Accelerated Optimization in the PDE Framework

Anthony Yezzi

Georgia Institute of Technology

School of ECE

Presenting joint work with:

Minas Beyamin, Jeff Calder, Ganesh Sundaramoorthi

- Infinite dimensional flat spaces[a] [b]
  - Denoising
  - Inpainting
  - Deblurring
  - Obstacle Problem
- More general infinite dimensional manifolds[c] [d]
  - Active Contours (visual tracking and image segmenation)
  - Active Surfaces (multiview stereo and radar reconstruction)
  - Diffeomorphisms (nonrigid image registration)

---

[a] "Accelerated variational PDE's for efficient solution of regularized inversion problems," (M. Benyamin, J. Calder, G. Sundaramoorthi, A. Yezzi), *J. Mathematical Imaging and Vision*, vol. 62, Jan. 2020, pp. 10–36.

[b] "PDE Acceleration: A convergence rate analysis and applications to obstacle problems," (J. Calder and A. Yezzi), *Research in Mathematical Sciences*, vol. 6, Dec. 2019, pp. 1–35.

[c] "PDE Acceleration for Active Contours," (A. Yezzi, G. Sundaramoorthi, and M. Benyamin), *IEEE Conf. Computer Vision and Pattern Recognition*, June 2019, pp. 12318–12329, Long Beach, CA.

[d] "Variational PDEs for Acceleration on Manifolds and Application to Diffeomorphisms," (G. Sundaramoorthi and A. Yezzi), *IEEE Conf. Neural Information Processsing Systems* Dec. 2018, pp. 3797–3807.

Early motivation: myocardial segmentation with trained 3D shapes

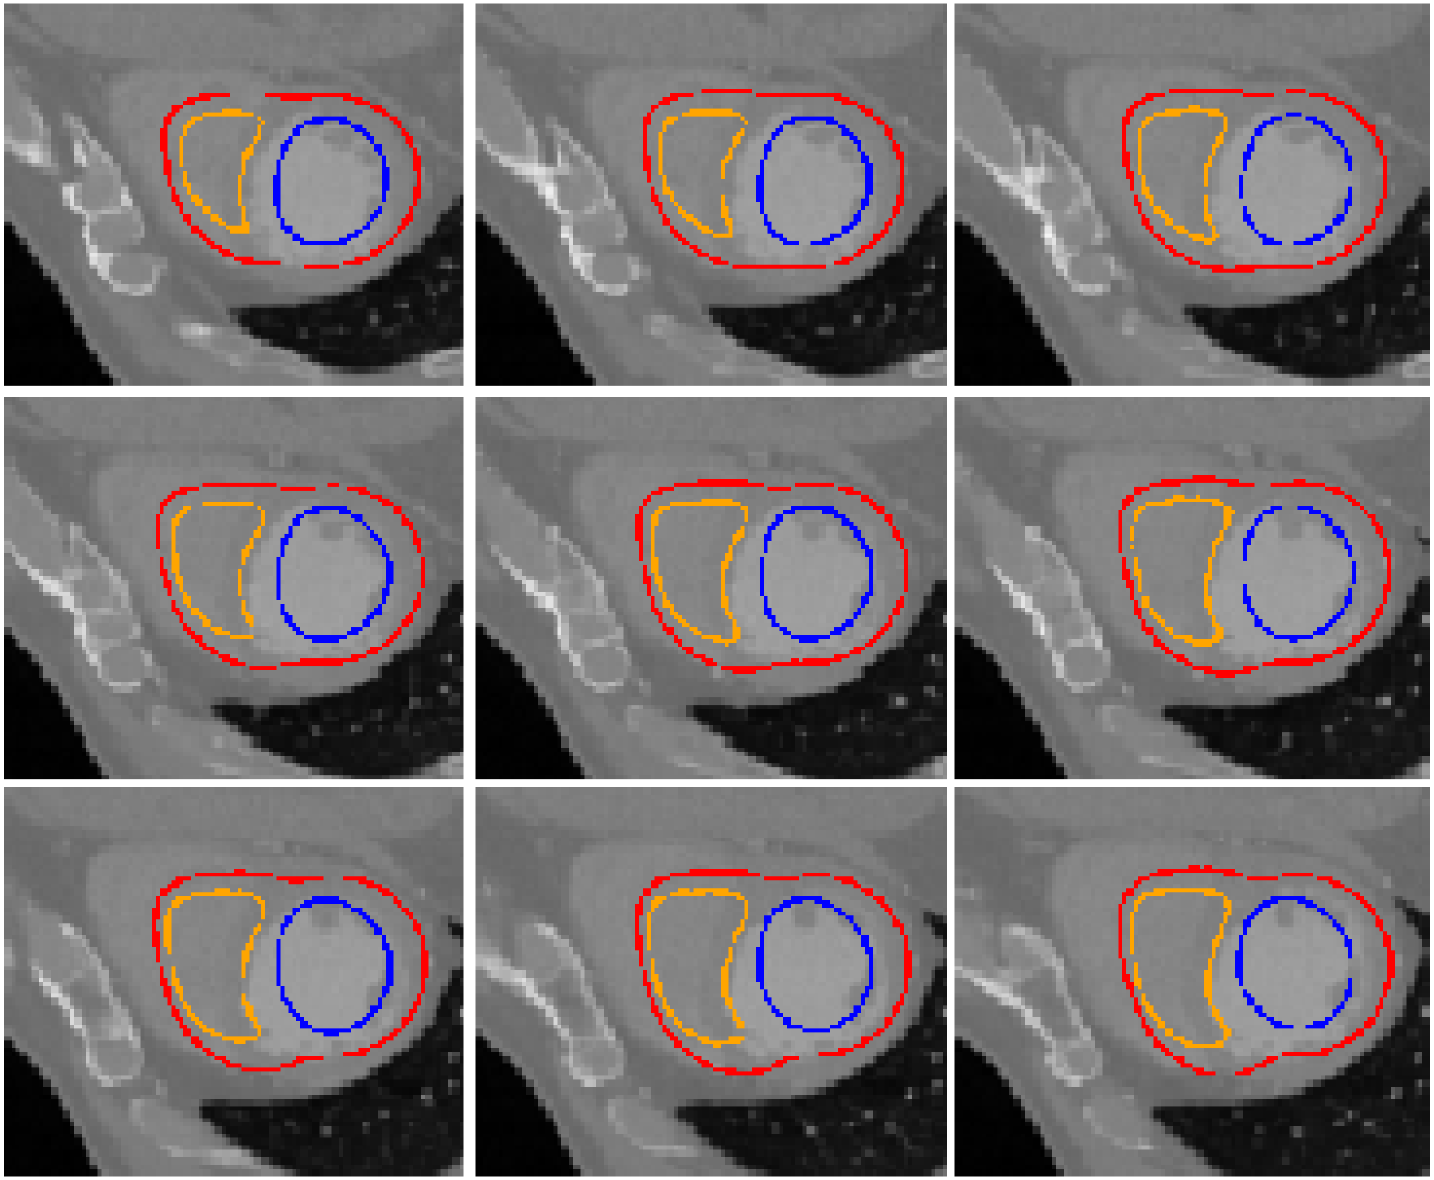

Initialization via training shape averages

Nonconvex optimization model

$$\text{Image Fitting Cost} = w_{LV} \int_{LV} (I - \mu_{LV})^2 dx + w_{RV} \int_{RV} (I - \mu_{RV})^2 dx$$

$$+ w_{myo} \int_{myo} (I - \mu_{myo})^2 dx + w_{BG} \int_{BG} \min\left((I - \mu_{lo})^2, (I - \mu_{hi})^2\right) dx$$

Subject to inequality constraints: $\mu_{lo} \le \mu_{myo} \le \mu_{RV} \le \mu_{LV} \le \mu_{hi}$

$$\text{Region Overlap Penalty} = \int_{LV \cap RV} dx + \int_{LV \cap BG} dx + \int_{RV \cap BG} dx$$

Optimization result

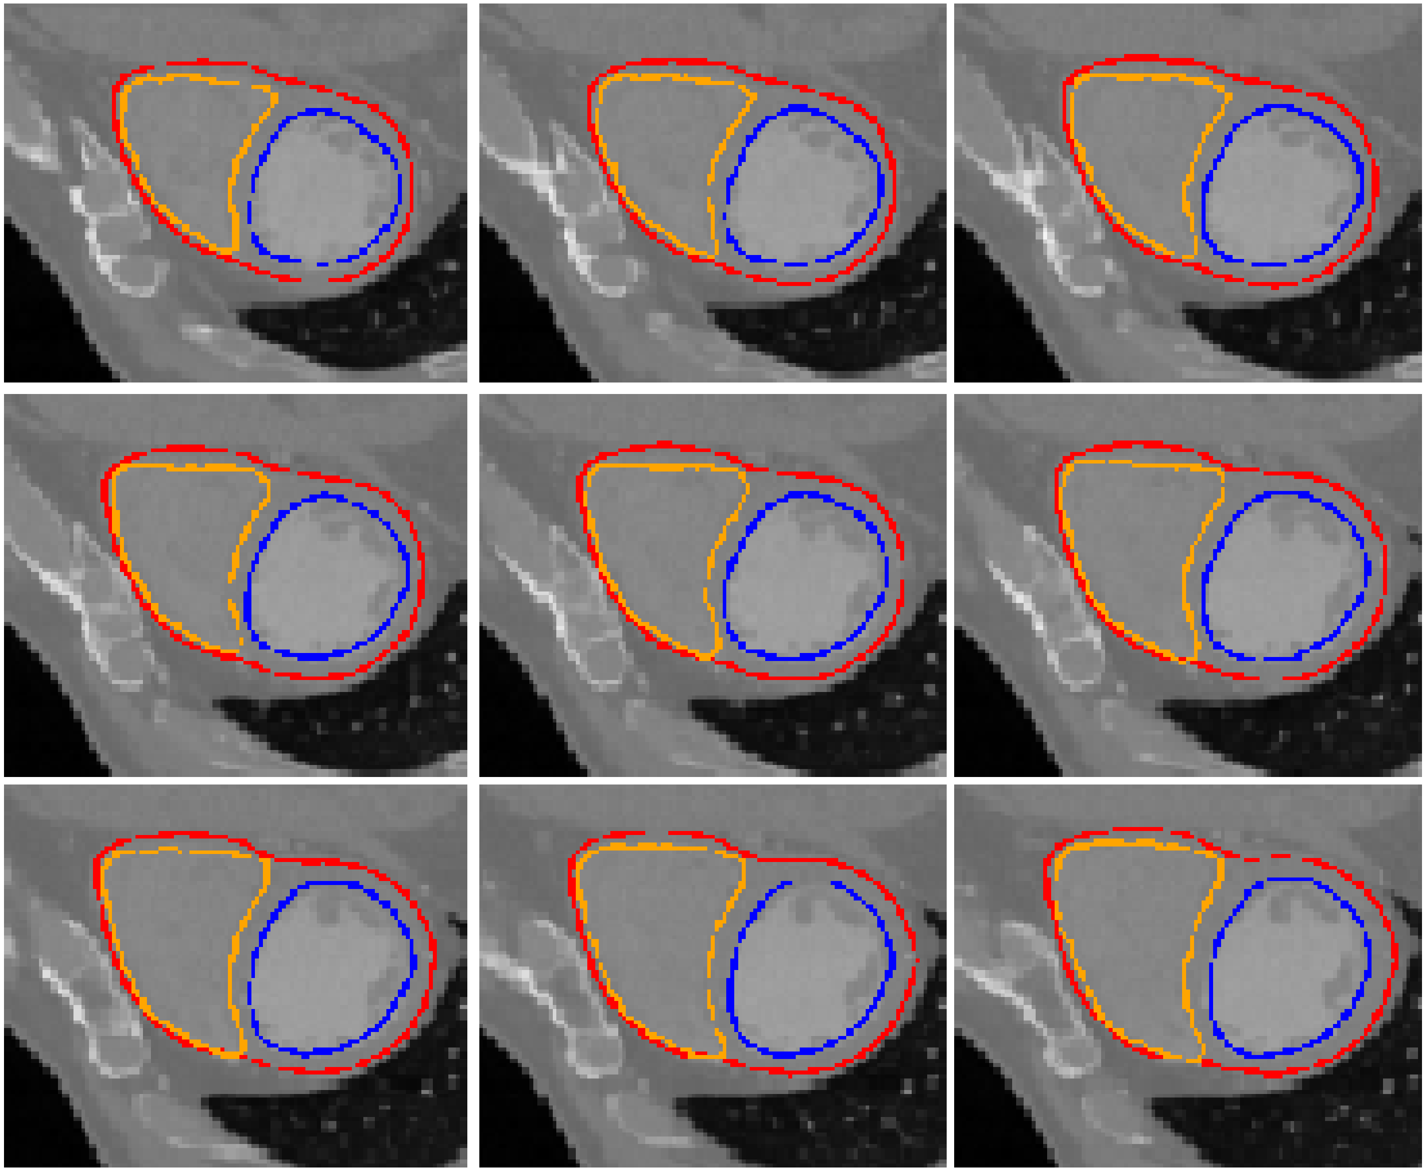

Manual versus Automatic

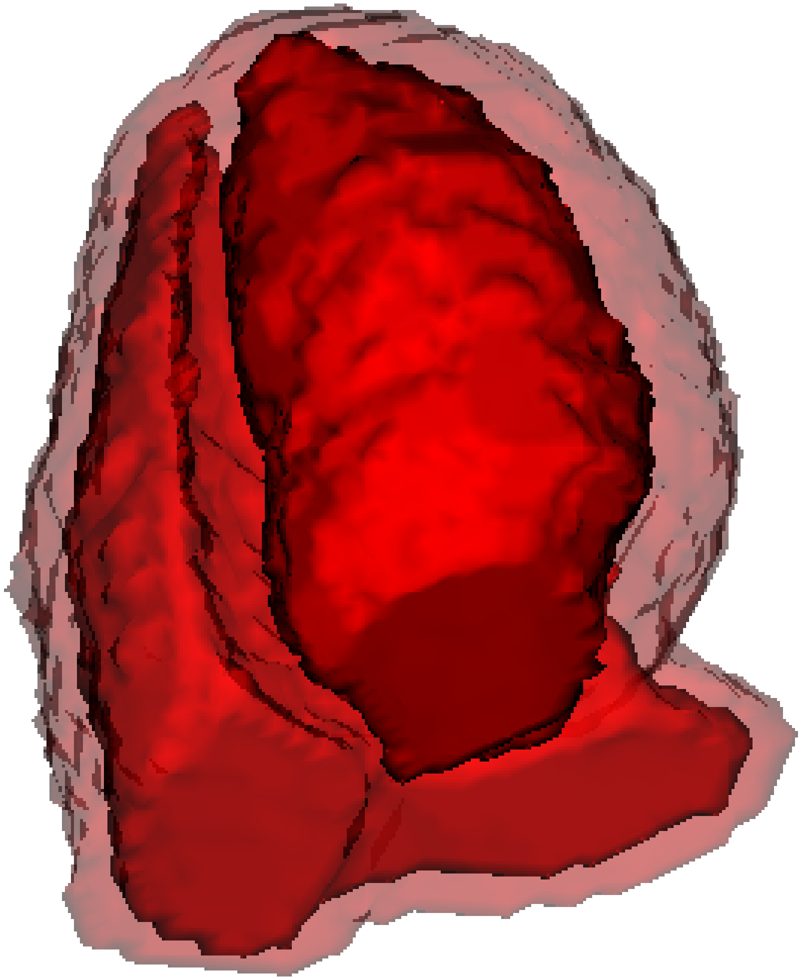

Cardiologist manual-traced result

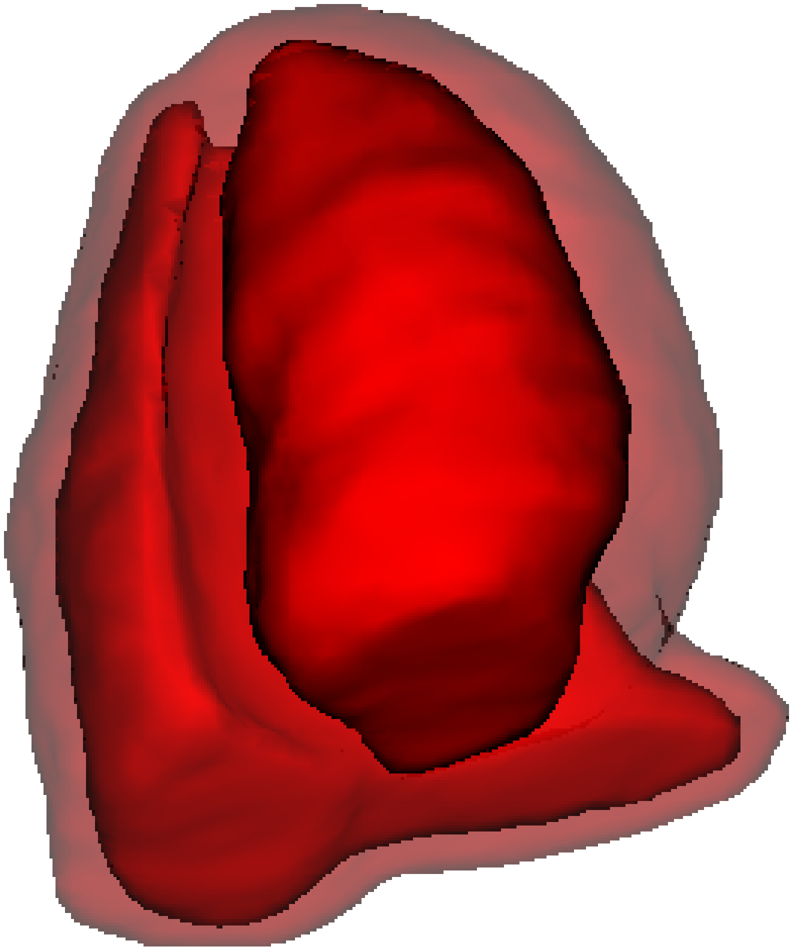

Automated algorithm result

Challenges faced by gradient descent

- Unwanted local minimizers
- Bouncing within narrow descent valleys.

Myocardial segmentation example

Gradient descent with momentum

When gradient descent iterations start bouncing back and forth across a descending narrow valley in the energy function, then the "average" of consecutive parameter increments will better approximate the descent direction **(for the same step size)**.

$$
\begin{aligned}
\text{search direction} \;=\;& \quad \lambda \quad\; \text{(prior search direction)} \\
+\;& (1-\lambda) \; \text{(gradient direction)}
\end{aligned}
$$

Myocardial segmentation example

Effect in myocardial segmentation scheme

Converged 3D shape-trained myocardial segmentation results (2D cross-sections shown) both with and without gradient momentum.

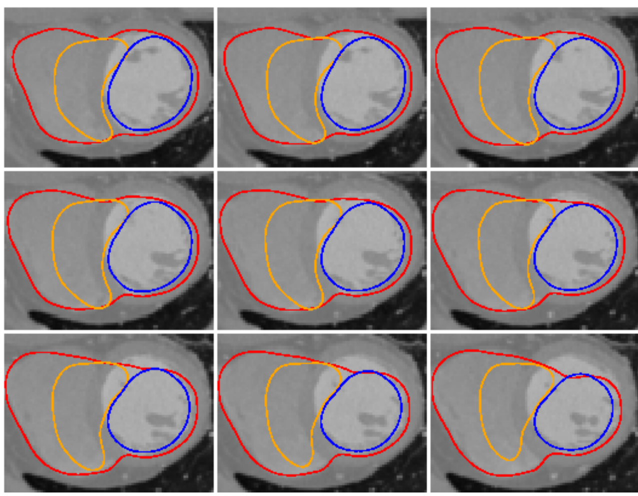

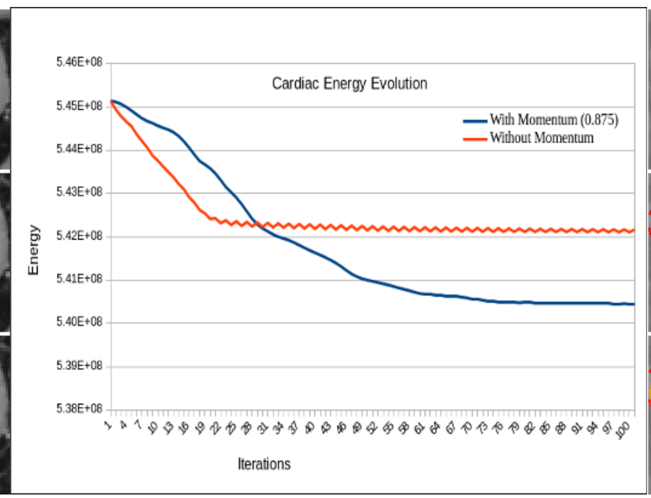

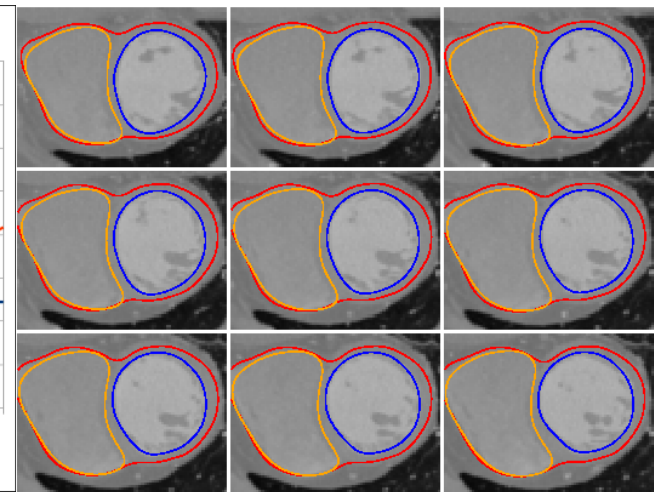

Without momentum | Compared energy descent | With momentum (0.875)

Nesterov's Accelerated Gradient Descent

Strategic, dynamically changing weights on the momentum term can further boost the descent process. In the case of a smooth, convex function $E(x)$, Nesterov put forth the following scheme which attains a rate of order $\frac{1}{t^2}$: $E(y_k) - E(x^*) \leq \frac{2\beta\|x_1 - x^*\|^2}{t^2}$

$$
\begin{aligned}
y_{k+1} &= x_k - \frac{1}{\beta}\nabla E(x_k) \\
x_{k+1} &= \underbrace{(1-\gamma_k)}_{\text{overshoot}} y_{k+1} + \gamma_k y_k
\end{aligned}
$$

where

$$
\begin{aligned}
\gamma_k &= \frac{1-\lambda_k}{\lambda_k + 1} \qquad [\text{always} \leq 0] \\
\lambda_k &= \frac{1 + \sqrt{1 + 4\lambda_{k-1}^2}}{2} \qquad [\lambda_0 = 0]
\end{aligned}
$$

Variational framework for accelerated gradient

Wibisono, Wilson, Jordan [NAS 2016] generalize Nesterov's and other momentum based gradient descent schemes in $\mathbb{R}^n$ based on the Bregman divergence of a convex distance generating function h

$$D(y,x) = h(y) - h(x) - \langle \nabla h(x), y - x \rangle$$

$\left(\textit{when } h(x) = \frac{1}{2}\|x\|^2 \textit{ this becomes the Euclidean distance } \frac{1}{2}\|y-x\|^2 \right)$

and the Euler-Lagrange equation for the action integral associated with the following Bregman Lagrangian

$$\mathcal{L}(X,V,t) = e^{a(t)+\gamma(t)} \left[D(X + e^{-a(t)}V, X) - e^{b(t)}\mathbf{U}(X) \right]$$

where the potential energy $\mathbf{U}$ represents the cost to be minimized.

In the Euclidean case this simplifies to

$$\textit{generalized action} = e^{\gamma(t)} \left[e^{-a(t)} \underbrace{\frac{1}{2}\|V\|^2}_{\mathbf{T}} - e^{a(t)+b(t)}\mathbf{U}(X) \right]$$

where $\mathbf{T}$ models the kinetic energy of a unit mass particle in $\mathbb{R}^n$.

Nesterov's methods belong to a subfamily of Bregman Lagrangians with the following choice of parameters, indexed by $k > 0$.

$$a = \log k - \log t$$

$$b = k \log t + \log \lambda$$

$$\gamma = k \log t$$

- $k = 2$: Nesterov's accelerated mirror descent
- $k = 3$: Nesterov's cubic-regularized Newton method

Nesterov Generalized Action Integral

$$\int \frac{t^{k+1}}{k} \left(\mathbf{T} - \lambda k^2 t^{k-2} \mathbf{U} \right) \, dt$$

Note that the explicit time dependence in the Lagrangian means that we do not have conservation of energy (the Hamiltonian is not conserved). The resulting Euler-Lagrange equation will yield, among other terms, **frictional forces** that monotonically dissipate energy, thereby guaranteeing convergence.

Motivation from Physics *(classical action)*

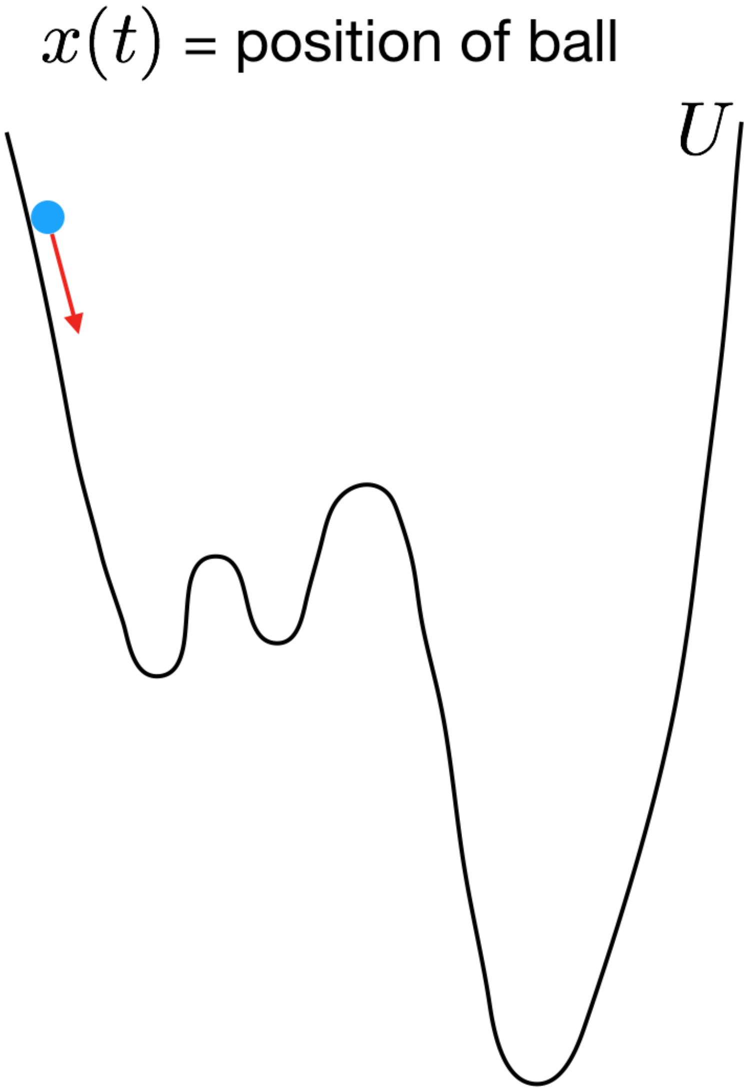

- Newton's law for conservative forces
  $F = \frac{d}{dt}\big(mv(t)\big), \quad F = -\nabla \mathbf{U}$
- Critical path for Action Integral
  $\int \frac{1}{2}mv^2 - \mathbf{U}$
- As the ball accelerates it
  - reaches a minimum faster
  - gains momentum to travel over shallow dips in the potential

Generalization to PDE Framework

We consider the calculus of variations problem

$$\min_u E[u] := \int_\Omega \Phi(x, \nabla u) + \Psi(x, u)\, dx.$$

The Euler-Lagrange equation satisfied by minimizers is

$$\nabla E[u] := \Psi_z(x, u) - \nabla \cdot (\nabla \Psi(x, \nabla u)) = 0, \tag{1}$$

where $\Phi = \Phi(x, p)$, $\nabla\Phi = \nabla_p \Phi$ and $\Psi = \Psi(x, z)$. We note that the L^2-gradient $\nabla E[u]$ satisfies

$$\left.\frac{d}{d\varepsilon}\right|_{\varepsilon=0} E[u + \varepsilon v] = \int_\Omega \nabla E[u]\, v\, dx \tag{2}$$

for all v smooth with compact support.

Kinetic Energy and Action Integral

We define the action integral

$$J[u] = \int_{t_0}^{t_1} k(t) \left(\frac{1}{2} \int_\Omega \rho u_t^2 \, dx - b(t) E[u] \right) \, dt, \tag{3}$$

where $k(t)$ and $b(t)$ are time dependent weights, $\rho = \rho(x)$ represents a mass density, and $u = u(x,t)$.

Therefore, the PDE accelerated descent equations are

$$\frac{\partial}{\partial t}(k(t)\rho u_t) = -k(t)b(t)\nabla E[u].$$

It is more convenient to define $a(t) = k'(t)/k(t)$ are rewrite

$$\boxed{u_{tt} + a(t)u_t = -b(t)\rho(x)^{-1}\nabla E[u].} \tag{4}$$

Gradient Descent = Infinite friction limit

Numerical Schemes for Accelerated PDE's

If we consider the explicit forward Euler discretization of the continuous gradient descent PDE we obtain

$$\frac{u(x,t+\Delta t)-u(x,t)}{\Delta t} = -\nabla E$$

This leads to the following simple discrete iteration

$$\begin{aligned} \Delta u^n(x) &= -\Delta t \, \nabla E^n \\ u^{n+1}(x) &= u^n(x) + \Delta u^n(x) \end{aligned} \tag{5}$$

Gradient descent stability constraint

Applying Von Nuemann analysis to the linearized version yields the update

$$
\begin{aligned}
U^{n+1}(\omega) &= \underbrace{(1 - \Delta t\, z(\omega))}_{\xi(\omega)} U^n(\omega) \\
z(\omega) &\doteq \frac{\mathrm{DFT}\left(\text{linearized\_homogeneous\_part\_of}\left(\nabla E^n\right)\right)}{\mathrm{DFT}\left(u^n\right)}
\end{aligned}
\qquad (6)
$$

which will be stable as long as the overall update amplification factor $\xi(\omega)$ does not have complex amplitude exceeding unity

$$\Delta t \leq \frac{z(\omega) + z^*(\omega)}{z(\omega) z^*(\omega)} = \frac{1}{z(\omega)} + \frac{1}{z^*(\omega)} = 2\Re\left(\frac{1}{z(\omega)}\right)$$

It is common in regularized optimization for the gradient amplifier to be real and non-negative: $z(\omega) \geq 0$, yielding the CFL condition

$$\Delta t \leq \frac{2}{z_{\max}} \tag{7}$$

where $z_{\max} \doteq \max_\omega z(\omega)$.

Accelerated descent stability constraint

Using central difference approximations for both time derivatives gives a second order discretization in time

$$\frac{u(x,t+\Delta t)-2u(x,t)+u(x,t-\Delta t)}{\Delta t^2}+a\frac{u(x,t+\Delta t)-u(x,t-\Delta t)}{2\Delta t}=-\nabla E(x,t)$$

which leads to the following update.

$$u^{n+1}(x)=\frac{2u^n(x)-\left(1-\frac{a\Delta t}{2}\right)u^{n-1}(x)-\Delta t^2\nabla E^n(x)}{1+\frac{a\Delta t}{2}}$$

In the case of real $z(\omega)$ we obtain the following CFL condition

$$\Delta t \leq \frac{2}{\sqrt{z_{\max}}}$$

Application to regularized inversion

Here we consider a very general class of variational regularized inversion problems in the accelerated PDE framework. In particular, we assume energy functions with the form

$$E(u) = \int_\Omega \underbrace{f\left(|\mathcal{K}u - g|\right)}_{\text{fidelity}} + \underbrace{r(\|\nabla u\|)}_{\text{regularity}} \; dx, \qquad \text{with } \dot{f}, \dot{r}, \ddot{r} > 0$$

where f is a monotonically increasing penalty on the residual error between data measurements g and a forward in the form of linear operator $\mathcal{K}$ applied to the reconstructed signal u, while r is a monotonically increasing penalty on the gradient of the reconstruction.

General case (nonlinear wave equation)

The continuum gradient of E has the form

$$\nabla E(u) = \frac{\dot{f}\left(|\mathcal{K}u - g|\right)}{|\mathcal{K}u - g|}\mathcal{K}^*\left(\mathcal{K}u - g\right) - \dot{r}(\|\nabla u\|)\,\nabla\cdot\left(\frac{\nabla u}{\|\nabla u\|}\right) - \ddot{r}(\|\nabla u\|)\,\frac{\nabla u^T\,\nabla^2 u\,\nabla u}{\|\nabla u\|^2}$$

$$= \underbrace{\frac{\dot{f}\left(|\mathcal{K}u - g|\right)}{|\mathcal{K}u - g|}}_{\lambda(u,x)>0}\mathcal{K}^*\left(\mathcal{K}u - g\right) - \underbrace{\frac{\dot{r}(\|\nabla u\|)}{\|\nabla u\|}}_{c(\nabla u)>0}\left(\nabla\cdot\nabla u - u_{\eta\eta}\right) - \underbrace{\ddot{r}(\|\nabla u\|)}_{d(\nabla u)>0}\,u_{\eta\eta}$$

This gives rise to the following class of accelerated flows which take the form of a nonlinear wave equation.

$$u_{tt} - c(\nabla u)\,\left(\nabla\cdot\nabla u - u_{\eta\eta}\right) - d(\nabla u)\,u_{\eta\eta} + a u_t = \lambda(u,x)\,\mathcal{K}^*\left(g - \mathcal{K}u\right)$$

Optimal damping in the linear case

When we have a linear PDE acceleration equation

$$u_{tt} + a u_t + Lu + \lambda u = f \quad \text{in } \Omega \times (0, \infty)$$

where L is a second order elliptic operator, a Fourier analysis leads to the optimal choice

$$a = 2\sqrt{\lambda_1 + \lambda}$$

where λ_1 is the first Dirichlet eigenvalue of L and the optimal convergence rate

$$|u(x, yt) - u^*(x)| < C\, e^{-at}$$

Quadratic regularization

The easiest special case to consider would be that of quadratic fidelity and regularity penalties without any forward model (more precisely, with $\mathcal{K}$ as the identity operator).

$$E(u) = \int_\Omega \frac{\lambda}{2}(u-g)^2 + \frac{c}{2}\|\nabla u\|^2 \, dx$$

The accelerated descent PDE therefore takes the form of a damped inhomogenous linear wave equation.

$$u_{tt} - c\,\nabla \cdot \nabla u + a u_t = \lambda(g-u) \tag{8}$$

Click Here for Demo

Beltrami regularization

$$E(u) = \int_\Omega \frac{\lambda}{2}\left(K * u - g\right)^2 + \underbrace{\frac{1}{\beta}\sqrt{1 + \|\beta \nabla u\|^2}}_{\sqrt{\epsilon^2 + \|\nabla u\|^2},\ \epsilon = \frac{1}{\beta}}\, dx \tag{9}$$

In this case the variational gradient is non-linear and the accelerated PDE takes the quasilinear form.

$$u_{tt} - \nabla \cdot \left(\frac{\beta \nabla u}{\sqrt{1 + \|\beta \nabla u\|^2}} \right) + a u_t = \lambda K^T * K * (g - u)$$

Conservative optimal damping estimate:

$$a = 2\sqrt{\beta \pi^2 + \lambda}$$

Total Variation regularization

If we consider the limit as $\beta \to \infty$, the Beltrami regularization penalty converges to the total variation penalty.

$$E(u) = \int_\Omega \frac{\lambda}{2} (K * u - g)^2 + \|\nabla u\| \, dx$$

with a non-linear variational gradient that decomposes as follows.

$$\nabla E = \lambda K^T * K * (u - g) - \nabla \cdot \left(\frac{\nabla u}{\|\nabla u\|} \right)$$

The accelerated PDE now takes the form

$$u_{tt} - \nabla \cdot \left(\frac{\nabla u}{\|\nabla u\|} \right) + a u_t = \lambda K^T * K * (u - g)$$

TV denoising

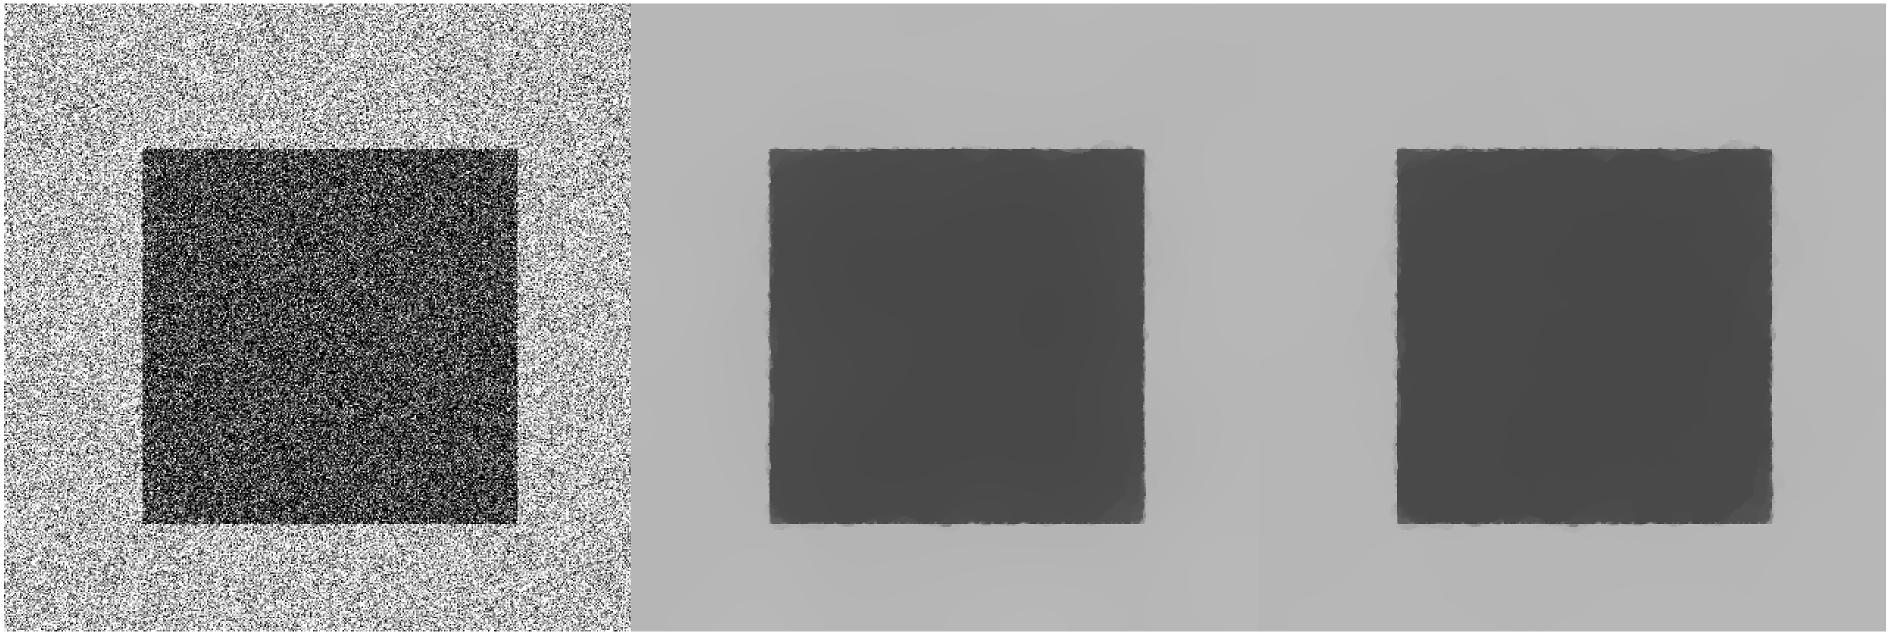

Noisy Square | Split Bregman | PDE acceleration

Figure 1: Denoising of a synthetic image with total variation restoration with $\lambda = 1000$ via (b) Split Bregman and (c) PDE acceleration. In PDE acceleration we used $\Delta t = \Delta x/2$ and $a = 6\sqrt{\lambda}$.

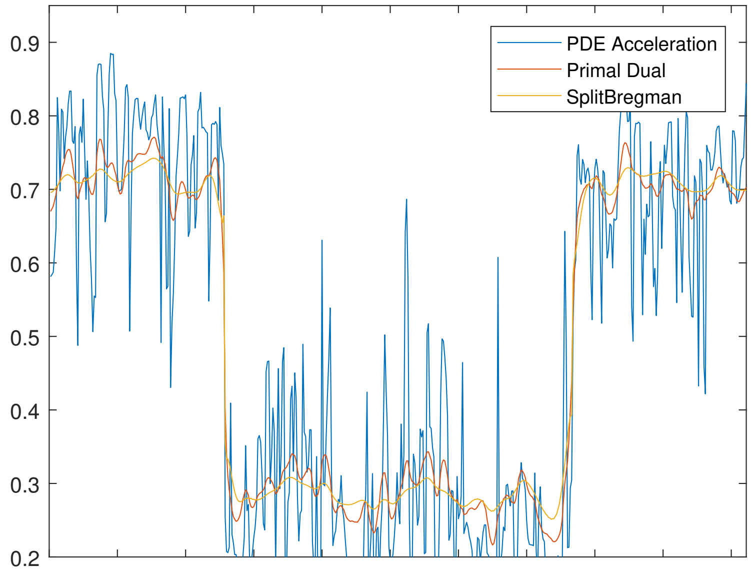
PDE Acceleration
Primal Dual
SplitBregman
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2

$t = 1$

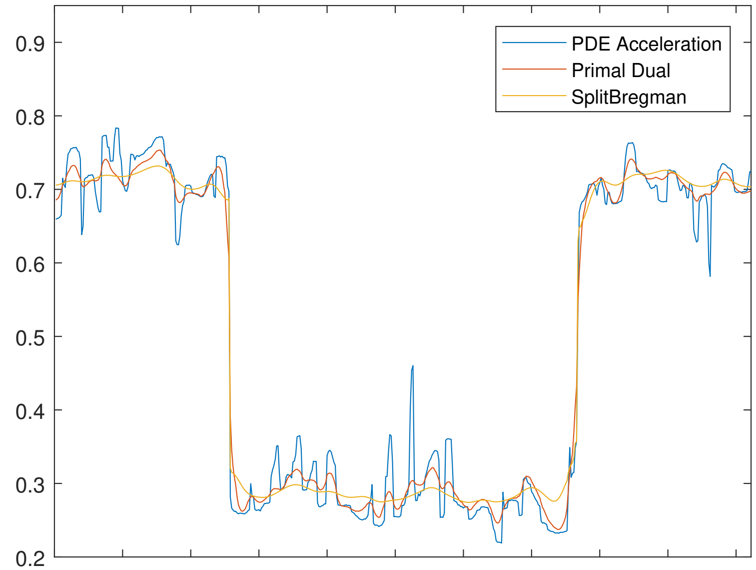
PDE Acceleration
Primal Dual
SplitBregman
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2

$t = 2$

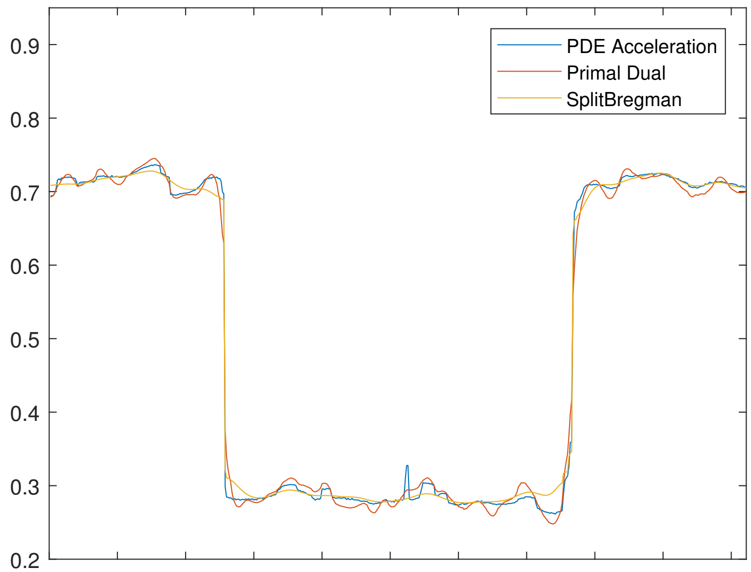
PDE Acceleration
Primal Dual
SplitBregman
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2

$t = 3$

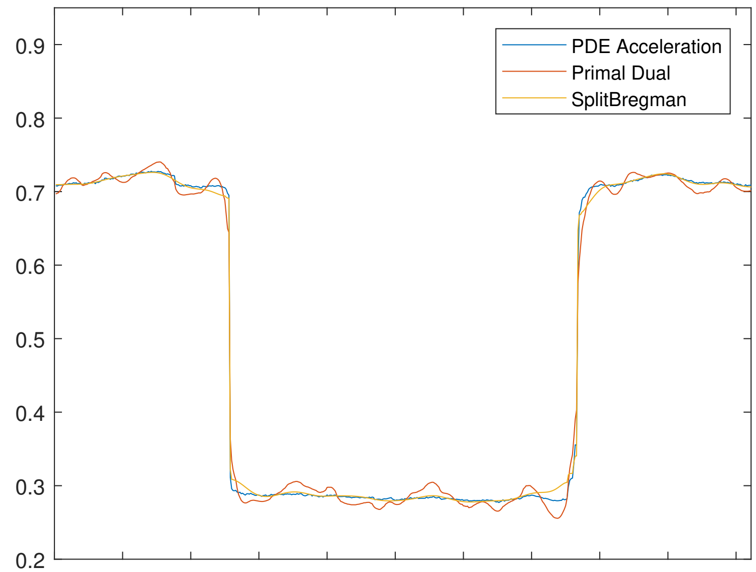
PDE Acceleration
Primal Dual
SplitBregman
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2

$t = 4$

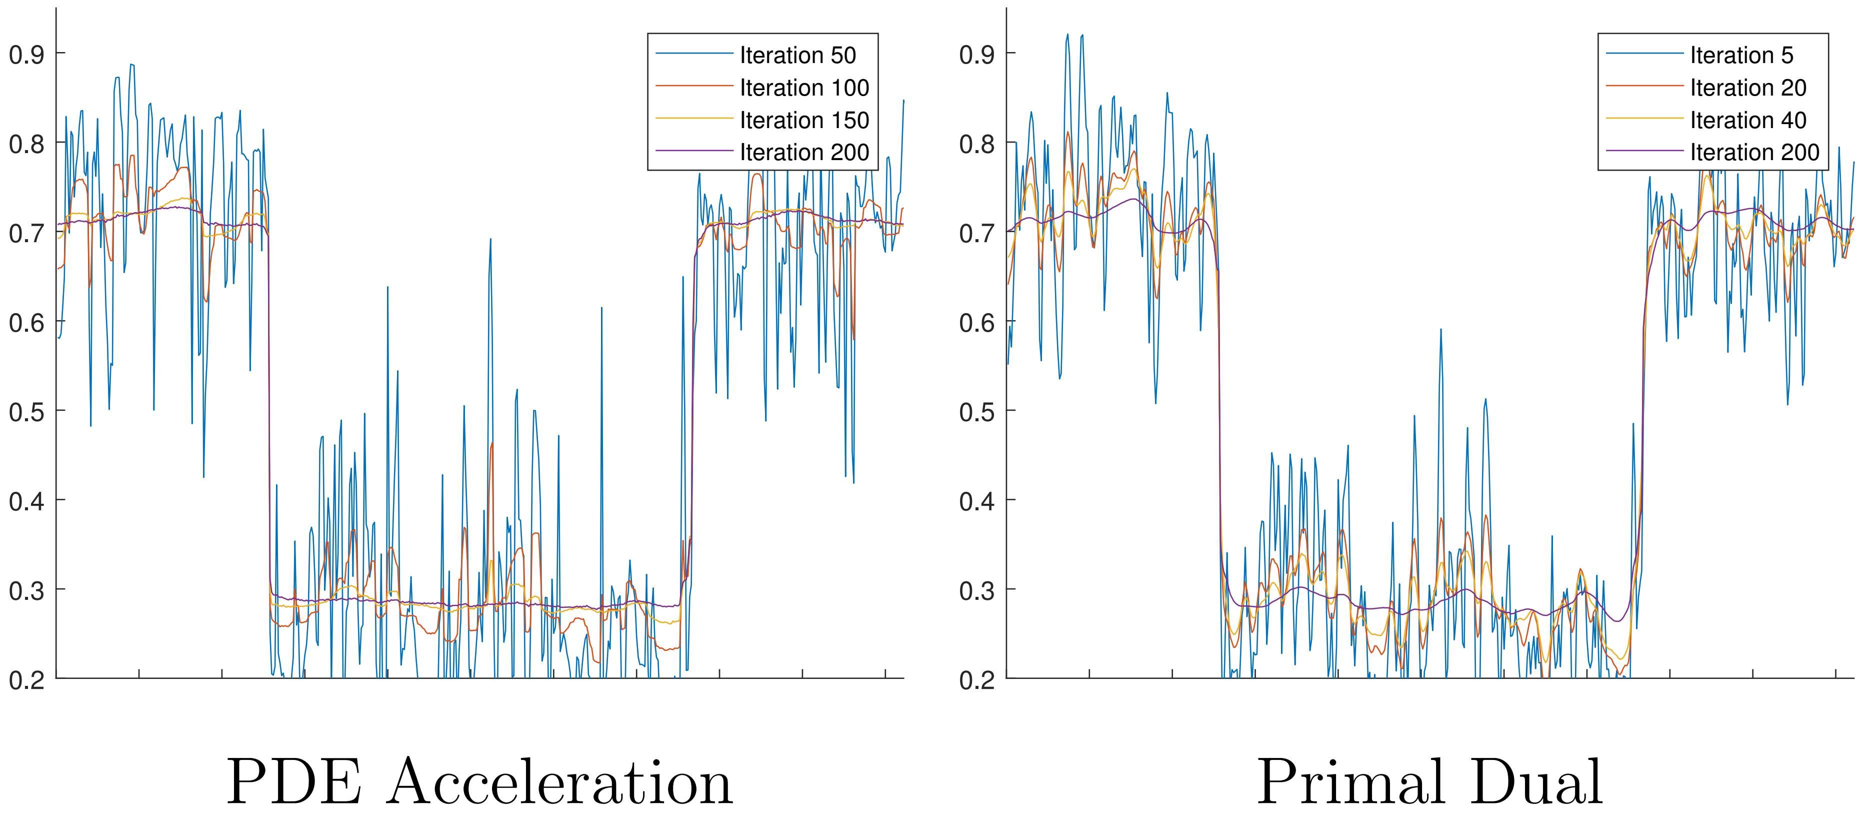

Figure 2: Comparison of flows generated by (a) PDE Acceleration and (b) Primal Dual for solving the TV restoration problem on the noisy square image. Notice the edges are better preserved in PDE acceleration earlier in the flow.

Active Contours

. . . deformable curves which evolve to capture object boundaries within images or other spatially distributed data.

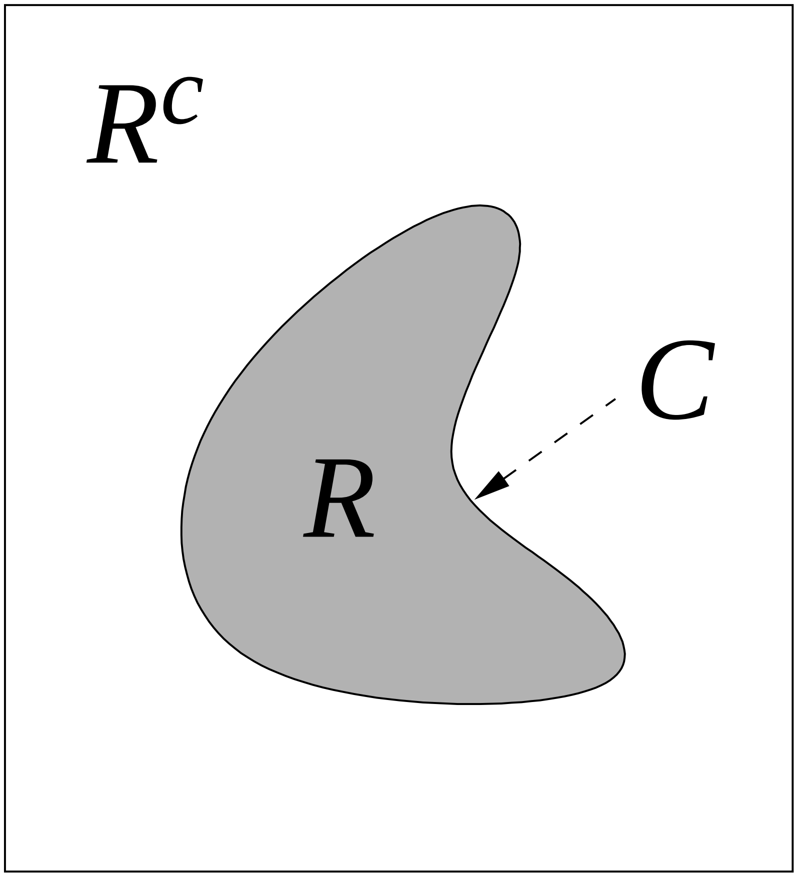

Edge-Based Active Contours evolve in response to image measurements *along* the contour.

Region-Based Active Contours evolve in response to image statistics *inside/outside* the contour.

EXAMPLES

FanBone Cardiac Knee

Level Set Implementations

Level set methods, introduced by Osher and Sethian, greatly facilitate the numerical implementation of active contours.

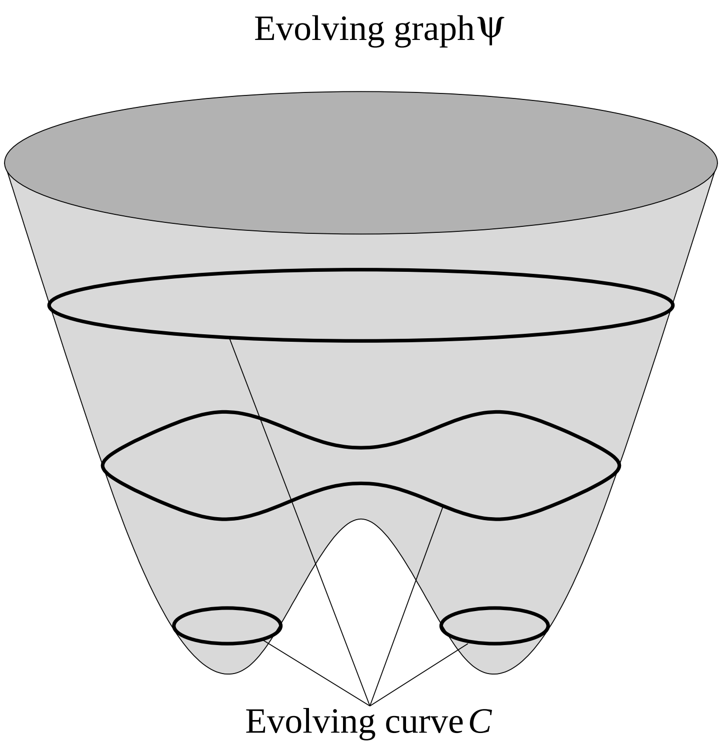

- Embed the curve C as the zero level set of a graph ψ:

$$\psi(C) = 0$$

- Evolve ψ such that its zero level set follows the desired motion for C:

$$\frac{\partial \psi}{\partial t} = -\nabla\psi \cdot \frac{\partial C}{\partial t}$$

Splitting and Merging

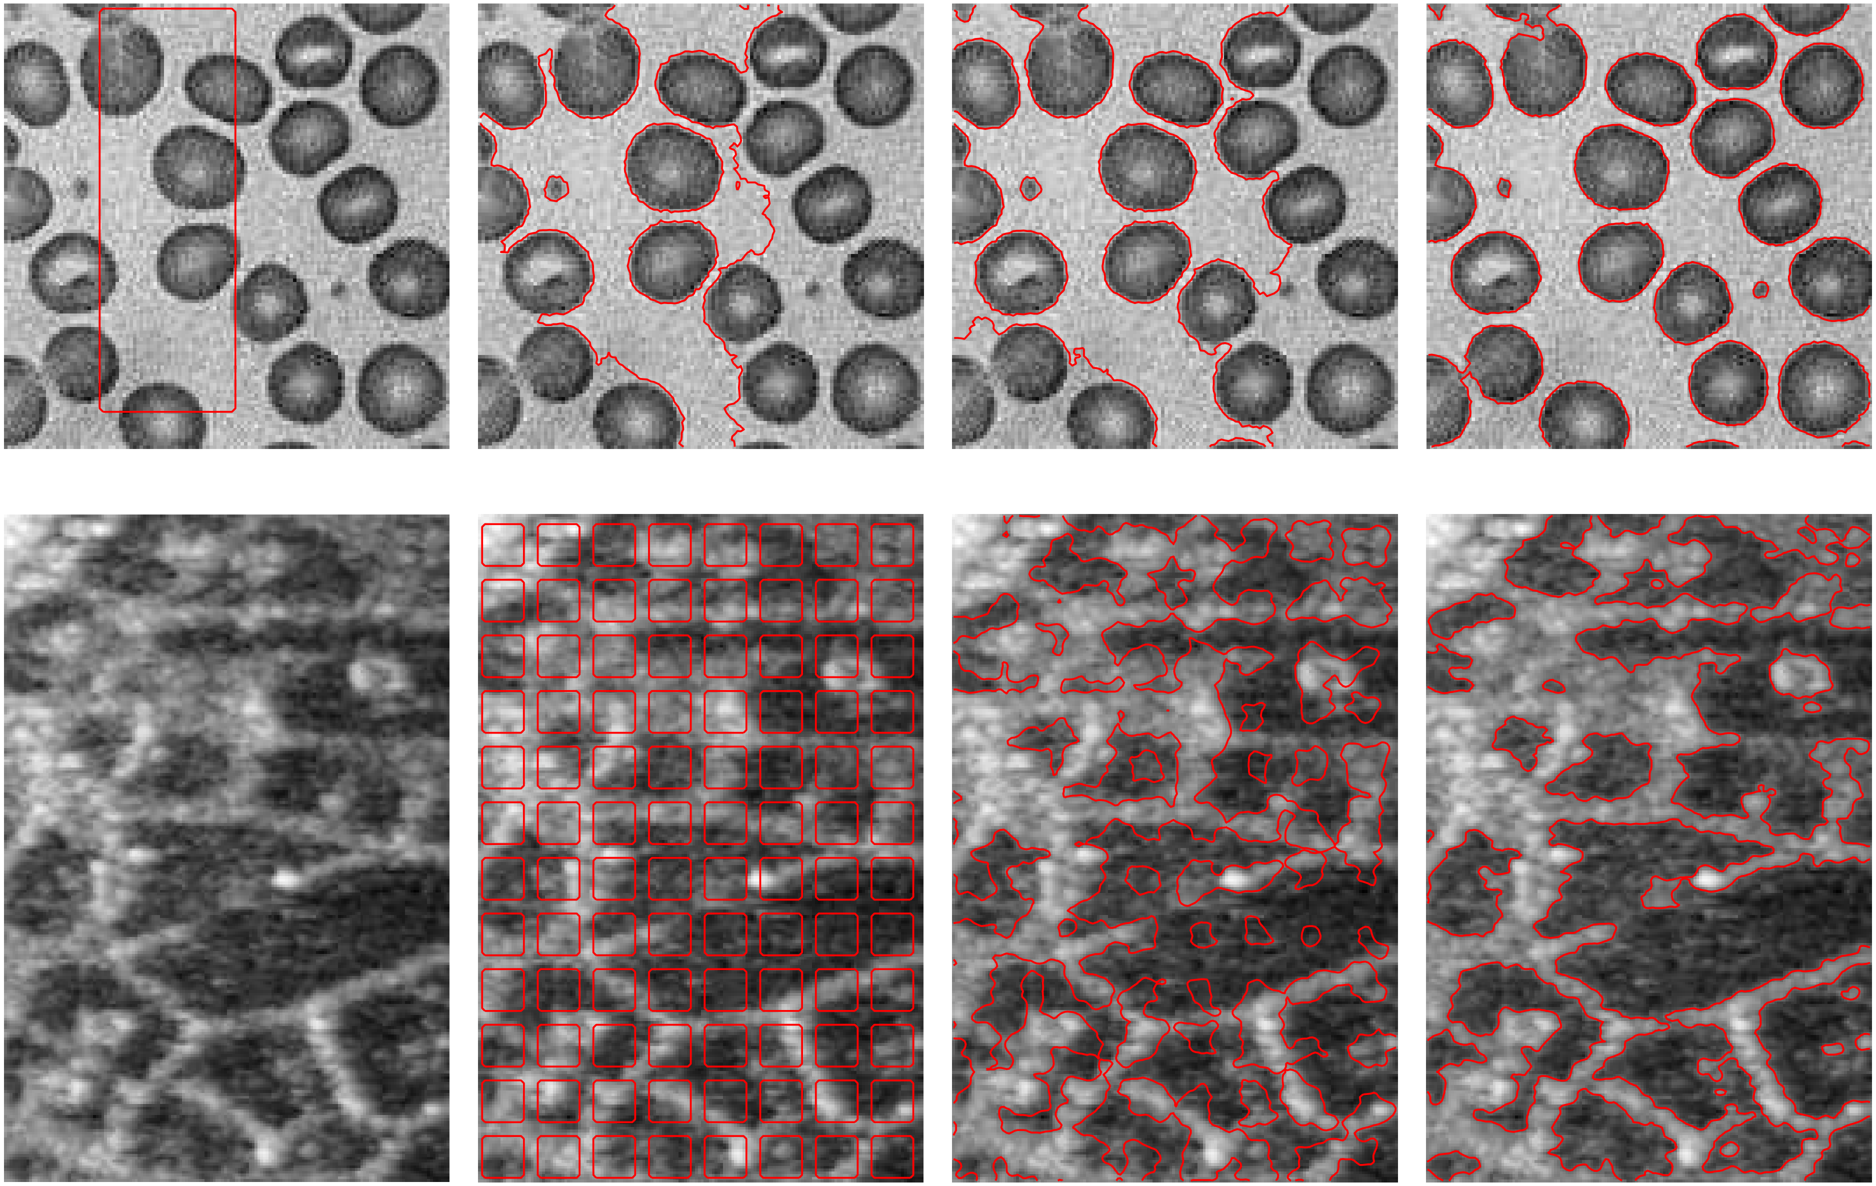

PDE Acceleration for Geometric Active Contours

$$\mathbf{T} \;=\; \frac{1}{2}\rho\int_C \left\|\frac{\partial C}{\partial t}\right\|^2 ds, \qquad \text{where } \rho = \text{constant mass per unit length}$$

$$\mathbf{U} \;=\; E(C), \qquad \text{chosen so } \delta E = \int_C \delta C \cdot \underbrace{fN}_{\text{gradient}} \, ds$$

The Euler-Lagrange equation of the generalized action yields

$$\underbrace{\frac{\partial^2 C}{\partial t^2}}_{\text{acceleration}} = \frac{\lambda k^2 t^{(k-2)}}{\rho} \underbrace{fN}_{\text{gradient}} - \underbrace{\frac{k+1}{t}\frac{\partial C}{\partial t}}_{\text{friction}}$$

$$\underbrace{-\left(\frac{\partial^2 C}{\partial s \partial t} \cdot \frac{\partial C}{\partial s}\right)\frac{\partial C}{\partial t} - \frac{\partial}{\partial s}\left(\frac{1}{2}\left\|\frac{\partial C}{\partial t}\right\|^2 \frac{\partial C}{\partial s}\right)}_{\text{advection/reaction}}$$

Accelerated Contour Evolution Model

Let α and β denote the tangential and normal evolution speeds

$$C_t = \alpha T + \beta N$$

The Euler-Lagrange flow may be written as the coupled system

$$\begin{aligned}
\alpha_t &= -(\alpha^2)_s + 2\alpha\beta\kappa - \frac{3}{t}\alpha \\
\beta_t &= -\left(\alpha\beta\right)_s + \left(\frac{1}{2}\beta^2 - \frac{3}{2}\alpha^2\right)\kappa + \frac{f}{\rho} - \frac{3}{t}\beta
\end{aligned}$$

If we start with $\alpha = 0$ then α remains zero and we obtain

$$\begin{aligned}
\beta_t &= \frac{1}{2}\beta^2\kappa + \frac{f}{\rho} - \frac{3}{t}\beta \\
C_t &= \beta N
\end{aligned}$$

Coupled First Order PDE system

If we start with zero initial velocity we can decompose this nonlinear second-order PDE into a simpler coupled system of nonlinear first order PDE's as follows

$$\frac{\partial C}{\partial t} = \beta N, \qquad \frac{\partial \beta}{\partial t} = \frac{\lambda k^2 t^{(k-2)}}{\rho} f + \frac{1}{2}\beta^2 \kappa - \frac{k+1}{t}\beta$$

<u>Implicit level set implementation</u>:

$$\frac{\partial \psi}{\partial t} = -\hat{\beta}\|\nabla \psi\|, \qquad \frac{\partial \hat{\beta}}{\partial t} = \frac{\lambda k^2 t^{(k-2)}}{\rho} \hat{f} + \nabla \cdot \left(\frac{1}{2}\hat{\beta}^2 \frac{\nabla \psi}{\|\nabla \psi\|} \right) - \frac{k+1}{t}\hat{\beta}$$

where $\hat{f}(x,t)$ and $\hat{\beta}(x,t)$ denote spatial extensions of f and β.

Conserved flowable mass model

If we allow the density ρ to evolve (and flow along the curve):

$$v = \text{internal flow speed}, \qquad \text{total mass velocity} = \frac{\partial C}{\partial t} + v\frac{\partial C}{\partial s}$$

We could therefore consider a more general kinetic energy as

$$\mathbf{T} = \int_C \frac{1}{2}\rho \left\| \frac{\partial C}{\partial t} + v\frac{\partial C}{\partial s} \right\|^2 ds$$

with the variable density evolution governed by the following continuity equation (local mass conservation)

$$\underbrace{\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial s}(\rho v)}_{\text{mass change}} + \rho \underbrace{\left(\frac{\partial^2 C}{\partial s \partial t} \cdot \frac{\partial C}{\partial s} \right)}_{\text{length change}} = 0$$

which can be imposed as a PDE Lagrange multiplier constraint.

The Euler-Lagrange equation of the generalized action yields the following coupled evolution of C, v, and ρ

$$\overbrace{\frac{\partial V}{\partial t}}^{\text{acceleration}} = \frac{\lambda k^2 t^{k-2}}{\rho} \overbrace{f\,N}^{\text{gradient}} - \overbrace{\left(V \cdot \frac{\partial C}{\partial s}\right) \frac{\partial V}{\partial s}}^{\text{advection}} - \overbrace{\frac{k+1}{t} V}^{\text{friction}} \qquad (10)$$

$$\frac{\partial C}{\partial t} = (V \cdot N) N, \qquad \frac{\partial \rho}{\partial t} = -\left(V \cdot \frac{\partial C}{\partial s}\right) \frac{\partial \rho}{\partial s} - \rho \frac{\partial V}{\partial s} \cdot \frac{\partial C}{\partial s}$$

where the velocity field V describes both the tangential flow of the mass as well as the normal flow of the curve itself.

(Note connection to optimal mass transport)

The region based flowable mass for diffeomorphism evolution (joint work led by Sundaramoorthi) is even more strongly connected!

Adding spatial evolution regularity

<u>Heuristic approach</u>: added velocity diffusion term

$$\overbrace{\frac{\partial V}{\partial t}}^{\text{acceleration}} = \frac{\lambda k^2 t^{k-2}}{\rho} \overbrace{f\,N}^{\text{gradient}} - \overbrace{\left(V \cdot \frac{\partial C}{\partial s}\right) \frac{\partial V}{\partial s}}^{\text{advection}} - \overbrace{\frac{k+1}{t} V}^{\text{friction}} + \overbrace{\tau \frac{\partial^2 V}{\partial s^2}}^{\text{diffusion}}$$

More principled approach: mass potential energy

If we imaging that the flowable density function ρ represents the "height" of an incompressible fluid contained within the evolving interface, we may model the potential energy of the mass configuration as

$$\mathbf{U}_{mass} = g\, L \int_C \frac{1}{2}\rho^2 ds$$

where g represents a tunable gravitational constant and L represents the total length of the evolving curve (this scaling makes the minimum potential energy value independent of the curve).

This will favor (but not constrain) deformations that keep the density constant (translations, rescalilngs, rotations, etc).

Incorporating stochastic acceleration terms

Finally, the accelerated framework for active contours, surfaces, and other PDE evolution schemes, offers a numerical opportunity to introduce random noise into the evolution process without destroying the continuity of the evolution or the evolving object.

$$\overbrace{\frac{\partial V}{\partial t}}^{\text{acceleration}} = \frac{\lambda k^2 t^{k-2}}{\rho} \overbrace{f\,N}^{\text{gradient}} - \overbrace{\left(V \cdot \frac{\partial C}{\partial s}\right) \frac{\partial V}{\partial s}}^{\text{advection}} - \overbrace{\frac{k+1}{t}}^{\text{friction}} V + \overbrace{\tau \mathcal{W}}^{\text{noise}}, \qquad \frac{\partial C}{\partial t} = (V \cdot N) N$$

Since the noise is added to the acceleration, it gets twice integrated in the construction of the evolving geometry and therefore does not immediately interfere with its continuity nor its first order differentiability (maintains a continous unit normal).

Preliminary results

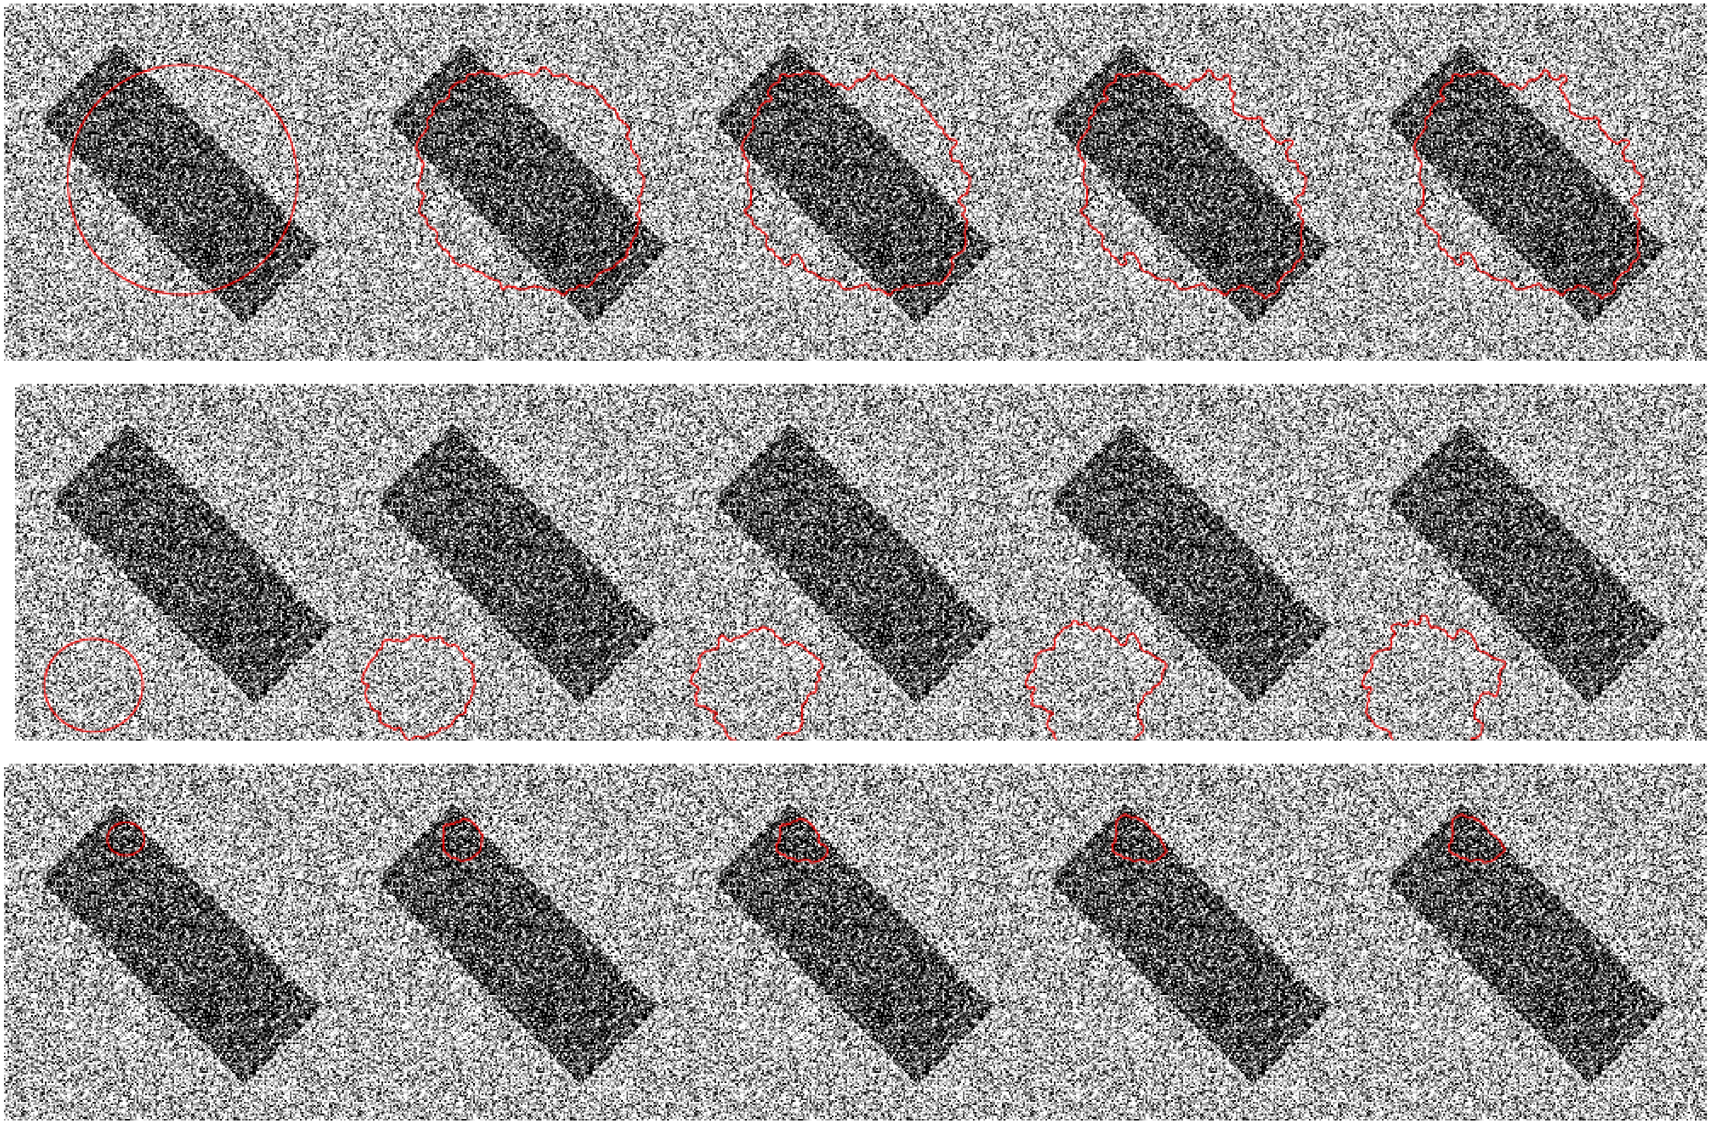

Three active contours getting stuck in different local minima

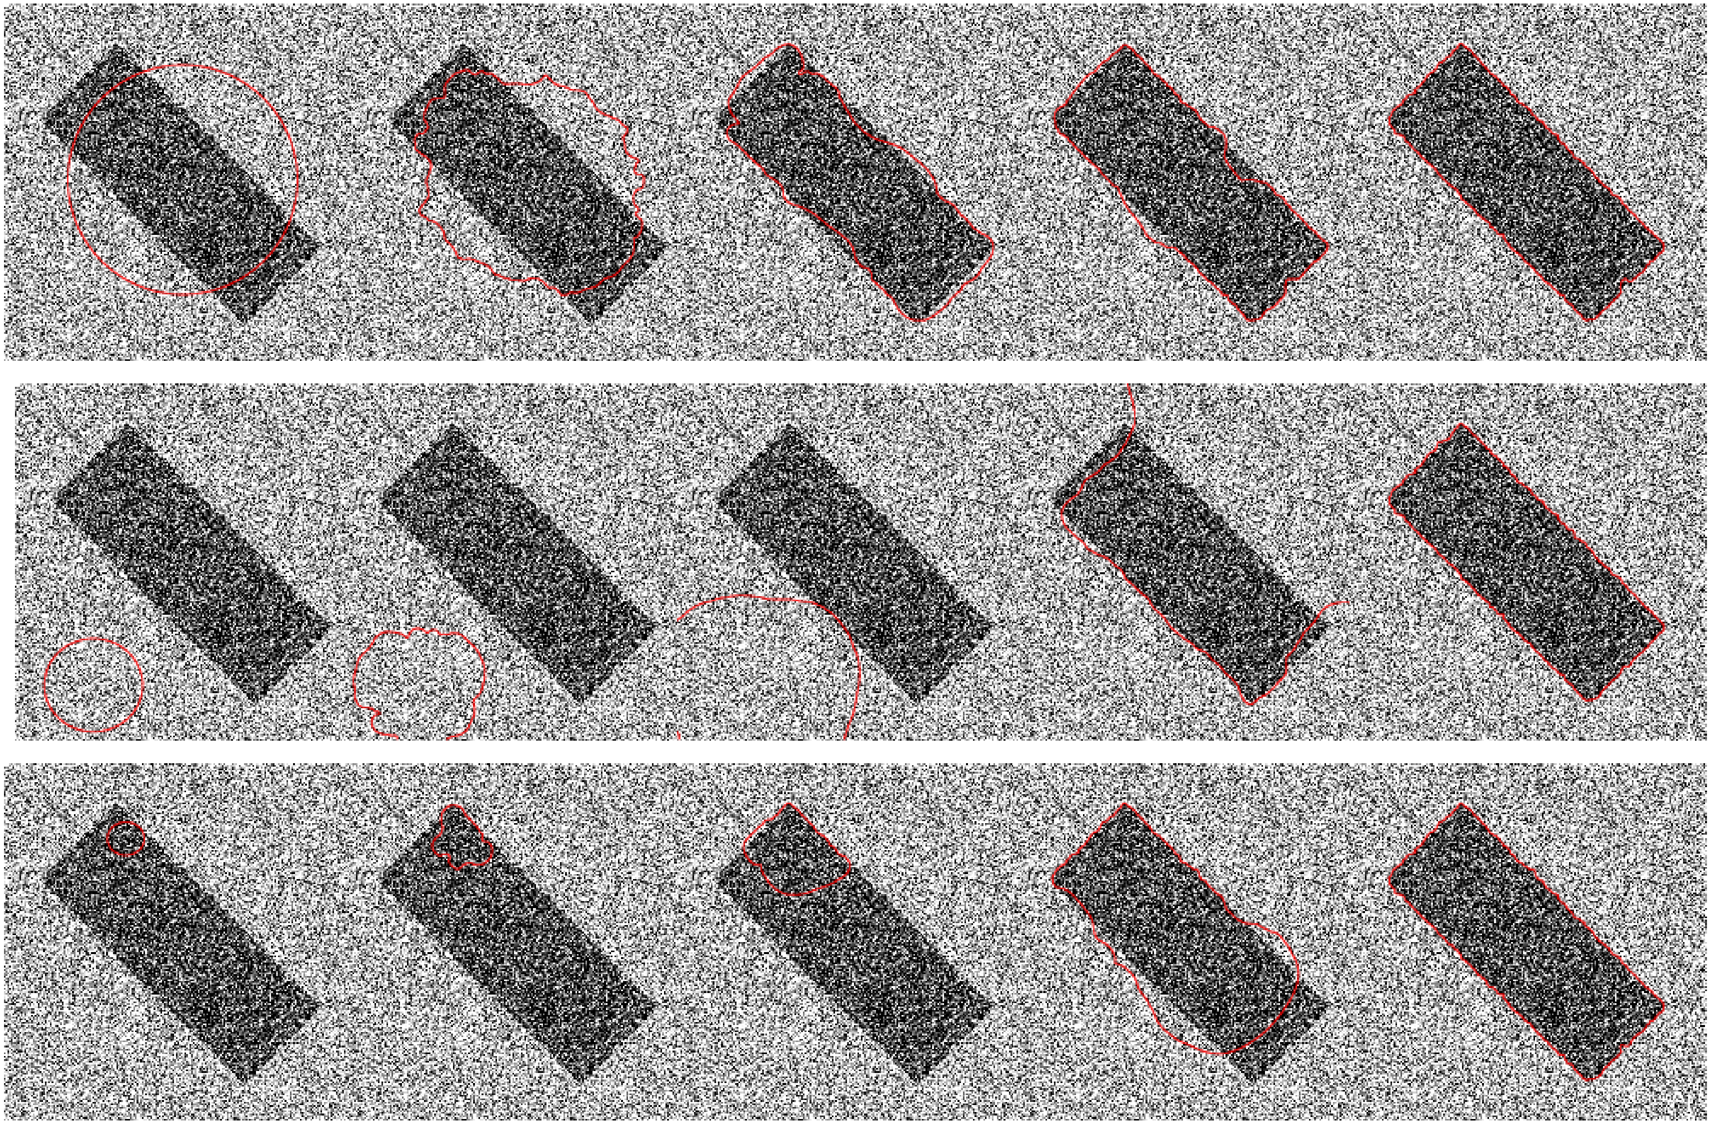

Accelerated active contours all converging to same minimzer

Demos: Let's play a bit!

Accelerated Active Contour Demos

Accelerated Active Surface Demo