Mathematisches Forschungsinstitut Oberwolfach

Report No. 07/2011

DOI: 10.4171/OWR/2011/07

Trends in Mathematical Imaging and Surface Processing

Organised by
Vicent Caselles, Barcelona
Martin Rumpf, Bonn
Guillermo Sapiro, Minneapolis
Peter Schröder, Pasadena

January 30th – February 5th, 2011

Abstract. Motivated both by industrial applications and the challenge of new problems, one observes an increasing interest in the field of image and surface processing over the last years. It has become clear that even though the applications areas differ significantly the methodological overlap is enormous. Even if contributions to the field come from almost any discipline in mathematics, a major role is played by partial differential equations and in particular by geometric and variational modeling and by their numerical counterparts. The aim of the workshop was to gather a group of leading experts coming from mathematics, engineering and computer graphics to cover the main developments.

Mathematics Subject Classification (2000): 35Q80, 49Q10, 49Q20, 65D18, 65K10, 65M60, 65M32, 68U10, 68U05.

Introduction by the Organisers

Variational methods and partial differential equations form a core of various methodologies in image analysis, computer vision and surface processing. *Geometric concepts* naturally appear in the modeling, processing or animation of surfaces for instance in computer graphics. Furthermore, the space of shapes - either explicitly given as parametric surfaces or implicitly encoded in 2D and 3D images - has an interesting geometric structure, the understanding of which is crucial for many applications. This workshop brought together researchers from computer vision as well as computer graphics and mathematicians with expertise in the calculus of variations, in PDE analysis and numerics, in discrete geometry and sparse approximation. The aim of the workshop was to dovetail the strengths of geometry,

analysis and numerics in order to get insight into new models.

Different from the previous workshop in 2007 with a particular focus on image decomposition, different types of surface representation, feature preserving methods, and the proper discretization of associated PDE models, this time the perspective was more global. In fact, instead on the processing of single geometries or shapes the focus was more on the geometry of the space of shapes, or instead of continuous energy descent methods, relaxation techniques and associated global optimization approaches have been presented. Some areas of particular relevance in the workshop have been

- *Shape space analysis*, with the rigorous definition of shape space as a Riemannian manifold, the efficient computation of the geodesic distance, the Fréchet or Karcher mean, the principal component analysis of shapes, a suitable embedding into the context of shape statistics, and different types of underlying distances such as the Wasserstein distance, a distanced based on viscous dissipation along motion paths, or the Gromov Hausdorff distance,

- *Global optimization approaches* in image labeling or segmentation, image impainting, and motion estimation, where convex relaxation is applied in the context of TV and L^1 minimization problems and the appropriate use of primal and dual optimization strategies leads to highly efficient algorithms,

- *Concepts from discrete geometry* and the modeling of discrete free form surfaces in architecture, where methods from discrete exterior calculus lead to effective discretization and fabrication restrictions on glass roofs pose challenging design problems and give new impulses to discrete geometry and integrable systems.

Let us finally mention that in monday afternoon we had a special get–to–know session with 3 minute statements of almost all participants on open problems, new methodologies, and demanding applications prepared the ground for intensive discussion between the participants from diverse disciplines over the whole week.

Workshop: Trends in Mathematical Imaging and Surface Processing

Table of Contents

Abstracts

Robust Geometry Processing

Pierre Alliez

Surface reconstruction from raw geometric data has received increasing attention due to the ever broadening range of geometric sensors and vision algorithms that provide little to no reliable attributes. A common approach to this problem involves filtering out outliers and inferring attributes before resorting to a reconstruction method. Typically, the data is first oriented (normals are computed) or, equivalently, signed (an inside/outside function is constructed based on the pointset). However, outlier removal often requires an interactive adjustment of parameters. Similarly, finding and orienting normals from raw geometric data can be as hard as reconstructing the whole surface itself: while there are several options to reliably estimate normal directions, robust normal orientation is considerable harder as recently reminded in Huang et al. [1]. In parallel to the reconstruction literature, the design of approximate unsigned distance functions which are robust to noise and outliers has recently made significant advances [2]. However, it has not yet benefited reconstruction as the resulting robust function does not lend itself to reliable surface reconstruction: contouring of an unsigned distance is unreliable, creating numerous geometric and topological artifacts. We propose to sign the unsigned distance in order to obtain an implicit function suitable for contouring. We leverage the fact that signing a function, rather than the data, can be made more robust by exploiting the property that signing a distance function makes the function smoother.

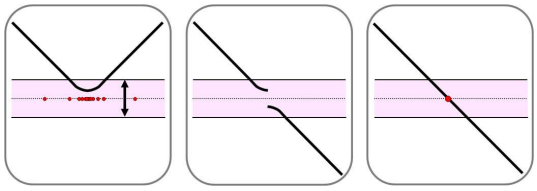

Figure 1. Reconstruction pipeline in 1D. Stage 1 (left): We construct an unsigned distance function (robust to noise and outliers, shown in red), and a threshold for the width of the ϵ-band (in pink). Stage 2 (center): We estimate the sign (± 1) of the function, along with a confidence of each estimate. Stage 3 (right): We construct the final signed function through smoothing, taking into account the unsigned function, estimated sign, and confidence.

Overview. Our main contribution is a practical method for efficient reconstruction of closed surfaces from pointsets that are potentially noisy, outlier-ridden, and undersampled [3]. Its simplicity and modularity facilitate customization and extensibility. Starting with a pointset possibly containing both noise and outliers,

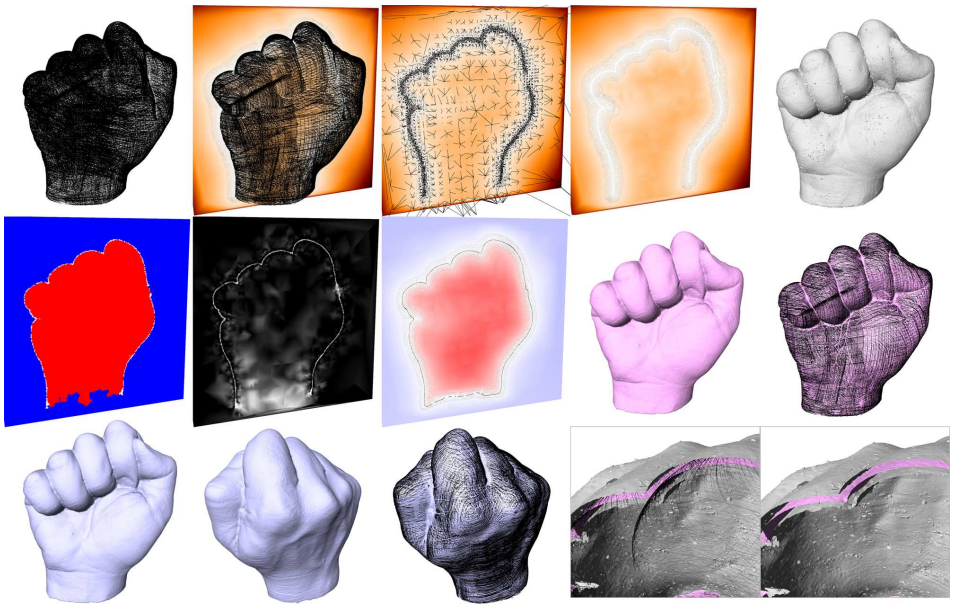

FIGURE 2. **Plaster Hand.** Data scanned with a Kreon laser scanner mounted on an articulated arm; the 18M point sampling is very anisotropic as it was obtained by manual sweeping of a 1D contact sensor. Top: input point set (with a big hole at the bottom and others due to occlusions between the fingers), point set and 2D cut of unsigned function, same 2D cut with nearby edges of the coarse mesh, same 2D cut alone, and full ϵ-band. Middle: 2D cuts of sign guess (red for inside, blue for outside and white uncertain), confidence (which decreases in the holes), signed function after smoothing, isosurface of the robust unsigned function obtained by marching tetrahedra in the lattice mesh, and same isosurface superimposed with input points. Bottom: views of the reconstructed surface obtained by Delaunay refinement without and with points added, and cut view of the ϵ-band with the reconstructed isosurface of the signed function inside, with and without the input points.

our surface reconstruction pipeline involves three main stages (see Figure 1): Computing an unsigned distance function to the input data. Robustness to outliers and noise is obtained by leveraging the advances in the design of Wasserstein-like metrics, allowing us to reliably identify an -band containing the densely sampled areas. Computing a global, stochastic sign estimation of the distance, first outside the -band where rays are traced against the -band to infer inside vs. outside, then inside the -band by propagating the sign estimates inward. The output of this step is a sign guess for the unsigned distance, along with its confidence ranging from zero to one. Smoothing the estimate to compute the signed distance through

a linear solve to reconstruct a smooth, closed surface. This last step also serves to repair holes. An example of the complete reconstruction pipeline is shown by Figure 2.

1

References

[1] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or, *Consolidation of unorganized point clouds for surface reconstruction*, ACM Transactions on Graphics **28** (2009), 176:1–7.

[2] F. Chazal, D. Cohen-Steiner, and Q. Mrigot, *Geometric inference for measures based on distance functions*, Report 00383685, INRIA, (2009), 1–24.

[3] P. Mullen, F. de Goes, M. Desbrun, D. Cohen-Steiner, and P. Alliez. *Signing the Unsigned: Robust Surface Reconstruction from Raw Pointsets*, Computer Graphics Forum, special issue 7th Annu. Sympos. Geometry Processing **29** (2010), 1733–1741.

Shape spaces, with particular attention to Riemannian Manifolds of Curves

Andrea Menucci

In the first part of the seminar, we ponder on what we use Shape Spaces for, how we design/operate on Shape Spaces, and briefly explore some mathematical toolsets available.

In a Shape Space M, the shape is a point, *i.e.* a variable $x \in M$. There are some natural questions, operations and actions that we may perform on it, related to vague notions of *calculation*, *nearness*, *optimization*, *cataloging*, etc etc. There are moreover actual uses of shape spaces for *shape analysis* and *shape optimization*, that (often implicitly) assume that certain mathematical properties of the underlying shape space are satisfied. We confront these with the available mathematical toolsets, and consequently state a precise set of goals: that the shape space should be a complete Riemannian manifold, where geodesics exist (both as *minimal geodesics* and as *geodesic shooting*, *i.e.* the exponential map is defined and surjective); we moreover ask that the [5] mean shape is defined.

In the second part of the seminar, we briefly see some examples of Shape Spaces from the literature (see bibliographical listing). Then we present a particular Shape Space, a Sobolev-type Riemannian space of curves, presented in [11]. In this model, the manifold M of all smooth immersed curves is endowed with a Riemannian metric, so that the space can be isometrically mapped to $\mathbb{R} \times \mathbb{R}^2 \times M_d$ where $\mathbb{R}$ represents the log-scaling, $\mathbb{R}^2$ the translation, and M_d the curve up to scaling and translation. This implies that in this space the uniform log-scaling and translations of curves are geodesics, and more in general that geodesics can be computed independently in the three components of "log-scaling", "translation", and "deformation" (this latter in M_d). We then concentrate on the space M_d; by a result in [13] this space is isometric to an infinite dimensional Stiefel Manifold. Up to completion, this Stiefel manifold is a complete Riemannian manifold, a smooth submanifold of $L^2 \times L^2$ (where $L^2 = L^2([0, 2\pi])$ is the usual Hilbert space of functions). Subsequently, adapting to the infinite dimensional case a result in

[2], we obtain that the geodesics equation in the Stiefel manifold has a closed-form solution; moreover it was later proven in [4] that any two points can be connected by a minimal geodesic. At the same time it is possible (and computationally light) to solve the gradient problem w.r.t. this metric. So in this Shape Space it is quite convenient to address problems that combine many diverse tools: we eventually present a tracking/filtering approach based on a first order dynamical system in M, again from [11].

The complete seminar is available from *http://mennucci.sns.it/11MFO*.

References

[1] A. Duci and Andrea Mennucci, *Banach-like metrics and metrics of compact sets*, 2007.

[2] A. Edelman, T. Arias, and S. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J Matrix Analy Appl **20** (1998), 303–353.

[3] J. Glaunès, A. Trouvé, and L. Younes,*Modeling planar shape variation via Hamiltonian flows of curves*, In Anthony Yezzi and Hamid Krim, editors, Analysis and Statistics of Shapes, Modeling and Simulation in Science, Engineering and Technology, Birkhäuser - Springer - Verlag, 2004, chapter 14.

[4] P. Harms and A. Mennucci, *Geodesics in infinite dimensional Stiefel and Grassmann manifolds*, *http://cvgmt.sns.it/papers/harmen10/*, 2010.

[5] H. Karcher, *Riemannian center of mass and mollifier smoothing*, Communications on Pure and Applied Mathematics, **30** (1977), 509–541.

[6] A. Mennucci, A. Yezzi, and G. Sundaramoorthi, *Properties of Sobolev Active Contours*, Interf. Free Bound. **10** (2008), 423–445.

[7] P. W. Michor and D. Mumford, *Vanishing geodesic distance on spaces of submanifolds and diffeomorphisms*, Documenta Math **10** (2005), 217–245.

[8] P. W. Michor and D. Mumford, *Riemannian geometris of space of plane curves*, J. Eur. Math. Soc. (JEMS), **8** (2006), 1–48.

[9] P. W. Michor and D. Mumford, *An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach*, Applied and Computational Harmonic Analysis **23** (2007), 76–113.

[10] G. Sundaramoorthi, A. Yezzi, A. Mennucci, and G. Sapiro, *New possibilities with Sobolev active contours*, Intn. Journ. Computer Vision **84** (2008), 113–129.

[11] G. Sundaramoorthi, A. Mennucci, S. Soatto, and A. Yezzi. *A new geometric metric in the space of curves and applications to tracking deforming objects by prediction and filtering*, SIAM Journal on Imaging Sciences **4** (2011), 109-145.

[12] A. Trouvé and L. Younes, *Local geometry of deformable templates*, SIAM J. Math. Anal. **37** (2005), 17–59.

[13] L. Younes, P. W. Michor, J. Shah, and D. Mumford, *A metric on shape space with explicit geodesics*, Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei (9) Mat. Appl. **19** (2008), 25–57.

On Variational Methods in Computer Vision

Christoph Schnörr

The talk was organized into three parts.

(1) First, three decades of variational modeling in computer vision were briefly overviewed, starting with landmark papers on a smooth convex variational

model for low-level vision [5], a non-smooth convex model for image decomposition and denoising [12], a variational model for the segmentation of piecewise smooth and constant images [10], and a probabilistic approach in terms of a variational model to image classes beyond piecewise smooth functions, and to stochastic inference [4].

Notable developments include computationally tractable relaxations and convexity. The latter plays a key role not only in modeling but also in establishing connections between computer vision and related research in applied mathematics. Last but not least, while variational models were considered as exotic in the beginning, they nowadays also have an impact on industrial projects.

(2) In the main part of the talk, novel results of the *continuous multiclass image labeling problem*

$$\min_{\ell} \Big\{ \int_{\Omega} s\big(x, \ell(x)\big) dx + J(\ell) \Big\}, \qquad \ell \colon \Omega \to L = \{1, 2, \ldots, l\}, \tag{1}$$

were presented. Specifically, based on the convex relaxation of (1) suggested in [7] ($\Delta_l \subset \mathbb{R}^l$ denotes the standard simplex),

$$\inf_{u \in C} \Big\{ \int_{\Omega} \big\langle s(x), u(x) \big\rangle dx + J(u) \Big\}, \tag{2a}$$

$$C = \{u \in \mathrm{BV}(\Omega)^l \colon u(x) \in \Delta_l \text{ for a.e. } x \in \Omega\}, \tag{2b}$$

well-posedness was shown in [8], and a complete characterization of regularizers J in terms of metric interaction potentials $d \colon L \times L \to \mathbb{R}_{++}$ was given in [6],[8] such that, for any partition of the image domain $\Omega = S \cup \overline{S}$ with $\mathrm{Per}(S) < \infty$,

$$J(e^i \chi_S + e^j \chi_{\overline{S}}) = d(i,j) \mathrm{Per}(S), \qquad \forall i, j \in L. \tag{3}$$

In the sense of (3), the approach (2) constitutes a relaxation of the piecewise-constant Mumford-Shah model, utilizing arbitrary image features in terms of an application-specific function $s(x)$ that only depends on given image data in some arbitrary way.

From the computational point of view, problem (2) after discretization is a specific instance of the class of non-smooth convex programs

$$\min_{u \in C} \max_{v \in D} \big\{ \langle s, u \rangle + \langle Lu, v \rangle - \langle b, v \rangle \big\}, \qquad D = D_1 \cap \cdots \cap D_r, \tag{4}$$

that also includes a slightly tighter relaxation [2] with different dual constraint sets D_i in (4), as well as problem formulations based on functional lifting [1], like a convex relaxation of the piecewise-*smooth* Mumford-Shah approach [1, 11].

In [9], a specific operator splitting techniques for the dual problem of (4) was devised that efficiently computes a global optimum in terms of iterations, that only use explicit orthogonal projections onto D_i, $i = 1, \ldots, r$, and a linear system that can be diagonalized by the discrete cosine transform.

(3) Finally, the difficulty was pointed out to incorporate shape prior knowledge into image labeling approaches without compromising computational tractability. In [13], a preliminary approach is presented to approximately solve this problem by coupling *binary* continuous cuts (i.e. $l = 2$ in (2); cf. [3]) with convex polyhedral relaxations of MRF-based shape priors.

References

[1] G. Alberti, G. Bouchitté, and G. Dal Maso, *The calibration method for the mumford-shah functional and free-discontinuity problems*, Calc. Var. Part. Diff. Eq. **16** (2003), 299–333.

[2] A. Chambolle, D. Cremers, and T. Pock, *A convex approach for computing minimal partitions*, R.I. 649, CMAP, CNRS UMR 7641, Ecole Polytechnique, 2008.

[3] T.F. Chan, S. Esedoglu, and M. Nikolova, *Algorithms for finding global minimizers of image segmentation and denoising models*, SIAM J. Appl. Math. **66** (2006), 1632–1648.

[4] S. Geman and D. Geman, *Stochastic relaxation, gibbs distributions, and the bayesian restoration of images*, IEEE Trans. Patt. Anal. Mach. Intell. **6** (1984), 721–741.

[5] B.K.P. Horn and B.G. Schunck, *Determining optical flow*, Artif. Intelligence **17** (1981), 185–203.

[6] J. Lellmann, F. Becker, and C. Schnörr, *Convex optimization for multi-class image labeling with a novel family of total variation based regularizers*, In IEEE 12th Int. Conf. Computer Vision, ICCV 2009, pages 646–653.

[7] J. Lellmann, J. Kappes, J. Yuan, F. Becker, and C. Schnörr, *Convex multi-class image labeling by simplex-constrained total variation*, In X.-C. Tai, K. Mórken, M. Lysaker, and K.-A. Lie, editors, *Scale Space and Variational Methods in Computer Vision (SSVM 2009)*, volume 5567 of *LNCS*, pages 150–162. Springer, 2009.

[8] J. Lellmann and C. Schnörr, *Continuous multiclass labeling approaches and algorithms*, Preprint 2011.

[9] L. Lellmann, D. Breitenreicher, and C. Schnörr, *Fast and exact primal-dual iterations for variational problems in computer vision*, In Computer Vision - ECCV 2010, volume 6312 of *LNCS*, pages 494–505. Springer, 2010.

[10] D. Mumford and J. Shah, *Optimal approximations by piecewise smooth functions and associated variational problems*, Comm. Pure Appl. Math. **42** (1989), 577–685.

[11] T. Pock, D. Cremers, H. Bischof, and A. Chambolle, *An algorithm for minimizing the mumford-shah functional*, In Proc. ICCV, IEEE Comp. Soc., 2009.

[12] L. Rudin, S. Osher, and E. Fatemi, *Nonlinear total variation based noise removal algorithms*, Physica D **60** (1992), 259–268.

[13] B. Schmitzer and C. Schnörr, *Weakly convex coupling continuous cuts and shape priors*, In Proc. 3rd Int. Conf. Scale Space and Variational Methods in Computer Vision (SSVM'11), *LNCS*. Springer, 2011, to appear.

Anatomically Complex Representation of Functional Anatomy

Michael Miller

Functional anatomy is the study of structure and function in curved anatomical coordinate systems. In this presentation we discuss representations of functional anatomy in the human brain. In part, our talk will focus on probability models of subcortical neuroanatomy which is as complex as the curved coordinate systems(caudate, putamen, hippocampus, amygdala, ventricles ...) of the brain itself. We will show the results from studies of human neuroanatomy, including several disease sets.

Regularized Reconstruction of M-Rep Shapes with Statistical A Priori Knowledge

Otmar Scherzer
(joint work with Matthias Fuchs)

The reconstruction of geometry or, in particular, the shape of objects is a common issue in image analysis. Starting from a variational formulation of such a problem on a shape manifold we introduce a regularization technique incorporating statistical shape knowledge. The key idea is to consider a Riemannian metric on the shape manifold which reflects the statistics of a given training set.

We expect the surface to be reconstructed the minimizing argument of an energy functional: In case of image segmentation we use for instance the Mumford-Shah functional [3] or the "Snakes" energy [2]. These energies incorporate some kind of regularization to ensure the well-posedness of the corresponding variational problems.

A second basic problem in shape analysis and recovery is the right representation of the hypersurface in implementations. There exists no canonical approach to model such objects, but the right choice depends on the expected topology and regularity of the solution.

We propose the use of *intelligent* shape models to represent the geometry we want to detect and to regularize the underlying detection problem. These shape models share the following two characteristics:

- A shape model is associated with a finite dimensional parameter manifold. An element of this manifold corresponds to an instance of the shape model.
- A shape model can be associated with statistical data which describes how frequently individual instances of the shape model occur.

We employ the two properties above to define a regularization technique which takes into account a priori information on the expected solution. As illustrated by some examples this enables us to detect geometries even if the original data is perturbed (e.g. some parts of it are missing or occluded), or allows us to significantly reduce the complexity of reconstruction methods.

In our work we combine the idea of adding a statistically motivated regularization term to a variational problem with the use of advanced shape models on manifolds. In particular, it provides a framework to use M-Reps [4] as a starting point for a Mahalanobis regularization of variational segmentation problems. To our knowledge this ([1]) is the first paper which combines the idea of M-Reps with statistical segmentation functionals. An extension of M-Reps to a shape space with a purely *geometric* metric is also discussed.

This work summarizes results from [1].

References

[1] M. Fuchs and O. Scherzer, *Regularized Reconstruction of Shapes with Statistical a priori Knowledge*, International Journal of Computer Vision **79** (2008), 119-135.

[2] M. Kass, A. Witkin, and D. Terzopoulos, *Snakes active contour models*, International Journal of Computer Vision **1** (1988), 321–331.
[3] D. Mumford and J. Shah, *Boundary detection by minimizing functionals*, Proc. IEEE Conf. Computer Vision Pattern Recognition, pages 22–26, 1985.
[4] K. Siddiqi and S. Pizer, *Medial Representations*, Springer, Computational Imaging Series, 2008.

Reconstruction of a 3D shape from its apparent contour

Giovanni Bellettini
(joint work with V. Beorchia and M. Paolini)

I have described[1] various recent results that we have obtained in the effort of reconstructing a three dimensional smooth (compact solid) scene[2] $E \subset \mathbb{R}^3$ starting from informations on its apparent contour[3] $\mathrm{AC}(\partial E)$, endowed with a Huffman labeling. Our original motivation was the study of a two-dimensional variational model[4]; however, the problem became almost immediately a topological one.

Theorem 1. *(see* [3]*). Let V be an oriented planar graph with terminal points, none of which is exterior, and with T-junctions, such that the exterior region lies on the left. Then there exists an oriented planar graph A endowed with a Huffman labelling, such that V is the visible part of A.*

This theorem does not specify a "preferred" completion of V; it says that there is at least one completion of V. The proof is constructive[5] and may require the introduction of several additional cusps and crossings, so that A may look like very complicated. It is therefore natural to ask whether A can be "simplified". This will be made more clear by the next results (see in particular Theorem 4 and Remark 2).

[1] See the references for precise definitions and statements.

[2] The boundary ∂E of E is not connected, in general.

[3] The apparent contour lies in the retinal plane $\mathbb{R}^2$, and is the image through the orthogonal projection $\pi : \mathbb{R}^3 \to \mathbb{R}^2$ of the critical curve (this is the curve where the rank of the differential of $\pi_{|\partial E}$ is not maximal). ∂E is supposed to be in generic position with respect to π. $\mathrm{AC}(\partial E)$ is an oriented graph with cusps and crossing singularities. The Huffman labeling is a pair of integer valued functions (f, d), the function f defined on $\mathbb{R}^2 \setminus \mathrm{AC}(\partial E)$, and the function d defined on $\mathrm{AC}(\partial E)$ itself, satisfying a list of compatibility conditions. f is constant on the connected components of $\mathbb{R}^2 \setminus \mathrm{AC}(\partial E)$: $f(x)$ counts the number of points on ∂E that project on $x \in \mathbb{R}^2$, and it can be recovered from the orientation of the graph. d is constant on the (relatively open) arcs of $\mathrm{AC}(\partial E)$: $d(x)$ counts the number of folds of ∂E which are in front of the unique point on the critical curve that projects on x. The visible part of the apparent contour is the set where $d = 0$, and is an oriented graph with T-junctions and terminal points singularities (the orientation is such that the exterior region lies on the left, and no terminal point is exterior).

[4] See [1]: this variational model could maybe be considered as a sort of refinement of the Nitzberg-Mumford model [8], [9].

[5] The proof is implemented inside the code appcontour developed by Maurizio Paolini (Università Cattolica di Brescia, Italy), see also footnote 7.

Theorem 2. *(see* [1]*). Let A be a planar oriented graph endowed with a Huffman labeling. Then there exists a 3D shape E such that $A = \mathrm{AC}(\partial E)$. Moreover, E is unique up to a continuous transformation of $\mathbb{R}^3$ which is strictly monotone in the view direction.*

Coupling this existence and uniqueness result with Theorem 1, we obtain a 3D shape E starting only from the graph V.

Remark 1. Finding a "preferred" 3D shape in the class of all shapes differing one each other by a transformation of $\mathbb{R}^3$ mentioned in Theorem 2 is a problem that can be settled in variational form and possibly implemented on a computer, and it is under investigation.

We remark that the Euler-Poincaré characteristic $\chi(\partial E)$ of ∂E (and hence of E and of $\mathbb{R}^3 \setminus E$) can be inferred from the Huffman labeling of A. Instead, we do not provide an explicit formula for counting the number of connected components of ∂E.

Definition 3. *We say that two 3D shapes E and F are equivalent if there exists a smooth path of isotopies of $\mathbb{R}^3$ which takes E into F.*

Theorem 4. *(see* [4]*). Let E and F be two 3D shapes. Then E and F are equivalent if and only if $\mathrm{AC}(\partial E)$ and $\mathrm{AC}(\partial F)$ can be joined by a finite number of elementary moves (and of their inverses), taken in an explicit (finite) list L.*

Such a kind of result can probably be extracted from more general results (see for instance [5]). The proof of [4] is given in our framework of closed embedded surfaces in $\mathbb{R}^3$ and it is reasonably[6] self-contained. The list L contains (strictly) the Reidemeister's moves for knots, which can be obtained when restricting to a smooth tubular surface around a knot.

Remark 2. It is also possible to prove, in the spirit of [7], but with the restrictions given by our embedded context, that an apparent contour can always be deformed using planar isotopies and moves (and their inverses) taken from L, so that all cusps disappear (but the number of crossings may possibly increase).

The software appcontour[7] is a code developed by Maurizio Paolini, in the effort of topologically reconstructing ∂E starting from an apparent contour endowed with a Huffman labeling. It needs an encoded description of the graph A[8] and, at the moment, it returns the number of connected components of ∂E, and the Euler-Poincaré characteristic of each of these connected components (and in particular $\chi(\partial E)$)[9]. Appcontour returns also some informations on the hierarchical inclusions between the various connected components of ∂E.

[6]Theorem 4 is based on a crucial result of J. Mather on the classification of singularities of stable maps from a three-dimensional manifold to a three-dimensional manifold.

[7]Freely downloadable at `http://appcontour.sourceforge.net`

[8]The description of A is topological in character.

[9]It also provides various invariants of apparent contours such as the Bennequin invariant, see for instance [6] and [2].

Remark 3. It would be interesting to find other topological invariants of ∂E which can be deduced looking at the graph A and at its labeling. In particular, suppose that ∂E_1 is the knotted torus, and take another connected surface ∂E_2, inequivalent to ∂E_1, obtained as follows: remove from a sphere two small holes, one around the north pole and the other one around the south pole. Then connect the holes with a knotted tube interior to the sphere (the "tunnel" inside the sphere is part of $\mathbb{R}^3 \setminus E_2$), see the cube with knotted holes as in Figure on page 31 in [10]. We have $\chi(\partial E_1) = \chi(\partial E_2)$. We do not know it there is a quantity readable from the corresponding apparent contours which allows to deduce that ∂E_1 and ∂E_2 are not equivalent.

I have concluded the discussion by pointing out some problems and generalizations that at the moment seem to be open.

References

[1] G. Bellettini, V. Beorchia, M. Paolini, *Topological and variational properties of a model for the reconstruction of three-dimensional transparent images with self-occlusions*, Int. J. Imaging Vision **32** (2008), 265–291.

[2] G. Bellettini, V. Beorchia, M. Paolini: *An explicit formula for a Bennequin-type invariant for apparent contours*, Topology Appl. **156** (2009), 747–760.

[3] G. Bellettini, V. Beorchia, M. Paolini: *Completion of visible contours*, SIAM J. Imaging Sci. **2** (2009), 777-799.

[4] G. Bellettini, V. Beorchia, M. Paolini: *Completeness of Reidemeister-type moves for surfaces embedded in three-dimensional space*, 2009, submitted.

[5] J.S. Carter, J.H. Rieger, M. Saito: *A combinatorial description of knotted surfaces and their isotopies*, Adv. Math. **127** (1997), 1–51.

[6] Ko Honda, 3-*Dimensional Methods in Contact Geometry,* 47–86 in S.K. Donaldson, Y. Eliashberg, M. Gromov Eds., Different Faces of Geometry, Internat. Math. Ser., Kluwer Acad. Publ. 2004.

[7] H.I. Levine, *Elimination of cusps*, Topology **3** (1965), 263–295.

[8] M. Nitzberg and D. Mumford, *The* 2.1-*D sketch*, in Proc. of the Third International Conference on Computer Vision, Osaka (1990).

[9] M. Nitzberg, D. Mumford and T. Shiota, *Filtering, Segmentation and Depth*, Lecture Notes in Computer Science **662**, Springer-Verlag, Berlin, 1993.

[10] D. Rolfsen, *Knots and Links*, AMS Chelsea Publishing, 2003.

Shape Space Exploration of Planar Quad Meshes

Niloy J. Mitra

(joint work with Yong-Liang Yang, Yi Jun Yang, and Helmut Pottmann)

Given a single planar quad (PQ) mesh, we define a *PQ mesh manifold* as the space of all PQ meshes having the same combinatorics as the input mesh. Navigation in this manifold of usually high dimension and co-dimension is computationally accessed through first and second order approximations, namely tangent spaces and quadratically parameterized osculant surfaces. Only those parts of the manifold are useful that correspond to aesthetically pleasing meshes. Aesthetics and other constraints of practical interest are cast into energy functions, for which we derive second order approximantions that are intrinsic to the PQ mesh manifold

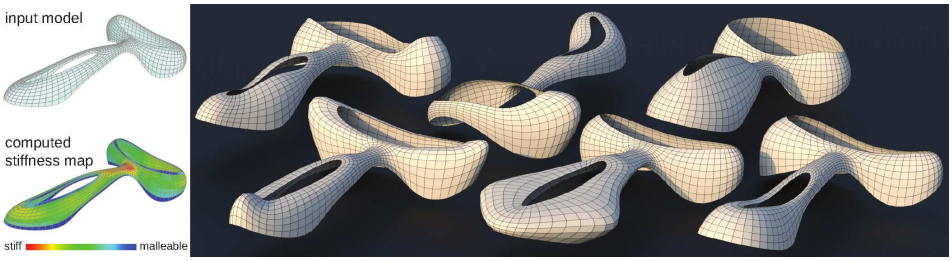

FIGURE 1. Starting from a single planar quad (PQ) mesh, our geometric framework allows navigation and exploration of the shape space of PQ meshes, enabling easy creation of aesthetic model variants. A global stiffness map informs the user of relative flexibility across the model (reference geometry from Yas Island Marina Hotel, Abu Dhabi, by Asymptote Architecture).

and computable in the parameter space of the osculant surface. Subsequently, we demonstrate our proposed geometric framework for navigation, design exploration, and optimization on a variety of architectural designs, especially in the context of form finding and fabrication-aware shape exploration (see Figure 1). Thus, we unify two traditionally separate phases in freeform architecture, namely (i) shape design and (ii) rationalization in view of the actual fabrication. Our mathematical formulation is general and suited for design exploration of other types of nonlinearly constrained geometric models.

Future research directions in multi-resolution PQ shape space exploration, and locally supported PQ deformations were discussed. We also elaborated about potential applications to other optimized or constrained meshes, specially for the purpose of fabrication-aware design.

Geometric discrete Laplacians on polygonal meshes

MAX WARDETZKY
(joint work with Marc Alexa)

We consider the construction of geometric Laplacians on discrete surfaces made up of polygonal faces, encompassing non-planar and non-convex polygons. Our construction is guided by closely mimicking structural properties of the smooth Laplace-Beltrami operator, leading to an extension of the widely employed cotan formula from triangle to polygonal meshes. Our approach sheds light on the geometry behind a mimetic finite difference scheme suggested by Brezzi and coworkers.

From smooth to discrete Laplacians. We consider an *oriented* 2-manifold mesh Σ, possibly with boundary, with vertex set V, edge set E, and face set F. We allow for faces that are *simple*, but possibly non-planar, polygons in $\mathbb{R}^3$. In our setup, faces are solely determined by their boundary edge loops. By *simple* we mean that each polygon forms a closed non self-intersecting loop. By *oriented*

we mean that all faces are equipped with a consistent orientation in such a way that any two adjacent faces induce opposite orientations on their common edge. In order to distinguish between the resulting two orientations of each inner edge, we work with oriented half-edges. In analogy to standard notation in the smooth setting, we let Ω^k, $k \in \{0,1\}$, be the real vector space of discrete k-forms on Σ. In particular, 0-forms are real values associated with vertices, while 1-forms are real values associated with half-edges. We require that $\alpha(e_{ij}) = -\alpha(e_{ji})$ for any 1-form α and any oriented half-edge e_{ij} from vertex i to vertex j.

In order to define discrete Laplacians acting on 0-forms (functions), we build on the smooth setting, where

$$\Delta = d^* d \ . \tag{1}$$

Here, as usual, $d : \Omega^0 \to \Omega^1$ denotes Cartan's exterior derivative with formal adjoint d^*, i.e., $(d^*\alpha, u)_{L^2} = (\alpha, du)_{L^2}$ for all $\alpha \in \Omega^1$ and all $u \in \Omega^0$.

In order to mimic the smooth setting, we are confronted with two tasks: (i) to provide a discrete version $\delta : \Omega^0 \to \Omega^1$ of Cartan's exterior derivative and (ii) to provide positive inner products on the spaces Ω^k of discrete k-forms for $k \in \{0,1\}$. Indeed, these two ingredients uniquely pin down the adjoint operator $\delta^* : \Omega^1 \to \Omega^0$ and hence—using (1) as a definition—lead to a family of discrete Laplacians.

The solution to the first task is naturally provided by considering the coboundary operator $\delta : \Omega^0 \to \Omega^1$, defined by $(\delta u)(e_{ij}) = u(j) - u(i)$. The second task requires choosing inner products on the spaces Ω^0 and Ω^1, which we represent by symmetric positive definite matrices M_0 and M_1, respectively. *Any* concrete choice of M_0 and M_1 then yields a discrete Laplacian of the form

$$\mathcal{L} = \delta^* \delta = M_0^{-1} L \quad \text{with} \quad L = \delta^T M_1 \delta \ . \tag{2}$$

Notice that $\mathcal{L}$ corresponds to the *strong* (or pointwise) formulation of the Laplacian, whereas L represents the *weak* (or integrated) version. Indeed, L gives rise to the Dirichlet energy of a function u defined on the vertices via $E_D(u) = \frac{1}{2} u^T L u$.

Desiderata. While any choice of inner products M_0 and M_1 would in principle be conceivable, not all choices are created equal in terms of resembling and maintaining structural properties of the smooth setting. In order to narrow the set down to "good" inner products, we are guided by requiring a number of important properties that hold for smooth Laplacians.

Locality (Loc). The smooth Laplacian is a differential operator; therefore, its definition is local. In the discrete case, we maintain locality by only working with diagonal matrices M_0 and by requiring that M_1 be defined *per face*, i.e.,

$$\alpha^T M_1 \beta = \sum_{f \in F} \alpha_{|f}^T M_f \beta_{|f} \tag{3}$$

for any pair (α, β) of discrete 1-forms. The sum runs over all faces and $\alpha_{|f}$ and $\beta_{|f}$, respectively, denote the restrictions of α and β to the boundary ∂f of f. For each $f \in F$, we work with a symmetric matrix M_f.

Symmetry (Sym). The symmetry of M_f (and thus M_1) reflects the fact that on a Riemannian manifold without boundary the Laplacian is formally self-adjoint.

Positive semi-definiteness (Psd). The smooth Laplacian on a Riemannian manifold without boundary is positive semi-definite with one-dimensional kernel equal to the constants. In the discrete setting we require that M_0 and each M_f (and therefore M_1) be positive definite. The kernels of $\mathcal{L}$ and L are then automatically one-dimensional since δ has a one-dimensional kernel given by the constants.

Linear precision(Lin). In the smooth case, if $\Sigma \subset \mathbb{R}^2$ is a planar domain and $u : \mathbb{R}^2 \to \mathbb{R}$ is a linear function, then $\Delta u = 0$. In the discrete case we require that if all vertices lie in a single plane then $(\mathcal{L}u)_i = 0$ at each interior vertex i and for each linear function u over the plane.

Scale invariance (Sca). In dimension two, the (smooth) Dirichlet energy, E_D, is a conformal invariant. In particular, E_D is invariant under uniformly rescaling a smoothly embedded surface. Therefore, we require that the weak discrete Laplacian, L, remains unchanged under uniformly rescaling a mesh by requiring that each M_f be scaling invariant.

Mean curvature & area gradient. The intrinsic Laplace-Beltrami operator of a surface that is smoothly embedded into $\mathbb{R}^3$ is intimately connected to the mean curvature vector via $\mathbb{H} = \Delta \mathbf{x}$, where $\mathbf{x} : \Sigma \to \mathbb{R}^3$ denotes the embedding.

Moreover, in the smooth case the mean curvature vector, $\mathbb{H}$, equals the L^2-gradient of the area functional. Extending this property to polygonal meshes, we consider *vector area* instead of surface area. The vector area, $\mathbf{A}(\gamma)$, of a simple closed (and sufficiently regular) curve $\gamma \subset \mathbb{R}^3$ is given by the surface integral of the unit normal over a (sufficiently regular) surface with boundary γ. $\mathbf{A}(\gamma)$ only depends on the boundary curve, *not* on the choice of a particular surface spanning this curve. Indeed,

$$\mathbf{A}(\gamma) = \frac{1}{2} \oint_\gamma \mathbf{x} \times d\mathbf{x} \ ,$$

where $\mathbf{x}$ denotes the position vector of γ.

Consider now all orthogonal projections of γ to planes in $\mathbb{R}^3$, together with the resulting signed areas of the planar regions enclosed by the images of γ. Then $|\mathbf{A}(\gamma)|$ is equal to the largest such area, and we call the corresponding orthogonal projection *maximal* with image $\bar{\gamma}$. In particular, $\mathbf{A}(\gamma) = \mathbf{A}(\bar{\gamma})$.

If f is a (possibly non-planar) polygon in $\mathbb{R}^3$ with vector area of magnitude $|\mathbf{A}(f)|$, then the gradient with respect to varying a vertex i of f satisfies

$$\nabla_i |\mathbf{A}(f)| = \left(\widetilde{L}_f X_f \right)_i \ , \quad \text{where } \widetilde{L}_f := \delta^T \widetilde{M}_f \delta \quad \text{and } \widetilde{M}_f := \frac{1}{|\mathbf{A}(f)|} B_f B_f^T \ .$$

Here B_f denotes the matrix whose rows correspond to the 3D positions of edge midpoints of f in a cyclic ordering of the edges.

A family of discrete Laplacians. Using (3), the matrices $\widetilde{M}_f$ yield positive semi-definite inner products on 1-forms, which by (2) give *pre-Laplacians* L that turn out to satisfy (Loc), (Sym), (Lin), and (Sca), but unfortunately *not* (Psd) for general polygonal meshes.

To recover (PSD) without giving up the other properties, we extend the construction in [1] from the case of planar to non-planar polygons. To this end, consider a spatial polygon f with k_f vertices, and let $\bar{f}$ denote the (planar) maximal projection of f. Let $E_{\bar{f}}$ be the $(k_f \times 3)$ matrix with rows representing the (cyclically ordered) edge vectors of $\bar{f}$. Let $C_{\bar{f}}$ be a $k_f \times (k_f - 2)$ matrix with columns spanning the null space of $E_{\bar{f}}^T$, and let $U_{\bar{f}}$ be a symmetric positive definite $(k_f - 2) \times (k_f - 2)$ matrix. Then:

Theorem. *The matrices* $M_f := \widetilde{M}_f + C_{\bar{f}} U_{\bar{f}} C_{\bar{f}}^T$ *give rise to positive definite inner products on* 1*–forms and, using* (2)*, yield Laplacians* L *that satisfy* (LOC)*,* (SYM)*,* (LIN)*, and* (PSD)*. Additionally, if* $C_{\bar{f}}$ *and* $U_{\bar{f}}$ *are chosen such that they are invariant under uniformly rescaling* f*, then* L *also satisfies* (SCA)*. Moreover, for triangle meshes, any choices of* $C_{\bar{f}}$ *and* $U_{\bar{f}}$ *yield the cotan Laplacian in* [2]*.*

As a consequence, our result generalizes the cotan formula from triangle to polygonal meshes and extends the construction in [1] from planar to non-planar polygons, while—via gradients of vector area—shedding light on the geometry underlying the corresponding mimetic finite difference construction.

References

[1] F. Brezzi, K. Lipnikov, and V. Simoncini, *A family of mimetic finite difference methods on polygonal and polyhedral meshes*, In Math. Models Methods Appl. Sci. **15** (2005), 1533–1553.

[2] U. Pinkall and K. Polthier, *Computing discrete minimal surfaces and their conjugates*, In Experim. Math. **2** (1993), 15–36.

Variational Imaging with TGV^2-regularization

Kristian Bredies

Most mathematical formulations of inverse problems are cast in the form of minimizing a Tikhonov functional, i.e.,

$$\min_u \; F(u) + \alpha R(u)$$

where F represents the fidelity with respect to the measured data, R is a regularization functional and $\alpha > 0$ a parameter. For mathematical imaging problems, letting R the *total variation* is the most common choice due to its edge-preserving properties [6]. However, for piecewise smooth images, TV tends to undesired piecewise constant solutions, the "staircasing artifacts".

In [1], the *total generalized variation* (TGV) of order k, defined as

$$\text{(1)}\quad \mathrm{TGV}_\alpha^k(u) = \sup\Big\{ \int_\Omega u \mathrm{div}^k v \,\mathrm{d}x \;\Big|\; v \in \mathcal{C}_\mathrm{c}^k(\Omega, \mathrm{Sym}^k(\mathbb{R}^d)), \; \|\mathrm{div}^l v\|_\infty \le \alpha_l,\ l = 0, \dots, k-1 \Big\},$$

has been proposed and analyzed. It constitutes a new image model which can be interpreted to incorporate smoothness from the first up to the k-th derivative.

Here, $\mathrm{Sym}^k(\mathbb{R}^d)$ denotes the space of symmetric tensors of order k with arguments in $\mathbb{R}^d$ and $\alpha_l > 0$ are fixed parameters. Choosing $k = 1$ and $\alpha_0 = 1$ yields the usual total variation functional. For $k = 2$, the functional reads as follows:

Definition 5. *Let $\Omega \subset \mathbb{R}^d$ be a bounded domain and $\alpha = (\alpha_0, \alpha_1) > 0$. For $u \in \mathrm{L}^1_{\mathrm{loc}}(\Omega)$, the functional*

$$\mathrm{TGV}^2_\alpha(u) = \sup \left\{ \int_\Omega u \mathrm{div}^2 v \, \mathrm{d}x \;\middle|\; v \in \mathcal{C}^2_{\mathrm{c}}(\Omega, S^{d\times d}), \|v\|_\infty \le \alpha_0, \; \|\mathrm{div}\, v\|_\infty \le \alpha_1 \right\} \tag{2}$$

is called the total generalized variation *of second order.*

Here, $S^{d\times d}$ is the set of symmetric matrices, $\mathcal{C}^2_{\mathrm{c}}(\Omega, S^{d\times d})$ the vector space of compactly supported, twice continuously differentiable $S^{d\times d}$-valued mappings and $divv \in \mathcal{C}^1_{\mathrm{c}}(\Omega, \mathbb{R}^d)$, $div^2 v \in \mathcal{C}_{\mathrm{c}}(\Omega)$ is defined by

$$(\mathrm{div}\, v)_i = \sum_{j=1}^d \frac{\partial v_{ij}}{\partial x_j},$$

$$\mathrm{div}^2 v = \sum_{i=1}^d \frac{\partial^2 v_{ii}}{\partial x_i^2} + 2 \sum_{i<j} \frac{\partial^2 v_{ij}}{\partial x_i \partial x_j}.$$

The norms of $v \in \mathcal{C}_{\mathrm{c}}(\Omega, S^{d\times d})$, $\omega \in \mathcal{C}_{\mathrm{c}}(\Omega, \mathbb{R}^d)$ are given by

$$\|v\|_\infty = \sup_{x\in\Omega} \Big(\sum_{i=1}^d |v_{ii}(x)|^2 + 2 \sum_{i<j} |v_{ij}(x)|^2 \Big)^{1/2},$$

$$\|\omega\|_\infty = \sup_{x\in\Omega} \Big(\sum_{i=1}^d |\omega_i(x)|^2 \Big)^{1/2}.$$

Basic results about this functional obtained in [1] can be summarized as follows.

Theorem 6. *Total generalized variation of second order enjoys the following properties:*

1. *TGV^2_α is a semi-norm on the Banach space $BGV^2_\alpha(\Omega)$,*
2. *$TGV^2_\alpha(u) = 0$ if and only if u is a polynomial of degree less than 2,*
3. *TGV^2_α and $TGV^2_{\tilde\alpha}$ are equivalent for $\tilde\alpha = (\tilde\alpha_0, \tilde\alpha_1) > 0$,*
4. *TGV^2_α is rotationally invariant,*
5. *TGV^2_α satisfies, for $r > 0$, $\rho_r u(x) = u(rx)$ and $(\tilde\alpha_0, \tilde\alpha_1) = (\alpha_0 r^2, \alpha_1 r)$, the scaling property*
$$TGV^2_\alpha \circ \rho_r = r^{-d} TGV^2_{\tilde\alpha}(u),$$
6. *TGV^2_α is proper, convex and lower semi-continuous on each $\mathrm{L}^p(\Omega)$, $1 \le p < \infty$.*

Moreover, as shown in [2], one can use Fenchel-Rockafellar duality to obtain the following dual representation:

Theorem 7. *For each $u \in \mathrm{L}^1(\Omega)$ we have*

$$TGV^2_\alpha(u) = \min_{w \in BD(\Omega)} \alpha_1 \|Du - w\|_{\mathcal{M}} + \alpha_0 \|\mathcal{E}w\|_{\mathcal{M}} \tag{3}$$

where $BD(\Omega)$ denotes the space of vector fields of Bounded Deformation [7]*, i.e., $w \in \mathrm{L}^1(\Omega, \mathbb{R}^d)$ such that the distributional symmetrized derivative $\mathcal{E}w = \frac{1}{2}(\nabla w + \nabla w^{\mathrm{T}})$ is a $S^{d\times d}$-valued Radon measure.*

The minimum characterization (3) yields an intuitive, informal interpretation of TGV^2_α: If, locally, u is smooth, we have $|\nabla^2 u| \ll |\nabla u|$, so it might be favorable for (3) to locally choose $w = \nabla u$ leading to the penalization of $|\nabla^2 u|$. On the other hand, if, locally, u possesses an edge, ∇u will not be approximated by any w, hence choosing $w = 0$ locally might be favorable for (3), leading to the penalization of $|\nabla u|$. This way, TGV^2_α reflects an optimal balancing between the first and the second derivative.

Total generalized variation of second order can immediately be applied to image denoising problems. Numerical experiments carried out in [1] show that TGV^2_α-regularized L^2-denoising produces visually appealing results with almost no staircase effect present in the solution (see Figure 1). When solving the ill-posed inverse problem $Ku = f$ for $K : \mathrm{L}^2(\Omega) \to \mathrm{L}^2(\Omega)$ with TGV^2 regularization, however, coercivity is needed. This follows from an topological equivalence result for norms in $\mathrm{BV}(\Omega)$:

Theorem 8. *Let $\Omega \subset \mathbb{R}^d$ be a bounded Lipschitz domain. Then there exist constants $0 < c < C < \infty$ such that for each $u \in BGV^2_\alpha(\Omega)$ there holds*

$$c\|u\|_{\mathrm{BV}} \le \|u\|_1 + TGV^2_\alpha(u) \le C\|u\|_{\mathrm{BV}}.$$

Sketch of proof. Setting $w = 0$ in (3) immediately implies that for each $u \in \mathrm{BGV}^2_\alpha(\Omega)$ we have $\mathrm{TGV}^2_\alpha(u) \le \alpha_1 \mathrm{TV}(u)$, hence we can set $C = \max(1, \alpha_1)$.

For the converse estimate, we first show that there is a $C_1 > 0$ such that for each $u \in \mathrm{BV}(\Omega)$ and $\bar{w} \in \ker \mathcal{E}$ there holds

$$\|Du\|_{\mathcal{M}} \le C_1\big(\|Du - \bar{w}\|_{\mathcal{M}} + \|u\|_1\big). \tag{4}$$

Next, recall that a Sobolev-Korn inequality is valid for $\mathrm{BD}(\Omega)$ [7]: There is a $C_2 > 0$ such that for each $w \in \mathrm{BD}(\Omega)$ there exists a $\bar{w} \in \ker \mathcal{E}$ such that $\|w - \bar{w}\|_1 \le C_2\|\mathcal{E}w\|_{\mathcal{M}}$. For this $\bar{w}$, we have, for some $C_3 > 0$,

$$\begin{aligned}\|Du - \bar{w}\|_{\mathcal{M}} &\le \|Du - w\|_{\mathcal{M}} + \|w - \bar{w}\|_1 \\ &\le C_3\big(\alpha_1\|Du - w\|_{\mathcal{M}} + \alpha_0\|\mathcal{E}w\|_{\mathcal{M}}\big)\end{aligned}$$

Plugged into (4) and adding $\|u\|_1$ on both sides, it follows that the inequality

$$\|u\|_{\mathrm{BV}} \le C_4\big(\|u\|_1 + \alpha_1\|Du - w\|_{\mathcal{M}} + \alpha_0\|\mathcal{E}w\|_{\mathcal{M}}\big)$$

holds for some $C_4 > 0$ independent of u and w. Taking the minimum over all $w \in \mathrm{BD}(\Omega)$ and choosing $c = C_4^{-1}$ finally yields the result by virtue of (3). □

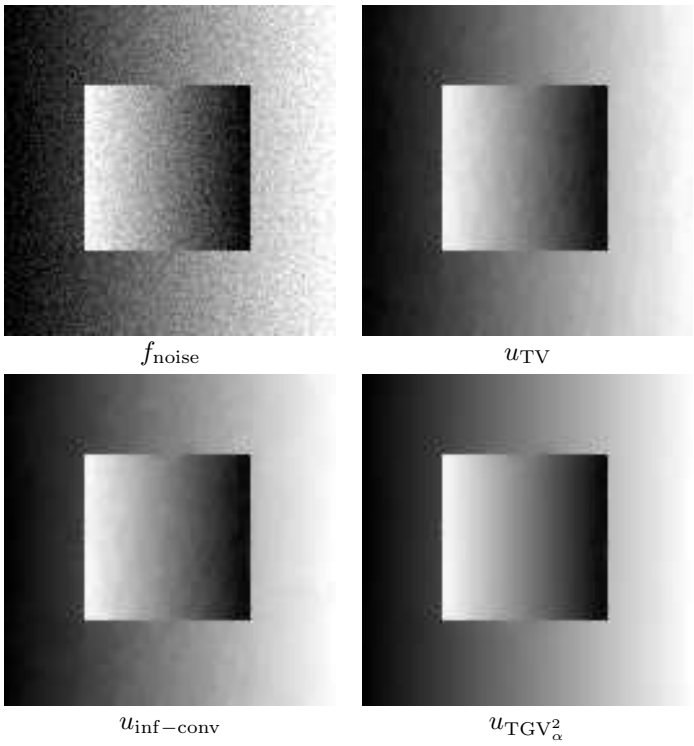

FIGURE 1. Denoising with Total Variation, infimal convolution and Total generalized variation.

The norm equivalence implies that $\text{BGV}^2_\alpha(\Omega) = \text{BV}(\Omega)$ in the sense of equivalent Banach spaces. In particular, the known (compact) embeddings into $\mathcal{L}^p$-spaces hold and one easily obtains [2]:

Corollary 9. *Let $1 < p < \infty$ such that $p \leq d/(d-1)$ and $P : \text{L}^p(\Omega) \to \mathcal{P}^1(\Omega)$ a linear projection onto the space of affine functions $\mathcal{P}^1(\Omega)$. Then, there is a $C > 0$ such that*

$$\|u\|_p \leq C\, TGV^2_\alpha(u) \qquad \forall u \in \ker P \subset \text{L}^p(\Omega). \tag{5}$$

The stated coercivity is then the main ingredient to apply the direct method for showing existence of minimizers for the Tikhonov functionals in order to solve $Ku = f$ (confer again [2]):

Theorem 10. *Let $1 < p < \infty$, $p \leq d/(d-1)$, Y be a Hilbert space, $K \in \mathcal{L}\big(\text{L}^p(\Omega), Y\big)$ a linear and continuous operator which is injective on $\mathcal{P}^1(\Omega)$ and $f \in Y$. Then, the problem*

$$\min_{u \in \text{L}^p(\Omega)} \ \tfrac{1}{2}\|Ku - f\|^2 + TGV^2_\alpha(u) \tag{6}$$

admits a solution.

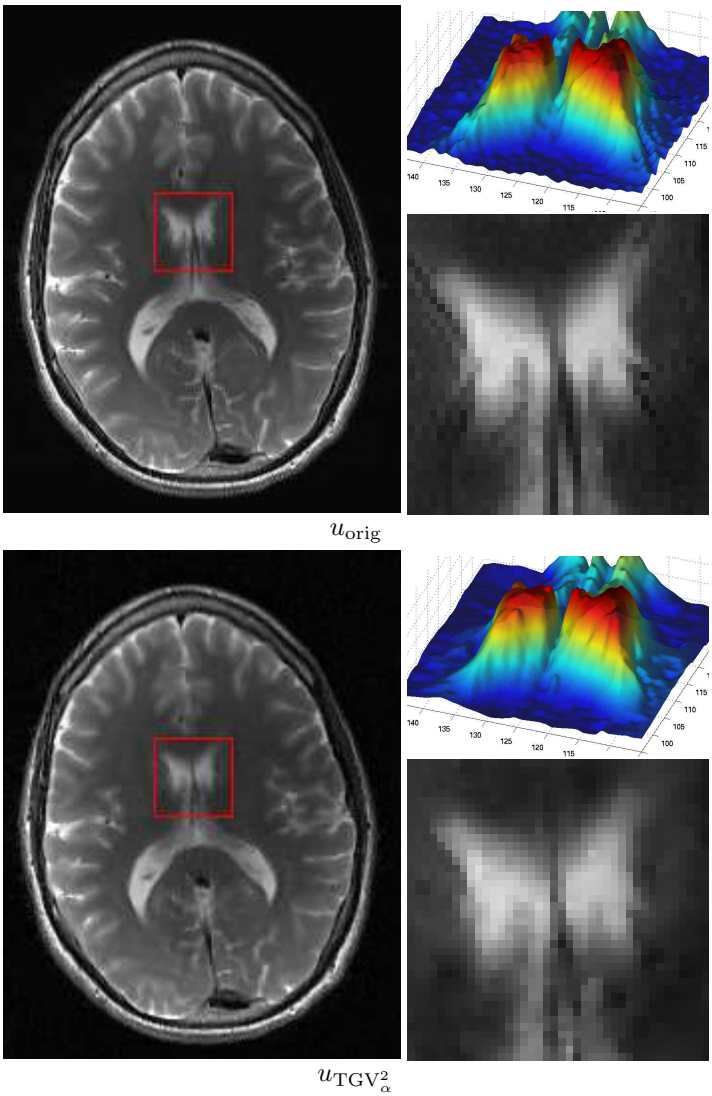

FIGURE 2. MRI undersampling reconstruction from $\approx 6\%$ of the Fourier points with TGV^2_α.

This result can now immediately be applied to a variety of inverse imaging problems. One of them is undersampling reconstruction in magnetic resonance imaging (MRI). Since in usual MRI, acquisition times are high, one way to accelerate imaging is recording only a small subset of the necessary data using multiple,

FIGURE 3. Deconvolution example. The original image u_{orig} [4] has been blurred and contaminated by noise resulting in f, u_{TV} and $u_{\text{TGV}^2_\alpha}$ are the regularized solutions recovered from f.

independent receiver coils. Given the data $f_1, \ldots, f_n$, the TGV^2_α-MRI reconstruction problem reads as:

$$\min_{u \in \text{L}^2(\Omega)} \frac{1}{2} \sum_{i=1}^{n} \|\widehat{\sigma_i u} - f_i\|^2_{2,\Sigma} + \text{TGV}^2_\alpha(u)$$

where u is the image to reconstruct, $\sigma_1, \ldots, \sigma_n$ are the *coil sensitivities*, $\widehat{\sigma_i u}$ denotes the Fourier transform of $\sigma_i u$ and $\Sigma \subset \mathbb{R}^2$ is the *undersampling pattern*. Numerical experiments performed in [5] confirm that this approach leads to high-quality reconstructions (see Figure 2) which depict human tissue more faithful than linear reconstructions or TV-regularized reconstructions.

Finally, another application of TGV^2_α-regularization is deblurring a noisy image. For a blurring kernel $k \in L^1(\Omega_0)$ satisfying $\bar{k} = \int_{\Omega_0} k \text{d}x \neq 0$, $\Omega \subset \mathbb{R}^2$ a Lipschitz domain, $\Omega' \subset \mathbb{R}^2$ a domain with $\Omega' - \Omega_0 \subset \Omega$, $f \in \text{L}^2(\Omega')$ and the convolution operator

$$(Ku)(x) = \int_{\Omega_0} u(x-y)k(y)\text{d}y, \quad x \in \Omega',$$

existence of a solution for the problem

$$\min_{u \in \mathrm{L}^2(\Omega)} \frac{1}{2} \int_{\Omega'} |(Ku)(x) - f(x)|^2 \mathrm{d}x + \mathrm{TGV}_\alpha^2(u)$$

can be shown by Theorem 10. A saddle-point formulation can be numerically solved by, for instance, a primal-dual algorithm [3]. Figure 3 shows the effect of TGV_α^2 compared to TV for deblurring a sample image.

References

[1] K. Bredies, K. Kunisch, and T. Pock. *Total Generalized Variation*, SIAM J. on Imaging Sciences, **3** (2010), 492–526.

[2] K. Bredies and T. Valkonen. *Inverse problems with second-order total generalized variation constraints*, In 9th International Conference on Sampling Theory and Applications (SampTA 2011), 2011.

[3] A. Chambolle and T. Pock. *A first-order primal-dual algorithm for convex problems with applications to imaging*, Technical report, Graz University of Technology, 2010.

[4] ˆ@ˆina@Flickr. Alina's eye. Licenced under CreativeCommons-by-2.0, 2009. http://www.flickr.com/photos/angel_ina/3201337190/.

[5] F. Knoll, K. Bredies, T. Pock, and R. Stollberger. *Second order total generalized variation (TGV) for MRI*, Magnetic Resonance in Medicine, **65** (2011), 480–491.

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

[7] R. Temam. *Mathematical problems in plasticity*, Gauthier-Villars, (1985).

The structure of isometric maps and symmetries

Leonid Guibas

Approximate isometries or distance preserving transformations of surfaces model many realistic deformations of geometric objects, yet their mathematical structures are not well understood. In this talk we examine the intrinsic structure of isometric and near-isometric maps, starting with Killing (isometric) vector fields and then go on to examine isometric-preserving multiscale descriptors and their applications in feature detection and shape analysis, concluding with algorithms for partial and approximate isometry detection. The use of Laplace-Betrami operator and the associated heat Kernel play a fundamental role in these developments.

Shape Spaces for Computational Anatomy

Laurent Younes

A fundamental question in Computational Anatomy addresses how the shapes of organs (anatomy) are affected by disease. For example, most cognitive disorders result in selective atrophy of various structures in the brain. Similarly, heart diseases generally induce significant remodeling of the cardiac muscle. Describing how and where shape changes occur can provide clinicians with essential information on the nature of the disease.

The following two families of statistical problems in shape spaces are typically addressed in Computational anatomy:

- Problem 1: Assume that a population of N subjects is given, belonging to two (or more) groups. Each subject is represented by an image or shape, $S_1, \ldots, S_N$. The question is whether the shapes are group-dependent? If yes, can one describe where they mostly affected?
- Problem 2: Assume also a population of N subjects belonging to two (or more) groups, but let each subject be represented by a pair of images. Refer to them as baselines $(B_1, \ldots, B_N)$ and follow-ups $(F_1, \ldots, F_N)$. Question: is the deformation (or shape variation) between F and B group-dependent? If yes, where is it affected?

Solving these problems relies on an suitable representation of shapes, or, more generally, of deformable objects. This representation will in turn rely of a specific construction of a "Riemannian" shape space, $\mathcal{S}$. Given this, Problem 1 can be addressed as follows. Compute an average shape (or template), using, for example, a Karcher mean. Call it $\bar{S}$. Compute minimizing geodesics between $\bar{S}$ and each S_k so that

$$S_k = \exp_{\bar{S}}(D_k)$$

for some $D_k \in T_{\bar{S}}\mathcal{S}$, where exp refers to the Riemannian exponential (the endpoint at time $t = 1$ of the geodesic starting from $\bar{S}$ with velocity D_k). One can then apply "standard" statistical methods to $D_1, \ldots, D_N$, which are vectors belonging to the same vector space.

For Problem 2, a possible approach is as follows. Start by computing minimizing geodesics between each baseline, B_k, and the corresponding follow-up, F_k, so that

$$F_k = \exp_{B_k}(\Delta_k)$$

with $\Delta_k \in T_{S_k}\mathcal{S}$. One then transports $\Delta_1, \ldots, \Delta_N$ to a single tangent space, say $T_{\bar{B}}\mathcal{S}$, where $\bar{B}$ is an average baseline, using parallel transport along geodesics linking B_k and $\bar{B}$. This results in vectors $D_1, \ldots, D_N$ to which "standard" statistical methods can be applied.

It therefore remains to build this shape space, and make sure that geodesics and parallel transport can be computed on it. Our construction, which is based on the deformable template paradigm, works as follows. Let G be the group of diffeomorphisms, acting transitively on a differential manifold $\mathcal{S}$ (shapes), so that $\mathcal{S}$ is a single orbit: $\mathcal{S} = G \cdot S_0$ for some reference shape S_0. One can define a right-invariant Riemannian metric on G by first defining an inner product on its Lie algebra, denoted V, and then letting, for $\varphi \in G$:

$$\|v\|_\varphi = \|v \circ \varphi^{-1}\|_V.$$

One then considers the projection $\pi : G \to \mathcal{S}$ defined by $\pi(\varphi) = \varphi \cdot S_0$, and defines a metric on $\mathcal{S}$ that turns π into a *Riemannian submersion*, i.e., one defines

$$\|\xi\|_S = \|v^\xi\|_V$$

where v^ξ is the horizontal lift of ξ on the Lie algebra, defined by

$$\begin{cases} v^\xi \cdot S = \xi \\ v^\xi \perp \{w \in V : w \cdot S = 0\} \end{cases}$$

where $w \cdot S$ denotes the infinitesimal action of w on S.

With this construction, geodesics on $\mathcal{S}$ can be deduced from geodesics on G, in the sense that, if $S = \pi(\phi)$, then $\exp_S(\xi) = (\exp_{\mathrm{id}}(v^\xi)) \circ \phi$. The exponentials in the group are well-known since Arnold's seminal work, and are solution of the so-called EPDiff equation. Parallel transport equations can also be written explicitly, even if the analysis is somewhat more involved computationally. The interested reader can refer to the referenced textbooks and to their bibliography.

References

[1] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, Springer, 1978.
[2] U. Grenander and M.I. Miller, *Pattern theory: from representation to inference*, Oxford University Press, USA, 2007.
[3] D. D. Holm. *Geometric Mechanics*, Imperial College Press, 2008.
[4] J. E. Marsden and T. S. Ratiu. *Introduction to Mechanics and Symmetry*, Springer, 1999.
[5] L. Younes. *Shapes and diffeomorphisms (series: applied mathematical sciences)*, Springer, 2010.

A variational framework for exemplar-based image inpainting

Pablo Arias

(joint work with Gabriele Facciolo, Vicent Caselles, Guillermo Sapiro)

Non-local methods for image denoising and inpainting have gained considerable attention in recent years. This is in part due to their superior performance in textured images, a known weakness of purely local methods. Local methods on the other hand have demonstrated to be very appropriate for the recovering of geometric structures such as image edges. The synthesis of both types of methods is a trend in current research. Variational analysis in particular is an appropriate tool for a unified treatment of local and non-local methods. In this work we propose a general variational framework for non-local image inpainting, from which important and representative previous inpainting schemes can be derived, in addition to leading to novel ones.

Following the tradition of non-local denoising and regularization [1, 2] we encode the image redundancy and self-similarity (measured as patch similarity) as a non-local weight function $w : \widetilde{O} \times \widetilde{O}^c \to \mathbb{R}$, which serves as a fuzzy correspondence (see Figure 1). Additionally a correspondence map can also be obtained as a limit of our model. As a result, although the focus of this work lies on inpainting, the framework we are introducing can be adapted for its application to other contexts [4]. This allows us to provide intuitive interpretations of the energy and to relate existing models for non-local regularization with exemplar-based inpainting. We provide a detailed discussion of these relations and interpretations in the text.

The proposed formulation is rather general and different inpainting schemes can be derived naturally from it, via the selection of the appropriate patch similarity criterion. In this work we present four of them, corresponding to similarity criteria based on L^2- and L^1-norms between patches or their gradients. The use of the L^2-norm between patches is related to the inpainting methods of [3, 5]

Gradient-based methods, combine the exemplar-based interpolation with PDE-based diffusion schemes. This results in a smoother continuation of the information across the boundary and inside the inpainting domain, and in a better propagation of structures. Furthermore, the inclusion of gradients in the patch similarity criterion allows to handle additive brightness changes.

The results obtained yield a partition of the inpainting domain into arbitrarily shaped segments which show an exact copy (of image or gradient values) of some region in the hole's complement. Transitions between the copied segments take place in a band around the boundary between the segments. The width of the band is the size of the patch. The four inpainting schemes differ in the way this blending is done (and in the partition found). We provide a comprehensive empirical comparison on real and synthetic problems, showing the benefits and limitations of each variation of the proposed formalism.

FIGURE 1. **Inpainting problem.** Left: on a rectangular image domain Ω, missing data u in a region O has to be reconstructed using the available image $\hat{u}$ over $O^c := \Omega \setminus O$. The set of centers of incomplete patches is $\widetilde{O} := O + \Omega_p$, being the latter the patch domain. The center image shows a completion found with one of the one of the schemes derived from the framework. Right: the resulting completion is a patch-work built by copying arbitrarily shaped regions from O^c. The red curves show the boundaries between the copied regions.

REFERENCES

[1] S. P. Awate and R. T. Whitaker, *Higer-order image statistics for unsupervised, information-theoretic, adaptive image filtering for image restoration*, In Proc. of the IEEE Conf. on CVPR **28** (2005), 44–51.

[2] A. Buades, B. Coll, and J.-M. Morel, *A non local algorithm for image denoising*, In Proc. of the IEEE Conf. on CVPR **2** (2005), 60–65.

[3] N. Kawai, T. Sato, and N. Yokoya. *Image inpainting considering brightness change and spatial locality of textures and its evaluation*, In Ad. in Image and Video Tech., pages 271–82. Springer Berlin Heidelberg, 2009.

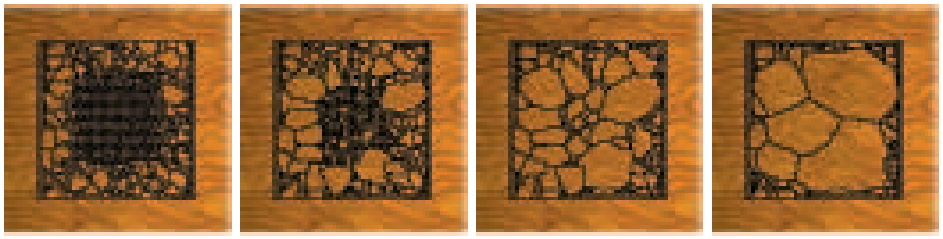

FIGURE 2. **Evolution of the minimization.** First 3 iterations and the final result (right) are shown for one of the proposed schemes. The black lines correspond to the boundaries of copied segments.

[4] G. Peyré, S. Bougleux, and L. D. Cohen. *Non-local regularization of inverse problems*, Preprint Hal-00419791 (2009).
[5] Y. Wexler, E. Shechtman, and M. Irani, *Space-time completion of video*, IEEE Trans. on PAMI **29** (2007), 463–476.

Spin Transformations of Discrete Surfaces

KEENAN CRANE
(joint work with Ulrich Pinkall, Peter Schröder)

Immersions of a surface M can be expressed via maps $f : M \to \mathrm{Im}\mathbb{H}$ into the imaginary part of the quaternions. Two immersions f and $\tilde{f}$ are called *spin equivalent* [1] if there exists some function $\lambda : M \to \mathbb{H}$ such that

$$d\tilde{f} = \bar{\lambda} df \lambda, \tag{1}$$

and it is clear that spin equivalence implies conformal equivalence since the induced Riemannian metrics are related by a positive scaling: $|d\tilde{f}|^2 = |\lambda|^4 |df|^2$. In digital geometry processing, preservation of conformal structure provides a valuable tool for maintining signal integrity (e.g., aspect ratio of mesh elements). It is therefore natural to seek an analog of spin transformations in the discrete setting.

Kamberov, Pedit, and Pinkall observed that, as a condition on λ, equation (1) is equivalent to

$$0 = d(\bar{\lambda} df \lambda) = d\bar{\lambda} \wedge df \lambda - \bar{\lambda} df \wedge d\lambda = -2\mathrm{Im}(\bar{\lambda} df \wedge d\lambda)$$

whenever M is simply-connected [1]. In other words, $\bar{\lambda} df \wedge d\lambda$ must be equal to some real function ρ, which leads to the integrability condition $\bar{\lambda} df \wedge d\lambda = \rho |df|^2$. The function ρ represents the change in *mean curvature half-density* between an initial immersion f and its spin transform $\tilde{f}$, i.e., $\tilde{H}|d\tilde{f}| = H|df| + \rho|df|$.

For discretization we consider the alternate formulation

$$D\lambda = \rho\lambda \tag{2}$$

where D is the self-adjoint elliptic operator given by

$$D\lambda := -\frac{df \wedge d\lambda}{|df|^2},$$

which we call the *quaternionic Dirac operator*, since it is locally equivalent to the standard Dirac operator for a spin-1/2 particle in the plane [2]. We can now specify a surface (up to isometry and uniform scaling) by prescribing a change $\rho|df|$ in mean curvature half-density and solving the eigenvalue problem

$$(D - \rho)\lambda = \gamma\lambda \tag{3}$$

for the smallest eigenvalue γ. The resulting pair $(\lambda, \rho+\gamma)$ satisfies our integrability condition (2), with a small constant shift in the prescribed curvature change.

Consider any surface M composed of piecewise smooth faces σ with linear edges e_{ij}. If λ is also linear along edges, then we have

$$\int_\sigma D\lambda|df|^2 = \int_\sigma d(df\lambda) = \sum_{e_{ij}\in\partial\sigma} \int_{e_{ij}} df\lambda = \sum_{e_{ij}\in\partial\sigma} (f_j - f_i)\frac{\lambda_i + \lambda_j}{2},$$

which gives us one way to discretize D. In particular, when M is simplicial we end up with the sparse linear operator

$$\mathsf{D}_{ij} = -\frac{1}{2A_i} e_j$$

for each edge e_j of each 2-simplex σ_i. In a similar fashion, the operator ρ can be discretized as $\mathsf{R} = \mathsf{PB}$ where $\mathsf{B}_{ij} = 1/3$ for each vertex v_j of face σ_i and $\mathsf{P}_{ii} = \rho_i$ for some prescribed value ρ_i on each face. The resulting discrete operator $\mathsf{A} = \mathsf{D} - \mathsf{R}$ is rectangular, and so there are several options for formulating the eigenvalue problem (3). The system $\mathsf{B}^*\mathsf{A} = \mathsf{B}^*\mathsf{B}$ obtained by averaging values from faces back to vertices is problematic because this local averaging artificially places high-frequency modes in the null space of the system, distorting solutions. An alternative is to note that the generalized eigenvalue problem $D^2\lambda = \gamma D\lambda$ shares solutions with our original problem (3), and can be discretized as $\mathsf{D}^*\mathsf{D}\lambda = \gamma\mathsf{B}^*\mathsf{D}\lambda$. In practice the spectrum of eigenvalues computed via this discretization closely matches the spectrum of the smooth operator, and solutions appear to converge linearly with respect to the mean edge length of simplices [2].

One might also consider discretizations that capture the essential structure of the smooth theory. For instance, Springborn, Schröder, and Pinkall consider equivalence classes of discrete immersions where edge lengths are related by positive scale factors at vertices [3]. This notion of *discrete conformal equivalence* leads to a theory that closely mimics the smooth setting [4], but in terms of constructive algorithms is limited to immersions into the plane or sphere. Therefore, a natural question to ask is whether this same notion of conformal equivalence can be extended to all spin transformations.

References

[1] G. Kamberov, F. Pedit, U. Pinkall *Bonnet Pairs and Isothermic Surfaces*, Duke Mathematical Journal **92, 3** (1998), 637–644.

[2] K. Crane, U. Pinkall, P. Schröder *Spin Transformations of Discrete Surfaces*, submitted.
[3] B. Springborn, P. Schröder, U. Pinkall *Conformal Equivalence of Triangle Meshes*, ACM Transactions on Graphics **27, 3** (2008), 77:1–77:11.
[4] A. Bobenko, U. Pinkall, B. Springborn *Discrete Conformal Maps and Ideal Hyperbolic Polyhedra*, arXiv:1005.2698v1

Sparse and Low Rank Recovery

Holger Rauhut

Compressive sensing (sparse recovery) is a new area in mathematical image and signal processing that predicts that sparse signals can be recovered from what was previously believed to be highly incomplete measurement [3, 5, 7, 12]. Recently, the ideas of this field have been extended to the recovery of low rank matrices from undersampled information [6, 8]; most notably to the matrix completion problem [4, 11].

A vector $x \in \mathbb{C}^N$ is called s-sparse if $\|x\|_0 := \#\{\ell, x_\ell \neq 0\} \leq s$. In practice, vectors will usually not be exactly s-sparse, but can be well-approximated by a sparse vector. In order to quantify this notion one introduces the best s-term approximation error in ℓ_p by $\sigma_s(x)_p := \inf_{z\in\mathbb{C}, \|z\|_0\leq s} \|x-z\|_p$. Informally, a vector x is called compressible if $\sigma_s(x)_p$ decays quickly in s.

The basic task of compressive sensing is to recover a sparse (or compressible) vector from undersampled linear information, that is, from

$$y = Ax \in \mathbb{C}^m, \quad A \in \mathbb{C}^{m\times N}$$

where $m \ll N$ and A is a suitable matrix, the so called measurement matrix. Linear algebra tells us that the above system has infinitely many solutions even if A has full rank, so traditional wisdom is that it is impossible to single out the original x. However, under the additional assumption that x is sparse reconstruction becomes indeed possible, as will be outlined.

The first approach for the reconstruction of x that probably comes to mind, is to consider the solution of the ℓ_0-minimization problem

$$\min_{z\in\mathbb{C}^N} \|z\|_0 \quad \text{subject to } Az = y\,.$$

Unfortunately, this optimization problem is NP-hard, so that tractable alternatives have to be found. A by-now well-understood approach is ℓ_1-minimization,

$$\min_{z\in\mathbb{C}^N} \|z\|_1 \quad \text{subject to } Az = Ax\,. \tag{1}$$

This convex optimization problem can be solved with efficient methods. Greedy algorithms, and iterative methods for the compressive sensing problem have been introduced as well. A very useful tool for the analysis of ℓ_1-minimization (as well as for some of the other algorithms) is the restricted isometry property (RIP). The restricted isometry constant δ_s is the smallest number such that

$$(1-\delta_s)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1+\delta_s)\|x\|_2^2 \quad \text{for all } s\text{-sparse } x.$$

It can be shown [3, 1, 9, 7] that if $\delta_{2s} \leq 0.46$ then every s-sparse vector is the unique solution to (1). Moreover, stability with respect to passing from sparse to compressible signals as well as to noise on the measurements holds as well.

Surprisingly, it is extremely hard and so far open to come up with deterministic constructions of matrices A which have small δ_{2s} with an optimal relation of s, m and N. A common way out, is to consider random matrices. A very popular choice of random matrices are Bernoulli and Gaussian random matrices, that is, all entries are independent and take either the values $\pm 1/\sqrt{m}$ with equal probability (Bernoulli), or are $\mathcal{N}(0, 1/m)$-normal distributed (Gaussian). For such matrices, $\delta_s \leq \delta$ with probability at least $1-\varepsilon$ provided $m \geq C\delta^{-2}(s \log(eN/s) + \log(2\epsilon^{-1}))$, where C is a universal constant [3, 7, 12]. In particular, we can recover every s-sparse vector from $y = Ax$ with high probability using ℓ_1-minimization whenever

$$m \geq C' s \log(eN/s).$$

This bound is optimal as follows from lower bounds of Gelfand widths of ℓ_p-balls with $0 < p \leq 1$ [10].

These ideas can be extended to the recovery of low rank matrices $X \in \mathbb{C}^{n \times q}$. Assume we take measurements $y = \mathcal{A}(X)$, where $\mathcal{A} : \mathbb{C}^{n \times q} \to \mathbb{C}^m$ is a linear map. Again the task is to reconstruct X from y in the interesting regime $m \ll nq$. The naive approach of solving the rank minimization problem

$$\min_{Z \in \mathbb{C}^{n \times q}} \operatorname{rank}(Z) \quad \text{subject to } \mathcal{A}(Z) = \mathcal{A}(X)$$

is again NP-hard. A tractable alternative is nuclear norm minimization,

$$\min_{Z \in \mathbb{C}^{n \times q}} \|Z\|_* \quad \text{subject to } \mathcal{A}(Z) = \mathcal{A}(X)\ .$$

Here, the nuclear norm is defined by $\|Z\|_* = \sum_\ell \sigma_\ell(Z)$, where the $\sigma_\ell(Z)$ are the singular values of Z. (Note that $\|\cdot\|_*$ is the dual of the operator norm.) Similarly as above, one introduces the restricted isometry constant of $\mathcal{A}$ as the smallest constant δ_r such that

$$(1-\delta_r)\|X\|_F^2 \leq \|\mathcal{A}(X)\|_2^2 \leq (1+\delta_r)\|X\|_F^2 \quad \text{for all } X \text{ with } \operatorname{rank}(X) \leq r\ ,$$

see [2, 6]. Here $\|X\|_F$ denotes the Frobenius norm of X. If $\delta_{4r} \leq \sqrt{2}-1$ then nuclear norm minimization recovers every X of rank at most r from $y = \mathcal{A}(X)$, see [6, 8]. Here again, optimal measurement maps are provided by Gaussian maps, that is, all coefficients $\mathcal{A}_{i,j,k}$ in the representation $\mathcal{A}(X)_i = \sum_{j,k} \mathcal{A}_{i,j,k} X_{j,k}$ are independent and $\mathcal{N}(0, 1/m)$-distributed. For such maps, $\delta_r \leq \delta$ with probability at least $1-\varepsilon$ provided

$$m \geq C\delta^{-2} r(n+q)\ .$$

Note that the term $r(n+q)$ is essentially the number of parameters required to describe a rank-r matrix in dimension $n \times q$ (up to constants).

For practical purposes it is essential to have structured measurement matrices and measurement maps. In particular, structure can speed up matrix vector multiplication, which is essential for recovery algorithms to work for large scale

problems. In the low rank recovery setup, special structure is provided in the matrix completion problem, where one samples entries of the matrix to be recovered, see [11, 4] for details.

For the compressive sensing problems several types of structured random matrices have been studied so far [12, 14, 15, 13]; in particular, sampling matrices arising in random sampling of bounded orthonormal systems and partial random circulant matrices [13]. A particular instance of the former arises when taking m randomly selected rows of the $N \times N$ discrete Fourier matrix. Then $\delta_s \leq \delta$ with high probability provided

$$m \geq C\delta^{-2} s \log^3(s) \log(N) \ .$$

We refer to [12] for further details.

References

[1] E. Candès, *The restricted isometry property and its implications for compressed sensing*, C. R. Acad. Sci. Paris S'er. I Math. **346** (2008), 589–592.

[2] E. Candès and Y. Plan, *Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements*, IEEE Trans. Information Theory, to appear.

[3] E. Candès and T. Tao, *Near optimal signal recovery from random projections: universal encoding strategies?*, IEEE Trans. Inf. Theory **52** (2006), 5406–5425.

[4] E.J. Candès and T. Tao, *The power of matrix completion: near-optimal convex relaxation*, IEEE Trans. Information Theory **56** (2010), 2053–2080.

[5] D. L. Donoho. *Compressed sensing*, IEEE Trans. Inform. Theory **52** (2006) 1289–1306.

[6] B. Recht, M. Fazel, and P. Parrilo. *Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization*, SIAM Rev. **52** (2010), 471–501.

[7] M. Fornasier, H. Rauhut, *Compressive Sensing*, In: O. Scherzer (ed.), Handbook of Mathematical Methods in Imaging, pages 187-228. Springer, 2011.

[8] M. Fornasier, H. Rauhut, and R. Ward. *Low-rank matrix recovery via iteratively reweighted least squares minimization*, Preprint, 2010.

[9] S. Foucart. *A note on guaranteed sparse recovery via ℓ_1-minimization*, Appl. Comput. Harmon. Anal. **29** (2010), 97–103.

[10] S. Foucart, A. Pajor, H. Rauhut, and T. Ullrich. *The Gelfand widths of ℓ_p-balls for $0 < p \leq 1$*, J. Complexity **26** (2010), 629–640.

[11] D. Gross. *Recovering low-rank matrices from few coefficients in any basis*, preprint, 2009.

[12] H. Rauhut, *Compressive sensing and structured random matrices*, In: M. Fornasier (ed.), Theoretical Foundations and Numerical Methods for Sparse Recovery, 1–92, de-Gruyter, 2010.

[13] H. Rauhut, J. Romberg, J. Tropp, *Restricted isometries for partial random circulant matrices*, Preprint, 2010.

[14] H. Rauhut, R. Ward, *Sparse Legendre expansions via ℓ_1-minimization*, preprint, 2010.

[15] M. Rudelson and R. Vershynin, *On sparse reconstruction from Fourier and Gaussian measurements*, Comm. Pure Appl. Math. **61** (2008), 1025–1045.

Shape Processing via Linear Programming Relaxation

Daniel Cremers

1. Toward Optimal Solutions to Shape Analysis

A large number of shape analysis and shape processing tasks have been formulated by means of energy minimization where solutions to the processing task correspond to configurations of minimal energy. Particular examples that we focus on in this line of work are:

- Curvature-regularized image segmentation [6],
- Curvature-regularized image inpainting [6],
- Willmore-type surfaces [7],
- Elastic matching of surfaces in 3D space [8].

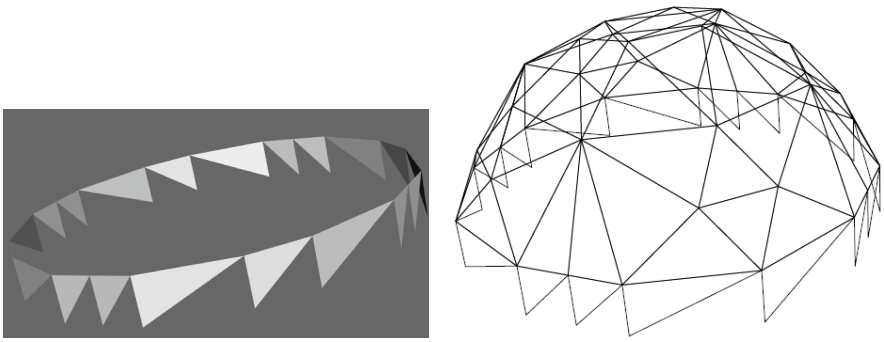

Figure 1. Specified boundary condition and discrete Willmore surface computed via linear programming relaxation. Source:[7].

Unfortunately most existing approaches to shape analysis and shape processing tasks focus on gradient flow solutions where respective solutions are computed by means of gradient descent. Since the underlying energies are generally not convex, solutions will invariably depend on the choice of initialization and may be aribtrarily bad – this dependency on the initialization becomes particularly undesirable in image processing tasks, where for example curvature regularity is added to some data term that drives respective segmentations to strong intensity edges: In noisy and cluttered images there are numerous spurious edges which all attract a local curve evolution scheme. As a consequence it would be highly desirable to compute solutions which do not depend on an initialization and which come with some optimality or quality guarantee.

For a limited number of shape optimization and shape analysis problems respective energies can be minimized globally in computation times that are polynomial in the number of basic elements discretizing the respective solution space. For example, the elastic matching of planar shapes can be solved by means of *dynamic*

FIGURE 2. User-labeled input image and globally optimal curvature regularized segmentation computed via linear programming relaxation. Source:[6].

programming techniques, where the optimal matching is equivalent to a monotonous cyclic path on a torus [3, 4]. Similarly length-regularized image segmentation can be solved by means of minimal s-t-cuts through respective graphs [2, 1, 5].

Unfortunately these efficient solutions often do not generalize well: The elastic matching of surfaces in three-dimensional space no longer corresponds to a shortest path on a torus, but rather to a codimension-two surface in four dimensional space. Similarly the segmentation with curvature regularization no longer corresponds to a graph cut problem.

Nevertheless, such problems can be addressed in a unifying framework. The key insight is that both the shortest path and the graph cut problem are specific instances of integer linear programs. While the shortest path formulation and the graph cut formulation do not generalize to the respective higher-dimensional problems, the paradigm of the integer linear program does.

2. Shape Inference by Linear Programming

A central aspect in efficiently computing solutions to shape optimization challenges lies in selecting the appropriate mathematical representation of the solutions. In this line of work, we will consistently revert to implicit representations. As a consequence, for problems like segmentation or shape matching curves and shapes no longer evolve locally according to some gradient direction starting from some initialization, but they rather emerge in specific locations independent of initialization.

In reformulating the above problems by means of integer linear programs there are two key ideas:

- Firstly, we represent all feasible solutions by a (potentially huge) set of N binary variables $x \in \{0,1\}^N$. For the case of curvature-regularized boundaries these variables indicate whether a set of consistently oriented adjacent edges is part of the segmentation or not. To additionally represent cost terms encoding integrals over the inside or outside area separated by the boundary, we can introduce additional binary variables for each pixel of the image. For the case of Willmore-type surfaces we can represent the

set of feasible solutions in terms of pairs of adjacent triangles, where each pair is either part of the solution or not. And for elastically matching two closed surfaces X and Y by means of a diffeomorphism $f : X \to Y$ we can discretize the graph

$$\Gamma = \{(x, f(x)) \mid x \in X\} \subset X \times Y, \tag{1}$$

which is a two-dimensional closed surface in four dimensions. With respect to these discrete representations of solutions spaces, respective energies for shape inference and matching amount to integer linear programs.

- Secondly, the consistency of solutions must be enforced by additional linear constraints. These assure that the selected basic elements x_i actually do form a coherent surface. As a consequence, the overall shape optimization problem amounts to an integer linear program of the form

$$\begin{aligned} &\min_{x \in \{0,1\}^N} w^\top x, \\ &s.\ t.\ Ax = b \end{aligned} \tag{2}$$

Upon relaxation, i.e. neglecting the integer constraint and allowing intermediate values in the interval $[0, 1]$ for all variables x_i, the resulting linear program can be solved in polynomial time. Subsequent binarization leads to solutions which are within an energetic bound from the optimum.

Figure 1, 2 and 3 show image segmentations with curvature regularity, discrete Willmore surfaces and an elastic surface matching computed in this manner. For further details we refer the reader to [6, 7, 8].

Acknowledgment

The above line of works was initiated by two of my former students Thomas Schoenemann and Frank R. Schmidt and worked out in collaboration with Simon Masnou, Fredrik Kahl, Thomas Windheuser and Ulrich Schlickewei.

References

[1] Y. Boykov and V. Kolmogorov, *An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision*, IEEE Trans. on Patt. Anal. and Mach. Intell. **26** (2004), 1124–1137.

[2] D. M. Greig, B. T. Porteous, and A. H. Seheult, *Exact maximum a posteriori estimation for binary images*, J. Roy. Statist. Soc., Ser. B. **51** (1989), 271–279.

[3] M. Maes, *Polygonal shape recognition using string-matching techniques*, Pattern Recognition **24** (1991), 433–440.

[4] F. R. Schmidt, Dirk Farin, and D. Cremers, *Fast matching of planar shapes in sub-cubic runtime*, In IEEE Int. Conf. on Computer Vision, Rio de Janeiro, October 2007.

[5] F. R. Schmidt, E. Toeppe, and D. Cremers, *Efficient planar graph cuts with applications in computer vision*, In Int. Conf. on Computer Vision and Pattern Recognition, 2009.

[6] T. Schoenemann, F. Kahl, and D. Cremers, *Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation*, In IEEE Int. Conf. on Computer Vision, Kyoto, 2009.

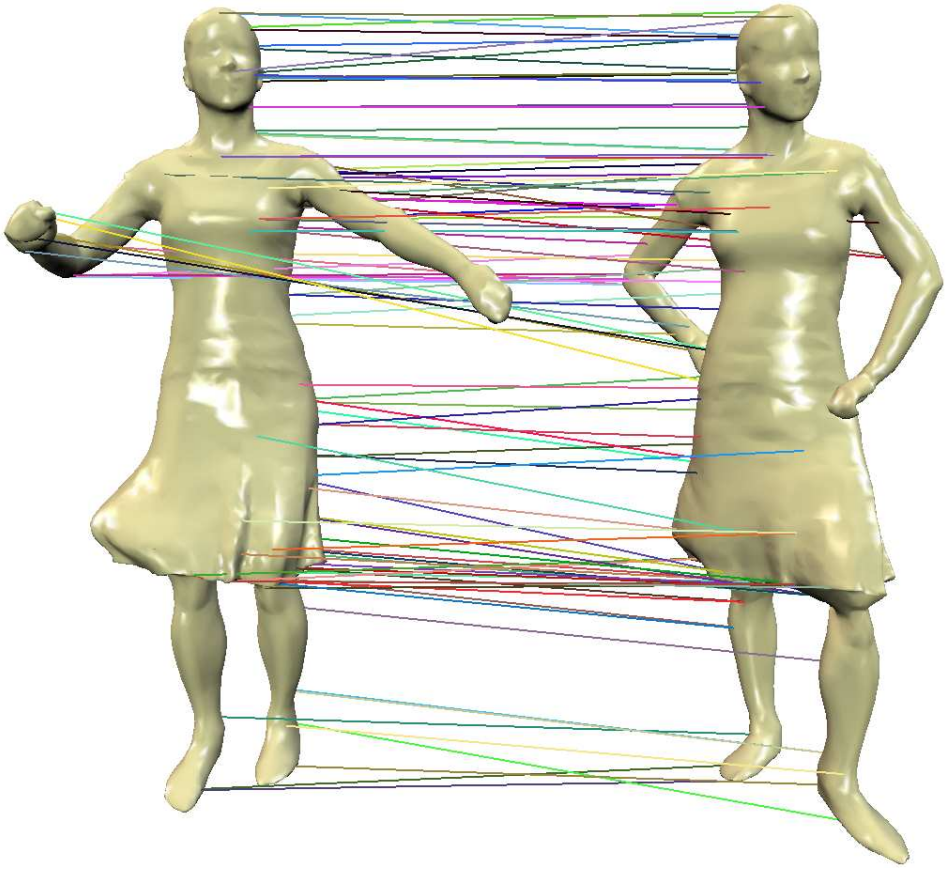

FIGURE 3. Elastic matching of shapes in 3D by means of linear programming relaxation. Source:[8].

[7] T. Schoenemann, S. Masnou, and D. Cremers, *On a linear programming approach to the discrete Willmore boundary value problem and generalizations*, In Curves and Surfaces, 2011.

[8] T. Windheuser, U. Schlickewei, F. R. Schmidt, and D. Cremers, *Minimal surfaces for 3d shape matching: A linear programming solution*, Technical Report TUM-I1101, Technische Universitaet Muenchen, January 2011.

Smooth Weights for Real-Time Deformation

Olga Sorkine
(joint work with Alec Jacobson, Ilya Baran and Jovan Popovic)

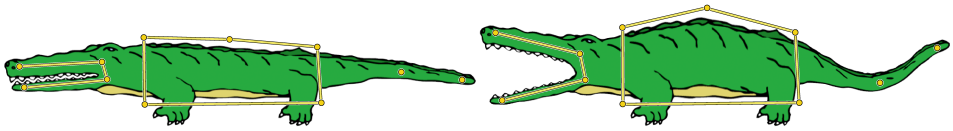

Figure 1. Bounded biharmonic blending supports points, bones, and cages arranged in an arbitrary configuration. This versatility makes it possible to choose the right tool for each subtask: bones to control rigid parts, cages to enlarge areas and exert precise control, and points to transform flexible parts.

Object deformation with linear blending dominates practical use as the fastest approach for transforming raster images, vector graphics, geometric models and animated characters. Given a shape $\Omega \in \mathbb{R}^2$ or $\mathbb{R}^3$ and control handles $H_1, \ldots, H_m \subset \Omega$, where each handle is a point, a straight line segment (i.e., skeletal bone) or a vertex of a polygonal cage, transformation by linear blending is formulated as

$$\forall \mathbf{x} \in \Omega,\ \mathbf{x}' = \sum_{j=1}^{m} w_j(\mathbf{x})\, T_j \mathbf{x}, \tag{1}$$

where T_j is an affine transformation and $w_j : \Omega \to \mathbb{R}$ is a weight function associated with control handle H_j. See Fig. 1 for an example of all three control handle types. The user is only required to supply the transformations T_j, and the whole shape Ω is deformed by the above simple equation, which is embarrassingly parallelizable and is routinely implemented in graphics hardware. An alternative to linear blending of affine transformation is, e.g., Dual Quaternions [4], which provide more intuitive interpolation of rotations at a small computational cost.

Unfortunately, blending schemes for skeletons or cages are not always easy to use because they may require manual weight painting or modeling closed polyherdal cages around objects. Our goal is to make the design and control of deformations simpler by allowing the user to work freely with the most convenient combination of handle types. We develop blending weights w_j that produce smooth and intuitive deformations for points, bones and cages of arbitrary topology.

The properties of the weight functions w_j determine the quality and the intuitiveness of the resulting deformation. Here are the properties we would like our weight functions to satisfy:

Smoothness: Lack of smoothness at the handles causes visible artifacts in 2D textured shapes and prevents placing handles directly *on* 3D shapes. Many popular weight definition schemes, such as Harmonic Coordinates [3], suffer from

this problem. We would like w_j to be at least C^1 at the control handles and C^∞ everywhere else.

Non-negativity: Negative weights lead to unintuitive handle influences, because regions of the shape with negative weights move in the opposite direction to the prescribed transformation. For instance, biharmonic functions [2], although smooth, may attain negative values.

Shape-awareness: Informally, shape-awareness implies intuitive correspondence between the handles and the domain Ω. The influence of the handles should conform to the features of the shape and fall off with shape-aware (as opposed to Euclidean) distance. The best shape-aware behavior one can hope for is when the weights w_j depend on the metric of Ω alone and do not change for any possible embedding of Ω.

Partition of unity: This classical property (also seen in, e.g., Bézier curves or NURBS) ensures that if the same transformation T is applied to all handles, the entire object will be transformed by T.

Locality and sparsity: Each handle should mainly control a shape region or feature in its vicinity, and each point in Ω should be influenced only by a few closest handles. Specifically, if any locally shortest path (in a shape-aware sense) from a point $\mathbf{x}_0$ to H_j passes near some other handle, then H_j is "occluded" from $\mathbf{x}_0$ and $w_j(\mathbf{x}_0)$ should be zero.

No local maxima: Each w_j should attain its global maximum (i.e., value of 1) on H_j and should have no other local maxima. This property provides monotonic decay of a handle's influence and guarantees that no unexpected influences occur away from the handle.

We propose to formulate the weights satisfying the properties above via a variational problem, namely as minimizers of the Laplacian energy:

$$\underset{w_j,\ j=1,\dots,m}{\arg\min} \sum_{j=1}^{m} \frac{1}{2} \int_\Omega \|\Delta w_j\|^2 dV \tag{2}$$

$$\text{subject to:} \quad w_j|_{H_k} = \delta_{jk} \tag{3}$$

$$w_j|_F \text{ is linear} \qquad \forall F \in \mathcal{F}_\mathcal{C} \tag{4}$$

$$\sum_{j=1}^{m} w_j(\mathbf{x}) = 1 \qquad \forall \mathbf{x} \in \Omega \tag{5}$$

$$0 \le w_j(\mathbf{x}) \le 1, \quad j = 1, \dots, m, \qquad \forall \mathbf{x} \in \Omega, \tag{6}$$

where $\mathcal{F}_\mathcal{C}$ is the set of all cage faces and δ_{jk} is Kronecker's delta. The weights can be found by discretizing the problem using, e.g., linear finite elements and employing a fast QP solver such as Mosek [1].

The above weights, which we call bounded biharmonic weights, spread the influences of the controls in a shape-aware and localized manner, even for objects with complex and concave boundaries, as can be seen in Fig. 2. The weights are C^1 smooth at the controls and C^∞ everywhere else, such that point controls can be placed directly onto the shape. The only exception to the smoothness property are

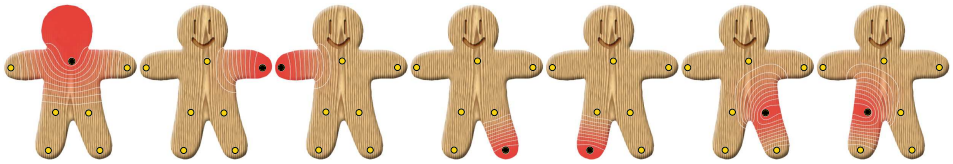

FIGURE 2. Bounded biharmonic weights. The blending weight intensity for each handle is shown in red with white isolines. Each handle has the maximum effect on its immediate region and its influence disappears in distant parts of the object.

skeletal joints where two or more bones meet, and cage vertices: in the former case, since joints are never "torn apart", the value of the weights at the joint actually does not matter; in the latter case, since cages are typically placed outside of the actual shape being deformed, C^0 smoothness there does not pose a problem visually. The properties of non-negativity, partition of unity and shape-awareness are satisfied by construction (the latter thanks to the fact that the Laplace operator solely depends on the metric on Ω).

The properties of locality, sparsity and lack of local maxima have been experimentally observed. We are currently looking into ways to analyze the shape of the active sets in order to formally prove or disprove these properties and understand the behavior better. The MFO Workshop "Trends in Mathematical Imaging and Surface Processing" has been extremely instrumental to the progress on this problem, facilitating our contact with math researchers working on similar variational problems. We are currently investigating notions developed for general obstacle problems [5] to help the analysis.

REFERENCES

[1] E. D. Andersen and K. D. Andersen, *The* MOSEK *interior point optimizer for linear programming: an implementation of the homogeneous algorithm*, In High Performance Optimization, pages 197–232. Kluwer Academic Publishers, 2000.

[2] M. Botsch and L. Kobbelt, *An intuitive framework for real-time freeform modeling*, ACM Trans. Graph. **23** (2004), 630–634.

[3] P. Joshi, M. Meyer, T. DeRose, B. Green, and T. Sanocki. *Harmonic coordinates for character articulation*, ACM Trans. Graph. **26** (2007), 71:1–71:9.

[4] L. Kavan, S. Collins, J. Zara, and C. O'Sullivan. *Geometric skinning with approximate dual quaternion blending*, ACM Trans. Graph. **27** (2008), 105:1–105:23.

[5] G. M. Troianiello, *Elliptic Differential Equations and Obstacle Problems*, Springer, 1987.

Scale space for point clouds and applications

JULIE DIGNE

Processing high precision point clouds surfaces is a very challenging problem. Indeed, standard methods are not designed to preserve high frequencies and details. Thanks to the definition of a new scale space for point clouds, methods can be

adapted so that they preserve the high precision while still extracting robust information.

Interestingly, several methods have been proposed to compute the differential properties of a surface by computing covariance matrices of local neighborhoods. We analyzed mathematically those methods in a unified framework by computing their asymptotic form when the size of the neighborhood tends to zero. All considered local schemes are of two kinds: either they perform a polynomial local regression, or they compute directly local moments. But the polynomial regression of order 1 is demonstrated to play a special role, because its iterations yield a *scale space*.

Using this scale space formulation, a lot of problems arising during the processing of raw point clouds can be solved. In particular, we developed a high precision meshing method that allows for the exact visualization of raw point clouds. Using the same scale space framework, we are able to suppress aliasing-like artifacts that are due to the superposition of scans. Merging of low frequencies deduced from the scale space and re-adding high frequencies yields the final shape with all details and without the artifacts.

Generalized Reflexion Operator for multi-impact

Etienne Vouga

To simulate multiple simultaneous impacts, one of two existing algorithms is typically used: pairwise propagation (Gauss-Seidel), or a method based on linear complementarity. As can be seen by examining the model systems-Newton's Cradle and Bernoulli problem- each method has its advantages and pitfalls. A new approach, inspired by the propagation of shock waves through elastic media, is the first to correctly simulate both of these problems, by guaranteeing symmetry preservation, energy conservation, and no artificial sticking.

On a first-order primal-dual algorithm

Thomas Pock
(joint work with Antonin Chambolle)

In this talk, we have discussed new results for a first-order primal-dual algorithm for non-smooth convex optimization problems with known saddle-point structure. The algorithm has been recently proposed in [3] for computing solutions of a convex formulation of the Mumford-Shah functional. In subsequent work, we have shown in [1] that the algorithm converges to a saddle-point with rate $O(1/N)$ in finite dimensions for the complete class of problems. Furthermore, the algorithm yields improved rates on problems with some degree of smoothness. In particular it achieves $O(1/N^2)$ convergence on problems, where the primal or the dual objective is uniformly convex, and yields linear convergence, i.e. $O(\omega^N)$ for some $\omega \in (0, 1)$, on smooth problems.

We have also discussed a simple pre-conditioning technique to additionally accelerate the convergence of the algorithm. We proposed a simple and easy to compute class of pre-conditioners for which convergence of the algorithm is guaranteed while keeping the computational complexity of the iterations unchanged. As a by-product, we can show that for a certain instance of the pre-conditioning, the proposed algorithm is equivalent to the old alternating step method for monotropic programming [2]. We demonstrate the improved performance of the algorithm by applying it to standard linear programming test problems and a few standard computer vision problems such as image restoration, inpainting and optical flow.

References

[1] A. Chambolle and T. Pock. *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision, 2010.
[2] J. Eckstein. *Splitting Methods for Monotone Operators with Application to Parallel Optimization*, PhD thesis, MIT, 1989.
[3] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. *An algorithm for minimizing the Mumford-Shah functional*, In International Conference on Computer Vision (ICCV), 2009.

Collaborative source separation and identification in images and signals via hierarchical sparse modeling

Pablo Sprechman

The talk presents a collaborative structured sparse model, C-Hilasso, that aims to add stability and prior information to the sparse representation of mixed signals. The application of C-Hilasso to source separation and identification problems will be discussed. The method that we propose goes as follows. First we build a structured dictionary to describe mixed signals by concatenating a set of subdictionaries, each of them learned to sparsely model one of a set of possible classes then. The coding of the mixed signal is performed by efficiently solving a convex optimization problem that combines standard sparsity with group sparsity. The present sources are identified by looking at the subdictionary selected in the coding. The collaborative filtering in C-Hilasso takes profit of the tenmporal/spatial redundancy present in the signals and shows to be critical to further stabilize the sparse representation.

Circular Arc Structures

Johannes Wallner
(joint work with Pengbo Bo, Helmut Pottmann, Martin Kilian, and Wenping Wang)

This research has been motivated by applications in freeform architecture (see e.g. the survey paper [4]). An important guiding principle in computational methods in this area is the balance between cost efficiency on the one hand, and adherence to the design intent on the other. Key issues are the simplicity of supporting

and connecting elements as well as repetition of costly parts [3]. This leads us to consider, in [1], the so-called circular arc structures as a means to faithfully realize freeform designs without giving up smooth appearance. In contrast to non-smooth meshes with straight edges where geometric complexity is concentrated in the nodes, we stay with smooth surfaces and rather distribute complexity in a uniform way by allowing edges in the shape of circular arcs. We study remarkable special cases of circular arc structures which possess simple supporting elements and whose faces may be realized by Dupin cyclides [2].

Definition. We consider mesh combinatorics (V, E, F) which are either regular triangular, or regular quadrangular, or regular hexagonal, with vertices $\{v_i\}_{i\in V}$ in Euclidean space, and circular arcs $\{e_{ij}\}_{(ij)\in E}$ as edges. This constitutes a circular arc structure, if the following geometric conditions are fulfilled:

(vii) For all $(i, j) \in E$ we have $v_i, v_j \in e_{ij}$ (i.e., edges connect adjacent vertices);
(viii) the tangent vectors of edges adjacent to v_j span a plane and possess a common normal vector n_j;
(ix) the configuration of these tangent vectors is the same for all vertices, up to Euclidean congruences;
(x) if for some edge (ij) the Euclidean reflection in the edge bisector $(v_i - v_j)^\perp$ exchanges normal vectors n_i, n_j, then this edge is called symmetric.

Without loss of much generality we restrict ourselves to the case that the angles between arcs e_{ij} emanating from vertex v_i assume the values $\frac{2\pi}{3}$, $\frac{2\pi}{3}$, $\frac{2\pi}{3}$ (in the hexagonal case), or α, $\pi - \alpha$ α, $\pi - \alpha$ (in the quadrilateral case), or $\frac{\pi}{3}, \ldots, \frac{\pi}{3}$ (triangular case).

Returning to the applications in freeform architecture mentioned in the introduction, we see that by approximating a smooth shape by a circular arc structure we may 'rationalize' it (i.e., decompose it into buildable parts) in a way which retains smoothness but uses the simplest edge elements possible, and and also leads to congruent nodes. This approximation task can be performed by introducing vertex positions and edges' tangent vectors as variables, and building up an appropriate nonlinear least squares optimization problem. Its initialization and solution (e.g. by means of level set methods) is discussed in [1].

It is elementary that edges which are symmetric are contained in a right circular cone: This fact is again useful in applications, since it can be directly used for supporting elements of simple shape. In case of quadrilateral combinatorics, if all boundary edges of a face are symmetric, then it is known that there is a unique Dupin cyclide such that the edges are principal curvature lines [2]. This property can be used to convert a circular arc structure into a smooth union of so-called principal patches, each of which is taken from a cyclide. Such geometric objects are the topic of [2].

From the viewpoint of discrete differential geometry, the circular arc structures of hexagonal combinatorics possess nice properties, since in this case there exists a

C^2 surface containing the union $\bigcup_{(ij)\in E} e_{ij}$ of edges (this is because the three normal curvatures in each vertex uniquely determine the second fundamental form). We have the following interesting feature:

Prop. *For any curvature-continuous surface which contains a hexagonal circular arc structure, the mean curvature in the vertex v_i equals*

$$H_i = \frac{1}{k} \sum_{j:(ij)\in E} 2\Big\langle \frac{v_j - v_i}{\|v_j - v_i\|^2}, n_i \Big\rangle \quad (k = 3). \tag{1}$$

The proof uses a geometric trick: Each arc e_{ij} is contained in its Meusnier sphere Σ_{ij}; and inversion ι_i with center v_i maps Σ_{ij} to a plane whose distance from the tangent plane $v_i + n_i^\perp$ equals half the normal curvature. Now all we have to do is read off that distance via inversion of v_j which leads to $\langle n_i, \iota_i(v_j) - v_i\rangle$, and to compute an average of the normal curvatures obtained in this way. □

Def. If a vertex v_i is adjacent to k edges which enclose successive angles of $\frac{2\pi}{k}$, then the mean curvature in this vertex is given by Equation (1).

Prop. *The mean curvature H_i as defined by Equation* (1) *for $k \geq 3$ edges is a Möbius invariant if it is encoded as the inverse radius of a sphere which touches the edges in the vertex v_i.*

The proof uses the fact that the pencil of spheres tangent to a plane in the vertex v_i is an affine line via the inverse radius as affine coordinate, and this affine structure is preserved by Möbius transformations. Thus all affine-invariant constructions with inverse radii of Meusnier sphere are Möbius invariant. This in particular applies to the average of normal curvatures computed by (1). □

Presumably circular arc structures have further applications. One is in the old and computationally difficult problem of milling of surfaces by a 5-axis cylindrical tool, which could be solved if we can approximate the surface under consideration by a dense union of arcs of constant radius [1].

References

[1] P. Bo, H. Pottmann, M. Kilian, W. Wang, J. Wallner, *Circular arc structures*, preprint 2011.
[2] A. I. Bobenko, E. Huhnen-Venedey, *Curvature line parametrized surfaces and orthogonal coordinate systems. Discretization with Dupin cyclides.* preprint, arXiv:1101.5955.
[3] M. Eigensatz, M. Kilian, A. Schiftner, N. Mitra, H. Pottmann, M. Pauly, *Paneling architectural freeform surfaces.* ACM Trans. Graphics **29**, article # 45.
[4] J. Wallner, H. Pottmann, *Geometric computing for freeform architecture.* J. Math. Industry **1** (2011).

Osmosis – A New Methodology for Visual Computing

Joachim Weickert

Discrete formulations of diffusion filters lead to symmetric Markov chain models with a constant steady state [2]. They are mainly applied in data smoothing and regularization tasks. The goal of this talk is to introduce novel, nonsymmetric Markov chains that we call osmosis models [3]. Instead of using diffusivities between two adjacent pixels, osmosis employs forward and backward osmoticities that may differ and lead to nontrivial steady states. A fully discrete 1-D linear osmosis process is given by

$$u_i^{k+1} \;=\; \left(1 - r\, g^+_{i+\frac{1}{2}} - r\, g^-_{i-\frac{1}{2}}\right) u_i^k \;+\; r\, g^-_{i+\frac{1}{2}}\, u_{i+1}^k \;+\; r\, g^+_{i-\frac{1}{2}}\, u_{i-1}^k$$

where u_i^k approximates the grey value in pixel i at time level k, $g^+_{i+\frac{1}{2}}$ is the forward osmoticity from pixel i to $i{+}1$, $g^-_{i+\frac{1}{2}}$ is the backward osmoticity from pixel $i{+}1$ to i, and $r = \frac{\tau}{h^2}$ denotes the mesh ratio between the time step size τ and the grid size h.

We present a discrete theory for linear osmosis processes in arbitrary dimensions where osmoticities are variant in space but constant in time. It has structural similarities to the theory for discrete diffusion filtering [2], and it includes preservation of the average grey value, positivity preservation, and convergence to a nontrivial steady state. Using the Perron-Frobenius theory, the steady state is given by a rescaled version of the dominant eigenvector of the transition matrix of the Markov chain.

A continuous formulation of osmosis leads to a partial differential equation of drift-diffusion type. In the 1-D case this is given by

$$\partial_t u \;=\; \partial_x \left(g\, \partial_x u\right) \;-\; \partial_x \left(d\, u\right)$$

with a diffusivity $g := \frac{g^+ + g^-}{2}$ and a drift coefficient $d := \frac{g^+ - g^-}{h}$.

We analyse the data integration qualities of osmosis processes in the compatible and incompatible case, and we discuss relations to gradient domain methods. While osmosis models are as easy to implement as diffusion filters, their potential goes far beyond the limitations of diffusion filtering: Our prototypical applications include data clustering, compact image approximations, image editing, image fusion, and novel shock-capturing schemes for hyperbolic conservation laws [1]. Figure 1 shows an example where osmosis is used for seamless image cloning.

Joint work with Kai Hagenburg (Saarbrücken), Michael Breuß (Saarbrücken), Oliver Vogel (Saarbrücken), and Peter Ochs (Freiburg).

References

[1] K. Hagenburg, M. Breuß, J. Weickert, and O. Vogel, *Novel schemes for hyperbolic PDEs using osmosis filters from visual computing*, To appear in Proc. Third International Conference on Scale Space and Variational Methods in Computer Vision (29 May – 2 June 2011, Ein-Gedi, Israel, Springer Lecture Notes in Computer Science, 2011.

[2] J. Weickert, *Anisotropic Diffusion in Image Processing*, Teubner, Stuttgart, 1998.

FIGURE 1. Seamless image cloning with osmosis. From left to right: **(a)** Original painting of Euler. **(b)** Original drawing of Lagrange (with to-be-cloned face selected). **(c)** Direct cloning on top of Euler's head. **(d)** Cloning with osmosis image editing.

[3] J. Weickert, K. Hagenburg, M. Breuß, O. Vogel, and P. Ochs, *Osmosis methods for visual computing*, Technical Report, Dept. of Mathematics, Saarland University, 2011.

On the sharp interface limit of joint Ambrosio–Tortorelli segmentation and phase-field matching

BENEDIKT WIRTH

The tasks of image segmentation and registration are of high importance in various areas, ranging from computer vision (*e.g.*stereo reconstruction) to medicine (*e.g.*automatic construction of anatomical atlases). Image segmentation means partitioning an image into different regions such as foreground and background, using for example its edge information, while image registration is the task of finding a mapping from one image to a second one which maps parts from one image to similar parts in the other. Image registration is often applied in combination with segmentation, in particular if only the edge sets of two images are to be matched [3]. Here, it is known that robustness can be ensured by simultaneous segmentation and registration [5]: A good segmentation is obviously required for a high quality registration, but a registration can also help to find the segmentation of one image, exploiting the one-to-one correspondence between both images and the segmentation of the second image. We will here analyze a generic, variational model for joint image segmentation and registration which is based on Ambrosio–Tortorelli phase fields, a smoothed representation of image edges. In particular, we will investigate the model behaviour in the sharp interface limit.

Joint segmentation and registration via phase fields. Given two images $y_0^0, y_1^0 : \Omega \to \mathbb{R}$ on some bounded connected Lipschitz domain $\Omega \subset \mathbb{R}^d$, their segmentation can be accomplished by minimizing the well-known Mumford–Shah

functional [4],

$$\mathcal{E}_{\mathrm{MS},y_i^0}[y_i,\mathcal{S}_i] = \int_\Omega (y_i - y_i^0)^2 \,\mathrm{d}x + a\int_{\Omega\setminus\mathcal{S}_i} |\nabla y|^2 \,\mathrm{d}x + \nu\mathcal{H}^{d-1}(\mathcal{S}_i)\,, \quad i = 0,1\,,$$

which yields a piecewise smooth approximation y_i of image y_i^0 as well as its edge or discontinuity set $\mathcal{S}$. Here, a and ν represent positive parameters, and $\mathcal{H}^{d-1}$ is the $(d-1)$-dimensional Hausdorff measure. Furthermore, a registration of two images with edge sets $\mathcal{S}_0, \mathcal{S}_1$ can be performed by minimizing the functional

$$\mathcal{E}_{\mathrm{reg},\mathcal{S}_0,\mathcal{S}_1}[\phi] = \gamma\mathcal{F}[\phi^{-1}(\mathcal{S}_0),\mathcal{S}_1] + \mathcal{W}[\phi] := \gamma\mathcal{H}^{d-1}(\phi^{-1}(\mathcal{S}_0)\Delta\mathcal{S}_1) + \mathcal{W}[\phi]$$

to obtain a matching deformation $\phi : \Omega \to \mathbb{R}^d$ with $\phi(\mathcal{S}_1) \approx \mathcal{S}_0$. Here, $\gamma > 0$ is a large penalty parameter, $A\Delta B = A \setminus B \cup B \setminus A$ denotes the symmetric difference of two sets, and

$$\mathcal{W}[\phi] = \int_\Omega c_1\|\mathcal{D}\phi\|^p + c_2\|\mathrm{cof}\mathcal{D}\phi\|^q + c_3 \det\mathcal{D}\phi^r + c_4 \det\mathcal{D}\phi^{-s} + c_5 \,\mathrm{d}x$$

for positive constants $c_1, \ldots, c_4$, $p > d$, $r, s > 0$, is a generic hyperelastic energy, needed to regularize the deformation [2] (one might just as well consider different hyperelastic deformation energies, the above one has only been chosen for simplicity). A simultaneous segmentation and registration of the images y_0^0 and y_1^0 can thus be achieved by finding the minimum of

$$\begin{aligned}\mathcal{E}[y_0,y_1,\mathcal{S}_0,\mathcal{S}_1,\phi] &= \mathcal{E}_{\mathrm{MS},y_0^0}[y_0,\mathcal{S}_0] + \mathcal{E}_{\mathrm{MS},y_1^0}[y_1,\mathcal{S}_1] + \mathcal{E}_{\mathrm{reg},\mathcal{S}_0,\mathcal{S}_1}[\phi] \\ &= \mathcal{E}_{\mathrm{MS},y_0^0}[y_0,\mathcal{S}_0] + \mathcal{E}_{\mathrm{MS},y_1^0}[y_1,\mathcal{S}_1] + \gamma\mathcal{H}^{d-1}(\phi^{-1}(\mathcal{S}_0)\Delta\mathcal{S}_1) + \mathcal{W}[\phi]\,.\end{aligned}$$

Since an explicit representation of edge sets $\mathcal{S}_i$ is numerically demanding, different approximations of the above segmentation energies have been developed. In particular, we will here consider the approximation of the Mumford–Shah functional by Ambrosio and Tortorelli [1],

$$\mathcal{E}^\varepsilon_{\mathrm{AT},y^0}[y,u] = \int_\Omega (y-y^0)^2 \,\mathrm{d}x + a\int_\Omega \left(u^2 + \frac{k_\varepsilon}{a}\right)|\nabla y|^2 \,\mathrm{d}x + \nu\mathcal{L}^\varepsilon_{\mathrm{AT}}[u]\,,$$

in which the edge set $\mathcal{S}$ is replaced by a phase field function $u : \Omega \to \mathbb{R}$ which is close to one and smooth everywhere except at the image edges, where it takes the value 0. The width of this diffused edge representation scales with the small parameter $\varepsilon > 0$. The term

$$\mathcal{L}^\varepsilon_{\mathrm{AT}}[u] = \frac{1}{2}\int_\Omega \varepsilon|\nabla u|^2 + \frac{1}{\varepsilon}(1-u)^2 \,\mathrm{d}x$$

is known to be an approximation of the total length of phase-field-encoded edges. In particular,

$$\left(\Gamma((L^1(\Omega))^2) - \lim_{\varepsilon\to 0}\mathcal{E}^\varepsilon_{\mathrm{AT},y^0}\right)[y,u] =: \mathcal{E}^0_{\mathrm{AT},y^0}[y,u] = \begin{cases}\mathcal{E}_{\mathrm{MS},y^0}[y,\mathcal{S}_y] & \text{if } u = 1 \text{ a.e.},\\ \infty & \text{else},\end{cases}$$

where $\mathcal{S}_y$ shall denote the edge set of the image y. The registration of two phase fields u_0, u_1 instead of edge sets $\mathcal{S}_0, \mathcal{S}_1$ can now be performed by minimizing the functional

$$\mathcal{E}_{\mathrm{reg},u_0,u_1}[\phi] = \gamma \mathcal{F}^\varepsilon[u_0 \circ \phi, u_1] + \mathcal{W}[\phi] := \gamma \frac{1}{\varepsilon^\rho} \int_\Omega (u_0 \circ \phi - u_1)^2 \,\mathrm{d}x + \mathcal{W}[\phi]$$

(where $\rho \geq 1$), and the corresponding joint segmentation and registration energy is given by

$$\begin{aligned} &\mathcal{E}^\varepsilon[y_0, y_1, u_0, u_1, \phi] \\ &= \mathcal{E}^\varepsilon_{\mathrm{AT},y_0^0}[y_0, u_0] + \mathcal{E}^\varepsilon_{\mathrm{AT},y_1^0}[y_1, u_1] + \mathcal{E}^\varepsilon_{\mathrm{reg},u_0,u_1}[\phi] \\ &= \mathcal{E}^\varepsilon_{\mathrm{AT},y_0^0}[y_0, u_0] + \mathcal{E}^\varepsilon_{\mathrm{AT},y_1^0}[y_1, u_1] + \gamma \frac{1}{\varepsilon^\rho} \int_\Omega (u_0 \circ \phi - u_1)^2 \,\mathrm{d}x + \mathcal{W}[\phi]\,. \end{aligned}$$

We are interested in the relation between this matching energy and the corresponding sharp interface energy $\mathcal{E}$, for which purpose we analyse the Γ-limit of $\mathcal{E}^\varepsilon$ for $\varepsilon \to 0$.

Γ-convergence analysis for $\mathcal{E}^\varepsilon$. Naïvely one would expect the energy $\mathcal{E}^\varepsilon$ to be an approximation of $\mathcal{E}$. However, this is not the case due to different phenomena which will be briefly explained below. In 1D we have the following.

Theorem 11. *Let $\Omega \subset \mathbb{R}$,* i.e.*$d = 1$. For $\rho = 1$,*

$$\Gamma((L^1(\Omega))^4 \times w\text{-}W^{1,p}(\Omega)) - \liminf_{\varepsilon \to 0} \mathcal{E}^\varepsilon = \mathcal{E}^0$$

where $w\text{-}W^{1,p}(\Omega)$ denotes the space $\{\phi \in W^{1,p}(\Omega) | \phi(\Omega) \subset \Omega\}$ equipped with the weak $W^{1,p}$-topology and where

$$\mathcal{E}^0[y_0, y_1, u_0, u_1, \phi] = \mathcal{E}^0_{\mathrm{AT},y_0^0}[y_0, u_0] + \mathcal{E}^0_{\mathrm{AT},y_1^0}[y_1, u_1] + \nu Z(\gamma/\nu) \mathcal{H}^{d-1}(\mathcal{S}_{y_1} \setminus \phi^{-1}(\mathcal{S}_{y_0})) + \mathcal{W}[\phi]$$

for a smooth, monotonously increasing concave function $Z : [0, \infty) \to [0, 1]$ with $Z(0) = 0$, $\lim_{z \to \infty} Z(z) = 1$.

Obviously, we only have a one-sided matching in the limit, and the corresponding penalty parameter is bounded by ν, no matter how large γ was chosen. The reason lies in strong deformations that develop close to the edges: If y_0 has an edge but y_1 does not, then the minimizing u_0 develops a typical phase field profile which however gets squeezed by the deformation (Fig. 1, left). If y_1 has an edge but y_0 does not, phase field u_0 locally develops a ghost edge, which becomes more and more pronounced the larger γ is (Fig. 1, middle).

Remark 12. *In 2D, even the one-sided matching term vanishes in the limit $\varepsilon \to 0$ so that $\mathcal{E}^\varepsilon[y_0, y_1, u_0, u_1, \phi]$ approximates the decoupled segmentation problem*

$$\mathcal{E}^0_{\mathrm{AT},y_0^0}[y_0, u_0] + \mathcal{E}^0_{\mathrm{AT},y_1^0}[y_1, u_1] + \mathcal{W}[\phi]\,.$$

The reason lies in the development of small ghost edge segments which get elongated by the deformation ϕ to form a complete ghost edge (Fig. 1, right).

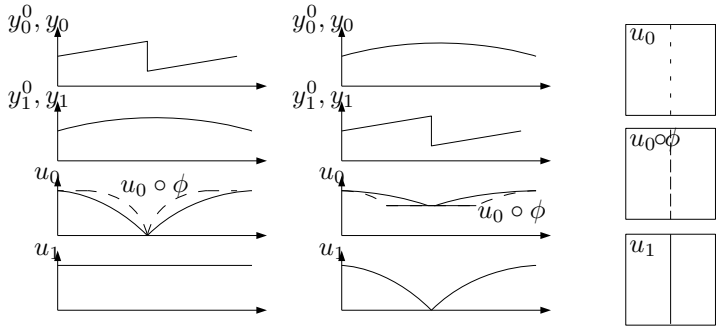

FIGURE 1. Left: In 1D, where image y_0 has an edge but y_1 does not, phase field u_0 gets locally squeezed to $u_0 \circ \phi$ so that $\mathcal{F}^\varepsilon[u_0 \circ \phi, u_1]$ vanishes in the limit $\varepsilon \to 0$. Middle: In 1D, where image y_1 has an edge but y_0 does not, phase field u_0 locally develops a ghost edge which is counted in $\mathcal{E}^\varepsilon_{\mathrm{AT},y_0^0}[y_0, u_0]$. Right: In 2D, where image y_1 has an edge but y_0 does not, phase field u_0 locally develops a sequence of small ghost edge segments that get elongated in $u_0 \circ \phi$ so that $\mathcal{F}^\varepsilon[u_0 \circ \phi, u_1]$ vanishes in the limit $\varepsilon \to 0$.

As a remedy of the above effects, one can increase the exponent ρ in $\mathcal{E}^\varepsilon$. This will prohibit a decrease of the mismatch penalty term via locally strong deformations. Thereby the increased ρ will enforce the formation of fully developed ghost edges so that in the limit $\varepsilon \to 0$ the mismatch term becomes zero, but the accumulated length of all ghost edges is counted with weight ν.

Theorem 13. *Let* $\Omega \subset \mathbb{R}$, i.e.$d = 1$. *For* $\rho > 1 + \frac{1}{\max\{p,r\}}$,

$$\Gamma((L^1(\Omega))^4 \times w\text{-}W^{1,p}(\Omega)) - \liminf_{\varepsilon \to 0} \mathcal{E}^\varepsilon = \mathcal{E}^0,$$

$$\mathcal{E}^0[y_0, y_1, u_0, u_1, \phi] = \mathcal{E}^0_{\mathrm{AT},y_0^0}[y_0, u_0] + \mathcal{E}^0_{\mathrm{AT},y_1^0}[y_1, u_1] + \nu \mathcal{H}^{d-1}(\phi^{-1}(\mathcal{S}_{y_0}) \Delta \mathcal{S}_{y_1}) + \mathcal{W}[\phi]\,.$$

In 2D, a similar result is to be expected, where the choice of ρ additionally has to take into account the other exponents of the hyperelastic deformation regularization.

REFERENCES

[1] L. Ambrosio and V. M. Tortorelli, *On the approximation of free discontinuity problems*, Bollettino dell'Unione Matematica Italiana, Sezione B, **6** (1992), 105–123.

[2] M. Droske and M. Rumpf, *A variational approach to non-rigid morphological registration*, SIAM Journal on Applied Mathematics **64** (2004), 668–687.

[3] Jingfeng Han, Benjamin Berkels, Marc Droske, Joachim Hornegger, Martin Rumpf, Carlo Schaller, Jasmin Scorzin, and Horst Urbach, *Mumford-Shah model for one-to-one edge matching*, IEEE Transactions on Image Processing **16** (2007), 2720–2732.

[4] David Mumford and Jayant Shah, *Optimal approximation by piecewise smooth functions and associated variational problems*, Communications on Pure and Applied Mathematics **42** (1989), 577–685.

[5] A. Yezzi, L. Zöllei, and T. Kapur. *A variational framework for integrating segmentation and registration through active contours*, Medical Image Analysis **7** (2003), 171–185.

Statistical analysis of surfaces for computational anatomy

Xavier Pennec

Over the last 30 years, there was an explosion of imaging modalities allows observing both the anatomy *in vivo* and *in situ* at multiple spatial scales (from cells to the whole body), multiple time scales (beating heart, growth, aging, evolution of species), and on multiple subjects. The combination of these new observation means and of the computerized methods is at the heart of *computational anatomy*, an emerging discipline at the interface of geometry, statistics and image analysis which aims at developing algorithms to model and analyze the biological shape of tissues and organs. The goal is to estimate representative organ anatomies across diseases, populations, species or ages, to model the organ development across time (growth or aging), to establish their variability, and to correlate this variability information with other functional, genetic or structural information (e.g. fiber bundles extracted from diffusion tensor images). From an applicative point of view, a first objective is to understand and to model how life is functioning at the population level, for instance by classifying pathologies from structural deviations (taxonomy). A second application objective is to better drive the adaptation of generic models of the anatomy (atlas) into patient-specific data (personalization) in order to help therapy planning (before), control (during) and follow-up (after).

Understanding and modeling the shape of organs is made difficult by the absence of physical models for comparing different subjects, the complexity of shapes, and the high number of degrees of freedom implied. The general method is to identify anatomically representative geometric features (points, tensors, curves, surfaces, volume transformations), and to describe and compare their statistical distribution in different populations. As these geometric features most often belong to manifolds that have no canonical Euclidean structure, we have to rely on more elaborated algorithmic basis. The Riemannian structure proves to be a powerful and consistent framework for computing simple statistics on finite dimensional manifolds [7, 8] and can be extend to a complete computing framework on manifold-valued images [9]. For instance, the choice of a convenient Riemannian metric on the space of positive define symmetric matrices (tensors) allows to generalize consistently to tensor fields many important geometric data processing algorithms such as interpolation, filtering, diffusion and restoration of missing data. This framework is particularly well suited to the statistical estimation of Diffusion Tensor Images [5], and can also be used for modeling the brain variability from sulcal lines drawn at the surface of the cerebral cortex [4].

To move from simple point-wise features to curves, surfaces and deformations, we believe that the embedding vector space of currents provide an interesting computational environment. It was introduced in the field by Glaunes and extended by Durrleman [2, 1] to a generative shape model which combines a random diffeomorphic deformation model a la Grenander & Miller, that encodes the geometric variability of the anatomical template, with a random residual shape variability model (a la Kendall) on the deformed template. We applied the efficient algorithmic toolbox developed for handling statistics on currents to the analysis of

the shape of the right ventricle of the heart in a population of Tetralogy of Fallot patients The resulting statistical model of the remodeling of the ventricle during growth turns out to have an anatomically meaningful interpretation [6]. The extensions of this type of methodology to longitudinal evolution estimations in populations is currently one of the most active topic in computational anatomy. We present here a simple model where we combine a static inter-subject change of coordinate system with a time-warp to transform the generic scenario of deformation at the population level to the subject specific longitudinal observations. When applied to different species (here bonobos vs chimpanzees) or to diseases (autism vs control), this model suggests that the change in the speed of evolution might be more important than the shape differences [3].

References

[1] S. Durrleman, P. Fillard, X.r Pennec, A. Trouvé, and N. Ayache, *Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents*, NeuroImage, 2011 (In press).

[2] S. Durrleman, X. Pennec, A. Trouvé, and N. Ayache, *Statistical models on sets of curves and surfaces based on currents*, Med. Image Analysis **13** (2009), 793–808.

[3] S. Durrleman, X. Pennec, A. Trouvé, N. Ayache and J. Braga, *Comparison of the endocast growth of chimpanzees and bonobos via temporal regression and spatiotemporal registration*, Proc. of Miccai STIA Workshop, Beijing, Sep. 2010.

[4] P. Fillard, V. Arsigny, X. Pennec, K.M. Hayashi, P.M. Thompson, and N. Ayache, *Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines*, Neuroimage **34** (2007), 639–650.

[5] P. Fillard, X. Pennec, V. Arsigny, and N. Ayache, *Clinical DT-MRI estimation, smoothing and fiber tracking with log-euclidean metrics*, IEEE Transactions on Medical Imaging **26** (2007), 1472–1482.

[6] T. Mansi, S. Durrleman, B. Bernhardt, M. Sermesant, H. Delingette, I. Voigt, P. Lurz, A.M Taylor, J. Blanc, Y. Boudjemline, X. Pennec, and N. Ayache, *A statistical model of right ventricle in tetralogy of Fallot for prediction of remodelling and therapy planning*, In Proc. of MICCAI'09, LNCS 5761, pages 214–221, London, UK, Sep. 2009. Springer.

[7] X. Pennec, *Probabilities and statistics on Riemannian manifolds: Basic tools for geometric measurements*, In Proc. of Nonlinear Signal and Image Processing (NSIP'99), vol. 1, p.194–198, June 20-23, Antalya, Turkey, 1999. IEEE-EURASIP.

[8] X. Pennec, *Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements*, J. of Mathematical Imaging and Vision **25** (2006), 127–154.

[9] X. Pennec, P. Fillard, and N. Ayache, *A Riemannian framework for tensor computing*, International Journal of Computer Vision **66** (2006), 41–66.

Shape Analysis via Metric Geometry

Ron Kimmel

Natural objects can be subject to various transformations yet still preserve properties that we refer to as invariants. In the talk I described a shape analysis framework that treats shapes as metric spaces constructed so as to be invariant or insensitive to desired quantities. One example is invariance to bending of an object in 3D space. Such deformations are known to conserve the intrinsic geometry in its

explicit sense, in which geodesic distances, defined by the shortest paths between points on the surface, are preserved.

Different geometric structures could be defined by diffusion geometry which is an umbrella term referring to geometric analysis of diffusion or random walk processes. Such methods, first introduced in theoretical geometry by Berard et al. in the 90's have matured into practical applications in the fields of manifold learning and shape analysis. The power of diffusion geometry is its insensitivity to topology changes in the shape we try to characterize. Other alternative geometries include volumetric measures, in which distances between points take into consideration the interior of the surface and the shortest path is determined by the shortest fly-through trajectory in the interior of the shape. Affine invariants could also play a role in this game, in which case an arclength on the surface, invariant to affine transformations of the embedding space, could be used to define either geodesic distances or a new diffusion geometry.

The above measures would be used within our context of shape analysis. We demonstrate how the right selection of a metric and a corresponding geometry enables us to use, for example, the Gromov-Hausdorff distance for comparing surfaces. The different choice of arclength and the resulting geometry enable us to develop computational tools for dealing with various shape analysis problems.

In the shape analysis community, diffusion geometry methods were used to define low-dimensional representations for manifolds, build intrinsic distance metrics and construct shape distribution descriptors, define spectral signatures like shape-DNA, and local descriptors. Diffusion embeddings were used for finding correspondence between shapes and detecting intrinsic symmetries. In many settings, the construction of diffusion geometry boils down to the definition of a Laplacian operator.

In this talk we show how to define an affine invariant arclength for surfaces in $R3$ in order to extend the set of existing non-rigid shape analysis tools. In a recent paper with Raviv, Bronstein, Bronstein, and Sochen) we showed that by re-defining the surface metric as its equi-affine version, the surface with its modified metric tensor can be treated as a canonical Euclidean object on which most classical Euclidean processing and analysis tools can be applied. The new definition of a metric can used to extend the fast marching method technique for computing geodesic distances on surfaces, where now, the distances are defined with respect to an affine invariant arclength. Applications of the proposed framework demonstrate its invariance, efficiency, and accuracy in shape analysis.

Using an (equi-)affine invariant diffusion geometry surfaces that go through squeeze and shear transformations can still be properly analyzed. The definition of an affine invariant metric enables us to construct an invariant Laplacian from which local and global geometric structures are extracted. Applications of the proposed framework demonstrate its power in generalizing and enriching the existing set of tools for shape analysis. All these properties were reviewed in the talk.

Modeling 3D shapes as Riemannian manifold is a ubiquitous approach in many shape analysis applications. In particular, in the recent decade, shape descriptors based on geodesic distances induced by a Riemannian metric have become popular. Notable examples of such methods are the canonical forms and the Gromov-Hausdorff, and the Gromov-Wasserstein frameworks, used in shape comparison and correspondence problems. Such methods consider shapes as metric spaces endowed with a geodesic distance metric, and pose the problem of shape similarity as finding the minimum-distortion correspondence between the metrics. The advantage of the geodesic distances is their invariance to inelastic deformations (bending) that preserve the Riemannian metric, which makes them especially appealing for non-rigid shape analysis. A particular setting of finding shape self-similarity can be used for intrinsic symmetry detection in non-rigid shapes.

The flexibility in the definition of the Riemannian metric allows extending the invariance of the aforementioned shape analysis algorithms by constructing a geodesic metric that is also invariant to global transformations of the embedding space. A particularly general and important class of such transformations are the affine transformations, which play an important role in many applications in the analysis of images and 3D shapes. Many approaches have been suggested to cope with the action of the affine group in a global manner, trying to undo the affine transformation in large parts of a shape or a picture. While the theory of affine invariance is known for many years and used for curves and flows, no numerical constructions applicable to manifolds have been proposed.

At the end of the talk I presented the explicit way to construct an (equi-)affine-invariant Riemannian geometry for 3D shapes. By defining an affine-invariant Riemannian metric, we can in turn define affine invariant geodesics, which result in a metric space with a stronger class of invariance.

New metrics allow us to develop efficient computational tools that handle non-rigid deformations as well as affine transformations. We demonstrate the usefulness of our construction in a range of shape analysis applications, such as shape processing, construction of shape descriptors, correspondence, and symmetry detection.

Reporter: Vicent Caselles

Participants

Prof. Dr. Marc Alexa
Fakultät f. Elektrotechnik u. Informatik
TU Berlin
Einsteinufer 17
10587 Berlin

Dr. Pierre Alliez
INRIA Sophia Antipolis
B.P. 93
2004 Route des Lucioles
F-06902 Sophia Antipolis Cedex

Pablo Arias
C/Tanger, 122-140
E-08018 Barcelona

Prof. Giovanni Bellettini
Dipartimento di Matematica
Universita di Roma Tor Vergata
Via della Ricerca Scientif. 1
I-00133 Roma

Markus Bördgen
Institut für Informatik
TU München
Boltzmannstr. 3
85748 Garching

Dr. Kristian Bredies
Institut für Mathematik
Karl-Franzens-Universität Graz
Heinrichstr. 36
A-8010 Graz

Prof. Dr. Vicent Caselles
Departament de Tecnologies de la
Informacio i les Comunicacions
Universitat Pompeu Fabra
Roc Boronat, 138
E-08018 Barcelona

Dr. Antonin Chambolle
Centre de Mathematiques Appliquees
Ecole Polytechnique
F-91128 Palaiseau Cedex

Keenan Crane
Department of Computer Science
California Institute of Technology
1200 East California Boulevard
Pasadena CA 91125
USA

Prof. Dr. Daniel Cremers
Lehrstuhl für Informatik IX:
Bildverstehen & wissensbasierte Systeme
Boltzmannstr. 3
85748 Garching

Julie Digne
Centre de Mathématiques et de Leurs
Applications (CMLA)
Ecole Normale Superieure de Cachan
61, Avenue du President Wilson
F-94235 Cachan Cedex

Prof. Dr. Selim Esedoglu
Department of Mathematics
University of Michigan
530 Church Street
Ann Arbor , MI 48109-1043
USA

Gabriele Facciolo
Departament de Tecnologies de la
Informacio i les Comunicacions
Universitat Pompeu Fabra
Roc Boronat, 138
E-08018 Barcelona

Carina Geldhauser
Hausdorff Center for Mathematics
Institute for Numerical Simulation
Endenicher Allee 60
53115 Bonn

Prof. Dr. Peter Giblin
Dept. of Mathematical Sciences
University of Liverpool
Peach Street
GB-Liverpool L69 7ZL

Prof. Dr. Eitan Grinspun
Department of Computer Science
Columbia University
Seeley W. Mudd Building
New York , NY 10027
USA

Prof. Dr. Leonidas J. Guibas
Computer Science Department
Stanford University
Gates Bldg 3B-374
Stanford CA 94305-9035
USA

Behrend Heeren
Hausdorff Center for Mathematics
Institute for Numerical Simulation
Endenicher Allee 60
53115 Bonn

Prof. Dr. Sung Ha Kang
School of Mathematics
Georgia Institute of Technology
686 Cherry Street
Atlanta , GA 30332-0160
USA

Prof. Dr. Ron Kimmel
Department of Computer Sciences
Technion
Israel Institute of Technology
Haifa 32000
ISRAEL

Prof. Dr. Reinhard Klein
Institut für Informatik
Universität Bonn
Beringstr. 1
53115 Bonn

Prof. Dr. Leif Kobbelt
Lehrstuhl für Informatik VIII
RWTH Aachen
Ahornstr. 55
52074 Aachen

Prof. Dr. Andrea C.G. Mennucci
Scuola Normale Superiore di Pisa
Piazza dei Cavalieri 7
I-56126 Pisa

Prof. Dr. Michael I. Miller
Center for Imaging Science
Clark 301
The Johns Hopkins University
3400 N. Charles Street
Baltimore , MD 21218-2686
USA

Prof. Dr. Niloy Mitra
Computer Science
4700 King Abdullah University of
Science and Technology
Thuwal 23955-6900
SAUDI ARABIA

Prof. Dr. Matteo Novaga
Dipartimento di Matematica
Universita di Padova
Via Trieste, 63
I-35121 Padova

Dr. Xavier Pennec
INRIA Sophia Antipolis
B.P. 93
2004 Route des Lucioles
F-06902 Sophia Antipolis Cedex

Dr. Gabriel Peyre
Universite Paris Dauphine
Place du Marechal DeLattre de Tassigny
F-75016 Paris

Prof. Dr. Ulrich Pinkall
Fakultät II-Institut f. Mathematik
Technische Universität Berlin
Sekr. MA 3-2
Straße des 17. Juni 136
10623 Berlin

Dr. Thomas Pock
Institute for Software Technology
University of Technology Graz
Inffeldgasse 16b/II
A-8010 Graz

Prof. Dr. Helmut Pottmann
Visualization Research Centre
4700 King Abdullah University of
Science and Technology
Thuwal 23955-6900
SAUDI ARABIA

Dr. Paola Pozzi
Abteilung für Angewandte Mathematik
Universität Freiburg
Hermann-Herder-Str. 10
79104 Freiburg

Prof. Dr. Holger Rauhut
Hausdorff Center for Mathematics
Institute for Numerical Simulation
Endenicher Allee 60
53115 Bonn

Prof. Dr. Martin Rumpf
Hausdorff Center for Mathematics
Institute for Numerical Simulation
Endenicher Allee 60
53115 Bonn

Prof. Dr. Otmar Scherzer
Computational Science Center
Universität Wien
Nordbergstr. 15
A-1090 Wien

Dr. Ulrich Schlickewei
Institut für Informatik
TU München
Boltzmannstr. 3
85748 Garching

Prof. Dr. Christoph Schnörr
Universität Heidelberg
Image & Pattern Analysis Group
Speyrer Str. 6
69115 Heidelberg

Prof. Dr. Peter Schröder
306 Annenberg Center for Information
Science & Technology MC 305-16
California Institute of Technology
1200 E. California Blvd.
Pasadena , CA 91125-5000
USA

Prof. Dr. Alla Sheffer
Department of Computer Science
University of British Columbia
2366 Main Mall
Vancouver , B.C. V6T 1Z4
CANADA

Justin Solomon
Computer Science Department
Stanford University
Stanford , CA 94305-2095
USA

Prof. Dr. Olga Sorkine
Department of Computer Science
ETH Zürich
Universitätstr. 6
CH-8092 Zürich

Pablo Sprechmann
Dept. of Electrical Engineering
University of Minnesota
200 Union Street
Minneapolis , MN 55455
USA

Prof. Dr. Alain Trouve
Ecole Normale Superieure de Cachan
61, Av. du President Wilson
F-94235 Cachan

Prof. Dr. Etienne Vouga
Department of Computer Science
Columbia University
New York NY 10027-7003
USA

Prof. Dr. Johannes Wallner
Institut für Geometrie
TU Graz
Kopernikusgasse 24
A-8010 Graz

Prof. Dr. Max Wardetzky
Institut für Numerische
und Angewandte Mathematik
Universität Göttingen
Lotzestr. 16-18
37083 Göttingen

Prof. Dr. Joachim Weickert
Fakultät für Mathematik & Informatik
Universität des Saarlandes
Campus E1.1
66041 Saarbrücken

Prof. Dr. Ross T. Whitaker
Department of Computer Science
University of Utah
Salt Lake City , Utah 84112
USA

Dr. Benedikt Wirth
Hausdorff Center for Mathematics
Institute for Numerical Simulation
Endenicher Allee 60
53115 Bonn

Prof. Dr. Laurent Younes
Dept. of Appl. Mathematics & Statistics
The Johns Hopkins University
34th and Charles Streets
Baltimore , MD 21218-2682
USA

Prof. Dr. Denis Zorin
Media Research Laboratory
Courant Institute of Mathematical
Sciences, New York University
719 Broadway 12th Fl.
New York NY 10003
USA