A New Model and Simple Algorithms for Multi-Label Mumford-Shah Problems

Byung-Woo Hong
Chung-Ang University
hong@cau.ac.kr

Zhaojin Lu and Ganesh Sundaramoorthi
King Abdullah University of Science and Technology (KAUST)
{zhaojin.lu,ganesh.sundaramoorthi}@kaust.edu.sa

Abstract

In this work, we address the multi-label Mumford-Shah problem, i.e., the problem of jointly estimating a partitioning of the domain of the image, and functions defined within regions of the partition. We create algorithms that are efficient, robust to undesirable local minima, and are easy-to-implement. Our algorithms are formulated by slightly modifying the underlying statistical model from which the multi-label Mumford-Shah functional is derived. The advantage of this statistical model is that the underlying variables: the labels and the functions are less coupled than in the original formulation, and the labels can be computed from the functions with more global updates. The resulting algorithms can be tuned to the desired level of locality of the solution: from fully global updates to more local updates. We demonstrate our algorithm on two applications: joint multi-label segmentation and denoising, and joint multi-label motion segmentation and flow estimation. We compare to the state-of-the-art in multi-label Mumford-Shah problems and show that we achieve more promising results.

1. Introduction

Many problems in computer vision are formulated as Mumford-Shah problems, i.e., *joint estimation* problems where two variables must be estimated from the image(s), each variable depending on the other. For example, the image denoising problem can be formulated using the framework of Mumford and Shah [12] (see also [2]): the two variables to be estimated are image edges and the denoised image, the idea being that to denoise the image one must have the edges (so as not to destroy the edges when smoothing), but edges can be detected reliably only when one has the denoised image. The estimate of one variable depends on the other, but both variables are *unknown*, and the estimation should be setup as a *joint* estimation problem where the variables are estimated simultaneously. Some other examples of joint estimation problems are segmentation and denoising [18, 20], 3D reconstruction [9], tomography [14], image registration [19], and optical flow [17].

Despite the more than 25 year history of Mumford-Shah problems, the problem remains a challenge to this day: the high dimensional nature of the problem implies most methods are sensitive to initialization and lock into undesirable estimates, they are computationally expensive, and further the implementation of existing methods are relatively sophisticated, making it difficult for a non-specialist to quickly implement.

In this paper, we address *multi-label* Mumford-Shah problems, in which labels (representing regions in the image) and functions must be estimated simultaneously. We wish to emphasize that our method applies to *any* multi-label Mumford-Shah problem, although our experiments are illustrated on joint segmentation and denoising, and motion segmentation. Our approach leads to a simple implementation, faster convergence, and a better estimation of the underlying variables, when compared to existing methods. Our method is less sensitive to initialization, and many times a random initialization is effective, making full automation a possibility, unlike existing methods which require hand initialization.

Our approach is based on a slightly different underlying statistical model than the original Mumford-Shah problem; nevertheless, our model has the same generative capabilities as the original statistical model. The advantage is that the proposed model leads to a decoupling of the variables to be estimated, and therefore the optimization problem becomes simpler. Further, the modified statistical model leads to better and easier-to-implement optimization algorithms in which the user is able to choose the degree of locality of label updates (from fully global to fully local).

1.1. Related Work

The original Mumford-Shah problem for joint edge detection and denoising was formulated in [12] (see also [2, 7]) and was setup as a variational problem. In the case that the edges in the image are restricted to be the boundaries of regions, forming a partitioning of the domain of the image, the problem was reformulated as a joint segmentation and denoising problem, and level set methods [13] have been employed to implement a gradient descent for

1063-6919/13 $26.00 © 2013 IEEE
DOI 10.1109/CVPR.2013.161

optimizing the functional [4, 20, 18]. Multiple regions are considered in [20] approaches by using logical combinations of multiple level sets. While the methods of [20, 18] were ground breaking in providing a numerical solution, the methods are sensitive to initialization and are usually trapped in undesirable local minima, and the problem is exacerbated with an increasing number of regions.

A more global approach to the two-label piecewise smooth Mumford-Shah functional has been considered in [8] using graph cuts [3] and it is shown to experimentally yield lower energy solutions and converge faster than [18]. Our work is based on the idea in [8] that one can provide better solutions to Mumford-Shah by looking at extensions of the functions into the entire domain of the image. We extend that work by first showing it comes from a different statistical model, and that the algorithm considered comes from an underlying energy that is being optimized. Further, we extend it to multiple labels, extend the algorithm to any desired locality of the label updates, and our algorithm does not rely on graph cuts. In [5], a discrete approach to [20] (multiple phases and a piecewise constant model) is implemented using advanced discrete optimization techniques. However, the case of a non-piecewise constant model was not considered in [5], which is of interest in this work.

Recently, convex optimization techniques have been applied to various problems in computer vision (e.g., [10]), the original Mumford-Shah problem of edge-detection and image denoising is no exception. While the Mumford-Shah functional is not convex, one can consider a relaxation, and optimize the relaxed functional. This approach is considered by [15, 16] where the functional is lifted to a higher dimensional space. The method does not to apply to the case of multi-label piecewise smooth Mumford-Shah where a partitioning of the domain is desired.

2. Mathematical Formulation

The multi-label Mumford-Shah functional is in the form:

$$E(\{R_i, f_i\}_{i=1}^N) = \sum_{i=1}^{N} \int_{R_i} D(I(x), f_i(x))\,\mathrm{d}x + \alpha \mathrm{Reg}(f_i) + \beta \mathrm{Len}(\partial R_i), \quad (1)$$

subject to the constraint that R_i are mutually exclusive and the union is the entire domain Ω of the image. Here $I : \Omega \to \mathbb{R}^k$ is the image (with k-channels), $f_i : R_i \to \mathbb{R}^k$ are approximations of the image in the regions R_i, ∂R_i denotes the boundary of R_i. Discrepancy of the data $I(x)$ to $f_i(x)$ is measured through D (e.g., $D(I(x), f_i(x)) = |I(x) - f_i(x)|^2$ for the original Mumford-Shah problem), $\mathrm{Reg}(\cdot)$ will impose spatial regularity of f_i in the region R_i, one possible choice is

$$\mathrm{Reg}(f_i) = \int_{R_i} |\nabla f_i(x)|^2\,\mathrm{d}x, \quad (2)$$

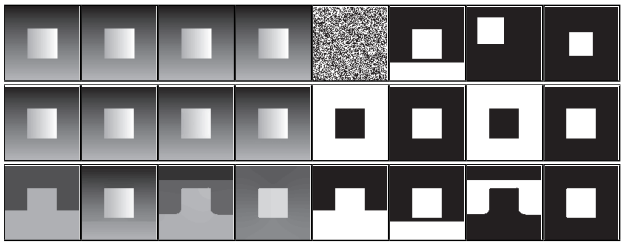

Figure 1. Results of segmentation and reconstruction by our method and the multiphase Mumford-Shah level sets method. Even on the simple two-label case, the classical Mumford-Shah functional has many local minima. [top] same input image (left 4 images) and different initializations (right 4 images). [middle] reconstruction (left 4 images) and segmentation (right 4 images) by the proposed method with the 4 different initializations given above. [bottom] reconstruction (left 4 images) and segmentation (right 4 images) by the multiphase Mumford-Shah level sets method with the 4 different initialization given above. Parameters are picked to obtain reasonable results for multiphase.

and the length penalty $\mathrm{Len}(\cdot)$ is to ensure spatial regularity of the region R_i (since $\mathrm{Len}(\partial R_i) = \int_\Omega |\nabla \mathbf{1}_{R_i}(x)|\,\mathrm{d}x$, where $\mathbf{1}_{R_i}$ indicates the characteristic function of R_i).

2.1. Why is the Mumford-Shah Problem Difficult?

The fact that the Mumford-Shah problem is a high-dimensional (i.e., f_i and R_i are infinite dimensional) and the non-convexity of the space of variables on which the functional is defined implies that the energy has many local minima. For example, Figure 1 shows that even in the case of simple synthetic images, the Mumford-Shah functional has many local minima. The reason that these problems are difficult to optimize is due to the dependence of each of the variables in the optimization on the other. Indeed, to define a function f_i, the region R_i must be known, but the region R_i is *unknown*. Similarly, if the functions f_i are known (i.e., the image reconstruction), then the discontinuities form the boundaries of the regions, and the regions can be defined, however the functions are *unknown*. Hence, there is a constraint that must be maintained between the function and the regions.

2.2. Generative Model

We wish to undue the dependence of the functions on the regions by considering a different underlying model for image formation. We illustrate the idea on the joint image denoising and segmentation problem, but it can be generalized in other Mumford-Shah (MS) problems. The underlying model in the original MS denoising/segmentation problem is

$$I(x) = f_i(x) + \eta(x),\ x \in R_i,\ f_i : R_i \to \mathbb{R}^k \quad (3)$$

$$p(f_i|R_i) \propto \exp\left(-\alpha \int_{R_i} |\nabla f_i(x)|^2\,\mathrm{d}x\right) \quad (4)$$

where η is a noise process, and the probability is the *conditional* prior probability of the function f_i. The generative model implies that the regions R_i must be first chosen according to a prior distribution, and the functions may then be chosen according to a conditional distribution (conditioned on the region). In the proposed formulation, we wish to remove this conditional dependence, while still maintaining the ability to generate all the images that the original model can generate. Our model is the following:

$$I(x) = f_i(x) + \eta(x),\ x \in R_i,\ f_i : \Omega \to \mathbb{R}^k \tag{5}$$

$$p(f_i) \propto \exp\left(-\alpha \int_\Omega |\nabla f_i(x)|^2 \,\mathrm{d}x\right) \tag{6}$$

where we have defined the functions on the *entire* domain Ω (and defined the regularity on Ω) rather than R_i, and therefore, the prior probability of the functions is *no longer conditional* on the regions. Hence to generate the image, the sampling of the regions and functions are done *independently*. The benefit of the proposed model is that in optimization, the regions can be updated more globally than in the original model. The cost is a higher dimensional problem. However, the proposed formulation leads to a better local optimum, and is even faster to converge than the traditional model as we show in experiments.

2.3. Optimization Algorithm

The proposed model leads to the following regularity term in the energy:

$$\mathrm{Reg}(f_i) = \int_\Omega |\nabla f_i(x)|^2 \,\mathrm{d}x \tag{7}$$

rather than (2), and the other terms in the energy (1) remain the same:

$$E(\{R_i, f_i\}_{i=1}^N) = \sum_{i=1}^N \int_{R_i} D(I(x), f_i(x)) \,\mathrm{d}x + \alpha \int_\Omega |\nabla f_i(x)|^2 \,\mathrm{d}x + \beta \mathrm{Len}(\partial R_i) \tag{8}$$

Since the conditional dependence of the function on the region is now removed, the update of the regions can be computed given the functions in a more global fashion, and the optimal update of the functions can be computed given the regions. Convex partitioning (e.g., [10]) techniques can be applied, but we wish to apply a simpler method (with global updates) to ensure easy implementation.

Given estimates of the functions $f_i : \Omega \to \mathbb{R}^k$, we describe how to obtain optimal estimates for the regions R_i. Notice that, unlike the standard model, when optimizing in R_i for a fixed f_i, the second term in (8) can be ignored as it does not depend on the region. For the moment, setting $\beta = 0$, the globally optimal estimate is (given the functions f_i are fixed) is

$$R_i = \{x\ :\ i = \arg\min_j D(I(x), f_j(x))\}. \tag{9}$$

The above optimum does not take into account the length term that implies spatial regularity of the regions, and thus we show how to integrate that next. We first approximate the length term with the following region integral:

$$\int_{R_i} W_{R_i}(x) \,\mathrm{d}x, \text{ where } W_{R_i}(x) = \frac{1}{\sigma}(G_\sigma * \mathbf{1}_{R_i^c})(x), \tag{10}$$

where G_σ indicates a Gaussian smoothing kernel of standard deviation σ, and $\mathbf{1}_{R_i^c}$ is the indicator function on R_i^c. It can be shown from the co-area formula and the Lebesgue differentiation theorem that the integral above converges to the length as $\sigma \to 0$. Therefore, when $\beta > 0$ we seek to optimize

$$\sum_{i=1}^N \int_{R_i} (D(I(x), f_i(x)) + \beta W_{R_i}(x)) \,\mathrm{d}x. \tag{11}$$

Note that we like to apply the same argument as before, and write the optimal solution as a simple thresholding step as in (9); however, part of the integrand depends on the region R_i (when $\beta \neq 0$), which is unknown. Fortunately, if one has an approximation of R_i, R_i', then substituting W_{R_i}, with $W_{R_i'}$, and optimizing (11) with respect to R_i yields

$$R_i = \{x\ :\ i = \arg\min_j D(I(x), f_j(x)) + \beta W_{R_j'}(x)\}. \tag{12}$$

Notice this requires two parameters, β and σ; if desired to avoid an additional parameter, one can get similar results by simply replacing $D(I(x), f_j(x)) + \beta W_{R_j'}(x)$ with $G_\sigma * D(I(\cdot), f_j(\cdot))$, i.e., enforce spatial regularity by smoothing the data term. As an initial approximation, we choose R_i' to be (9). The equation (12) can then be iterated with the previous estimate replacing R_i'. In practice, the iteration is done only once, which yields an accurate segmentation, as we show. As $\sigma \to 0$ and $D = 0$, the iterative scheme was considered in [11] and approximates mean curvature flow as shown in [1, 6], and thus minimizes the length of the boundary of the regions. In our numerical scheme, we consider a fixed σ and only one iteration, the effects of a large number of iterations are attained by a larger σ.

Also, given that regions R_i are fixed, we can compute a global minimum for the f_i (in the case of a convex D). For the case that $D(x, y) = |x - y|^2$, we have that

$$-\alpha \Delta f_i(x) = (I(x) - f_i(x)) \mathbf{1}_{R_i}(x). \tag{13}$$

We now have the following simple iterative algorithm to obtain an approximate solution to the Mumford-Shah problem:

1. Initialize $\{R_i\}_{i=1}^N$
2. Solve for the f_i :

$$\begin{cases} -\alpha\Delta f_i(x) = (I(x) - f_i(x))\mathbf{1}_{R_i}(x) & x \in \text{int}(\Omega) \\ \frac{\partial f_i}{\partial \mathcal{N}} = 0 & x \in \partial\Omega \end{cases}$$

where $\mathcal{N}$ denotes the normal to the boundary (Neumann boundary conditions), and $\text{int}(\cdot)$ denotes interior. The equations can be solved efficiently with a conjugate gradient solver.

3. Compute R_i':

$$R_i' = \{x \in \Omega \,:\, i = \arg\min_j |I(x) - f_j(x)|^2\},$$

4. Update R_i:

$$R_i = \{x \in \Omega \,:\, i = \arg\min_j F_j(x)\},$$

where $F_j(x) = |I(x) - f_j(x)|^2 + \beta W_{R_j'}(x)$.

5. Repeat 2-4 until R_i converges

3. From Global to More Local Updates

In the previous section, we have derived a more *global* optimization algorithm than traditional local approaches for optimizing Mumford-Shah-type problems. Our approach is not a fully global method, however, each update of the regions approximates a global optimum of the energy conditioned on the functions, and each update of the functions is a global optimizer conditioned on the regions. In many cases, a global approach yields undesirable estimates that capture undesired clutter (see Fig. 9). A local approach is less affected by clutter if the initialization is close to the desired solution, but is susceptible to local features (e.g., noise). In this section, we show that the algorithm derived in the previous section can be generalized to an algorithm that is a trade-off between the global updates and local updates of traditional methods, yielding advantages of both.

Suppose first that one has an initial estimate of the labels to be determined. Our global update of the regions given the current functions implies that a pixel far from a region can change to the same label as the far region. In contrast, in local approaches, only pixels in one region that border another region may switch to the bordering region. In order to optimize the underlying energy in a way that is not too global nor too local, we consider an algorithm in which a pixel of a region may switch to another region only if the pixel is sufficiently close to the region, the possibility of a switch decreasing with distance to the candidate region. The simple intuition above can be formalized mathematically as follows. Define a band of a region R_i as follows:

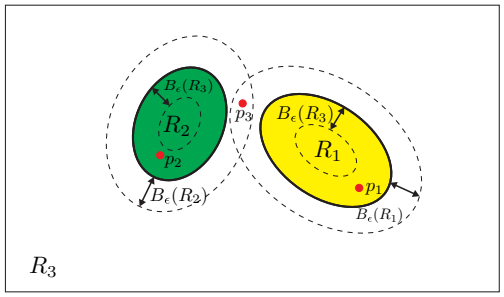

Figure 2. This figure illustrates quantities involved in the global-local tradeoff algorithm for optimizing the multi-label Mumford-Shah-type problem. The regions R_i with different labels are represented in solid colors and bands $B_\varepsilon(R_i)$ are represented by the dashed lines. p_i is allowed to transition to R_j if $p_i \in B_\varepsilon(R_j)$ (p_1 can to transition to R_3, p_2 to R_3, and p_3 to R_1 or R_2).

$B_\varepsilon(R_i) = \{x \in \Omega \,:\, 0 < d_{R_i}(x) < \varepsilon\}$ where d_{R_i} is the Euclidean distance map of R_i, and then define the inverse prior likelihood of pixel switching from region R_i, $s_i : \Omega \to \mathbb{R}^+$, as any non-decreasing function of the distance map that becomes infinite outside the band; the simple example we choose is

$$s_i(x) = \begin{cases} 1 & x \in B_\varepsilon(R_i) \cup R_i \\ +\infty & x \in \Omega \backslash (B_\varepsilon(R_i) \cup R_i) \end{cases}. \quad (14)$$

For the moment, as in the previous section, we ignore the effects of the length regularization in the energy ($\beta = 0$), then given the functions f_i, an update of the regions that decreases the energy is

$$R_i' = \{x \in \Omega : i = \arg\min_j s_j(x) D(I(x), f_j(x))\}. \quad (15)$$

Notice that the update does not increase the energy (when $\beta = 0$) and the functions are fixed: suppose that a point $x \in \Omega$ switches labels from i to j, then it is clear that $D(I(x), f_j(x)) \leq s_j(x) D(I(x), f_j(x)) < s_i(x) D(I(x), f_i(x)) = D(I(x), f_i(x))$, hence the contribution of x to the energy is less than or equal before the switch so the energy is not increased. If there is no switch, then the contribution of x to the energy remains the same, not increasing the energy. As $s_i(x) = +\infty$ for x that is sufficiently far from the region R_i, it is clear that x cannot switch to label i, thus more global transitions in the algorithm of the previous section are avoided. See Figure 2 for an illustration of the $B_\varepsilon(R_i)$ and transitions possible. To induce the effects of regularity (when $\beta \neq 0$), the update derived in the previous section is applied:

$$R_i = \{x \,:\, i = \arg\min_j s_j(x) \left[D(I(x), f_j(x)) + \beta W_{R_j'}(x)\right]\}. \quad (16)$$

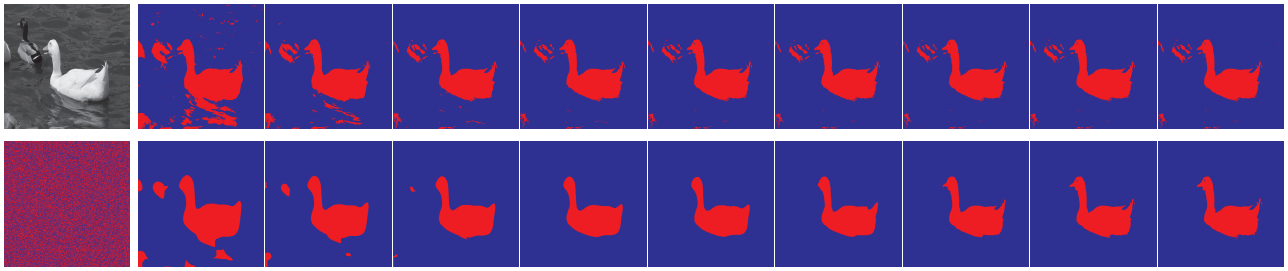

Figure 3. This figures shows intermediate steps in the coarse-to-fine algorithm for achieving an estimate in multi-label Mumford-Shah-type problems. [top-left] a given image to segment. [bottom-left] random initialization. [top] segmentation results by multiphase level sets method [4]. [bottom] segmentation results by our proposed method.

The functions f_i are now updated using the exact same scheme as in the previous section (13), although f_i need only be solved in $B_\varepsilon(R_i)$, which leads to a considerable speed increase.

It is clear that as $\varepsilon \to +\infty$, the algorithm defaults to the algorithm that yields a global update. In the case that $\varepsilon \to 0$, the possible points that can be updated are located only on the boundaries of the regions, where they can switch to an adjacent region. Therefore, as $\varepsilon \to 0$, the method defaults to region competition [21].

We now address the issue of initialization in the derived global-local tradeoff algorithm. We propose to use the above algorithm in a coarse-to-fine scheme. The parameter σ is a scale parameter that controls the coarse-ness of the estimate achieved: larger σ produces a coarse solution - with details smoothed out, and smaller σ produces an estimate that is fine - and is effected by fine scale structures in the data. Our final estimation scheme is then:

1. Choose a random label for each $x \in \Omega$.
2. Choose $\varepsilon = +\infty$ and σ large, and run the algorithm derived in this section.
3. Decrease σ and ε, and run the algorithm again.
4. Repeat Step 3 until the regions have converged.

Figure 3 shows intermediate steps of the iterative procedure of the global-local tradeoff algorithm. Details of the particular steps of the algorithm (e.g., how to decrease σ) is given in the next section.

4. Experiments

In the first experiment, we test the performance of the proposed global algorithm described in Section 2.3 for solving Mumford-Shah-type problems. We start with a case where a fully non-parametric model for the functions f_i are needed: segmentation of brain MRI into various substructures in the presence of a bias field. Tilts of the patient's head in the MR machine causes gradients in the image, which makes sub-parts of the brain hard to recover without smooth, non-constant functions f_i. The data is obtained from BrainWeb, which has ground truth segmentations. We test our algorithm on various slices with varying degrees of bias, from 0 to 100% bias. We compare our algorithm to the current state-of-the-art in multi-label Mumford-Shah (segmentation and denoising) - multiphase Mumford-Shah (MPMS) using level sets [20]. In order to get a comprehensive picture of the performance of the algorithms, we test the algorithms over a large number of parameters that are for the region smoothness and the function regularization on different initializations. We use the same number of parameters for the common parameters used in all the methods. The number of labels is also given to the algorithms and it ranges from 4 to 8. The initializations chosen are 5 different random initializations (each pixel is assigned a random number between 1 and the number of labels), and another initialization where we simply quantize the image into the number of labels.

We wish to test *both* the accuracy of the new statistical model and the optimization scheme, and therefore, it is appropriate to measure the results with respect to closeness to the ground truth. Therefore, we measure accuracy of results with respect to ground truth in terms of the F-measure. The brain image contains 12 different labels, but the prominent labels that are used for clinical purposes are the gray and white matter, and the other labels are hard to capture without additional training data, which we avoid, since we would like to test performance of our new statistical formulation of Mumford-Shah-type problems. Thus, we measure accuracy only in terms of the gray and white matter. Results of this experiment are summarized in Figures 4,5,6. The optimal (in terms of F-measure) segmentation and reconstruction are presented for each method over all the parameters in Figure 4. In order to show that the full-non parametric functions f_i are needed, we have included results for the case of constant functions based on the Chan-Vese model optimized using our algorithm (we refer to this as "ours CV"). The results indicate that multiphase Mumford-Shah (MPMS) level sets are extremely sensitive to initialization, yielding very inaccurate solutions for ran-

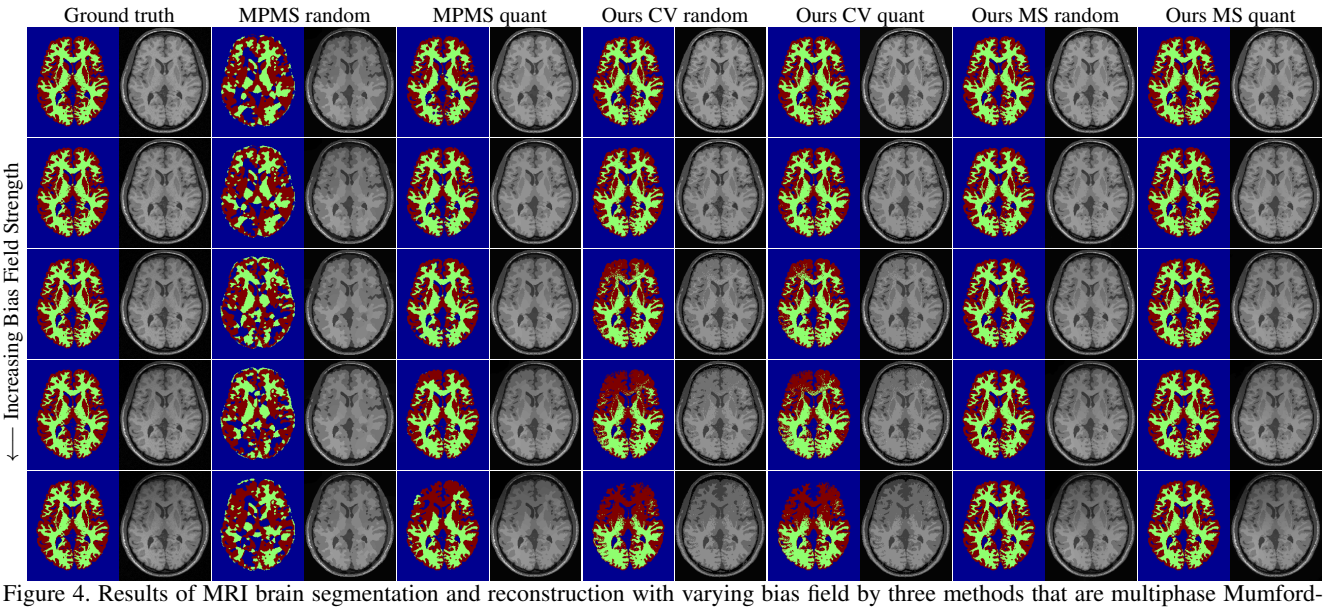

Figure 4. Results of MRI brain segmentation and reconstruction with varying bias field by three methods that are multiphase Mumford-Shah level sets, ours Chan-Vese, and ours Mumford-Shah using two different initializations that are random and equally quantized images. In each block of 5×2 images, segmentation is shown on the left and reconstruction on the right for the images with increasing level of bias field from top to bottom. From left to right, the image blocks show the ground truth, MPMS with random and equally quantized initialization, ours CV with random and equally quantized initialization, ours MS with random and equally quantized initialization. The results show high sensitivity of multiphase methods to initialization and local minima, while the proposed approach consistently and accurately recovers the desired segmentation.

dom initializations to fairly accurate solutions for the quantized initialization. The quality of MPMS degrades with increasing bias. The assumption of constant functions using our approach yields highly accurate results for small bias regardless of initialization, but quickly degrades with increasing bias. Our global approach to fully non-parametric Mumford-Shah yields the most accurate results and is stable to initialization and bias. Figure 5 summarizes the results over parameters measured in terms of F-measure. The graph shows that our method is the most accurate and remains accurate even with varying initialization. In fact, the median F-measure (over all parameters) is highest for the proposed method. The performance of MPMS is significantly worse than the proposed method, and the results vary drastically with respect to parameters.

The last graph of the first experiment (Figure 6) shows the algorithm speed of the proposed method (piecewise smooth) versus MPMS. Even though the functions must be computed on the whole domain using our approach, it is faster than MPMS in which the function computation is restricted to regions, even for a significantly large number of labels (8). Since, it is not true that the results are the same for both MPMS and the proposed approach, we compared times in the case of zero bias and quantized initialization, where both methods yield similar results, and therefore, it makes sense to compare speed.

In the second experiment we demonstrate the effect of the global-local tradeoff scheme in a multi-label segmentation problem. The proposed multi-label segmentation algorithms with both the global optimization scheme and the global-local tradeoff optimization scheme are performed on a number of natural scenery images. Since the objective of this experiment is to highlight the difference between the segmentations resulted by the global-local tradeoff scheme and the global scheme of our algorithm, a simple piecewise constant image model is assumed. The proposed global algorithm takes the region smoothness parameter σ that is related to the feature scale of object and the number of labels N. In this experiment, we set $N = 4$ for all the test images. The global-local tradeoff algorithm is described in Section 3 and σ is reduced by a factor of 2 and the band size is $\varepsilon = \text{ceil}(2\sigma)$ for Step 3 of the algorithm described in Section 3. The final σ is chosen to be the same as the given σ for the global-update algorithm (i.e., without band). Results of the two methods are presented in Figure 7 where the images to segment are shown on the top row, the segmentation results by the global algorithm on the middle row, and the segmentation results by the global-local tradeoff algorithm on the bottom row. Notice that the results with the global-local approach preserve the same level of detail as the ones with the global approach while discarding irrelevant clutter. For some cases it is observed that the global-local tradeoff approach yields less number of final labels than the result labels with the purely global approach, which happens when

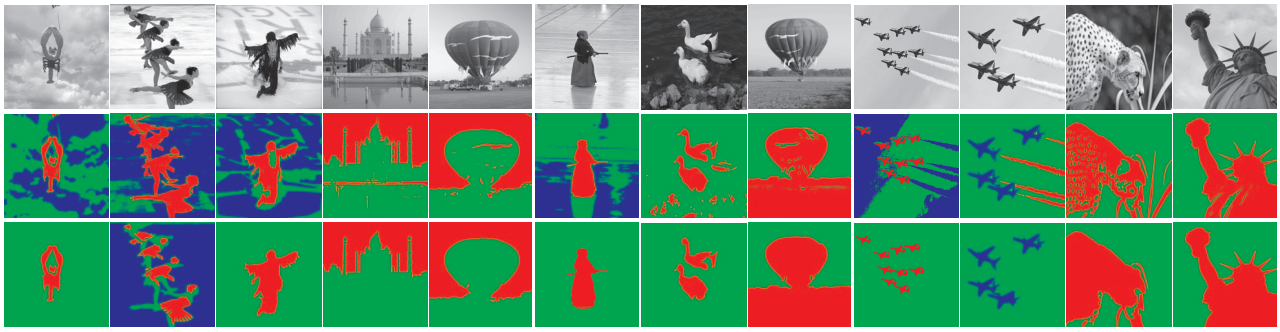

Figure 7. Results of global-local tradeoff algorithm. The multi-label segmentation is performed on a number of natural scenery images (top) by our proposed method with the global optimization scheme (middle) and the global-local tradeoff scheme (bottom).

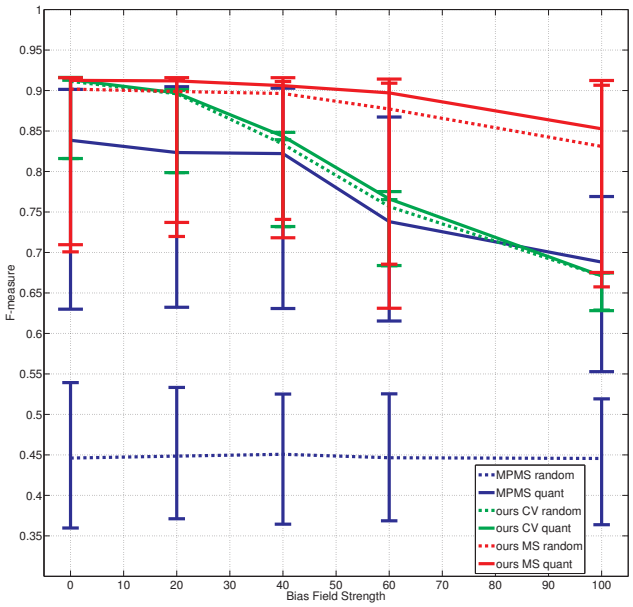

Figure 5. This figure shows the optimal results (with respect to F-measure) of segmentation of images with varying bias field with multiphase Mumford-Shah (MPMS) level sets, the proposed method with constant functions (ours CV), and the proposed method with smooth functions (ours MS). For each method, random and equally quantized (quant) images are used for initialization. The bars indicate minimum and maximum F-measure, and the left-right lines are drawn through median values. The proposed method has high F-measure regardless of the initialization and the intensity of the bias field. MPMS perform poorly with under a bias field and varies considerably with bias field intensity.

the labels corresponding to insignificant clutter disappear at a coarse scale in the coarse-to-fine optimization procedure.

In the last experiment, we demonstrate our method on the task of motion segmentation where it is desired to segment the image into regions of smooth velocities. The energy is $E(v_i, R_i) = \sum_i \int_{R_i} |I_1(x + v_i(x)) - I_2(x)|^2 \, \mathrm{d}x + \int_{R_i} |\nabla v_i(x)|^2 \, \mathrm{d}x + \mathrm{Len}(\partial R_i)$, where I_1 and I_2 are images. We apply a coarse-to-fine scheme for fully global, fully local, and the global-local tradeoff algorithm (e.g., large σ until convergence, and then σ is lowered, etc...). All schemes have the same maximum and minimum σ. While fully global updates result in capturing clutter, and fully local captures irrelevant fine structures, the global/local tradeoff approach achieves an accurate solution. Our method also converges much quicker than both a fully global and fully local approach. The initialization is simply a quantization of the optical flow field using a standard optical flow method. The band size and the σ are chosen with the same scheme described in the previous paragraph. Figure 8 shows some steps in the evolution, and the current region and the reconstructed velocity are shown at each step. Results for the final motion and segmentation on more data is given in Figure 9.

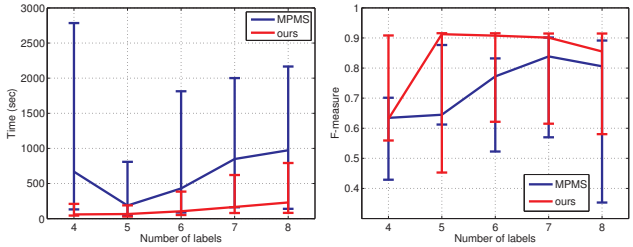

Figure 6. The left plot shows that the proposed method converges faster than multiphase Mumford-Shah, even for a large number of labels. The right plot shows the segmentation accuracy of MPMS is less than the proposed method, suggesting that MPMS converged at an undesirable local minima instead of a more distant and better solution. Even when MPMS quickly converges to a nearby undesirable local minimum, the proposed method is faster, converging to a better solution. Note that bars indicate max/min values and left-right lines are drawn through median values.

5. Conclusion

We have proposed a new statistical model for multi-label Mumford-Shah problems, slightly modify the standard one. This leads to new algorithms where the function updates can be done in many ways: from fully global updates to fully local updates (the user is free to choose). The algorithms are

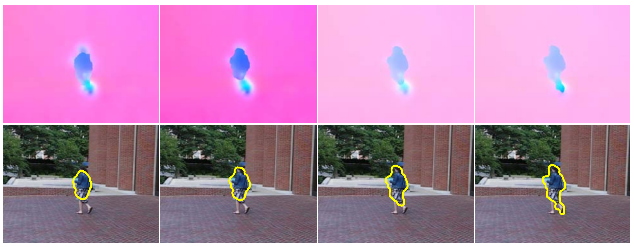

Figure 8. Intermediate steps in the evolution of the velocity (top) and the segmenting curve (bottom) for the motion segmentation based on the global-local tradeoff algorithm.

Figure 9. Results of global-local algorithm (third column) in motion segmentation in comparison to fully global (left column), and fully local algorithms (second column). The fourth column is the optical flow field determined by the global-local algorithm (using the standard color scheme). The global-local tradeoff algorithm avoids clutter, and is less susceptible to irrelevant fine detail.

easy to implement for a non-specialist, and yield results less susceptible to initialization and local irrelevant details than traditional approaches. A possible coarse-to-fine algorithm was derived that lends itself well to full automation. Experiments suggest that the proposed algorithm has increased the performance of traditional multi-label Mumford-Shah problems both in terms of accuracy and speed.

Acknowledgement

B. W. H was supported by Korea NRF-2010-220-D00078.

References

[1] G. Barles and C. Georgelin. A simple proof of convergence for an approximation scheme for computing motions by mean curvature. *SIAM Journal on Numerical Analysis*, pages 484–500, 1995. 3

[2] A. Blake and A. Zisserman. *Visual reconstruction*. 1987. 1

[3] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *PAMI, IEEE Transactions on*, 23(11), 2001. 2

[4] T. Chan and L. Vese. Active contours without edges. *Image Processing, IEEE Transactions on*, 10(2), 2001. 2, 5

[5] N. El-Zehiry and L. Grady. Discrete optimization of the multiphase piecewise constant mumford-shah functional. In *Energy Minimization Methods in CVPR*. Springer, 2011. 2

[6] L. Evans. Convergence of an algorithm for mean curvature motion. *Indiana Mathematics Journal*, 42(2), 1993. 3

[7] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *PAMI, IEEE Transactions on*, (6), 1984. 1

[8] L. Grady and C. Alvino. Reformulating and optimizing the mumford-shah functional on a grapha faster, lower energy solution. *ECCV*, 2008. 2

[9] H. Jin, S. Soatto, and A. Yezzi. Multi-view stereo reconstruction of dense shape and complex appearance. *International Journal of Computer Vision*, 63(3):175–189, 2005. 1

[10] J. Lellmann and C. Schnörr. Continuous multiclass labeling approaches and algorithms. *SIAM Journal on Imaging Sciences*, 4(4):1049–1096, 2011. 2, 3

[11] B. Merriman, J. Bence, and S. Osher. Diffusion generated motion by mean curvature. *UCLA CAM Report*, 1992. 3

[12] D. Mumford and J. Shah. Boundary detection by minimizing functionals. In *CVPR*, 1985. 1

[13] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *J. of Computational Physics*, 79(1), 1988. 1

[14] J. O'Sullivan and J. Benac. Alternating minimization algorithms for transmission tomography. *Medical Imaging, IEEE Transactions on*, 26(3):283–297, 2007. 1

[15] T. Pock, D. Cremers, H. Bischof, and A. Chambolle. An algorithm for minimizing the mumford-shah functional. In *IEEE Conference on ICCV*, 2009. 2

[16] E. Strekalovskiy, A. Chambolle, and D. Cremers. A convex representation for the vectorial mumford-shah functional. In *IEEE Conference on CVPR*, 2012. 2

[17] D. Sun, E. Sudderth, and M. Black. Layered segmentation and optical flow estimation over time. In *CVPR*, 2012. 1

[18] A. Tsai, A. Yezzi Jr, and A. Willsky. Curve evolution implementation of the mumford-shah functional for image segmentation, denoising, interpolation, and magnification. *Image Processing, IEEE Trans. on*, 10(8), 2001. 1, 2

[19] G. Unal and G. Slabaugh. Coupled pdes for non-rigid registration and segmentation. In *CVPR*, 2005. 1

[20] L. Vese and T. Chan. A multiphase level set framework for image segmentation using the mumford and shah model. *IJCV*, 50(3), 2002. 1, 2, 5

[21] S. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *PAMT, IEEE Trans. on*, 18(9), 1996. 5