Dr. Dominique Zosso **Summary of Research Interests**

From Sparsity in Images and Information Science to Efficient PDE Solvers

Images are ubiquitous today. Nearly every mobile phone has an integrated camera and selfies are *de rigueur*. More importantly, images have become essential in remote sensing, medicine, biology, and computer vision. Indeed, images provide a rich description about the reality around us. In the visible spectrum, they tell us something about our environment and allow us to navigate safely or reach for objects. Images, however, extend also to other spectra and allow us to see through objects or measure quantities at very small scales. In medical imaging we extract information from images about a physical reality that is hardly accessible to inspection otherwise. *No matter what imaging modality one uses, automated image analysis is key.*

Images are always a projection of a more complicated underlying physical reality which is of interest; image understanding describes an ensemble of individual problems such as denoising, segmentation, classification, decomposition, aiming at getting information about this hidden physical layer out of images. Many associated research questions have not been answered to full satisfaction, yet. Mathematically speaking, most image-related tasks are instances of inverse problems, where one looks for an underlying, hidden layer of information given a set of derived measurements. New applications, types of images, and questions to be answered through imaging come up everyday. On the other hand,

> “ by combining the power of modern computing with the plentiful data of the digital era, it [= Big Data] promises to solve virtually any problem—crime, public health, the evolution of grammar, the perils of dating—just by crunching the numbers. ”
>
> (G. Marcus and E. Davis, 2014 New York Times Op-Ed)

Number and size of these challenges keep increasing at a startling rate. The key to tackling the relentlessly growing model complexity and sheer amount of data lies in the *smart exploitation of problem structure and sparsity*. In general, there is a strong convergence between problems and methods in imaging on the one side, and data science and machine learning on the other side. Maybe more unexpectedly, an important class of partial differential equations equally bears strong similarities with these imaging and data problems. Indeed, elliptic PDE, which are commonly found in scientific computing, mathematical physics, material science, computer graphics, and consumer electronics, correspond to convex optimization problems of the same form as those encountered in imaging and data science. Solving these PDE more efficiently greatly increases the range of problems that we are able to handle computationally. My research is located at the confluence of these three major topics:

In a nutshell, my research interests are *mathematical problems at the cutting-edge frontier of image and data understanding*, typically formulated as variational models, and the development of *efficient algorithms for the numerical solution of these models and related PDE*.

An overview of my recent research accomplishments in these areas is given below: §1 focuses on algorithms, §2 focuses on variational image decomposition models, and §3 summarizes earlier work, in particular my Ph.D. thesis.

Keywords: variational methods, numerical PDE and PDE on graphs, convex optimization, mathematical imaging, data science, machine learning, minimal surfaces.

1 Efficient Algorithms for Images, Graphs, and Elliptic PDE

> “ Behind any great model, there is a great algorithm. ”

1.1 Primal-dual hybrid gradients

While ADMM- and split-Bregman-type methods have been successful for some time, more recently the primal-dual hybrid gradients method [1, 2] has been revived from earlier works in economics [3]. The idea is as follows: Let X and Y be two finite dimensional convex and closed sets; consider a continuous linear operator $K\colon X \to Y$ with finite norm, and two proper, convex, lower-semicontinuous functions, $f\colon Y \to \mathbb{R}$ and $g\colon X \to \mathbb{R}$. Using the *convex conjugate* (Legendre-Fenchel transform), $f^* : Y \to \mathbb{R}$, $f(y) := \sup_{s\in Y} \langle y, s\rangle - f(s)$, the minimization (left) and saddle point problem (right) are equivalent:

$$\min_{x\in X} \; f(Kx) + g(x) \qquad\qquad \min_{x\in X} \; \max_{y\in Y} \; \langle Kx, y\rangle - f^*(y) + g(x). \tag{1}$$

Now, the primal-dual hybrid gradients (PDHG) algorithm for the saddle point problem consists of three steps [2]: The first two steps are the proximal update of the dual and primal variable, respectively. The third, extra-gradient step is an

overrelaxation of the primal update in order to overcome the stepsize shortening typical of first order methods:

$$y^{n+1} = \arg\min_{y \in Y} \; -\langle K\bar{x}^n, y\rangle + f^*(y) + \frac{1}{2r_1}\|y - y^n\|^2, \tag{2a}$$

$$x^{n+1} = \arg\min_{x \in X} \; \langle x, K^* y^{n+1}\rangle + g(x) + \frac{1}{2r_2}\|x - x^n\|^2, \tag{2b}$$

$$\bar{x}^{n+1} = x^{n+1} + \theta(x^{n+1} - x^n), \tag{2c}$$

for $0 \le \theta \le 1$. In general, $O(1/n)$ convergence has been shown for fixed r_1, r_2, such that $r_1 r_2 \|K\|^2_{X\to Y} \le 1$.

Contributions: I have successfully applied this relatively new (revived) technique to many convex optimization problems presented below, achieving great numerical performance and important speed-ups. Examples include the obstacle problem §1.2, Beltrami regularization §1.3, and clustering §1.5. For the clustering/graph partitioning problem, we are now able to solve 3D spectral space partitioning problems that could not previously be handled without significant parallelization. A project of even broader significance is the translation of the algorithm to the realm of elliptic PDE, discussed in §1.6.

1.2 Obstacle problem [4]

The obstacle problem is a classic motivating example in the mathematical study of variational inequalities and free boundary problems with broad applications, *e.g.*, in fluid filtration in porous media, constrained heating, elasto-plasticity, optimal control, or financial mathematics [5, 6, 7, 8]. The classical formulation is motivated by the equilibrium position of a membrane described by the graph of $u\colon \Omega \subset \mathbb{R}^n \to \mathbb{R}$, with Dirichlet boundary conditions, and restricted from below by the obstacle $\varphi\colon \Omega \to \mathbb{R}$ [6]:

$$\min_u \quad \int_\Omega \sqrt{1 + |\nabla u|^2} \qquad \text{s.t.} \quad u \ge \varphi, \quad u(\partial\Omega) = 0 \tag{3}$$

For an example, see figure 1. Using the convex-conjugate of the surface functional, we obtain the primal-dual problem:

$$\min_{\substack{u \ge \varphi \\ u(\partial\Omega)=0}} \quad \max_{\substack{y \\ |y|<1}} \quad \int_\Omega \langle \nabla u, y\rangle + \sqrt{1 - |y|^2} \tag{4}$$

Contributions: We solve this saddle point problem very efficiently using primal-dual hybrid gradients (see §1.1), resulting in a considerable speed up of 1-2 orders of magnitude over most recent state of the art methods, *e.g.*, [9].

1.3 A Primal-Dual hybrid gradients algorithm for efficient Beltrami regularization [10]

Sochen, Kimmel and Malladi introduced the powerful Beltrami framework for image denoising and enhancement, based on the Polyakov model in string theory [11]. Adopting this pure geometric point of view amounts to seeing objects such as images, shapes, or vector fields as geodesics in some higher dimensional space. The Beltrami framework associates the spatial coordinates along with the features by defining a Monge surface $\mathcal{M}$, say $X\colon (x, y) \to (x, y, I)$ for a 2D scalar image. The Beltrami energy then measures the area of this image membrane, and its minimization amounts to regularizing (smoothing) the image in a feature preserving way.

Beltrami-regularization compares favorably to other popular methods in terms of quality, thanks to its feature-awareness without staircaising (as opposed to TV and H^1). Consider for example the simple Beltrami denoising problem:

$$\min_I \quad \int \sqrt{1 + \beta^2 |\nabla I|^2} + \frac{1}{\alpha}(I - I_0)^2, \tag{5}$$

where β is a parameter controlling the aspect ratio between spatial and feature coordinates. The relation to the minimal surface problem is obvious. A direct implementation of the Beltrami flow is obtained by using an Euler explicit forward scheme. However, this scheme is heavily limited by the CFL condition and is therefore slow. We used variable splitting and augmented Lagrangian techniques for optimization, before. Here we use a *primal-dual reformulation* of the Beltrami energy and the associated convex data-term, corresponding to the saddle-point formulation in §1.1. Using the Fenchel-Legendre transform of the Beltrami energy, we rewrite the primal image denoising problem as a primal-dual problem:

$$\min_I \quad \max_{\phi,\ |\phi|^2 \le \beta^2} \quad \int \langle \nabla I, \phi\rangle + \frac{\sqrt{\beta^2 - |\phi|^2}}{\beta} + \frac{1}{\alpha}(I - I_0)^2, \tag{6}$$

where ϕ is the dual variable. Thanks to the adjoint equality, $\int \langle \nabla I, \phi\rangle = -\int I \operatorname{div}\phi$, the resulting primal-dual problem can be solved efficiently using iterative projections of alternate directions gradient descent [3, 2].

Contributions: Our algorithm converges considerably faster than classical Beltrami implementations, and extends directly to problems other than simple denoising, such as deblurring, inpainting and compressed sensing, see [10]. The algorithm was successfully used in clinical applications [12] and cross-fertilized speed-ups in other minimal surface problems (see below).

1.4 Non-local Beltrami and Beltrami on graphs

Non-local regularization currently produces very promising results in diverse applications. I am working on rendering the Beltrami-energy fully non-local, by extending its definition to non-local operators. Let $w(\mathbf{x}, \mathbf{y})$ be the non-negative weights between points $\mathbf{x}, \mathbf{y} \in \Omega$. We define the non-local gray-scale Beltrami functional as

$$E_{NL-B}(u) := \int_\Omega \sqrt{1 + \beta^2 \|\nabla_w u(\mathbf{x})\|_2^2}, \tag{7}$$

where $(\nabla_w u)(\mathbf{x}, \mathbf{y}) := \sqrt{w(\mathbf{x}, \mathbf{y})}(u(\mathbf{y}) - u(\mathbf{x}))$, $\mathbf{x}, \mathbf{y} \in \Omega$, defines the non-local gradient and $\|\nabla_w u(\mathbf{x})\|_2$ its L^2-norm. This functional is largely unstudied so far, and in particular I'm interested in deriving the functional as a hyper-surface measure from a non-local embedding that formally looks like $X \colon (\mathbf{x} \in \Omega, \Omega \setminus \mathbf{x}) \mapsto (\mathbf{x}, \Omega \setminus \mathbf{x}, u(\mathbf{x}))$.

Discretization leads to an equivalent definition on graphs. Let $G(V, E)$ be a graph with edge-weights $\{w_{ij}\}_{(i,j)\in E}$. We can then define the convex graph-Beltrami energy of $u \colon V \to \mathbb{R}$:

$$E_{G-B}(u) := \sum_{i \in V} \sqrt{1 + \beta^2 \sum_{j|(i,j)\in E} w_{ij}(u(j) - u(i))^2}. \tag{8}$$

Contributions: This graph-based definition makes the benefits of Beltrami regularization available for data defined on graphs, and applies *e.g.*, to clustering or segmentation, with numerous applications in machine learning.

1.5 A minimal surface criterion for graph partitioning and image segmentation [13, 14]

One such application of the Beltrami regularizer on graphs is related to the data clustering or graph partitioning problem. Here, we group the vertices of a graph (*e.g.*, individuals in a social network, documents in an archive, pixels of an image) such that each group minimizes an energy term related to its coherence. We study the graph k-partitioning problem,

$$\min_{V = \sqcup_i V_i} \quad \sum_{i=1}^{k} J_\beta^*(V_i), \tag{9}$$

where $J_\beta^*(V_i)$ is the minimal constrained Beltrami energy of the vertex set V_i, defined as:

$$J_\beta^*(S \subset V) := \min_u \quad E_{G-B}(u) \quad \text{s.t.} \quad u \colon V \to \mathbb{R}, \quad \|u\|_{V,p} = 1, \quad u|_{S^c} = 0. \tag{10}$$

The inner energy (10) represents the ground state of a soap-bubble defined as the graph of a function over the subset S and having prescribed volume 1. The complete model (9) thus represents the *total-surface area of a type of foam*, where bubbles have prescribed volume and compete through surface tension. To solve this bi-level optimization problem efficiently, we propose a primal-dual hybrid gradients method for the computation of the ground state of each partition with relaxed boundary conditions (inner problem), and a rearrangement algorithm to find the optimal partitions (outer problem). The inner problem is convergent and we prove that the outer iterations do not increase the objective value.

Natural applications are omnipresent clustering problems in data science. Visually intuitive are manifold partitions, such as partitioning a 2D torus into 16 disjoint regions, illustrated in figure 2. A simple graph partitioning model corresponding to region-based image segmentation can be obtained by extending the square-lattice graph with edges between spatially distant pixels, weighted by feature similarity (see figure 3).

Contributions: The graph-based Beltrami energy considered here is an interesting generalization and interpolation between the total variation (TV, $\beta \to \infty$) and Dirichlet ($\beta \to 0$) energies that have previously been used for graph partitioning. The image segmentation application is based on a graph construction with both short- and long-range connections, reproducing both edge- and region-based segmentation criteria, as well as natural hybrids that can deal, *e.g.*, with illumination bias in images. The model naturally extends to more interesting features such as texture and arbitrary numbers of partitions (phases).

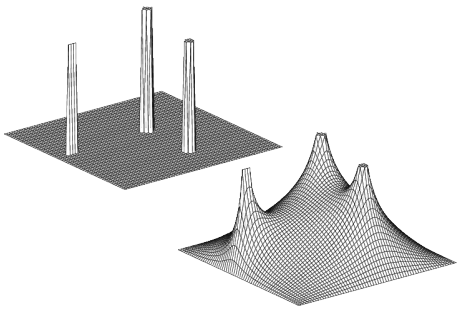

Figure 1: **Obstacle problem**. Given obstacles (left), minimal surface (right).

Figure 2: **Partition of a 2D torus**. Double-periodic partition into 16 disjoint regions with minimal surface.

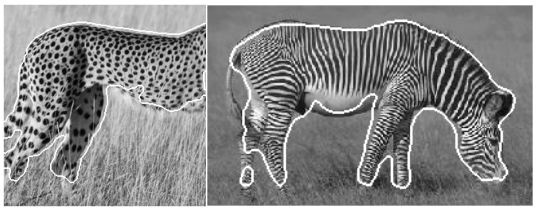

Figure 3: **Image segmentation as graph partitioning**. For graph construction long-distance edges are included, weighted by feature similarity (here: inverse of Wasserstein distance), generalizing region-based image segmentation to texture features, multiple phases, and images with bias.

1.6 PDHG for elliptic PDE [15]

Elliptic PDE describe the minimizers of associated convex minimization problems. For example, consider the Poisson problem $-\Delta u = f$, which is the minimizer of $\min_u \ \frac{1}{2}\|\nabla u\|^2 - \langle u, f\rangle$. We can identify this problem with the general algorithm in §1.1. Thanks to the particular structure, the dual variable can be eliminated, and the resulting scheme has the two equivalent forms:

$$u^{n+1} = u^n + r_1 r_2 \sum_{i=0}^{n} \Delta\bar{u}^i/(1+r_1)^{n-i+1} \qquad (u^{n+1} - 2u^n + u^{n-1}) + r_1(u^{n+1} - u^n) = r_1 r_2 \Delta\bar{u}^n, \tag{11}$$

where the LHS is a boosted version of explicit heat diffusion, and the RHS is reminiscent of a damped wave equation: $u^n_{tt} + \frac{r_1}{h}u^n_t = \frac{r_1 r_2}{h^2}\Delta\bar{u}^n$. Compare the LHS to the regular explicit heat diffusion update:

$$u^{n+1} = u^n + r_1 r_2 \Delta u^n. \tag{12}$$

The scheme in (12) is well-known to be critical regarding time-step selection and therefore dreadfully slow. Why would the very similar LHS scheme be of any interest, then? Thanks to the sum in (11), previous heat updates continue to be taken into account, the impact of which is essential for the efficiency and convergence of the scheme. Consider the following two cases: *Away from the solution*, subsequent heat updates point towards the same direction and the sum lets the history of updates interfere constructively, effectively greatly increasing the time step of the explicit heat update. Assuming $\Delta\bar{u}^{n-i} \approx \Delta u^n$ constant for at least some small i (recent iterations), we can approximate (11) by $u^{n+1} \approx u^n + r_2\Delta u^n$. The proposed scheme is thus heat diffusion with a $1/r_1$-fold greater effective time step than (12). *Close to the solution*, u oscillates about the optimum and subsequent heat steps have opposing sign; here, the summation leads to destructive interference of subsequent updates, dampens the oscillations, and effectively stabilizes the explicit heat update.

Contributions: I am currently exploring the implications of the resulting update schemes (11). The updates are explicit, and do not require inversion of the Laplacian (in general, PDHG only requires operator and adjoint, as opposed to ADMM/split-Bregman). The scheme is stable for $r_1 r_2|\Delta| \leq 1$, yet thanks to the boosting of the sum it performs orders of magnitudes faster than standard explicit heat diffusion schemes. The method is related to the "momentum method" and Nesterov's gradient acceleration. Preliminary results show better than off-the-shelf performance for various problems, including the Poisson and obstacle problem, MBO, elliptic PDE on graphs, and even Schrödinger-operator eigenproblems with non-convex constraints.

2 Variational Models for Image Decomposition

Typically driven by practical needs expressed by collaborators in the experimental sciences, I have been working on the development of novel variational models for fundamental inverse problems in imaging.

2.1 Variational mode decomposition [16, 17, 18]

Many real world signals are a superposition of different underlying components of different frequency. Timeseries data of traffic density or climate measurements, for example, contain superposed oscillations of daily, weekly, seasonal, and longer periods. To analyze such data, it is often useful to unmix and isolate these different modes. The empirical mode decomposition (EMD) algorithm [19] is commonly used but suffers from severe drawbacks.

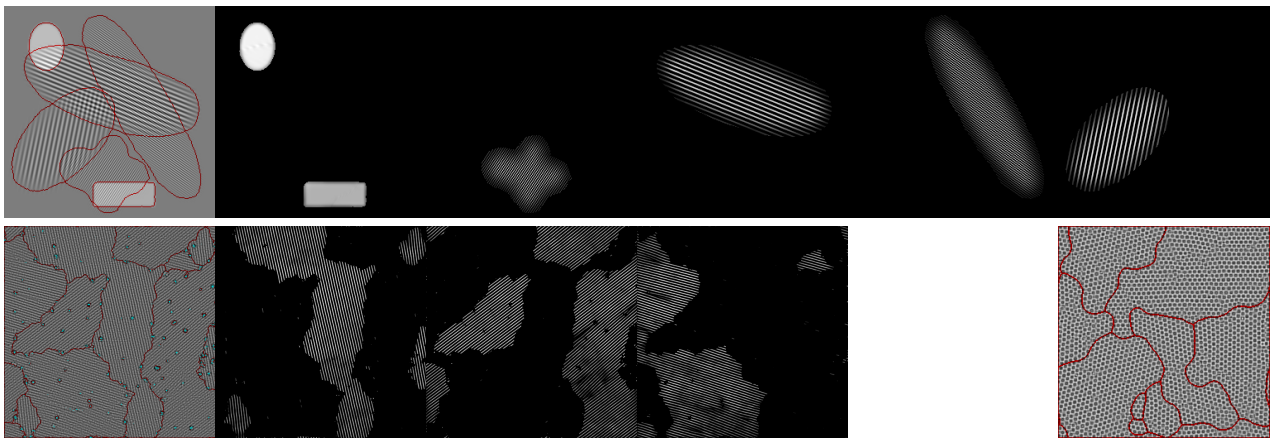

Figure 4: **Variational mode decomposition**. Top: composite 2D input image with detected mode boundaries (red); the 5 modes. Bottom left: Atomic force microscopy image of peptide β-sheets input, with detected grain boundaries (red) and artifacts (cyan); the 3 modes. Bottom right: segmentation boundaries (red) of colloidal crystal grain image by coupling of three sub-modes.

Here, we propose an entirely non-recursive variational mode decomposition model, where these modes are extracted concurrently. The model looks for an ensemble of K modes $u_k \colon \Omega \subset \mathbb{R} \to \mathbb{R}$ and their respective center frequencies ω_k, such that the modes collectively reproduce the input signal $f \colon \Omega \to \mathbb{R}$:

$$\min_{u_k,\ \omega_k} \quad \sum_k \left\| \partial_t \left[\left(\delta(t) + \tfrac{j}{\pi t}\right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 \quad \text{s.t.} \quad \forall t \colon \sum_k u_k(t) = f(t). \tag{13}$$

We then generalize this decomposition model to higher signal dimensions $\Omega \subset \mathbb{R}^n$ as follows:

$$\min_{u_k,\ \omega_k,\ A_k} \quad \sum_k \alpha_k \left\| \nabla [u_{AS,k}(\mathbf{x}) e^{-j\langle \omega_k, \mathbf{x} \rangle}] \right\|_2^2 + \lambda \sum_k TV(A_k) \quad \text{s.t.} \quad \forall \mathbf{x} \colon \sum_k A_k(\mathbf{x}) u_k(\mathbf{x}) = f(\mathbf{x}). \tag{14}$$

The analytic signal in 1D has single-sided Fourier spectrum, and we want to generalize this property to higher dimensions by defining $\hat{u}_{AS,k}(\omega) := (1 + \operatorname{sgn}(\langle \omega, \omega_k \rangle)) \hat{u}_k(\omega)$. Further, we include the binary functions $A_k \colon \Omega \to \{0, 1\}$, to allow for sparse spatial support of the bandlimited modes u_k. Coupling several submodes with a single support function enables more complex patterns, such as crystal lattices, to be segmented.

Contributions: If we constrain the modes to have mutually exclusive support, *i.e.*, $\forall \mathbf{x} \colon \sum_k A_k(\mathbf{x}) = 1$, then we obtain an image segmentation model based on textural coherence of the regions in Fourier-spectral terms—piecewise-constant frequency—, and through the co-area formula the TV-term effectively acts as the boundary length regularizer. This effectively implements Chan-Vese segmentation [20] based on Fourier features. To optimize this model, we make use of efficient threshold dynamics and MBO-like strategies [21, 22, 23]. Examples of variational mode decomposition problems are shown in figure 4.

2.2 Variational image segmentation with artifact detection and bias correction [24]

Intensity-based image segmentation has essentially been solved by the famous Chan-Vese model [20]:

$$E_{CV}(c_1, c_2, \phi) = \int_\Omega H(\phi)(I - c_1)^2 + (1 - H(\phi))(I - c_2)^2 + \mu \int_\Omega |\nabla H(\phi)|, \tag{15}$$

where ϕ is the evolving level set function, c_1 and c_2 are the mean intensities in each region (updated during iterations), and the last term favors regular boundaries. However, this model fails when images are affected by artifacts (scars, scratches) and illumination bias that outweigh the actual image contrast [25, 26]. Here, we introduce a new model for segmenting such damaged images: First, we introduce a dynamic artifact class $X \colon \Omega \to \{0, 1\}$ with cost γ, which prevents intensity outliers from skewing the segmentation. Second, in Retinex-fashion, we decompose the image into a flattened structural and a smooth bias part, $I = B + S$. Based on the phase-field approximation, $u \colon \Omega \to [0, 1]$ instead of $H(\phi)$, the energy functional is

$$\begin{aligned} E_{CV-X}(c_1, c_2, u, S, B, X) = \int_\Omega (1 - X) \left(u(S - c_1)^2 + (1 - u)(S - c_2)^2 \right) + \mu \int_\Omega |\nabla u| \\ + \alpha \int_\Omega |\nabla B|^2 + \beta \int_\Omega (I - B - S)^2 + \langle \lambda, I - B - S \rangle + \gamma \int_\Omega X. \end{aligned} \tag{16}$$

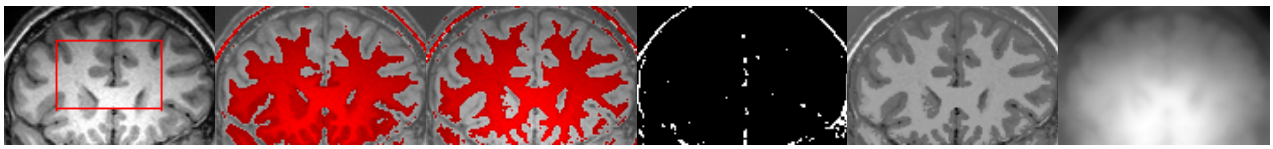

(a) Input (b) Classic CV (c) Proposed (d) Artifacts X (e) Structure S (f) Bias B

Figure 5: **Artifact detection and bias correction**. (a) Input image affected by strong illumination bias, and initial segmentation (red frame). (b) Classic CV result: bias causes over-segmentation of white matter in the ventral part, and under-segmentation in the dorsal part. (c-f) Proposed model: The structure S is clearly bias-corrected ("flat"), and the resulting segmentation of much better quality.

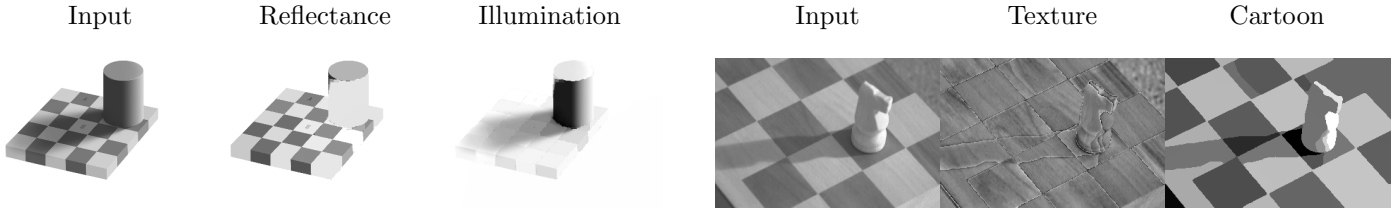

Figure 6: **Unifying non-local Retinex model**. Retinex is a family of models to decompose illumination and reflectance. Our unifying model performs, for example, shadow-removal (left) as well as cartoon-texture decomposition (right).

Contributions: To solve the proposed model, we devise an efficient alternate direction minimization scheme using time-split MBO-like threshold dynamics for the TV-term $|\nabla u|$ [23]. The resulting algorithm has very competitive performance, and provides a great improvement of segmentation quality over existing region-based segmentation models. See figure 5.

2.3 Non-local Retinex [27, 28]

Retinex is a theory on the human visual perception attempting to explain how the human visual system is capable of adaptively coping with spatially varying illumination [29, 30]. In image processing, the retinex theory has been implemented in various different flavors, each adapted to specific tasks, including color balancing, contrast enhancement, dynamic range compression and shadow removal in consumer electronics and imaging, bias field correction in medical imaging or even illumination normalization, *e.g.*, for face detection [31, 32, 33, 34, 35, 36, 37, 38, 39].

In this project, I develop a unifying framework for retinex that is able to reproduce many of the existing retinex implementations within a single model, including all gradient-fidelity based models, variational models [31], and kernel-based models [40, 41]. The fundamental assumption, as shared with many retinex models, is that the observed image I is the product of illumination B and underlying reflectance R of the object, $I = B \cdot R$. Taking the logarithm $i = \log(I)$, etc., splits the observed image into a sum of contributions, $i = b + r$. In Morel's PDE Retinex-model [42], illumination is supposed to vary smoothly and the reflectance is recovered from a hard-thresholded Laplacian of the observed image, $\Delta \hat{r} \approx \delta_\tau \Delta i$.

Here, we define our unifying Retinex model in similar but more general two steps: We reinterpret the gradient thresholding as variational models with sparsity constraints. First, we look for a filtered non-local gradient $(\nabla_{w,f} i)(\mathbf{x}, \mathbf{y}) := \sqrt{w(\mathbf{x}, \mathbf{y})} f(i(\mathbf{y}) - i(\mathbf{x}))$, $\mathbf{x}, \mathbf{y} \in \Omega$, that is the solution of an optimization problem consisting of two terms: a sparsity prior of the reflectance and a fidelity prior of the reflectance gradient to the observed image gradient. Second, since this filtered gradient almost certainly is not a consistent image gradient, we then fit an actual reflectance gradient to it, subject to further sparsity and fidelity priors:

$$\hat{r} = \arg\min_r \quad \|\nabla_w r - \nabla_{w,f} i\|_p^p + \alpha \|r\|_2^2 + \beta \|r - i\|_2^2 . \tag{17}$$

Contributions: This generalized formulation includes Morel's PDE model, but more importantly so it allows making connections with other, variational or kernel-based Retinex implementations. We provide simple algorithms for the optimization of our functional. In the quadratic case $p = 2$, we can link our model to a plausible neural mechanism through Wilson-Cowan equations. We derive novel Retinex flavors by using more interesting non-local versions for the sparsity and fidelity prior. Eventually, we define within a single framework new Retinex applications to shadow detection and removal, nonuniformity correction, cartoon-texture decomposition, as well as color and hyperspectral image enhancement, see examples in figure 6.

3 Earlier Work

3.1 Geometric image registration [43, 44, 45, 46, 47, 48]

In my Ph.D. project, I developed a novel geometric framework called *geodesic active fields* (GAF) for image registration. In image registration, one looks for the underlying deformation field that best maps one image onto another.This is a classic ill-posed inverse problem, which is usually solved by adding a regularization term. Here, I proposed a *multiplicative* coupling between the fidelity $f(\mathbf{u})$ with the uncanny Beltrami energy of the embedded deformation field $\mathbf{u}$ as regularizer:

$$\min_{\mathbf{u}} \quad \left\{ E_{GAF}(\mathbf{u}) := \int_\Omega f(\mathbf{u})\sqrt{g(\mathbf{u})} \right\} \qquad \text{with} \qquad \begin{cases} X\colon \Omega \subset \mathbb{R}^2 \to \Omega \times \mathbb{R}^2,\ (x,y) \mapsto (x,y,u_x,u_y) \\ h_{ij} := \mathrm{diag}(1,1,\beta^2,\beta^2) \\ g(\mathbf{u}) = 1 + \beta^2(|\nabla u_x|^2 + |\nabla u_y|^2) + \beta^4 |\nabla u_x \times \nabla u_y|^2 \end{cases} \tag{18}$$

where $f(\mathbf{u})$ measures the current local misalignment between the two images under $\mathbf{u} = (u_x, u_y)$, *e.g.*, based on pixel intensity mismatch. For example we consider $f(\mathbf{u}) := (F(\mathbf{x}) - M(\mathbf{x}+\mathbf{u}))^2$, where F and M denote the pixel values of the fixed and moving image, respectively. We estimate the underlying deformation field by minimizing this energy functional. It is easy to see that thanks to this construction, data fidelity f and deformation field regularity $\sqrt{g}$ compete locally, which is in contrast to all additive regularization schemes, where the tradeoff happens globally.

Contributions: This is a *weighted minimal surface problem* similar to the geodesic active contours model in image segmentation [49]. It is the only re-parametrization invariant registration scheme, and applies to images defined on arbitrary Riemannian manifolds, including omnidirectional and multi-scale images.

3.2 Geometric image segmentation with Harmonic Active Contours [50, 51]

As a more fancy application of the Beltrami framework, we propose a segmentation method purely based on the geometric representation of images as surfaces embedded in a higher dimensional space, enabling us to naturally work with multi-channel images. The segmentation is based on an active contour, described as the zero level set of the level set function $\phi\colon \Omega \to \mathbb{R}$, embedded in the image manifold along with image features $f^1, \ldots, f^k$:

$$\min_{\phi} \quad \left\{ E_{HAC}(\phi) := \int_\Omega \sqrt{g(\phi)} \right\} \qquad \text{with} \qquad \begin{cases} X\colon \quad \Omega \to \Omega \times \mathbb{R}^{k+1},\ (x,y) \mapsto (x,y,f^1,\ldots,f^k,\phi) \\ h_{ij} := \mathrm{diag}(1,1,\alpha_1^2,\ldots,\alpha_k^2,\beta^2) \\ g(\phi) = 1 + \sum \alpha_i |\nabla f^i|^2 + \beta^2 |\nabla \phi|^2 \\ \qquad\quad + \tfrac{1}{2}\sum \alpha_i^2 \alpha_j^2 |\nabla f^i \times \nabla f^j|^2 + \sum \alpha_i^2 \beta^2 |\nabla f^i \times \nabla \phi|^2 \end{cases} \tag{19}$$

This approach is purely geometrical and does not require additional weighting of the energy functional to drive the segmentation task to the image contours. Indeed, the joint embedding of the image features and the level set both regularizes the level set and couples it to the image features (last term in $g(\phi)$).

Contributions: This constitutes a more general framework for image segmentation by defining the optimal segmentation function as the harmonic map minimizing the hyper-surface of the manifold. The proposed technique is therefore called *harmonic active contour* (HAC).

3.3 Efficient algorithm for the level set method [52]

The level set method (LSM) [53] is a popular technique for tracking moving interfaces *e.g.*, in fluid dynamics and computer vision. A general problem, that includes image segmentation, can be written as minimization $\phi\colon \Omega \to \mathbb{R}$, as follows: [54]:

$$\min_{\phi} \quad \int_\Omega w_b(\mathbf{x})|\nabla H(\phi)| + w_r(\mathbf{x}) H(\phi), \tag{20}$$

where w_b and w_r are boundary and region terms. The LSM is limited by the CFL condition and does not implicitly preserve a signed distance function. Here, we split variables and include the signed distance function property, $|\nabla\phi| = 1$ a.e., as a constraint:

$$\min_{\phi,\ \varphi,\ u,\ \mathbf{q},\ \mathbf{p}} \quad \int_\Omega w_b(\mathbf{x})|\mathbf{q}| + w_r(\mathbf{x})u \quad \text{s.t.} \quad u = H(\varphi),\ \ \phi = \varphi,\ \ \mathbf{q} = \nabla u,\ \ \mathbf{p} = \nabla\phi,\ \ |\mathbf{p}| = 1 \tag{21}$$

Contributions: We solve by alternating direction gradient descent, enforcing constraints through augmented Lagrangians (ADMM). Our algorithm is faster than redistancing and produces better results than quadratic penalization alone [55, 56].

4 Future Research: Vision and Objectives

My research is driven by the quest for better models and tools for image and data understanding. *At different levels of abstraction, I want to work on a broad spectrum of mathematical questions and problems related to image processing, computer vision, and data science, that all have their useful application in science and society.* Moreover, many machine learning tasks have a structure very similar to image processing problems, such as clustering or classification. It is therefore a logical choice to explore the mutual similarities between mathematical imaging and machine learning problems.

Specifically, I want to perform my research along the following directions (presented in arbitrary order):

1. **Primal-dual methods for PDE**: As shown in §1.6, the primal-dual optimization methods are also applicable to such elliptic PDE problems, and we expect to make a significant contribution to the state-of-the-art by exploring this idea. Indeed, these primal-dual methods yield neatly structured algorithms in a very disciplined way, and the resulting computational steps are of stupefying simplicity, with excellent computational performance. I want to explore the use of such primal-dual methods for various PDE problems; this includes also eigenproblems of the Schrödinger operator and compressed modes, where the constraints are non-convex. Equally, I'm interested in the extension of the method to non-convex functionals in general, so as to greatly enlarge the domain of applicability. Finally, I am exploring the use of "primal-double-dual" schemes in cases where both functionals $f(Ax)$ and $g(Bx)$ involve linear operators. Preliminary experiments suggest that this can have tremendous implications for compressed sensing and image deconvolution.

2. **Geometric graph partitioning**: I want to analyze some simple problems, and work on numerical and computational solutions to larger problems. There is an exciting connection of our rearrangement algorithm with existing schemes for mean curvature motion (such as the MBO scheme). We want to explore our graph-partitioning method to revisit problems in image segmentation, potentially unifying region- and contour-based segmentation, and seamlessly extending to multiphase segmentation. We expect contributions to image segmentation, document classification, topic modeling, or threat/community detection in social networks.

3. **Non-local minimal surfaces and surfaces on graphs**: The generalization of the Beltrami energy to graphs still holds many secrets to be unveiled. In particular, I am interested in making a generalization to graphs not merely by rewriting the functional in terms of non-local differential operators, but through an analogous derivation from a theory of "non-local differential geometry."

4. **Unconventional image segmentation problems**: I want to continue exploring models for "unconventional image segmentation." 2D-extensions of the Variational Mode Decomposition will be useful for texture-based segmentation, with relations to holography. Preliminary results on both segmentation models are exciting. I want to further investigate the possibilities of threshold dynamics in the optimization algorithms, and study existence, uniqueness, and convergence. In combination with graph-partitioning, I want to explore different uncanny texture descriptors to define appropriate patch equivalence classes.

5. **Vesicles tracking in bio-microscopy**: Vesicles are substructures of cells and part of an intricate cellular transportation system. Understanding their dynamics through microscopy image sequences is key to numerous research questions. I want to develop a vesicle tracking method by exploring the similarity to imaging through turbulence: vesicles can be modeled as single-pixel seed blurred by a smoothing kernel. The seeds are sparse in the image sequence, modulo motion. Once trajectories extracted, we analyze their properties and infer characteristics about the underlying motion.

Due to their translational nature, my research interests in biomedical imaging, computer vision and machine learning are demonstrably fundable by various agencies, including NSF, NIH, and DoD, and I am looking forward to start new ventures and collaborations. Indeed, I am constantly engaged in more interdisciplinary, application-driven research, in collaboration with different partners in the field of biomedical imaging at large, and often involving undergraduate and graduate student researchers. In all these applications, my mathematical imaging contributions enable ground-breaking research and high-profile publications. I value these collaborations because they are always a source of "good problems," because the interaction and dialog with researchers from different fields is always enriching and inspiring, and because I know that the solutions I develop in this context will eventually benefit science and society.

Bottom line: Answering upcoming imaging problems still requires a lot of original research in applied mathematics, mathematical modeling intuition, and the synthesis of powerful existing and novel mathematical concepts. I want to continue working on inverse problems in imaging, and related problems on graphs and PDE. I am strongly willing to further enlarge my toolbox in applied mathematics, and dedicate myself to research for new image understanding and data science solutions.

“ Taking images and collecting data is easy. Understanding them is the challenge. ”

References

[1] M. Zhu and T. Chan, "An efficient primal-dual hybrid gradient algorithm for total variation image restoration," UCLA CAM Report 08-34, Tech. Rep., 2008.

[2] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, Dec. 2011.

[3] K. J. Arrow, L. Hurwicz, and H. Uzawa, *Studies in linear and non-linear programming.* Cambridge Univ. Press, 1958.

[4] D. Zosso, M. M. Xia, B. Osting, and S. J. Osher, "The obstacle problem revisited (again)," *in preparation*, p. 20, 2015.

[5] H. Attouch, G. Buttazzo, and G. Michaille, *Variational analysis in Sobolev and BV spaces: applications to PDEs and optimization*, 2nd ed. SIAM, 2014.

[6] L. A. Caffarelli, "The obstacle problem revisited," *The Journal of Fourier Analysis and Applications*, vol. 4, no. 4-5, pp. 383–402, Jul. 1998.

[7] A. Friedmann, *Variational principles and free-boundary problems.* New York: Wiley, 1982.

[8] J.-F. Rodrigues, *Obstacle problems in mathematical physics.* Amsterdam: Elsevier Science Publishers, 1987.

[9] G. Tran, H. Schaeffer, W. M. Feldman, and S. J. Osher, "An L^1 penalty method for general obstacle problems," CAM Report 14-27, Tech. Rep., 2014.

[10] D. Zosso and A. Bustin, "A primal-dual projected gradient algorithm for efficient Beltrami regularization," UCLA CAM Report 14-52, Tech. Rep., 2014.

[11] N. Sochen, R. Kimmel, and R. Malladi, "A general framework for low level vision," *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 310–318, 1998.

[12] A. Bustin, M. A. Janich, A. C. Brau, F. Odille, S. D. Wolff, O. Shubayev, D. Stanley, and A. Menini, "Joint denoising and motion correction: initial application in single-shot cardiac MRI," *Journal of Cardiovascular Magnetic Resonance*, vol. 17, no. Suppl 1, p. Q29, 2015.

[13] D. Zosso and B. Osting, "A minimal surface criterion for graph partitioning," *submitted to: AIMS J. Inverse Problems and Imaging*, p. 30, 2015.

[14] D. Zosso, B. Osting, and S. J. Osher, "A Dirichlet energy criterion for graph-based image segmentation," in *submitted to: ICDM Workshop "Graph Algorithms for Imaging and Networks"*, 2015, p. 10.

[15] ——, "Elliptic PDE and primal-dual optimization methods," *in preparation*, 2015.

[16] K. Dragomiretskiy and D. Zosso, "Variational mode decomposition," *IEEE Transactions on Signal Processing*, vol. 62, no. 3, pp. 531–544, Feb. 2014.

[17] ——, "Two-dimensional variational mode decomposition," in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, ser. Lecture Notes in Computer Science, X.-C. Tai, E. Bae, T. F. Chan, and M. Lysaker, Eds., vol. 8932. Springer International Publishing, 2015, pp. 197–208.

[18] D. Zosso, K. Dragomiretskiy, A. L. Bertozzi, and P. S. Weiss, "Two-dimensional compact variational mode decomposition: spatially compact and spectrally sparse image decomposition and segmentation," *in preparation*, p. 24, 2015.

[19] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu, "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1971, pp. 903–995, Mar. 1998.

[20] T. F. Chan and L. A. Vese, "Active contours without edges," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 266–277, 2001.

[21] B. Merriman, J. K. Bence, and S. Osher, "Diffusion generated motion by mean curvature," UCLA CAM Report 92-18, Tech. Rep., 1992.

[22] ——, "Diffusion generated motion by mean curvature," *AMS Selected Letters, Crystal Grower's Workshop*, pp. 73–83, 1993.

[23] S. Esedoglu and Y. H. R. Tsai, "Threshold dynamics for the piecewise constant Mumford-Shah functional," *Journal of Computational Physics*, vol. 211, pp. 367–384, 2006.

[24] D. Zosso, J. An, J. Stevick, N. Takaki, M. Weiss, L. S. Slaughter, H. H. Cao, P. S. Weiss, and A. L. Bertozzi, "Image segmentation with dynamic artifacts detection and bias correction," *submitted to: AIMS J. Inverse Problems and Imaging*, p. 24, 2015.

[25] T. Brox and D. Cremers, "On local region models and a statistical interpretation of the piecewise smooth Mumford-Shah functional," *International Journal of Computer Vision*, vol. 84, no. 2, pp. 184–193, Jul. 2008.

[26] C. Li, C.-Y. Kao, J. C. Gore, and Z. Ding, "Minimization of region-scalable fitting energy for image segmentation," *IEEE Transactions on Image Processing*, vol. 17, no. 10, pp. 1940–1949, Oct. 2008.

[27] D. Zosso, G. Tran, and S. Osher, "A unifying Retinex model based on non-local differential operators," in *IS&T/SPIE Electronic Imaging*, C. A. Bouman, Ed., Feb. 2013, pp. 865 702–1–16.

[28] D. Zosso, G. Tran, and S. J. Osher, "Non-local Retinex—a unifying framework and beyond," *SIAM J. Imaging Sciences*, vol. 8, no. 2, pp. 787–826, 2015.

[29] E. H. Land, "The Retinex," *American Scientist*, vol. 52, no. 2, pp. 247–253,255–264, 1964.

[30] E. H. Land and J. J. McCann, "Lightness and Retinex theory," *Journal of the Optical Society of America*, vol. 61, no. 1, pp. 1–11, Jan. 1971.

[31] R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel, "A variational framework for Retinex," *International Journal of Computer Vision*, vol. 52, no. 1, pp. 7–23, 2003.

[32] Z.-u. Rahman, D. J. Jobson, and G. A. Woodell, "Retinex processing for automatic image enhancement," *Journal of Electronic Imaging*, vol. 13, no. 1, p. 100, 2004.

[33] E. Provenzi, M. Fierro, A. Rizzi, L. De Carli, D. Gadia, and D. Marini, "Random spray Retinex: A new Retinex implementation to investigate the local properties of the model," *IEEE Transactions on Image Processing*, vol. 16, no. 1, pp. 162–171, Jan. 2007.

[34] L. Lei, Y. Zhou, and J. Li, "An investigation of Retinex algorithms for image enhancement," *Journal of Electronics (China)*, vol. 24, no. 5, pp. 696–700, Sep. 2007.

[35] M. K. Ng and W. Wang, "A Total Variation model for Retinex," *SIAM Journal on Imaging Sciences*, vol. 4, no. 1, pp. 345–365, Jan. 2011.

[36] R. Sobol, "Improving the Retinex algorithm for rendering wide dynamic range photographs," *Journal of Electronic Imaging*, vol. 13, no. 1, p. 65, 2004.

[37] W. Ma, J.-M. Morel, S. Osher, and A. Chien, "An L1-based variational model for Retinex theory and its application to medical images," in *CVPR 2011*. IEEE, Jun. 2011, pp. 153–160.

[38] G. D. Finlayson, S. Hordley, and M. Drew, "Removing shadows from images using retinex," *10th Color Imaging Conference*, 2002.

[39] T. Chen, W. Yin, X. S. Zhou, D. Comaniciu, and T. S. Huang, "Total variation models for variable lighting face recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 9, pp. 1519–1524, 2006.

[40] D. J. Jobson, Z.-u. Rahman, and G. A. Woodell, "A multiscale retinex for bridging the gap between color images and the human observation of scenes," *IEEE Transactions on Image Processing*, vol. 6, no. 7, pp. 965–76, Jan. 1997.

[41] R. Palma-Amestoy, E. Provenzi, M. Bertalmío, and V. Caselles, "A perceptually inspired variational framework for color enhancement," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 3, pp. 458–74, Mar. 2009.

[42] J.-M. Morel, A.-B. Petro, and C. Sbert, "A PDE formalization of Retinex theory," *IEEE Transactions on Image Processing*, vol. 19, no. 11, pp. 2825–2837, May 2010.

[43] D. Zosso and J.-P. Thiran, "A scale-space of cortical feature maps," *IEEE Signal Processing Letters*, vol. 16, no. 10, pp. 873–876, 2009.

[44] ——, "Geodesic active fields on the sphere," in *International Conference on Pattern Recognition, ICPR2010*, 2010.

[45] D. Zosso, X. Bresson, and J.-P. Thiran, "Geodesic active fields—a geometric framework for image registration," *IEEE Transactions on Image Processing*, vol. 20, no. 5, pp. 1300–1312, Nov. 2011.

[46] ——, "Geodesic active fields—a geometric framework for image registration," in *8th ICNAAM 2010, International Conference of Numerical Analysis and Applied Mathematics*, ser. AIP Conference Proceedings, T. Simos, Ed., vol. 1281, 2010, pp. 1027–1030.

[47] D. Zosso, "Geodesic active fields: A geometric framework for image registration," PhD Thesis, École Polytechnique Fédérale de Lausanne, 2011. [Online]. Available: http://library.epfl.ch/theses/?nr=5175

[48] D. Zosso, X. Bresson, and J.-P. Thiran, "Fast geodesic active fields for image registration based on splitting and augmented Lagrangian approaches," *IEEE Transactions on Image Processing*, vol. 23, no. 2, pp. 673–683, Mar. 2014.

[49] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. J. Comput. Vis.*, vol. 22, no. 1, pp. 61–79, 1997.

[50] V. Estellers, D. Zosso, X. Bresson, and J. P. Thiran, "Harmonic active contours for multichannel image segmentation," in *Proceedings - International Conference on Image Processing, ICIP*, Sep. 2011, pp. 3141–3144.

[51] V. Estellers, D. Zosso, X. Bresson, and J.-P. Thiran, "Harmonic active contours," *IEEE Transactions on Image Processing*, vol. 23, no. 1, pp. 69–82, Jan. 2014.

[52] V. Estellers, D. Zosso, R. Lai, S. Osher, J.-P. Thiran, and X. Bresson, "Efficient algorithm for level set method preserving distance function," *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp. 4722–4734, Dec. 2012.

[53] S. Osher and J. A. Sethian, "Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, vol. 79, no. 1, pp. 12–49, Nov. 1988.

[54] M. Burger and S. J. Osher, "A survey on level set methods for inverse problems and optimal design," *European journal of applied mathematics*, vol. 16, no. 02, pp. 263–301, 2005.

[55] C. Li, C. Xu, C. Gui, and M. D. Fox, "Level set evolution without re-initialization: a new variational formulation," in *IEEE CVPR*, vol. 1. IEEE, 2005, pp. 430–436.

[56] ——, "Distance regularized level set evolution and its application to image segmentation," *IEEE Transactions on Image Processing*, vol. 19, no. 12, pp. 3243–3254, Dec. 2010.