Total Variation Geometry I: Concepts and Motivation

Paul Burchard

January 2, 2002

Abstract

Classical differential geometry is the study of smooth manifolds and smooth maps between them. In that setting, the most important geometric invariant of maps—and the primary tool for their regularization—is the *harmonic energy*. In many applications, however, it is essential to broaden the notion of map, to consider maps with discontinuities along lower-dimensional *features* in the domain (edges, creases, shocks, etc.).

The nascent subject of *total variation geometry* addresses this need by replacing harmonic energy with the most interesting geometric invariant that is neutral to discontinuities of codimension r in the map; this *geometric total variation energy* is the integral of the r-th elementary symmetric function of the singular values of the derivative of the map. The resulting regularization, which replaces the usual heat flow, is unique in that it does not destroy discontinuities along features, but instead regularizes their shape through mean-curvature motion—including proper handling of junctions between features.

Total variation geometry proposes the most natural extension of previous work on total variation to cover vector-valued functions, features with junctions, features of higher codimension, and arbitrary domain and range geometry. As a new research program, this subject offers many open problems in analysis, geometry, numerics, and applications.

1 Introduction.

A prominent feature of many applications is the existence of *features*: lower-dimensional sets in a domain, along which the functions or fields involved have discontinuities or very thin boundary layers. These features are not static, typically, but instead display interesting dynamics influenced by their geometry. This dynamical behavior may, moreover, involve interesting interactions between different features when they meet at *junctions*. Examples of features include edges in images and video; shocks, vortices, and phase boundaries in fluid flow; defects in liquid crystals; creases and corners on object surfaces; and so on.

Alternatively, in some applications, one begins with a description of the dynamics of certain lower-dimensional geometric features, and wishes to create a more convenient numerical representation in which the feature is implicitly represented as the discontinuity (or level set) of an auxiliary function, in the spirit of the level set method [O-S]. Of course, such feature dynamics are typically derived as an approximation to the dynamics of discontinuities or boundary layers of a more complete field model; but as long as the dynamics are known, it is not necessary for the fields used to represent geometric features to have any additional physical significance.

In either case, the nascent subject of *total variation geometry* aims to provide the necessary enhancements to analysis, geometry, and numerics to handle functions with meaningful discontinuities of various types. In all of these areas of enhancement, *regularization* (smoothing) is a key tool, providing manageable approximations for convergence proofs, a priori biases to make inverse problems solvable, and stabilization for numerical methods. Now, however, the regularization cannot be permitted to smooth away the meaningful discontinuities; instead, it should only geometrically smooth their shapes. Our main goal in this paper is to identify, and initiate the study of, the required regularizations.

In classical differential geometry, which considers only smooth maps, the appropriate regularization is of course the well-known *heat flow*. This flow results from dissipating as rapidly as possible the *harmonic energy* of the map (the L^2 norm of its derivative). Our approach to identifying the correct regularizations will be to explore a broader spectrum of geometric energies, and the properties of the *geometric motions* that result from their dissipation.

The singular case that has already received extensive attention is that of scalar functions, whose discontinuities (or level sets) are restricted to representing features of codimension one without junctions. The correct geometric energy in this case is known to be the *total variation* [R-O-F], whose corresponding regularization is the *mean-curvature motion* of the features [E-S]. Most of the existing work is based on the use of level sets [O-S] rather than actual discontinuities as features; indeed, the primary analytic theory used to justify the mean-curvature regularization requires continuous functions [C-I-L]. Recently, however, approaches to the analysis of the fully discontinuous case have emerged [Yi].

Total variation geometry proposes the most natural generalization of this previous work on feature regularization to cover

- features with junctions (requires vector-valued functions for generic representation);
- features of higher codimension (again, requires vector-valued functions for generic representation); and
- arbitrary domain and range geometry.

The unique choice of energy that meets these requirements is the *geometric total variation energy* introduced here, namely, the integral of the r-th elementary symmetric function of the singular values of the derivative of the map.

In the case of codimension one, for example, the geometric total variation energy may be written

$$E_{TV}^{(1)}[u] = \int_X \operatorname{tr} |\nabla u| \, d\mathrm{vol}_X$$

for a map $X \xrightarrow{u} Y$ between manifolds (where $|A| := \sqrt{AA^*}$). The corresponding replacement for the heat flow would be the geometric motion

$$u_t = |\nabla u| \; \nabla \cdot \left(|\nabla u|^{-1} \nabla u \right),$$

a degenerate-elliptic system of nonlinear partial differential equations generalizing the implicit mean-curvature flow.

As a new research program, total variation geometry offers many interesting open problems in analysis, numerics, and applications. Implicit motion of features with junctions has been studied previously [M-B-O, Z-C-M-O], but using an explicit region decomposition, as opposed to the vector-valued approach proposed here which is more general. Some intial work on higher codimension [B-C-M-O] and general domain geometry [B-C-O-S] has already appeared, but much more remains to be done.

2 Geometric Motions.

Definitions. A *geometric evolution* is an evolution process for maps $X \xrightarrow{u} Y$ between manifolds that is both local and an invariant of the geometry of the manifolds. For smooth maps, a local process is essentially a differential equation, but since we are specifically interested in discontinuities, we must allow for more general processes (e.g., a differential equation plus an entropy rule). A *geometric motion* is a geometric evolution that is *advectional*—i.e., one that can be achieved, in the weak sense, by transporting Y-values of the map along some locally-determined flow in the domain X.

We will search for the correct regularizations for total variation geometry among *variational geometric motions of second order*. Such a motion is, by definition, the advectionally-constrained steepest descent of a first-order geometric functional

$$E[u] := \int_X e(\nabla u)\, d\mathrm{vol}_X. \tag{1}$$

(Note to differential geometers: the "gradient" ∇ simply denotes the ordinary derivative D, that is, the tangent functor $TX \xrightarrow{Tu} TY$. Similarly, the "divergence" $\nabla\cdot$ should be understood as referring to the adjoint D^* of the derivative with respect to the metrics on X and Y.)

Let us explain the components of preceding definition:

According to a standard result of linear algebra, the functional E will be *geometric* iff $e(\nabla u)$ is a symmetric function of the singular values of the derivative ∇u of u. Recall that the *singular values* of a rectangular matrix A are the eigenvalues of the self-adjoint square matrix $|A| := \sqrt{AA^*}$.

To say that the descent is *advectionally-constrained* means that we only allow variations δu of the map that can be written in the form $\delta u = -\nabla u \cdot v$ for some vector field v on X.

Finally, the notion of *steepest descent* depends on how we measure the rate of change of $E[u]$ with respect to u; that is, on how we measure the size of variations δu of u. In order to produce a geometric motion as the steepest descent of the geometric functional $E[u]$, it turns out that the size of (and angle between) variations should be measured using the *bilinearization* of the density e as the inner product. More precisely, for two advectional variations $\delta_1 u$ and $\delta_2 u$ of a map u, their inner product should be measured as

$$\langle \delta_1 u,\, \delta_2 u\rangle_u = \int_X (\delta_1 u)^*\, Q\, (\delta_2 u)\, d\mathrm{vol}_X, \tag{2}$$

where Q bilinearizes the density e via the formula

$$\langle A,\, Q\, A\rangle = \chi e(A), \tag{3}$$

with χ denoting the homogeneity operator. (These geometric aspects of variational calculus are more fully explained in Section 4.)

Diffusion and Advection Forms. The resulting variational geometric motion, corresponding to a geometric functional E, can be written as a differential equation in two interesting forms.

The *diffusion form* of the geometric motion is

$$u_t = Q^{-1}\, \nabla \cdot (Q\, \nabla u). \tag{4}$$

This form shows that the evolution can be viewed as the standard heat flow $u_t = \nabla \cdot (\nabla u)$ twisted by the (possibly degenerate) quadratic form Q. Because of the possible degeneracy of Q, the Q^{-1} must be interpreted advectionally, that is, zero eigenvalues (directions not in the range of ∇u) are not inverted. The *advection form* of the geometric motion is

$$u_t + \nabla u\, \left(-{e'}^{-1}\, \nabla \cdot e' \right) = 0. \tag{5}$$

Here $-{e'}^{-1}\, \nabla \cdot e'$ is a (possibly singular) flow field on X that we interpret as advecting the values of u. The variational calculations needed to derive the above forms (4) and (5) of the geometric motion are deferred to Section 4.

The ability to use both forms for the motion shows these are quite special differential equations in the general vector-valued case ($\dim Y > 1$).

Example 1: Heat Flow. For scalar functions $X \xrightarrow{u} \mathbb{R}$, the *harmonic energy* is

$$E_{\mathrm{harm}}[u] = \int_X |\nabla u|^2 \, d\mathrm{vol}_X.$$

Since e is already bilinear, the inner product between variations is just the usual L^2 inner product, independent of u:

$$\langle \delta_1 u, \delta_2 u \rangle_u = \int_X \delta_1 u^* \, \delta_2 u \, d\mathrm{vol}_X.$$

For scalar functions, the advectional constraint has no generic effect. The resulting steepest descent is therefore the usual heat flow

$$u_t = \nabla \cdot (\nabla u) = \Delta u.$$

Example 2: Scalar TV. For scalar functions $X \xrightarrow{u} \mathbb{R}$, the *total variation* energy is the total area of all of its level sets:

$$E_{TV}[u] = \int_X |\nabla u| \, d\mathrm{vol}_X.$$

To get a geometric motion, E_{TV} must be considered as functional on the nonlinear *space of level set functions*, whose geometry is defined by an L^2 inner product, not on the variations δu themselves, but on the resulting *normal motions* $-\frac{\delta u}{|\nabla u|}$ of the level sets of u, integrated up over each level set, and then over the set of levels:

$$\langle \delta_1 u, \, \delta_2 u \rangle_u = \int_X \left(-\frac{\delta_1 u}{|\nabla u|} \right) \left(-\frac{\delta_2 u}{|\nabla u|} \right) |\nabla u| \, d\mathrm{vol}_X.$$

It is easy to check that this geometric definition of the inner product in terms of level sets coincides with the bilinearization of the total variation energy density, as prescribed by our general procedure.

The resulting steepest descent is the well-known *mean curvature flow* of the level sets of u:

$$u_t = |\nabla u| \; \nabla \cdot \left(\frac{\nabla u}{|\nabla u|} \right).$$

3 The Geometric Total Variation Energy.

Our goal is to show that the unique variational geometric motion of second order suitable for regularizing features of codimension r, including their junctions, is the one based on taking the energy density e to be the r-th elementary symmetric function of the singular values of ∇u, which we call the *geometric total variation* energy density of degree r.

Codimension and Homogeneity. A simple scaling argument shows that the energy density e must be homogeneous of degree r in order to be neutral to smooth features of codimension r in the domain.

Consider the scaling $s_\varepsilon(x_1, \ldots, x_n) = (\frac{x_1}{\varepsilon}, \ldots, \frac{x_r}{\varepsilon}, x_{r+1}, \ldots, x_n)$ on $\mathbb{R}^n$. For a smooth function u on $\mathbb{R}^n$, this scaling causes $u \circ s_\varepsilon$ to acquire a singularity along the submanifold $\{0\} \times \mathbb{R}^{n-r}$ of codimension r as $\varepsilon \to 0$, and conversely, to be smoothed out normal to that submanifold as $\varepsilon \to \infty$. If

e is homogeneous of degree h in ∇u, and u depends only on the variables $x_1, \dots, x_r$ normal to the submanifold, then

$$E[u \circ s_\varepsilon] = \varepsilon^{r-h} E[u].$$

Thus, it is energetically favorable for the flow

- to instantly smooth discontinuities of codimension $r < h$;
- to be neutral (to highest order) to discontinuities of codimension $r = h$;
- to spontaneously create discontinuities of codimension $r > h$;

and we must have $h = r$ to obtain the desired behavior.

Junctions and Convexity. Vector-valued functions can robustly represent, by means of their codimension-r levelest sets, not merely topologically-closed features of codimension r in the domain, but also junctions between up to $\binom{n+1}{n-d+1}$ of them, where $n = \min(\dim X, \dim Y)$ is the number of singular values, and $d = \dim X - r$ is the dimension of the features (see Figure 1). Indeed, a generic vector-valued function will have critical levelest sets of codimension r that meet at such junctions, and so if a geometric motion were to smear out such junctions, it would also eventually smear out all the features that meet there.

In the scalar case, where there is only one nonzero singular value, and the codimension can only be $r = 1$, the homogeneity requirement is already sufficient to determine a unique geometric energy density, namely the total variation density $e(\nabla u) = |\nabla u|$.

In the vector-valued case, homogeneity (or in geometric terms, codimension) is generally insufficient to determine the energy. The extra algebraic freedom to choose a symmetric homogeneous function of the singular values corresponds to the extra geometric freedom to choose what happens to the junctions between features of the given codimension. The geometric goal should obviously be maximum preservation of these junctions, and since convexity in the energy corresponds precisely to dissipation of geometric information in the corresponding steepest descent, the corresponding algebraic goal would be to find the *minimally convex* geometric energy density of the correct homogeneity.

(To be more technically correct, in the vector setting, the algebraic goal must be to find the minimally *quasiconvex* geometric energy density of the

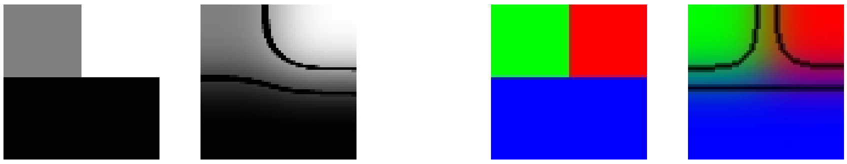

Figure 1: Unlike scalar functions, vector-valued functions can robustly represent features that meet in junctions. In the first image, edge discontinuities of a scalar function appear to meet in a triple junction. The second image, however, with sample level sets marked, shows that a small perturbation (here, blurring) can destroy the apparent junction; there are really only two edges, which had partially merged in the first image to form the apparent third leg of the junction. In the third image, edge discontinuities of a vector-valued function meet in a true triple junction. As suggested by the fourth image, with sample levelest sets marked, this triple point cannot be removed by small perturbations.

correct homogeneity. Quasiconvexity is the natural replacement for convexity when dealing with functions of matrices like ∇u.)

The unique choice selected by this criterion is the *geometric total variation* energy density, namely, the elementary symmetric function of degree r in the singular values of the derivative ∇u. Figure 2 visually demonstrates the correctness of this choice in the case of codimension one, where the geometric total variation energy density is the *absolute trace* energy—the linear sum of the singular values:

$$e^{(1)}_{TV}(\nabla u) := \operatorname{tr} |\nabla u| = \sigma_1 + \cdots + \sigma_n. \tag{6}$$

(As before, $|A| := \sqrt{AA^*}$ denotes the matrix whose eigenvalues are the singular values of the rectangular matrix A.) For general codimension r, we can write the elementary symmetric function of singular values as

$$e^{(r)}_{TV}(\nabla u) := \sum_{i_1 < \cdots < i_r} \sigma_{i_1} \cdots \sigma_{i_r}. \tag{7}$$

Sketch of Proof. The correspondence cited above between preservation of junctions and the minimal quasiconvexity of the geometric energy can be explained in more detail by looking at how averaging of matrices effects their singular values. (For simplicity, we will consider only the case of homogeneous degree one.)

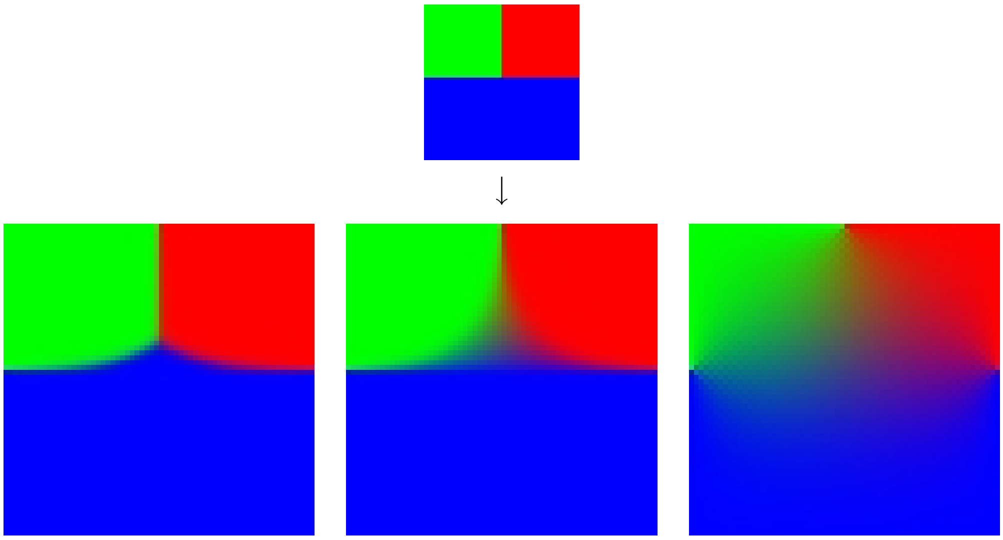

Figure 2: Energy density must be minimally quasiconvex to avoid smearing features and junctions, and produce only geometric regularization. Correct homogeneity only avoids smearing features without junctions; incorrect homogeneity smears both features and junctions.

Let e be an energy density which is convex and homogeneous of degree one; then for any nice bump function ϕ (smooth, non-negative, unit mass), we know that $E[\phi * u] \leqslant E[u]$ because $\nabla(\phi * u) = \phi * \nabla u$. That is, smearing by ϕ is always either neutral or favorable. We want to investigate when equality—neutrality to smearing—is possible.

It will be sufficient to consider piecewise constant functions u with polygonal boundaries composed of faces F_i. For such functions,

$$\nabla u = \sum_i \tau_i \, \delta_{F_i}(x),$$

where τ_i is the rank-one matrix formed as the tensor product of the difference of colors across face F_i with the conormal to F_i. Now,

$$E[\phi * u] \quad = \quad \int_X e\left(\sum_i \tau_i(\phi * \delta_{F_i})\right) \, \mathrm{dvol}_X$$

$$\leqslant \sum_i e(\tau_i) \int_X (\phi * \delta_{F_i})\, d\mathrm{vol}_X$$
$$= \sum_i e(\tau_i)\, \mathrm{area}(F_i),$$

with equality possible in general only if e is linear on the positive cone spanned by the τ_i. Thus, any kind of strict convexity makes smearing energetically favorable (vs. doing nothing), and therefore leads to *some* kind of smoothing (though the smearing by some particular ϕ, or even any ϕ, may not be the most energetically-favorable form of the smoothing).

As a first example, consider a straight edge between two colors. There, we always get neutrality, because any homogeneous degree-one density is linear on the positive cone spanned by a single matrix τ.

Next, consider a curved edge between two colors, which we will simplify to a kinked edge consisting of two polygonal faces:

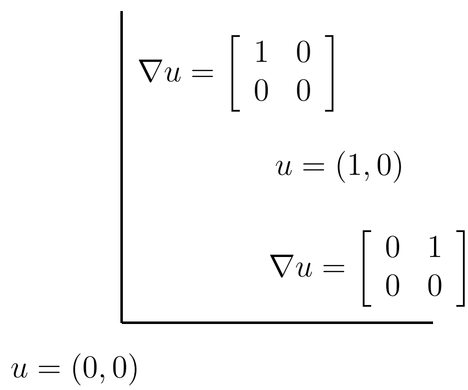

For this example, we have two matrices, both with singular values $(1, 0)$, but whose average has singular values $(\frac{1}{\sqrt{2}}, 0)$. Thus, no matter which homogeneous degree-one e is chosen, it must have strictly convex values on the positive cone spanned by these two matrices (whose values, up to constant, start at 1, dip to $\frac{1}{\sqrt{2}}$, and return back to 1). Thus, any homogeneous degree-one e will smooth curved edges (typically by curvature motion—not by smearing).

Finally, the most interesting example: a straight, symmetric triple junction separating equidistant colors. The faces and matrices τ_i are shown below:

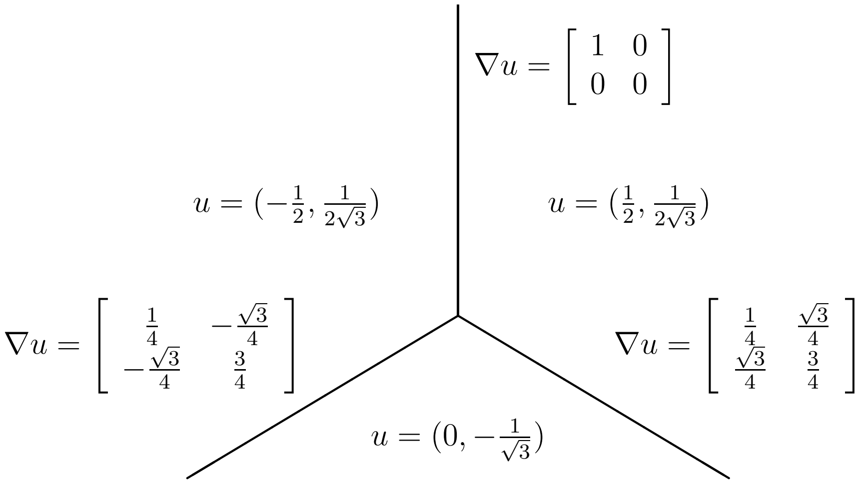

Here we finally see the distinction between different candidates for the geometric energy: it is only the absolute trace $e^{(1)}_{TV}$ that is linear on the positive cone spanned by these three matrices, each of which has singular values $(1, 0)$. In particular, note that the average of all three matrices is $\frac{1}{2}I$, with singular values $(\frac{1}{2}, \frac{1}{2})$.

4 Variational Calculus for Geometric Motions.

Motivation. Variational formulations of dynamic processes have many advantages. Analytically, the resulting a priori energy bounds can be an important part of proving existence and uniqueness of solutions. Practically, an energy formulation often lends itself to more obvious generalization to new settings and applications.

Nevertheless, an energy functional does not, by itself, determine any dynamics. To make sense of "steepest descent" in energy values, we need some way to measure rates of change of energy with respect to variations in the function. That is, we must impose some kind of *geometry* on the space of functions—a way to measure the size of infinitesimal variations within the function space.

Consider a simple example from finite-dimensional calculus: the quadratic energy function $E(x, y) = x^2 + y^2$ on the plane. Using the ordinary Euclidean length measurement $ds^2 = dx^2 + dy^2$, the direction at (x, y) in which E decreases most rapidly is just the inward radial direction $(-x, -y)$. But if instead, we use the line element $ds^2 = 4\,dx^2 + dy^2$ to measure distances in the plane (and thus rates of change of E), then the direction at (x, y) in which E decreases most rapidly becomes $(-x, -4y)$.

These geometrical aspects of descent are, in fact, at the heart of the well-known *conjugate gradients* algorithm for minimizing quadratic functions $f(x) = \frac{1}{2}x^*Ax - b^*x$ on a vector space (or equivalently, solving symmetric linear systems $Ax = b$). Ideally, we could perform direct "radial" descent toward the solution, which is nothing but steepest descent in the metric $ds_A^2 = dx^*A\,dx$ adapted to the system. Unfortunately, we do not have sufficient information about ds_A^2 to immediately determine its steepest descent direction; instead, we must start out with steepest descent in the (already diagonalized) Euclidean metric $ds^2 = dx^*dx$. As the algorithm proceeds, however, it builds up an increasing amount of knowledge of the metric ds_A^2, and its descent makes a transition from being steepest for the original metric ds^2 to being steepest for the new metric ds_A^2—thereby giving the conjugate gradients algorithm its well-known efficiency.

Variational Formalism. Variational calculus is concerned with locating critical points (particularly minima or maxima) of a real-valued energy function E on a (possibly infinite-dimensional) manifold M. Assuming this function is sufficiently smooth at $x \in M$, it has a natural differential $dE(x)$ there, which for any tangent vector v at x gives the corresponding directional derivative $dE(x)(v)$ of E at x.

To get a flow, we need a vector field; but since dE is dual to the vectors at each point, it is actually a field of *covectors* on M. Converting covectors to vectors requires a Riemannian metric on the manifold—that is, an inner product $\langle \cdot, \cdot \rangle$ on each tangent space. Given such a metric, the natural one-to-one correspondence $v \mapsto \langle v, \cdot \rangle$ between vectors and covectors always corresponds a vector to a covector for which it is a *steepest ascent* vector (that is, if $\lambda(\cdot) := \langle v, \cdot \rangle$ and $\langle w, w \rangle = \langle v, v \rangle$, then $\lambda(w) \leqslant \lambda(v)$).

For our variational problem, we would like to convert the covector field $dE(x)$ into a vector field, and then reverse it to obtain a steepest descent vector field. The resulting vector field $\dot{x}$ is related to dE via

$$dE(x)(\cdot) = -\langle \dot{x}, \cdot \rangle . \tag{8}$$

In the infinite-dimensional setting of energy functionals, the inner product will usually take the form of an integral. If our manifold is now an infinite-dimensional space of functions on some finite-dimensional manifold X, then the inner product between variations at u typically takes the form

$$\langle \delta_1 u, \delta_2 u \rangle_u = \int_X \delta_1 u \, Q(x, u, u_x, \ldots) \, \delta_2 u \, dx.$$

The calculation to write dE in the dual form (8) typically involves integration by parts. However, the measure in the integral for the inner product might involve geometric factors, so care must be taken to perform the integration by parts correctly.

Example: Curve Shortening. As a warm-up, we perform curve-shortening. Here the geometric constraints are more obvious.

The function space is the space of parameterized curves $[0,1] \xrightarrow{u} Y$ in the manifold Y (obviously a nonlinear function space when Y is nonlinear). The tangent space of this function space at a particular curve u is the space of vector fields on Y along u. Our energy is the curve length,

$$E(u) = \int_{[0,1]} |u_x| \, dx$$

(which is the Total Variation energy).

The (invariantly defined) differential of the energy is

$$dE(v) = -\int_{[0,1]} \left(\frac{u_x}{|u_x|} \right)_x v \, dx = -\left\langle \left(\frac{u_x}{|u_x|} \right)_x, v \right\rangle_{L^2(dx)}$$

Thus, the naive variational calculation, performed implicitly using the coordinate-dependent inner product $L^2(dx)$ on the tangent space at u of our function space, would conclude the flow to be

$$u_t = \left(\frac{u_x}{|u_x|} \right)_x$$

But this flow has no invariant geometric meaning—it depends on the parameterization, not just the curve itself (the right-hand side is the derivative of the unit tangent vector of the curve with respect to the (artificial) parameter x).

As we know, the correct curve-shortening flow, which moves the curve according to the curvature times the unit normal, can be expressed as the derivative of the unit tangent vector with respect to the (natural) arclength parameter ds.

To obtain the correct result variationally, the only thing we have to do is to use the natural inner product on each tangent space: $L^2(ds) = L^2(|u_x|dx)$. Note that this inner product depends on the current position u in the function space – just like the inner product in arbitrary coordinates for any nonlinear manifold.

$$dE(v) = -\int_{[0,1]} \frac{1}{|u_x|}\left(\frac{u_x}{|u_x|}\right)_x v\,(|u_x|\,dx) = -\left\langle \frac{1}{|u_x|}\left(\frac{u_x}{|u_x|}\right)_x, v\right\rangle_{L^2(ds)}$$

so that

$$u_t = \frac{1}{|u_x|}\left(\frac{u_x}{|u_x|}\right)_x$$

as desired.

General Geometric Steepest Descent. Consider a geometric energy

$$E = \int_X e(\nabla u)\,d\mathrm{vol}_X, \tag{9}$$

for maps $X \xrightarrow{u} Y$, that is, an energy for which the energy density e is solely a symmetric function of the singular values of its matrix argument.

Although the energy is a geometric invariant, in order to obtain a geometric motion by constrained steepest descent on this energy, it is essential to work on a nonlinear function space which carries a geometrically natural metric, so that the size of variations can be correctly measured. In the general case, the most natural function space metric on advectional variations is

$$\langle \delta_1 u, \delta_2 u\rangle_u = \int_X \delta_1 u^*\, Q\, \delta_2 u\, d\mathrm{vol}_X \tag{10}$$

where the quadratic form Q is expressed in spectral terms as

$$Q := \sum_i \frac{e_{\sigma_i}}{\sigma_i}\,\widehat{u}_i \otimes \widehat{u}_i, \tag{11}$$

and where the singular value decomposition of ∇u has been written as

$$\nabla u = \sum_i \sigma_i\,\widehat{u}_i \otimes \widehat{n}_i \tag{12}$$

($\widehat{n}_i$ is the normalized source singular vector, and $\widehat{u}_i$ the normalized target singular vector, corresponding to the singular value σ_i).

One way to see the naturality of the advectional form Q is to note that, when e is homogeneous in the singular values, Q is just a bilinearization of e. Specifically,

$$\langle \nabla u,\, Q\,\nabla u\rangle = \chi e$$

where $\chi := \sum_i \sigma_i \frac{\partial}{\partial \sigma_i}$ is the homogeneity operator. Even if e is not homogeneous, we will see that the corresponding bilinearized (but non-geometric) energy density $f(\nabla w) := \langle \nabla w,\, Q\,\nabla w\rangle$ can be used to define the gradient flow.

We will show that the variation of any geometric matrix function e can be expressed in the form

$$\delta\,[e(A)] = \langle e'(A),\, \delta A\rangle, \tag{13}$$

where the matrix-valued derivative $e'(A)$ is computed in spectral terms as

$$e'(A) := \sum_i e_{\sigma_i}\, \widehat{u}_i \otimes \widehat{n}_i, \tag{14}$$

and the angle brackets denote the inner product on matrices

$$\langle A, B\rangle := \mathrm{tr}(A^* B)$$

defined by the source and target metrics.

Assuming the variational formula (13), we may integrate by parts and use the identity $e'(\nabla u) = Q\,\nabla u$ to obtain

$$\begin{aligned} -\delta E &= -\int_X \langle e'(\nabla u),\, d\delta u\rangle\, d\mathrm{vol}_X \\ &= \int_X \nabla\cdot(e'(\nabla u))\cdot \delta u\, d\mathrm{vol}_X \\ &= \int_X \big(Q^{-1}\,\nabla\cdot(Q\,\nabla u)\big)\cdot Q\,\delta u\, d\mathrm{vol}_X. \end{aligned}$$

(where $\nabla\cdot()$ is more generally the divergence operator $d^* = *d*$ in the metric on X, and in the constrained case, Q^{-1} first projects into the constraint).

Thus, the constrained gradient flow is a diffusion equation

$$u_t = Q^{-1}\, d^*(Q\,\nabla u) = \Delta_{X,Q}(u), \tag{15}$$

where $\Delta_{X,Q}$ is the Laplace operator corresponding to the given metric on X and the u-dependent metric Q on Y, or more exactly, on the trivial Y-bundle over X. In other words, this flow is instantaneously the same as the gradient flow of the quadratic—but in the original metrics, non-geometric—energy energy density $f(\nabla w) = \langle \nabla w,\, Q\,\nabla w\rangle$, with respect to the L^2 function-space metric corresponding to the distorted metric on Y. (In the scalar TV case, the factor Q becomes the $\frac{1}{|\nabla u|}$ factor that turns value variations into normal velocities of level sets.)

Alternatively, it is apparent that this constrained gradient flow can be rewritten as an advection equation

$$u_t + \nabla u \left(-e'^{-1}\, d^* e'\right) = 0. \tag{16}$$

Now to check the variational formula (13). For any singular value σ of A, we have the identity

$$0 = \det\left(A^* A - \sigma^2 I\right).$$

Since the determinant is multilinear in the rows of a matrix, the variation of this identity gives

$$0 = \sum_k \det(\delta_k A)$$

where

$$(\delta_k A)_{ij} := \begin{cases} A_i \cdot A_j - \sigma^2 \delta_{ij} & \text{if } i \neq k, \\ \delta A_k \cdot A_j + A_k \cdot \delta A_j - 2\sigma\, \delta\sigma\, \delta_{kj} & \text{if } i = k, \end{cases}$$

and A_i are the columns of A. Expanding the $\det(\delta_k A)$ by minors and solving for $\delta\sigma$ yields

$$\delta\sigma = \sum_{k,l} \left(\frac{(-1)^{k+l}\, \mathrm{minor}_{kl}\left(A^* A - \sigma^2 I\right)}{\sum_m \mathrm{minor}_{mm}\left(A^* A - \sigma^2 I\right)} \right) \frac{A_k}{\sigma} \cdot \delta A_l.$$

The parenthesized fraction of determinants in the foregoing can be interpreted in an invariant manner. After isometric coordinate change at the point of interest, which makes the first coordinate the direction of the source singular vector $\widehat{n}$ corresponding to the singular value σ, then $(1, 0, \ldots, 0)$ is in both the kernel and cokernel of the self-adjoint matrix $(A^* A - \sigma^2 I)$, whose first row and first column and therefore both zero. In particular, in these coordinates, the only non-zero minor is the $(0, 0)$ minor, making the parenthesized factor in question equal to $\delta_{k0}\,\delta_{l0}$, or in invariant terms, the projection operator $\widehat{n}\otimes\widehat{n}$ to the span of the source singular vector $\widehat{n}$. Thus,

$$\delta\sigma = \sum_{k,l} (\widehat{n}\otimes\widehat{n})_{kl}\,\frac{A_k}{\sigma}\cdot\delta A_l = \langle \widehat{u}\otimes\widehat{n},\, \delta A\rangle.$$

Finally, this gives

$$\delta\,[e(A)] = \sum_i e_{\sigma_i}\,\delta\sigma_i = \sum_i e_{\sigma_i}\,\langle \widehat{u}_i\otimes\widehat{n}_i,\, \delta A\rangle = \langle e'(A),\, \delta A\rangle$$

as claimed.

References

[B-C-M-O] Burchard, P., Cheng, L.-T., Merriman, B., Osher, S.: Motion of curves in three spatial dimensions using a level set. J. Comp. Phys. **170**(2), 720–741 (2001)

[B-C-O-S] Bertalmio, M., Cheng, L.-T., Osher, S., Sapiro, G.: Variational problems and partial differential equations on implicit surfaces. J. Comp. Phys. **174**, 759–780 (2001)

[C-I-L] Crandall, M. G., Ishii, H., Lions, P.-L.: User's guide to viscosity solutions of second order partial differential equations. Bull. Amer. Math. Soc. **27**, 1–67 (1992)

[E-S] Evans, L. C., Spruck, J.: Motion of level set via mean curvature. J. Diff. Geom. **33**, 635–681 (1991)

[M-B-O] Merriman, B., Bence, J., Osher, S.: Motion of multiple junctions, a level set approach. J. Comp. Phys. **112**, 334–363 (1994)

[O-S] Osher, S., Sethian, J.: Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. J. Comp. Phys. **79**, 12–49 (1988)

[R-O-F] Rudin, L. I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. Physica D **60**, 259–268 (1992)

[Yi] Yip, A.: Stochastic motion by mean curvature. Archive for Rational Mechanics and Analysis (to appear)

[Z-C-M-O] Zhao, H.-K., Chan, T., Merriman, B., Osher, S.: A variational level set approach to multiphase motion. J. Comp. Phys. **127**, 179–195 (1996)