Asymmetric Cuts : Joint Image Labeling and Partitioning

Thorben Kroeger[1], Jörg H. Kappes[2], Thorsten Beier[1], Ullrich Koethe[1] and Fred A. Hamprecht[1,2]

[1] Multidimensional Image Processing Group, Heidelberg University
[2] Heidelberg Collaboratory for Image Processing, Heidelberg University

Abstract. For image segmentation, recent advances in optimization make it possible to combine noisy region appearance terms with pairwise terms which can not only *discourage*, but also *encourage* label transitions, depending on boundary evidence. These models have the potential to overcome problems such as the shrinking bias. However, with the ability to encourage label transitions comes a different problem: strong boundary evidence can overrule weak region appearance terms to create new regions out of nowhere. While some label classes exhibit strong internal boundaries, such as the background class which is the pool of objects. Other label classes, meanwhile, should be modeled as a single region, even if some internal boundaries are visible.
We therefore propose in this work to treat label classes asymmetrically: for some classes, we allow a further partitioning into their constituent objects as supported by boundary evidence; for other classes, further partitioning is forbidden. In our experiments, we show where such a model can be useful for both 2D and 3D segmentation.

1 Introduction

Image segmentation methods typically rely on two complementary sources of information: object appearance and boundary evidence. For example, in semantic labeling tasks [14] a set of object classes of interest is given. Each image can contain one or more of these instances, but might also contain many objects of unknown classes ("background"). One approach for semantic segmentation is to make use of (noisy) local object class probabilities – as obtained from learned appearance models – which can be regularized using local boundary cues.

On the other hand, pure partitioning problems, as in the Berkeley Segmentation Dataset [30], do not specify any object classes but rely on boundary evidence alone [5,3,37].

In this work, we propose a combined semantic labeling *and* partitioning, called *Asymmetric Segmentation*, which can naturally deal with object classes which are known to have strong internal boundaries and jointly optimizes the region labeling, the boundaries between classes and the boundaries within classes. Furthermore we present a novel algorithm called Asymmetric Multi-Way Cut (AMWC) for solving those problems.

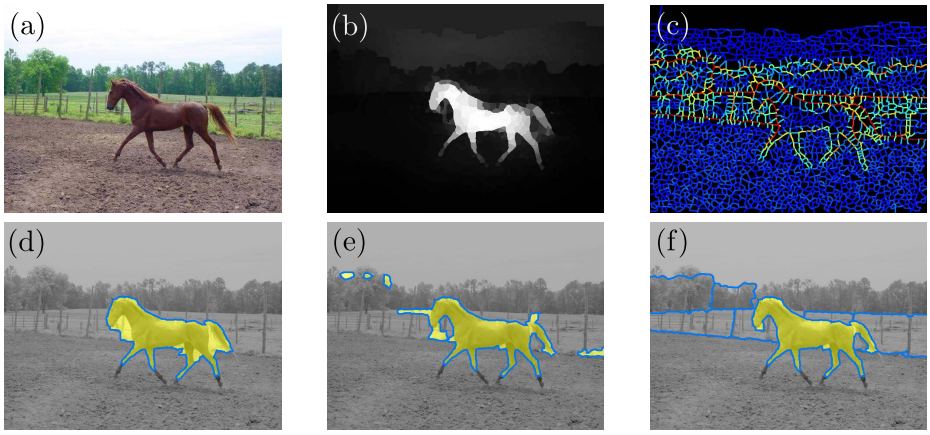

Fig. 1: Segmentation of image (a) can combine information from both region appearance terms (b) and boundary probabilities (c). We examine different variants of pairwise Conditional Random Field models with Potts potentials.
Graph cut (d) uses positive coupling strengths only, which leads to shrinking bias. Multi-way cut (e) uses both negative and positive coupling strengths, such that the creation of boundaries can be actively encouraged. However, this leads to some spurious labelings, induced by strong boundary evidence. Our proposed variant, (f), can yield a better segmentation by allowing boundaries within the background class.

Many segmentation algorithms, including AMWC, are formulated as second-order Conditional Random Fields [17] over a discrete set of labels, in which the unary potentials transport local evidence for each object class. The pairwise potentials are usually chosen to be Potts functions (2) with varying coupling strengths $w \in \mathbb{R}^+$, which may depend on boundary evidence. The optimal labeling can then be found by minimizing the associated energy function.

Graph cut based algorithms have been extremely influential in the last decade [12,23,33], because they allow to find the optimal solution for binary labeling problems and approximate solutions for multi-label problems with non-negative coupling strengths in polynomial time. They regularize noisy detections by penalizing boundary length. Unfortunately, this leads to "shrinking bias" [36], i.e. thin, elongated objects are cut off (Fig 1d). As a countermeasure, the coupling strength can be chosen as an inverse function of boundary evidence, making label transitions *less* costly when strong boundary evidence exists and *more* costly when boundary evidence is weak. However, the general problem remains: positive coupling strengths cannot actively *encourage* label transitions.

Since negative coupling strengths w *encourage* label transitions and positive w *discourage* label transitions, a model with no restriction on the sign of w may be more expressive: with strong boundary evidence (resulting in $w < 0$) along a thin, elongated object, shrinking bias can be overcome.

Recently, Kappes et al. [18] presented a method that, contrary to others [22,24], is able to find the globally optimal solution for these more general models. This allows us to evaluate the models without any error introduced by approximate optimization.

However, besides their increased computational hardness, models which can *encourage* label transitions also have a major drawback: spurious label transitions are provoked in highly cluttered background (Fig 1e) or within textured objects when strong boundary evidence overrules homogeneous region appearance terms.

In this work, we investigate a new class of models, which do allow intra-category boundaries. The energy function can still be expressed as a pairwise Conditional Random Field. Fig. 1f shows the result of our new model, in which we allow internal edges in the "background" class, but disallow internal edges in the "foreground" class.

Contributions. (i) A novel class of segmentation models, where some classes may have internal boundaries and others may not. (ii) An exact solver for AMWC problems based on a formulation using binary edge indicator variables, based on [19]. (iii) Experiments that show when such a formulation is useful and when it is not.
The C++-code of our method will be made available within OpenGM [2].

2 Related Work

Let $G = (\mathcal{V}, \mathcal{E})$ be a given pixel (or superpixel) adjacency graph. Each node $i \in \mathcal{V}$ can be assigned one of k discrete labels: $l_i \in \mathcal{L} = \{0, \dots, k-1\}$. Many common pixel or superpixel labeling problems [12,33,35,36,6,17] are then formulated as an energy minimization over the sums of unary and pairwise terms:

$$\underset{\boldsymbol{l} \in \mathcal{L}^{|\mathcal{V}|}}{\operatorname{argmin}} \left\{ \sum_{i \in \mathcal{V}} E_i(l_i) + \sum_{(i,j) \in \mathcal{E}} E_{ij}(l_i, l_j) \right\}. \tag{1}$$

The *unary terms* are functions $E_i : \mathcal{L} \to \mathbb{R}$ and indicate the local preference of node i to be assigned a label. The *pairwise terms* are functions $E_{ij} : \mathcal{L} \times \mathcal{L} \to \mathbb{R}$ that express the local joint preferences of two adjacent nodes i and j. A common choice for the binary term is a *Potts* function [12]:

$$E_{ij}^{\boldsymbol{w}}(l_i, l_j) = \begin{cases} 0 & \text{if } \; l_i = l_j \\ \boldsymbol{w}_{ij} & \text{if } \; l_i \neq l_j \end{cases}. \tag{2}$$

Depending on the weight $\boldsymbol{w}_{ij}$ a Potts function can either *encourage* ($\boldsymbol{w}_{ij} < 0$) or *discourage* ($\boldsymbol{w}_{ij} > 0$) label transitions.

Binary labeling problems with $\mathcal{L} = \{0, 1\}$ and pairwise potentials with $\forall (i,j) \in \mathcal{E} : \boldsymbol{w}_{ij} \geq 0$ (*graph cut problems*) can be solved in polynomial time with a max-flow algorithm [12,23,11]. Graph cut has been ubiquitous in image segmentation

[33], but penalizes a weighted sum of cut edges which leads to the problem of shrinking bias [36], for which sophisticated countermeasures were developed.

For $k > 2$ labels, (1) becomes a multi-label energy minimization problem that is NP-hard in general, even for non-negative weights $\boldsymbol{w}$. We will refer to the general problem of (1) with $2 \leq k \ll |\mathcal{V}|$ and $\boldsymbol{w} \in \mathbb{R}^{|\mathcal{E}|}$ as the *multi-way cut problem* because the optimal solution can be found as a multicut in a graph with special structure (reviewed in Sec. 3). Approaches to solve this problem approximately are, amongst others, move-making algorithms [12,6] and linear programming [21,25,19,17]. An integer linear program to which violated constraints are added in a cutting-plane fashion [18,19] is able to find the globally optimal solution on many problem instances, see Sec. 3.

A special case of the labeling problem with $E_i(l_i) \equiv 0$, a virtually unlimited set of labels $k = |\mathcal{V}|$ and no restriction on the sign of $\boldsymbol{w}_{ij}$ is called the *correlation clustering* [7] or *multicut problem* [13]. The multicut formulation has recently become popular for unsupervised image segmentation [3,4,6,20,37,1,27,9] where the weights $\boldsymbol{w}$ are either learned in a supervised fashion [3,6,4,20,27], or derived from boundary detectors such as gPb [29] as in [37]. Multicut models have an inherent model-selection ability [6] such that they recover the optimal number of regions needed for an accurate segmentation automatically, based only on boundary evidence. However as region appearance is not taken into account, the resulting segments are not given a class label, which is left as a post-processing step, e.g. done in [20]. Note that we review only unified formulations with a specified objective function here and omit workflows that chain several processing steps.

3 Multi-Way Cut Formulation

Before we will define Asymmetric Multi-way Cut in Sec. 4 we first review the Multiway Cut representation of the labeling problem (1) with Potts potentials (2), based on [18,19].

Given a graph $G = (\mathcal{V}, \mathcal{E})$ for (1), we call the nodes $\mathcal{V}$ the *internal nodes* and edges $\mathcal{E}$ the *internal edges*. We then define a new graph $G' = (\mathcal{V}', \mathcal{E}')$, where a set of *terminal nodes* $T = \{t_0, \ldots, t_{k-1}\}$, representing k labels, has been added (Fig. 2, left). Furthermore, *terminal edges* are introduced between each pair of internal and terminal nodes as well as between all pairs of terminal nodes:

$$\mathcal{V}' = \mathcal{V} \cup T \tag{3a}$$

$$\mathcal{E}' = \mathcal{E} \cup \{(t, v) \mid t \in T, v \in \mathcal{V}\} \cup \{(t_i, t_j) \mid 0 \leq i < j < k\}\,. \tag{3b}$$

Problem (1) with potentials (2) can be written using indicator variables $\boldsymbol{y}$ for the edges $\mathcal{E}'$ and weights $\boldsymbol{w}'$, derived from the potentials $E_i(\cdot)$ and $E_{ij}(\cdot,\cdot)$ [18]:

$$\underset{\boldsymbol{y} \in \{0,1\}^{|\mathcal{E}'|}}{\operatorname{argmin}} \left\{ \sum_{(i,j) \in \mathcal{E}'} \boldsymbol{w}'_{ij} \cdot \boldsymbol{y}_{ij} \right\} \quad \text{s.t. } \boldsymbol{y} \in \mathrm{MWC}_G, \tag{3c}$$

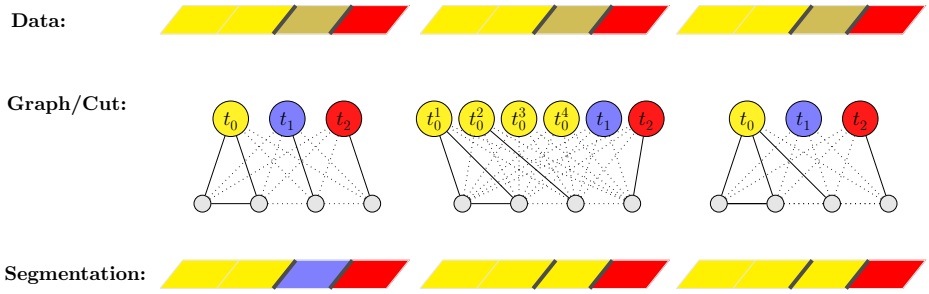

Fig. 2: Illustration of the graph representation of (1), as given by (3a)-(3b). Shown is a 4×1 pixel image strip ("Data" row). There are three possible classes $\mathcal{L} = \{\text{red}, \text{blue}, \text{yellow}\}$. From the left to the right pixel, the region appearance terms indicate strong preference for the yellow class (pixels 1, 2), a preference for yellow over the blue class (pixel 3, p_3), and a strong preference for the red class for the last pixel. Furthermore, there are strong boundaries to the left and right of p_3. **Left:** Multi-way cut solution. Due to the strong boundary evidence, p_3 is assigned the blue label contrary to its unary potential. **Middle:** The AMWC model with $\mathcal{A} = \{\text{yellow}\}$ is formulated as a MWC by duplicating the terminal nodes for the yellow class $|\mathcal{V}| = 4$ times. **Right:** With the modified constraint (4c′), the number of terminal nodes does not have to be increased.

where MWC_G is the multi-way cut polytope defined by linear constraints [19]:

$$\sum_{(i,j)\in P} \boldsymbol{y}_{ij} \geq \boldsymbol{y}_{uv} \quad \forall\, (u,v) \in \mathcal{E},\ P \in \text{Path}(u,v) \subseteq \mathcal{E} \tag{4a}$$

$$\boldsymbol{y}_{tt'} = 1 \quad \forall\, (t,t') \in T, t \neq t' \tag{4b}$$

$$\boldsymbol{y}_{tu} + \boldsymbol{y}_{tv} \geq \boldsymbol{y}_{uv} \quad \forall\, (u,v) \in \mathcal{E}, t \in T \tag{4c}$$

$$\boldsymbol{y}_{tu} + \boldsymbol{y}_{tv} \geq \boldsymbol{y}_{tv} \quad \forall\, (u,v) \in \mathcal{E}, t \in T \tag{4d}$$

$$\boldsymbol{y}_{tv} + \boldsymbol{y}_{uv} \geq \boldsymbol{y}_{tu} \quad \forall\, (u,v) \in \mathcal{E}, t \in T. \tag{4e}$$

For internal nodes, the *cycle constraint* (4a) [13,3] intuitively forbids dangling boundaries. Constraint (4b) ensures that all terminals are always separated. Finally, (4c)-(4e) constitute cycle constraints for all cycles of three nodes involving one terminal. In particular, (4c) says that, if there is a label transition ($\boldsymbol{y}_{uv} = 1$), label t cannot belong to both u and v (which would be the case for $\boldsymbol{y}_{tu} = 0$ and $\boldsymbol{y}_{tv} = 0$).

Although a complete description of MWC_G needs a possibly exponential number of constraints, in practice only a small set of active constraints is needed to find the (valid) globally optimal solution. The cutting plane method [18,19] first formulates an unconstrained integer linear program and then identifies violated constraints in the solution, which are subsequently added to the problem. This is repeated until a solution does not violate any of the constraints in (4a)-(4e), yielding the globally optimal solution.

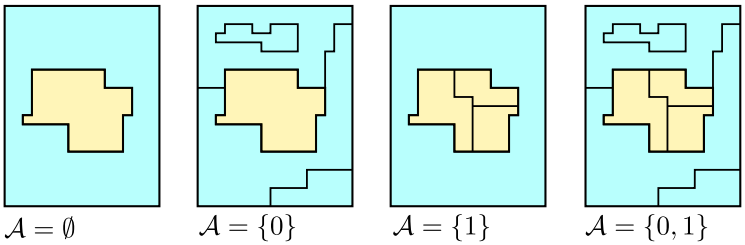

Fig. 3: Illustration of allowed cuts, depending on $\mathcal{A}$. Blue denotes background (label 0) and yellow foreground (label 1). From left to right we demand no label transitions within foreground and background regions ($\mathcal{A} = \emptyset$), we allow label transitions within the background class ($\mathcal{A} = \{0\}$), or only within the foreground class ($\mathcal{A} = \{1\}$) or in both classes ($\mathcal{A} = \{0, 1\}$). In all cases, we admit only closed contours (black boundaries).

4 The Asymmetric Multi-way Cut Formulation

In the proposed asymmetric multiway cut (AMWC) model, we want to allow internal boundaries within regions labeled as $l \in \mathcal{A}$, and disallow internal boundaries in all regions labeled $l \in (\{0, \dots, k-1\} \setminus \mathcal{A})$, as illustrated in Fig. 3.

4.1 Formulation within the Binary Edge Labeling Framework

We first give the advantageous formulation of AMWC in terms of the binary labeling of $\mathcal{E}'$ in (3c). One way to formulate the model as a multiway cut is to replace every terminal node $t_a \in \mathcal{A}$ with a set $T_a = \{t_a^0, \dots, t_a^{|\mathcal{V}|-1}\}$, as shown in Fig. 2, middle and Fig. 4, for which edge weights are copied from the existing terminal edges. These new nodes can represent a partitioning of class a into sub-classes, which are separated by salient boundaries in the image. Similar to the multicut formulation, the label space has to be increased dramatically. By setting $|T_a| = |\mathcal{V}|$, solutions where every node is assigned a different label from the set T_a are made possible.

However, instead of adding $|\mathcal{A}| \cdot (|\mathcal{V}| - 1)$ additional terminal nodes, the same effect can be achieved by simplifying a single constraint in the binary edge labeling formulation of Sec. 3. We relax constraint (4c) to be

$$\boldsymbol{y}_{tu} + \boldsymbol{y}_{tv} \geq \boldsymbol{y}_{uv} \quad \forall\, (u, v) \in \mathcal{E}, t \in (T \setminus \mathcal{A}). \tag{4c$'$}$$

In practice, we can extend the implementation of [19] such that constraint (4c) is only added in the cutting-plane procedure when $t \notin \mathcal{A}$ holds.

With this simple change, we have a method that is able to solve our new AMWC-type models to global optimality.

4.2 Formulation as a Node Labeling Problem

Our AMWC model can still be formulated as a second-order Conditional Random Field (1) with Potts potentials. Starting from the edge labeling formulation

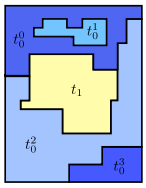

Fig. 4: In this example with $\mathcal{A} = 0$, the background class (label zero, shown in blue) actually consists of several distinct segments t_0^0, t_0^1, t_0^2 and t_0^3 (different shades of blue), all separated by boundaries (black). Another region, $t_1 \notin \mathcal{A}$ cannot be split into sub-segments.

from the previous section, we construct the corresponding model by choosing a new set of labels $\mathcal{L}'$ as well as constructing appropriate unary and pairwise potentials $E'_i(\cdot)$ and $E'_{ij}(\cdot,\cdot)$. Let the original set of labels be $\mathcal{L} = \{t_0, \dots, t_{k-1}\}$. Then we set

$$\mathcal{L}' = \bigcup_{a \in \mathcal{A}} \left(\{t_a^0, \dots, t_a^{|\mathcal{V}|-1}\} \right) \cup (\mathcal{L} \setminus \mathcal{A}) \tag{5a}$$

$$E'_i(l'_i) = \begin{cases} E_i(l'_i) & \text{if} \quad l'_i \in (\mathcal{L} \setminus \mathcal{A}) \\ E_i(a) & \text{if} \quad l'_i = t_a^j \end{cases} \tag{5b}$$

$$E'_{ij}(l'_i, l'_j) = E_{i,j}(a,b) \quad \text{with } t_a^u = l_i, t_b^v = l_j \qquad \forall u, v. \tag{5c}$$

In (5a), we introduce $|\mathcal{V}| - 1$ additional labels for each label class for which internal boundaries are allowed. These extra labels are assigned the same weights in the new unary terms (5b) as the original label class. Finally, the Potts terms do not change (5c).

The inflated label space makes this formulation unwieldy for practical optimization methods: similar to multicut models, a large number of different labelings, obtained by label permutations, have the same energy.

4.3 Labeling of regions and boundaries

At first sight, it may seem that an optimal solution of AMWC can also be obtained by first running the *Multi-way cut* algorithm using the original label set $\mathcal{L} = \{0, \dots, k-1\}$ to obtain a segmentation into regions $R_1, \dots, R_n$ and then to run the *multicut* algorithm for each region R_i separately to obtain an internal partitioning. We give two toy examples in the supplementary material which show that this decomposition is not possible.

5 Experiments

Here, we show qualitatively when the AMWC model is useful and when it is not.

2D Segmentation. We consider foreground/background segmentation problems with $\mathcal{L} = \{0, 1\}$ and $\mathcal{A} = \{0\}$. For all images, we first compute an oversegmentation into superpixels using a seeded watershed algorithm on an elevation map combining gradient magnitude and the output of the generalized

probability of boundary (gPb) detector from [29]. Then, the superpixel adjacency graph $G = (\mathcal{V}, \mathcal{E})$ defines the structure of the Conditional Random Field (1). For each edge $(i, j) \in \mathcal{E}$ which represents the shared boundary between superpixels i and j, we compute the mean boundary probability, $\mathbf{f}_{ij}$, as given by the gPb detector. Weights $\boldsymbol{w}_{ij} \in \mathbb{R}$ are then obtained as follows

$$\boldsymbol{w}_{ij} = \log \frac{1 - \mathbf{f}_{ij}}{\mathbf{f}_{ij}} + \log \frac{1 - \beta}{\beta}, \tag{6}$$

where $\beta \in [0, 1]$ is a hyper-parameter giving the prior boundary probability. As region appearance terms, we use the output of the object saliency detector [32], or region appearance terms derived from manually placed object bounding boxes. Again, there is a bias hyper-parameter α which gives the prior foreground probability. Finally, we write the energy function (1) as

$$\underset{\boldsymbol{l} \in \mathcal{L}^{|\mathcal{V}|}}{\operatorname{argmin}} \left\{ \gamma \cdot \sum_{i \in \mathcal{V}} E_i(l_i) + \sum_{(i,j) \in \mathcal{E}} E_{ij}^{\boldsymbol{w}}(l_i, l_j) \right\}, \tag{7}$$

where the hyper-parameter γ weights unary and pairwise terms.

Fig. 5 gives examples where the AMWC formulation can help and where it cannot. Column (a) shows the original images, taken from benchmark datasets [31,30,10]. Column (b) shows foreground maps either obtained using [32] or given as manual bounding box annotations (rows 1-4). The boundary probability for superpixel edges is visualized in column (c). The Multicut algorithm, column (d), ignores the region appearance terms and gives a decomposition into regions which are shown with random colors. Column (e) shows the visually best solution of a standard graph cut model. Finally, columns (f) and (g) show results for both the Multi-way cut model as well as the AMWC model with $\mathcal{A} = \{\text{background}\}$.

For each row, hyper-parameters α, β, γ – shared among Multicut, Multi-way cut and Asymmetric Multi-way Cutmodels – were chosen to give reasonable and comparable results for these three algorithms.

For the pedestrian detection (Fig. 5, rows 1–4) both the local appearance and edge detection terms are weak. The latter leads to regions which "leak" into the background when using multicut segmentation. Classical graph cut methods suffer severely from the very rough local data terms and show many artifacts caused by shrinking bias. While Multi-way cut generates foreground-artifacts due to strong edges present in the *background*, the proposed AMWC model can handle these by introducing closed contours in the background at these locations. Here, AMWC performs best, even though strong within-foreground contours produce some artifacts (row 2 and 4).

For the examples using saliency detection as the region appearance model (Fig. 5, rows 5–8), both the region and edge terms are more confident. However, graph cut still shows shrinking artifacts and MWC sometimes "hallucinates" foreground-regions in the background, though this is no longer as significant as with the weaker data terms in rows 1–4.

The run times for Multicut, Multi-way cut and AMWC are comparable with less than 10 seconds per image.

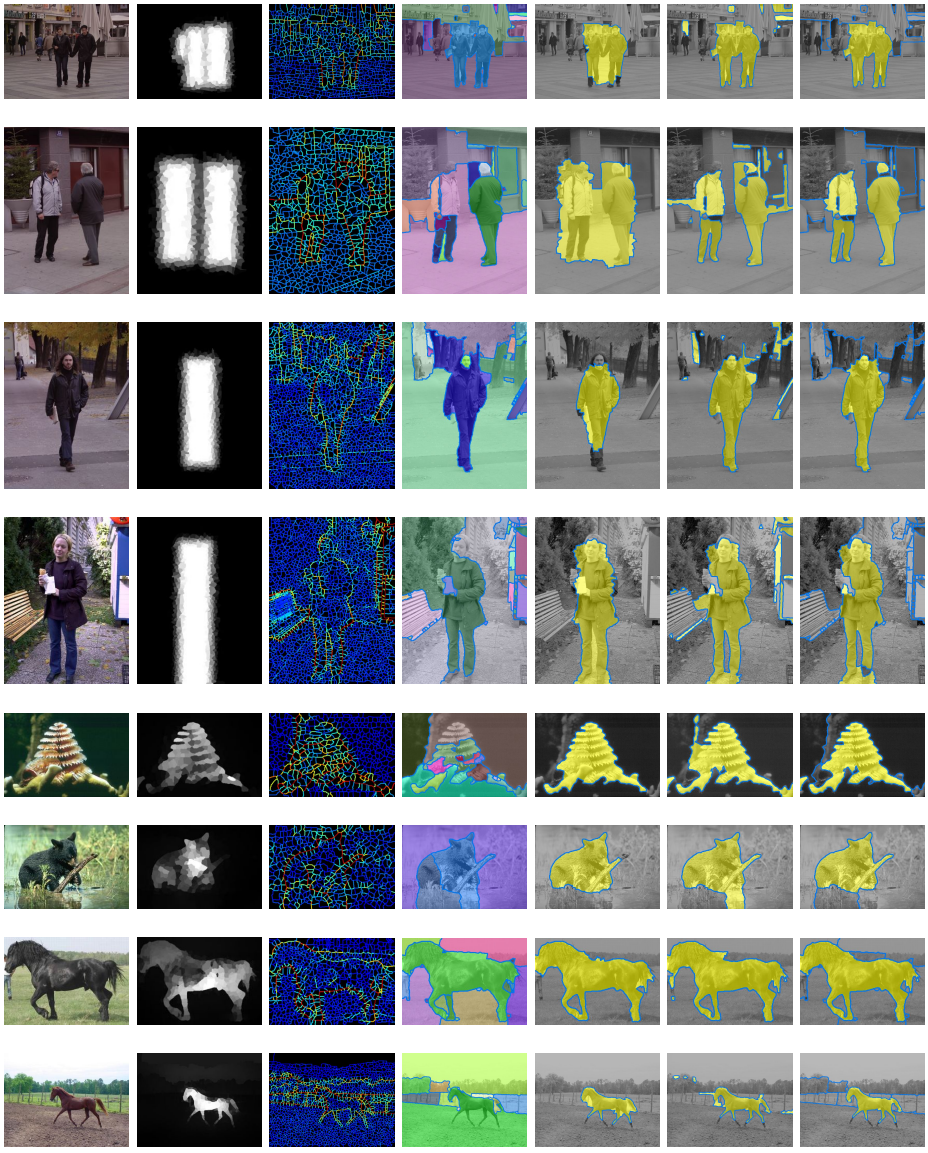

(a) image (b) region terms (c) boundary terms (d) MC (e) GC (f) MWC (g) AMWC

Fig. 5: Example segmentations of various pairwise Conditional Random Fields Models. Multicut (MC) is uninformative as regards to category predictions. Graphcut (GC) suffers from shrinking bias. Multi-way cut (MWC) may produce spurious regions induced by strong boundary evidence. AMWC is a joint semantic labeling and partitioning methods that suffers from none of the above. However, it requires specification of the asymmetry (which classes can be partitioned).

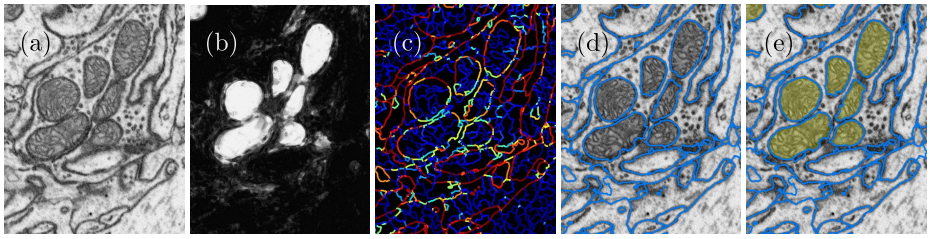

Fig. 6: From the raw volume image (a), a mitochondria versus background probability map (b) was obtained using ilastik [34]. Boundary probabilities are obtained from a random forest classifier using local edge features (c). The multicut algorithm (d) yields a decomposition of the volume image into segments, relying only on the boundary evidence in (c). Finally, Asymmetric Multi-way Cut is the first method that can jointly find a similar segmentation to the one in (d), while at the *same* time labeling regions by their appearance.

3D Data. To better understand the functioning principles of the brain, researchers use electron microscopy techniques to obtain ultra high-resolution isotropic volumetric images of brain tissue. These exquisite images contain densely packed neurons, which can only be distinguished by their separating membranes. In addition, intra-cellular structures such as mitochondria are visible. In the past, automatic methods to segment mitochondria [28], synapses [26,8], as well as the segmentation of the volume into distinct neurons, e.g. [15,4,16] have been considered separately. The AMWC allows to combine both problems into a single, joint model. We consider two classes, mitochondrion m and cytoplasm c. We set $\mathcal{A} = \{c\}$. Fig. 6 shows results on the FIBSEM dataset from [4], where we combine learned pixel-wise mitochondrion probabilities with learned membrane (boundary) probabilities.

This is the first time that a *joint* partitioning into neurons and labeling of intracellular structures becomes possible.

6 Conclusion

We have introduced a new sub-class of non-submodular pairwise multi-label Conditional Random Fields with Potts potentials in which (i) label transitions can be both discouraged as well as encouraged and (ii) some labels, such as background, are allowed to have internal boundaries. As a consequence, strong boundaries within these classes can be naturally accommodated by a further partitioning. The proposed model can be solved exactly using an extension of an existing Multi-way cut solver. We expect this model to be most useful in a regime where regional appearance terms and boundary evidence are both noisy, but (and this is crucial) complementary. In this setting, the present paper offers a principled unified approach to simultaneous labeling and partitioning.

References

1. Alush, A., Goldberger, J.: Break and conquer: Efficient correlation clustering for image segmentation. In: 2nd International Workshop on Similarity-Based Pattern Analysis and Recognition (2013) 4
2. Andres, B., Beier, T., Kappes, J.H.: OpenGM: A C++ library for discrete graphical models. ArXiv e-prints (2012), http://hci.iwr.uni-heidelberg.de/opengm2 3
3. Andres, B., Kappes, J.H., Beier, T., Kothe, U., Hamprecht, F.A.: Probabilistic image segmentation with closedness constraints. In: ICCV. pp. 2611–2618. IEEE (2011) 1, 4, 5
4. 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: ECCV, pp. 778–791. Springer (2012) 4, 10
5. Arbelaez, P.: Boundary extraction in natural images using ultrametric contour maps. In: Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on. pp. 182–182. IEEE (2006) 1
6. Bagon, S., Galun, M.: Large scale correlation clustering optimization. CoRR abs/1112.2903 (2011) 3, 4
7. Bansal, N., Blum, A., Chawla, S.: Correlation clustering. Machine Learning 56(1-3), 89–113 (2004) 4
8. Becker, C., Ali, K., Knott, G., Fua, P.: Learning context cues for synapse segmentation in em volumes. In: MICCAI. Springer (2012) 10
9. Beier, T., Kroeger, T., Kappes, J.H., Koethe, U., Hamprecht, F.: Cut, glue & cut: A fast, approximate solver for multicut partitioning. In: CVPR (2014) 4
10. Borenstein, E., Ullman, S.: Class-specific, top-down segmentation. In: Computer Vision—ECCV 2002, pp. 109–122. Springer (2002) 8
11. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. Pattern Analysis and Machine Intelligence 26(9), 1124–1137 (2004) 3
12. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. Pattern Analysis and Machine Intelligence 23(11), 1222–1239 (2001) 2, 3, 4
13. Chopra, S., Rao, M.: The partition problem. Mathematical Programming 59(1-3), 87–115 (1993), http://dx.doi.org/10.1007/BF01581239 4, 5
14. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html 1
15. Jain, V., Bollmann, B., Richardson, M., Berger, D.R., Helmstaedter, M.N., Briggman, K.L., Denk, W., Bowden, J.B., Mendenhall, J.M., Abraham, W.C., et al.: Boundary learning by optimization with topological constraints. In: CVPR. IEEE (2010) 10
16. Jurrus, E., Watanabe, S., Giuly, R.J., Paiva, A.R., Ellisman, M.H., Jorgensen, E.M., Tasdizen, T.: Semi-automated neuron boundary detection and nonbranching process segmentation in electron microscopy images. Neuroinformatics 11(1), 5–29 (2013) 10
17. Kappes, J.H., Andres, B., Hamprecht, F.A., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Lellmann, J., Komodakis, N., Rother, C.: A comparative study of modern inference techniques for discrete energy minimization problems. In: CVPR (2013) 2, 3, 4

18. Kappes, J.H., Speth, M., Andres, B., Reinelt, G., Schnörr, C.: Globally optimal image partitioning by multicuts. In: EMMCVPR. pp. 31–44. Springer (2011) 3, 4, 5
19. Kappes, J.H., Speth, M., Reinelt, G., Schnörr, C.: Higher-order segmentation via multicuts. CoRR abs/1305.6387 (2013) 3, 4, 5, 6
20. Kim, S., Nowozin, S., Kohli, P., Yoo, C.: Task-specific image partitioning. Transactions on Image Processing (2012) 4
21. Kohli, P., Shekhovtsov, A., Rother, C., Kolmogorov, V., Torr, P.: On partial optimality in multi-label MRFs. In: ICML. pp. 480–487. ACM (2008) 4
22. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. Pattern Analysis and Machine Intelligence, IEEE Transactions on 28(10), 1568–1583 (2006) 3
23. Kolmogorov, V., Zabin, R.: What energy functions can be minimized via graph cuts? Pattern Analysis and Machine Intelligence 26(2), 147–159 (2004) 2, 3
24. Komodakis, N., Paragios, N., Tziritas, G.: MRF energy minimization and beyond via dual decomposition. Pattern Analysis and Machine Intelligence, IEEE Transactions on 33(3), 531–552 (2011) 3
25. Komodakis, N., Tziritas, G.: Approximate labeling via graph cuts based on linear programming. Pattern Analysis and Machine Intelligence, IEEE Transactions on 29(8), 1436–1453 (2007) 4
26. Kreshuk, A., Straehle, C.N., Sommer, C., Koethe, U., Cantoni, M., Knott, G., Hamprecht, F.A.: Automated detection and segmentation of synaptic contacts in nearly isotropic serial electron microscopy images. PloS one 6(10), e24899 (2011) 10
27. Kröger, T., Mikula, S., Denk, W., Köthe, U., Hamprecht, F.A.: Learning to Segment Neurons with Non-local Quality Measures. In: MICCAI. Springer (2013) 4
28. Lucchi, A., Smith, K., Achanta, R., Knott, G., Fua, P.: Supervoxel-based segmentation of mitochondria in EM image stacks with learned shape features. Transactions on Medical Imaging 31(2), 474–486 (2012) 10
29. Maire, M., Arbeláez, P., Fowlkes, C., Malik, J.: Using contours to detect and localize junctions in natural images. In: CVPR. pp. 1–8. IEEE (2008) 4, 7
30. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: ICCV. vol. 2, pp. 416–423. IEEE (2001) 1, 8
31. Opelt, A., Pinz, A., Fussenegger, M., Auer, P.: Generic object recognition with boosting. Pattern Analysis and Machine Intelligence 28(3), 416–431 (2006) 8
32. Perazzi, F., Krähenbühl, P., Pritch, Y., Hornung, A.: Saliency filters: Contrast based filtering for salient region detection. In: CVPR. pp. 733–740. IEEE (2012) 8
33. Rother, C., Kolmogorov, V., Blake, A.: Grabcut: Interactive foreground extraction using iterated graph cuts. In: Transactions on Graphics. vol. 23, pp. 309–314. ACM (2004) 2, 3, 4
34. Sommer, C., Straehle, C., Kothe, U., Hamprecht, F.A.: Ilastik: Interactive learning and segmentation toolkit. In: Symposium on Biomedical Imaging: From Nano to Macro. pp. 230–233. IEEE (2011) 10
35. Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C.: A comparative study of energy minimization methods for Markov random fields. In: ECCV, pp. 16–29. Springer (2006) 3
36. Vicente, S., Kolmogorov, V., Rother, C.: Graph cut based image segmentation with connectivity priors. In: CVPR. pp. 1–8. IEEE (2008) 2, 3, 4
37. Yarkony, J., Ihler, A., Fowlkes, C.C.: Fast planar correlation clustering for image segmentation. In: ECCV, pp. 568–581. Springer (2012) 1, 4