces

A particularly interesting example is obtained when we have unit vectors defined on the surface. That is, we have data of the form $u : S \rightarrow S^{n-1}$. When $n = 3$ our unit vectors lie on the sphere. Examples of this data include principal directions (or general directional fields on 3D surfaces) and chromaticity vectors (normalized RGB vectors) for color images defined on the surface. This is also one of the most studied cases of the theory of harmonic maps due to its physical relationship with liquid crystals, and it was introduced in [54] for the regularization of directional data, unit vectors, on the plane (see also [11, 45, 52]). This framework of harmonic maps was used in computer graphics for texture mapping and surface parameterization, as pointed out earlier.

We still want to minimize an energy of the form

$$\int_S \|\nabla_S u\|^p dS,$$

though in this case ∇_S is the vectorial gradient and the minimizer is restricted to be a unit vector. It is easy to show, e.g., [8, 53], that the gradient

descent of this energy is given by the coupled system of PDE's

$$\frac{\partial u_i}{\partial t} = \operatorname{div}_{\mathcal{S}} (\| \nabla_{\mathcal{S}} u \|^p \nabla_{\mathcal{S}} u_i) + u_i \| \nabla_{\mathcal{S}} u \|^p, \quad 1 \leq i \leq n.$$

This flow guarantees that the initial unit vector $u(x, y, z, 0)$ remains a unit vector $u(x, y, z, t)$ all the time, thereby providing an equation for isotropic ($p = 2$) and anisotropic ($p = 1$) diffusion and regularization of unit vectors on a surface.

We can now proceed as before, and embed the surface \mathcal{S} into the zero level-set of ψ , obtaining the following gradient descent flows (see [4] for the derivation):

$$\frac{\partial u_i}{\partial t} = \frac{1}{\| \nabla \psi \|} \nabla \cdot \left(\frac{P_{\nabla \psi} \nabla u_i}{\| P_{\nabla \psi} \nabla u \|^2} \| \nabla \psi \| \right) + u_i \| P_{\nabla \psi} \nabla u \|^p. \quad (1.11)$$

Note once again that although the regularization is done intrinsically on the surface, this equation only contains Cartesian gradients. An example of this flow for anisotropic diffusion of principal direction vectors is given in Figure 3. On the left, we see the surface of a bunny with its correspondent vector field for the major principal direction. Any irregularity on the surface produces a noticeable alteration of this field, as can be seen in the details *a* and *b*. In the details *a'* and *b'*, we see the result of applying the flow (1.11). Once again, the implementation of this flow with our framework is straightforward, while it would require very sophisticated techniques on triangulated surfaces (techniques that, in addition, are not supported by theoretical results).

Following also the work [54, 55] for color images defined on the plane, we show in Figure 4 how to denoise a color image painted on an implicit surface.

The basic idea is to normalize the RGB vector (a three dimensional vector) to a unit vector representing the chroma, and diffuse this unit vector with the harmonic maps flow (1.11).² The corresponding magnitude, representing the brightness, is smoothed separately via scalar diffusion flows as those presented before (e.g., the intrinsic heat flow or the intrinsic anisotropic heat flow). That is, we have to regularize a map onto S^2 (the chroma) and another one onto \mathbb{R} (the brightness).

3.3 Pattern formation on surfaces via reaction-diffusion flows

The use of reaction-diffusion equations for texture synthesis became very popular in computer graphics following the works of Turk [60] and Witkin

²We re-normalize at every discrete step of the numerical evolution to address deviations from the unit norm due to numerical errors [16]. We could also extend the framework in [1] and apply it to our equations.




FIGURE 3. Intrinsic vector field regularization. Top: original field of major principal direction of the surface. Details a and b : original field. Details a' and b' : after anisotropic regularization.




FIGURE 4. Intrinsic vector field regularization. Left: heavy noise has been added to the 3 color channels. Right: color image reconstructed after 20 steps of anisotropic diffusion of the chroma vectors.

and Kass [64]. These works follow original ideas by Turing [59], who showed how reaction diffusion equations can be used to generate patterns. The basic idea in these models is to have a number of “chemicals” that diffuse at different rates and that react with each other. The pattern is then synthesized by assigning a brightness value to the concentration of one of the chemicals. The authors in [60, 64] used their equations for planar textures and textures on triangulated surfaces. By using the framework here described, we can simply create textures on (implicit/implicitized) surfaces, without the elaborate schemes developed in those papers.

Assuming a simple isotropic model with just two chemicals u_1 and u_2 , we have

$$\frac{\partial u_1}{\partial t} = F(u_1, u_2) + D_1 \Delta u_1,$$

$$\frac{\partial u_2}{\partial t} = G(u_1, u_2) + D_2 \Delta u_1,$$

where D_1 and D_2 are two constants representing the diffusion rates and F and G are the functions that model the reaction.

Introducing our framework, if u_1 and u_2 are defined on a surface \mathcal{S} implicitly represented as the zero level set of ψ we have

$$\frac{\partial u_1}{\partial t} = F(u_1, u_2) + D_1 \frac{1}{\|\nabla \psi\|} \nabla \cdot (P_{\nabla \psi} \nabla u_1 \parallel \nabla \psi \parallel), \quad (1.12)$$

$$\frac{\partial u_2}{\partial t} = G(u_1, u_2) + D_2 \frac{1}{\|\nabla \psi\|} \nabla \cdot (P_{\nabla \psi} \nabla u_2 \parallel \nabla \psi \parallel). \quad (1.13)$$

For simple isotropic patterns, Turk [60] selected

$$F(u_1, u_2) = s(16 - u_1 u_2),$$

$$G(u_1, u_2) = s(u_1 u_2 - u_2 - \beta),$$

where s is a constant and β is a random function representing irregularities in the chemical concentration. Examples of this, for implicit surfaces, are given in Figure 5 (the coupled PDE's shown above are run until steady state is achieved).




FIGURE 5. Texture synthesis via intrinsic reaction-diffusion flows on implicit surfaces. Left: isotropic. Right: anisotropic. Pseudo-color representation of scalar data is used. The numerical values used in the computations were $D_1 = 1.0$, $D_2 = 0.0625$, $s = 0.025$, $\beta = 12.0 \pm 0.1$, $u_1(0) = u_2(0) = 4.0$.

To simulate anisotropic textures, instead of using additional chemicals as in [60], we use anisotropic diffusion, as suggested in [64]. For this purpose, we replace eq. (1.12) with:

$$\frac{\partial u_1}{\partial t} = F(u_1, u_2) + D_1 \frac{1}{\|\nabla \psi\|} \nabla \cdot ((\vec{d} \cdot P_{\nabla \psi} \nabla u_1) \vec{d} \|\nabla \psi\|), \quad (1.14)$$

where \vec{d} is a vector field tangent to the surface, e.g., the field of the major principal direction (which for our examples has been also accurately computed directly on the implicit surface, using the technique proposed in [38]). Note how this particular selection of the anisotropic reaction-diffusion flow direction provides a texture that helps on the shape perception of the object. Additional patterns can be obtained with different combinations of the reaction and diffusion parts of the flow.

3.4 Flow visualization on 3D surfaces

Inspired by the work on line integral convolution [9] and that on anisotropic diffusion [46], the authors of [61] suggested to use anisotropic diffusion to visualize flows in 2D and 3D. The basic idea is, starting from a random

image, anisotropically diffuse it in the directions dictated by the flow field. The authors presented very nice results both in 2D (flows on the plane) and 3D (flows on a surface), but once again using triangulated surfaces which introduce many computational difficulties. In a straightforward fashion we can compute these anisotropic diffusion equations on the implicit surfaces with the framework here described, and some results are presented in Figure 6. Note the complicated topology and how both the inside and outside parts of the surfaces are easily handled with our implicit approach. Also note that, when we choose the vector field to be that of one of the principal directions, the result emphasizes the surface shape.



FIGURE 6. Flow visualization on implicit 3D surfaces via intrinsic anisotropic diffusion flows. Left: flow aligned with the major principal direction of the surface. Right: flow aligned with the minor principal direction of the surface. Pseudo-color representation of scalar data is used.

4 Concluding remarks

In this chapter, we have described a framework for solving variational problems and PDE's for data defined on surfaces. The technique borrows ideas from the level set theory and the theory of intrinsic flows via harmonic maps. The surface is embedded in a higher dimensional function, and the Euler-Lagrange flow or PDE is solved in the Cartesian coordinate system of this embedding function. The equations are simultaneously intrinsic to the implicit surface and defined on the embedding Cartesian space. With this framework we enjoy accuracy, robustness, and simplicity, as expected from the computation of differential characteristics on iso-surfaces (implicit

surfaces) and the use of classical and well established numerical techniques in Cartesian grids. In addition to presenting the general approach, we have exemplified it with equations arising in image processing and computer graphics.

We believe this new framework opens up a large number of theoretical and practical questions. In the theoretical arena, the large amount of results available for harmonic maps (see for example [54] for a review on this) need to be extended to the “implicit harmonic maps” equations presented in this chapter. The effect of perturbations on the surface (zero-level set) on the solutions of the intrinsic PDE should be investigated. This is crucial to understand the desired accuracy of surface implicitation algorithms. We expect that as with the level set theory (e.g., [14, 23]), these theoretical results will follow. On the practical side, it is of interest to address other related equations that appear in the mathematical physics, image processing, and computer graphics literature. For example, how to extend the use of harmonic maps for texture mapping (and not just texture synthesis). This was done for triangulated surfaces in [3, 19, 27]. Also to investigate threshold dynamics and convolution generated motions [31, 37, 49] for implicit surfaces. Other PDE’s, like those for image inpainting [40] or image segmentation [10] can be extended to work on implicit surfaces following the theory described in this chapter, expecting the same quality of results that were obtained on the plane. We could also use this framework to experimentally study results as those in [30]. Finally, the use of the approach here presented for regularization in inverse problems, e.g., [24], is of interest as well.

To conclude, we should note that it is natural to ask about the target manifold for the most general form of harmonic maps, when this target is not just the Euclidean space or a unit ball, but a general hypersurface. In [36] we have shown how to extend the framework described here to arbitrary *implicit* target surfaces. The key idea is again to implicitly represent the target manifold as the level-set of a higher dimensional function, and then implement the equations in the Cartesian coordinate system of this new embedding function. In the case of a variational problem, the search for the minimizing map is restricted to the class of maps whose target is the level-set of interest. In the case of partial differential equations, all the equation characteristics are implicitly represented. A set of equations that while defined on the whole Euclidean space, are intrinsic to the implicit target manifold and map into it, is then obtained. See figure 7 for an example of denoising of a texture map from \mathbb{R}^2 onto a “teapot” surface. This result was obtained with the following evolution equation, that guarantees that the diffusion of the map, though performed in a 3D Cartesian grid, is intrinsic to the implicit surface that is the teapot:

$$\frac{\partial \vec{u}}{\partial t} = \Delta \vec{u} - (\Delta \vec{u} \cdot \nabla \psi) \nabla \psi, \quad (1.15)$$




FIGURE 7. Diffusion of a texture map for an implicit teapot (noisy on the top and regularized on the bottom). A checkerboard texture is mapped.

where \vec{u} is the map and ψ is the signed distance function to the teapot.

Note also that general motion of curves on implicit surfaces is studied in [5, 15]. These works, together with the one here presented, provide then the basic framework for solving generic PDE's on implicit surfaces.

Acknowledgements

The implicitation software was provided by Santiago Betelú and H. Zhao. This work was partially supported by grants from the Office of Naval Research ONR-N00014-97-1-0509 and ONR-N00014-97-1-0027, the Office of Naval Research Young Investigator Award, the Presidential Early Career Awards for Scientists and Engineers (PECASE), a National Science Foundation CAREER Award, the National Science Foundation Learning and Intelligent Systems Program (LIS), NSF-DMS-9706827, NSF DMS 0074735, ARO-DAAG-55-98-1-0323 and IIE-Uruguay.

5 REFERENCES

- [1] F. Alouges. An energy decreasing algorithm for harmonic maps. In J.M. Coron et al., editor, *Nematics*, Nato ASI Series, pages 1–13. Kluwer Academic Publishers, Netherlands, 1991.
- [2] L. Alvarez, P. L. Lions, and J. M. Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, (29):845–866, 1992.
- [3] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. Laplace-beltrami operator and brain flattening. Technical report, University of Minnesota ECE Report, 1998.
- [4] M. Bertalmío, L.T. Cheng, S. Osher, and G. Sapiro. Variational problems and partial differential equations on implicit surfaces. *Journal of Computational Physics*, (174):759–780, December 2001.
- [5] M. Bertalmío, G. Sapiro, and G. Randall. Region tracking on level set methods. *IEEE Transactions on Medical Imaging*, (18):448–451, 1999.
- [6] M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion. *IEEE Transactions on Image Processing*, 7(3):421–432, 1998.
- [7] J. Bloomenthal. *Introduction to Implicit Surfaces*. Morgan Kaufmann Publishers, Inc., 1997.
- [8] H. Brezis, J. M. Coron, and E. H. Lieb. Harmonic maps with defects. *Communications in Mathematical Physics*, (107):649–705, 1986.
- [9] B. Cabral and C. Leedom. Imaging vector fields using line integral convolution. In *ACM Computer Graphics (SIGGRAPH '93)*, number 27 in 4, pages 263–272, 1993.
- [10] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contour. *International Journal of Computer Vision*, 22(1):61–79, 1997.
- [11] T. Chan and J. Shen. Variational restoration of non-flat image features: Models and algorithms. Technical Report 9920, UCLA-CAM, June 1999.
- [12] S. Chen, B. Merriman, S. Osher, and P. Smereka. A simple level set method for solving stefan problem. *Journal of Computational Physics*, (135):8, 1995.
- [13] Y. Chen, M.C. Hong, and N. Hungerbuhler. Heat flow of p-harmonic maps with values into sphere. *Math. Z.*, (205):25–35, 1994.

- [14] Y. G. Chen, Y. Giga, and S. Goto. Uniqueness and existence of viscosity solutions of generalized mean curvature flow equation. *J. Diff. Geom.*, (33), 1991.
- [15] L. T. Cheng. *The Level Set Method Applied to Geometrically Based Motion, Material Science, and Image Processing*. PhD thesis, UCLA, 2000.
- [16] R. Cohen, R. M. Hardt, D. Kinderlehrer, S. Y. Lin, and M. Luskin. Minimum energy configurations for liquid crystals: Computational results. In *Theory and Applications of Liquid Crystals*, pages 99–121. IMA Volumes in Mathematics and its Applications, Springer-Verlag, New York, 1987. J. L. Ericksen and D. Kinderlehrer.
- [17] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Discrete differential-geometry operators in *nd*. Technical report, Multi-res modeling group TR, Caltech, September 2000.
- [18] J. Dorsey and P. Hanrahan. Digital materials and virtual weathering. *Scientific American*, 282(2):46–53, 2000.
- [19] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multi-resolution analysis of arbitrary meshes. In *Computer Graphics (SIGGRAPH '95 Proceedings)*, pages 173–182, 1995.
- [20] M. Eck and H. Hoppe. Automatic reconstruction of b-spline surfaces of arbitrary topological type. In *Computer Graphics 1996*, 1996.
- [21] J. Eells and L. Lemarie. A report on harmonic maps. *Bull. London Math. Soc.*, 10(1):1–68, 1978.
- [22] J. Eells and L. Lemarie. Another report on harmonic maps. *Bull. London Math. Soc.*, 20:385–524, 1988.
- [23] L. C. Evans and J. Spruck. Motion of level sets by mean curvature, i. *J. Differential Geometry*, (33):635–681, 1988.
- [24] O. Faugeras, F. Clément, R. Deriche, R. Keriven, T. Papadopoulos, J. Gomes, G. Hermosillo, P. Kornprobst, D. Lingrad, J. Roberts, T. Viéville, and F. Devernay. The inverse eeg and meg problems: The adjoint state approach i: The continuous case. Technical Report RR-3673, INRIA, June 1999.
- [25] S. F. Frisken, R. N. Perry, A. Rockwood, and T. Jones. Adaptively sampled fields: A general representation of shape for computer graphics. In *Computer Graphics (SIGGRAPH) 2000*, 2000.
- [26] M. Giaquinta, G. Modica, and J. Soucek. Variational problems for maps of bounded variation with values in s^1 . *Cal. Var.*, (1):87–121, 1993.

- [27] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle. Conformal surface parameterization for texture mapping. Technical Report 1611, University of Minnesota, IMA, April 1999.
- [28] R. M. Hardt. Singularities of harmonic maps. *Bulletin of the American Mathematical Society*, 34(1):15–34, 1997.
- [29] G. Huiskamp. ‘difference formulas for the surface laplacian on a triangulated surface. *Journal of Computational Physics*, (95):477–496, 1991.
- [30] T. Ilmanen. Generalized flows of sets by mean curvature on a manifold. *Indiana Journal of Math.*, (4):671–705, 1992.
- [31] H. Ishii. A generalization of bence, merriman, and osher algorithm for motion by mean curvature. In A. Damlamian, J. Spruck, and A. Visintin, editors, *Curvature Flows and Related Topics*, pages 111–127. Gakkōtoshō, 1995.
- [32] R. Kimmel. Intrinsic scale space for images on surfaces: The geodesic curvature flow. *Graphical Models and Image Processing*, (59):365–372, 1997.
- [33] J. J. Koenderink. The structure of images. *Biological Cybernetics*, (50):363–370, 1984.
- [34] V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. *Computer Graphics*, pages 313–324, 1996.
- [35] S. Mauch. Closest point transform. Technical report, Caltech, December 2000. www.ama.caltech.edu/~seanm/software/cpt/cpt.html.
- [36] F. Mémoli, G. Sapiro, and S. Osher. Solving variational problems and partial differential equations mapping into general target manifolds. Technical Report 1827, IMA, January 2002.
- [37] B. Merriman, J. Bence, and S. Osher. Diffusion generated motion by mean curvature. In J. E. Taylor, editor, *Computational Crystal Growers Workshop*, pages 73–83. American Mathematical Society, Providence, Rhode Island, 1992.
- [38] O. Monga, S. Benayoun, and O. Faugeras. From partial derivatives of 3d density images to ridge line. In *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recogniti*, pages 354–359, 1992.
- [39] H. P. Moreton and C. H. Séquin. Functional minimization or fair surface design. In *Computer Graphics (SIGGRAPH)*, pages 167–176, 1992.

- [40] M. Bertalmío, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Computer Graphics (SIGGRAPH)*, July 1992.
- [41] H. Zhao and S. Osher, B. Merriman, and M. Kang. Implicit, non-parametric shape reconstruction from unorganized points using a variational level set method. *Comp. Vision and Image Understanding*, (80):295–314, 2000.
- [42] S. J. Osher and R. P. Fedkiw. Level set methods: An overview and some recent results. *J. Comput. Phys.*, (169):463–502, 2001.
- [43] S. J. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, (79):12–49, 1988.
- [44] D. Peng, B. Merriman, S. Osher, H. Zhao, and M. Kang. A pde-based fast local level set method. *J. Comput. Phys.*, (155):410–438, 1999.
- [45] P. Perona. Orientation diffusion. *IEEE Trans. Image Processing*, (7):457–467, 1998.
- [46] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern. Anal. Machine Intell.*, (12):629–639, 1990.
- [47] E. Praun, A. Finkelstein, and H. Hoppe. Lapped textures. In *ACM Computer Graphics (SIGGRAPH)*, July 2000.
- [48] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, (60):259–268, 1992.
- [49] S. J. Ruuth, B. Merriman, , and S. Osher. Convolution generated motion as a link between cellular automata and continuum pattern dynamics. *J. of Computational Physics*, (151):836–861, 1999.
- [50] G. Sapiro. *Geometric Partial Differential Equations and Image Processing*. Cambridge University Press, New York, 2001.
- [51] L. Simon. *Lectures on Geometric Measure Theory*. Australian National University, 1984.
- [52] N. Sochen, R. Kimmel, and R. Malladi. A general framework for low level vision. *IEEE Trans. Image Processing*, (7):310–318, 1998.
- [53] M. Struwe. On the evolution of harmonic mappings of riemannian surfaces. *Comment. Math. Helvetici*, (60):558–581, 1985.
- [54] B. Tang, G. Sapiro, and V. Caselles. Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case. *Int. Journal Computer Vision*, 36(2):149–161, 2000.

- [55] B. Tang, G. Sapiro, and V. Caselles. Chromaticity diffusion. *IEEE Trans. Image Processing*, (1), May 2001.
- [56] G. Taubin. Estimation of planar curves, surfaces, and nonplanar space curves defined by implicit equations with applications to edge and range image segmentation. *IEEE Trans. PAMI*, 13(11):115–1138, 1991.
- [57] G. Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. In *IEEE International Conference Computer Vision, Boston, MA*, pages 902–907. IEEE, 1995.
- [58] J. N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, (40):1528–1538, 1995.
- [59] A. Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society B*, (237):37–72, 1952.
- [60] G. Turk. Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics*, 25(4):289–298, July 1991.
- [61] andPreufer U. Diwald and M. Rumpf. Anisotropic diffusion in vector field visualization on euclidean domains and surfaces. *IEEE Trans. Visualization and Computer Graphics*, (6):139–149, 2000.
- [62] G. Winkenbach and D. H. Salesin. Rendering parametric surfaces in pen and ink. In *Computer Graphics (SIGGRAPH 96)*, pages 469–476. SIGGRAPH, 1996.
- [63] A. Witkin and P. Heckber. Using particles to sample and control implicit surfaces. In *Computer Graphics (SIGGRAPH)*, pages 269–278. SIGGRAPH, 1994.
- [64] A. Witkin and M. Kass. Reaction-diffusion textures. In *Computer Graphics (SIGGRAPH)*, number 25 in 4, pages 299–308. Computer Graphics (SIGGRAPH), July 1991.
- [65] A. P. Witkin. Scale-space filtering. In *Int. Joint. Conf. Artificial Intelligence*, number 2, pages 1019–1021, 1983.
- [66] G. Yngve and G. Turk. Creating smooth implicit surfaces from polygonal meshes. Technical report, Graphics, Visualization, and Usability Center. Georgia Institute of Technology, 1999. www.cc.gatech.edu/gvu/geometry/publications.html.
- [67] D. Zhang and M. Hebert. Harmonic maps and their applications in surface matching. In *Proc. CVPR '99*. CVPR, 1999.

- [68] H. K. Zhao, T. Chan, B. Merriman, and S. Osher. A variational level set approach to multi-phase motion. *J. of Computational Physics*, (127):179–195, 1996.