

Segmenting Planar Superpixel Adjacency Graphs w.r.t. Non-planar Superpixel Affinity Graphs

Bjoern Andres^{1,*}, Julian Yarkony^{2,*}, B.S. Manjunath², Steffen Kirchhoff¹,
Engin Turetken³, Charless C. Fowlkes⁴, and Hanspeter Pfister¹

¹ Harvard University

² UC Santa Barbara

³ EPFL

⁴ UC Irvine

Abstract. We address the problem of segmenting an image into a previously unknown number of segments from the perspective of graph partitioning. Specifically, we consider minimum multicuts of superpixel affinity graphs in which all affinities between non-adjacent superpixels are negative. We propose a relaxation by Lagrangian decomposition and a constrained set of re-parameterizations for which we can optimize exactly and efficiently. Our contribution is to show how the planarity of the adjacency graph can be exploited if the affinity graph is non-planar. We demonstrate the effectiveness of this approach in user-assisted image segmentation and show that the solution of the relaxed problem is fast and the relaxation is tight in practice.

1 Introduction

The formalization of the image segmentation problem as a multicut problem has recently attracted considerable attention [2,3,6,9,10,15,17]. This problem consists in finding a partition of a weighted superpixel adjacency graph into connected components (segments) such that the set of edges that straddle different segments (the multicut) has minimum total weight. A positive edge weight penalizes and a negative edge weight rewards all segmentations in which the connected superpixels are in different segments. The weights depend on the image in a way this is typically learned from data. They are called an *affinity*.

With the notable exception of [10], recent work has focused on the problem described so far in which affinities are only defined for pairs of *adjacent* superpixels. This has been motivated by three reasons. First and primary is that affinities between adjacent superpixels can be estimated reasonably well with methods such as the global probability of boundary [4]. Second is that the number of affinities grows linearly in the number of superpixels which is beneficial for modeling and optimization. Third is that superpixel adjacency graphs are *planar* which affords a relaxation of the multicut problem that can be solved efficiently and is tight in practice [17], although the multicut problem for planar graphs remains NP-hard [5].

* Authors contributed equally.




Fig. 1. a) In natural images, we often observe pixels at different locations that clearly belong to different segments but no clear position where the segment boundary should be located. This occurs whenever there is a smooth transition from the color of one region to another color of another region. In this case, affinities between adjacent superpixels do not provide strong cues which leads to errors in the segmentation. We show in Section 5 that negative affinities between non-adjacent superpixels can overcome this problem. **b)** At the same time, *positive* affinities between non-adjacent superpixels do not convey strong cues because separate segments can have similar texture.

In this paper, we study the generalization of the problem that allows for negative affinities between non-adjacent superpixels. The objective is to find a segmentation of the superpixel adjacency graph such that the sum of affinities between superpixels in different segments is minimal. We propose a relaxation by Lagrangian decomposition and a constrained set of re-parameterizations for which we can optimize exactly and efficiently. We demonstrate the effectiveness of this approach in user-assisted image segmentation and show that the solution of the relaxed problem is fast and the relaxation is tight in practice.

Our concentration on *negative* non-local affinities is motivated by the following observations: In natural images, we often observe pixels at different locations that clearly belong to different segments but no clear position where the segment boundary should be located. This occurs whenever there is a smooth transition from the color of one region to another color of another region (Fig. 1a). In this case, affinities between adjacent superpixels do not provide strong cues which leads to errors in the segmentation. We show in Section 5 that negative affinities between non-adjacent superpixels can overcome this problem. At the same time, *positive* affinities between non-adjacent superpixels do not convey strong cues because separate segments can have similar texture (Fig. 1b). A conceptual difficulty of non-local positive affinities is discussed in Appendix A.1.

2 Related Work

The multicut problem is known to be NP-complete, [7], even for planar graphs [5]. In computer vision, this combinatorial optimization problem has been used

to formalize image segmentation. To date, multicut of superpixel adjacency graphs [2,17] are among the closest, in terms of partition metrics, to the man-made segmentations in the Berkeley Segmentation Benchmark [4].

Recent work on the multicut problem for computer vision application has afforded an exact cutting plane algorithm for general graph that is applicable if edge weights are strong [3], efficient greedy algorithms suitable for large problems and problems where edge weights are weak [6], as well as diverse linear programming (LP) relaxations [9,10,12,15,17]. The Lagrangian decomposition we propose in this paper is built on [17] which considers optimal multicut of planar graphs. Our decomposition extends [17] by allowing for the introduction of negative affinities between non-adjacent superpixels.

3 Image Segmentation by Multicuts

We now introduce the mathematical framework for our discussion. We consider a graph (V, E) in which vertices indicate superpixels and weighted edges quantify an affinity between superpixels. A partition of the node set is encoded as a labeling $x \in \{0, 1\}^E$ of edges. We use $x_e = 0$ to indicate that the pair of vertices connected by edge e are in the same segment and use $x_e = 1$ to indicate that the pair are in separate segments. We call edge e *uncut* if $x_e = 0$ and *cut* if $x_e = 1$. The multicut of a partition is the set of all edges that are cut.

Edge weights $\theta \in \mathbb{R}^E$ define an objective function over partitions. Here, θ_e is called the affinity of the superpixels connected by the edge e . Negative values indicate a reward for superpixels to be in separate segments and positive values indicate a cost for superpixels to be in separate segments. A partition that minimizes the sum of the affinities of cut edges is called a minimum partition. The corresponding multicut is called a minimum multicut. We write the objective value of a multicut (indicated by) x as the inner product $\theta^T x$.

Not all binary labelings x correspond to partitions. Consider a graph of three superpixels a, b, c which are connected by edges ab, ac, bc . Consider also the labeling $x_{ab} = 1, x_{ac} = 0, x_{bc} = 0$. This labeling states that superpixels a and b are in separate segments while saying that both a and b are in the same segment as c , which is contradiction.

A necessary and sufficient condition for a labeling to define a partition is that there are no edges e such that $x_e = 1$ within a connected component [8]. This condition can be written as a set of inequalities as follows: For every cycle of edges c and every edge f in c , if $x_f = 1$, then at least one other edge in the cycle is cut. These inequalities are called the *cycle inequalities*. The multicut problem can be written as an integer program with cycle inequalities.

$$\begin{aligned} \min_{x \in \mathbb{R}^E} \quad & \theta^T x & (1) \\ \text{subject to} \quad & \sum_{e \in c \setminus f} x_e \geq x_f \quad \forall c \in \text{cycles}(V, E) \quad \forall f \in c & (2) \\ & x_e \in \{0, 1\} \quad \forall e \in E . \end{aligned}$$




Fig. 2. We propose a Lagrangian decomposition of the multicut problem that consists of the following sub-problems: 1. the independent edge sub-problem, 2. the planar multicut sub-problem, 3. a set of cycle sub-problems.

4 Lagrangian Decomposition

We now propose a Lagrangian decomposition of the multicut problem that consists of the sub-problems depicted in Fig. 2. In the independent edge sub-problem, edges can be cut or uncut independently. In the planar multicut sub-problem, any solution must be a partition of the vertices and thus obey the cycle inequalities. Affinities between non-adjacent superpixels are not included in the planar multicut sub-problem. In a cycle sub-problem an edge between non-adjacent superpixels is associated with a path between the corresponding vertices of that edge. A cycle sub-problem enforces that the number of edges cut in that cycle is not equal to one, thus enforcing the cycle inequalities for that cycle.

We use θ^{ie} , θ^{pmc} , and θ^c to denote the affinities of the independent edge, planar multicut and the c -th cycle sub-problem, respectively. The corresponding solutions are x^{ie} , x^{pmc} , x^c . For short hand, we refer to edges between adjacent superpixels as planar edges and edges between non-adjacent superpixels as non-planar edges. We use θ^{NP} to denote the affinities of the non-planar edges and θ^P to denote the affinities of planar edges. We index non-planar edges with \hat{e} and planar edges with e . Given the sub-problems, we write the Lagrangian decomposition below.

$$\max_{[\theta^{ie}, \theta^{pmc}, \theta^c \forall c]} \min_{x^{ie}} (\theta^{ie})^T x^{ie} + \min_{x^{pmc}} (\theta^{pmc})^T x^{pmc} + \sum_{c \in C} \min_{x^c} (\theta^c)^T x^c \quad (3)$$

subject to

$$\theta^{ie} + \theta^{pmc} + \sum_{c \in C} \theta^c = \theta \quad (4)$$

4.1 The Planar Multicut Sub-problem

Each of the sub-problems is tractable, except for the planar multicut sub-problem which is NP-hard [5]. To make this sub-problem tractable, we constrain the re-parameterization. Specifically, we enforce that the minimal objective value of

the planar multicut sub-problem is zero¹. Consequently, the empty multicut is optimal and thus, optimization is tractable. In order for the objective value of the planar multicut sub-problem to be zero, it is sufficient to enforce that every 2-colorable segmentation² has a non-negative value [17].

We introduce a matrix Z in which each row corresponds to a 2-colorable segmentation. We index 2-colorable segmentations by r . For every index r and every edge e , $Z_{re} = 1$ indicates that e is cut in the r -th 2-colorable segmentation. The constraint that every 2-colorable partition has a non-negative value is formalized below.

$$Z\theta^{pmc} \geq 0 \quad (5)$$

4.2 The Independent Edge Sub-problem

The affinities θ^{ie} in the independent edge sub-problem have important properties. First, we establish $\theta^{ie} \leq 0$ for every $e \in E$. Suppose there exists an edge e such that $\theta_e^{ie} > 0$. Then, $x_e^{ie} = 0$ in any optimal solution of the independent edge sub-problem. Notice that setting $\theta_e^{pmc} := \theta_e^{pmc} + \theta_e^{ie}$ and $\theta_e^{ie} := 0$ does not loosen the bound of Lagrangian decomposition and does not decrease the optimal value of the planar multicut sub-problem. Thus, we can restrict the re-parameterization to allocate non-positive affinity to every edge in the independent edge sub-problem. Furthermore, this means we can take the value of the independent edge sub-problem to be the sum of the terms in θ^{ie} as cutting every edge in the independent edge sub-problem is an optimal solution.

Another important property is that the affinity in the independent edge sub-problem for a given edge e is lower bounded by $\min(\theta_e, 0)$. This property is established in Appendix A.2

4.3 The Cycle Sub-problems

The affinities θ^c in the cycle sub-problem have two important properties at the optimal re-parameterization. These properties are established in Appendix A.3.

The first property is that the values of the affinities of the planar edges in a given cycle sub-problem are exactly opposite of the value of the affinity of the corresponding non-planar edge. Thus, the affinities in each cycle sub-problem are defined by one parameter. We denote the parameter associated with the c -th cycle as ψ^c . The affinity of the non-planar edge is $-\psi^c$ and the affinity of each planar edges is ψ^c . We use ψ to denote the concatenation of all ψ^c . Second that the affinity of the non-planar edges is strictly non-positive $\psi^c \geq 0$. Notice also that the optimal value of each cycle sub-problem is zero so the term corresponding to them in (3) can be removed.

¹ Notice that we impose a requirement on θ^{pmc} not a requirement on x^{pmc} .

² A 2-colorable segmentation is a partition of the set of superpixels into connected subsets (segments) such that every segment can be given a different color from all adjacent segments, using only two colors overall.

The second property is that the sum of the affinities of the copies of a given non-planar edge \hat{e} across cycle sub-problems can be no less than $\theta_{\hat{e}}^{\text{NP}}$. We use a matrix W to define this constraint. We index W by \hat{e}, c where $W_{\hat{e}c} = 1$ indicates that non-planar edge \hat{e} is present in the c -th cycle sub-problem. The constraint is written formally below.

$$-W\psi \geq \theta^{\text{NP}} \quad (6)$$

4.4 Value of a Re-parameterization

We now discuss the value of a given re-parameterization. Since the optimal values of the cycle sub-problems and the planar multicut sub-problem are zero, the value of a given re-parameterization is the optimal value of the independent edge sub-problem. If we re-write the re-parameterization condition in (4), we notice that $\theta^{ie} = \theta - \theta^{pmc} - \sum_c \theta^c$. We now discuss the mapping of ψ values to $\sum_c \theta^c$. This is done via a matrix operation $Y\psi$. Matrix Y is indexed by e, c where $Y_{ec} = 1$ indicates that planar edge e is present in cycle c . We write the value of a re-parameterization that obeys all of the constraints discussed in this section below.

$$\min_{X^{ie}} (\theta^{ie})^t X^{ie} = 1^T \theta^{ie} = 1^T (\theta^{\text{P}} - \theta^{pmc} - Y\psi) + 1^T (\psi + \theta^{\text{NP}}) . \quad (7)$$

The term $1^T (\theta^{\text{P}} - \theta^{pmc} - Y\psi)$ corresponds to the value of the contribution to the objective from planar edges. The term $1^T (\psi + \theta^{\text{NP}})$ corresponds to the value of the contribution to the objective from non-planar edges. We now write the constraint that each planar edge have non-positive affinity in the independent edge sub-problem.

$$\theta^{\text{P}} - \theta^{pmc} - Y\psi \leq 0 \quad (8)$$

We use (6) to ensure that each non-planar edge has non-positive affinity in the independent edge sub-problem. We write the constraint that the affinity in the single edge sub-problem for a given edge e is lower bounded by $\min(\theta_e, 0)$ as follows.

$$\min([0, \theta^{\text{P}}]) \leq \theta^{\text{P}} - \theta^{pmc} - Y\psi \quad (9)$$

In summary, the lower bound of Lagrangian decomposition is the optimal value of the linear program below.

$$\max_{\theta^{pmc}, \psi} 1^T (\theta^{\text{P}} - \theta^{pmc} - Y\psi) + 1^T (\psi + \theta^{\text{NP}}) \quad (10)$$

$$\text{subject to } Z\theta^{pmc} \geq 0 \quad (11)$$

$$-W\psi \geq \theta^{\text{NP}} \quad (12)$$

$$\theta^{\text{P}} - \theta^{pmc} - Y\psi \leq 0 \quad (13)$$

$$\min([0, -\theta^{\text{P}}]) \leq -\theta^{pmc} - Y\psi \quad (14)$$

$$\psi \geq 0 \quad (15)$$

4.5 Optimization of the Lower Bound

Since we can not enumerate all of the constraints of the form $Z\theta^{pmc} \geq 0$ or enumerate all the cycle sub-problems, we adopt a cutting plane approach in the dual. Given a set of cycle sub-problems and a set of 2-colorable partitions Z , we solve the LP in (10)–(15), then compute the most violated constraint corresponding to a 2-colorable partition. This is done by minimizing the following objective over the set of 2-colorable partitions which we denote as P_2 .

$$\min_{\bar{X} \in P_2} \theta^{pmc} \bar{X} \tag{16}$$

We can compute the minimum value 2-colorable segmentation \bar{X} in time $O(N^{3/2} \log(N))$ where N is the number of planar edges [14]. This computation is fast in practice (less than a second in all our applications). For every segment in \bar{X} , we add the constraint corresponding to the cut that separates this segment from the other segments. We repeat this until the minimum value 2-colorable partition has nearly 0 value. At any point in optimization, we can lower bound the optimal value of the planar multicut sub-problem by 3/2 times the value of the optimal 2-colorable partition as described in [17].

Next, we solve the dual of the Lagrangian decomposition, which provides a non-integral partition. We then find violated cycle inequalities in the non-integral partition and add the corresponding sub-problems to Lagrangian decomposition. The addition of cycle sub-problems corresponds to adding columns to Y and W . The dual of Lagrangian decomposition is written below. It is derived in the supplement. We use ϕ to simplify notation where $\phi = \min(0, -\theta^P)$.

$$\min_{\gamma \geq 0, \omega \geq 0, \delta \geq 0, \beta \geq 0} \begin{aligned} & 1^T(\theta^P + \phi + \theta^{NP}) - \phi^T Z^T \gamma \\ & - (\theta^{NP})^T \omega + ((-\theta^P)^T - \phi^T) \beta \end{aligned} \tag{17}$$

$$\text{subject to } (Z^T \gamma + \beta) = 1 + \delta \tag{18}$$

$$W^T \omega \geq Y^T \beta - Y^T 1 + 1 - Y^T \delta \tag{19}$$

The vector $Z^T \gamma$ defines a non-integral partition. The vector ω determines which non-planar edges are cut. Here $\omega_{\hat{e}} = 1$ indicates that edge \hat{e} is uncut and $\omega_{\hat{e}} = 0$ indicates that edge \hat{e} is cut (not a typo). Violated cycle inequalities are then found in this partition and the corresponding cycle sub-problems are added to Lagrangian decomposition. Violated cycle inequalities only correspond to cycles involving a non-planar edge. Violated cycle inequalities can be found using Dijkstra’s algorithm to find a path between the superpixels connected by any non-planar edge \hat{e} for which the fractional number of cut edges is minimal. If the fractional number of cut edges on this path is less than $1 - \omega_{\hat{e}}$, we add the cycle sub-problem containing this path to the Lagrangian decomposition. The iteration terminates when no violated cycle inequalities are found in $Z^t \gamma$.

4.6 Constructing a Segmentation

We now consider two approaches for converting of $Z^T\gamma$ to an integral partition. The first approach, which is called simple rounding, is based on rounding $Z^T\gamma$ to a partition. We cut every edge e such that $(Z^T\gamma)_e \geq T$ where T is a parameter. We uncut all cut edges within connected components. To select the optimal T we can try all unique values of $Z^T\gamma$ or a fixed number of values uniformly spaced over the interval $[0, 1]$. Notice that $Z^T\gamma \geq T$ defines which non-planar edges are cut.

For a more principled approach, we solve an integer program that constructs a partition using a weighted sum of the 2-colorable segmentations with non-zero γ . In practice, the number of 2-colorable segmentations with non-zero γ is small (tens in our applications). We now discuss the integer program explicitly.

We use \ddot{X} to describe the integral partition that we are constructing. We use a vector $\check{\gamma}$ to describe a fractional multicut. We index $\check{\gamma}$ with r where $\check{\gamma}_r$ defines how much of the r -th 2-colorable partition is included in X . Vector $\check{\gamma}$ has only indices corresponding to non-zero indices of γ . We use $\check{\omega}$ to define which of the non-planar edges are uncut. We use $\check{\omega}$ to penalize non-planar edges being uncut.

Our ILP is built on the following constraints. We enforce that a planar edge with negative affinity e can only be cut if the multicut defined by $\check{\gamma}$ cuts this edge at least once. We enforce that a planar edge with positive affinity e can only be uncut if, for every partition r , $Z_{r\hat{e}}\check{\gamma}_r = 0$. We enforce that a non-planar edge \hat{e} can only be cut in X if the cycle inequalities from the primal LP are satisfied. We write this optimization as an ILP below. The solution of this ILP does not make up a substantial portion of the computation time in practice.

$$\min_{\ddot{X}, \check{\omega}, \check{\gamma}} (\theta^P)^T \ddot{X}_e - (\theta^{NP})^T \check{\omega} + 1^T \theta^{NP} \tag{20}$$

$$\text{s.t. } \ddot{X}_e \leq \sum_{r \text{ s.t. } Z(r,e)=1} \check{\gamma}_r \quad \forall e \text{ s.t. } \theta_e < 0 \tag{21}$$

$$\check{\gamma}_c \leq \ddot{X}_e \quad \forall r \text{ s.t. } Z(r,e) = 1; \quad \forall e \text{ s.t. } \theta_e \geq 0 \tag{22}$$

$$\check{\omega}_{\hat{e}} \geq 1 - \sum_{r \text{ s.t. } [\exists e \in c \quad Z(r,e)=1]} \check{\gamma}_e; \quad \forall c \text{ s.t. } W(\hat{e}, c) = 1 \tag{23}$$

$$\ddot{X}_e \in \{0, 1\} \quad \forall e \tag{24}$$

$$\check{\omega}_{\hat{e}} \in \{0, 1\} \quad \forall \hat{e} \tag{25}$$

$$\check{\gamma}_r \in [0, 1] \quad \forall r \tag{26}$$

Once \ddot{X} is constructed, cut edges within connected components are uncut and the partition is returned. In practice, the objective value of the segmentation obtained by solving the ILP is nearly the same as the lower bound (cf. Section 5.2). The entire Lagrangian decomposition algorithm is detailed in Algorithm 1.

Algorithm 1. Semi-Planar Multicuts

```

Initialize  $Z=\[], W=\[], Y=\[];$ 
while true do
  while true do
    Solve LP in (10)
     $\bar{x} \leftarrow \arg \min_{x^{2c}} \sum_e -\lambda_e x_e$ 
     $Z \leftarrow [Z \cup \bar{x}]$ 
    if  $\sum_e -\lambda_e \bar{x}_e \sim 0$  then
      break
    end
  end
   $\gamma \leftarrow$  solve LP in (17)
  for  $\hat{e}$  do
    Use Dijkstra's algorithm to find a path between the superpixels
    connected by edge  $\hat{e}$  for which the fractional number of cut edges is
    minimal. If this number is less than  $1 - \omega_{\hat{e}}$ , add a cycle sub-problem
    corresponding to this path to the Lagrangian decomposition.
  end
  end
   $\gamma \leftarrow$  solve LP in (17)
  round  $Z^T \gamma$  using ILP in (20) or apply simple rounding on  $Z^T \gamma$  to produce a
  final segmentation  $X$ .

```

5 Application

5.1 Effectiveness

To demonstrate the effectiveness of Algorithm 1, we implement a workflow for inter-active image segmentation that requires only trivial user input for both the joining and cutting of segments. In the beginning, the user is presented with an optimal multicut of a superpixel segmentation. This multicut is computed using only affinities between adjacent superpixels. Subsequently, any combination of two types of input is accepted to correct errors in the segmentation: 1. Scribbles on the image to indicate pixels that belong to the same segment, 2. Pairs of points in the image to indicate pixels that belong to different segments.

The evidence provided by scribbles is incorporated into the multicut problem by constraining variables to 0. The evidence provided by pairs of points is incorporated as a non-local affinity with sufficiently large negative weight. Alternatively, these affinities could also be incorporated as hard constraints. Hard constraints are processed similarly to non-local affinities with finite negative weight, except that the corresponding variable is fixed to 1 and the constant contribution to the objective function, $-\infty$, is ignored. Either way, every pair of points excludes from the solution set all segmentations in which the points are in different segments.

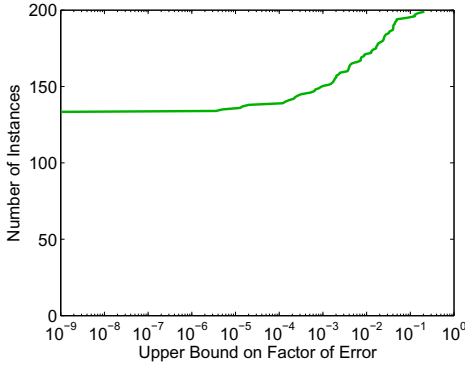


Fig. 3. The fraction by which the value of a segmentation output by Alg. 1 deviates from the global optimum (first axis) is small for a large number of instances (second axis). 132 of 200 instances are solved to optimality. For 170 of the 200 instances, the error is less than one percent. Depicted is the upper bound on the error provided by the gap between the value of a segmentation and the LP lower bound.




Fig. 4. Alg. 1 terminates after 4.73 s in the mean, 4.16 s in the median and 28.05 s in the worst case of the 200 instances. This is significantly faster than the an exact algorithm for general graphs [3] that takes 51.32 s in the mean, 7.79 s in the median and 3307.42 s (55 minutes) in the worst case.

At any point in time, the user can request that the augmented problem with additional constraints be solved to obtain an updated segmentation. There is no requirement for the user to focus on one or the other type of errors, nor to prefer one over the other type of input. Results that confirm the effectiveness of the algorithm and the consistency of our implementation are depicted in Fig. 5.

5.2 Quality of Bounds and Runtime

To examine the bounds provided by Alg. 1 and to compare its runtime to that of the exact algorithm for general graphs [3], we set up segmentation problems with negative non-local affinities for the 200 test images of the BSD500 benchmark. To ensure that experiments are unbiased, i.e. not subjective to specific user input, we construct these problems automatically as follows.

Starting with the same superpixels and affinities between adjacent superpixels as in [17], we search the proximity of every superpixel for the superpixel whose mean color differs maximally from that of the first superpixel. For every image, we add the 50 strongest non-local affinities. Both, our implementation of Alg. 1 and the optimized C++ implementation [1] of [3], use IBM ILOG Cplex as an LP and ILP solver, respectively. We use Blossom V [11] for computing the optimal 2-colorable partitions.

The fraction by which the value of a segmentation output by Alg. 1 deviates from the global optimum is depicted on the first axis in Fig. 3. More precisely, it is the upper bound on this deviation provided by the gap between the value

of a segmentation and the LP lower bound. This gap is zero for 132 of the 200 benchmark images, indicating that these problems are solved to optimality by Alg. 1. For 170 of the 200 instances, the error is less than one percent.

The distribution of runtimes for the 200 problems is depicted in Fig. 4. Unlike the exact algorithm for general graphs whose worst-case runtime is prohibitive, Alg. 1 is fast enough for inter-active applications.

6 Conclusions

We motivated the use of negative affinities between non-adjacent superpixels in the multicut formulation of the image segmentation problem. For the resulting NP-hard combinatorial optimization problem, we proposed a relaxation by Lagrangian decomposition and a constrained set of re-parameterizations for which we can optimize exactly and efficiently by solving an LP. We implemented a cutting plane approach in the dual, along with an algorithm to construct segmentation from fractional solutions. This algorithms finds optimal to near optimal solutions of segmentation problems with non-local negative affinities for natural images. In our applications, it is more than 10 times faster on average than an exact algorithm based on integer programming.

A Appendix

A.1 Difficulty of Positive Non-local Terms

We now consider why positive non-local terms are difficult. Let v and w be superpixels connected by a positive non-local term whose corresponding edge is x_k . Let $S(v, w)$ the set of all rings such that if all edges in a particular ring are cut then the positive edge x_k must also be cut. The corresponding addition to the LP relaxation is written below.

$$\forall r \in S(p_1, p_2) \quad x_k + \sum_{e \in r} x_e \leq |r| . \quad (27)$$

These constraints that appear elegantly in the form of separator inequalities in [13] are remarkably fragile and almost always result in fractional solutions to even trivial multicut problems. Moreover, these fractional solutions are uninformative with regards to what the optimal solution is. As an example, consider five superpixels s_a, s_b, s_c, s_d, s_e in the plane. Let s_a be surrounded on all sides by s_b, s_c, s_d , let s_e be a neighbor of s_b, s_c, s_d and disconnected from s_a . Also, let s_a , and s_e be paired by an infinitely strong positive term. The remaining terms are described as follows:

$$\theta_{s_a, s_b} = \theta_{s_a, s_c} = \theta_{s_a, s_d} = \infty \quad (28)$$

$$\theta_{s_e, s_b} = \theta_{s_e, s_c} = \theta_{s_e, s_d} = -1 \quad (29)$$

$$\theta_{s_b, s_c} = \theta_{s_c, s_d} = \theta_{s_b, s_d} = \infty \quad (30)$$

Clearly, the optimal solution is to put each superpixel in the same connected component, and the cut cost is zero. However, the optimal fractional solution is to cut each repulsive edge with value $-2/3$, and objective value is -2 .




Fig. 5. To demonstrate the effectiveness of Algorithm 1, we implement a workflow for inter-active image segmentation that requires only trivial user input for both the joining and cutting of segments. Presented with an initial and typically imperfect segmentation, the user can scribble on the image (red) to indicate pixels that belong to the same segment and click on pairs of points (green) that belong to different segments. The latter are incorporated as non-local negative affinities.

A.2 Lower Bound on Affinities in the Independent Edge Sub-problem

We now establish that $\min([0, -\theta^P]) \leq -\theta^{pmc} - Y\psi$. To do this, we rely on a theorem established in [16] which states that given affinities θ for an instance of the planar multicut problem, the affinities in the planar multicut sub-problem $-\theta^{pmc}$ are lower bounded by $\min([-\theta, 0])$. In [16], no non-planar edges are considered.

Given optimal parameters ψ , we have a description for a planar multicut problem instance. Thus, the parameters of the independent edge sub-problem and the planar multicut sub-problem satisfy the following inequality.

$$\min([- \theta^P + Y\psi, 0]) \leq -\theta^{pmc} . \tag{31}$$

Suppose that, at the optimal re-parameterization, the constraint

$$\min([0, -\theta^P]) \leq -\theta^{pmc} - Y\psi \tag{32}$$

is not satisfied. Thus, there must be a planar edge e such that the constraint is unsatisfied for. In this case, there must exist a cycle sub-problem c containing e such that $\psi_e^c \geq 0$. We now define α and σ as follows.

$$\alpha = -\min([- \theta_e, 0]) - \theta_e^{pmc} - Y\psi \tag{33}$$

$$\sigma = \max([0, \psi^c + \alpha]) - \psi^c . \tag{34}$$

Now set $\psi^c = \max([0, \psi^c + \alpha])$. For each edge $f \neq e$ in c set $\theta_f^{pmc} = \theta_f^{pmc} - \sigma$. Notice that the bound remains constant, no additional edges are unsatisfied and the constraint of the form $\min([0, -\theta_e^P]) \leq -\theta_e^{pmc} - (Y\psi)_e$ is $-\sigma$ closer to being satisfied. This process can be repeated until the constraint $\min([0, -\theta^P]) \leq -\theta^{pmc} - Y\psi$ is satisfied for all edges.

A.3 Properties of Cycle Sub-problems

We now consider the constraints (6) on the parameters of the cycle sub-problems discussed in Section 4.3. We establish that they do not loosen the lower bound of Lagrangian decomposition as follows. Assume that we have maximized the lower bound of Lagrangian decomposition and that these constraints are violated. We then alter the re-parameterization without loosening the bound so as to enforce the relevant constraints for a single cycle. This process can be repeated for every cycle in which the constraints are violated.

Consider a given cycle sub-problem c and no structure in the affinities. Suppose that there are zero or two or more edges with non-positive affinity. If this is the case, move all negative affinity to the independent edge sub-problem and all positive affinity to the planar multicut sub-problem. If the non-planar edge has positive affinity, move the positive affinity to other copies of that edge in cycle sub-problems and the independent edge sub-problem. Preserve that the affinity in the independent edge sub-problem for each edge is non-positive. Finally, remove cycle sub-problem c from the decomposition. Notice that these alterations to the re-parameterization do not loosen the lower bound of the decomposition.

Now consider the case where there is one negative affinity edge in cycle sub-problem c . Set the affinity of the edge with negative affinity to the opposite of the affinity of the edge with the smallest non-negative affinity. We denote the value of the smallest non-negative affinity as ϕ^c . Now, set all affinities on edges with non-negative affinity to ϕ^c . This last step is done because no optimal configuration of cycle sub-problem c cuts a non-negative affinity edge e unless

that edge has affinity equal to ϕ^c . We send the “excess” positive affinity from cycle sub-problem c to the planar multicut sub-problem, which does not loosen the lower bound of the decomposition.

Since the addition of cycle sub-problems induces non-planar edges to be uncut or induces planar edges along a path connecting the ends of a non-planar edge to be cut, it is clear that the affinity associated with the non-planar edge in each given cycle sub-problem is non-positive.

References

1. Andres, B.: <https://github.com/bjoern-andres/graph>
2. Andres, B., Kappes, J.H., Beier, T., Köthe, U., Hamprecht, F.A.: Probabilistic image segmentation with closedness constraints. In: ICCV (2011)
3. Andres, B., Kroeger, T., Briggman, K.L., Denk, W., Korogod, N., Knott, G., Koethe, U., Hamprecht, F.A.: Globally optimal closed-surface segmentation for connectomics. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part III. LNCS, vol. 7574, pp. 778–791. Springer, Heidelberg (2012)
4. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. TPAMI 33(5), 898–916 (2011)
5. Bachrach, Y., Kohli, P., Kolmogorov, V., Zadimoghaddam, M.: Optimal coalition structures in graph games. arXiv ePrint (2011)
6. Bagon, S., Galun, M.: Large scale correlation clustering optimization. arXiv ePrint, abs/1112.2903 (2011)
7. Bansal, N., Blum, A., Chawla, S.: Correlation clustering. Machine Learning 56, 89–113 (2004)
8. Chopra, S., Rao, M.R.: The partition problem. Math. Program. 59, 87–115 (1993)
9. Kappes, J.H., Speth, M., Andres, B., Reinelt, G., Schnörr, C.: Globally optimal image partitioning by multicuts. In: Boykov, Y., Kahl, F., Lempitsky, V., Schmidt, F.R. (eds.) EMMCVPR 2011. LNCS, vol. 6819, pp. 31–44. Springer, Heidelberg (2011)
10. Kim, S., Nowozin, S., Kohli, P., Yoo, C.D.: Higher-order correlation clustering for image segmentation. In: NIPS (2011)
11. Kolmogorov, V.: Blossom V: a new implementation of a minimum cost perfect matching algorithm. Mathematical Programming Computation 1(1), 43–67 (2009)
12. Nowozin, S., Jegelka, S.: Solution stability in linear programming relaxations: graph partitioning and unsupervised learning. In: ICML, pp. 769–776 (2009)
13. Nowozin, S., Lampert, C.H.: Global interactions in random field models: A potential function ensuring connectedness. SIAM J. Img. Sci. 3(4), 1048–1074 (2010)
14. Shih, W.-K., Wu, S., Kuo, Y.S.: Unifying maximum cut and minimum cut of a planar graph. IEEE Trans. Comput. 39(5), 694–697 (1990)
15. Vitaladevuni, S.N.P., Basri, R.: Co-clustering of image segments using convex optimization applied to EM neuronal reconstruction. In: CVPR (2010)
16. Yarkony, J.: MAP inference in Planar Markov Random Fields with Applications to Computer Vision. PhD thesis, University of California, Irvine (2012)
17. Yarkony, J., Ihler, A., Fowlkes, C.C.: Fast planar correlation clustering for image segmentation. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012, Part VI. LNCS, vol. 7577, pp. 568–581. Springer, Heidelberg (2012)