Lecture 12:

Variational Methods III: TV Denoising, Equivalence Results

Contents

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Motivation

- focus on one of the most important variational methods:
  total variation (TV) regularisation
- present analytical solutions for the space-discrete 1-D case
- establish 1-D equivalence results to other methods:
  - TV diffusion
  - soft Haar wavelet shrinkage

Total Variation (TV) Regularisation

(Rudin et al. 1992, Acar / Vogel 1994)

- Variational method with penaliser $\Psi(|\nabla u|^2) := 2|\nabla u|$.
- in 1-D: filtered version $u(x)$ of noisy signal $f(x)$ as minimiser of

$$E_f(u) \; := \; \int_a^b \left((f-u)^2 + 2\alpha\, |u_x|\right) dx$$

- nice in theory:
  - piecewise constant filtered signal in discrete 1-D case (ideal segmentation, Nikolova 2000)
  - no additional parameters (besides α)

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

- difficult in practice:
  - regulariser $2|u_x|$ not differentiable in 0
  - Euler-Lagrange equation given by

$$\frac{u-f}{\alpha} \; = \; \partial_x \left(\frac{1}{|u_x|}\, u_x \right)$$

  - Its unbounded diffusivity $\frac{1}{|u_x|}$ creates numerical problems.
  - usually approximated by a model with bounded diffusivity:

$$\frac{u-f}{\alpha} \; = \; \partial_x \left(\frac{1}{\sqrt{\varepsilon^2 + u_x^2}}\, u_x \right)$$

  - This approximation may not reproduce some of the nice theoretical properties.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

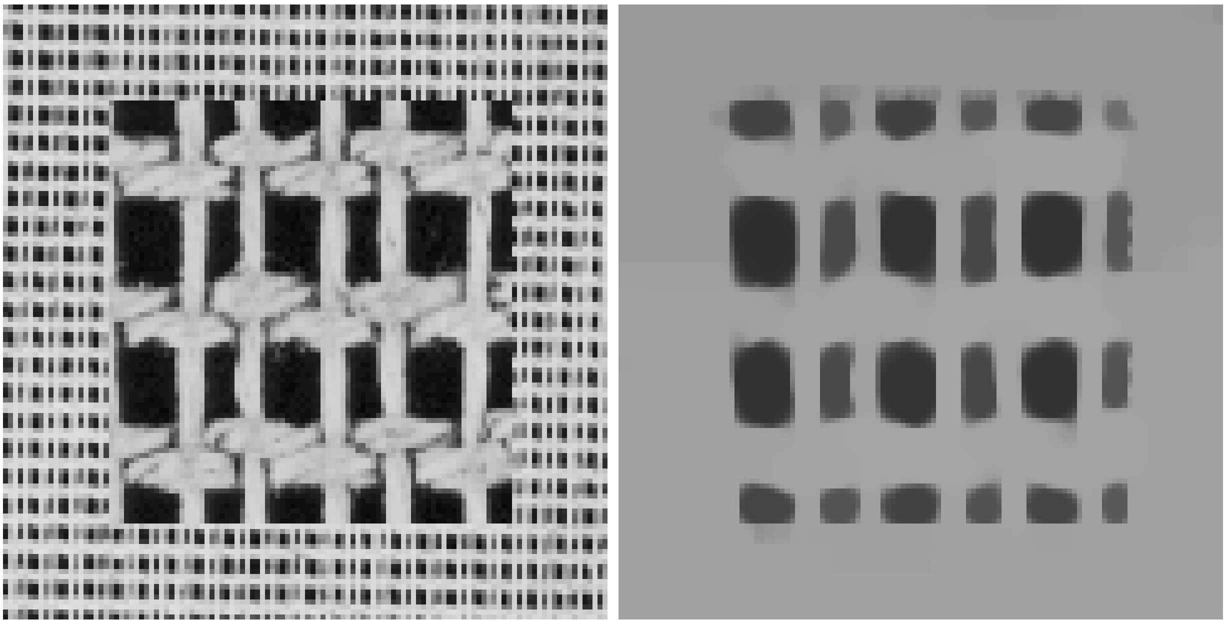

Left: Original image, 128×128 pixels. **Right:** TV regularisation with $\alpha = 100$ using a fixed point scheme with Gauß-Seidel iterations and a regularised diffusivity $\frac{1}{\sqrt{\varepsilon^2+u_x^2}}$ with $\varepsilon = 0.001$. **Author:** T. Brox (2003).

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Total Variation (TV) Diffusion

(Andreu et al. 2001)

- diffusion filter corresponding to TV regularisation: uses diffusivity $g(|\nabla u|^2) = \frac{1}{|\nabla u|}$.
- in 1-D: simplify f under the diffusion equation

$$\partial_t u \;=\; \partial_x \left(\frac{1}{|u_x|} \, u_x \right)$$

  with reflecting boundary conditions
- nice in theory:
  - adaptive diffusivity $1/|u_x|$ becomes 0 at ideal edges and ∞ in ideal flat regions
  - no parameters besides diffusion time t
  - flattens every signal in finite (!) time (Andreu et al. 2001)
  - preserves shape of some objects (Bellettini et al. 2002)

TV Diffusion (2)

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

- difficult in practice:
  - unbounded diffusivity requires infinitesimally small time steps (explicit case) or creates highly ill-conditioned linear systems of equations (semi-implicit case)
  - usually approximated by a model with bounded diffusivity:

$$\partial_t u \;=\; \partial_x \left(\frac{1}{\sqrt{\varepsilon^2 + u_x^2}} \, u_x \right)$$

  - This modified model does not reproduce the nice theoretical properties exactly.

TV Diffusion (3)

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

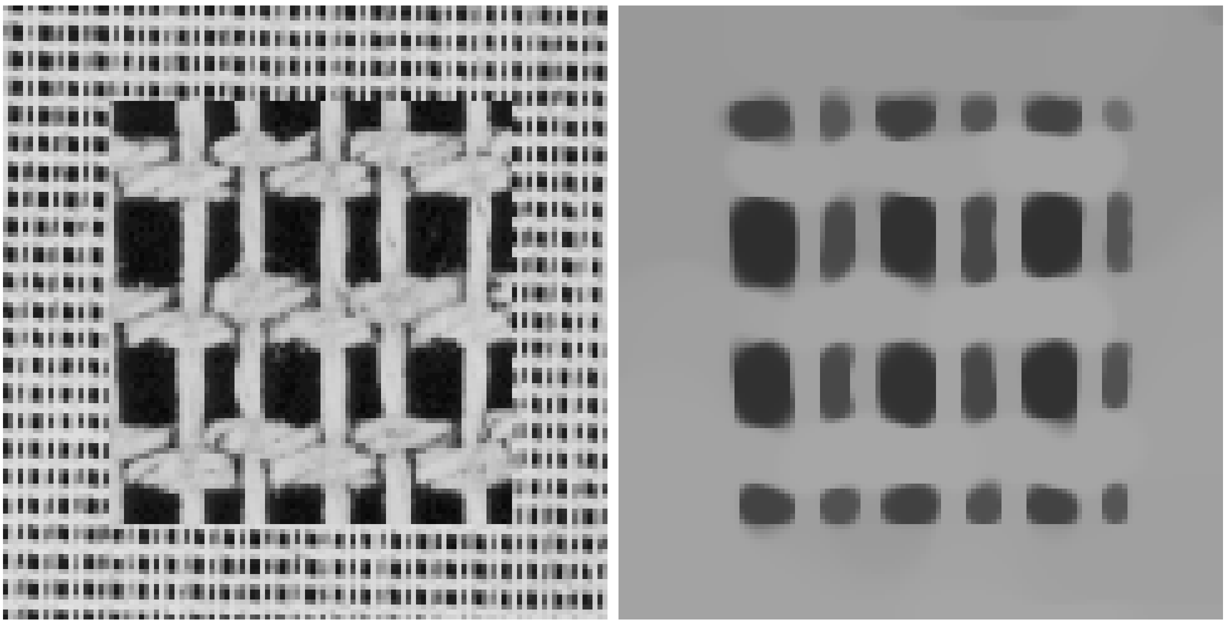

Left: Original image, 128×128 pixels. **Right:** TV diffusion with stopping time $t = 100$ using an explicit scheme with $\tau = 0.00025$ and a regularised diffusivity $\frac{1}{\sqrt{\varepsilon^2+u_x^2}}$ with $\varepsilon = 0.001$. **Author:** T. Brox (2003).

Soft Wavelet Shrinkage

(Donoho 1995)

- Given: discrete 1-D signal $\mathbf{f} = (f_i)_{i\in\mathbb{Z}}$ (piecewise constant function)
- Represent $\mathbf{f}$ by shifted and translated versions of a scaling function φ and a wavelet function ψ:

$$\mathbf{f} = \sum_{i\in\mathbb{Z}} c_i^n \varphi_i^n + \sum_{j=-\infty}^{n} \sum_{i\in\mathbb{Z}} d_i^j \psi_i^j,$$

where $\psi_i^j(s) := 2^{-j/2}\,\psi(2^{-j}s - i)$, j denotes scale, and i shift.

- Shrink the wavelet coefficients d_i^j towards 0:

$$S_\theta(d_i^j) := \begin{cases} d_i^j - \theta \operatorname{sgn} d_i^j & \text{if} \quad |d_i^j| > \theta, \\ 0 & \text{if} \quad |d_i^j| \le \theta. \end{cases}$$

- Reconstruct the filtered signal $u(\theta)$ from the modified coefficients:

$$\mathbf{u}(\theta) = \sum_{i\in\mathbb{Z}} c_i^n \varphi_i^n + \sum_{j=-\infty}^{n} \sum_{i\in\mathbb{Z}} S_\theta(d_i^j) \psi_i^j.$$

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

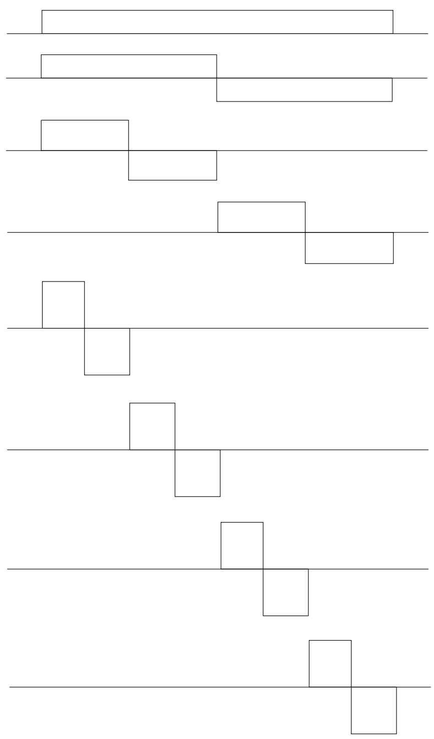

Scaling function and Haar wavelet basis for signals with $2^3 = 8$ pixels. **From top to bottom:** Scaling function $\phi_{3,0}$, wavelets $\psi_{3,0}$, $\psi_{2,0}$, $\psi_{2,1}$, $\psi_{1,0}$, $\psi_{1,1}$, $\psi_{1,2}$, $\psi_{1,3}$. These 8 orthonormal vectors allow to represent any discrete signals of length 8. **Author:** S. Zimmer (2002).

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Soft Wavelet Shrinkage (3)

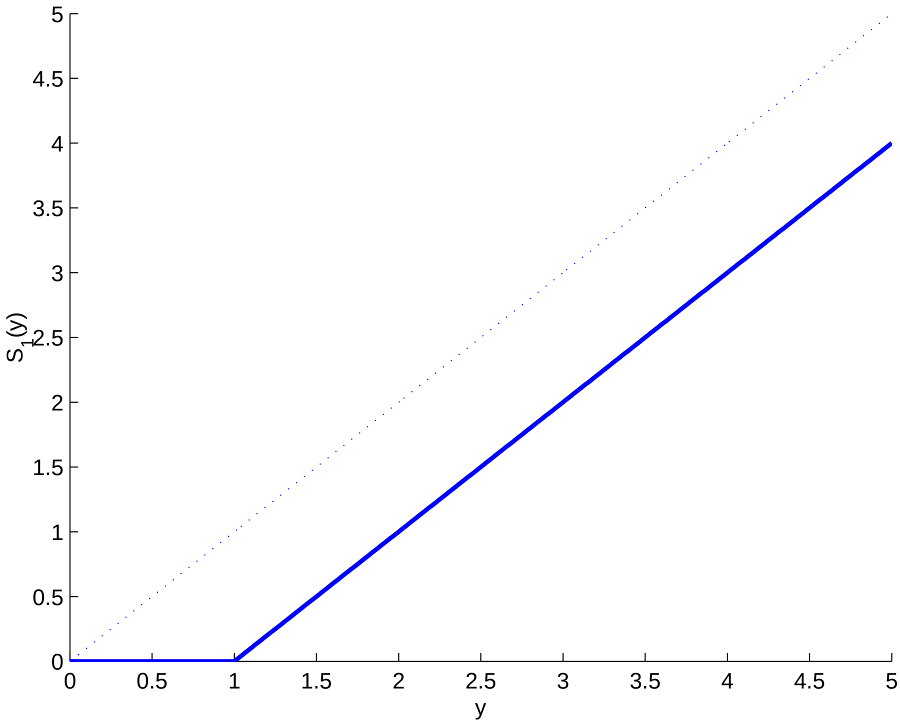

Shrinkage function for soft wavelet shrinkage with $\theta = 1$. **Author:** P. Mrázek (2003).

Soft Wavelet Shrinkage (4)

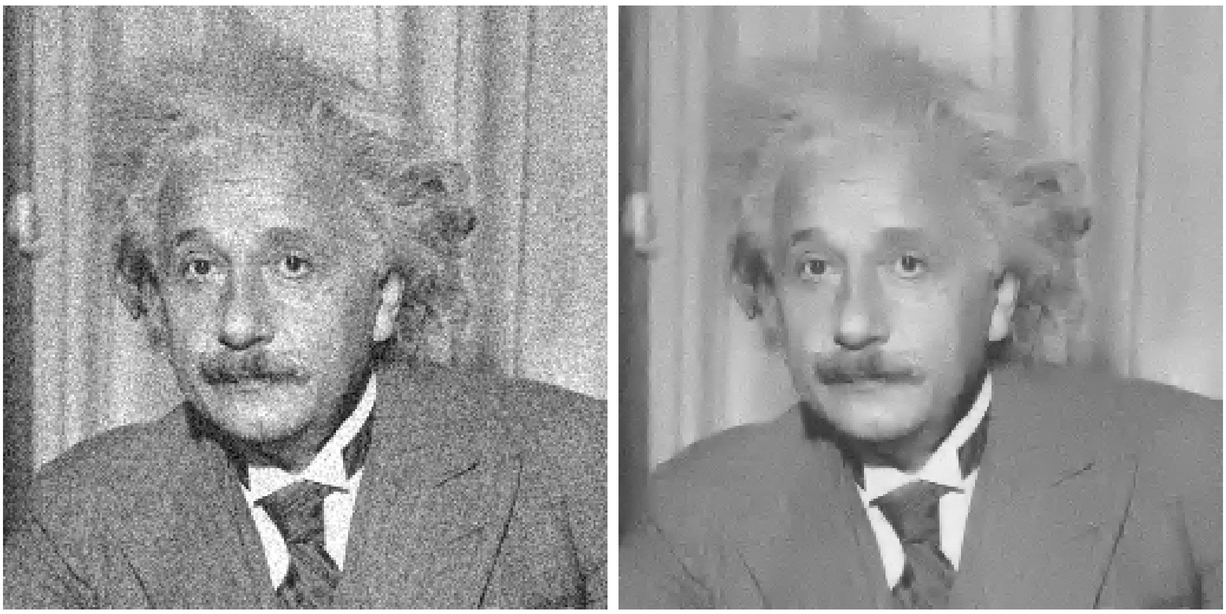

Image denoising with wavelet shrinkage. **Left:** Image of Einstein (256×256 pixels), degraded by Gaussian noise with $\sigma = 17.4$. **Right:** Result after using shift invariant soft wavelet shrinkage with Haar wavelets and threshold $\theta = 30$.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Equivalence Results for 2-Pixel-Signals

Haar Wavelet Shrinkage for 2-Pixel-Signals

- **analysis step of signal (f_0, f_1) in the Haar basis:**

$$\text{coefficient for } \phi = (\tfrac{1}{\sqrt{2}}, \tfrac{1}{\sqrt{2}}): \quad c = f^\top \phi = \frac{f_0 + f_1}{\sqrt{2}}$$

$$\text{coefficient for } \psi = (\tfrac{1}{\sqrt{2}}, \tfrac{-1}{\sqrt{2}}): \quad d = f^\top \psi = \frac{f_0 - f_1}{\sqrt{2}}$$

- **soft thresholding of the wavelet coefficient:**

$$S_\theta(d) = \begin{cases} d - \theta \operatorname{sgn} d & \text{if} \quad |d| > \theta, \\ 0 & \text{if} \quad |d| \le \theta. \end{cases}$$

- **synthesis step:**
  coefficients c and $S_\theta(d)$ yield filtered signal $\mathbf{u}(\theta) = c\phi + S_\theta(d)\psi$ with

$$u_0(\theta) = \begin{cases} f_0 + \frac{\theta}{\sqrt{2}} \operatorname{sgn}(f_1 - f_0) & \text{if} \quad \theta < |f_1 - f_0|/\sqrt{2}, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

$$u_1(\theta) = \begin{cases} f_1 - \frac{\theta}{\sqrt{2}} \operatorname{sgn}(f_1 - f_0) & \text{if} \quad \theta < |f_1 - f_0|/\sqrt{2}, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

Equivalence Results for 2-Pixel-Signals (3)

TV Diffusion for 2-Pixel-Signals

- space discretisation with grid size 1 and reflecting boundary conditions gives:

$$\begin{aligned} \dot u_0 &= \operatorname{sgn}(u_1 - u_0) \\ \dot u_1 &= -\operatorname{sgn}(u_1 - u_0) \end{aligned}$$

  with initial conditions $u_0(0) = f_0$ and $u_1(0) = f_1$.

- has the analytical solution

$$u_0(t) = \begin{cases} f_0 + t \operatorname{sgn}(f_1 - f_0) & \text{if } \quad t < |f_1 - f_0|/2, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

$$u_1(t) = \begin{cases} f_1 - t \operatorname{sgn}(f_1 - f_0) & \text{if } \quad t < |f_1 - f_0|/2, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

- equivalent to soft Haar wavelet shrinkage with threshold $\theta = \sqrt{2}t$.
- finite extinction time is obvious in the two-pixel model
- no numerical problems with unbounded diffusivities

Equivalence Results for 2-Pixel-Signals (4)

TV Regularisation for 2-Pixel-Signals

- filtered signal (u_0, u_1) of (f_0, f_1) minimises

$$E_f(u_0, u_1) = (f_0 - u_0)^2 + (f_1 - u_1)^2 + 2\alpha\, |u_1 - u_0|$$

- minimiser given by (forthcoming assignment)

$$u_0(\alpha) = \begin{cases} f_0 + \alpha \operatorname{sgn}(f_1 - f_0) & \text{if } \quad \alpha < |f_1 - f_0|/2, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

$$u_1(\alpha) = \begin{cases} f_1 - \alpha \operatorname{sgn}(f_1 - f_0) & \text{if } \quad \alpha < |f_1 - f_0|/2, \\ (f_0 + f_1)/2 & \text{else.} \end{cases}$$

- identical to TV diffusion with stopping time $t = \alpha$
- equivalent to soft Haar wavelet shrinkage with threshold $\theta = \sqrt{2}\alpha$

MIA
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Equivalence Results for N-Pixel-Signals

Haar Wavelet Shrinkage in the Shift Invariant N-Pixel Case

- We perform a wavelet decomposition on the finest scale only.
- Haar wavelets create natural two-pixel pairings.
- However, the wavelet shrinkage is not shift invariant.
- Remedy to create translation invariance:
  cycle spinning (Coifman / Donoho 1995)
  - perform shrinkage on the original signal
  - shift signal by 1 pixel, perform shrinkage, shift back
  - average both filtered signals

M I A

A Wavelet-Inspired Numerical Algorithm for TV Diffusion

- use analytical solution of the two-pixel model for constructing a numerical scheme for TV diffusion
- construct a numerical scheme that is equivalent to shift invariant Haar wavelet shrinkage on a single scale:
  - perform TV diffusion on all even–odd pixel pairs (u_{2i}, u_{2i+1})
  - perform TV diffusion on all odd–even pixel pairs (u_{2i-1}, u_{2i})
  - average both results
- leads after some calculations to the following scheme:

$$\begin{aligned}\frac{u_i^{k+1} - u_i^k}{\tau} &= \frac{1}{h}\,\mathrm{sgn}\,(u_{i+1}^k - u_i^k)\,\min\left(1, \frac{h}{4\tau}|u_{i+1}^k - u_i^k|\right)\\ &- \frac{1}{h}\,\mathrm{sgn}\,(u_i^k - u_{i-1}^k)\,\min\left(1, \frac{h}{4\tau}|u_i^k - u_{i-1}^k|\right).\end{aligned}$$

Equivalence Results for N-Pixel-Signals (3)

- stabilisation of the naive scheme for the TV diffusion $u_t = \partial_x(\mathrm{sgn}(u_x))$:

$$u_i^{k+1} = u_i^k + \frac{\tau}{h}\,\mathrm{sgn}\left(u_{i+1}^k - u_i^k\right) - \frac{\tau}{h}\,\mathrm{sgn}\left(u_i^k - u_{i-1}^k\right)$$

- explicit, but *absolutely stable*: satisfies extremum principle

$$\min_j f_j \le u_i^{k+1} \le \max_j f_j$$

- $O(\tau + h^2)$ approximation to the continuous TV diffusion for

$$\tau \le \frac{h}{4} \min\left(|u_{i+1}^k - u_i^k|,\, |u_i^k - u_{i-1}^k|\right).$$

  for large τ: approximates a linear diffusion process

- does not require to consider regularised TV diffusion

$$\partial_t u = \partial_x \left(\frac{1}{\sqrt{\varepsilon^2 + u_x^2}}\, u_x \right)$$

  and performs competitively

Equivalence Results for N-Pixel-Signals (4)

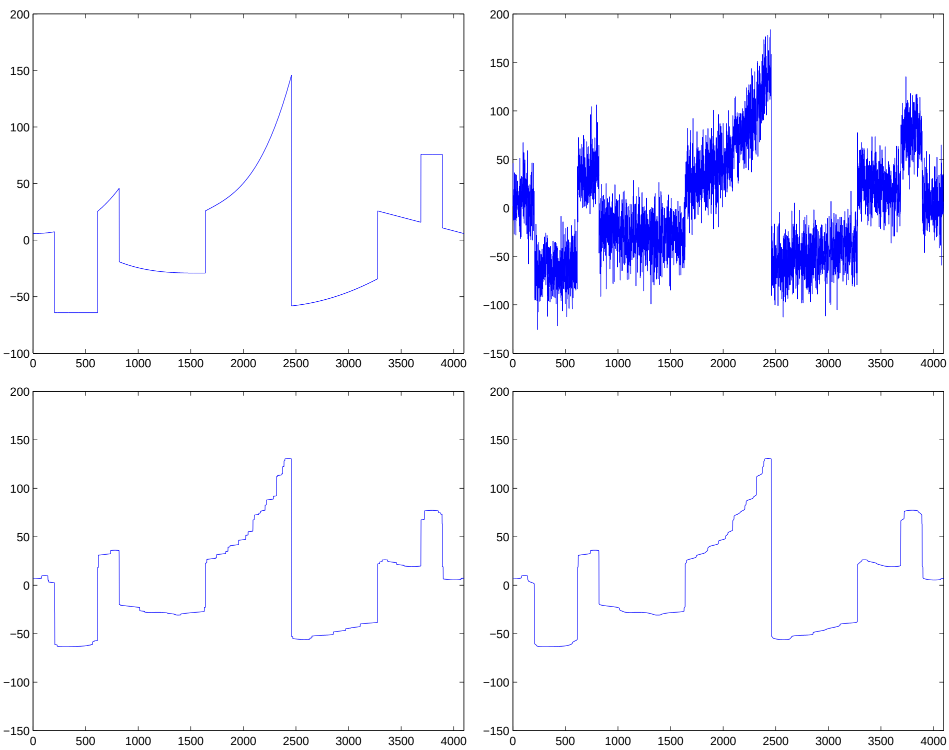

Top left: original signal without noise. **Top right:** with additive Gaussian noise. **Bottom left:** result with two-pixel scheme for TV diffusion. **Bottom right:** result with classical regularised scheme for TV diffusion.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Equivalence Results Between TV Diffusion and TV Regularisation

- Space-discrete TV diffusion and TV regularisation are also equivalent in the N pixel case!
- An analytical solution exists with simple properties:
  - only extrema move
  - movement in the direction of contrast reduction
  - pixels merge in finite time and remain merged
  - each segment moves linearly in time with a speed that is inversely proportional to its number of pixels

Consequences:

- In 1-D, TV diffusion and TV regularisation lead to simple region merging!
- Advantages over classical region merging:
  - merging order is not ad hoc
  - optimality interpretations

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

The Two-Dimensional Case

TV Diffusion in 2-D

- Consider
$$\partial_t u = \operatorname{div}\left(g(|\nabla u|^2)\,\nabla u\right)$$
with the TV diffusivity
$$g(|\nabla u|^2) \;=\; \frac{1}{|\nabla u|}$$
- Unfortunately, TV diffusion and TV regularisation are no longer equivalent in 2-D.
- However, one can generalise the 2-pixel schemes in 1-D to 4-pixel schemes in 2-D (Welk et al. 2005).

The Two-Dimensional Case (2)

Analytical Solution for Four-Pixel Images

- consider 2×2 pixel image $(f_{i,j})$ with $i, j \in \{0, 1\}$
- specific space discretisation leads to the dynamical system

$$\begin{aligned} \frac{du_{i,j}}{dt} &= 4\, g(n^2(u))\, (\mu - u_{i,j}), \qquad i, j \in \{0, 1\}, \\ u_{i,j}(0) &= f_{i,j} \end{aligned}$$

μ:	average grey value
$n(u)$:	approximates $\lvert\nabla u\rvert$ in the midpoint
$g(n^2(u))$:	approximates the diffusivity $g(\lvert\nabla u\rvert^2)$

- has analytical solution that evolves towards μ:

$$u_{i,j}(t) = \begin{cases} \mu + \left(1 - \frac{4t}{n(f)}\right)\,(f_{i,j} - \mu), & 0 \le t < \frac{n(f)}{4}, \\ \mu & t \ge \frac{n(f)}{4}. \end{cases}$$

The Two-Dimensional Case (3)

A Locally Analytic Scheme (LAS) for Images of Arbitrary Size

- for each time step:
  - consider the four 2×2 cells containing some pixel (i, j)
  - compute their analytical solutions
  - average the four results
- does not use any regularisation
- absolutely stable: satisfies extremum principle
- conditionally consistent
- good sharpness at edges, even for large time steps
- very simple algorithm that can even be used on mobile phones, PDAs, ...

Comparison of schemes for TV diffusion. **Left:** Original image, 93×93 pixels. **Middle:** Standard explicit scheme for regularised TV diffusion ($\varepsilon = 0.01$, $\tau = 0.0025$, 40000 iterations). **Right:** Same with LAS scheme without regularisation ($\tau = 0.1$, 1000 iterations). Note that 40 times larger time steps are used.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Summary

- Considered two-pixel ideas to analyse relations between 1-D wavelet shrinkage, TV diffusion and TV regularisation.
- Equivalence results when only 2-pixel pairs are considered.
- Correct result on finite extinction time.
- In the shift invariant N-pixel case, wavelet shrinkage gives a novel algorithm for TV diffusion.
- This algorithm is unconditionally stable and does not require to regularise TV diffusion.
- Discrete TV regularisation and TV diffusion are also equivalent in the N-pixel case.
- Analytical solution shows that classical region merging results.
  This gives a novel theoretical justification of region merging.
- TV diffusion and TV regularisation are not equivalent in 2-D.
- 2-pixel schemes for 1-D TV diffusion can be generalised to 4-pixel schemes for 2-D TV diffusion.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

References

- L. I. Rudin, S. Osher, E. Fatemi: Nonlinear total variation based noise removal algorithms. *Physica D*, Vol. 60, 259–268, 1992.
  (introduced TV regularisation)
- F. Andreu, V. Caselles, J. I. Diaz, J. M. Mazón: Qualitative properties of the total variation flow. *Journal of Functional Analysis*, Vol. 188, No. 2, 516–547, Feb. 2002.
  (establishes important properties of TV diffusion such as finite extinction time)
- D. L. Donoho: De-noising by soft thresholding. *IEEE Transactions on Information Theory*, Vol. 41, 613–627, 1995.
  (investigates soft wavelet shrinkage)
- G. Steidl, J. Weickert, T. Brox, P. Mrázek, M. Welk: *On the equivalence of soft wavelet shrinkage, total variation diffusion, total variation regularization, and SIDEs*. SIAM Journal on Numerical Analysis, Vol. 42, No. 2, 686-713, 2004.
  An extended version appeared as Technical Report No. 94, Department of Mathematics, Saarland University, August 2003.
  (`http://www.mia.uni-saarland.de/publications.shtml`)
  (presents proofs for the equivalence results)
- M. Welk, J. Weickert, G. Steidl: A four-pixel scheme for singular differential equations. In R. Kimmel, N. Sochen, J. Weickert (Eds.): *Scale Space and PDE Methods in Computer Vision.* Springer LNCS 3459, pp. 610–621, 2005.
  (`http://www.mia.uni-saarland.de/publications.shtml`)
  (analysis of the 2-D case)

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Assignment P3 – Programming

Download the file `Ex03.tar` from

`http://www.mia.uni-saarland.de/Teaching/dic08.shtml`

and unpack it by `tar xvf Ex03.tar`.

Problem 1 (Anisotropic Diffusion) (6 points)

The object code `ced.o` has all routines on board in order to become an explicit scheme for coherence-enhancing anisotropic diffusion, except for routines for calculating the diffusion tensor.

(a) Supplement the file `diff_tensor.c` with the missing code. You may use the included routines for principal axis transformation and backtransformation. Compile the programme with

```
gcc -O2 -o ced ced.o diff_tensor.c -lm
```

(b) Use the programme ced for enhancing the fingerprint image `finger.pgm` with the parameters $C = 1$, $\sigma = 0.5$, $\rho = 4$, $\alpha = 0.001$, $\tau = 0.2$, 40 iterations. You will observe that the extremum principle is violated by the standard discretisation that is used in the algorithm.

(c) Use ced for creating your own Christmas postcards. It's easy: just take `xmas.pgm` and filter it with the same parameters as for the fingerprint.

(d) You can use ced to visualise *all* stripes of `fabric.pgm` at different scales. Use the standard parameters and increase the number of iterations.

M I A
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30

Assignment P3 (2)

Problem 2 (Diffusion-Reaction Methods) (6 points)

The programme `diffreact.c` is an almost complete implementation of a modified explicit scheme (Lecture 11, Page 20) for a diffusion-reaction method using the Charbonnier diffusivity.

(a) Supplement it with the missing code for the reaction step (if $\alpha > 0$) and compile it with

```
gcc -O2 -o diffreac diffreac.c -lm
```

(b) Perform some comparisons between the diffusion-reaction method with data weight α and the pure diffusion method with stopping time α. The pure diffusion method results from using `diffreac` with $\alpha = 0$. You may use the images `einstein.pgm` and `sbrain.pgm`, and the programme `difference` from P1 – Assignment 1.

Assignment P3 (3)

For assessment:

Use `tar cvzf P3_`*yourname*`.tgz` *file1 file2 …* to pack the following files into an archive:

- the supplemented `diff_tensor.c`,
- a CED-enhanced version of `finger.pgm`,
- a representative CED-processed version of either `xmas.pgm` or a series (select maximal 5 images) of CED-processed versions of `fabric.pgm` visualising stripes at different scales,
- the supplemented `diffreac.c`,
- two filtering results for variational image filtering with different α,
- the corresponding filtering results from pure diffusion and the difference images.

Include in the archive also a short `README` file stating the parameters used for each of the submitted images. Please answer in the `README` file also the following questions:

- How could you see in task 1 that the extremum principle was violated?
- What were your reasons in task 2 to choose the stopping times (iteration counts) as you did in the diffusion-reaction method?

Send the archive by e-mail to the address specified by your tutor.

Deadline for electronic submission: Friday, May 30, 10 am.