A Restricted-Domain Dual Formulation for Two-Phase Image Segmentation

Jack Spencer

Department of Mathematics,
University of Liverpool, UK.

Abstract

In two-phase image segmentation, convex relaxation has allowed global minimisers to be computed for a variety of data fitting terms. Many efficient approaches exist to compute a solution quickly. However, we consider whether the nature of the data fitting in this formulation allows for reasonable assumptions to be made about the solution that can improve the computational performance further. In particular, we employ a well known dual formulation of this problem and solve the corresponding equations in a restricted domain. We present experimental results that explore the dependence of the solution on this restriction and quantify imrovements in the computational performance. This approach can be extended to analogous methods simply and could provide an efficient alternative for problems of this type.

Keywords: Image Processing, Segmentation, Total Variation, Convex Relaxation, Dual Formulation.

1 Introduction

Image segmentation is the meaningful partitioning of an image based on certain characteristics. In two-phase segmentation this consists of determining the foreground and background of a domain $\Omega \in \mathbb{R}^2$, i.e. find a closed boundary separating subregions Ω_1 and Ω_2. This is distinct from multiphase approaches, where more than two separate regions are determined. Our work concerns the continuous setting, which we will briefly discuss in the next section. Equivalent problems in the discrete setting have been well studied with details found in [Boykov and Kolmogorov, 2004]. A comprehensive background behind the following functional can also be found in [Chambolle and Pock, 2016]. Briefly, the aim is to determine an indicator function, $u(x)$, that labels the foreground and background by minimising the following energy:

$$\min_{u\in\{0,1\}} \left\{ \int_\Omega |\nabla u(x)| \, dx + \lambda \int_\Omega f(x)u(x)dx \right\}. \tag{1}$$

The function $f(x)$ is typically referred to as the fitting term, determining how the segmentation solution corresponds to the data. It is balanced by a regularisation term, in this case the total variation (TV) semi-norm which penalises the length of the segmentation boundary. For large λ the following will hold precisely:

$$\begin{aligned} f(x) < 0, &\quad \text{foreground,} \\ f(x) > 0, &\quad \text{background.} \end{aligned}$$

When λ is varied the solution for u will fit the data with more regularity, i.e. some areas where $f < 0$ will be background and some areas where $f > 0$ will be foreground. However, this is most likely where f is close to 0. With that in mind our work considers what improvements can be made by concentrating on regions in the domain where $f(x) \approx 0$. An exception to this concerns cases where f has strong noise, which we will return to

later. A number of choices for f exist depending on the application, including piecewise-constant segmentation based on the work of [Chan and Vese, 2001]]:

$$f(x) = (z - c_1)^2 - (z - c_2)^2, \tag{2}$$

where c_1 and c_2 are intensity constants indicating average foregound and background intensitites of the image $z(x)$, respectively. We also consider a selection based fitting term such as [Spencer and Chen, 2015]

$$f(x) = (z - c_1)^2 - (z - c_2)^2 + \gamma P(x), \tag{3}$$

where $\gamma P(x)$ is a distance selective term based on user input. In Section 4 we present results using equations (2) and (3). This approach is not limited to the fitting terms mentioned above, and can be extended to any segmentation problem in this framework. Alternatives for future consideration include bias field segmentation [Chen et al., 2013] and interactive convex active contours [Nguyen et al., 2012].

A restricted-domain approach is analogous to banded segmentation methods such as [Rommelse et al., 2003] and [Zhang et al., 2014], among many others. This work differs in the sense that the restriction is based on the values of the fitting term, rather than the location of the boundary at an iteration, which is potentially simpler computationally. We compute an approximation of the global minimiser of the energy (1), with the accuracy determined by the level of domain restriction. This is based on the following initialisation of the indicator function: $u^{(0)} = H(-f)$, where $H(\cdot)$ is the Heaviside function.

In the following we will briefly introduce existing methods for finding the global minimisers of two-phase segmentation problems, introducing the dual formulation of [Bresson et al., 2007] based on [Chambolle, 2004]. We then discuss the proposed approach where a restricted domain based on the fitting term is considered, before detailing the method and how it relates to [Bresson et al., 2007]. Finally, we present some results for three examples for various restrictions on the domain, quantifying the accuracy and computational performance in comparison to the original method. We then offer some concluding remarks.

2 Convex Relaxation for Two-Phase Segmentation

We now introduce the details of the approach we consider in this work. Again, a comprehensive background of this work can be found in [Chambolle and Pock, 2016] and many others. Essentially, convex relaxation in this case involves relaxing the binary constraint in the original functional (1), i.e. $u \in [0, 1]$. The seminal work here is [Chan et al., 2006] who found global minimisers of the two-phase piecewise-constant Mumford-Shah model [Mumford and Shah, 1989] (assuming fixed intensity constants, c_1, c_2). Therefore, the problem considered here is:

$$\min_{u \in [0,1]} \left\{ \int_\Omega |\nabla u(x)| \, dx + \lambda \int_\Omega f(x) u(x) dx \right\}. \tag{4}$$

In [Chan et al., 2006] they introduce a penalty function to enforce the constraint on u and solve using time marching. It is also possible to use additive operator splitting [Spencer and Chen, 2015], split Bregman [Goldstein et al., 2010], and Chambolle-Pock [Chambolle and Pock, 2011] among many others. However, initially we intend to implement our restricted-domain approach on the dual formulation used in [Bresson et al., 2007]. We breifly detail this approach next.

2.1 Dual Formulation

The dual formulation of this problem was first introduced by [Bresson et al., 2007], based on the work of [Chambolle, 2004], [Aujol et al., 2006] and the references therein. The idea is to introduce a new variable, $v(x)$, and minimise the following functional alternately:

$$\min_{u,v} \left\{ \int_\Omega |\nabla u(x)| \, dx + \frac{1}{2\theta} \int_\Omega (u(x) - v(x))^2 \, dx + \int_\Omega \lambda f(x) v(x) + \alpha \psi(v) \, dx \right\}, \tag{5}$$

where $\psi(v) = \max\{0, 2|v-1/2|-1\}$. By splitting the variables in this way, the minimisation of u concentrates on the TV term, and the minimisation of v satisifes the fitting and constraint requirements. In [Bresson et al., 2007] the regularisation term is weighted, however here we concentrate on the original problem (4). The parameter α ensures the constraints on the indicator function $u(x)$ in (5) are met, and can be set automatically [Chan et al., 2006]. The minimistion of u and v can be achieved iteratively by the following steps. With fixed v, the solution of u is given by

$$u(x) = v(x) - \theta\nabla \cdot \rho(x),$$

where $\rho = (\rho^1, \rho^2)$ is the solution of

$$\nabla(\theta\nabla \cdot \rho - v) - |\nabla(\theta\nabla \cdot \rho - v)|\rho = 0,$$

which can be solved by a fixed point method. With fixed u, the solution for v is given as

$$v(x) = \min\{\max\{u(x) - \theta\lambda r(x), 0\}, 1\}.$$

This is repeated until convergence. Further details can be found in [Bresson et al., 2007], including the definition of the discrete gradient and divergence operators from [Chambolle, 2004].

3 Proposed Approach: Segmentation in a Restricted Domain

We now introduce our approach for reducing the computation time for this type of problem. We begin by assuming the solution in certain parts of the discretised domain based on the values of $f(x)$ for a given problem. The indicator function is then fixed at 1 or 0 at these points for the foreground (Fg) and background (Bg) respectively. We solve the equation in the remaining region, which we call the restricted domain (RD). Let us define $q \in [0, 1]$ such that the following thresholding of the fitting function holds. We define a value $\hat{q} \in \mathbb{R}$ such that the percentage of nodes in the restricted domain is q ($\times 100$):

$$\begin{cases} Fg &= \{x : x \in \Omega,\ f(x) \leq -\hat{q}\} \\ Bg &= \{x : x \in \Omega,\ f(x) \geq \hat{q}\} \\ RD &= \{x : x \in \Omega \setminus Fg \setminus Bg\}. \end{cases}$$

The value of $\hat{q}$ is initially 0 and is increased until the selected value of q is satisfied. In other words, for $q = 0$, $RD = \emptyset$ and the solution is a zero-thresholding of the fitting function, or equivalent to selecting a large λ in the original problem (4). For $q = 1$, $Fg = \emptyset$ and $Bg = \emptyset$ and we consider the problem in a conventional manner with no restriction on the domain. For $q \in (0, 1)$ we consider a restricted domain of varying degrees as illustrated in Figure 1, where the corresponding region of interest is given in grey. This means that we need to minimise the energy in a restricted domain which, when combined with the dual formulation of [Bresson et al., 2007] and [Chambolle, 2004], means that we are solving the following equation for ρ at certain points:

$$\nabla(\theta\nabla \cdot \rho - v) - |\nabla(\theta\nabla \cdot \rho - v|\rho = 0. \tag{6}$$

This involves making certain assumptions about the solution for ρ. In [Bresson et al., 2007], it is initialised as $\rho^{(0)} = 0$ and the solution of (6) is clearly dependent on u and v. However, for the initialisation $u^{(0)} = -H(f)$ the solution of ρ has a predictable form. In particular, $\rho \approx 0$ where $|f|$ is largest and $\rho \in [-1, 1]$ where $|f|$ is closer to 0. If q is selected sensibly (we will return to this later), when $u^{(0)} = -H(f)$ it is reasonable to assume that the solution of (6) for $x \in Fg \cap Bg$ is $\rho^*(x) = 0$. Clearly, the larger q is the less reliable this assumption becomes and the corresponding solution for ρ will be less accurate. However, part of this work concerns what consists of a sensible selection for q and whether it is possible to make reasonable assumptions about f that can improve the efficiency of minimising the original formulation (4).

We now ellaborate on the details behind minimising (4) with a dual formulation in a restricted domain. We first consider the following minimisation problem:

$$\min_u \left\{ \int_\Omega |\nabla u(x)|\ dx + \frac{1}{2\theta} \int_\Omega (u(x) - v(x))^2\ dx \right\}.$$

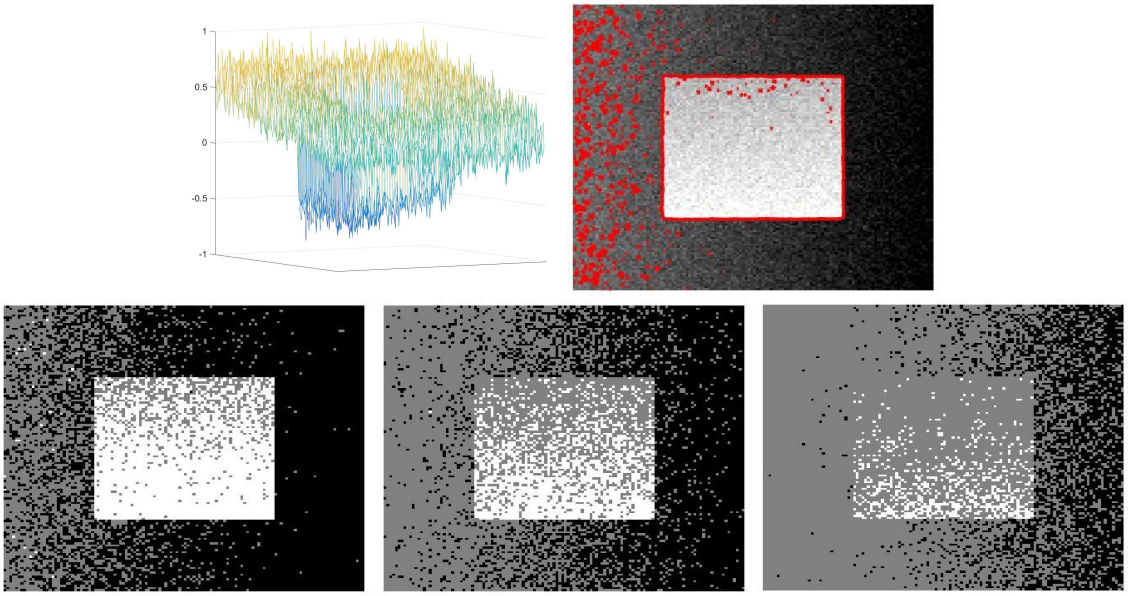

Figure 1: Restricted Domain: The top row shows the fitting term (left) and the image (right) with the zero level-set of $f(x)$ given in red. The bottom row indicates the corresponding regions Fg (white), Bg (black), and RD (grey) for values of $q = 0.25, 0.5, 0.75$ from left to right, respectively.

The solution, based on our approach, is given by

$$u(x) = \begin{cases} 1, & \text{for } x \in Fg \\ 0, & \text{for } x \in Bg \\ v - \theta \nabla \cdot \rho, & \text{for } x \in RD, \end{cases} \tag{7}$$

where $\rho = (\rho_1, \rho_2)$ satisfies

$$\rho = \begin{cases} 0, & \text{for } x \in Fg \cap Bg \\ \nabla(\theta \nabla \cdot \rho - v) - |\nabla(\theta \nabla \cdot \rho - v)|\rho = 0, & \text{for } x \in RD. \end{cases} \tag{8}$$

For $x \in RD$ the following fixed point method, with time step τ, will solve the equation for ρ:

$$\rho^{n+1} = \frac{\rho^n + \tau \nabla(\nabla \cdot \rho^n - v/\theta)}{1 + \tau |\nabla(\nabla \cdot \rho^n - v/\theta)|}$$

As before, the following minimisation problem is then solved with u fixed:

$$\min_v \left\{ \frac{1}{2\theta} \int_\Omega (u(x) - v(x))^2 \, dx + \int_\Omega \lambda f(x) v(x) + \alpha \psi(v) \, dx \right\}.$$

We combine our assumptions about u and ρ with the work of [Bresson et al., 2007] to give the corresponding solution as

$$v(x) = \begin{cases} 1, & \text{for } x \in Fg \\ 0, & \text{for } x \in Bg \\ \min\{\max\{u(x) - \theta\lambda f(x), 0\}, 1\}, & \text{for } x \in RD. \end{cases} \tag{9}$$

As with [Bresson et al., 2007], as discussed in the previous section, u and v are minimised alternately until convergence. The main advantage of this approach concerns finding the solution of ρ at each iteration with the fixed point method detailed above. Based on the choice of q it is possible that significant advantages exist in terms of computation time with minimal compromise on the quality of the solution. We will discuss some exceptions to this, as well as future considerations in the following sections.

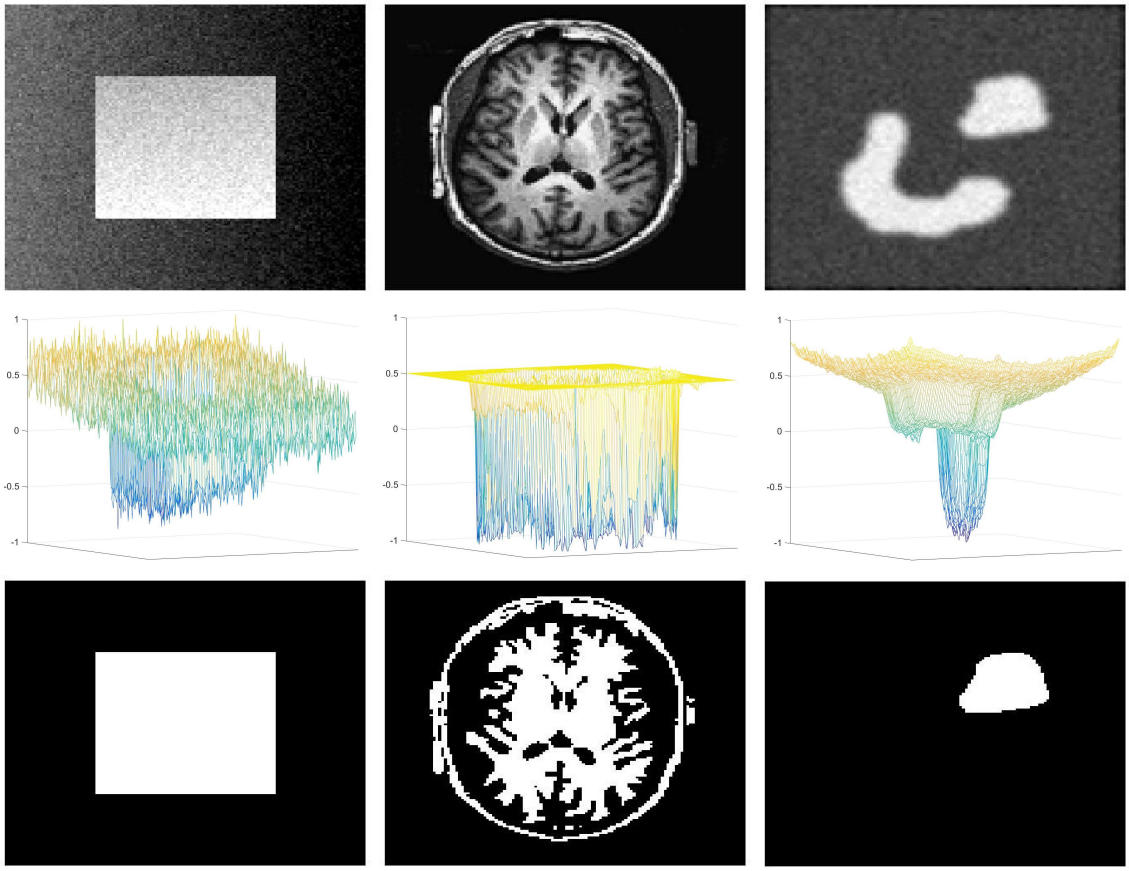

Figure 2: Test Problems: From left to right are Examples 1-3. From top to bottom are the image (z), fitting term (f), and thresholded segmentation result (GT) using the original method of [Bresson et al., 2007].

q	Example 1		Example 2		Example 3	
	E_1	E_2	E_1	E_2	E_1	E_2
0	0.876	5.64×10^{2}	0.943	2.49×10^{2}	0.920	5.59×10^{1}
0.1	0.961	1.59×10^{2}	0.946	1.47×10^{2}	0.962	1.19×10^{1}
0.2	0.987	5.16×10^{1}	0.964	8.90×10^{1}	0.969	1.05×10^{1}
0.3	0.995	1.94×10^{1}	0.982	4.94×10^{1}	0.974	8.02×10^{0}
0.4	0.999	6.76×10^{0}	0.987	3.47×10^{1}	0.989	3.45×10^{0}
0.5	1.000	3.24×10^{0}	0.988	2.98×10^{1}	0.995	7.55×10^{-1}
0.6	1.000	2.65×10^{0}	0.989	2.90×10^{1}	0.999	6.03×10^{-1}
0.7	1.000	1.24×10^{0}	0.989	2.90×10^{1}	1.000	4.78×10^{-1}
0.8	1.000	5.34×10^{-1}	0.989	2.90×10^{1}	0.999	4.24×10^{-1}
0.9	1.000	2.49×10^{-1}	0.989	2.73×10^{1}	1.000	3.30×10^{-1}
1	1.000	2.10×10^{-5}	1.000	3.28×10^{-2}	1.000	2.10×10^{-1}

Table 1: Results: For Examples 1-3 we vary $q \in [0, 1]$ and provide E_1 and E_2.

4 Experimental Results

In this section we introduce some results for the test problems shown in Figure 2, using (2) for f in Examples 1 and 2 and (3) for f in Example 3. The focus of these results is to determine the dependence on q, i.e. to what extent can we restrict the domain for problems of this type? As a comparison, we use a result from [Bresson et al., 2007] (for a particular choice of λ in each case). Specifically, we iterate until the following stopping criterion is met at the ℓ^{th} iteration:

$$\max\left\{\| u^{(\ell)} - u^{(\ell-1)} \|, \| v^{(\ell)} - v^{(\ell-1)} \|\right\} \leq \delta.$$

For $\delta = 10^{-10}$ we set $u^{GT}(x) = u^{(\ell)}$ in the original dual formulation. We also use the thresholded result:

$$GT(x) = \begin{cases} 1, & \text{for } x \in u^{GT}(x) > \epsilon \\ 0, & \text{for } x \in u^{GT}(x) \leq \epsilon. \end{cases}$$

Following convention we set $\epsilon = 0.5$. We refer to the solution for the proposed method (with $\delta = 10^{-2}$) as u^* and its corresponding thresholded result as Ω_1^*. This allows us to define the two error measurements that we use in discussing the results when varying the parameter in the proposed method, q. The first is the Tanimoto Coefficient between the thresholded results, and the second is the L^2 difference between the proposed solution and the original solution:

$$E_1 = \frac{N(GT \cap \Omega_1^*)}{N(GT \cup \Omega_1^*)}, \qquad E_2 = \int_\Omega (u^* - u^{GT})^2 \, dx.$$

Here $N(\cdot)$ refers to the number of nodes in the enclosed region and $E_1 \in [0,1]$, with $E_1 = 1$ indicating a perfect result. With the second error measurement, clearly we expect E_2 to approach 0 as q increases. We don't necessarily expect E_2 to be 0 for $q \leq 1$ as u^{GT} is not binary precisely and in the tests we use $\delta = 10^{-2}$, so there is likely to be a minor difference. Whilst E_2 is a useful measure to demonstrate the correspondence between the value of q and the original method, we are primarily interested in E_1 as the crucial indicator of a successful segmentation result.

In Table 1 we present the main results for $q \in [0,1]$. We include $q = 0$ (i.e. a completely thresholded result) and $q = 1$ (i.e. the original method) to demonstrate the full effect of the choice of q. Both error measurements are included (to 3 s.f.) and we can see that increasing q to 0.4 is enough to produce a very good result ($E_1 > 0.98$) in all examples. For Example 2 we can see that for $q < 1$, E_1 does not reach 1 meaning that restricting the domain of the dual formulation always changes the segmentation result for this fitting term. However, there is a very close correspondence between the results even for small values of q, which is encouraging. As expected E_2 tends to decrease as q increases, and the solution in the restricted domain is reasonably close to the original solution. To put these results in context the size of all images tested here are 128×128. These results demonstrate that using a dual formulation in a restricted domain is a viable approach for problems of this type.

In Figure 3 we include the computation time (in seconds, to 1 d.p.) for different choices of $q \in [0,1]$. Clearly the expectation is that as q increases t should also increase, but this does not quite hold absolutely. This is likely to be down to features in these fitting terms that mean minor increases to q, perhaps counterintuitively, slightly simplify the problem. For the original method of [Bresson et al., 2007] the computation time was $t = 6.0s$ for Example 1, $t = 14.0s$ for Example 2, and $t = 8.2s$ for Example 3 (which we will refer to as t_1, t_2, and t_3 respectively). For these results, and in the following, we set $\delta = 10^{-2}$ as a stopping criterion, and use them as a benchmark in each case. From Figure 3 it can be seen that for Examples 1-3 t is only greater than t_1, t_2, or t_3 for $q > 0.9$. For smaller values of q it is possible to make significant gains in terms of computation time. For Example 1, when $E_1 > 0.98$ the average time is $t = 2.8s$ for $q \leq 0.9$. Similarly for Examples 2 and 3, the average times are $t = 8.7s$ and $t = 5.5s$ respectively. This corresponds to a time saving of 53%, 38%, and 33% for Examples 1-3, for cases with an accurate segmentation. Extending these tests to a wider choice of fitting terms, and investigating the effect of changing λ and δ, would help further determine the effectiveness of the proposed approach.

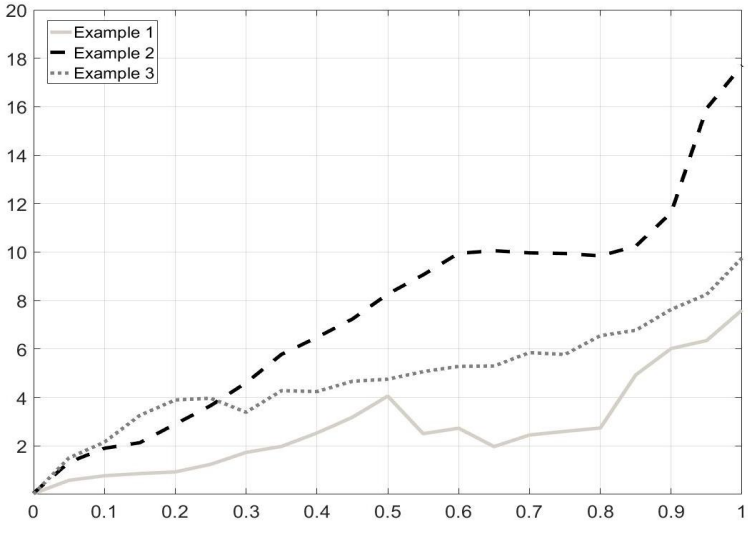

Figure 3: Computation Time, $t(q)$.

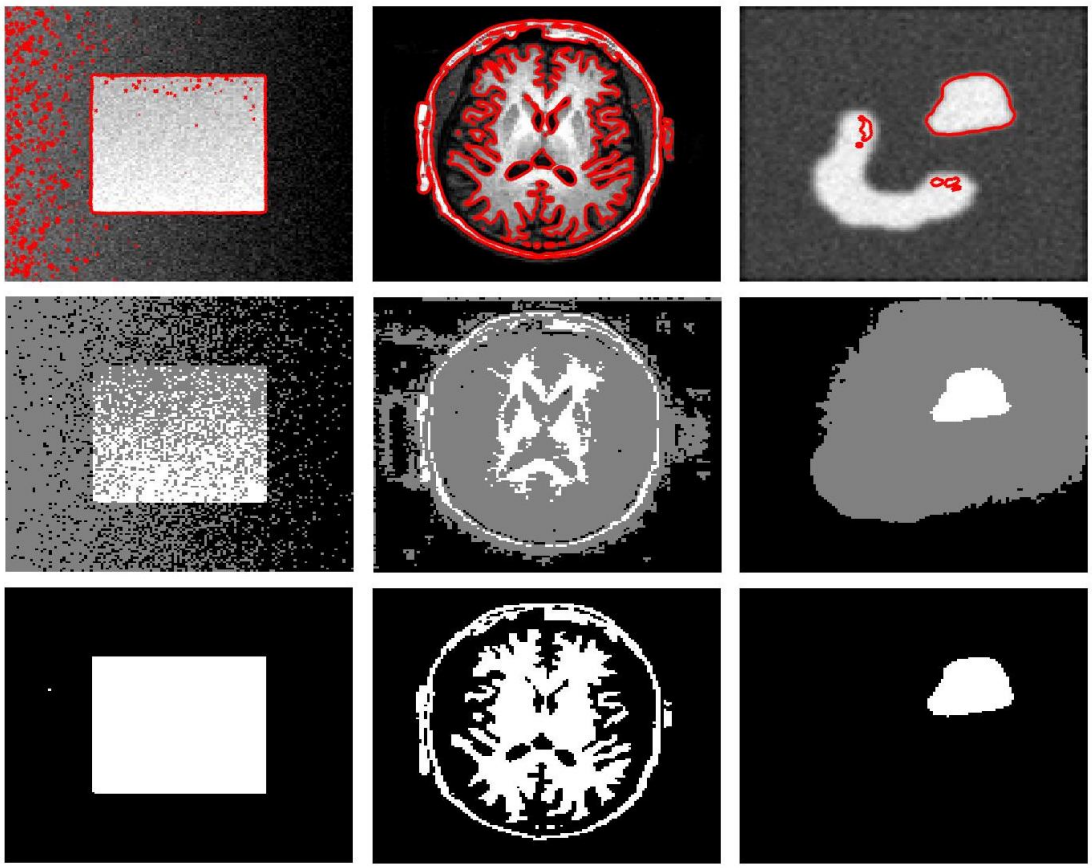

Figure 4: Results: From left to right are Examples 1-3, respectively. From top to bottom are the image with the zero level-set of f in red, the regions Fg (white), Bg (black), RD (grey), and Ω_1^* for $q = 0.5$.

In Figure 4 we present some example results for $q = 0.5$. For Example 1, $E_1 = 1.000$ and $t = 4.1s$. For Example 2, $E_1 = 0.988$ and $t = 8.3s$. For Example 3, $E_1 = 0.995$ and $t = 4.8s$. Compared to the original method, $t_1 = 6.0s$, $t_2 = 14.0s$, and $t_3 = 8.2s$. In each case a significant improvement can be made in terms of computation time with minimal compromise on the quality of the result as measured by E_1.

5 Conclusion

The results presented support the idea that the domain can be restricted for problems of this type, without compromising the quality of the result. This allows for significant gains in terms of computation time. Additional testing to verify these findings would be beneficial, particularly for a wider variety of fitting terms. An example would be where f contains high levels of noise. Further considerations might be necessary to restrict the domain in a robust way, and developing the required understanding would help develop this approach further. For $q = 1$ the computation times for Examples 1-3 are $t = 7.6s, 17.7s$, and $9.8s$ respectively, which corresponds to approximately a 25% increase for the proposed method when no restriction of the domain is considered. If the efficiency of the domain restriction could be improved this would allow for higher values of q to be selected for a reduced cost, which could be particularly beneficial in cases of high noise in the fitting term.

The results presented are for images of size 128×128 in order to explore the viability of restricting the domain for this problem. Improvements in t for larger images, or 3D problems, could be particularly valuable. Extending this approach to these cases is of interest, and would help support the proposed approach further. We are also considering the extension of this approach beyond the dual formulation of [Bresson et al., 2007], such as split-Bregman [Goldstein et al., 2010] and additive operator splitting [Spencer and Chen, 2015]. Following the framework introduced here, assumptions about the solution of (4) can be adapted to other methods in a similar way. However, the initial results presented here are encouraging.

Acknowledgements

The author would like to acknowledge the support of the EPSRC grant EP/N014499/1.

References

[Aujol et al., 2006] Aujol, J. F., Gilboa, G., Chan, T., and Osher, S. (2006). Structure-texture decomposition–modeling, algorithms, and parameter selection. *International Journal of Computer Vision*, 67(1):111–136.

[Boykov and Kolmogorov, 2004] Boykov, Y. and Kolmogorov, V. (2004). An experimental comparison of min-cut/max-flow algorithms for minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137.

[Bresson et al., 2007] Bresson, X., Esedoglu, S., Vandergheynst, P., Thiran, J. P., and Osher, S. (2007). Fast global minimization of the active contour/snake model. *Journal of Mathematical Imaging and Vision*, 28(2):151–167.

[Chambolle, 2004] Chambolle, A. (2004). An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20:89–97.

[Chambolle and Pock, 2011] Chambolle, A. and Pock, T. (2011). A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40:120–145.

[Chambolle and Pock, 2016] Chambolle, A. and Pock, T. (2016). An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319.

[Chan et al., 2006] Chan, T., Esedoglu, S., and Nikolova, M. (2006). Algorithms for finding global minimizers of image segmentation and denoising models. *SIAM Journal on Applied Mathematics*, 66(5):1632–1648.

[Chan and Vese, 2001] Chan, T. and Vese, L. (2001). Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277.

[Chen et al., 2013] Chen, D., Yang, M., and Cohen, L. (2013). Global minimum for a variant Mumford-Shah model with application to medical image segmentation. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 1(1):48–60.

[Goldstein et al., 2010] Goldstein, T., Bresson, X., and Osher, S. (2010). Geometric applications of the split bregman method. *Journal of Scientific Computing*, 45(1-3):272–293.

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

[Nguyen et al., 2012] Nguyen, T., Cai, J., Zhang, J., and Zheng, J. (2012). Robust interactive image segmentation using convex active contours. *IEEE Transactions on Image Processing*, 21:3734–3743.

[Rommelse et al., 2003] Rommelse, J., Lin, H., and Chan, T. (2003). A robust level set algorithm for image segmentation and its parallel implementation. *UCLA CAM Report*, 03-05.

[Spencer and Chen, 2015] Spencer, J. and Chen, K. (2015). A convex and selective variational model for image segmentation. *Communications in Mathematical Sciences*, 13(6):1453–1472.

[Zhang et al., 2014] Zhang, J., Chen, K., Yu, B., and Gould, D. (2014). A local information based variational model for selective image segmentation. *Inverse Problems and Imaging*, 8(1):293–320.