

Random Walks for Image Cosegmentation

Maxwell D. Collins¹ Jia Xu¹ Leo Grady² Vikas Singh¹
¹University of Wisconsin-Madison ²Siemens Corporate Research
Madison, WI Princeton, NJ

{mcollins, jiaxu}@cs.wisc.edu, leo.grady@siemens.com, vsingh@biostat.wisc.edu

June 16, 2012

Abstract

We recast the *Cosegmentation* problem using Random Walker (RW) segmentation as the core segmentation algorithm, rather than the traditional MRF approach adopted in the literature so far. Our formulation is similar to previous approaches in the sense that it also permits Cosegmentation constraints (which impose consistency between the extracted objects from ≥ 2 images) using a nonparametric model. However, several previous nonparametric cosegmentation methods have the serious limitation that they require adding one auxiliary node (or variable) for every pair of pixels that are similar (which effectively limits such methods to describing only those objects that have high entropy appearance models). In contrast, our proposed model completely eliminates this restrictive dependence – the resulting improvements are quite significant. Our model further allows an optimization scheme exploiting quasiconvexity for model-based segmentation with no dependence on the scale of the segmented foreground. Finally, we show that the optimization can be expressed in terms of linear algebra operations on sparse matrices which are easily mapped to GPU architecture. We provide a highly specialized CUDA library for Cosegmentation exploiting this special structure, and report experimental results showing these advantages.

1 Introduction


The problem of Cosegmentation [1], has attracted much attention from the community in the last few years [2–5]. The basic goal in Cosegmentation is to segment a common salient foreground object from two or more images, as shown in Fig. 1. Here, consistency between the (extracted) object regions is accomplished by imposing a *global* constraint which penalizes variations between the objects’ respective histograms or appearance models. The idea has been adopted for obtaining concise measures of image similarity [1], discovering common appearance patterns in image sets [6], medical imaging [7], and building 3D models from community photo collections [2, 8]. Motivated by the spectrum of these applications, some recent papers have sought to better understand the optimization-specific aspects of this problem – in particular, issues such as sub-modularity [1], linear programming relaxations [4], dual-decomposition



Figure 1: A representative application for cosegmentation.

based strategies [5], network flow constructions [7], and maximum-margin inspired interpretations [9]. Most of these works provide good insights on cosegmentation, but *only* in the context of the traditional Markov Random Field (MRF) based segmentation objective (referred to as *graph-cuts* [10]). This may be partly because the first work on Cosegmentation [1] presented means for including global constraints within segmentation but was designed *specifically* for the MRF function. The present paper complements this body of research, and provides an end-to-end analysis of the Cosegmentation problem in terms of the Random Walker segmentation function [11] – these results show that in many cases, well-known advantages of the Random Walker model extend nicely to the Cosegmentation setting as well.

A Toy example. An important aspect of our formulation is that it is possible to employ a *nonparametric* appearance model for *arbitrary* distributions but *without* incurring rather substantial additional computational costs. When there are significant regions of homogeneity in the foreground (inline figure), we clearly want a distribution which has a corresponding peak. In Fig. 1, the distribution should capture the fact that the bear is “brown and furry”, and *not* try to differentiate one patch of fur from another across multiple images. To illustrate why this point is relevant, let us analyze the overhead of some existing methods for cosegmentation by considering a simple toy example (see inline image pair) where we should identify the common blue circle (in distinct backgrounds). Several approaches for cosegmentation with a nonparametric model require that an auxiliary node (or variable) be introduced into the graph whenever two pixels share the *same* bin [4, 7] (i.e., the two pixels are similar). While the segmentation aspect for the blue circles by itself is easy, the cost of introducing an auxiliary node for perceptually similar pixels can grow *very* quickly – counting just the blue foreground pixels for a 196×196 image pair, one must introduce 42 million additional variables, and the associated cost is infeasible even for moderately sized images. As a result, these previous models are limited to feasibly cosegmenting only those image pairs that have a relatively high entropy distribution (i.e., each bin is shared by only a few pixels). Our formulation has *no such limitation*, since auxiliary nodes are *not* needed to perform the optimization. Consequently, it is possible to perform cosegmentation in general settings where the target foreground is summarized with an arbitrary appearance model (color, texture), but with no associated restriction on its entropy. Further, several nonparametric coseg-



mentation models (except [9]) are somewhat limited to segmenting foregrounds which are at the same scale within each image. Our method compares histograms *independent* of scale, and the objective is shown to be *quasiconvex*. This is leveraged to develop an optimization scheme which can solve this *nonconvex* problem with optimality guarantees. Finally, all optimization steps in the core method can be expressed as sparse linear algebra operations, which makes GPU implementations possible.

Related Work. The initial model for Cosegmentation in [1] provided means for including global constraints to enforce consistency among the two foreground histograms in *addition* to the MRF segmentation terms for each image. The objective function incorporating these ideas was expressed as follows

$$E_{\text{coseg}} = \text{MRF}_{\text{image 1}} + \text{MRF}_{\text{image 2}} + \lambda E^{\text{global}}(h_1, h_2), \quad (1)$$

where $E^{\text{global}}(\cdot, \cdot)$ was assumed to penalize the ℓ_1 variation between each h_1 and h_2 (the two foreground histograms obtained *post* segmentation). The appearance model for the histogram was assumed to be generative (i.e., Gaussian), and a novel scheme called trust region graph cuts was presented for optimizing the resulting form of (1). Subsequently, [4] argued in favor of using an ℓ_2^2 -distance for $E^{\text{global}}(\cdot, \cdot)$ whereas [7] developed a reward based model. A scale-free method is presented in [12], which biases the segmentation towards histograms belonging to a low-rank subspace. Batra et al. [2] suggested exploiting user interaction if available for cosegmentation (again, using MRF terms) and [3, 9] have adopted a clustering based approach to the problem. Recently, [5] compared several existing MRF-based models, and presented a new posterior-maximizing scheme which was solved using dual decomposition. Other recent ideas include modulating the graph-cuts objective *a priori* by finding similar patterns across images [3], generating multiple segmentations of each image and identifying which segmentation pair is similar [13], and identifying salient regions followed by a filtering step to leave out distinct regions *not* shared across images [14]. One reason for these varied strategies for the problem is that when a histogram difference term is added to the segmentation, the resultant objective is no longer *submodular* (therefore, not easy to optimize). So, the focus has been on improved approximate algorithms for different choices of E^{global} .

A commonality among these existing works has been the preference for MRF (i.e., graph-cuts) based terms for segmentation. Part of the reason is that combinatorial methods such as graph-cuts are extensively used in vision, and are known to be efficient. On the other hand, graph-partitioning methods such as Random Walker [11] also work well for image segmentation and are widely used. Our formalization here suggests that it is also well suited for the Cosegmentation problem and offers efficiency benefits (e.g., issues identified in the blue circles example) in the nonparametric setting.

The **contributions** of this paper are: (1) We derive a cosegmentation model with the Random Walker segmentation at its core. The model finally reduces to a Box-QP problem (convex program with box constraints). Based on this structure, we propose a specialized (and efficient) gradient projection based procedure which finds a global real-valued optimum of the model (which preserves many advantages of Random Walker [11]). (2) Our model allows for a nonparametric representation of the foregrounds (e.g., using distributions over texture words), but one which permits *any* distribution of features without incurring additional computational costs. This provides a substantial

advantage over the existing nonparametric cosegmentation approaches which are limited only to regions described by a high entropy model (i.e., object features must be spread evenly across bins). **(3)** We extend this model to a scale-independent penalty. This paper presents a novel optimization method for a class of objectives based on quasiconvex functions. We prove correctness, and demonstrate it for model-based image segmentation. These theoretical results are of independent interest. **(4)** Our optimization consists of linear algebra operations (BLAS Level 1, 2) on sparse matrices. Consequently, the algorithm is easy to implement on a GPU architecture. We give a specialized open-source CUDA library for Cosegmentation.

2 Random Walker and its properties

The Random Walker segmentation algorithm has been studied extensively in the computer vision literature. Essentially, the method simulates a random walk from each pixel in the image to a set of user specified seed points where the walk is biased by image intensity gradients. The eventual assignment of pixels to foreground or background is determined by whether the pixel-specific walk reaches a foreground (or background) seed first. Observe a subtle yet important property of how the Random Walker model is specified, and what the solution actually denotes. Because of direct analogues in circuit theory and physics, the formalization, even in its original form, seeks a solution in reals (not integers). What is eventually solved is therefore *not* a relaxation because the variables have a clear probabilistic meaning. As a result, thresholding these probabilities at 0.5 is statistically sound; conceptually, this is different from solving a binary linear program in reals and recovering a $\{0, 1\}$ solution by rounding. In practice, Random Walker is optimized by recasting segmentation as the solution to a combinatorial Dirichlet problem. Random Walker derived segmentations offer some benefits with respect to boundary length regularization, number of seeds, metrication errors, and shrinking bias [15].

3 Random Walker for Cosegmentation

We begin our presentation by rewriting the Random Walker algorithm for a single image as a quadratic minimization problem (also see [16]). As is common, we assume a 4-connected neighborhood over the image, weighted according to a Gaussian function of normalized Euclidean distances between pixel intensities, $w_{ij} = \exp(-\beta\|p_i - p_j\|)$. The Laplacian L of the graph is then

$$L_{ij} = \begin{cases} \sum_k w_{ik} & \text{if } i = j \\ -w_{ij} & \text{if } i \neq j \text{ and } (i, j) \in \text{neighborhood graph} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The Laplacian is diagonally dominant and so $L \succeq 0$; we can derive the following convex quadratic program,

$$\min_x x^T L x \quad \text{subject to} \quad x^{(s)} = m^{(s)}, \quad (3)$$

where $x^{(s)}$ are the values for certain seed pixels, and $m^{(s)}$ is the known value of those seeds (i.e., foreground or background). Each component of the solution x^* will then be a pixel's probability of being assigned to the foreground. To output a $\{0, 1\}$ segmentation, we may threshold x^* at $\frac{1}{2}$, to obtain a hard $x \in \{0, 1\}^n$ segmentation which matches the solution from [11].

Pre-processing. Cosegmentation methods [7] use a pre-processing procedure to determine inter- (and intra-) image pixel similarity. This is generated by tessellating the RGB color space (i.e., pixel distribution) into clusters or by using SIFT (or color pattern models, edge-profiles, textures etc) based correspondence methods, see [17]. We can derive a matrix H such that

$$H_i^{kj} = \begin{cases} 1 & \text{if pixel } j \text{ is in histogram bin } k \text{ in image } i \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Here, pixels are assigned to the same bin if they are similar. With an appropriate H , the global term E^{global} from (1) requires that at the level of individual histogram bins k , the algorithm assign approximately the same number of pixels to each foreground region (the objective incurs a penalty proportional to this difference). This ensures that the appearance models of the two foregrounds based on the features of interest are similar, and has been used very successfully in object recognition [18]. Observe that this difference only serves as a *regularizer* for the main segmentation task, and does not drive it on its own. This is relevant because as with any global measure, such models (and the measurement of their variations) may not be perfect. But existing literature suggests that when derived from good context-based features [18], such appearance model based differences provide a meaningful global bias for Cosegmentation [2].

Cosegmentation for 2+ images. Given a segmentation for n pixels, $x \in \{0, 1\}^n$, one may use the H matrix from (4) to write the histogram of only the foreground pixels as $h = Hx$. The expression gives the form of constraints needed for Cosegmentation. Let L_1, \dots, L_m be the Laplacian matrices of graphs constructed using each of the images, and H_1, \dots, H_m be the histogram assignment matrices from (4), with the property that $H_i^{kj} = H_i^{kj'} = 1$ for pixels j and j' if and only if j and j' are similar. Here, if one uses SIFT matches, then the matrix entry may reflect the confidence of the match. Now, we seek to segment the two images simultaneously, under the constraint that their histograms match. For this purpose, it suffices to consider the following optimization model

$$\begin{aligned} \min_{x_i, h_i, \bar{h}} \quad & \sum_i x_i^T L_i x_i + \lambda \|h_i - \bar{h}\|_2^2 \\ \text{s.t.} \quad & x_i \in [0, 1]^{n_i} \\ & x_i^{(s)} = m_i^{(s)} \quad i = 1 \dots m \\ & H_i x_i = h_i \end{aligned} \quad (5)$$

The second term in the objective above corresponds to $E^{\text{global}}(h_1, h_2)$ in (1), and the last constraint sets up the foreground histograms, h_i using H_i and (bin k in H_1 corresponds to bin k in H_2, \dots, H_m which makes a direct comparison between histograms possible). Instead of comparing the histograms to each other, we compare

them to a common *global* histogram \bar{h} which at the optimum will be the centroid of the h_i 's. The resulting inter-image matching penalty will then be the trace of the covariance between foreground histograms across the image set. This model additionally extends to multiple labels (i.e., multiple objects) by adding additional columns to the optimization variables identically to [11]. The resulting problem can be easily decomposed into separate segmentations for each object class. Existing cosegmentation methods, on the other hand, mostly tackle figure-ground labeling.

Each Laplacian matrix L_i is positive semidefinite, so along with the histogram distances the objective function is convex. Further, the feasible region is the intersection of bound constraints and linear equalities. We directly have:

Theorem 3.1. For $\lambda \geq 0$, (5) is a convex problem.

3.1 Deriving an equivalent Box-QP

The model in (5) can already be solved using widely available convex programming methods, and provides the desired solution to the Cosegmentation problem using the Random Walker segmentation function. Next, we will derive an equivalent model (but with a much nicer structure) that will allow the design of specialized solvers and thereby lead to far more efficient algorithms.

Consider the left hand side of the equality constraint on each h_i , substituted into the objective function with a penalty $\|h_1 - h_2\|$. Further, let us choose bounds to limit x to the unit box as well as suitably enforce the seed constraints. This process gives a quadratic problem of the form

$$\begin{aligned} \min_{x_1, x_2} \quad & \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}^T \begin{bmatrix} L_1 + \lambda H_1^T H_1 & -\lambda H_1^T H_2 \\ -\lambda H_2^T H_1 & L_2 + \lambda H_2^T H_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\ \text{s.t.} \quad & l_i \leq x_i \leq u_i \quad x_i \text{ is of size } [0, 1]^{n_i} \quad i = 1, 2, \end{aligned} \quad (6)$$

where the 2-tuple (l_i, u_i) is $(1, 1)$ for foreground seeds, $(0, 0)$ for background seeds, and $(0, 1)$ otherwise. For $m > 2$ images we optimize over x_1, \dots, x_m, \bar{h} with quadratic objective matrix

$$\begin{bmatrix} L_1 + \lambda H_1^T H_1 & & & -\lambda H_1 \\ & \ddots & & \vdots \\ & & L_m + \lambda H_m^T H_m & -\lambda H_m \\ -\lambda H_1^T & \dots & -\lambda H_m^T & \lambda m I \end{bmatrix}$$

It can be verified that (6) is *equivalent* to (5). The difference is that it is now expressed as a bound-constrained quadratic problem (or *box-QP* due to the box constraints). Like (5), the model in (6) also permits general purpose convex programming methods. However, we can design means to exploit its special structure since the model is *nearly* an unconstrained quadratic problem.

4 Scale-Free Cosegmentation

A limitation of previous cosegmentation methods is their sensitivity to the scale of the target object, since histogram-based priors are dependent on scale. For example, if an

otherwise identical object appears in the second image such that it occupies twice as many pixels as in the first image, then $h_2 = 2h_1$. Consequently, $\|h_2 - h_1\| > 0$, meaning that the larger scale is penalized in traditional formulations. We show here how our formulation may be made scale-invariant. Formally, our goal is to modify the cosegmentation term to satisfy

$$E(sh_i, \bar{h}) = E(h_i, \bar{h}) \quad \forall s \in \mathbb{R}_+. \quad (7)$$

This property may be satisfied by a normalization step,

$$E(h_i, \bar{h}) = \left\| \frac{h_i}{\|h_i\|} - \frac{\bar{h}}{\|\bar{h}\|} \right\|^2 = 2 - 2 \frac{h_i^T \bar{h}}{\|h_i\| \|\bar{h}\|}. \quad (8)$$

Here we further note that under minimization and choice of the parameter λ in the combined model this is equivalent to minimizing

$$- \frac{h_i^T \bar{h}}{\|h_i\| \|\bar{h}\|} = -\cos(\angle h_i \bar{h}) \quad (9)$$

Substituting these normalized histograms in (5) leads to a function that cannot be efficiently optimized. However, in the Random Walker setting, we *can* optimize this function when the model histogram \bar{h} is a fixed unit vector. The resulting problem is related to model-based segmentation, where we are imposing a known histogram distribution in segmenting images. For image i we solve the problem

$$\begin{aligned} \min_{x_i, h_i} \quad & x_i^T L_i x_i - \lambda \frac{h_i^T \bar{h}}{\|h_i\|} \\ \text{s.t.} \quad & h_i = H_i x_i \\ & l_i \leq x_i \leq u_i \end{aligned} \quad (10)$$

where we name

$$\hat{g}_{\bar{h}}(h) = - \frac{\bar{h}^T h}{\|h\|}. \quad (11)$$

In order to proceed with the minimization of our scale-invariant energy, we must first establish some properties of (11).

Definition 4.1 ([19]). We define a function f to be *quasiconvex* if its sublevel sets are convex subsets of the domain. Equivalently, for any x, x' in the domain of f and $\lambda \in [0, 1]$,

$$f((1 - \lambda)x + \lambda x') \leq \max\{f(x), f(x')\} \quad (12)$$

We call (12) the ‘‘Jensen’s Inequality for Quasiconvexity.’’

Theorem 4.2. $\hat{g}_{\bar{h}}$ as defined in (11) is quasiconvex.

Proof. Consider any h_1, h_2 wlog chosen to satisfy

$$\begin{aligned} \hat{g}_{\bar{h}}(h_2) &\leq \hat{g}_{\bar{h}}(h_1), \\ \frac{\bar{h}^T h_2}{\|h_2\|} &\geq \frac{\bar{h}^T h_1}{\|h_1\|}, \end{aligned}$$

Multiply by $(1 - \lambda)\|h_1\|\|h_2\| \geq 0$

$$(1 - \lambda)\bar{h}^T h_2 \|h_1\| \geq (1 - \lambda)\bar{h}^T h_1 \|h_2\|,$$

Add $\lambda\bar{h}^T h_1 \|h_1\|$,

$$\begin{aligned} \lambda\bar{h}^T h_1 \|h_1\| + (1 - \lambda)\bar{h}^T h_2 \|h_1\| &\geq \lambda\bar{h}^T h_1 \|h_1\| + (1 - \lambda)\bar{h}^T h_1 \|h_2\| \\ (\lambda\bar{h}^T h_1 + (1 - \lambda)\bar{h}^T h_2) \|h_1\| &\geq \bar{h}^T h_1 (\lambda\|h_1\| + (1 - \lambda)\|h_2\|) \end{aligned}$$

taking any $\lambda \in [0, 1]$.

Since all these vectors are in the nonnegative orthant, $\bar{h}^T h_1, \bar{h}^T h_2 \geq 0$, and $-\hat{g}_{\bar{h}}(h_2) \geq -\hat{g}_{\bar{h}}(h_1)$ this inequality is equivalent to:

$$\frac{\lambda\bar{h}^T h_1 + (1 - \lambda)\bar{h}^T h_2}{\lambda\|h_1\| + (1 - \lambda)\|h_2\|} \geq \frac{\bar{h}^T h_1}{\|h_1\|} \quad (13)$$

Using this expression with the triangle inequality to show Jensen's inequality for quasiconvexity (12) gives

$$\begin{aligned} \hat{g}_{\bar{h}}(\lambda h_1 + (1 - \lambda)h_2) &= -\frac{\lambda\bar{h}^T h_1 + (1 - \lambda)\bar{h}^T h_2}{\|\lambda h_1 + (1 - \lambda)h_2\|} \\ &\leq -\frac{\lambda\bar{h}^T h_1 + (1 - \lambda)\bar{h}^T h_2}{\lambda\|h_1\| + (1 - \lambda)\|h_2\|} \end{aligned}$$

Using (13),

$$\begin{aligned} &\leq -\frac{\bar{h}^T h_1}{\|h_1\|} \\ &= \hat{g}_{\bar{h}}(h_1) = \max\{\hat{g}_{\bar{h}}(h_1), \hat{g}_{\bar{h}}(h_2)\} \end{aligned}$$

□

Corollary 4.3. *The scale-free energy on x , $g_{\bar{h}}(x) = \hat{g}_{\bar{h}}(Hx)$ is quasiconvex for histogram assignment matrix H .*

Proof. This G represents the outer composition of E with an affine function. This operation preserves quasiconvexity (see [20]). □

The next section exploits these properties of $g_{\bar{h}}$ to solve the segmentation problem using this penalty.

Theorem 4.4. *$g_{\bar{h}}(x)$ is Lipschitz-smooth when $\|Hx\| > 0$.*

Proof. Let $\hat{\cdot} = \frac{\cdot}{\|\cdot\|}$. Assume wlog $\|h_1\| \geq \|h_2\|$.

$$\begin{aligned} \|h_1 - h_2\| &\geq \left\| \frac{\|h_2\|}{\|h_1\|} h_1 - h_2 \right\| \\ &= \|h_2\| \times \|\hat{h}_1 - \hat{h}_2\| \\ \frac{1}{\|h_2\|} \|h_1 - h_2\| &\geq \|\hat{h}_1 - \hat{h}_2\| \end{aligned} \quad (14)$$



Figure 2: Segmentation using the model of section 4 on a set of images from the iCoseg dataset with differences in scale. \bar{h} from an image in the same set was applied as a prior.

Thus for any function f which is L -smooth,

$$\begin{aligned} \|f(\hat{h}_1) - f(\hat{h}_2)\| &\leq L\|\hat{h}_1 - \hat{h}_2\| \\ &\leq \frac{L}{\|h_2\|}\|h_1 - h_2\| \end{aligned} \quad (15)$$

And if we lower-bound $\|h_1\| \geq \|h_2\| \geq C > 0$, $f(\cdot)$ is $\frac{L}{C}$ -smooth.

In our case, $g_{\bar{h}}(h)$ is an affine function of \hat{h} , and any affine function will be Lipschitz. \square

4.1 Nonconvex Sum Minimization

For the following, we consider the setting of minimizing $h = f + g$, such that f is convex, g is quasiconvex and both are bounded below. Note that under these conditions, h is not necessarily quasiconvex and may have multiple local minima. Nonetheless, our method proposed below can optimize our segmentation objective with $f(x) = x^T Lx$ and g as defined in Corollary 4.3. Let $x_f^* = \operatorname{argmin}_x f(x)$, similarly define x_g^* and x_h^* .

Theorem 4.5. *Define*

$$\mathcal{P}(\alpha) = \begin{pmatrix} \operatorname{argmin}_{x \in X} f(x) \\ \text{s.t. } g(x) \leq \alpha \end{pmatrix} \quad (16)$$

For x any solution to $\mathcal{P}(g(x_h^))$, $h(x) = h(x_h^*)$.*

Proof. Since x is feasible for $\mathcal{P}(g(x_h^*))$, $g(x) \leq g(x_h^*)$. With x, x_h^* both feasible for the problem, x is a solution iff $f(x) \leq f(x_h^*)$.

We can add these inequalities to show $h(x) \leq h(x_h^*)$, which by definition of x_h^* gives equality. \square

Theorem 4.6. $\mathcal{P}(\alpha)$ has no solutions for $\alpha < g(x_g^*)$, and x is a solution to $\mathcal{P}(g(x_f^*))$ iff x is a solution $\forall \alpha > g(x_f^*)$.

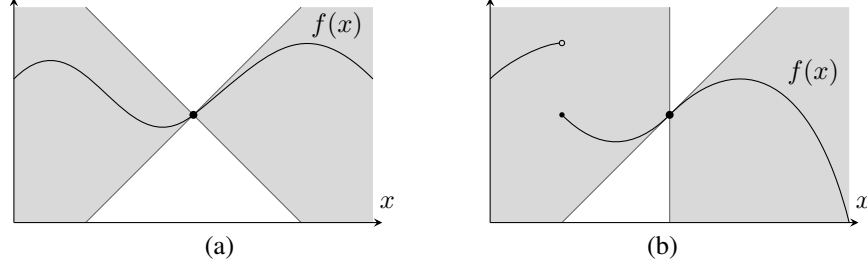


Figure 3: Given a pair $(x, f(x))$ for a L -Lipschitz function f , the Lipschitz condition guarantees that the function will always lie between two lines of slope L and $-L$ through this point and in the shaded region shown in (a). The *one-sided Lipschitz condition* in Definition 4.7 only bounds the rate of increase, and the function must lie in the shaded region shown in (b).

Proof. For $\alpha < g(x_g^*)$ the feasible set of $\mathcal{P}(\alpha)$ is the empty set, by the definition of x_g^* .

Take any $\alpha \geq g(x_f^*)$. The feasible region of $\mathcal{P}(\alpha)$ is a superset of the feasible region of $\mathcal{P}(g(x_f^*))$. Let x^* be a solution to $\mathcal{P}(\alpha)$. Since x_f^* lies in the feasible region of this problem and $f(x_f^*) \leq f(x)$ for any x , then $f(x^*) = f(x_f^*)$. As long as the minimum of f lies in the feasible region of $\mathcal{P}(\alpha)$, this will be a solution. \square

Definition 4.7 ([21]). A function $f : [a, b] \rightarrow \mathbb{R}$ is *one-sided Lipschitz* if for any $x_1, x_2 \in [a, b]$

$$(x_1 - x_2)(f(x_1) - f(x_2)) \leq m(x_1 - x_2)^2 \quad (17)$$

for some m .

Intuitively, in the case that f is continuously differentiable, the Lipschitz condition bounds $|f'(x)|$ while this condition only guarantees an upper bound on $f'(x)$. This is illustrated in Figure 3.

Lemma 4.8. Let \circ denote the composition operator: $(g \circ \mathcal{P})(\alpha) = g(\mathcal{P}(\alpha))$. We claim:

$$(g \circ \mathcal{P})(\alpha) = \alpha \text{ for any } \alpha \in [g(x_g^*), g(x_f^*)] \quad (18)$$

Proof. Let $x = \mathcal{P}(\alpha)$. x satisfies the KKT conditions for the $\mathcal{P}(\alpha)$ problem and either:

Case 1) x is the unconstrained minimum of f (i.e. $0 \in \partial f(x)$)

Case 2) x lies on the boundary of the feasible region (so the $g(x) \leq \alpha$ constraint is active and $-\partial f(x) \cap \partial g(x) \neq \emptyset$).¹

As previously used for Theorem 4.6, case 1 will only be true for $\alpha \geq g(x_f^*)$, with equality in the range we consider here. In case 2, we have the theorem simply from the fact that the constraint is active. \square

Lemma 4.9. $f \circ \mathcal{P}$ is monotonically non-increasing.

¹ Here ∂f denotes the subdifferential of f .

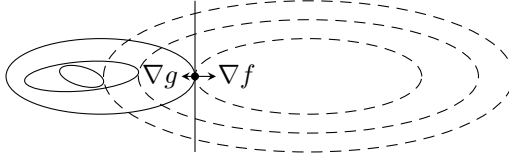


Figure 4: A series of expanding sublevel sets of f (solid ellipses) and the corresponding minimal intersecting sublevel sets of g (dashed ellipses). The gradient at a boundary is normal to a supporting line, and each pair of sublevel sets intersect only along this line.

Proof. Take any $\alpha_1 < \alpha_2$. Let $x_1 = \mathcal{P}(\alpha_1)$ and $x_2 = \mathcal{P}(\alpha_2)$. Since $g(x_1) = \alpha_1 < \alpha_2$ we know x_1 is feasible for the $\mathcal{P}(\alpha_2)$ problem. Thus, since x_2 is a minimizer for the $\mathcal{P}(\alpha_2)$ problem, $f(x_1) \geq f(x_2)$. \square

Theorem 4.10. $h \circ \mathcal{P}$ is one-sided Lipschitz.

Proof. Take any $\alpha_1, \alpha_2 \in [g(x_g^*), g(x_f^*)]$.

Since $f \circ P$ is monotonically non-increasing,

$$(\alpha_1 - \alpha_2) ((f \circ P)(\alpha_1) - (f \circ P)(\alpha_2)) < 0$$

So if we let $m \geq 1$, and because $(\alpha_1 - \alpha_2)^2 \geq 0$,

$$\begin{aligned} (\alpha_1 - \alpha_2) ((f \circ P)(\alpha_1) - (f \circ P)(\alpha_2)) &\leq (m - 1)(\alpha_1 - \alpha_2)^2 \\ (\alpha_1 - \alpha_2) ((f \circ P)(\alpha_1) - (f \circ P)(\alpha_2) + (\alpha_1 - \alpha_2)) &\leq m(\alpha_1 - \alpha_2)^2 \end{aligned}$$

Using Lemma 4.8,

$$(\alpha_1 - \alpha_2) ((h \circ P)(\alpha_1) - (h \circ P)(\alpha_2)) \leq m(\alpha_1 - \alpha_2)^2$$

So that h is one-sided Lipschitz with constant m . \square

With the result of Theorem 4.10, by simply sampling $h(P(\alpha))$ densely enough, we can get an estimate of this function to *arbitrary* precision over the entire interval $\alpha \in [g(x_g^*), g(x_f^*)]$. If we select a global minimum of this estimated function, we can derive a point with objective *arbitrarily* close to the true minimum of h . The bounds are given from the following theorem:

Lemma 4.11. For any interval $[\alpha_1, \alpha_2] \subseteq [g(x_g^*), g(x_f^*)]$ and any $\alpha \in [\alpha_1, \alpha_2]$,

$$\begin{aligned} (h \circ \mathcal{P})(\alpha) &\geq (h \circ \mathcal{P})(\alpha_2) - (\alpha_2 - \alpha) \\ &\geq (h \circ \mathcal{P})(\alpha_2) - (\alpha_2 - \alpha_1) \end{aligned} \tag{19}$$

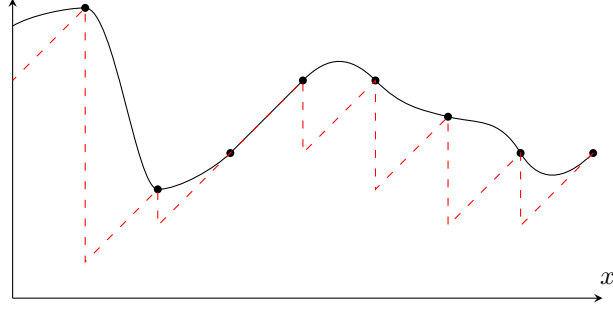


Figure 5: Illustration of the lower bound used in Lemma 4.11. If we only have the sampled points shown, we can nonetheless guarantee that a function which is one-sided Lipschitz will lie above the dashed lines.

Proof. Take $m = 1$ in Lemma 4.11, applied to this particular α, α_2 :

$$(\alpha - \alpha_2)((h \circ P)(\alpha) - (h \circ P)(\alpha_2)) \leq (\alpha - \alpha_2)^2$$

By interval construction $\alpha - \alpha_2 \leq 0$, so

$$\begin{aligned} (h \circ P)(\alpha) - (h \circ P)(\alpha_2) &\geq (\alpha - \alpha_2) \\ (h \circ P)(\alpha) &\geq (h \circ P)(\alpha_2) - (\alpha_2 - \alpha) \end{aligned}$$

□

Theorem 4.12. Take an increasing sequence $\alpha_1, \dots, \alpha_n$ from the domain of $h \circ P$ with $\alpha_1 = g(x_g^*)$ and $\alpha_n = g(x_f^*)$. Let τ be the maximum gap $\alpha_{i+1} - \alpha_i$. Denote the minimum sample by α^* , chosen such that

$$(h \circ P)(\alpha^*) = \min_i (h \circ P)(\alpha_i) \quad (20)$$

Then the function h has lower bound:

$$h(x) \geq (h \circ P)(\alpha^*) - \tau \quad \forall x \quad (21)$$

Proof. By Theorem 4.5, we can instead show

$$(h \circ P)(\alpha) \geq (h \circ P)(\alpha^*) - \tau \quad \forall \alpha \quad (22)$$

Consider any α . By our choice of $\alpha_1, \dots, \alpha_n$ there is an α_i, α_{i+1} such that $\alpha \in [\alpha_i, \alpha_{i+1}]$. By Lemma 4.11,

$$(h \circ P)(\alpha) \geq (h \circ P)(\alpha_i) - (\alpha_{i+1} - \alpha_i)$$

by the construction of α^* ,

$$\geq (h \circ P)(\alpha^*) - (\alpha_{i+1} - \alpha_i)$$

and since $\tau \geq \alpha_{i+1} - \alpha_i$

$$\geq (h \circ P)(\alpha^*) - \tau$$

□

4.2 Application to Scale-Free Cosegmentation

In the scale-free cosegmentation setting our \mathcal{P} problem consists of

$$\begin{aligned} \min_x \quad & x^T L x \\ \text{s.t.} \quad & \lambda \bar{h}^T H x \geq -\alpha \|H x\| \\ & l \leq x \leq u \end{aligned} \tag{23}$$

With this case f is convex and g is quasiconvex. The \mathcal{P} problem can therefore be solve efficiently using ordinary methods. We can calculate x_f^* as the solution to the ordinary random-walker segmentation and x_g^* is the projection of \bar{h} onto the feasible set under seed constraints.

5 Optimization

This section describes our strategy for solving (6) to optimality in a highly efficient manner. We use a projected gradient-based method which would additionally form the key component if we were to use any variation (*i.e.* augmented Lagrangian) as stated explicitly in [22].

Identifying a Sparse Structure: Expressing our model as a purely bound-constrained problem as in (6) requires the formation of the $H_i^T H_i$ products which form *dense* $n \times n$ matrices. Consequently, our optimization method must be careful not to explicitly form these matrices. Fortunately, we observe that explicit formulation of these matrices may be avoided by gradient projection methods, which are a simple and efficient class of algorithms in which it is only necessary to be able to calculate matrix-vector products. Here, this product can be distributed over the *sparse* components,

$$(L + H_1^T H_1)x_1 = Lx_1 + H_1^T (H_1 x_1). \tag{24}$$

With this modification, we can solve our Box-QP in (6) by adapting the Gradient Projection/Conjugate Gradient (GPCG) algorithm of [23]. We describe this strategy next.

5.1 GPCG

GPCG solves quadratic problems with a rectilinear feasible region $\Omega = \{x : l \leq x \leq u\}$. The algorithm alternates between two main phases (GP and CG): these correspond to alternately estimating the active set at the minimum and finding the minimum while keeping the active set fixed.

A) Gradient Projection (GP). Gradient projection coupled with the projected line search allows fast searching of a wide range of Ω . As a result, we can rapidly estimate which face of the feasible region the optimum lies on.

A.1) Search Direction for GP: In this phase, we search along a *projected gradient* $\nabla_{\Omega}\mathcal{O}(x)$ which has been modified so the corresponding search direction $d = -\nabla_{\Omega}\mathcal{O}$ does not point outside the feasible region. Specifically, this search direction is constructed to satisfy $\exists \epsilon > 0$ such that $x + \epsilon d \in \Omega$. We then use a *projected line search* (described later) to arrive at a step length α and update $x \leftarrow x + \alpha d$.

A.2) Phase Switch for GP: We switch to the conjugate gradient phase if *either* of the following conditions are satisfied: **(a)** The *active set* $\mathcal{A} = \{i : x_i = l_i \text{ or } x_i = u_i\}$ remains unchanged between two iterations; note that the active set corresponds to the minimal face (w.r.t. inclusion) of Ω to which x belongs; **(b)** GP is making little progress.

B) Conjugate Gradient (CG). We search a given face of the feasible region of our model using the conjugate gradient phase described below.

B.1) Search Direction for CG: Given the active set, our algorithm calculates a search direction conjugate to the previous direction (under the projection on to the free variables). Note that this method of generating a search direction is the same as applying ordinary conjugate gradient descent to a restriction of the QP to the current minimal face.

B.2) Phase Switch for CG: If the projected gradient points *out* of the current face (or if the iterations are making little progress), we switch back to the gradient projection (GP) phase. Formally, this is true if $\exists i \in \mathcal{A}(x)$ and either

$$x_i = l_i \text{ and } \partial_i \mathcal{O}(x) < 0, \quad \text{or} \quad x_i = u_i \text{ and } \partial_i \mathcal{O}(x) > 0.$$

Note that these “phase switch” conditions will never be satisfied for the face which contains the global minimum for our model. Thus, when the gradient projection phase has found the correct active set, conjugate gradient iterations suffice to explore the final face.

C) Projected Line Search: The projected line search modifies the active set by an arbitrary number of elements, and helps our GPCG process. Given a starting point x and search direction d , the line search finds $\alpha > 0$ which produces a *sufficient decrease* under the Armijo rule of the function $\phi(\alpha) = \mathcal{O}(P[x + \alpha d])$, where P describes the projection function which maps its input to the closest point in Ω . This can be thought of as a “line search” along a 1-manifold which bends to stay within Ω (thus, not a ray). Rather than directly finding all the piecewise quadratic segments of $\phi(\alpha)$, we efficiently produce a sufficient decrease using an estimate of ϕ by sampling one point at a time (as in [23]). It can be verified that *all* operations above can be expressed as Level 1, 2 BLAS operations. This allows a highly parallel implementation, described next.

D) GPU Implementation. Graphical Processing Units (GPUs) have gained attention from a wide range of the scientific computing community for their ability to efficiently address parallel problems [24]. These architectures operate by running multiple instances of a piece of code simultaneously, operating on different parts of a dataset. While this approach is not well-suited to all algorithms, Level 1 and 2 BLAS operations used in our algorithm are known to fit well with this architecture and can therefore

exhibit a significant speedup. The linear algebra operations comprising our GPCG algorithm for Cosegmentation may be easily parallelized using high-level languages as CUDA. In fact, the CUSPARSE toolkit (used here), supports Level 2 and 3 BLAS operations on sparse matrices as well. Further, the control flow of our procedure relies only on the results of *accumulations*, so the standard bottleneck of transferring data between main and GPU memory is not a major factor, and entry-level hardware (found on many workstations already) is sufficient.

6 Histogram Construction

In this section we discuss the fact that a wide variety of *different* preprocessing steps can take advantage of the *same* core method to produce a cosegmentation result.

At the heart of cosegmentation is some notion of similarity between pixels. The key property of cosegmentation over independent cosegmentation is that the foregrounds should contain *similar* distributions of pixels. In the group of cosegmentation methods we consider here, we rely on *histograms*, in which similar pixels are placed in the same histogram bin and dissimilar pixels are placed in different histogram bins. This is done through a *histogram assignment matrix* H as in (4).

In order to apply a cosegmentation method to a pair of images, we must first construct this matrix. One aspect of the problem is to assign a *bin* to some subset (or all) of the pixels between the two images. This needs to be done in such a way that similar foregrounds have similar histograms so that it matches a common-sense idea of when two foregrounds are similar. This includes some invariance to lighting, rotation, etc.

The bin assignment problem can be divided into two steps:

1. Assign local descriptors [25] to each pixel independently.
2. Find sets of “similar” pixels.

6.1 Texture-Based Histograms

We empirically found high-quality dense histograms were most consistently produced by clustering in a descriptor space based on texture. We derive the descriptors by applying the 17-filter bank of texture filters proposed in [18] shown in Figure 6. The responses at each pixel can then be clustered using nonparametric methods such as mean-shift or the greedy clustering of [26] with modified stopping conditions, avoiding explicitly specifying the number of clusters k (as we found the optimal k varies somewhat between images).

As an alternative to naïve clustering, the authors of [18] further present a method of training a texture dictionary from an incompletely labeled dataset. They present an agglomerative clustering algorithm based on maximizing a probability

$$\tilde{P}(\hat{\mathbf{c}}|\{\mathbf{H}_n\}) = \frac{P(\{\mathbf{H}_n\}|\hat{\mathbf{c}})}{P(\{\mathbf{H}_n\}|\hat{\mathbf{c}}) + P(\{\mathbf{H}_n\}|\mathbf{c}^{\text{same}})} \quad (25)$$




Figure 6: Winn filters from [18] used in our histogram generation. The first three blur filters are applied to each channel of the Lab colorspace, with the rest applied only to the lightness channel.




Figure 7: Clustering (a) and matching (b) results generated using VLFeat [27].

for histograms \mathbf{H}_n and training labels $\hat{\mathbf{c}}$. This clustering can be leveraged in texture-based cosegmentation algorithms, producing a binning method optimized to handle differentiating between object classes in the training data.

As described above, the first step produces local contextual *descriptors* of each image pixel. The clustering step finds those pixels which are similar under the given descriptor (similar *color* and *texture* will be in the same bin). We also consider SIFT-based features in step one, and an optical-flow-based pipeline, discussed below.

6.2 SIFT-based

A popular class of descriptors are those based on gradient binning, including SIFT, GLOH, etc. The high dimension of these descriptors makes it relatively difficult to find high-quality dense matches between the feature vectors of different pixels. In our setting (which does not allow *i.e.* assuming a homography), we found the matching to either be sparse or overly specific. Representative matching and clustering results are shown in Figure 7, which are reasonable but did not lead to better segmentation results. This additional module was thus not utilized further.

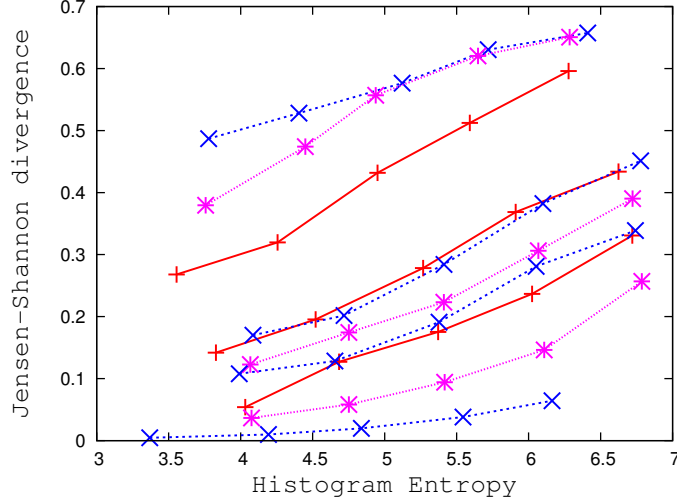


Figure 8: Comparing the entropy of the constructed histogram (x axis), with a symmetric KL-divergence between the true foreground histograms (y axis), as we vary the number of clusters. Each line shows one image pair. We consistently see higher-entropy histograms producing greater divergence for the target segmentation.

6.3 Histograms from Optical Flow

The flexibility of our method with respect to histogram construction allows the use of *application-specific* preprocessing steps. There is a natural parallel between cosegmentation and the task of segmenting *video*. When the images are temporally related, we have a relationship between foreground pixels determined by the movement shown in the video. This is a well-studied problem in computer vision, allowing us a class of *optical flow* methods.

We can find corresponding pixels using optical flow, placing pixels i and i' in the same bin if the position of i maps to i' in the next image. Applying this scheme to the frames of a video sequence suggests the model

$$\min_{x \in X} \sum_i x_i^T L_i x_i + \lambda \sum_i \|H_i x_i - H'_{i+1} x_{i+1}\|_2^2 \quad (26)$$

where the second sum compares subsequent frames of an image sequence.

In Fig. 16 we perform correspondences based on optical flow to map superpixel-based bins in order to segment frames of a video sequence. Other correspondence determination methods can be used and no change to our algorithm is needed.

6.4 Histogram Entropy

As discussed above, using low-entropy histograms allows for more accurate matching between images for some images. Intuitively, recall the example from Fig. 1 where a high-entropy histogram tries to differentiate between different patches of fur on a

brown bear. This cannot be done in a consistent manner across realistic images, so erroneous matches are introduced. By contrast, our histogram matching technique better describes the true texture description of the bear’s fur as a combinations of very few homogenous textures. We verify this experimentally in Fig. 8 which plots the statistical distance between the histogram of the true foreground for a sample of images from our dataset. In these cases we found that as the entropy increases, so does the JS divergence measure between the histograms for the true segmentations. Low entropy histograms, however, relate to smaller divergences, which impose the global Cosegmentation constraint more tightly.

7 Experiments

Our experimental evaluations included comparisons of our implementation (on a Nvidia Geforce 470) to another Cosegmentation method [7], and the Random Walker algorithm [11] (run independently on both images). We also performed experiments using the methods in [4] and [9], but due to the problem of solving a large LP and incorporation of foreground/background seeds respectively, results could not be obtained for the entire dataset described below. To assess *relative* advantages of the specialized GPCG procedure, we also compared it with a stand-alone implementation of (5) linked with a commercial QP solver (using multiple cores). We provide a qualitative and quantitative overview of the performance w.r.t. state of the art. Additional experiments demonstrate the efficacy of the multiple-image and scale-free segmentation models.²

7.1 Weak Boundary Experiment

For an initial probe of the behavior of the Random Walker as the engine of a cosegmentation method, we replicate the experiment of [11] in a cosegmentation setting. The setup in [11] used a synthetic image consisting of two white triangles which were separated by a black line that contained a large hole, such that the triangles appear to be touching (see Figure 9). This image is challenging to segment into the two triangles because each triangle exhibits exactly the same appearance (white), and there is a substantial portion of the boundary between them for which there is zero contrast. Consequently, segmentation algorithms which rely exclusively on appearance priors or boundary contrast are unable to separate the two triangles. In particular, [11] showed that a segmentation driven by a boundary length prior (e.g., graph cuts) was unable to separate these triangles when the inputs consisted of only a single seed in each triangle.

We may investigate the same scenario in a cosegmentation setting. Specifically, if we let an image pair consist of two replicas of the “two triangles” image from [11], we may supply one seed in each of the triangles of only *one* of the images. Further, we set up a synthetic histogram in such a way that the desired (i.e., perceptual) partition incurs zero histogram variation penalty relative to E^{global} in (1). An example of such a histogram pair is shown in Fig. 9 (rightmost column) where common colors indicate similarity among pixels in both images (based on some arbitrary feature). Given this

²The full set of segmentation results and code is available at <http://pages.cs.wisc.edu/~mcollins/pubs/cvpr2012.html>




Figure 9: Figure showing RW cosegmentation’s performance on weak boundaries. Columns (left to right): (1) Identical pair of images, with seeds placed only on the top image; (2) Output potentials of random walker; (3) Output of graph-cuts based cosegmentation (4) Identical histogram bin assignments for both images (common colors indicate similar pixels, considered both inter and intra image);

input image pair, our cosegmentation algorithm based on the Random Walker is able to successfully segment both images into two triangles (see Fig. 9, second column), despite the fact that each triangle exhibits the same appearance and it not well-separated by the boundary. In contrast, if we apply the traditional cosegmentation model which penalizes boundary length to the same problem, these two images are poorly segmented as a result of the difficulty of the problem instance (see Fig. 9, third column). Note that while this experiment is artificially constructed, it nonetheless suggests that Cosegmentation based on the Random Walker will retain the beneficial properties of the Random Walker segmentation algorithm and some of its advantages over boundary length regularization.

7.2 Datasets

In order to leverage all available test data we aggregated all images provided by the authors in iCoseg [2], Momi-coseg [3] and Pseudoflow-coseg [7], and further supplemented them with a few additional images from the web and [28]. In order to compare with algorithms that only handle image *pairs* we selected a dataset with 50 pairs of images (100 total). For the > 2 image case we used the iCoseg dataset from [2]. Since a number of image sets from (from [2]) share only semantically similar foreground objects and are unsuitable for cosegmentation with common appearance models, we selected 88 subsets comprising 389 of the 643 iCoseg images (also observed in [13]).

7.3 Running time complexity

We now discuss what is the strongest advantage of this framework. We show an example in Fig. 10 of the running time of the proposed model relative to [7], as a function of

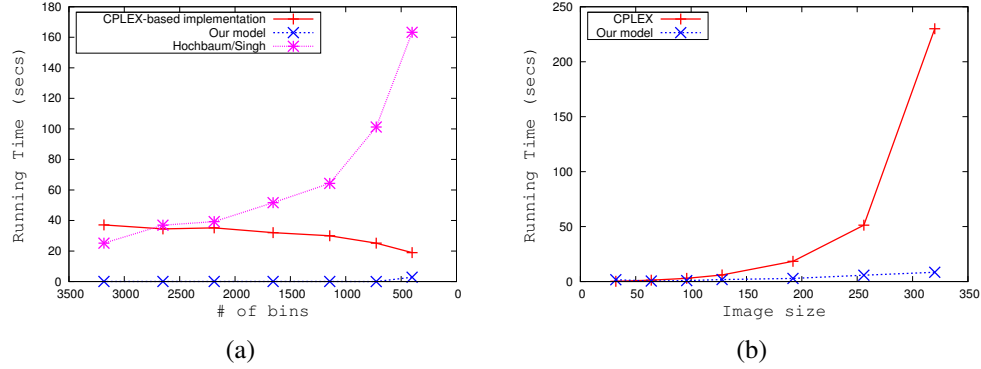


Figure 10: Variation in run-time with histogram granularity (a) and image size (b) relative to a Cplex-based implementation and Pseudoflow-coseg [7].

decreasing entropy (number of bins). The plot suggests that for a realistic range, our implementation has a negligible computation time relative to [7] and the CPLEX-based option – in fact, our curve almost coincides with the x -axis (and even this cost was dominated primarily by the overhead from problem setup done on the CPU). For the most expensive data point shown in Fig. 10, the model from [7] generated 10^7 auxiliary nodes (about 12 GB of memory). Due to the utility of low entropy histograms, these experiments show a salient benefit of our framework and its immunity to the ‘coarseness’ of the appearance model. Over all 128×128 images, the wall-clock running time of our CUDA-based model was 10.609 ± 5.230 seconds (a significant improvement over both [7, 9]). The time for a Cplex driven method (*utilizing four cores* in parallel) was 17.982 ± 5.07 seconds, but this increases sharply with greater problem size.

An artificial image shown in Figure 11 was used in order to allow for isolation of specific variables, and mitigate variability (in optimization time) as a function of the specific problem instance.

With an increase in the image size, the running time of the model is expected to increase because of two reasons. First, the number of pixels to segment determines the size of the input problem and the dimensionality of the decision variables in the optimization. Second, if the histogram is generated the same way across different sizes, the number of pixels in each histogram bin also increases. An analysis of this behavior is presented in Figure 10. The image shown in Fig. 11 was generated for various sizes, with the number of pixels along each side plotted along the x axis in Fig. 10. Our specialized GPU-based Cosegmentation library and an implementation based on the Cplex solver were used for Random Walker Cosegmentation. In this result, we see only marginal increase in the overall running time of our model (the curve stays close to the x -axis); on the other hand, the running time of the Cplex-based implementation increases quickly as we increase the size of the images. Notice how this difference is substantially magnified for larger images. The reason is that our algorithm distributes each individual BLAS operation over the GPU’s computational units;

as a result, higher-dimensional operations take better advantage of the parallelization in the model presented in this paper. We believe our method is thus especially suited for large images where computation time is a consideration.

7.4 Effect of λ parameter

The images in Fig. 12 demonstrate the role of the histogram consistency bias, λ in equation (5). For a very small λ , the segmentation probabilities are diffuse (compare to the independent random walker results from Fig. 15); as the influence of the bias grows, the consistency between the histograms makes the partitions more pronounced.

7.5 Performance w.r.t. pair methods

We evaluated the quality of segmentations (0/1 loss) on the 50 image pairs described above, relative to Pseudoflow-coseg [7], LP [4] (only partial), and discriminative clustering [9]. As in [11], a few seeds were placed to specify foreground and background regions, and given to all methods. Representative samples on the images from [2] using the method in [7, 9] are shown in Fig. 13. Averaged over the pair dataset, the segmentation accuracy of our method was $93.6 \pm 2.9\%$, where as the gross values for the algorithms from [7] and [9] were 89.1% and 84.1% respectively.

7.6 Performance on 2+ images

Across the iCoseg dataset we achieved an accuracy of 93.7% with seeds provided by 5 different users. The algorithm of [9] achieves an accuracy of 82.2% across the dataset (excluding some for which the implementation provided by the authors did not complete). Representative image sets and accuracy comparisons appear in Table 1 and Figure 14.

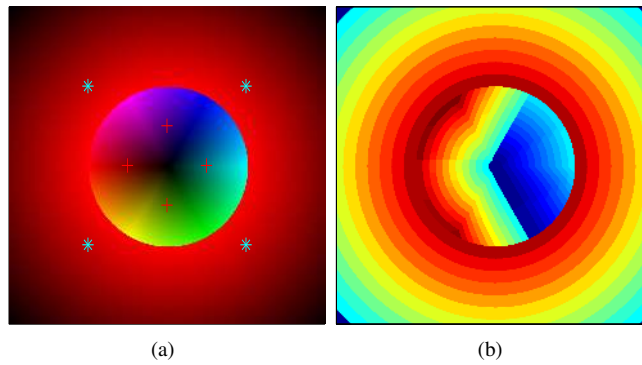


Figure 11: Artificial image used in computation time experiments (a) and corresponding histogram bins (b). Seed points were placed as shown in both foreground and background. To create a cosegmentation problem, the same image was used twice.




Figure 12: Effect of varying λ parameter on an example image for $\lambda = 10^{-8}$ in Column 2, and $\lambda = 10^{-6}$ in Column 3. A segmentation potential biased towards matching the histograms makes the partitions more pronounced.

We note that these results must be interpreted with the caveat that different methods respond differently to seed locations and the level of discrimination offered by the underlying appearance model. Since most cosegmentation models (including this paper) share the same basic construction at their core (i.e., image segmentation with an appearance model constraint), variations in performance are in part due to the input histograms. The purpose of our experiments here is to show that at the very least one can expect similar (if not better) qualitative results with our model, but with more flexibility and quite significant computational advantages.





Figure 13: Comparison results on example images (columns 1,2) of the the Pseudoflow based method of [7] (columns 3,4) and the discriminative clustering approach of [9] (columns 5,6), with segmentation from RW-based cosegmentation (columns 7,8).



	Our Method	[9]
Airshow (8/21)	99.7	80.5
Alaskan bear (5/18)	92.2	76.6
Christ (6/14)	95.6	97.3
Ferrari (4/10)	96.4	59.4
Goose (13/31)	97.3	97.5
Helicopter (9/11)	96.9	98.1
Kendo (10/31)	91.4	93.4
Lobster Kite (4/11)	97.0	88.5
Monk (4/17)	83.3	87.1
Soccer (5/36)	93.9	79.9
Speed Skater (9/13)	80.8	77.3
Liberty (11/41)	93.2	86.1

Table 1: Segmentation accuracy for some iCoseg image sets. Subsets were chosen which have similar appearance under a histogram model.



Figure 14: Interactive segmentation results using the multi-image model across five images from the “Kendo” set of iCoseg. As the number of images increases, *less* user input is required, as seen by the lack of such for the middle image.

7.7 Comparisons to independent Random Walker runs

In Fig. 15, we present qualitative results from our algorithm, and from independent runs of Random Walker (both with up to two seeds per image). A trend was evident on all images – the probabilities from independent runs of Random Walker on the two images were diffuse and provide poorer boundary localization. This is due to the lack of global knowledge of the segmentation in the other image. Random walker based Cosegmentation is able to leverage this information, and provides better contrast and crisp boundaries for thresholding (a performance boost of up to 10%).

8 Conclusions

We present a new framework for the cosegmentation problem based on the Random Walker segmentation approach. While almost all other cosegmentation methods view the problem in the MRF (graph-cuts) setting, our algorithm translates many of the advantages of Random Walker to the task of simultaneously segmenting common foregrounds from related images. Significantly, our formulation completely eliminates a practical limitation in current (nonparametric model based) Cosegmentation methods that requires the overall image histogram to be approximately flat (in order to restrict the number of auxiliary nodes added). Our model extends nicely to the multi-image setting using a penalty with statistical justification. A further extension allows model-based segmentation which is independent of the relative scales of the model and target




Figure 15: Columns (1–2): Input images; Columns (3–4) segmentation potentials from independent random walker runs on the two images; Columns (5–6) Segmentation potentials from Random Walker based cosegmentation. Note that the object boundaries have become significantly more pronounced because of the histogram constraint.



Figure 16: Segmentation using correspondences from optical flow on video sequence from [28]. Shows outline of segmented foreground in red, with foreground and background indications. Our algorithm achieves 99.3% accuracy.

foregrounds. We discuss its optimization specific properties, give a state of the art GPU based library, and show quantitative and qualitative performance of the method.

9 Acknowledgments

This work is funded via NIH 5R21AG034315-02 and NSF 1116584. Partial support was provided by UW-ICTR through an NIH CTSA Award 1UL1RR025011 and NIH grant 5P50AG033514 to W-ADRC. Funding for M.D.C. was provided by NLM 5T15LM007359 via the CIBM Training Program. The authors would like to thank Petru M. Dinu for consultations on GPGPU development.

References

- [1] C. Rother, V. Kolmogorov, T. Minka, and A. Blake. Cosegmentation of image pairs by histogram matching - Incorporating a global constraint in MRFs. In *CVPR*, 2006.
- [2] D. Batra, A. Kowdle, D. Parikh, J. Luo, and T. Chen. Interactive co-segmentation with intelligent scribble guidance. In *CVPR*, 2010.
- [3] W. Chu, C. Chen, and C. Chen. MOMI-cosegmentation: Simultaneous segmentation of multiple objects among multiple images. In *ACCV*, 2010.

- [4] L. Mukherjee, V. Singh, and C. Dyer. Half-integrality based algorithms for cosegmentation of images. In *CVPR*, 2009.
- [5] S. Vicente, V. Kolmogorov, and C. Rother. Cosegmentation revisited: Models & optimization. In *ECCV*, 2010.
- [6] Y.J. Lee and K. Grauman. Collect-cut: Segmentation with top-down cues discovered in multi-object images. In *CVPR*, 2010.
- [7] D.S. Hochbaum and V. Singh. An efficient algorithm for co-segmentation. In *ICCV*, 2009.
- [8] A. Kowdle, D. Batra, W. C. Chen, and T. Chen. iModel: Interactive co-segmentation for object of interest 3D modeling. *ECCV Workshop*, 2010.
- [9] A. Joulin, F. Bach, and J. Ponce. Discriminative clustering for image co-segmentation. In *CVPR*, 2010.
- [10] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *PAMI*, 23(11):1222–1239, 2001.
- [11] L. Grady. Random walks for image segmentation. *PAMI*, 28(11):1768–1783, 2006.
- [12] Lopamudra Mukherjee, Vikas Singh, and Jiming Peng. Scale invariant cosegmentation for image groups. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2011.
- [13] S. Vicente, C. Rother, and V. Kolmogorov. Object cosegmentation. In *CVPR*, 2011.
- [14] K. Chang, T. Liu, and S. Lai. From co-saliency to co-segmentation: An efficient and fully unsupervised energy minimization model. In *CVPR*, 2011.
- [15] A. K. Sinop and L. Grady. A seeded image segmentation framework unifying graph cuts and random walker which yields a new algorithm. In *ICCV*, 2007.
- [16] A. Levin, D. Lischinski, and Y. Weiss. A closed-form solution to natural image matting. *PAMI*, 30(2):228–242, 2008.
- [17] J. Cui, Q. Yang, F. Wen, Q. Wu, C. Zhang, L. Van Gool, and X. Tang. Transductive object cutout. In *CVPR*, 2008.
- [18] J. Winn, A. Criminisi, and T. Minka. Object categorization by learned universal visual dictionary. In *ICCV*, 2005.
- [19] M.S. Bazaraa, H.D. Sherali, and C.M. Shetty. *Nonlinear Programming*. John Wiley & Sons, Inc., 2003.
- [20] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [21] Reinhard Frank, Josef Schneid, and Christoph W. Ueberhuber. Stability properties of implicit Runge-Kutta methods. *SIAM J. on Numerical Analysis*, 22(3):497–514, 1985.
- [22] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer, USA, 2nd edition, 2006.
- [23] J.J. More and G. Toraldo. On the solution of large quadratic programming problems with bound constraints. *SIAM J. on Optimization*, 1(1):93–113, 1991.
- [24] S. Lee and S. J. Wright. Implementing algorithms for signal and image reconstruction on graphical processing units. Technical report, University of Wisconsin Madison, 2008.
- [25] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(10):1615–1630, 2005.
- [26] T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293 – 306, 1985.
- [27] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision

algorithms. <http://www.vlfeat.org/>, 2008.

- [28] D. Tsai, M. Flagg, and J. M. Rehg. Motion coherent tracking with multi-label mrf optimization. In *British Machine Vision Conference (BMVC)*, 2010.