

# UNSUPERVISED SEGMENTATION OF DEFORMING 3D MESHES VIA DEFORMATION-AWARE GRAPH CUTS

Yu Su

*Beijing City International School, Beijing 100000, China.*

*Corresponding Email: [susu071025@icloud.com](mailto:susu071025@icloud.com)*

**Abstract:** We propose an unsupervised method for segmenting deforming 3D mesh sequences using a deformation-aware graph cut. Our approach constructs a spatiotemporal graph and formulates segmentation as a minimum s-t cut problem. Unary costs, derived from per-vertex deformation energy, separate near-rigid parts from deforming regions, while pairwise costs enforce spatial smoothness. By iteratively applying max-flow/min-cut, the algorithm greedily extracts coherent parts without supervision. Results show the automatic partitioning of complex animations into meaningful, temporally-consistent components, validating our approach.

**Keywords:** Graph cut; 3D mesh segmentation; Unsupervised learning

## 1 INTRODUCTION

Mesh segmentation is a fundamental problem in computer graphics[1-2]. While traditional methods partition static geometry based on cues like curvature and concavity[3], emerging needs in animation and 4D modeling require the segmentation of deforming mesh sequences. This task aims to identify near-rigid components that move coherently through a non-rigid animation, thereby partitioning the object into its functional or anatomical substructures.

Segmenting deforming meshes introduces the key challenge of maintaining temporal consistency, rendering purely geometric criteria insufficient. Segmentation must therefore account for temporal deformation patterns. Prior works have addressed this by detecting regions of large deformation or clustering motion trajectories[2]. Our work builds on this concept by leveraging per-vertex deformation magnitude as a primary cue to distinguish between highly dynamic and relatively static regions of a mesh.

This paper introduces an unsupervised, deformation-aware graph cut method for segmenting deforming mesh sequences. We formulate the problem as a binary vertex labeling task, solved via max-flow/min-cut optimization on a specially constructed graph. The graph's unary terms encode temporal deformation energy, while its pairwise terms encode the mesh's spatial adjacency. This formulation allows a minimum cut, corresponding to a globally optimal solution for this energy, to cleanly separate a single component from the deforming mesh.

Our approach uses a greedy strategy: each iteration applies a binary graph cut to isolate and remove the most rigid component, repeating on the remaining mesh. Relying solely on vertex deformation, it partitions shapes into meaningful segments without manual annotation or training data. Validation on a complex animation shows consistent identification of natural anatomical divisions across ~135 frames, indicating the promise of deformation-based graph cut segmentation for unsupervised dynamic shape analysis.

## 2 RELATED WORK

**Static 3D Mesh Segmentation:** Segmenting static 3D shapes often relies on surface geometry and concavity; for example, Katz and Tal used fuzzy clustering and minimum cuts at deep concavities for hierarchical mesh decomposition[4]. Many methods use shape-aware metrics for segmentation—e.g., spectral clustering on mesh graphs segments a mesh based on eigenfunctions of an affinity matrix[5], and shape diameter function (SDF) based segmentation uses a volumetric shape function to consistently partition the mesh[6]. Other methods cluster faces or vertices by geometric similarity or primitive fitting[7], aiming to segment objects into meaningful parts, often evaluated by alignment with human perception.

Another line of work has explored segmentation with graph-based algorithms. Golovinskiy and Funkhouser proposed a randomized hierarchical segmentation using multiple random cuts on a mesh graph[4]. Their method uses repeated randomized min-cuts to find segmentation boundaries, merging them into a final partition and demonstrating the effectiveness of graph cuts. Similarly, some methods use normalized cuts or other graph partitioning techniques[8]. Markov Random Field (MRF) models further refine segmentation by treating it as a mesh labeling problem with pairwise smoothness constraints.

**Spatiotemporal and Dynamic Segmentation:** Few studies address segmentation of deforming or animated meshes. Notably, Lee et al. segment mesh animations into near-rigid components by analyzing inter-frame face deformations[2]. They segment at high-deformation regions to approximate rigid parts, separating joints or bends. Later methods cluster vertices or faces by motion using space-time graphs. Similarly, our approach builds a spatiotemporal graph with deformation features.

**Flow-Based Segmentation Methods:** Graph cuts (e.g., Boykov-Jolly, Boykov-Kolmogorov) are widely used for global image segmentation. We extend this to deforming 3D meshes by formulating segmentation as energy minimization with

deformation and Potts-model costs, allowing direct use of max-flow/min-cut solvers. Unlike prior 3D methods, our approach automatically extracts deformation features from the sequence.

**Supervised and Data-Driven Segmentation:** While our method is unsupervised, learning-based methods like Kalogerakis et al. use Conditional Random Fields and labeled data for semantic segmentation, but require extensive training and are limited to known classes. In contrast, our approach requires no training or labels, discovering segments purely from motion, and is suitable when labeled data is unavailable.

### 3 METHODOLOGY

Our method segments a sequence of 3D meshes with full vertex correspondence, as commonly found in animations or simulations with fixed topology. This allows us to measure each vertex's motion over time and use it for segmentation.

### 4 GRAPH CONSTRUCTION

We construct a spatiotemporal graph  $G = (V, E)$  where each node  $v \in V$  corresponds to a vertex of the mesh (and by extension, the trajectory of that vertex across all frames). We add undirected edges between nodes to encode spatial adjacency and enforce smoothness of the segmentation:

- **Spatial Edges:** For any two vertices that are connected by an edge in the mesh's triangular connectivity, we add an edge in  $G$ . These edges connect immediate neighbors on the mesh surface. Let  $E_{\text{spatial}}$  denote this set of edges. They form a graph isomorphic to the mesh's vertex graph for each frame (since connectivity is constant over time).
- **Temporal Connections:** We *do not* add explicit temporal edges between the same vertex across frames; instead, each vertex is a single node capturing its temporal behavior via the unary term. Essentially, each node already represents the vertex across all frames. (If the correspondence were not given, one could create per-frame vertex copies and link them temporally[2], but that is not needed here.)

Next, we define an s-t graph structure for performing a binary cut. In addition to the nodes representing mesh vertices, we introduce two special terminal nodes: the Source (S) and Sink (T). Each mesh vertex can connect to S or T with a weighted edge, encoding a preference for that vertex to be in one of two labels (which we can think of as Label 1 = "part of the segment" vs Label 0 = "not part of the segment" in the current cut).

### 5 DEFORMATION ENERGY AND UNARY COSTS

The key driver for segmentation is the amount each vertex *deforms* over time. We compute a deformation energy for each vertex  $i$  as the total variance of its 3D position over the sequence. Denote the position of vertex  $i$  at frame  $t$  as  $(x_{i,t}, y_{i,t}, z_{i,t})$ . We calculate:

$$d_i = \mathrm{Var}(x) + \mathrm{Var}(y) + \mathrm{Var}(z),$$

i.e., sum of coordinate variances over all frames. This scalar  $d_i \geq 0$  measures how much vertex  $i$  moves;  $d_i = 0$  would mean the vertex is static, while larger values indicate more motion. We then determine a rigidity threshold  $\tau$  – chosen as the median of all  $d_i$  values across the mesh. The intuition is that approximately half the vertices move less than  $\tau$  (more rigid) and half move more (more deformable). Vertices with  $d_i < \tau$  are likely part of relatively rigid components (e.g., core body parts), whereas vertices with  $d_i \geq \tau$  often belong to highly deforming parts (e.g., limbs, extremities).

Using this deformation measure, we assign unary costs (terminal edge weights) for each vertex  $i$ :

- Connect vertex  $i$  to the Source with weight  $\lambda$  if  $d_i < \tau$ , else weight 0.
- Connect vertex  $i$  to the Sink with weight  $\lambda$  if  $d_i \geq \tau$ , else weight 0.

Here  $\lambda$  is a tunable constant that represents the strength of the unary bias. In practice we set  $\lambda$  to a moderate value (e.g. comparable to or slightly larger than typical pairwise costs) to enforce the deformation cue while still allowing the cut to deviate if needed. The above scheme means: low-deformation vertices have a cost  $\lambda$  if they are cut to the Sink side (i.e., not included in the part), whereas high-deformation vertices have a cost  $\lambda$  if cut to the Source side. In other words, the energy will be lower if low-motion vertices stay with the Source (label 1) and high-motion vertices stay with the Sink (label 0). This encourages the cut to separate the graph such that mostly low-motion vertices are on the Source side of the cut (forming the segment), and high-motion vertices are on the Sink side (excluded from that segment). The threshold  $\tau$  thus biases the cut to prefer dividing the mesh along deformation lines.

### 6 PAIRWISE SMOOTHNESS TERM

To ensure that the resulting segment is spatially coherent, we add a pairwise smoothness cost on each edge between neighboring vertices. For each edge  $(i,j) \in E_{\text{spatial}}$  connecting vertices  $i$  and  $j$  on the mesh, we assign a penalty if  $i$  and  $j$  end up in different labels. Specifically, the pairwise cost is:

$$E_{\text{pair}}(i,j) = w \cdot \mathbf{1}(L_i \neq L_j),$$

where  $L_i \in \{0,1\}$  is the binary label (Sink or Source) of vertex  $i$ , and  $\mathbf{1}(\cdot)$  is an indicator function that is 1 if the argument is true (labels differ) and 0 if labels are the same. We choose a constant weight  $w$  (e.g.  $w=150$  in our experiments) for all adjacency edges. This is essentially a Potts model smoothness prior: it penalizes cut edges and thus encourages the segment to be a contiguous region on the mesh. A high  $w$  means the algorithm will prefer

to cut fewer edges (creating larger, smoother segments) unless strongly compelled by the unary terms. If  $w$  is too low, the segmentation could become noisy or fragmented. In our setting, we found that a relatively large  $w$  (on the order of the median  $d_i$  values) works well, since we expect only a few major parts rather than a very intricate segmentation boundary.

Putting it together, the energy function for a binary labeling  $L = \{L_i\}$  can be written as:

$$E(L) = \sum_{i \in V} U_i(L_i) + \sum_{(i,j) \in E_{\text{spatial}}} w \cdot \mathbf{1}(L_i \neq L_j),$$

where  $U_i(L_i)$  is the unary cost for assigning label  $L_i$  to vertex  $i$ , determined by  $d_i$  and  $\tau$  as described above. This energy is submodular (it has the form of a graph cut with non-negative weights), so it can be minimized exactly by solving a minimum s-t cut (max-flow) problem. We use the Boykov-Kolmogorov algorithm via the PyMaxflow library to compute the min-cut efficiently.

## 7 ITERATIVE SEGMENTATION PROCEDURE

The graph and energy define a binary cut (Source: low-deformation core; Sink: remainder), but our goal is multi-part segmentation without pre-setting the number of segments. We use an iterative, greedy strategy: extract one segment at a time, remove it, and repeat binary segmentation on the rest, yielding a sequence of prominent parts.

Our algorithm is as follows (pseudocode form):

Input: mesh\_sequence (with  $N$  vertices per frame), max\_iterations, min\_part\_size

Output: segmentation\_labels for all vertices (each vertex gets an integer part ID)

1. Compute per-vertex deformation energy  $d_i$  for all vertices (as above).
2. Initialize all vertex labels as 0 (meaning “unsegmented”).
3. Set current\_label = 1.
4. For iter = 1 to max\_iterations:
  - a. Let  $U$  = set of vertices with label 0 (still unsegmented).
  - b. If  $|U| < \text{min\_part\_size}$ , break (stop if remaining region is too small to segment further).
  - c. Construct graph  $G$  on  $U$  with edges and weights as defined (using the same  $\tau$  and  $\lambda$ ,  $w$ ).
  - d. Run max-flow / min-cut on  $G$  to obtain an optimal binary split of  $U$ .
  - e. Let  $S \subset U$  be the subset of vertices that are on the Source side of the cut (i.e., selected as the new segment).
  - f. If  $|S| < \text{min\_part\_size}$ , then this cut is too small; mark those vertices as segmented (or continue) without assigning a new label and continue to next iteration.
  - g. Assign label = current\_label to all vertices in  $S$ .
  - h. Remove  $S$  from  $U$  (they are now segmented); increment current\_label by 1.
5. End For.
6. Return the segmentation labels.

This iterative procedure uses graph cuts to greedily extract the most coherent low-deformation region each time, starting with the largest rigid part and later segmenting more deformable regions. It stops when the maximum iterations are reached or too few vertices remain, avoiding spurious small segments.

We keep the deformation threshold  $\tau$  fixed as the global median of all  $d_i$ , ensuring a consistent unary bias. Although  $\tau$  could be updated per iteration, the global median proved effective and provides a stable criterion for “high” deformation.

The complexity of each graph cut is  $O(V \cdot E)$  in practice with Boykov-Kolmogorov, which for our mesh ( $N$  vertices,  $\sim N^3$  edges for a triangulated mesh) is manageable. Since we typically extract only a handful of parts (on the order of 3-6 in our examples), the iterative approach is efficient.

## 8 RESULTS

We evaluated our method on a 135-frame 3D mesh sequence of a dancing cartoon mouse, comprising  $N=5,000$  vertices with fixed connectivity. The animation includes varied motions such as arm waving and walking, causing significant deformation across different body parts (Figure 1).



**Figure 1** Example Frames from the Input Deforming Mesh Sequence

As a prerequisite for segmentation, we compute the per-vertex deformation energy  $d_i$ . This metric effectively highlights extremities like the hands and feet as high-motion areas, while the torso and head remain relatively rigid. This energy map, shown in Figure 2, serves as the primary cue for our graph cut formulation.



**Figure 2** Visualization of Per-Vertex Deformation Energy (Blue: Low, Red: High)

We set the rigidity threshold  $\tau$  to the median of the deformation values and ran our iterative segmentation with a unary weight  $\lambda$  and a pairwise smoothness weight of  $w=150$ . A minimum part size of 50 vertices was enforced to prevent spurious segments.

The iterative process begins by isolating the most dynamic regions. In the first iteration, the graph cut separates the hands and feet, which exhibit the highest deformation due to waving and stepping. This initial result is depicted in Figure 3.



**Figure 3** The First Cut Isolates the Highly Deformable Hands and Feet

After these extremities are removed, the second iteration segments the next most coherent part. The algorithm identifies a natural boundary at the neck and shoulders, cleanly partitioning the relatively rigid head and upper torso from the lower body, as illustrated in Figure 4.



**Figure 4** The Second Cut Separates the Head and Upper Torso at the Neckline

In the final main iteration, the algorithm operates on the remaining lower body and tail, isolating the tail as a distinct component based on its motion pattern. This is shown in Figure 5.



**Figure 5** The final Cut Isolates the Tail from the Lower Body

After three iterations, the algorithm yields four distinct segments: (1) hands and feet, (2) head and upper torso, (3) tail, and (4) the lower torso and legs. The final partition demonstrates smooth boundaries that align with natural joints and successfully recovers semantically meaningful regions, thanks to the pairwise smoothness term. The complete segmentation is visualized in Figure 6.



**Figure 6** The Final Four-Part Segmentation of the Mesh, with Each Component Colored Uniquely

Qualitative results show our method produces meaningful segmentations: high-deformation regions are isolated early, head and torso are grouped then separated, and boundaries align with articulation points. The method is robust to minor mesh noise and motion inconsistencies, as the global deformation metric reduces the impact of short-term jitter. Without anatomical priors, it segments parts based on motion, often matching semantic regions, showing potential for automatic rigging or articulation analysis.

## 9 CONCLUSION

In this paper, we introduce an unsupervised method for segmenting deforming 3D mesh sequences based on a deformation-aware graph cut framework. Our approach formulates the task as an energy minimization problem on a spatiotemporal graph, where a unary term derived from per-vertex deformation magnitude distinguishes near-rigid from dynamic regions, and a pairwise term enforces spatial coherence. Through an iterative application of the max-flow/min-cut algorithm, our method greedily extracts motion-consistent components without requiring training data or manual annotations, successfully partitioning complex animations into meaningful, temporally-consistent segments that align with natural anatomical parts. This work validates the adaptation of graph-cut techniques for robust dynamic shape analysis. Future research directions include exploring global multi-label optimization to replace the iterative binary scheme and conducting a comprehensive quantitative evaluation against established benchmarks.

## COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

## REFERENCES

- [1] Katz S, Tal A. Hierarchical Mesh Decomposition Using Fuzzy Clustering and Cuts. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 2003, 22(3): 954–961.
- [2] Golovinskiy A, Funkhouser T. Randomized Cuts for 3D Mesh Analysis. *ACM Transactions on Graphics*, 2008, 27(5): 145.
- [3] Kalogerakis E, Hertzmann A, Singh K. Learning 3D Mesh Segmentation and Labeling. *ACM Transactions on Graphics*, 2010, 29(4): 102.
- [4] Lee TY, Wang YS, Chen TG. Segmenting a Deforming Mesh into Near-Rigid Components. *The Visual Computer*, 2006, 22(9-11): 729–739.
- [5] Shapira L, Shamir A, Cohen-Or D. Consistent Mesh Partitioning and Skeletonisation using the Shape Diameter Function. *The Visual Computer*, 2008, 24(4): 249–259.
- [6] Liu R, Zhang H. Segmentation of 3D Meshes through Spectral Clustering. *Proceedings of Pacific Graphics*, 2004: 298–305.
- [7] Boykov Y, Jolly MP. Interactive Graph Cuts for Optimal Boundary & Region Segmentation of Objects in N-D Images. *Proceedings of ICCV*, 2001: 105–112.
- [8] Boykov Y, Kolmogorov V. An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2004, 26(9): 1124–1137.