

Compassionately Conservative Balanced Cuts for Image Segmentation

Nathan D. Cahill

Tyler L. Hayes Renee T. Meinhold

John F. Hamilton

Rochester Institute of Technology

{ndcsma, tlh6792, rtm9271, jfhsms}@rit.edu

Abstract

The Normalized Cut (NCut) objective function, widely used in data clustering and image segmentation, quantifies the cost of graph partitioning in a way that biases clusters or segments that are balanced towards having lower values than unbalanced partitionings. However, this bias is so strong that it avoids any singleton partitions, even when vertices are very weakly connected to the rest of the graph. Motivated by the Bühler-Hein family of balanced cut costs, we propose the family of Compassionately Conservative Balanced (CCB) Cut costs, which are indexed by a parameter that can be used to strike a compromise between the desire to avoid too many singleton partitions and the notion that all partitions should be balanced. We show that CCB-Cut minimization can be relaxed into an orthogonally constrained ℓ_{τ} -minimization problem that coincides with the problem of computing Piecewise Flat Embeddings (PFE) for one particular index value, and we present an algorithm for solving the relaxed problem by iteratively minimizing a sequence of reweighted Rayleigh quotients (IRRQ). Using images from the BSDS500 database, we show that image segmentation based on CCB-Cut minimization provides better accuracy with respect to ground truth and greater variability in region size than NCut-based image segmentation.

1. Introduction

The Normalized Cut (NCut) graph partitioning cost function was first introduced over two decades ago to tackle the perceptual grouping problem [19, 20], and its approximate minimization via continuous relaxation into a generalized eigenvector problem has emerged as one of the fundamental techniques of data clustering [14].

The key observation that motivated the normalization of graph cut cost functions is that when graph partitioning is performed by minimizing the sum of the weights of edges that are cut (the Cut cost), the resulting partitions will be *unbalanced*; vertices that are weakly connected to the rest of the graph are likely to be separated to form singleton partitions, while the rest of the graph vertices are likely to re-



Figure 1: (a) A graph having unit weight edges except for the two edges with weight $\alpha/2$. (b) Minimizing the Cut cost removes the weakly-connected vertex, whereas (c) minimizing the Normalized Cut (NCut) cost yields a more "balanced" partitioning.

main in large partitions. Normalizing the cut cost by the degrees of each partition (NCut [19, 20]), the size of each partition (Average Cut [18]), or the minimum of the size of each partition and its complement (Ratio or Cheeger Cut [7]) yield balanced partitions that have similar total degree or size when minimized.

Is it possible that these types of normalization go too far? Consider the example of Figure 1, in which we wish to partition the toy graph into two subgraphs. This graph (1a) contains seven vertices and nine edges; all but two of the edges have unit weight, and the two indicated edges have weight $\alpha/2$ for $\alpha \in (0,1)$. Since $\alpha < 1$, the degree of the vertex sharing these two edges is guaranteed to be smaller than the degree of any other vertex in the graph. Hence, as shown in (1b), minimizing the Cut cost will separate that weakly-connected vertex from the rest of the graph. Minimizing the NCut cost creates a more balanced partitioning as shown in (1c); however, this partitioning will always have a lower-cost NCut than removing the weaklyconnected vertex, even when α is infinitesimally small. Perhaps the Ratio Cut (RCut) is a better choice: for this graph, $\alpha = 2/3$ is a critical value above which (1c) minimizes the RCut cost, and below which (1b) yields the minimum. But is there really anything special about $\alpha = 2/3$?

If this graph represented an image to be segmented, one would want to use ground-truth segmentation maps from



Figure 2: Using the Bühler-Hein *p*-Normalized (a) and *p*-Ratio (b) Cut costs to partition the graph in Fig. (1a). Over all possible (p, α) , either the partitionings in Fig.'s (1b) or (1c) yield the minimum value. For both families of costs, there are substantial ranges of α for which only one partitioning emerges for all *p*.

training images to suggest an optimal critical value for α , above which the more "balanced" partitioning is desired, and below which the weakly-connected vertex should be separated. The ground truth might dictate that even the Ratio Cut cost provides too much normalization, or instead, that it does not provide enough.

One potential idea is to use the infinite families of balanced cost functions defined by Bühler and Hein [4] and indexed by a parameter $p \in (1, \infty)$. We refer to these families as the Bühler-Hein *p*-Normalized (BHN_p) and the Bühler-Hein *p*-Ratio (BHR_p) Cuts. BHN_p is equivalent to NCut for p = 2 and the Normalized Cheeger (NCh) Cut as $p \rightarrow 1^+$, and BHR_p is equivalent to RCut for p = 2 and the Ratio Cheeger (RCh) Cut as $p \rightarrow 1^+$, so both families of cut costs can be thought of as interpolating between different normalization factors. However, as shown in Fig. (2), when we minimize BHN_p Cut and BHR_p Cut for the graph in Fig. (1), there are substantial ranges of α for which no value of *p* will allow a bifurcation between solutions.

In this paper, we show that although the Bühler-Hein costs provide a mechanism for interpolating between Normalized (or Ratio) Cuts and Normalized (or Ratio) Cheeger Cuts, they do not provide a path for interpolating all the way to the unnormalized Cut cost. However, a different set of infinite families of balanced cost functions *can* be constructed that do enable such interpolation. Our proposed families are called the Compassionately Conservative Normalized (CCN) Cut and the Compassionately Conservative Ratio (CCR) Cut. Both are indexed by a parameter $\tau \in (0, \infty)$; CCN_{τ} is equivalent to NCut for $\tau = 2$ and the unnormalized Cut cost for $\tau = 0$, and CCR_{τ} is equivalent to RCut for $\tau = 2$ and the unnormalized Cut cost for $\tau = 0$. As we can see from Fig. (3), when we minimize CCN_{τ} and CCR_{τ} for the graph in Fig. (1), different ranges of τ can be selected that will allow a bifurcation between solutions for any value of α .

This paper makes four novel contributions to graph par-



Figure 3: Using the CCN_{τ} (a) and CCR_{τ} (b) Cut costs to partition the graph in Fig. (1a). Over all possible (p, α) , either the partitionings in Fig.'s (1b) or (1c) yield the minimum value. For both families of costs, for any choice of $\alpha \in (0, 1)$, ranges of τ exist that enable either partitioning.

titioning in computer vision:

- presenting Compassionately Conservative Balanced (CCB) Cuts: families of cut costs that enable normalizations ranging from Normalized/Ratio Cuts to unnormalized Cuts and can be naturally extended to k-way partitionings with k > 2;
- presenting a continuous relaxation of CCB Cut minimization and illuminating its connection to computing Piecewise Flat Embeddings (PFE) [24];
- presenting an efficient algorithm for solving the relaxed problem via minimizing a succession of reweighted Rayleigh quotients (IRRQ); and
- demonstrating empirical advantages of the CCB Cut costs and the IRRQ minimization algorithm for the application of image segmentation.

2. Balanced Cut Costs

Consider an undirected weighted graph $\mathcal{G} = (V, E)$ that we wish to partition into two disjoint subgraphs $\mathcal{G}_i = (V_i, E_i), i = 1, 2$, by removing the edges connecting V_1 and V_2 . A standard cost of partitioning \mathcal{G} is the *Cut* cost, defined as the total weight of the edges that have been removed:

$$\operatorname{Cut}(V_1, V_2) = \sum_{v_i \in V_1, v_j \in V_2} w_{i,j} , \qquad (1)$$

where the vertex set $V = \{v_1, v_2, \dots, v_n\}$, and **W** is the weighted adjacency matrix (or *affinity* matrix) of \mathcal{G} .

Minimizing the Cut cost is undesirable, however, as it can often yield partitionings that simply disconnect one vertex from the rest of the graph [22]. More balanced partitions emerge if the Cut cost is normalized by some function of the total sizes or total degrees (volumes) of the subgraphs. Such balanced cut costs can be expressed generally by:

$$BCut(V_1, V_2) = \frac{Cut(V_1, V_2)}{\Theta(V_1, V_2)} , \qquad (2)$$

where $\Theta(V_1, V_2)$ is a symmetric function that decreases with increasing differences in the sizes or degrees of V_1 and V_2 . A variety of balanced cut costs have been proposed in the literature, including the Normalized Cut [19, 20], Average Cut [18], Ratio Cut [8, 10], and Normalized and Ratio Cheeger Cuts [4, 7].

In Theorem 4.1 of [4], Bühler and Hein show that there exists an infinite family of balanced cut costs that contains Normalized, Average, Ratio, and Cheeger Cuts. If we consider the following function:

$$\Phi_p(v) = \frac{2^p}{\left(\left(\frac{1}{v}\right)^{\frac{1}{p-1}} + \left(\frac{1}{1-v}\right)^{\frac{1}{p-1}}\right)^{p-1}} , \qquad (3)$$

for p > 1 and $v \in (0,1)$, then $\operatorname{BCut}(V_1, V_2)$ with $\Theta(V_1, V_2) = \Phi_p(|V_1|/|V|)$ is the aforementioned BHR_p Cut, and it approaches the Ratio Cheeger Cut as $p \to 1^+$ and equals the Average/Ratio Cuts for p = 2. If $\Theta(V_1, V_2) = \Phi_p(\operatorname{Vol}(V_1)/\operatorname{Vol}(V))$, $\operatorname{BCut}(V_1, V_2)$ is the BHN_p Cut, and it approaches the Normalized Cheeger Cut as $p \to 1^+$ and equals the Normalized Cut for p = 2. (The volume of a subgraph is defined as $\operatorname{Vol}(V_\ell) = \sum_{v_j \in V_\ell} d_j$, where $d_j = \sum_m w_{j,m}$ is the degree of vertex v_j .) Note that this is a rescaling of the balanced cut costs in [4], defined so that the maximum value of $\Theta(V_1, V_2)$ is always unity.

In addition to providing a way to interpolate between Cheeger Cuts and Normalized/Ratio Cuts for 1 ,(3) enables cut costs with normalizations that are slightlymore conservative (closer to unity) than Normalized/RatioCuts by selecting <math>p > 2. It is straightforward to show this by noting that:

$$\Phi_{\infty}(v) = \lim_{p \to \infty} \Phi_p(v) = \sqrt{4v(1-v)} = \sqrt{\Phi_2(v)} , \quad (4)$$

and so $1 > \Phi_{\infty}(v) > \Phi_2(v)$ for all $v \in (0, 1), v \neq 1/2$, which can be seen in the left side of Figure 4.

The behavior of (3) as $p \to \infty$ suggests that there is room to define even more conservative normalization functions. We do so in this paper by considering the following transformation of $\Phi_p(v)$:

$$\Psi_{\tau}(v) = \Psi_{\frac{2}{p-1}}(v) = (\Phi_{p}(v))^{\frac{1}{p-1}}$$

$$= \frac{2^{p/(p-1)}}{\left(\frac{1}{v}\right)^{\frac{1}{p-1}} + \left(\frac{1}{1-v}\right)^{\frac{1}{p-1}}} = \frac{2^{1+\tau/2}}{\left(\frac{1}{v}\right)^{\tau/2} + \left(\frac{1}{1-v}\right)^{\tau/2}} ,$$
(5)

where $\tau = 2/(p-1)$. (The reason for switching parametrizations from p to τ will be addressed at the end



Figure 4: Left: the family $\Phi_p(v)$ that enables the balanced cut costs defined by [4], which includes Normalized/Ratio Cuts and Cheeger Cuts but limits the extent of more conservative normalizations. Right: the family $\Psi_{\tau}(v)$ that enables the proposed balanced cut costs; for $0 < \tau < 2$, the proposed costs interpolate from Normalized/Ratio Cuts to Cuts with no normalization at all.

of Section 3.) Clearly $\Psi_2(v) = \Phi_2(v)$, and so Normalized/Ratio Cuts arise from employing (5) in place of (3) for $\tau = 2$. More interestingly, however, is that $\Psi_0(v) = 1$, and so for $0 < \tau < 2$, defining BCut (V_1, V_2) with $\Theta(V_1, V_2) =$ $\Psi_{\tau}(|V_1|/|V|)$ or $\Psi_{\tau}(\text{Vol}(V_1)/\text{Vol}(V))$ provides a way to interpolate between Normalized/Ratio Cuts and Cuts with no normalization at all. This behavior of increasing $\Psi_{\tau}(v)$ as $\tau \to 0^+$ can be seen in the right side of Figure 4.

Of pedagogical interest is that $\lim_{\tau\to\infty} \Psi_{\tau}(v) = \delta(v-1/2)$, and so as $\tau \to \infty$, employing (5) to define $\Theta(V_1, V_2)$ would yield balanced cut costs that diverge whenever the sizes or degrees of V_1 and V_2 are unequal. Since our interest in this paper is to explore the impact of cut costs having more conservative normalizations than Normalized/Ratio Cuts, we will restrict our attention to the use of $\Psi_{\tau}(v)$ for $\tau \in (0, 2)$.

In addition to enabling a greater range of conservative normalizations, the use of (5) in place of (3) yields an added benefit: the resulting Balanced Cut costs can be directly extended to form costs of partitionings into k > 2 subgraphs. (As discussed in [4], it is not obvious whether such a direct extension to the Bühler-Hein cut costs exists.) If we now consider that we wish to partition \mathcal{G} into k disjoint subgraphs $\mathcal{G}_i = (V_i, E_i), i = 1, 2, \ldots, k$, then an unbalanced way to measure the cost of the partitioning is the *multiway* cut cost, defined in terms of pairwise cut costs as:

$$\operatorname{Cut}(V_1,\ldots,V_k) = \frac{1}{2} \sum_{\ell=1}^k \operatorname{Cut}(V_\ell,V \setminus V_\ell) \quad . \tag{6}$$

The family of multiway balanced cut costs that we propose is described generally by the *Compassionately*

Conservative Balanced Cut:

$$CCB_{\tau}(V_1, \dots, V_k) = \frac{1}{2} \sum_{\ell=1}^k \frac{Cut(V_{\ell}, V \setminus V_{\ell})}{2^{\tau/2} \left(\sum_{v_j \in V_{\ell}} \pi_j \right)^{\tau/2}} , \quad (7)$$

where $\boldsymbol{\pi} = [\pi_1, \ldots, \pi_n]^{\mathrm{T}}$ is a user-defined vector of positive weights. If $\boldsymbol{\pi} = \mathbf{1}_n$, then $\sum_{v_j \in V_\ell} \pi_j = |V_\ell|$, and we refer to (7) as the *Compassionately Conservative Ratio Cut* (CCR_{τ}). If $\boldsymbol{\pi} = \mathbf{d}$ (the degree vector), then $\sum_{v_j \in V_\ell} \pi_j = \operatorname{Vol}(V_\ell)$, and we refer to (7) as the *Compassionately Conservative Normalized Cut* (CCN_{τ}). With some algebraic manipulation, it is straightforward to see that (7) is a generalization of (2); when k = 2, (7) reduces to (2) with $\Theta(V_1, V_2) = \Psi_{\tau} \left(\sum_{v_j \in V_1} \pi_j / \sum_{v_j \in V} \pi_j \right)$. Furthermore, (7) generalizes multiclass cut costs presented in the research literature; CCB₂, CCR₂, and CCN₂ are scaled versions of the Multiclass Penalized Cut [25], the Multiclass Ratio Cut [21], and the Multiclass Normalized Cut [23, 21], respectively, and so choosing $\tau \in (0, 2)$ generates Multiclass Balanced Cut costs that interpolate between (6) and any of these previously developed costs.

3. Relaxing the CCB Cut

As with many balanced cut costs, minimizing the CCB Cut is NP-hard. However, we can formulate a continuous relaxation of CCB_{τ} for all $\tau > 0$. Unfortunately, the spectral clustering relaxation of BHN_p and BHR_p in terms of the second eigenvector of the graph *p*-Laplacian [4] is of no help; it is only tight as $p \to 0^+$ (and therefore not for $\tau \in (0, 2)$), and furthermore, CCB Cuts are not even members of the Bühler-Hein family of cuts except for when $\tau = p = 2$. Hence, we must identify a different relaxation.

To do so, we first reformulate CCB_{τ} to express it in terms of an $n \times k$ indicator matrix **X** so that $X_{i,j} = 1$ if $v_i \in V_j$ and $X_{i,j} = 0$ otherwise. If \mathbf{x}_i is the *i*th column of **X**, then $\sum_{v_j \in V_i} \pi_j$ can be written in terms of the diagonal matrix $\mathbf{\Pi} = \operatorname{diag}(\pi)$ as $\mathbf{x}_i^T \mathbf{\Pi} \mathbf{x}_i$, and the pairwise cut cost between V_i and $V \setminus V_i$ can be written as $\operatorname{Cut}(V_i, V \setminus V_i) = \mathbf{x}_i^T \mathbf{W} (1 - \mathbf{x}_i) = \mathbf{x}_i^T \mathbf{d} - \mathbf{x}_i^T \mathbf{W} \mathbf{x}_i =$ $\mathbf{x}_i^T \mathbf{D} \mathbf{x}_i - \mathbf{x}_i^T \mathbf{W} \mathbf{x}_i = \mathbf{x}_i^T \mathbf{L} \mathbf{x}_i$, where $\mathbf{L} = \mathbf{D} - \mathbf{W}$. This allows us to express (7) as:

$$\operatorname{CCB}_{\tau}(V_1,\ldots,V_k) = \frac{1}{2} \sum_{\ell=1}^k \frac{\mathbf{x}_{\ell}^{\mathbf{T}} \mathbf{L} \mathbf{x}_{\ell}}{2^{\tau/2} (\mathbf{x}_{\ell}^{\mathbf{T}} \mathbf{\Pi} \mathbf{x}_{\ell})^{\tau/2}} \quad . \tag{8}$$

Minimizing (7) is equivalent to minimizing (8) subject to the constraint that $\mathbf{X}^{T}\mathbf{X}$ is positive diagonal, which ensures that none of the V_i 's will collapse to the empty set.

To guide us towards an appropriate CCB_{τ} relaxation, we first consider how the Multiclass Penalized Cut (CCB_2) can be relaxed. Using an argument similar to Yu and Shi [23],

we see that if $\mathbf{Y} = \mathbf{X} (\mathbf{X}^T \mathbf{\Pi} \mathbf{X})^{-1/2}$, then $\mathbf{Y}^T \mathbf{\Pi} \mathbf{Y} = \mathbf{I}$ and (7) is equivalent to a scalar multiple of tr $(\mathbf{Y}^T \mathbf{L} \mathbf{Y})$. Hence, the solution to minimizing a relaxed version of CCB₂ is $\tilde{\mathbf{Y}} = \mathbf{U}\mathbf{Q}$, where \mathbf{U} is the $n \times k$ matrix whose columns are the orthonormal eigenvectors $\mathbf{u}_2, \ldots, \mathbf{u}_{k+1}$ corresponding to the smallest nontrivial eigenvalues of $\mathbf{\Pi}^{-1/2}\mathbf{L}\mathbf{\Pi}^{-1/2}$, and \mathbf{Q} is an arbitrary $k \times k$ orthogonal matrix. The optimal solution $\tilde{\mathbf{X}}$ to (7) with $\tau = 2$ can then be approximated by k-means clustering [16], nonmaximal suppression [23] or Procrustean rounding [25] on $\tilde{\mathbf{Y}}$.

Note that if we define $\hat{\mathbf{y}}_i$ to be the transpose of the i^{th} row of \mathbf{Y} , then minimizing the relaxed version of CCB₂ is equivalent to solving the constrained minimization problem:

$$\min_{\mathbf{Y}\in\mathbb{R}^{n\times k}} \quad \mathcal{J}_2(\mathbf{Y}) := \sum_{i=1}^n \sum_{j=1}^n w_{i,j} \|\hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j\|_2^2 \qquad (9)$$

subject to: $\mathbf{Y}^{T} \mathbf{\Pi} \mathbf{Y} = \mathbf{I}$, $\mathbf{Y}^{T} \mathbf{\Pi} \mathbf{1} = \mathbf{0}$,

which, when $\pi = \mathbf{d}$, is identical to the Laplacian Eigenmaps (LE) problem [3] for computing embeddings of data that are assumed to lie on a manifold. The *balance constraint* $\mathbf{Y}^{T}\mathbf{\Pi}\mathbf{1} = \mathbf{0}$ is necessary to avoid eigenvectors of $\mathbf{\Pi}^{-1/2}\mathbf{L}\mathbf{\Pi}^{-1/2}$ corresponding to the trivial eigenvalue.

Turning our attention now to the more general case where $\tau > 0$, if we define $\alpha_{\ell} = \left(\sum_{i} \pi_{i} x_{i,\ell}^{2}\right)^{-1/2}$ for $\ell = 1, \ldots, k$, and we note that $\left(\mathbf{X}^{T} \mathbf{\Pi} \mathbf{X}\right)^{-1/2} = \text{diag}(\boldsymbol{\alpha})$, we can write:

$$CCB_{\tau}(V_{1},...,V_{k}) = \frac{1}{2} \sum_{\ell=1}^{k} \frac{\sum_{i,j} w_{i,j} (x_{i,\ell} - x_{j,\ell})^{2}}{2^{\tau/2} \left(\sum_{i} \pi_{i} x_{i,\ell}^{2}\right)^{\tau/2}}$$
$$= \frac{1}{2^{1+\tau/2}} \sum_{\ell=1}^{k} \alpha_{\ell}^{\tau} \sum_{i,j} w_{i,j} ||\mathbf{x}_{i,\ell} - x_{j,\ell}|^{\tau}$$
$$= \frac{1}{2^{1+\tau/2}} \sum_{i,j} w_{i,j} || \left(\mathbf{X}^{T} \mathbf{D} \mathbf{X}\right)^{-1/2} (\hat{\mathbf{x}}_{i} - \hat{\mathbf{x}}_{j}) ||_{\tau}^{\tau}$$
$$= \frac{1}{2^{1+\tau/2}} \sum_{i,j} w_{i,j} || \hat{\mathbf{y}}_{i} - \hat{\mathbf{y}}_{j} ||_{\tau}^{\tau} .$$
(10)

Hence, the relaxation of (8) is obtained by dropping the condition that $\mathbf{Y} = \mathbf{X} (\mathbf{X}^T \mathbf{\Pi} \mathbf{X})^{-1/2}$ and solving the constrained minimization problem:

$$\min_{\mathbf{Y} \in \mathbb{R}^{n \times k}} \quad \mathcal{J}_{\tau}(\mathbf{Y}) := \sum_{i=1}^{n} \sum_{j=1}^{n} w_{i,j} \| \hat{\mathbf{y}}_{i} - \hat{\mathbf{y}}_{j} \|_{\tau}^{\tau} \quad (11)$$

subject to: $\mathbf{Y}^{\mathbf{T}} \mathbf{\Pi} \mathbf{Y} = \mathbf{I}$.

Note that the balance constraint $\mathbf{Y}^{T} \mathbf{\Pi} \mathbf{1} = \mathbf{0}$ is only necessary for the case $\tau = 2$.

Remarkably, in the special case where $\pi = d$ and $\tau = 1$, (11) is exactly the Piecewise Flat Embedding (PFE) problem [24], whose solution is an embedding in which the data are naturally clustered due to the promotion of sparsity in the differences between rows of **Y**. It is for this reason that we choose to parametrize CCB Cuts in terms of τ and not p: solutions to the relaxed versions of CCB_{τ} for $\tau \in (0, 1]$ can be expected to be piecewise constant, consistent with the sparse nature of solutions to ℓ_{τ} -minimization problems for $\tau \in (0, 1]$.

4. IRRQ Minimization Algorithm

While (11) with $\tau = 2$ has a solution that can be written in terms of the eigenvectors corresponding to the smallest k nontrivial eigenvalues of $\Pi^{-1/2} \mathbf{L} \Pi^{-1/2}$, no analytical form of the solution is known for $\tau \neq 2$. In [24], Yu et al. proposed approximating the solution to (11) with $\tau = 1$ via the Splitting Orthogonality Constraint (SOC) algorithm [15], which requires an initial estimate of **Y** and performs a nested iteration with parameters at each level that must be tuned. While it is possible to partially generalize this approach to handle (11) with $\tau > 1$, $\mathcal{J}_{\tau}(\mathbf{Y})$ is non-convex for $\tau \in (0, 1)$, and so the SOC algorithm is not applicable in this regime. Furthermore, even in the regime in which it is applicable, the use of splitting in the algorithm formulation means that the solutions will not strictly satisfy the orthogonality constraint $\mathbf{Y}^T \mathbf{DY} = \mathbf{I}$.

Here, we propose an alternative algorithm for solving (11) that can be applied for all $\tau \in (0,2]$, that does not require an initial estimate of Y, and that does strictly satisfy the orthogonality constraint. Our alternative algorithm is motivated from the Iteratively Reweighted Least Squares (IRLS) algorithms commonly used to solve ℓ_1 minimization problems [9]. In IRLS, ℓ_1 -minimization is performed by iteratively solving a succession of weighted least-squares (ℓ_2 -minimization) problems, with the weights updated at each iteration so as to decrease the impact of large residual errors. IRLS algorithms do require initialization, but it is the *weights* that must be initialized as opposed to the *solution*. Weights are typically initialized to unity (although they can be initialized differently by an expert user), and in specific cases [9], IRLS algorithms have provable convergence guarantees.

In our relaxed problem (11), the presence of the orthogonality constraint renders IRLS algorithms invalid. However, solutions to (11) can be approximated by iteratively solving a series of *constrained* weighted ℓ_2 -minimization problems, each of the form:

$$\min_{\mathbf{Y} \in \mathbb{R}^{n \times k}} \quad \mathcal{J}_2^{\mathbf{\Gamma}}(\mathbf{Y}) := \sum_{i=1}^n \sum_{j=1}^n w_{i,j} \gamma_{i,j} \| \hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j \|_2^2$$
(12)
subject to: $\mathbf{Y}^{\mathbf{T}} \mathbf{\Pi} \mathbf{Y} = \mathbf{I} , \quad \mathbf{Y}^{\mathbf{T}} \mathbf{\Pi} \mathbf{I} = \mathbf{0} ,$

where Γ is the $n \times n$ matrix of weights (with entries $\gamma_{i,j}$) that is updated at each iteration in a manner similar to IRLS.

To establish a connection between (11) and (12), we first eliminate the balance constraint from (12) using the result of the following Lemma, which is proved in Appendix C of [6]:

Lemma 4.1. Let

$$\begin{aligned} \mathcal{Y} &= \left\{ \mathbf{Y} \in \mathbb{R}^{n \times k} | \mathbf{Y}^{\mathbf{T}} \mathbf{\Pi} \mathbf{Y} = \mathbf{I}, \mathbf{Y}^{\mathbf{T}} \mathbf{\Pi} \mathbf{1} = \mathbf{0} \right\} \ , \\ \mathcal{G} &= \left\{ \mathbf{G} \in \mathbb{R}^{(n-1) \times k} | \mathbf{G}^{\mathbf{T}} \mathbf{G} = \mathbf{I} \right\} \ , \end{aligned}$$

and define $\aleph(\mathbf{G}) = \Pi^{-1/2} \mathbf{B} \mathbf{G}$ for $\mathbf{G} \in \mathcal{G}$, where $\mathbf{B} = \mathbf{M} (\mathbf{M}^T \mathbf{M})^{-1/2}$, $\mathbf{M} \in \mathbb{R}^{n \times (n-1)}$ is a full rank matrix with null $(\mathbf{M}^T) = span(\mathbf{q})$, and $\mathbf{q} = \Pi^{1/2} \mathbf{1} / \| \Pi^{1/2} \mathbf{1} \|$. Then \aleph is a bijection from \mathcal{G} to \mathcal{Y} .

Using this Lemma allows us to solve (12) by first solving:

$$\hat{\mathbf{G}} := \arg\min_{\mathbf{G}^{\mathbf{T}}\mathbf{G}=\mathbf{I}} \mathcal{J}_{2}^{\mathbf{\Gamma}}(\boldsymbol{\aleph}(\mathbf{G})) \quad , \tag{13}$$

and then computing $\mathbf{Y} = \aleph(\hat{\mathbf{G}})$.

An analogy to IRLS suggests that under the assumption that component-wise differences in the embedding do not vanish, the best choice of weights for (12) would be $\gamma_{i,j} = \|\hat{\mathbf{y}}_i^* - \hat{\mathbf{y}}_j^*\|_{\tau}^{\tau} / \|\hat{\mathbf{y}}_i^* - \hat{\mathbf{y}}_j^*\|_2^2$ for $i \neq j$, where \mathbf{Y}^* is the solution to (11). This choice of weights would yield $\mathcal{J}_2^{\Gamma}(\mathbf{Y}^*) = \mathcal{J}_{\tau}(\mathbf{Y}^*)$. In practice, however, \mathbf{Y}^* is unknown, and many of the component-wise differences in the embedding will vanish. Hence, we propose using regularized weights as suggested in [9]:

$$\gamma_{i,j} := \left[w_{i,j} \| \hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j \|_2^2 + \varepsilon^2 \right]^{-1 + \tau/2} \quad , \qquad (14)$$

and we update ε according to the schedule prescribed by [9], which suggests:

$$\varepsilon \leftarrow \min(\varepsilon, n^{-1}r(\mathbf{Y})_{\kappa+1})$$
, (15)

where $r(\mathbf{Y})_{\kappa}$ is the κ^{th} largest element of $\left\{ w_{i,j}^{1/2} \| \hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j \|_2, \forall i, j = 1, \dots, n \right\}.$

Combining the steps of solving (12) and updating (14)–(15) into a sequence of iterations yields Algorithm 1 for computing the solution to (11). Since (12) is equivalent to (13), and (13) can be transformed into an unconstrained minimization of a Rayleigh Quotient, we term this algorithm *Iteratively Reweighted Rayleigh Quotient* (IRRQ) Minimization.

Unlike SOC, IRRQ requires tuning of only a single hyperparameter κ , it can be used for any $\tau \in (0, 2)$, and it guarantees a solution in which the orthogonality constraint is strictly enforced. Furthermore, IRRQ does not

Algorithm 1 IRRQ Algorithm for Solving (11)
procedure IRRQ($\mathbf{W}, k, \tau, \kappa$)
$\gamma_{i,j}^{(0)} := 1, \varepsilon_0 := 1, m := 0, n := \operatorname{size}(\mathbf{W}, 1)$
while $\varepsilon_m > 0$ do
(a) $\mathbf{Y}^{(m+1)} := \arg\min_{\mathbf{Y}_{\mathbf{T}}^{\mathbf{T}}\mathbf{\Pi}\mathbf{Y}=\mathbf{I}} \mathcal{J}_{2}^{\mathbf{\Gamma}^{(m)}}(\mathbf{Y})$
$\mathbf{Y}^{\mathrm{T}} \mathbf{\Pi} 1 = 0$
(b) $\varepsilon_{m+1} := \min\left(\varepsilon_m, n^{-1}r(\mathbf{Y}^{(m+1)})_{\kappa+1}\right)$
(c) $\gamma_{i,j}^{(m+1)} := \left[w_{i,j} \ \hat{\mathbf{y}}_i - \hat{\mathbf{y}}_j \ _2^2 + \varepsilon_{m+1}^2 \right]^{-1+\tau/2}$
end while
return $\mathbf{Y}^{(m)}$
end procedure

require a preprocessing step to estimate an initial clustering; rather, only the weights $\gamma_{i,j}$ must be initialized, and they can all be initialized to unity. Interestingly, this initialization is equivalent to implicitly using an initial clustering that corresponds to the solution of the relaxed NCut problem. This is because $\mathcal{J}_2^{11^{\text{TT}}}(\mathbf{Y})$ is equivalent to the LE objective function $\mathcal{J}_2(\mathbf{Y})$ in (9). If different initializations are desired, for instance, by computing an initial embedding $\mathbf{Y}^{(0)}$ using a Gaussian Mixture Model as in [24], these can be incorporated by setting the initial weights to be $\gamma_{i,j}^{(0)} = \left\|\hat{\mathbf{y}}_i^{(0)} - \hat{\mathbf{y}}_j^{(0)}\right\|_{\tau}^{\tau} / \left\|\hat{\mathbf{y}}_i^{(0)} - \hat{\mathbf{y}}_j^{(0)}\right\|_2^2$ or $\gamma_{i,j}^{(0)} = \left[w_{i,j}\right] \left\|\hat{\mathbf{y}}_i^{(0)} - \hat{\mathbf{y}}_j^{(0)}\right\|_2^2 + \varepsilon_0^2 \right]^{-1+\tau/2}$.

Solving IRRQ Step (a)

Computing $\mathbf{Y}^{(m+1)}$ in step (a) of the IRRQ minimization algorithm is nontrivial. From the relationship between (12)–(13), we can see that computing $\mathbf{Y}^{(m+1)}$ is equivalent to solving

$$\hat{\mathbf{G}}^{(m+1)} := \arg\min_{\mathbf{G}^{\mathbf{T}}\mathbf{G}=\mathbf{I}} \mathcal{J}_2^{\mathbf{\Gamma}^{(m)}}(\boldsymbol{\aleph}(\mathbf{G})) \quad , \qquad (16)$$

and then computing $\mathbf{Y}^{(m+1)} = \boldsymbol{\aleph} \left(\hat{\mathbf{G}}^{(m+1)} \right)$. Problem (16) can be expressed as the Rayleigh quotient minimization:

$$\hat{\mathbf{G}}^{(m+1)} := \arg\min \operatorname{tr} \left(\mathbf{G}^{\mathbf{T}} \mathbf{B}^{\mathbf{T}} \mathbf{\Pi}^{-1/2} \hat{\mathbf{L}}^{(m)} \mathbf{\Pi}^{-1/2} \mathbf{B} \mathbf{G} \right. \\ \left. \cdot \left(\mathbf{G}^{\mathbf{T}} \mathbf{G} \right)^{-1} \right) , \qquad (17)$$

where $\hat{\mathbf{L}}^{(m)}$ is the Laplacian of the graph having weight matrix $\mathbf{W} \odot \mathbf{\Gamma}^{(m)}$ and \odot denotes Hadamard product. The solution to (17) is given by $\hat{\mathbf{G}}^{(m+1)} = \mathbf{U}\mathbf{H}$, where $\mathbf{U} \in \mathbb{R}^{(n-1)\times k}$ is the matrix whose columns are the orthonormal eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_k$ corresponding to the smallest eigenvalues of $\mathbf{B}^T \mathbf{\Pi}^{-1/2} \hat{\mathbf{L}}^{(m)} \mathbf{\Pi}^{-1/2} \mathbf{B}$, and $\mathbf{H} \in \mathbb{R}^{k \times k}$ is an arbitrary orthogonal matrix. (Note that by eliminating the balance constraint, we also eliminate the possibility of a trivial eigenvalue of $\mathbf{B}^{T} \mathbf{\Pi}^{-1/2} \hat{\mathbf{L}}^{(m)} \mathbf{\Pi}^{-1/2} \mathbf{B}$. Such an eigenvalue would have eigenvector \mathbf{p} for which $\mathbf{B}\mathbf{p}$ is in the direction of \mathbf{q} ; however, this is contradicted by the fact that range(\mathbf{B}) = range(\mathbf{M}^{T}).)

The ability of this solution to scale to large *n* depends critically on the structure of **M**, and on whether or not **B** must explicitly be constructed. Since we are free to choose *any* full-rank $n \times (n-1)$ matrix **M** such that $\mathbf{M}^{\mathbf{T}}\mathbf{q} = \mathbf{0}$, we make the particular choice $\mathbf{M}^{\mathbf{T}} = [\hat{\mathbf{q}}| - q_1\mathbf{I}_{n-1}]$, where $\hat{\mathbf{q}} = [q_2, q_3, \dots, q_n]^{\mathbf{T}}$ and \mathbf{I}_{n-1} is the $(n-1) \times (n-1)$ identity matrix. This particular choice of **M** is sparse, and therefore, as shown in Appendix D of [6], the multiplication of an arbitrary vector by $\mathbf{B}^{\mathbf{T}}\mathbf{\Pi}^{-1/2}\hat{\mathbf{L}}^{(m)}\mathbf{\Pi}^{-1/2}\mathbf{B}$ can be performed efficiently without explicitly constructing **B**.

A final note is that the solution to step (a) is not actually unique: $\mathbf{Y}^{(m+1)}$ can be postmultiplied by \mathbf{H}^{T} and still be a valid solution (recalling that \mathbf{H} is orthogonal). This is not a problem for steps (b) and (c) of Algorithm 1, because r and $\gamma_{i,j}$ are invariant to such transformations of $\mathbf{Y}^{(m+1)}$. As a consequence, IRRQ minimization could yield an entire family of solutions to the PFE problem. This could be problematic because the ℓ_{τ} -norms/pseudo-norms are not invariant under orthogonal transformations for $\tau \neq 2$. In practice, however, we have found that the ℓ_{τ} -norms/pseudo-norms are minimized for the choice $\mathbf{H} = \mathbf{I}$ and we therefore suggest this choice. Proof that this is the best choice remains an open problem.

Choosing κ for Rapid Convergence

In IRLS algorithms for ℓ_{τ} -minimization, linear convergence can typically be achieved for $\tau \geq 1$ and super-linear convergence for $\tau \in (0,1)$ if κ is chosen large enough so that if the resulting solution is θ -sparse, then $\kappa > \theta$. (Actually, there are much more sophisticated convergence results in [9], but this is a good rule-of-thumb.) Without attempting to prove convergence results for the IRRO algorithm, we use a similar strategy in choosing κ . In practice, the main difficulty in choosing κ is that θ is not known exactly until the problem is solved. To approximate θ , we use an estimate $\hat{\theta}$ equal to twice the number of graph edges that connect different clusters from a k-means clustering performed on the initial embedding $\mathbf{Y}^{(0)}$. For "scale-free" behavior, we introduce the hyperparameter $\tilde{\kappa} \in (0,1)$ that can then be mapped to κ by $\kappa = \hat{\theta} + \tilde{\kappa} \left(2 |E| - \hat{\theta} \right)$, where |E| is the total number of edges in the graph.

5. CCB Cuts for Segmentation

To investigate the performance of CCB Cuts for image segmentation, we segment the 200 test images in the BSDS500 dataset [1], each of which has a variety of manually-labeled segmentations with different numbers of



Figure 5: Example BSDS500 test image and segmentation results. First row: Lab affinities; second row: mPb affinities; third row: PMI affinities. *-M indicates multiway segmentation; *-H indicates hierarchical 2-way segmentation. k = 28 clusters specified for each case. Images appearing to have fewer segments actually have a number of small/singleton segments.



Figure 6: Hierarchical clustering results for NCuts (CCN_2 Cuts) and CCN_1 Cuts for different numbers of clusters k and Lab affinities. Top row: segmentation maps; bottom row: colorized segmentation maps.

segments that can be used as ground truth. We explore three different affinity constructions for each image:

- Lab: a Gaussian kernel on squared Euclidean distances in L*a*b* color space;
- **mPb**: an exponentiated negative maximum of the *multiscale probability of boundary* (mPb) using the code provided with the BSDS500 dataset [1]; and,
- **PMI**: an exponentiated version of a statistic related to *pointwise mutual information* (PMI) used in crisp boundary detection [13].

We use a 10-pixel radius to define neighborhoods for affinity construction. For efficiency, we downsample each affinity matrix by two scale levels (and subsequently upsample the computed embeddings) using the strategy in [2, 17].

Figure 5 illustrates segmentation results on a BSDS500 test image for both multi-way and hierarchical 2-way segmentation using various types of balanced cuts. Although different segmentations appear to have different numbers of regions, many of the regions are small/singleton, and all segmentation results have the same number of regions. Figure 6 shows how hierarchical 2-way segmentation using CCN Cuts with $\tau = 1$ does not exhibit the iterative "chopping" behavior common to hierarchical NCuts minimization.

Quantitative Validation

Drawing any conclusions based on visual assessment of segmentation maps can be problematic in the absence of quantitative validation. First, to validate that CCB segmentation yields less balanced partitions than NCut and Cheeger Cut-based segmentation, we show in Figure 7 the "degree spread" for segmentation results from various methods across all test images, as measured by the ratio of the standard deviation of the partition volumes to the mean partition volume.

Next, we evaluate segmentation performance with respect to ground truth using the three criteria discussed in [1] and used in [24]: Segmentation covering, Probabilistic Rand Index (PRI), and Variation of Information (VI). Covering and PRI increase and VI decreases as segmentations become closer to the ground truth. We segment each BSDS500 test image multiple times by minimizing CCB Cuts for various values of $\tau \in (0, 2]$, once corresponding to each value of k (number of segments) reflected in one of the ground-truth segmentations, and once for each of k = 5, 10, 15, 20, and 25, if any of these are not reflected in the ground truth. Each multi-way segmentation algorithm proceeds by solving/approximating the corresponding embedding and then subsequently performing k-means clustering on the result. We also perform segmentation using hierarchical 2-way cuts; in this case, clustering is performed at each step by identifying the threshold that best minimizes



Figure 7: Ratio of the standard deviation of partition volumes to the mean partition volume for CCN_{τ} , NCut and Normalized Cheeger Cut-based segmentation of BSDS images for various numbers of segments.

the balanced cut cost in question. For comparison purposes, we also apply the inverse power method of [11]; this method hierarchically approximates 2-way cuts that minimize the Normalized/Ratio Cheeger costs.

In CCB Cut minimization (for both multi-way and hierarchical 2-way segmentation), we iterate until the relative change in cost between two iterations falls below 1%, or until a maximum of 50 iterations are performed, whichever occurs first. The hyperparameter $\tilde{\kappa}$ is set to 10^{-4} for Lab affinities and to 0.2 for mPb and PMI affinities. Following the strategy of [24], we report results for both a *fixed* scheme, in which we run the algorithm repeatedly with kcorresponding to each number of segments in the groundtruth and average the performance results from the multiple runs, and the *dynamic* scheme, in which we choose the value of k from k = 5, 10, 15, 20, and 25 that yields the best performance for a particular image. Table 1 shows the performance results for a subset of the cut costs; to save space, we excluded results from smaller values of τ , which typically performed worse than $\tau = 1$ but better than $\tau = 2$, and we excluded results from ratio-costs, which are similar to those of normalized-costs.

A few interesting conclusions can be drawn. First, the use of simple Lab affinities surprisingly yields superior results compared to the more complicated mPb and PMI affinities. Second, PRI values appear to indicate that greater degrees of balance are better, whereas VI and Covering values seem to indicate that CCN with $\tau = 1$ provides the best results. Third, better results are found when hierarchical 2-way segmentation is performed as opposed to simultaneous multiway segmentation. However, in our MATLAB implementation, multiway segmentation is significantly faster than hierarchical 2-way segmentation, as shown in Table 2. In the future, it would be interesting to compare these results to those of hierarchical 2-way segmentation via minimizing ratios of differences of set functions [12, 5].

Affinity	Method	Covering		PRI		VI	
		fixed	dynamic	fixed	dynamic	fixed	dynamic
Lab	NCut-M	0.48	0.65	0.70	0.83	2.17	1.50
	$CCN_1 - M$	0.49	0.66	0.69	0.83	2.13	1.47
	NCut-H	0.51	0.67	0.78	0.87	1.98	1.43
	$CCN_1 - H$	0.57	0.69	0.69	0.78	1.81	1.38
	NCheeger	0.46	0.62	0.78	0.87	2.21	1.66
mPb	NCut-M	0.29	0.48	0.75	0.84	2.86	2.10
	$CCN_1 - M$	0.32	0.51	0.75	0.83	2.75	1.96
	NCut-H	0.30	0.48	0.83	0.89	2.61	2.00
	$CCN_1 - H$	0.38	0.52	0.73	0.78	2.35	1.83
	NCheeger	0.27	0.48	0.83	0.88	2.90	2.31
PMI	NCut-M	0.34	0.53	0.76	0.85	2.62	1.86
	$CCN_1 - M$	0.40	0.58	0.77	0.86	2.34	1.65
	NCut-H	0.37	0.54	0.87	0.90	2.37	1.75
	$CCN_1 - H$	0.47	0.60	0.78	0.81	1.99	1.55
	NCheeger	0.33	0.51	0.84	0.89	2.65	2.04

Table 1: Comparison of various segmentation methods on BSDS500 test set, averaged across images. For each affinity construction, the best results for each performance measure are highlighted in bold. *–M indicates multiway segmentation; *–H indicates hierarchical 2-way segmentation.

Cost Function	Quartile I	Median	Quartile III	Max
$CCN_{\tau} - Multiway$	2.0	2.9	4.2	14.9
CCN_{τ} – Hierarchical	13.0	23.4	41.6	258.2
NCheeger	2.4	4.0	9.4	434.6
CCR_{τ} – Multiway	1.0	1.4	1.9	8.0
CCR_{τ} – Hierarchical	10.0	18.6	31.6	180.7
RCheeger	1.4	2.6	9.5	378.1

Table 2: Statistics of computation times (minutes) required for segmentation, excluding affinity construction. Statistics are computed across all values of τ , all numbers of clusters, all images, and all affinity types.

6. Conclusion

Compassionately Conservative Balanced Cuts enable normalizations ranging from Normalized/Ratio Cuts to unnormalized Cuts, allowing a tradeoff between generating partitions that are too similar in size/degree and avoiding all singletons. They can be directly applied to the k-way partitioning problem for any $k \ge 2$, and their minimization can be relaxed into a problem that can be solved by minimizing a succession of Rayleigh quotients. Image segmentation experiments show that minimizing the CCB Cut yields more accurate results than minimizing NCuts/RCuts and generates regions having greater variability in size/degree.

A. Code

Prototype implementations of algorithms in this paper are available at the MATLAB Central File Exchange under File ID #66158.

B. Acknowledgements

NDC thanks Thomas Bühler and Matthias Hein for helpful discussions.

References

- P. Arbeláez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 33(5):898–916, May 2011. 6, 7
- [2] P. Arbeláez, J. Pont-Tuset, J. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping. In *Proc. Computer Vision and Pattern Recognition*, CVPR, pages 328– 335, 2014. 7
- [3] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003. 4
- [4] T. Bühler and M. Hein. Spectral clustering based on the graph p-Laplacian. In Proc. International Conference on Machine Learning, ICML, pages 81–88. ACM, 2009. 2, 3, 4
- [5] T. Bühler, S. S. Rangapuram, S. Setzer, and M. Hein. Constrained fractional set programs and their application in local clustering and community detection. In *Proc. International Conference on Machine Learning*, ICML, pages 624–632, 2013. 8
- [6] N. D. Cahill, T. L. Hayes, R. T. Meinhold, and J. H. Hamilton. Compassionately conservative balanced cuts for image segmentation. In *eprint arXiv:1803.09903*, March 2018. 5, 6
- [7] J. Cheeger. A lower bound for the smallest eigenvalue of the laplacian. In *Proceedings of the Princeton Conference in Honor of Professor S. Bochner*, pages 195–199, 1969. 1, 3
- [8] C.-K. Cheng and Y.-C. A. Wei. An improved two-way partitioning algorithm with stable performance. *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, 10(12):1502–1511, 1991. 3
- [9] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics*, 63(1):1–38, 2010. 5, 6
- [10] L. Hagen and A. B. Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, 11(9):1074– 1085, 1992. 3
- [11] M. Hein and T. Bühler. An inverse power method for nonlinear eigenproblems with applications in 1-spectral clustering and sparse PCA. In *Advances in Neural Information Processing Systems*, NIPS, pages 847–855, 2010. 8
- [12] M. Hein and S. Setzer. Beyond spectral clustering tight relaxations of balanced graph cuts. In Advances in Neural Information Processing Systems, NIPS, pages 2366–2374, 2011. 8
- [13] P. Isola, D. Zoran, D. Krishnan, and E. H. Adelson. Crisp boundary detection using pointwise mutual information. In *Proc. European Conference on Computer Vision*, ECCV, pages 799–814. Springer, 2014. 7
- [14] A. K. Jain. Data clustering: 50 years beyond K-means. Pattern Recognition Letters, 31(8):651–666, 2010. 1
- [15] R. Lai and S. Osher. A splitting method for orthogonality constrained problems. *Journal of Scientific Computing*, 58(2):431–449, 2014. 5

- [16] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems*, NIPS, pages 849–856, 2002. 4
- [17] J. Pont-Tuset, P. Arbeláez, J. T. Barron, F. Marques, and J. Malik. Multiscale combinatorial grouping for image segmentation and object proposal generation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 39(1):128–140, 2017. 7
- [18] S. Sarkar and P. Soundararajan. Supervised learning of large perceptual organization: Graph spectral partitioning and learning automata. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(5):504–525, 2000. 1, 3
- [19] J. Shi and J. Malik. Normalized cuts and image segmentation. In *Proc. Computer Vision and Pattern Recognition*, CVPR, pages 731–737, Jun 1997. 1, 3
- [20] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(8):888–905, Aug 2000. 1, 3
- [21] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007. 4
- [22] Z. Wu and R. Leahy. An optimal graph theoretic approach to data clustering: theory and its application to image segmentation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 15(11):1101–1113, Nov 1993. 2
- [23] S. X. Yu and J. Shi. Multiclass spectral clustering. In Proc. International Conference on Computer Vision, ICCV, pages 313–319, 2003. 4
- [24] Y. Yu, C. Fang, and Z. Liao. Piecewise flat embedding for image segmentation. In *Proc. International Conference on Computer Vision*, ICCV, pages 1368–1376, 2015. 2, 5, 6, 7, 8
- [25] Z. Zhang and M. I. Jordan. Multiway spectral clustering: A margin-based perspective. *Statistical Science*, 23(3):383– 403, 2008. 4