TV Denoising, Inverse Problems, and Blood Vessel Ultrasound Slices

Student: Mihails Kosmans, Supervisor: Dr Björn Stinner

University of Warwick

mihails.kosmans@warwick.ac.uk

Motivation

Imagine for a moment that you captured an image. You could have used a camera to take a photo, a CT scanner to reconstruct a cross-section of someone's abdomen, or a telescope array to image a distant galaxy. All three cases have one particular thing in common — noise.

Be it the subtle motion of a sensor, the discrete nature of photon counting, or the imperfections inherent in manufacturing — noise is a common issue in data capture and analysis. Image processing concerns itself with extracting useful information from images corrupted by noise, and in this poster we will have a look at one such approach in the context of medical imaging.

A monochrome 2D image can be considered as a function u:

$$u : \Omega \subset \mathbb{R}^2 \to \mathbb{R},$$

mapping points in some subset Ω of the plane $\mathbb{R}^2$ (a rectangle for example) into the reals $\mathbb{R}$. For $x \in \Omega$, $u(x)$ then represents the brightness or *intensity* at the point x. Suppose now, that you have a noisy image of something for which u is the *true underlying image* — we denote this *noisy image* u^δ and define it analogously to u. By itself it is impossible to re-construct u from u^δ — the problem is said to be *ill-posed*, and so we make an educated guess about some property the image should have. We pick a function which measures in some way "how little of this property does the input image have?" (if there is little then $\mathcal{R}$ is large) — this function is the *regulariser* $\mathcal{R} : U \to [0, \infty) \cup \{\infty\}$, where U is a space in which we expect our image to lie. We also pick some function $\rho : U \times U \to [0, \infty)$ which measures "how different are the two input images?". Taking some $\alpha > 0$, we formulate the problem of denoising u^δ as that of finding a function u^δ_α satisfying:

$$u^\delta_\alpha \in \underset{u \in U}{\arg\min}(\rho(u, u^\delta) + \alpha\mathcal{R}(u)), \tag{1}$$

that is — u^δ_α is an element of the space U, which *minimises* the function $\rho(u, u^\delta) + \alpha\mathcal{R}(u)$. The role of the so-called *regularisation parameter* α is to regulate how much $\mathcal{R}$ affects u^δ_α. A higher value of α forces u^δ_α to have "more" of the property encoded in $\mathcal{R}$, whilst a lower one favours u^δ_α that are "closer" to u^δ.

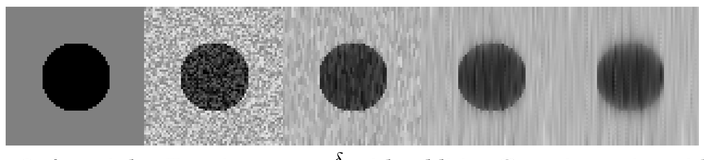

Left to right: True image u, u^δ with additive Gaussian noise with $\mu = 0$ and $\sigma = 0.114$, u^δ_α denoised via total variation regularisation using Chambolle's algorithm for $\alpha = 0.2, 1.0, 2.0$.

This setup happens to mimic the one of so-called *inverse problems*. Combined with the fact that ultrasound (US) scanning is in itself an inverse problem [2], this motivates the following section.

Inverse Problems

Given *data* $v \in V$, consider a function $F : \mathcal{D}(F) \subseteq U \to V$ with *domain of definition* (i.e. defined on) $\mathcal{D}(F)$, which given an image $u \in \mathcal{D}(F)$ has as its output v:

$$F(u) = v. \tag{2}$$

Given v we expect to recover u exactly by inverting F (assuming it's *injective*, that is — mapping distinct u to distinct v), but in practice, any recording of data is subject to noise, and we almost never actually have that exact v. *Noisy data* $v^\delta \in V$ could for instance represent recordings from microphones listening for sonar "pings", or it could be a sensor in an an MRI scanner recording radio emission, and we wish to obtain an approximation to u from it, be it a topographical map of the sea-floor or the shape of someone's heart. As before, the problem is ill-posed, and so again — we *regularise* it and search for some u^δ_α satisfying [4, Chap. 3]:

$$u^\delta_\alpha \in \arg\min_{u \in U} \mathcal{T}_{\alpha, v^\delta}(u),$$

where $\mathcal{T}_{\alpha, v^\delta}(u) := \rho(F(u), v^\delta) + \alpha\mathcal{R}(u)$ for all $u \in \mathcal{D}(F)$, and $\rho : V \times V \to [0, \infty)$ is now measuring the distance between elements of V instead of U as before. To make sure this actually makes sense, we *define* $\mathcal{T}_{\alpha, v^\delta}$ to be ∞ for any $u \notin \mathcal{D}(F)$.

We want to establish that unlike the previous problems, this one is *well-posed*, which for us will mean that:

- There really *exists* some minimiser $u^\delta_\alpha \in \arg\min_{u \in U} \mathcal{T}_{\alpha, v^\delta}(u)$.
- If we have two "similar" samples v^{δ_1} and v^{δ_2}, then the corresponding reconstructed images $u^{\delta_1}_\alpha$ and $u^{\delta_2}_\alpha$ are also somehow "similar". This property is called *stability*.
- Finally, if the v^δ is taken to approach v, i.e. $v^\delta \to v$ and we let the effects of the regulariser diminish by taking $\alpha \to 0$, then we expect that $u^\delta_\alpha \to u$, i.e. — that the solution tends towards the true underlying image. This property is *convergence*.

Two extra properties that would also be nice to know are — "how fast" does $u^\delta_\alpha \to u$ as $v^\delta \to v$ and $\alpha \to 0$ (*convergence rates*), and if we vary v^δ how much would u^δ_α vary (*stability estimates*).

It turns out [4, Chap. 3.1] that for spaces U and V with norms $\|\cdot\|_U$ and $\|\cdot\|_V$ and some suitable topologies τ_U and τ_V, whilst picking $\rho(\tilde{v}, v)$ to be $\|\tilde{v} - v\|^p_V$ for some $1 \leq p < \infty$, under sufficient assumptions on $U, V, \mathcal{R}$, and F (primarily about convexity and continuity properties) the three main results can be established. So, taking these assumptions as given [4, p. 63–66]:

Theorem 1 $\arg\min_{u \in U} \mathcal{T}_{\alpha, v^\delta}(u)$ *is nonempty.*

That is, minimisers from which we can pick u^δ_α do actually exist!

Theorem 2 *If $\{v_n\}_{n=1}^\infty \subseteq V$ is a sequence that converges to v^δ in norm, then for any sequence $\{u_n\}_{n=1}^\infty$ with:*

$$u_n \in \underset{u \in U}{\arg\min}\, \mathcal{T}_{\alpha, v_n}(u)$$

there is a subsequence, $\{u_{n_k}\}_{k=1}^\infty$, that is convergent in τ_U. The limit of any such τ_U-convergent subsequence is a minimiser $\tilde{u} \in \arg\min_{u \in U} \mathcal{T}_{\alpha, v^\delta}(u)$, and $\mathcal{R}(u_{n_k}) \to \mathcal{R}(\tilde{u})$ as $k \to \infty$.

Said differently — if $v_n \to v^\delta$ in norm, and if for each n we pick an arbitrary minimiser u_n of $\mathcal{T}_{\alpha, v_n}$, then as $n \to \infty$ the u_n are guaranteed to get closer and closer to $\arg\min_{u \in U} \mathcal{T}_{\alpha, v^\delta}(u)$.

For the following, note that an $\mathcal{R}$-minimising solution $u^\dagger$ of (2) is one that satisfies $\mathcal{R}(u^\dagger) = \min\{\mathcal{R}(u) : u \in \mathcal{D}(F), F(u) = v\}$.

Theorem 3 *If (2) has a solution u with $\mathcal{T}_{\alpha, v}(u) < \infty$, if as a function of δ, $\alpha : (0, \infty) \to (0, \infty)$ satisfies:*

$$\alpha(\delta) \to 0 \text{ and } \frac{\delta^p}{\alpha(\delta)} \to 0, \text{ as } \delta \to 0, \tag{3}$$

if a sequence $\{\delta_n\}_{n=1}^\infty$ converges to 0 with $||v - v^{\delta_n}||_V \leq \delta_n$ for all $n \in \mathbb{N}$, then every sequence $\{u_n\}_{n=1}^\infty$ of elements minimizing $\mathcal{T}_{\alpha(\delta_n), v^{\delta_n}}$, has a subsequence $\{u_{n_k}\}_{k=1}^\infty$ that is convergent with respect to τ_U. The limit, denoted $u^\dagger$, of each τ_U-convergent sequence $\{u_{n_k}\}_{k=1}^\infty$ is an $\mathcal{R}$-minimising solution of (2), and $\mathcal{R}(u_n) \to \mathcal{R}(u^\dagger)$ as $n \to \infty$. If in addition, the $\mathcal{R}$-minimising solution $u^\dagger$ is unique, then $u_n \to u^\dagger$ with respect to τ_U as $n \to \infty$.

In other words, if the "strength" of the noise δ_n tends to 0, and if we decrease α to 0 slowly enough for the fraction in (3) to also tend to 0, then if we pick an arbitrary minimiser u_n of $\mathcal{T}_{\alpha(\delta_n), v^{\delta_n}}$ for each n, then as $n \to \infty$ the u_n are guaranteed to get closer and closer to the $\mathcal{R}$-minimising solutions $u^\dagger$ of (2). If there is only one such $u^\dagger$ in existence, then the u_n approach that exact $u^\dagger$.

Given some extra assumptions, the two additional "nice to know" results can also be established.

Total Variation Regularisation

We now look at a practical example. First, we define the regulariser $\mathcal{R}_1$ for functions $u \in L^1(\Omega)$ as done in [3, p. 4]:

$$\mathcal{R}_1(u) := \sup\left\{\int_\Omega u\nabla \cdot \varphi dx : \varphi \in C^1_c(\Omega; \mathbb{R}^2), ||\varphi(x)||_{L^\infty(\Omega)} \leq 1\right\}.$$

We call it the "*total variation*" of u in Ω. Dropping some details, $\mathcal{R}_1$ can be thought of as a measure of how oscillatory a given function is within Ω. $u \in L^1(\Omega)$ just means that $\int_\Omega |u| dx < \infty$ (up to some equivalence).

Assuming Ω satisfies some additional conditions (which an open rectangle does), the following result can be established:

Theorem 4 *Suppose $L : U \to V$ is bounded and linear, with $U := L^2(\Omega)$ and V a Hilbert space. If $\mathcal{D}(F) \neq \emptyset$ is closed and convex, and $F(u) = L(u)$ for any $u \in \mathcal{D}(F)$, and if:*

$$\mathcal{T}^{TV}_{\alpha, v}(u) := ||F(u) - v||^2_{L^2(\Omega)} + \alpha\mathcal{R}_1(u) < \infty$$

for some $u \in \mathcal{D}(F)$, then the problem of minimising $\mathcal{T}^{TV}_{\alpha, v^\delta}$ satisfies theorems 1, 2, and 3.

Here, the space $L^2(\Omega)$ is the normed space of all functions (up to some equivalence) for which:

$$\|u\|_{L^2(\Omega)} := \left(\int_\Omega |u|^2 dx\right)^{\frac{1}{2}} < \infty.$$

The regulariser $\mathcal{R}_1$ integrates in some sense — the absolute value or "strength" of the "derivative" of u, and in doing so *penalises* any solutions to our problem that vary quickly in Ω. Additionally, $\mathcal{R}_1$ allows for discontinuities in the image data, in-particular *edges*. This means that $\mathcal{R}_1$ can be a good guess for a property we might want if we expect the feature we want to extract to be some relatively "flat" region of the image with a sharp edge. This leads us nicely into denoising.

Denoising

It turns out that choosing $V := L^2(\Omega) = U$, and $F := \mathrm{Id} : U \to U = V$ where $\mathrm{Id}(u) := u$ for all $u \in U$, satisfies theorem 4. $L^2(\Omega)$ is a known Hilbert space and Id is bounded and linear (i.e. $\|\mathrm{Id}(u)\|_V \leq C\|u\|_U$ and $\mathrm{Id}(\beta\tilde{u} + u) = \beta\,\mathrm{Id}(\tilde{u}) + \mathrm{Id}(u)$ for any $\beta \in \mathbb{R}$, any $u \in U$, and some $C > 0$). By substituting in our choice of F and noting that we now have $v = u$, we come full-circle to (1). Our problem is now to find some u^δ_α such that:

$$u^\delta_\alpha \in \underset{u \in U}{\arg\min}(\|u - u^\delta\|^2_V + \alpha\mathcal{R}_1(u)). \tag{4}$$

We will have a look at the results of denoising US scan data representing slices of a blood vessel in the forearm. US waves are scattered most strongly at the interfaces of materials with different densities, and the interface between blood and the tissues surrounding it does just that. This means that on a typical US scan slice, the blood vessel appears as a dark, and relatively flat region of the image. This is why we choose $\mathcal{R}_1$ as our regulariser.

Discretisation

Problem (4) is continuous, but computers only operate on discrete data. To match this restriction we have to transform our problem. Following [1, Chap. 6], we assume that our u is defined on a rectangular domain Ω with integer side-lengths, and break it up into unit squares (pixels). We assign each pixel a corresponding intensity value that represents how bright that individual pixel is. If each square is indexed by j running left to right (think "x-axis") and i running from top to bottom (think "y-axis"), from 1 to N_i and N_j respectively, then we can say that the i, j'th square has intensity $u_{i,j} \in [0, \infty)$ and think of it as somehow being at the centre of the square. This motivates the concept of a discrete gradient $(\nabla u)_{i,j} := ((\nabla_x u)_{i,j}, (\nabla_y u)_{i,j})$ where:

$$(\nabla_x u)_{i,j} := \begin{cases} u_{i+1,j} - u_{i,j} \text{ if } i < N_i, \\ 0 \text{ if } i = N_i, \end{cases} \quad \text{and,}$$

$$(\nabla_y u)_{i,j} := \begin{cases} u_{i,j+1} - u_{i,j} \text{ if } j < N_j, \\ 0 \text{ if } j = N_j. \end{cases}$$

The integral can be discretely approximated as a sum over the pixels, providing an analogue of the function in (4)'s RHS:

$$\sum_{i,j} |u_{i,j} - u^\delta_{i,j}|^2 + \alpha \sum_{i,j} |(\nabla u)_{i,j}|, \tag{5}$$

where $u^\delta_{i,j}$ is the discrete analogue of the noisy image u^δ, and $|(\nabla u)_{i,j}| := \sqrt{(\nabla_x u)^2_{i,j} + (\nabla_y u)^2_{i,j}}$. Minimisation of (5) can then be performed using *"Chambolle's Algorithm"* [1, Chap 3].

Computational Experiments

In our problem, an US probe moves along a patient's forearm, recording the intensity of scattered waves as they return to the probe. The resulting data is inverted to produce a sequence of images representing the densities of material in the cross-sections along the arm. These images are "stacked" to produce a 3D plot of densities, which may then be *segmented* on a computer to produce a 3D model of a blood vessel.

Noise impedes the segmentation process as it makes it hard to tell where boundaries between different materials lie. We apply Chambolle's algorithm to such a data sample, to see if it aids segmentation, which in this case is using the so-called "grow from seed" algorithm with identical starting parameters for each image set. Going from the top, pairs of images represent a sample US slice and an associated segmentation. The top image is noisy, whilst the rest were denoised with $\alpha = 0.2, 0.4, 1.0$ respectively:

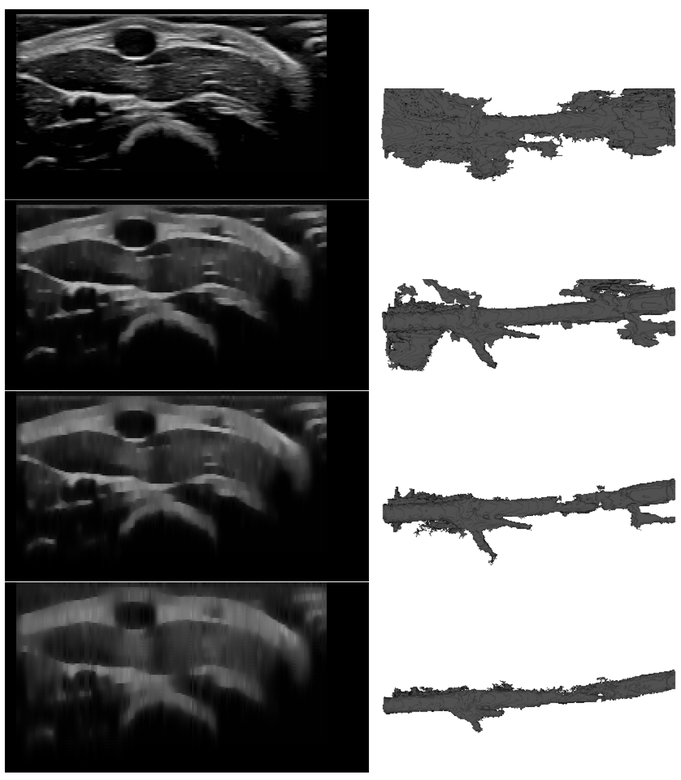

The first segmentation is quite poor, and it's hard to recognize it as a blood vessel at all. The second is closer to something we might expect, albeit still has stray elements. The third has fewer stray elements, but a section of what we might expect to be the blood vessel is missing. The last has few stray elements, but a section of the blood vessel at the bottom that was likely a fork where another vessel split off, is heavily truncated.

Although TV regularisation shows some promise, it nonetheless seems clear that TV alone is insufficient for this task. Approaches combining these classical methods from inverse problems with machine learning, and those taking into account the elliptical geometry of blood vessels may play a helpful role in producing a more faithful reconstruction of a blood vessel.

Acknowledgements

I wish to thank my supervisor Dr Björn Stinner for supervising this project, Risa Romy (PGR) for providing real-life context for the theory to be applied to, and the Royal Berkshire NHS Foundation Trust for providing sample image data for processing.

References

[1] Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.

[2] David L. Colton and Rainer Kress. *Inverse acoustic and electromagnetic scattering theory*, volume 93. Springer, Cham, 4th edition, 2019.

[3] Ivar Ekeland and Roger Temam. *Convex analysis and variational problems*. SIAM, 1999.

[4] Otmar Scherzer, Markus Grasmair, Harald Grossauer, Markus Haltmeier, and Frank Lenzen. *Variational Methods in Imaging*. Springer, New York, 2009.