

Higher-Order Minimum Cost Lifted Multicuts for Motion Segmentation

Margret Keuper
Data and Web Science Group
University of Mannheim, Germany
keuper@uni-mannheim.de

Abstract

Most state-of-the-art motion segmentation algorithms draw their potential from modeling motion differences of local entities such as point trajectories in terms of pairwise potentials in graphical models. Inference in instances of minimum cost multicut problems defined on such graphs allows to optimize the number of the resulting segments along with the segment assignment. However, pairwise potentials limit the discriminative power of the employed motion models to translational differences. More complex models such as Euclidean or affine transformations call for higher-order potentials and a tractable inference in the resulting higher-order graphical models. In this paper, we (1) introduce a generalization of the minimum cost lifted multicut problem to hypergraphs, and (2) propose a simple primal feasible heuristic that allows for a reasonably efficient inference in instances of higher-order lifted multicut problem instances defined on point trajectory hypergraphs for motion segmentation. The resulting motion segmentations improve over the state-of-the-art on the FBMS-59 dataset.

1. Introduction

Motion segmentation, *i.e.* the task of segmenting all moving objects visible in a video, is a long-standing task in computer vision [5, 25, 30, 26, 34]. On a low level, it requires an accurate estimation and adequate comparison of the point-wise observable motion. At the same time, the number of independently moving objects has to be inferred and all points have to be correctly assigned to a motion segment. While spectral clustering approaches have been used traditionally in this field [5, 34, 13], correlation clustering, also referred to as *minimum cost multicut* problem, has recently proven successful on this task [20].

In fact, correlation clustering differs in the objective function from spectral clustering approaches since it does not prefer balanced cuts. At the same time, it directly optimizes for the right number of objects, such that there is no need for a separate model selection. As spectral clustering,

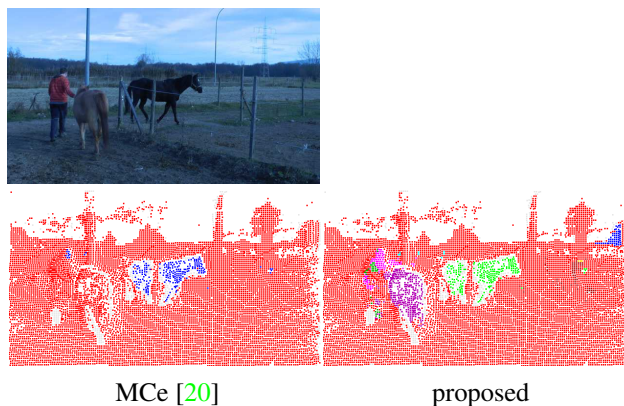


Figure 1. Motion segmentation result for a sequence with two consistently moving objects under scaling (the man and the horse next to him). The higher-order lifted multicut model allows for a more precise segmentation.

the objective function of the minimum cost multicut problem is defined by edge costs. Unlike the spectral clustering scenario, these edge costs can have any real value and thus explicitly model attractive and repulsive terms in the objective function.¹

The formulation of the minimum cost *lifted* multicut model in [21] facilitates the definition of such costs between any two nodes, even if they are not connected in the original graph. This generalization of the cost function allowed multicut formulations to match the state-of-the-art in image segmentation and has shown its benefits for multiple target tracking [35].

One of the key challenges in motion segmentation is to disambiguate between different objects moving according to the same motion model as illustrated in Fig. 1. This is comparable to the difficulty in image segmentation when different objects have similar color and texture. Lifted Multicuts can dissolve such ambiguities appropriately for image segmentation [21]. In this paper, we show that the same is

¹Variants of spectral clustering explicitly incorporating (soft) repulsive terms have been proposed in for example [8, 11].

true for motion segmentation. In Fig. 1, the man and the horse next to him can be segmented from the background and from each other although they move similarly.

However, pairwise potentials, such as the lifted multicut model from [21] allows to define, are very limited in their capacity to describe object motion. In fact, the Euclidean difference between two local motion descriptors such as optical flow vectors or point trajectories, measures how well the behavior of the two entities can be described by a single translational motion model. Depending on the motion recorded in the video, this simple model can yield good performance as shown *e.g.* in [31, 20]. However, when we want to segment objects from videos showing complex motion patterns caused for example by object or camera rotation, scaling motion or zooming, more complex motion models are needed. Transformations describing translation, rotation and scaling can be estimated from two motion vectors. Thus, for any three points, one can estimate how well their motion can be described by one Euclidean transformation. Edges that describe such motion differences are thus at least of order three. Affine motion differences can be described with edges of order four and to assign costs to differences in homographies, the minimum required edge-order is five.

In this paper, we propose the higher-order minimum cost *lifted* multicut model. We provide a rigorous definition along with a simple primal feasible heuristic that allows for practical applicability in computer vision. To the best of our knowledge, we are the first to look into this specific problem, which is specially suitable to formulate the motion segmentation objective. We demonstrate the benefit of the proposed model on the motion segmentation task using third-order edges and improve over the state-of-the-art on the FBMS-59 [31] dataset.

2. Related Work

Higher-order graph decompositions have been addressed prominently in computer vision since [2]. For example in the subsequent works [1, 38, 7, 30, 39, 10, 37], spectral approaches to the clustering of hypergraphs have been addressed. Specifically, [30, 39, 10] model higher-order motions. Zografos *et al.* [39] model 3D motions using group invariants and [10] model higher-order motion subspaces. The segmentation is then generated by projecting the resulting hyper-graph onto its primal graph and solving the spectral clustering problem there. Higher-Order Markov Random Field (MRF) and Conditional Random Field (CRF) models have been proposed for example in [12, 24, 33, 27].

Our formulation of higher-order minimum cost multicut problems is different from both previous approaches. In contrast to spectral clustering, the multicut formulation does not suppose any balancing criterion. Further, we directly infer segmentations from the hyper-graph without any projection onto its primal graph.

In contrast to MRFs, the proposed approach allows higher-order edges to connect vertices globally, violating the Markov property. Further, MRFs and CRFs aim at inferring a node labeling with labels given a priori while multicut approaches aim at inferring an edge labeling yielding an optimal number of segments.

However, most previous works on minimum cost multicut in computer vision focus on the definition of problems with pairwise potentials [17, 21, 20, 36]. The exception is the model first presented by [22, 23] and extensively studied in [18]. Therein, higher-order costs are used for image segmentation on superpixel graphs with pairwise neighborhood connectivity. In [18] an implementation of an ILP solver is proposed for small problem instances.

In contrast, we propose a higher-order *lifted* multicut model, which allows the definition of edge costs of any order in the lifted graph as well as in its connectivity defining subgraph. Further, both models differ in that the proposed model allows edges of any order to define the graph connectivity, while in [22, 23, 18], the graph connectivity is defined exclusively by pairwise edges.

Lifted multicut on simple graphs have been proposed in [21] along with a primal feasible heuristic to generate solutions. A different heuristic solver for the problem has been proposed in [4]. Here, we generalize the solver from [21] to facilitate the inference in higher-order problems.

In our setup, motion segmentation is cast as a point trajectory grouping problem. In a similar way, it has previously been addressed in [5, 25, 30, 26, 34, 31, 32, 16, 20]. From sparse motion segmentations, framewise dense segmentations can be computed for example by the variational approaches from [29, 28].

3. The Higher-Order Lifted Multicut Problem

Here, we define an optimization problem, the Higher-Order Minimum Cost Lifted Multicut Problem in analogy to the Lifted Multicut Problem proposed in [21]. Its feasible solutions correspond to the decompositions of a hypergraph and its objective function can assign, for any set e of nodes, a real valued cost to all decompositions for which all nodes $v \in e$ are in the same component. A component of a hypergraph is any non-empty subgraph that is node-induced and connected by edges of any order. A decomposition of a hypergraph is any partition Π of the node set such that, for every $V' \in \Pi$, the subgraph induced by V' is a connected component of the hypergraph. An instance of the problem is then defined with respect to an undirected hypergraph $G = (V, F)$ and an additional set of hyperedges F' connecting sets of nodes that are not necessarily neighbors in G . Practically, we can distinguish between pairwise edges $e \in F_p \cup F'_p$ with $|e| = 2$ and higher-order edges $e \in F_h \cup F'_h$ with $|e| > 2$. In our experiments, we limit F and F' to edges e of order $|e| = 2$ or 3. For every edge

$e \in E = F \cup F'$, a cost $c_e \in \mathbb{R}$ is assigned to all feasible solutions for which all nodes connected by e are in the same component.

Then, we can define a feasible set $Y_E \subseteq \{0, 1\}^E$ whose elements $y \in Y_E$ are 01-labelings of all edges $E = F \cup F'$. The feasible set is defined such that: (1) the feasible solutions $y \in Y_E$ relate one-to-one to the decompositions of the hypergraph G . (2) for every edge $e \in E$, $y_e = 1$ if and only if all nodes connected by e are in the same component of G . It can be defined rigorously by linear inequality constraints on higher-order edges [22, 23] (2) and (3), cycle inequalities [9] (4) and path and cut constraints (5) and (6) [21].

Definition 1 For any hypergraph $G = (V, F)$, any $F' \subseteq \bigcup_k \binom{V}{k} \setminus F$ and any $c : E := F \cup F' \rightarrow \mathbb{R}$, the 01 linear program written below is an instance of the *Higher-Order Minimum Cost Lifted Multicut Problem* w.r.t. G , E and c .

$$\min_{y \in Y_E} \sum_{e \in E} c_e y_e \quad (1)$$

with $Y_E \subseteq \{0, 1\}^E$ the set of all $y \in \{0, 1\}^E$ with

$$\forall e, e_h \in E \mid e \subset e_h : y_e \leq y_{e_h} \quad (2)$$

$\forall e_h \in E \mid \bar{G} = (e_h, e \in E \mid e \subset e_h)$ is connected :

$$(1 - y_{e_h}) \leq \sum_{e \in E \mid e \subset e_h} (1 - y_e) \quad (3)$$

$$\forall C \in \text{cycles}(G) \forall e \in C : (1 - y_e) \leq \sum_{e' \in C \setminus \{e\}} (1 - y_{e'}) \quad (4)$$

$\forall e \in F, \forall vw \subset e \forall P \in vw\text{-paths}(G) :$

$$(1 - y_e) \leq \sum_{e' \in P} (1 - y'_{e'}) \quad (5)$$

$\forall e \in F, \forall vw \subset e \forall C \in vw\text{-cuts}(G) :$

$$y_e \leq \sum_{e' \in C} (y'_{e'}) \quad (6)$$

Note that, since higher-order edges $e \in F$ are connectivity defining, these conditions need to hold on cycles, vw -paths (any path in G , connecting v and w) and vw -cuts (any cut in G assigning v and w to distinct components) of any order. In the higher-order model from [22], it is sufficient to ensure the cycle inequalities (4) for cycles on pairwise edges. For the motion segmentation application, it is reasonable to allow higher-order edges to define the graph connectivity since most motion models can only be estimated on at least two nodes. If, in this case, for a pair of nodes v and w all higher-order edges with $\{v, w\} \subseteq e$ have label $y_e = 0$, there was no motion model that supports v and w belonging to the same object. Consequently, any grouping of a subset of e , for which the motion model can not be evaluated, to the same component, is uninformed and should thus be avoided. However, if the connectivity is defined from the projection of e onto a primal graph, this connection is possible and cost-neutral.

Algorithm 1: Kernighan-Lin Algorithm. The function `higher_order_update_bipartition` is given in Algorithm 2.

Data: weighted undirected lifted higher-order graph $G = (V, F \cup F', c)$, starting 01-edge labeling $y^{|F \cup F'|}$

Result: 01-edge labeling y

```

1  $t \leftarrow 1$  while  $t < \text{max\_iter}$  and not  $y^t = y^{t-1}$  do
2   foreach  $(a, b) \in \text{adjacent\_partitions}(y^{t-1})$  do
3      $y^t \leftarrow \text{higher\_order\_update\_bipartition}(G, a, b)$ 
4   foreach  $a \in \text{partitions}(y^t)$  do
5      $y^t \leftarrow \text{higher\_order\_update\_bipartition}(G, a, \emptyset)$ 

```

The definition of lifted edges allows to ensure a consistent motion segment connectivity. For example background motion can often only be correctly estimated from points with large spatial distance. In this case, the constraints (6) ensure an existing path between these background points. At the same time, defining the connectivity on sets of points in a spatial neighborhood allows to cut apart distinct but consistently moving objects.

In the following, we describe a primal feasible heuristic that allows to solve instances of the Higher-Order Lifted Multicut Problem in practice.

4. Higher-Order Lifted Kernighan-Lin Algorithm

The Kernighan-Lin [19] heuristic is known to work well for the balanced set partitioning problem. The original heuristic has been modified in [3] to disregard the size of the resulting partitions as well as to optimize for the number of components. In [21], it has been further adapted to work with sparse and lifted graphs; an additional joining move has been introduced for a more stable behavior with different initializations. Here, we further extend the algorithm to work with *higher-order* lifted graphs. While the outer loop of the algorithm, running over all pairs of neighboring components is only marginally affected, the major change concerns the inner iterations, where the gain of moving vertices is computed and successively updated. The algorithm takes as input an instance of the higher-order lifted multicut problem and an initial decomposition of G and outputs a decomposition of G whose higher-order lifted multicut has an objective value lower than or equal to that of the initial decomposition. As the basic KLj-Algorithm [21], it maintains, throughout its execution, a decomposition of G which is encoded as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The nodes $a \in \mathcal{V}$ of \mathcal{G} are components of G and its edges $ab \in \mathcal{E}$ connect all components a and b of G which are connected in G by an edge of arbitrary order.

Algorithm 1 starts from an initial decomposition provided as input. As KLj [21], in each iteration, it tries to

Algorithm 2: Function `higher_order_update_bipartition`

Data: weighted undirected lifted higher-order graph
 $G = (V, F \cup F', c)$, a pair of partitions A and B

Result: 01 - edge labeling y

```
1  $D^{A \cup B} = \text{compute\_differences}(F, F', A, B)$ 
2  $BDY \leftarrow \text{compute\_boundary\_nodes}(F, A, B)$ 
3  $\Delta_{join} \leftarrow \text{compute\_gain\_from\_joining}(F, F', A, B)$ 
4  $S^{|A \cup B|} = 0$  // cumulative gain
5  $M = []$  // empty move vector
6 ; for  $i \leftarrow 1$  to  $|BDY|$  do
7    $v^* \leftarrow \text{argmax}_v (D_{(A \cup B) \cap BDY})$ 
8    $M.\text{push\_back}(v^*)$ 
9   foreach  $e \in G.\text{edges}(v^*)$  do
10    if  $\text{same\_partition}(G.\text{nodes}(e) \setminus v^*)$  then
11      foreach  $w \in G.\text{nodes}(e) \setminus v^*$  do
12        if  $\text{same\_partition}(v^*, w)$  and
13           $|G.\text{nodes}(w) \setminus v^*| = 1$  then
14           $D_w \leftarrow D_w - 2c_e$ 
15        else if  $\text{same\_partition}(v^*, w)$  and
16           $|G.\text{nodes}(w) \setminus v^*| > 1$  then
17           $D_w \leftarrow D_w - c_e$ 
18        else if not  $\text{same\_partition}(v^*, w)$  and
19           $|G.\text{nodes}(w) \setminus v^*| = 1$  then
20           $D_w \leftarrow D_w + 2c_e$ 
21        else if not  $\text{same\_partition}(v^*, w)$  and
22           $|G.\text{nodes}(w) \setminus v^*| > 1$  then
23           $D_w \leftarrow D_w + c_e$ 
24        else if not  $\text{same\_partition}(v^*, w)$  then
25           $D_w \leftarrow D_w - c_e$ 
26
27    $S_i = S_{i-1} + D_{v^*}$ 
28    $BDY \leftarrow \text{update\_boundary}(v^*, F, A, B)$ 
29  $k \leftarrow \text{argmax}_i S_i$  // best number of moves
30 if  $\Delta_{join} > S_k$  and  $\Delta_{join} > 0$  then
31    $\text{join\_partitions}(y, A, B)$ ;
32 else if  $S_k > 0$  then
33    $\text{move\_nodes}(y, A, B, k)$ 
```

improve the current decomposition by one of these three transformations: (1) moving a set of nodes between two neighboring components, (2) moving a connected set of nodes from one component to a new component, (3) joining two neighboring components. The main operation of Alg. 1 is called “higher_order_update_bipartition”. Its input is the current decomposition and a pair $ab \in \mathcal{E}$ of neighboring components of G . The neighborhood in G is defined with

respect to edges $e \in F$ of any order. It evaluates transformations (1) and (3) for ab . Transformation (2) is assessed by executing “higher_order_update_bipartition” for each component and \emptyset .

In the operation “higher_order_update_bipartition” a sequence M of elementary transformations of the components a and b is constructed greedily such that at every consecutive move-operation increases the cumulative gain S maximally (or decreases it minimally) while preserving a feasible solution. Therefore, the operation “compute_differences” computes, at the beginning of each execution of “higher_order_update_bipartition”, for every element $w \in a \cup b$ the difference D_w in the objective function when w is moved between a and b . Then, an element on the current boundary between a and b with maximal difference is added to the sequence M . The differences D are updated according to Alg. 2, ll. 9-26.

If the objective value can be decreased by executing either the first $k \in \mathbb{N}_0$ elementary transformations or by joining the components a and b the optimal of these two operations is carried out.

While components are defined with respect to the graph $G = (V, F)$, differences in objective value are computed with respect to the graph $G' = (V, E)$ with $E = F \cup F'$.

As in KLj [21], the number of outer iterations of Alg. 1 is not bounded by a polynomial and we cannot give any guarantee for convergence. However, in practice, the algorithm converged in less than 50 iterations for the experiments described in Sec. 6.

5. Graph Construction

5.1. Point Trajectories

Point trajectories are spatio-temporal curves that describe the trajectory of a single object point in the image plane. They build the basis for many motion segmentation methods such as [5, 13, 14, 31, 20]. Here, we use the method from [5] to generate dense long-term point trajectories from precomputed optical flow. While we are aware that more recent optical flow algorithms allow for better motion segmentations [15], we build our trajectories from large displacement optical flow [6] to allow for a direct comparison to previous work. For a video of length N , [5] yields n point trajectories p_i with the maximum length N , where n depends on the desired sampling rate. Due to occlusions and mistakes in the optical flow estimation, most trajectories are significantly shorter than N , and some trajectories start after frame 1 to ensure even point sampling throughout the sequence.

5.2. Higher-Order Motion Models

Although this is not sufficient to accurately describe object motion in a 3D environment recorded with a possibly

moving camera, we restrict ourselves to edge potentials of order two and three for practical reasons. This allows to measure the difference of point motions according to Euclidean motion models, *i.e.* from the group of transformations describing translation, rotation and scaling in the 2D plane. This is a subset of the group of similarity transformations in the 2D plane, where reflections are excluded.

We further argue that in any case, the easiest model that can explain the motion of a set of points with a single transformation should be used. If two points are moving according to the same translational motion model, we can assume that they belong to the same object without looking at further points around them. Only if their motion is different according to a purely translational model, looking at more complex motion models adds information. This results in a *motion-adaptive graph construction* strategy.

Motion-Adaptive Graph Construction We propose to construct the higher-order graph G depending on the pairwise costs computed from motion differences. The algorithm is described in Alg. 3. For any pair of trajectories, we compute their cost of belonging to the same translational motion model. Only if this cost is repulsive, we look at all further points to compute for every three-tuple the cost of belonging to the same motion model for translation, rotation and scaling. The respective third-order edges are inserted along with their costs.

This strategy allows to integrate second and third order potential without losing model capacity. Further, compared to generating the full graph with higher-order potentials, it yields a huge complexity reduction in practice.

Lifted Graph Construction To construct higher-order lifted graphs $G' = (V, F \cup F')$, we precompute for every trajectory the set of its 12 spatially nearest neighbors \mathcal{N} . The edge set F is the subset of the full edge set E , computed according to Alg. 3, that contains exactly all pairwise edges $e_{ij} \in E$ for which at least one of the following three conditions holds: (1) $p_i \in \mathcal{N}(p_j)$, (2) $p_j \in \mathcal{N}(p_i)$ (3) the maximum spatial distance between p_i and p_j is below 40 pixels.

Second Order Costs Second order costs are computed from pairwise differences on point trajectories. We compute such differences only for trajectories which have at least two frames in common. Since it has proven successful in previous work [20], we compute such differences based on motion, color and spatial distance cues. As suggested by [31], we define the pairwise motion difference of two trajectories at time t as

$$d_t^m(p_i, p_j) = \frac{\|\partial_t p_i - \partial_t p_j\|}{\sigma_t}. \quad (7)$$

Algorithm 3: Motion Adaptive Graph Construction.

Data: set of point trajectories V with $p_k \in V$ with $k \in \{1 \dots n\}$
Result: weighted undirected higher-order graph $G = (V, E)$, cost vector c

```

1  $G \leftarrow (V, E = \emptyset)$ 
2  $c = []$ 
3 foreach  $p_i$  and  $p_j \in V$  do
4    $c_{ij} \leftarrow \text{compute\_translational\_motion\_cost}(p_i, p_j)$ 
5    $E \leftarrow E \cup (p_i, p_j)$  // insert pairwise edge
6   if  $c < 0$  then
7      $c.\text{push\_back}(c_{ij})$ 
8   else
9      $c.\text{push\_back}(0)$ 
10    foreach  $p_k \in V \setminus \{p_i, p_j\}$  do
11       $c_{ijk} \leftarrow$ 
12         $\text{compute\_higher\_order\_motion\_cost}(p_i, p_j, p_k)$ 
13       $E \leftarrow E \cup (p_i, p_j, p_k)$  // insert
        higher-order edge
13       $c.\text{push\_back}(c_{ijk})$ 

```

Here, $\partial_t p_i$ and $\partial_t p_j$ are the partial derivatives of p_i and p_j with respect to the time dimension and σ_t is the variation of the optical flow as defined in [31]. The motion distance of two trajectories is defined by the maximum over time

$$d^m(p_i, p_j) = \max_t d_t^m(p_i, p_j). \quad (8)$$

As proposed in [20] color and spatial distances d^c and d^s are computed as average distances over the common lifetime of two trajectories. These three cues are combined non-linearly to compute the costs

$$c_{ij} = \begin{matrix} -\max(\bar{\theta}_0 & +\theta_1 d^m(p_i, p_j) + \\ & \theta_2 d^s(p_i, p_j) + \theta_3 d^c(p_i, p_j) \\ & \theta_0 & +\theta_1 d^m(p_i, p_j) \end{matrix} \quad (9)$$

with weights and intercept values θ as proposed in [20].

Third Order Costs We compute third order motion differences as proposed in [30]. For any two trajectories p_i and p_j coexisting from time t to t' , we estimate the Euclidean motion model $\mathcal{T}_{ij}(t)$, consisting of rotation R_α , translation $v := (v_1, v_2)^\top$ and scaling s as

$$\begin{aligned} \alpha &= \arccos \left(\frac{(p_i(t') - p_j(t'))^\top (p_i(t) - p_j(t))}{\|p_i(t') - p_j(t')\| \cdot \|p_i(t) - p_j(t)\|} \right) \\ s &= \frac{\|p_i(t') - p_j(t')\|}{\|p_i(t) - p_j(t)\|} \\ v &= \frac{1}{2} (p_i(t') + p_j(t') - s R_\alpha (p_i(t) + p_j(t))) \end{aligned} \quad (10)$$

The distance to any third trajectory p_k existing from t to t' can then be measured by $d_{ij}^t(p_k) = \|\mathcal{T}_{ij}(t)p_k(t) - p_k(t')\|$. For numerical reasons, $d_{ij}^t(p_k)$ is normalized by

$$\gamma_{ij}^t = \frac{1}{\sigma_t} \left(\frac{1}{2} \left(\frac{\|p_i(t) - p_j(t)\|}{\|p_i(t) - p_k(t)\|} + \frac{\|p_i(t) - p_j(t)\|}{\|p_j(t) - p_k(t)\|} \right) \right)^{\frac{1}{4}}, \quad (11)$$

with σ_t being the optical flow variation as in (7).

To render distances symmetric, [30] propose to consider the maximum $d_{\max}^t(i, j, k) = \max(\gamma_{ij}^t d_{ij}^t(p_k), \gamma_{ik}^t d_{ik}^t(p_j), \gamma_{jk}^t d_{jk}^t(p_i))$, which yields an over-estimation of the true distance. While this is unproblematic in a spectral clustering scenario, where distances are used to define positive point affinities, it can lead to problems in the multicut approach. Over-estimated distances lead to under-estimated join probabilities and thus eventually to switching the sign of the cost function towards repulsive terms. To avoid this effect, we compute both $d_{\max}^t(i, j, k)$ and, analogously, $d_{\min}^t(i, j, k)$. For both, we compute the maximum motion distance over the common lifetime of p_i, p_j and p_k as $d_{\max}(i, j, k) = \max_t d_{\max}^t(i, j, k)$ and $d_{\min}(i, j, k) = \max_t d_{\min}^t(i, j, k)$. We evaluate the costs $c(d_{\max}(i, j, k))$ and $c(d_{\min}(i, j, k))$ for both distances as $c(d) = \theta_0 + \theta_1 d$ and compute the final edge costs

$$c_{ijk} = \begin{cases} c(d_{\min}(i, j, k)) & \text{if } c(d_{\max}(i, j, k)) > 0 \\ c(d_{\max}(i, j, k)) & \text{if } c(d_{\min}(i, j, k)) < 0 \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

Thus, we make sure not to set any costs for edges whose underlying motion is controversial. Here, we set $\theta_0 = -1$ and $\theta_1 = 0.08$ manually.

Implementation Details In practice, we insert pairwise edges e_{ij} in G and G' only if the spatial distance between p_i and p_j is below 100 pixel. This is in analogy to [20] and due to the fact that for nearby points, the approximation of the true motion by a simplified model is usually better than for points at a large distance. Also, since the number of pairwise edges increases quadratically with the maximal spatial distance, this heuristic decreases the computational load significantly. For the same reason, we introduce an edge sampling strategy for third order edges. For every three-tuple of points, we compute the maximum pairwise distance d . From all three-tuples with $20 < d < 300$, we randomly sample $\frac{100}{d^2}\%$, while we insert all edges e_{ijk} with $d \leq 20$. This also prevents from a too strong imbalance of long range edges over short range edges.

6. Experiments

We evaluate the proposed higher-order lifted multicut model on the motion segmentation benchmark FBMS-59

| Set A (29 sequences) | P | R | F | O |
|----------------------|--------|--------|---------------|--------------|
| SC [31] | 85.10% | 62.40% | 72.0% | 17/65 |
| Higher-Order SC [30] | 81.55% | 59.33% | 68.68% | 16/65 |
| MC [20] | 84.94% | 71.22% | 77.48% | 23/65 |
| HO MC (ours) | 83.51% | 75.54% | 79.33% | 28/65 |
| Set B (30 sequences) | P | R | F | O |
| SC [31] | 79.61% | 60.91% | 69.02% | 24/69 |
| Higher-Order SC [30] | 82.11% | 64.67% | 72.35% | 27/69 |
| MC [20] | 82.87% | 69.89% | 75.83% | 27/69 |
| HO MC (ours) | 83.62% | 68.83% | 75.51% | 27/69 |

Table 1. Segmentation results on the FBMS-59 dataset on Set A (top) and Set B (bottom). We report **P**: average precision, **R**: average recall, **F**: F-measure and **O**: extracted objects with $F \geq 75\%$. All results are computed for sparse trajectory sampling at 8 pixel distance. Our result **HO MC** is computed on the non-lifted purely higher-order model to allow for a direct comparison to the listed competing methods.

[31], which is an extended version of the BMS-26 benchmark from Brox-Malik [5]. It contains 59 sequences of varying length (from 19 to 800 frames) and diverse content and motion. It provides manual annotations for all moving objects in the videos for every 20th frame. To allow for training, the dataset has been split into two subsets of 29 and 30 sequences for training and testing, respectively. While we agree that training all model parameters is highly desirable, we did not do so. This is due to the fact that (1) neither of the state-of-the-art methods [30, 31, 20] is training-based and (2), the training set, with 29 sparsely annotated sequences, is rather small. Thus, to avoid confusion, we hence denote the training split by *Set A* and the test split by *Set B*.

Evaluation To assess the capacity of our model components, we first evaluate a purely higher-order non-lifted version of our model. In this model, all pairwise costs are removed and all edges are connectivity defining. We compare this simple model to [31, 30] and the purely motion-based version of [20]. While [31] and [20] only consider translational motion, the affinities in [30] are defined most similarly to our higher-order costs. As the proposed approach, [20] formulate a multicut problem while [31, 30] follow a spectral clustering approach. The results are given in Tab. 1 in terms of precision, recall, f-measure and the number of extracted objects. Precision and recall are not directly comparable but they can serve as cues for under- or over-segmentation. The f-measure is an aggregate of both. From Tab. 1, we can observe that our higher-order lifted multicut model outperforms the higher-order spectral clustering method from [30] by about 10% on Set A and 3.5% on Set B. While there is a clear improvement on both sets,

| Set A (29 sequences) | P | R | F | O |
|----------------------|--------|--------|---------------|--------------|
| MCe [20] | 86.73% | 73.08% | 79.32% | 31/65 |
| HOPMC | 86.45% | 76.28% | 81.05% | 31/65 |
| AOMC | 83.91% | 76.44% | 80.00% | 33/65 |
| Lifted HOPMC | 88.19% | 77.10% | 82.27% | 31/65 |
| Lifted AOMC | 86.83% | 77.79% | 82.06% | 32/65 |
| Set B (30 sequences) | P | R | F | O |
| MCe [20] | 87.88% | 67.7% | 76.48% | 25/69 |
| HOPMC | 85.83% | 68.74% | 76.34% | 26/69 |
| AOMC | 84.19% | 72.64% | 77.99% | 26/69 |
| Lifted HOPMC | 85.88% | 69.24% | 76.68% | 26/69 |
| Lifted AOMC | 87.77% | 71.96% | 79.08% | 25/69 |

Table 2. Segmentation Results on FBMS-59 on Set A (top) and Set B (bottom). We report **P**: average precision, **R**: average recall, **F**: F-measure and **O**: extracted objects with $F \geq 75\%$. All results are computed for sparse trajectory sampling at 8 pixel distance.

the imbalance is remarkable. A similarly remarkable imbalance can be observed when comparing the performance of the two spectral clustering methods [31] and [30]. The higher-order model [30] yields a lower f-measure on Set A compared to [31]. Yet it outperforms [31] on Set B by about 3%. This indicates that the motion statistics in both splits are significantly different.

When we compare our higher-order model to the pairwise minimum cost multicut model from [20], we can observe an improvement on Set A. On Set B, both model perform almost equally in terms of f-measure.

In Tab. 2, we show the evaluation of our higher-order multicut model with the motion-adaptive order, denoted AOMC (compare Alg. 3). This model has access to similar pairwise cues as the motion and color-based version from [20], denoted MCe. It has as well access to the higher-order motion cues from Eq. (12).

As a sanity check for the motion-adaptive graph construction, we generate graphs that simply contain all pairwise costs c_{ij} as well as all third order edges with costs c_{ijk} without any adaptation w.r.t. the costs. We denote this additive model by HOPMC (higher-order + pairwise multicut). On Set A, both variants improve over MCe [20] while on Set B, HOPMC performs similarly as MCe [20] and is outperformed by AOMC by about 1.5% in f-measure.

Lifted versions of both types of proposed problems (HOPMC and AOMC) yield a further improvement on both sets w.r.t. the respective non-lifted model. However, on Set B, the segmentation quality of Lifted HOPMC does not significantly outperform the one of MCe. In contrast, the proposed Lifted AOMC consistently outperforms all competing methods and baselines.

Fig. 3 gives an example of the segmentation quality under scaling. In the *horses05* sequence, scaling is caused by

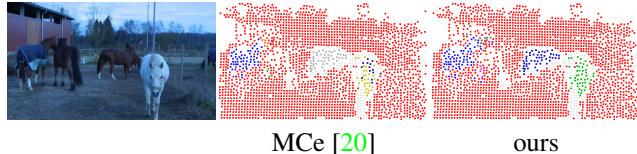


Figure 3. The scaling motion of the white horse moving towards the camera causes over-segmentation with a simple motion model [20]. With the proposed Lifted AOMC, this can be avoided.

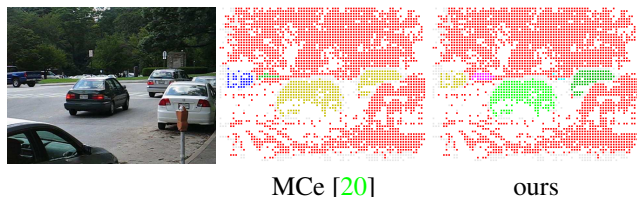


Figure 4. The two cars in front move to the same direction, leading an assignment to the same cluster with the non-lifted multicut approach [20]. The Lifted AOMC can assign the different cars to distinct segments.

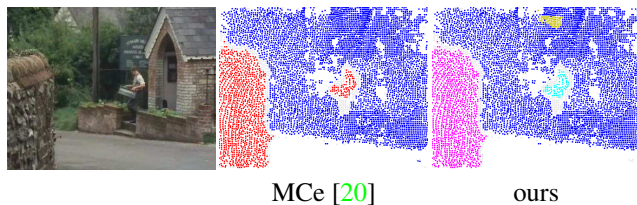


Figure 5. Due to camera motion the person and the wall are assigned to the same cluster with the non-lifted multicut approach [20]. The Lifted AOMC allows for correct segmentation.

the motion of the white horse towards the camera. This causes over-segmentation in the competing method MCe [20], which can not handle higher-order motion models. With the proposed Lifted AOMC, the segmentation can be improved.

Fig. 4 and 5 both show examples, where the same label is assigned to distinct objects that move similarly. In the *cars2* sequence in Fig. 4, this is indeed due to similar real world object motion, whereas, in the *marple10* sequence, the effect is due to camera motion and the scene geometry. In both cases, the formulation of the Lifted AOMC problem allows to tell the distinct objects apart. However, in the *marple10* sequence (Fig. 5), we can observe a spurious segment in the background, which is probably caused by unprecise flow estimation.

In Fig. 6, we show an example of the *goats01* sequence. Here, the head and body of the goat in front are segmented into a distinct components by the pairwise method [20], because of the expressed articulated motion. Although our third order model can not explicitly handle articulation, the over-segmentation can be fixed in this case.



Figure 2. Samples of our Lifted AOMC segmentation results densified by [29]. Even for articulated motion, our segmentations show little over-segmentation.

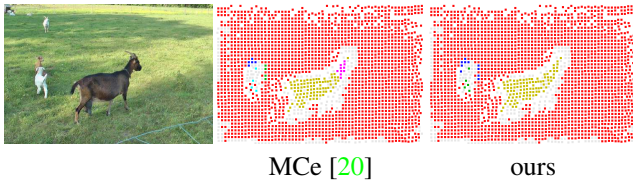


Figure 6. The articulated motion causes over-segmentation in [20]. The Lifted AOMC performs better.

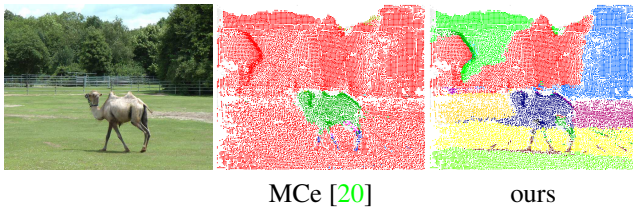


Figure 7. Failure case. The dominant camera motion causes strong over-segmentation with the proposed method. Here, our third order model can not model the motion appropriately.

In contrast, Fig. 7 shows a failure case of the proposed method. Due to the dominant camera motion in a scene with complex geometry, the Euclidean motion model fits particularly badly. Thus our model leads to the segmentation of the scene into its depth layers, and thus to strong over-segmentation.

Several examples of pixel-segmentations computed from our sparse segmentation using [29] are given in Fig. 2. The densified segmentations look reasonable. On the bear example, the articulated leg motion still causes some over-segmentation. One of the horses in the *horses05* sequence is missed. However, even small object such as the tray in the *marple12* sequence or the phone in the *marple13* sequence can be correctly segmented.

Scalability Last, we want to evaluate our proposed heuristic for the higher-order minimum cost lifted multicut problems (compare Alg. 1) in terms of computation times. Fig. 8 shows a plot of the computation times of our full

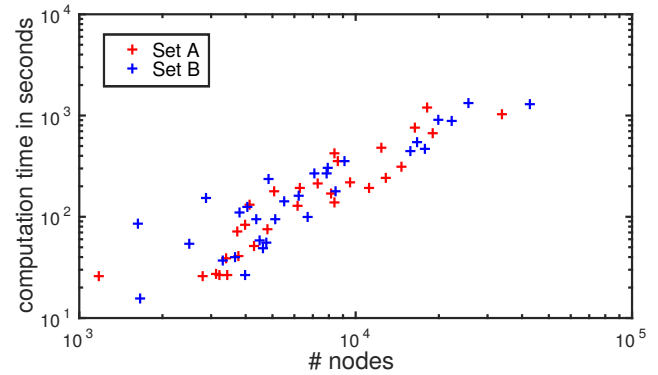


Figure 8. Computation times in logarithmic scale of the problem instances from Set A and Set B with respect to the number of nodes in logarithmic scale, *i.e.* clustered trajectories.

pipeline on Set A and B w.r.t. the number of clustered point trajectories in logarithmic scale. The runtime distribution indicates linear runtime behavior. However, the number of large problem instances is too small to make any claim. Yet, the plot shows that heuristic solutions can be generated in a few minutes for most instances.

7. Conclusion

In this paper, we have presented a generalization of the minimum cost lifted multicut problem to a higher-order minimum cost lifted multicut problem, with an application to motion segmentation. For this new problem class, we have proposed a heuristic solver that allows to generate solutions on dense point trajectory graphs. Further, for the motion segmentation application, we have proposed an algorithm that allows to insert higher-order edges adaptively w.r.t. pairwise motion differences. With this approach, we improve over the state-of-the-art in motion segmentation on the FBMS-59 benchmark. Since the proposed problem can model highly expressive cost functions, we hope that it is going to be useful for further computer vision applications.

References

- [1] S. Agarwal, K. Branson, and S. Belongie. Higher order learning with graphs. In *ICML*, 2006. 2
- [2] S. Agarwal, J. Lim, L. Zelnik-Manor, P. Perona, D. Kriegman, and S. Belongie. Beyond pairwise clustering. In *CVPR*, 2005. 2
- [3] B. Andres, T. Beier, and J. Kappes. OpenGM: A C++ library for discrete graphical models. *CoRR*, abs/1206.0111, 2012. 3
- [4] T. Beier, B. Andres, U. Köthe, and F. A. Hamprecht. An efficient fusion move algorithm for the minimum cost lifted multicut problem. In *ECCV*, 2016. 2
- [5] T. Brox and J. Malik. Object segmentation by long term analysis of point trajectories. In *ECCV*, 2010. 1, 2, 4, 6
- [6] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE TPAMI*, 33(3):500–513, 2011. 4
- [7] G. Chen and G. Lerman. Spectral curvature clustering scc. In *ICCV*, 2009. 2
- [8] S. E. Chew and N. D. Cahill. Semi-supervised normalized cuts for image segmentation. In *CVPR*, 2015. 1
- [9] S. Chopra and M. Rao. The partition problem. *Mathematical Programming*, 59(1–3):87–115, 1993. 3
- [10] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *ICCV*, 2013. 2
- [11] A. Eriksson, C. Olsson, and F. Kahl. Normalized cuts revisited: A reformulation for segmentation with linear grouping constraints. *JMIV*, 39(1), 2011. 1
- [12] A. Fix, A. Gruber, and B. Boros. A graph cut algorithm for higher-order markov random fields. In *ICCV*, 2011. 2
- [13] K. Fragkiadaki, P. Arbelaez, P. Felsen, and J. Malik. Learning to segment moving objects in videos. In *CVPR*, 2015. 1, 4
- [14] K. Fragkiadaki, G. Zhang, and J. Shi. Video segmentation by tracing discontinuities in a trajectory embedding. In *CVPR*, 2012. 4
- [15] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. Technical report, arXiv:1612.01925, Dec 2016. 4
- [16] P. Ji, H. Li, M. Salzmann, and Y. Dai. Robust motion segmentation with unknown correspondences. In *ECCV*. 2014. 2
- [17] J. H. Kappes, M. Speth, B. Andres, G. Reinelt, and C. Schnörr. Globally optimal image partitioning by multicuts. In *EMMCVPR*, 2011. 2
- [18] J. H. Kappes, M. Speth, G. Reinelt, and C. Schnörr. Higher-order segmentation via multicuts. *Computer Vision and Image Understanding*, 143c, 2016. 2
- [19] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell Systems Technical Journal*, 49:291–307, 1970. 3
- [20] M. Keuper, B. Andres, and T. Brox. Motion trajectory segmentation via minimum cost multicuts. In *ICCV*, 2015. 1, 2, 4, 5, 6, 7, 8
- [21] M. Keuper, E. Levinkov, N. Bonneel, G. Lavoue, T. Brox, and B. Andres. Efficient decomposition of image and mesh graphs by lifted multicuts. In *ICCV*, 2015. 1, 2, 3, 4
- [22] S. Kim, S. Nowozin, P. Kohli, and C. D. Y. Higher-order correlation clustering for image segmentation. In *NIPS*. 2011. 2, 3
- [23] S. Kim, C. D. Yoo, and S. Nowozin. Image segmentation using higher-order correlation clustering. *IEEE TPAMI*, 36(9):1761–1774, 2014. 2, 3
- [24] X. Lan, S. Roth, D. Huttenlocher, and M. J. Black. Efficient belief propagation with learned higher-order markov random fields. In *ECCV*, 2006. 2
- [25] J. Lezama, K. Alahari, J. Sivic, and I. Laptev. Track to the future: Spatio-temporal video segmentation with long-range motion cues. In *CVPR*, 2011. 1, 2
- [26] Z. Li, J. Guo, L. Cheong, and S. Zhou. Perspective motion segmentation via collaborative clustering. In *ICCV*, 2013. 1, 2
- [27] J. Liu, J. Jinglu Wang, T. Fang, C.-L. Tai, and L. Quan. Higher-order crf structural segmentation of 3d reconstructed surfaces. In *ICCV*, 2015. 2
- [28] S. Müller, P. Ochs, J. Weickert, and N. Graf. Robust interactive multi-label segmentation with an advanced edge detector. In *German Conference on Pattern Recognition (GCPR)*, volume 9796 of *LNCS*, pages 117–128. Springer, 2016. 2
- [29] P. Ochs and T. Brox. Object segmentation in video: a hierarchical variational approach for turning point trajectories into dense regions. In *ICCV*, 2011. 2, 8
- [30] P. Ochs and T. Brox. Higher order motion models and spectral clustering. In *CVPR*, 2012. 1, 2, 5, 6, 7
- [31] P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *IEEE TPAMI*, 36(6):1187 – 1200, Jun 2014. 2, 4, 5, 6, 7
- [32] H. Rahmati, R. Dragon, O. M. Aamo, L. V. Gool, and L. Adde. Motion segmentation with weak labeling priors. In *GCPR*, 2014. 2
- [33] K. Schelten and S. Roth. Mean field for continuous high-order mrfs. In *DAGM*, 2012. 2
- [34] F. Shi, Z. Zhou, J. Xiao, and W. Wu. Robust trajectory clustering for motion segmentation. In *ICCV*, 2013. 1, 2
- [35] S. Tang, B. Andres, M. Andriluka, and B. Schiele. Multi-person tracking by multicut and deep matching. *CoRR*, abs/1608.05404, 2016. 1
- [36] S. Tang, B. Andres, M. Andriluka, and B. Schiele. Multi-person tracking by multicut and deep matching, 2016. 2
- [37] T. Wu, A. R. Benson, and D. F. Gleich. General tensor spectral co-clustering for higher-order data. In *NIPS*, 2016. 2
- [38] D. Zhou, J. Huang, and B. Schölkopf. Learning with hypergraphs: clustering, classification, and embedding. In *NIPS*, 2007. 2
- [39] V. Zografos, R. Lenz, E. Ringaby, M. Felsberg, and K. Nordberg. Fast segmentation of sparse 3d point trajectories using group theoretical invariants. In *ACCV*, 2014. 2