

Multi-class Transductive Learning Based on ℓ^1 Relaxations of Cheeger Cut and Mumford-Shah-Potts Model

Xavier Bresson · Xue-Cheng Tai · Tony F. Chan · Arthur Szlam

Published online: 27 August 2013
© Springer Science+Business Media New York 2013

Abstract Recent advances in ℓ^1 optimization for imaging problems provide promising tools to solve the fundamental high-dimensional data classification in machine learning. In this paper, we extend the main result of Szlam and Bresson (Proceedings of the 27th International Conference on Machine Learning, pp. 1039–1046, 2010), which introduced an exact ℓ^1 relaxation of the Cheeger ratio cut problem for unsupervised data classification. The proposed extension deals with the multi-class transductive learning problem, which consists in learning several classes with a set of labels for each class. Learning several classes (i.e. more than two classes) simultaneously is generally a challenging problem, but the proposed method builds on strong results introduced in imaging to overcome the multi-class issue. Besides, the proposed multi-class transductive learning algorithms also benefit from recent fast ℓ^1 solvers, specifically designed for the total variation norm, which plays a central role in our approach. Finally, experiments demonstrate that the proposed

ℓ^1 relaxation algorithms are more accurate and robust than standard ℓ^2 relaxation methods s.a. spectral clustering, particularly when considering a very small number of labels for each class to be classified. For instance, the mean error of classification for the benchmark MNIST dataset of 60,000 data in \mathbb{R}^{784} using the proposed ℓ^1 relaxation of the multi-class Cheeger cut is 2.4 % when only one label is considered for each class, while the error of classification for the ℓ^2 relaxation method of spectral clustering is 24.7 %.

Keywords Data analysis · Clustering · Transductive learning · Multi-class · Ratio cut · Min cut · Potts and Mumford-Shah energies · Exact relaxation · Fast L1 optimization · Total variation

1 Introduction

Partitioning data into sensible groups is a fundamental problem in machine learning and science in general. One of the most popular approaches is to find the best (balanced) cut of a graph representing data, such as the normalized cut of Shi and Malik [24] or the Cheeger ratio cut [9]. However, solving balanced/ratio cut problems is NP-hard, which has lead people to compute approximate solutions. The most well-known approach to approximate the solution of a ratio cut is the spectral clustering method, which is based on a ℓ^2 relaxation of the original ratio cut. This ℓ^2 relaxation reduces to solving a generalized system of eigenvectors for the graph Laplacian, then selects the 2nd smallest eigenvector and finally partitions into two groups by thresholding (this requires testing multiple thresholds). Different normalizations of the graph Laplacian lead to different spectral clustering methods. These methods often provide good solutions but can fail on somewhat benign problems; for example see

X. Bresson (✉)
Department of Computer Science, City University of Hong Kong,
Kowloon Tang, Hong Kong
e-mail: xbresson@cityu.edu.hk

X.-C. Tai
Department of Mathematics, University of Bergen, Bergen,
Norway
e-mail: xue-cheng.tai@uib.no

T.F. Chan
Department of Mathematics and Computer Science, Hong Kong
University of Science and Technology, Kowloon, Hong Kong
e-mail: tonyfchan@ust.hk

A. Szlam
Department of Mathematics, The City College of New York, New
York, NY, USA
e-mail: aszlam@courant.nyu.edu

the two-moons example in Fig. 1. In this case, the relaxation leading to the spectral clustering methods is too weak. A stronger relaxation was introduced by Bühler and Hein in [7]. They described the p -spectral clustering method, which considers the ℓ^p relaxation of the Cheeger ratio cut, instead of the ℓ^2 relaxation. They showed that the relaxed solution of the p -spectral clustering problem tends asymptotically to the solution of the Cheeger cut problem when $p \rightarrow 1$. In [10, 26] (also see [25]), it was proved that the relaxation for $p = 1$ is actually exact, i.e. the solution of the ℓ^1 relaxation problem provides an exact solution of the Cheeger cut problem. Unfortunately, there is no algorithm that guarantees to find global minimizers of the ℓ^1 relaxation problem (we recall that the problem is NP-hard). However, the experiments in [7, 26] showed that good results can be obtained with these stronger relaxations; the works [3, 15, 16] have further strengthened the case for ℓ^1 relaxation methods and related ideas, and have charted a new and promising research direction for improving spectral clustering methods.

In this work, we propose to extend [26]. In particular, we are interested in extending to the challenging multi-class ratio cut problem, and adding label information to obtain a transductive problem. Standard approaches for the unsupervised learning problem usually proceed by recursive two-class clustering. In this paper, we will use results recently introduced in imaging science to solve the multi-class learning problem. The papers [1, 6, 8, 19, 20, 28] have proposed tight approximations of the solution of the multi-phase image segmentation problem based on ℓ^1 relaxation techniques. The main contribution of this paper is to develop efficient multi-class algorithms for the transductive learning problem. We will introduce two multi-class algorithms based on the ℓ^1 relaxation of the Cheeger cut and the piecewise constant Mumford/Shah or Potts models [22, 23]. Experiments show that these new multi-class transductive learning algorithms improve the classification results compared to spectral clustering algorithms, particularly in the case of a very few numbers of labels.

2 Unsupervised Data Classification with ℓ^1 Relaxation of the Cheeger Cut

2.1 The Model

In this section, we recall the main result of [26] and proposed a modified and improved version of the algorithm introduced there. Let $G = (V, E)$ be a graph where V is the set of nodes and E is the set of edges weighted by a function W_{ij} , $\forall (ij) \in E$. A classical method for clustering is to consider the Cheeger minimization problem [9]:

$$\min_{\Omega \subset V} \frac{\text{Cut}(\Omega, \Omega^c)}{\min(|\Omega|, |\Omega^c|)} \quad (1)$$

which partitions the set V of points into two sets Ω and Ω^c (the complementary set of Ω in V). The cut is defined as $\text{Cut}(\Omega, \Omega^c) := \sum_{i \in \Omega, j \in \Omega^c} w_{ij}$ and $|\cdot|$ provides the number of points in a given set. The Cheeger problem is NP-hard. However, it was shown in [10], and by the authors of this paper using a different argument in [26], that there exists an exact continuous relaxation of (1) as follows. Let us consider the minimization problem w.r.t. a function $u : V \rightarrow [0, 1]$:

$$\min_{u \in [0, 1]} \frac{\|Du\|_1}{\|u - m(u)\|_1} \quad (2)$$

where $\|Du\|_1 := \sum_{ij} w_{ij} |u_i - u_j|$ is the graph-based total variation of the function u , $m(u)$ is the median of u , and $\|u - m(u)\|_1 = \sum_i |u_i - m(u)|$. If a global minimizer u^* of (2) can be computed, then it can be shown that this minimizer would be the indicator of a set Ω^* (i.e. $u^* = 1_{\Omega^*}$) corresponding to a solution of the NP-hard problem (1). But there is no algorithm that guarantees to compute global minimizers of (2) as the problem is non-convex. However, experiments show that the proposed minimization algorithm in [26], which we will review below, produces good approximations of the solution.

Recent advances in ℓ^1 optimization offer powerful tools to design a fast and accurate algorithm to solve the minimization problem (2). First, observe that minimizing (2) is equivalent to:

$$\min_{u \in [0, 1]} \frac{\|Du\|_1}{\|u\|_1} \quad \text{s.t.} \quad m(u) = 0 \quad (3)$$

Indeed, the energy is not changed if a constant is added to u . So it is possible to restrict the minimization problem to functions u with zero median. Then, the ratio minimization problem (3) can be solved using the method of Dinkelbach [11] (also used in imaging problems s.a. [17, 18]) which introduces the minimax problem:

$$\min_{u \in [0, 1]} \max_{\lambda \in \mathbb{R}} \|Du\|_1 - \lambda \|u\|_1 \quad \text{s.t.} \quad m(u) = 0 \quad (4)$$

Then, we consider a standard two-step iterative algorithm:

- (i) Fix λ , compute the solution of the constrained minimization problem:

$$u^{n+1} = \operatorname{argmin}_{u \in [0, 1]} \|Du\|_1 - \lambda^n \|u\|_1 \quad \text{s.t.} \quad m(u) = 0 \quad (5)$$

- (ii) Fix u , compute the solution of the maximization problem:

$$\lambda^{n+1} = \operatorname{argmax}_{\lambda \in \mathbb{R}} \|Du^{n+1}\|_1 - \lambda \|u^{n+1}\|_1 \quad (6)$$

For the minimization problem (5), observe that the constraint zero median is not linear, but it can be replaced by the approximate linear constraint $\sum_i u_i \leq |V|/2$. Indeed,

suppose that $u_i \in \{0, 1\}$ then the median of u is zero if $\sum_i u_i \leq \sum_i (1 - u_i)$ which yields to $\sum_i u_i \leq |V|/2$. We will consider the notation $\mathbf{1} \cdot u := \sum_i u_i$ in the rest of the paper.

In order to deal efficiently with the non-differentiability of the ℓ^1 norm in (6), a splitting approach associated with an augmented Lagrangian method and the Alternating Direction Method of Multipliers [13] can be used along the same lines as [4, 14]. Hence, we consider the constrained minimization problem:

$$\begin{aligned} \min_{u,v \in [0,1], d} \quad & \|d\|_1 - \lambda \|v\|_1 \\ \text{s.t.} \quad & d = Du, \quad v = u, \quad \mathbf{1} \cdot v \leq |V|/2 \end{aligned} \tag{7}$$

whose linear constraints can be enforced with an augmented Lagrangian method as:

$$\left\{ \begin{aligned} & (u^{n+1}, v^{n+1}, d^{n+1}) \\ & = \operatorname{argmin}_{u,v \in [0,1], d} \|d\|_1 - \lambda \|v\|_1 \\ & \quad + \alpha_d \cdot (d - Du) + \frac{r_d}{2} |d - Du|^2 \\ & \quad + \alpha_v \cdot (v - u) + \frac{r_v}{2} (v - u)^2 + \alpha_m \cdot (\mathbf{1} \cdot v - |V|/2) \tag{8} \\ & \alpha_d^{n+1} = \alpha_d^n + r_d \cdot (d^{n+1} - Du^{n+1}) \\ & \alpha_v^{n+1} = \alpha_v^n + r_v \cdot (v^{n+1} - u^{n+1}) \\ & \alpha_m^{n+1} = \max(0, \alpha_m^n + r_m \cdot (\mathbf{1} \cdot v^{n+1} - |V|/2)) \end{aligned} \right.$$

Three sub-minimizations need to be solved. The minimization problem w.r.t. u :

$$\min_u \frac{r_d}{2} \left| Du - \left(d + \frac{\alpha_d}{r_d} \right) \right|^2 + \frac{r_v}{2} \left(u - \left(v + \frac{\alpha_v}{r_v} \right) \right)^2$$

whose solution u^* is given by a Poisson problem:

$$(r_v + r_d D^T D)u = r_d D^T \left(d + \frac{\alpha_d}{r_d} \right) + r_v \left(v + \frac{\alpha_v}{r_v} \right) \tag{9}$$

The solution of (9) can be estimated by a few steps of conjugate gradient descent as D is extremely sparse. The minimization problem w.r.t. v :

$$\min_{v \in [0,1]} -\lambda \|v\|_1 + \frac{r_v}{2} \left(v - \left(u - \frac{\alpha_v}{r_v} \right) \right)^2 + \alpha_m \cdot (\mathbf{1} \cdot v - |V|/2)$$

has an analytical solution given by unshrinkage [26] and truncated into $[0, 1]$:

$$\begin{aligned} v^* &= \Pi_{[0,1]} \left(f_v + \frac{\lambda}{r_v} \frac{f_v}{|f_v|} \right), \\ \text{with } f_v &:= u - \frac{\alpha_v}{r_v} - \frac{\alpha_m}{r_v} \end{aligned} \tag{10}$$

To avoid the constant trivial solution, we also apply the “renormalization” step: $v^* \leftarrow \frac{v^* - \min(v^*)}{\max(v^*) - \min(v^*)}$. The minimization problem w.r.t. d :

$$\min_d \|d\|_1 + \frac{r_d}{2} \left| d - \left(Du - \frac{\alpha_d}{r_d} \right) \right|^2$$

has also an analytical solution given by shrinkage [12]:

$$\begin{aligned} d^* &= \max \left(|f_d| - \frac{1}{r_d}, 0 \right) \frac{f_d}{|f_d|}, \\ \text{with } f_d &:= Du - \frac{\alpha_d}{r_d} \end{aligned} \tag{11}$$

For the maximization problem (6), the solution is as follows:

$$\lambda^{n+1} = \frac{\|Du^{n+1}\|_1}{\|u^{n+1}\|_1} \tag{12}$$

We will consider a steepest gradient descent method instead of (12) to get a smoother evolution of λ^{n+1} :

$$\lambda^{n+1} = \lambda^n - \delta_\lambda \cdot \left(\lambda^n - \frac{\|Du^{n+1}\|_1}{\|u^{n+1}\|_1} \right) \tag{13}$$

To summarize the algorithm introduced in this section, we write down the pseudo-code Algorithm 1.

Algorithm 1 Unsupervised learning with ℓ^1 relaxation of the Cheeger cut

$u^{n=0}$ given by random initialization

while outer loop not converged **do**

$\alpha_d^{q=0}, \alpha_v^{q=0}, \alpha_m^{q=0} \leftarrow 0$

while inner loop not converged **do**

$u^{n+1,q+1}$ given by (9)

$v^{n+1,q+1}$ given by (10)

$d^{n+1,q+1}$ given by (11)

α_v^{q+1} given by (8)

α_d^{q+1} given by (8)

α_m^{q+1} given by (8)

end while

λ^{n+1} given by (13)

end while

2.2 Experiments

In this section, we demonstrate results using the unsupervised classification Algorithm 1. For each experience, we build the weight matrix using the self-tuning construction of [29]. We use ten nearest neighbors, and the tenth neighbor determines the local scale. The universal scaling parameter is set to 1. For Algorithm 1, we set $r_d = 10$, $r_v = 100$, $r_m = 6K/N$, where N is the number of data points and K is the number of classes, and $\delta_\lambda = 0.4$. Figure 1 presents the well-known two-moon dataset [7]. Each moon has 1,000 data points in \mathbb{R}^{100} . This example shows that the solution of the ℓ^1 relaxation is tighter than the solution of the ℓ^2 relaxation (see caption for more details). In Table 1, we compare quantitatively our algorithm with the

Table 1 Unsupervised learning result for the two-moon dataset. *Column* (a) reports the *minimum* energy value (1) considered among 100 random initializations and the error is the percentage of misclassified data of the minimum energy. *Column* (b) reports the *average* energy

value (1) and the misclassification error for 100 random initializations. Note that the same random initialization was used for Algorithms 1 and [15]

	(a) Best energy (error)	(b) Mean energy (error)
Algorithm 1	0.4688 (1.25 %)	0.4693 (1.53 %)
Hein and Bühler [15]	0.4688 (1.25 %)	0.4700 (1.61 %)
Spectral clustering [24]	0.4705 (1.75 %)	0.4705 (1.75 %)

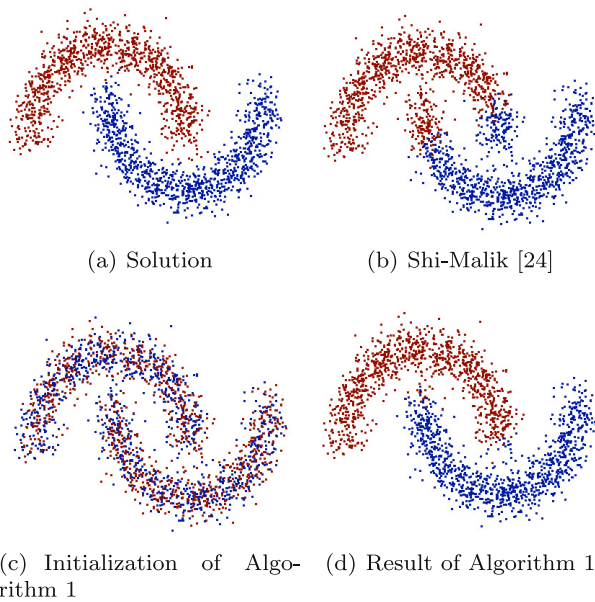


Fig. 1 Unsupervised classification of the two-moon dataset. Each moon has 1,000 data points in \mathbb{R}^{100} . Figure (b) is the result given by the spectral clustering method of Shi and Malik [24]. It fails to produce the correct result as the ℓ^2 relaxation is too weak. Figure (d) is the result of the ℓ^1 relaxation algorithm and figure (c) is the random initialization. The proposed algorithm succeeds to compute the correct result. This also shows that the solution of the ℓ^1 relaxation is tighter than the solution of the ℓ^2 relaxation (Note: it is a color figure)

spectral clustering method of Shi and Malik [24] and the related method of Hein and Bühler in [15], which is available at <http://www.ml.uni-saarland.de/code/oneSpectralClustering/oneSpectralClustering.html> ([16] is not yet available for comparison). Our method and [15] outperform the spectral clustering method.

In Fig. 2, we apply the standard recursive two-class partitioning approach to deal with more than two classes. Figure 2(b) shows the result by spectral clustering and Fig. 2(c) presents the result with our algorithm (see caption for more details).

On the right hand side of Fig. 3, we display a projection of the MNIST benchmark dataset, available at <http://yann.lecun.com/exdb/mnist/>, to 3 dimensions via PCA. This data set consists of 70,000 28×28 images of handwritten digits, 0 through 9, usually broken into a 60000 point

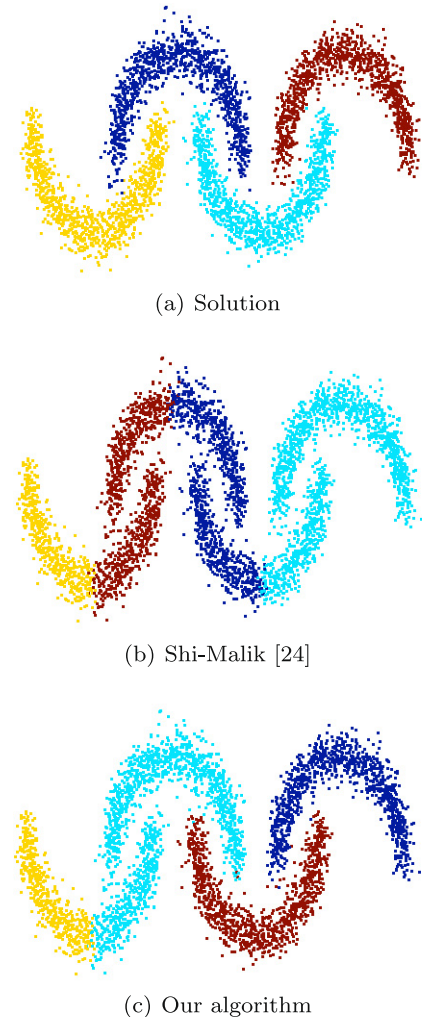


Fig. 2 Unsupervised classification for the four-moon dataset. The standard recursive two-class partitioning approach is applied. Figure (b) shows the result by spectral clustering [24] and figure (c) presents the result with Algorithm 1. Although our algorithm produces a better result than spectral clustering, it still fails to compute the solution. When more than two classes are considered then the quality of the results given by the recursive algorithm actually strongly depends on the choice of the initialization. In fact, for most initializations, the standard recursive two-class partitioning approach will not be able to give the solution (Note: it is a color figure)

training set and a 10000 point test set; thus the data is presented as 70000 points in \mathbb{R}^{784}). The data was preprocessed by projecting onto 50 principal components. Table 2 com-

Table 2 Unsupervised learning result for the MNIST dataset. *Column* (a) reports the *minimum* energy value (1) considered among 10 random initializations and the error is the percentage of misclassified data

of the minimum energy. *Column* (b) reports the *average* energy value (1) and the misclassification error for 10 random initializations. Note that the same random initialization was used for Algorithms 1 and [15]

	(a) Best energy (error)	(b) Mean energy (error)
Algorithm 1	0.6014 (11.6829 %)	0.6020 (11.6986 %)
Hein and Bühler [15]	0.6014 (11.6829 %)	0.6023 (11.7014 %)
Spectral clustering [24]	0.8546 (29.8833 %)	0.8546 (29.8833 %)




Fig. 3 Projection into a 3D space (via PCA) of the MNIST benchmark dataset. This data set consists of 60,000 28×28 images and 10,000 training images (each image is a data point in \mathbb{R}^{784}) of handwritten digits, 0 through 9 (Note: it is a color figure)

compares quantitatively our algorithm with the spectral clustering method of Shi and Malik [24] and the related method of Hein and Bühler in [15]. Our method and [15] outperform the spectral clustering method.

3 Transductive Data Classification with ℓ^1 Relaxation of the Multi-class Cheeger Cut

In this section, we extend the unsupervised two-phase Cheeger learning algorithm of Sect. 2 to a transductive multi-class Cheeger learning algorithm. The most natural extension of (1) to K classes is as follows:

$$\min_{\Omega_1, \dots, \Omega_K} \sum_{k=1}^K \frac{\text{Cut}(\Omega_k, \Omega_k^c)}{\min(|\Omega_k|, |\Omega_k^c|)}$$

$$\text{s.t. } \bigcup_{k=1}^K \Omega_k = V \text{ and } \Omega_i \cap \Omega_j = \emptyset \quad \forall i \neq j$$

The previous minimization problem is equivalent to the following problem:

$$\min_{\{u_k\}_{k=1}^K \in \{0,1\}} \sum_{k=1}^K \frac{\|Du_k\|_1}{\|u_k - m(u_k)\|_1} \tag{14}$$

$$\text{s.t. } \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V$$

The set of minimization used in the above minimization problem is not convex because binary functions do not make a convex set. Thus we consider the following relaxation:

$$\min_{\{u_k\}_{k=1}^K \in [0,1]} \sum_{k=1}^K \frac{\|Du_k\|_1}{\|u_k - m(u_k)\|_1} \tag{15}$$

$$\text{s.t. } \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V$$

In Sect. 2, we recall that the continuous ℓ^1 relaxation of the two-phase Cheeger minimization problem is exact, meaning that the (continuous) solution of (2) provides a (discrete) solution of the original Cheeger problem (1). We do not know if the ℓ^1 relaxation is still exact when multiple classes are considered, i.e. if the (continuous) solution of (15) provides a (discrete) solution of the original multi-class Cheeger problem (14). For the multi-class Cheeger-based learning problem considered in this paper, experiments show that the solutions $\{u_k\}_{k=1}^K$ are close to binary functions, but there is no theoretical guarantee of this observation.

As the transductive learning problem is considered here then a (small) set l_k of labels is given for each class Ω_k (i.e. $l_k \subset \Omega_k$, see Fig. 4) and the following minimization problem is thus considered:

$$\min_{\Omega_1, \dots, \Omega_K} \sum_{k=1}^K \frac{\text{Cut}(\Omega_k, \Omega_k^c)}{\min(|\Omega_k|, |\Omega_k^c|)}$$

$$\text{s.t. } \bigcup_{k=1}^K \Omega_k = V \text{ and } \Omega_i \cap \Omega_j = \emptyset \quad \forall i \neq j$$

$$\text{and given } \{l_k\}_{k=1}^K \tag{16}$$




Fig. 4 Illustration of the multi-class transductive problem. Given a set of labeled data points $\{l_k\}_{k=1}^K$ (the colored points), the objective is to find the data classes $\{\Omega_k\}_{k=1}^K$ that minimize the Cheeger energy (16)

which is equivalent to:

$$\begin{aligned} & \min_{\{u_k\}_{k=1}^K \in \{0,1\}} \sum_{k=1}^K \frac{\|Du_k\|_1}{\|u_k - m(u_k)\|_1} \\ & \text{s.t. } \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V \quad \text{and} \\ & u_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \end{aligned}$$

and which is relaxed to:

$$\begin{aligned} & \min_{\{u_k\}_{k=1}^K \in [0,1]} \sum_{k=1}^K \frac{\|Du_k\|_1}{\|u_k - m(u_k)\|_1} \\ & \text{s.t. } \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V \quad \text{and} \\ & u_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \end{aligned}$$

We now extend the two-phase algorithm designed in Sect. 2 to the multi-phase case:

$$\begin{aligned} & \min_{\{u_k\}_{k=1}^K \in [0,1]} \max_{\{\lambda_k\}_{k=1}^K \in \mathbb{R}} \sum_{k=1}^K \|Du_k\|_1 - \lambda_k \|u_k\|_1 \\ & \text{s.t. } m(u_k) = 0, \quad \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V, \quad \text{and} \\ & u_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \end{aligned}$$

The median constraint is relaxed to $\mathbf{1} \cdot u_k \leq |V|/K$. We again consider a standard two-step iterative algorithm:

- (i) Fix λ_k , compute the solution for the K minimization problems:

$$u_k^{n+1} = \operatorname{argmin}_{u_k \in [0,1]} \|Du_k\|_1 - \lambda_k^n \|u_k\|_1$$

$$\begin{aligned} & \text{s.t. } m(u_k) = 0, \quad \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V, \quad \text{and} \\ & u_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \end{aligned}$$

- (ii) Fix u_k , compute the solution of the K maximization problems:

$$\lambda_k^{n+1} = \operatorname{argmax}_{\lambda \in \mathbb{R}} \|Du_k^{n+1}\|_1 - \lambda \|u_k^{n+1}\|_1 \tag{17}$$

The minimization problems (17) are solved as follows:

$$\begin{cases} (u_k^{n+1}, v_k^{n+1}, d_k^{n+1}) \\ = \operatorname{argmin}_{u_k, v_k \in [0,1], d_k} \|d_k\|_1 \\ - \lambda \|v_k\|_1 + \alpha_{dk} \cdot (d_k - Du_k) + \frac{r_d}{2} |d_k - Du_k|^2 \\ + \alpha_{vk} \cdot (v_k - u_k) + \frac{r_v}{2} (v_k - u_k)^2 \\ + \alpha_{mk} \cdot (\mathbf{1} \cdot v_k - |V|/K) \\ \text{s.t. } \sum_{k=1}^K v_k = 1 \text{ and} \\ v_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \\ \alpha_{dk}^{n+1} = \alpha_d^n + r_d \cdot (d_k^{n+1} - Du_k^{n+1}) \\ \alpha_{vk}^{n+1} = \alpha_v^n + r_v \cdot (v_k^{n+1} - u_k^{n+1}) \\ \alpha_{mk}^{n+1} = \max(0, \alpha_m^n + r_m \cdot (\mathbf{1} \cdot v_k^{n+1} - |V|/K)) \end{cases} \tag{18}$$

The solution of the minimization problems w.r.t. u_k, v_k, d_k is the same as the solution given in the previous section. Finally, the projection onto the convex simplex set $\sum_{k=1}^K v_k = 1$ is given by [21, 28]. Observe that the final solution $\{u_k^*\}_{k=1}^K$ of (17) is not guaranteed to be binary. Hence, a conversion step is required to make $\{u_k^*\}_{k=1}^K$ binary. The most natural conversion is as follows:

$$\hat{u}_k^*(i) = \begin{cases} 1 & \text{if } k = \operatorname{argmax}_{p \in \{1, \dots, K\}} u_p^*(i) \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in V \tag{19}$$

where $\{\hat{u}_k^*\}_{k=1}^K$ are binary functions satisfying $\sum_{k=1}^K \hat{u}_k^* = 1$.

To summarize the algorithm introduced in this section, we write down the pseudo-code Algorithm 2. Eventually, Fig. 5 presents a simple illustration of the proposed multi-class Cheeger transductive model.

4 Transductive Data Classification with ℓ^1 Relaxation of the Multi-class Mumford-Shah-Potts Model

In this section, we develop an alternative to the multi-class Cheeger transductive classification algorithm introduced in the previous section. A successful multi-phase segmentation algorithm in imaging is the multiphase piecewise constant Mumford-Shah method [22] (continuous setting) or the Potts method [23] (discrete setting). These methods are well suited to solve the image segmentation problem and the idea

Algorithm 2 Transductive learning with ℓ^1 relaxation of the multi-class Cheeger cut

```

 $u_k^{n=0}$  given by a few steps of heat diffusion of the indicator
functions of labels
while outer loop not converged do
   $\alpha_{d_k}^{q=0}, \alpha_{v_k}^{q=0}, \alpha_{m_k}^{q=0} \leftarrow 0$ 
  while inner loop not converged do
     $u_k^{n+1,q+1}$  given by (9)
     $v_k^{n+1,q+1}$  given by (10) + simplex projection
    [21, 28] + labels given by (18)
     $d_k^{n+1,q+1}$  given by (11)
     $\alpha_{d_k}^{q+1}$  given by (18)
     $\alpha_{v_k}$  given by (18)
     $\alpha_{m_k}^{q=0}$  given by (18)
  end while
   $\lambda_k^{n+1}$  given by (13)
end while
    
```

in this section is to extend them to the transductive learning problem. Note that the piecewise constant Mumford-Shah/Potts models have been first implemented with the level set method [27, 30] and the graph cut method [5]. However, these methods are either too slow, not optimal, not accurate enough or the memory allocation can be important. Recent advances in ℓ^1 optimization algorithms provide efficient tool to solve the piecewise constant Mumford-Shah/Potts models [1, 6, 8, 19, 20, 28]. These recent improvements will be used to develop an efficient algorithm for the transductive Potts model:


$$\begin{aligned}
 & \min_{\Omega_1, \dots, \Omega_K} \sum_{k=1}^K \underbrace{\text{Cut}(\Omega_k, \Omega_k^c)}_{\simeq \text{Per}(\Omega_k)} \\
 & \text{s.t. } \bigcup_{k=1}^K \Omega_k = V \quad \text{and} \quad \Omega_i \cap \Omega_j = \emptyset \quad \forall i \neq j \\
 & \text{and given } \{l_k\}_{k=1}^K
 \end{aligned} \tag{20}$$

where Per stands for perimeter. The relationship between cut and perimeter comes from the coarea formula [25]. In the continuous setting, we have $\text{Per}(\Omega) = \int_{\mathcal{M} \subset \mathbb{R}^d} |\nabla f|$ when $f = 1_\Omega(x)$ is the indicator function of a geometric set Ω defined as $f(x) = 1 \ \forall x \in \Omega$ and 0 otherwise. Then, discretizing the total variation energy leads to $\int_{\mathcal{M} \subset \mathbb{R}^d} |\nabla f| \simeq \sum_{i,j \in V} w_{i,j} |f_i - f_j|$ and plugging the indicator function $f_i = 1_\Omega(i)$ finally gives $\text{Per}(\Omega) \simeq \sum_{i,j \in V} w_{i,j} |f_i - f_j| = \sum_{i \in \Omega, j \in \Omega^c} w_{i,j} = \text{Cut}(\Omega, \Omega^c)$. The minimization problem (20) is equivalent to the following problem:

$$\min_{\{u_k\}_{k=1}^K \in \{0,1\}} \sum_{k=1}^K \|Du_k\|_1$$



(a) Solution



(b) Spectral



(c) Algorithm 2

Fig. 5 Transductive classification of the four-moon dataset. The objective is to classify the four moons using 3 labels for each moon. Figure (b) presents the result with the spectral method (ℓ^2 relaxation) and figure (c) shows the result with the ℓ^1 relaxation of the multi-class Cheeger cut (Algorithm 2). The ℓ^1 relaxation produces a better classification result than the ℓ^2 relaxation (Note: it is a color figure)

$$\text{s.t. } \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V, \quad \text{and}$$

$$u_k(i) = \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases}$$

The set of minimization used in the above minimization problem is not convex because binary functions do not make a convex set. Thus we consider the following relaxation:

$$\min_{\{u_k\}_{k=1}^K \in [0,1]} \sum_{k=1}^K \|Du_k\|_1$$

$$\begin{aligned} \text{s.t. } & \sum_{k=1}^K u_k(i) = 1 \quad \forall i \in V, \quad \text{and} \\ u_k(i) &= \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \end{aligned}$$

The previous minimization problem is solved as:

$$\left\{ \begin{aligned} & (u_k^{n+1}, v_k^{n+1}, d_k^{n+1}) \\ & = \operatorname{argmin}_{u_k, v_k \in [0,1], d_k} \|d_k\|_1 \\ & \quad + \alpha_{d_k} \cdot (d_k - Du_k) + \frac{r_d}{2} |d_k - Du_k|^2 \\ & \quad + \alpha_{v_k} \cdot (v_k - u_k) + \frac{r_v}{2} (v_k - u_k)^2 \\ \text{s.t. } & \sum_{k=1}^K v_k = 1 \quad \text{and} \\ v_k(i) &= \begin{cases} 1 & \text{if } i \in l_p \text{ and } k = p \\ 0 & \text{if } i \in l_p \text{ and } k \neq p \end{cases} \\ \alpha_{d_k}^{n+1} &= \alpha_{d_k}^n + r_d \cdot (d_k^{n+1} - Du_k^{n+1}) \\ \alpha_{v_k}^{n+1} &= \alpha_{v_k}^n + r_v \cdot (v_k^{n+1} - u_k^{n+1}) \end{aligned} \right.$$

The solution of the minimization problems w.r.t. u_k, d_k is the same as the solution given in Sect. 2. The minimization w.r.t. v_k is simply given by:

$$v_k^* = \Pi_{[0,1]}(f_{v_k}) \quad \text{with } f_{v_k} := u_k - \frac{\alpha_{v_k}}{r_v} \quad (21)$$

Algorithm 3 Transductive learning with ℓ^1 relaxation of multi-class Mumford-Shah-Potts model

```

 $u_k^{n=0}$  given by a few steps of heat diffusion of the indicator
functions of labels
 $\alpha_{d_k}^{n=0}, \alpha_{v_k}^{n=0}, \alpha_m^{n=0} \leftarrow 0$ 
while outer loop not converged do
   $u_k^{n+1}$  given by (9)
   $v_k^{n+1}$  given by (21) + simplex projection [21, 28] +
  labels given by (18)
   $d_k^{n+1}$  given by (11)
   $\alpha_{d_k}^{n+1}$  given by (18)
   $\alpha_{v_k}^{n+1}$  given by (18)
end while

```

and project onto the convex simplex set $\sum_{k=1}^K v_k = 1$ using [21, 28]. Observe that the final solution $\{u_k^*\}_{k=1}^K$ of (17) is not guaranteed to be binary. Hence, a conversion step is required to make $\{u_k^*\}_{k=1}^K$ binary. Like in the previous section, the binary conversion is as follows:

$$\hat{u}_k^*(i) = \begin{cases} 1 & \text{if } k = \arg \max_{p \in \{1, \dots, K\}} u_p^*(i) \\ 0 & \text{otherwise} \end{cases} \quad \forall i \in V \quad (22)$$

Table 3 Transductive learning for the four-moon dataset. Column (I) reports the *minimum* energy value considered among 10 random tests and the associated error of misclassified data and computational time. Column (II) reports the *average* energy value, the misclassification er-

ror and the computational time for the 10 tests. The energy considered for the Cheeger model is (16), for the Mumford-Shah-Potts is (20) and for the spectral method is (23). Finally, n_l is the number of (randomly selected) labeled data for each class

	(I) Best energy (error/time)	(II) Mean energy (error/time)
Algorithm 2 (Cheeger)	1.981 (35.27 %/52.66 sec)	2.040 (33.53 %/52.54 sec)
Algorithm 3 (Mumford-Shah-Potts)	1046 (24.47 % /7.72 sec)	1418 (28.36 % /5.56 sec)
Spectral clustering [2]	0.000 (38.47 %/3.49 sec)	0.000 (40.96 %/3.67 sec)
(a) $n_l = 1$		
Algorithm 2 (Cheeger)	1.530 (0.47 % /55.47 sec)	1.621 (9.41 % /49.86 sec)
Algorithm 3 (Mumford-Shah-Potts)	1190 (43.97 % /7.53 sec)	1510 (21.86 %/6.33 sec)
Spectral clustering [2]	0.316 (23.72 %/2.57 sec)	0.703 (12.09 %/3.43 sec)
(b) $n_l = 3$		
Algorithm 2 (Cheeger)	1.530 (0.47 % /58.79 sec)	1.532 (0.45 % /55.17 sec)
Algorithm 3 (Mumford-Shah-Potts)	1533 (0.50 %/6.63 sec)	1567 (4.53 %/5.81 sec)
Spectral clustering [2]	0.778 (1.07 %/3.66 sec)	1.197 (2.73 %/3.86 sec)
(c) $n_l = 6$		
Algorithm 2 (Cheeger)	1.530 (0.47 %/31.39 sec)	1.530 (0.46 % /45.18 sec)
Algorithm 3 (Mumford-Shah-Potts)	1532 (0.45 % /4.61 sec)	1534 (0.47 %/5.35 sec)
Spectral clustering [2]	1.589 (0.57 %/3.48 sec)	2.378 (0.53 %/3.56 sec)
(d) $n_l = 10$		
Algorithm 2 (Cheeger)	1.530 (0.45 %/18.06 sec)	1.537 (0.41 %/16.32 sec)
Algorithm 3 (Mumford-Shah-Potts)	1534 (0.37 % /2.94 sec)	1541 (0.39 % /3.26 sec)
Spectral clustering [2]	8.996 (0.37 % /3.25 sec)	9.349 (0.42 %/3.37 sec)
(e) $n_l = 100$		

Table 4 Transductive learning for the MNIST dataset. *Column (I)* reports the *minimum* energy value considered among 10 random tests and the associated error of misclassified data and computational time. *Column (II)* reports the *average* energy value, the misclassification error and the computational time for the 10 tests. The energy considered for the Cheeger model is (16), for the Mumford-Shah-Potts is (20) and

for the spectral method is (23). Finally, 60,000 unlabeled data points are considered and $n_l \times 10$ labeled data points where n_l is the number of (randomly selected) labeled data for each class. Therefore, the total number of data points for each experiment is $60,000 + n_l \times 10$ data points

	(I) Best energy (error/time)	(II) Mean energy (error/time)
Algorithm 2 (Cheeger)	1.370 (2.451 % /51.42 min)	1.371 (2.454 % /56.06 min)
Algorithm 3 (Mumford-Shah-Potts)	7839 (12.198 % /4.91 min)	8371 (11.254 % /5.39 min)
Spectral clustering [2]	0.000 (35.89 % /1.05 min)	0.000 (29.346 % /1.02 min)
(a) $n_l = 1$		
Algorithm 2 (Cheeger)	1.370 (2.444 % /48.54 min)	1.370 (2.449 % /46.20 min)
Algorithm 3 (Mumford-Shah-Potts)	8176 (2.311 % /4.69 min)	8191 (2.417 % /4.51 min)
Spectral clustering [2]	1.022 (12.50 % /0.86 min)	1.538 (9.529 % /0.86 min)
(b) $n_l = 5$		
Algorithm 2 (Cheeger)	1.371 (2.469 % /48.81 min)	1.373 (2.456 % /50.46 min)
Algorithm 3 (Mumford-Shah-Potts)	8208 (2.529 % /3.19 min)	8235 (2.458 % /3.73 min)
Spectral clustering [2]	1.725 (5.988 % /0.87 min)	2.245 (4.254 % /0.87 min)
(c) $n_l = 10$		
Algorithm 2 (Cheeger)	1.370 (2.433 % /51.29 min)	1.372 (2.407 % /45.60 min)
Algorithm 3 (Mumford-Shah-Potts)	8307 (2.271 % /3.10 min)	8335 (2.305 % /2.79 min)
Spectral clustering [2]	5.833 (3.224 % /0.87 min)	6.239 (2.728 % /0.87 min)
(d) $n_l = 50$		
Algorithm 2 (Cheeger)	1.369 (2.355 % /52.08 min)	1.375 (2.365 % /44.29 min)
Algorithm 3 (Mumford-Shah-Potts)	8426 (2.177 % /2.52 min)	8454 (2.304 % /2.25 min)
Spectral clustering [2]	8.840 (2.177 % /0.88 min)	9.126 (2.480 % /0.89 min)
(e) $n_l = 100$		

where $\{\hat{u}_k^*\}_{k=1}^K$ satisfy $\sum_{k=1}^K \hat{u}_k^* = 1$.

To summarize the algorithm introduced in this section, we write down the pseudo-code Algorithm 3.

5 Experiments

In this section, we show classification results using the transductive algorithms developed in Sects. 3 and 4. We will work on the four moons and MNIST datasets described above. For both data sets, we build the weights matrix using the self-tuning construction of [29]. We use ten nearest neighbors, and the tenth neighbor determines the local scale. The universal scaling parameter is set to 1. For Algorithm 1, we set $r_d = 10$, $r_v = 100$, $r_m = 6K/N$, where N is the number of data points and K is the number of classes, and $\delta_\lambda = 0.4$. For Algorithm 3, we set $r_d = 10$ and $r_v = 100$. We choose the labeled points randomly, and fix a number of labeled points to draw from each class.

We compare Algorithm 2 and Algorithm 3 with a spectral transductive learning method from [2], which uses linear least squares on the eigenvectors of the normalized Laplacian to estimate the classes. That is, given the weight matrix W as before, we set $\mathcal{L} = I - S^{-1/2}WS^{-1/2}$, where S

is the diagonal matrix with the row sums on the diagonal, that is, $S_{ii} = \sum_j W_{ij}$. We compute the $l + 1$ lowest eigenvalue eigenvectors ϕ_0, \dots, ϕ_l of \mathcal{L} , and form the $N \times l$ matrix $\Phi = [\phi_1 \dots \phi_l]$; note that as usual we have omitted the density vector ϕ_0 . Each row of Φ corresponds to a data point. Next we form the matrix Φ_{lab} by extracting the rows of Φ corresponding to the labeled data points. Let L denote the number of classes, and p be the number of labeled data points. Given the $p \times L$ binary label matrix Y , we compute

$$A = (\Phi_{lab}^T \Phi_{lab})^{-1} \Phi_{lab}^T Y$$

which minimizes the least square energy:

$$\|Y - \Phi_{lab} A\|_2^2 \tag{23}$$

To compute the class labels of the unlabeled points, we set $R = \Phi A$, and let

$$y_j = \operatorname{argmax}_i R_{ji}$$

Tables 3 and 4 compare the proposed ℓ^1 relaxations of the multi-class Cheeger cut (Algorithm 2) and the Mumford-Shah-Potts (Algorithm 3) with the competitive spectral

method of [2] (by selecting the number l of eigenvectors which minimizes the error). We have tested different numbers of labels (n_l is the number of labeled data for each class) that are selected randomly. We repeat each experiment 10 times. For each experiment, the labeled points were chosen randomly and the same labeled points were used for the multi-class Cheeger cut model, the Mumford-Shah-Potts model and the spectral method. The ℓ^1 relaxations of the multi-class Cheeger cut and the Mumford-Shah-Potts outperform the spectral method in all cases, significantly so when a very small number of points are labeled.

6 Conclusion

The paper introduces new ℓ^1 relaxation methods for the multi-class transductive learning problem. These relaxation methods are inspired from recent advances in imaging science which offer fast, accurate and robust ℓ^1 optimization tools which allow to go beyond standard ℓ^2 relaxation methods, i.e. spectral clustering methods. Experiments demonstrate that the ℓ^1 relaxations of the multi-class Cheeger cut and the Mumford-Shah-Potts outperform the spectral clustering method, and even more significantly when a very small number of labels is considered.

Reproducible Research The code is available at <http://www.cs.cityu.edu.hk/~xbresson/codes.html#learningmulti>.

Acknowledgement Xavier Bresson is supported by the Hong Kong RGC under Grant GRF110311.

References

- Bae, E., Yuan, J., Tai, X.-C.: Global minimization for continuous multiphase partitioning problems using a dual approach. *Int. J. Comput. Vis.* **92**(1), 112–129 (2009)
- Belkin, M.: Problems of learning on manifolds. PhD thesis, University of Chicago (2003)
- Bertozzi, A., Flenner, A.: Diffuse interface models on graphs for classification of high dimensional data. UCLA CAM Report 11-27 (2011)
- Bioucas-Dias, J.M., Figueiredo, M.A.: A new TwIST: two-step iterative shrinkage/thresholding algorithms for image restoration. *IEEE Trans. Image Process.* **16**(12), 2992–3004 (2007)
- Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Trans. Pattern Anal. Mach. Intell.* **26**(9), 1124–1137 (2004)
- Brown, E.S., Chan, T.F., Bresson, X.: A convex relaxation method for a class of vector-valued minimization problems with applications to Mumford-Shah segmentation. UCLA CAM Report 10-43 (2010)
- Bühler, T., Hein, M.: Spectral clustering based on the graph p -Laplacian. In: *International Conference on Machine Learning*, pp. 81–88 (2009)
- Chambolle, A., Cremers, D., Pock, T.: A convex approach for computing minimal partitions. Technical Report TR-2008-05, Dept. of Computer Science, University of Bonn, Bonn (2008)
- Cheeger, J.: A lower bound for the smallest eigenvalue of the Laplacian. *Problems in Analysis*, 195–199 (1970)
- Chung, F.R.K.: *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, vol. 92 (1997). Published for the Conference Board of the Mathematical Sciences, Washington, DC
- Dinkelbach, W.: On nonlinear fractional programming. *Manag. Sci.* **13**, 492–498 (1967)
- Donoho, D.: De-noising by soft-thresholding. *IEEE Trans. Inf. Theory* **41**(33), 613–627 (1995)
- Glowinski, R., Le Tallec, P.: *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. SIAM, Philadelphia (1989)
- Goldstein, T., Osher, S.: The split Bregman method for L1-regularized problems. *SIAM J. Imaging Sci.* **2**(2), 323–343 (2009)
- Hein, M., Bühler, T.: An inverse power method for nonlinear eigenproblems with applications in 1-spectral clustering and sparse PCA. In: *Advances in Neural Information Processing Systems (NIPS)*, pp. 847–855 (2010)
- Hein, M., Setzer, S.: Beyond spectral clustering—tight relaxations of balanced graph cuts. In: *Advances in Neural Information Processing Systems (NIPS)* (2011)
- Kolev, K., Cremers, D.: Continuous ratio optimization via convex relaxation with applications to multiview 3D reconstruction. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009)
- Kolmogorov, V., Boykov, Y., Rother, C.: Applications of parametric maxflow in computer vision. In: *International Conference on Computer Vision*, pp. 1–8 (2007)
- Lellmann, J., Kappes, J., Yuan, J., Becker, F., Schnörr, C.: Convex multi-class image labeling by simplex-constrained total variation. In: *International Conference on Scale Space and Variational Methods in Computer Vision*, pp. 150–162 (2009)
- Lellmann, J., Schnörr, C.: Continuous multiclass labeling approaches and algorithms. Univ. of Heidelberg, Tech. Rep. (2010)
- Michelot, C.: A finite algorithm for finding the projection of a point onto the canonical simplex of \mathbb{R}^n . *J. Optim. Theory Appl.* **50**(1), 195–200 (1986)
- Mumford, D., Shah, J.: Optimal approximations of piecewise smooth functions and associated variational problems. *Commun. Pure Appl. Math.* **42**, 577–685 (1989)
- Potts, R.B., Domb, C.: Some generalized order-disorder transformations. *Math. Proc. Camb. Philos. Soc.* **48**, 106–109 (1952)
- Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(8), 888–905 (2000)
- Strang, G.: Maximal flow through a domain. *Math. Program.* **26**, 123–143 (1983)
- Szlám, A., Bresson, X.: Total variation and cheeger cuts. In: *Proceedings of the 27th International Conference on Machine Learning*, pp. 1039–1046 (2010)
- Vese, L.A., Chan, T.F.: A multiphase level set framework for image segmentation using the Mumford and Shah model. *Int. J. Comput. Vis.* **50**(3), 271–293 (2002)
- Zach, C., Gallup, D., Frahm, J.M., Niethammer, M.: Fast global labeling for real-time stereo using multiple plane sweeps. In: *Vision, Modeling, and Visualization*, pp. 243–252 (2008)
- Zelnik-Manor, L., Perona, P.: Self-tuning spectral clustering. In: *Advances in Neural Information Processing Systems 17 (NIPS 2004)* (2004)
- Zhao, H.K., Chan, T.F., Merriman, B., Osher, S.: A variational level set approach to multiphase motion. *J. Comput. Phys.* **127**, 179–195 (1996)



Xavier Bresson received a B.A. of Physics from University of Marseille, a M.Sc. in Electrical Engineering from Ecole Supérieure d'Electricité (SUPELEC) in Paris and a M.Sc. in Signal Processing from University of Paris XI. In 2005, he completed a Ph.D. at the Swiss Federal Institute of Technology (EPFL). In 2006–2010, he joined the Department of Mathematics at University of California, Los Angeles (UCLA) as a Postdoctoral Scholar. In 2010, he joined the Department of Computer Science at

City University of Hong Kong as an Assistant Professor. His area of research is computational mathematics applied to machine learning and imaging science. He has developed efficient and generic algorithms to solve a large span of problems in these research fields. He has published more than 40 papers in international journals and conferences. On Feb. 2013, he organized the workshop on “Convex Relaxation Methods for Geometric Problems in Scientific Computing” at the Institute for Pure and Applied Mathematics (IPAM) at UCLA. In Feb. 2014, he will participate to the Program on “Network Science and Graph Algorithms” at the Institute for Computational and Experimental Research in Mathematics (ICERM) at Brown University. He is also organizing the SIAM Conference on Imaging Science in Hong Kong on May 2014.



Xue-Cheng Tai received the licentiate degree and the Ph.D. degree in applied mathematics from Jyväskylä University, Jyväskylä, Finland in 1989 and 1991, respectively. After holding several research positions in Europe, he became an Associate Professor in 1994 at the University of Bergen, Bergen, Norway, and a Professor in 1997. He was also a faculty member of Nanyang Technological University of Singapore from 2007 to 2010. He has been a Member of the “Center for Mathematics for Appli-

cations” in Oslo and a Member of the “Center of Integrated Petroleum Research” in Bergen. His research interests include numerical PDE for image processing, multigrid and domain decomposition methods, iterative methods for linear and nonlinear PDE problems, and parallel computing. His recent research interests are main on image processing and computer vision. He has guided numerous master’s and Ph.D. students, and published more than 80 scientific papers. He has been a reviewer and an editor for a number of international journals.



Tony F. Chan received the B.S. and M.S. degrees in engineering from CalTech, Pasadena, CA, and the Ph.D. degree in computer science from Stanford University, Stanford, CA. He was with CalTech (applied math) as a Research Fellow and taught computer science at Yale University, New Haven, CT, before joining the faculty at the University of California, Los Angeles (UCLA), in 1986 as Professor of mathematics. He served as a Chair of the Department of Mathematics, UCLA, from 1997–2000, and as Director of

the Institute for Pure and Applied Mathematics (IPAM) from 2000 to 2001. From July 2001 to June 2006, he served as Dean of Physical Sciences, UCLA, and became an Assistant Director, Directorate for Mathematics and Physical Sciences at National Science Foundation (2006–2009). He has been the President at the Hong Kong University of Science and Technology since October 2009, holding positions in the Mathematics and Computer Science Departments. His current research interests include mathematical image processing, computer vision, and computer graphics, computational brain mapping, VLSI physical design optimization, multiscale computational methods, multigrid and domain decomposition algorithms, iterative methods, Krylov subspace methods, parallel algorithms, and computational linear algebra.



Arthur Szlam is currently assistant professor at the City College of New York. Before that he was a postdoc at NYU and UCLA, and a graduate student at Yale university.