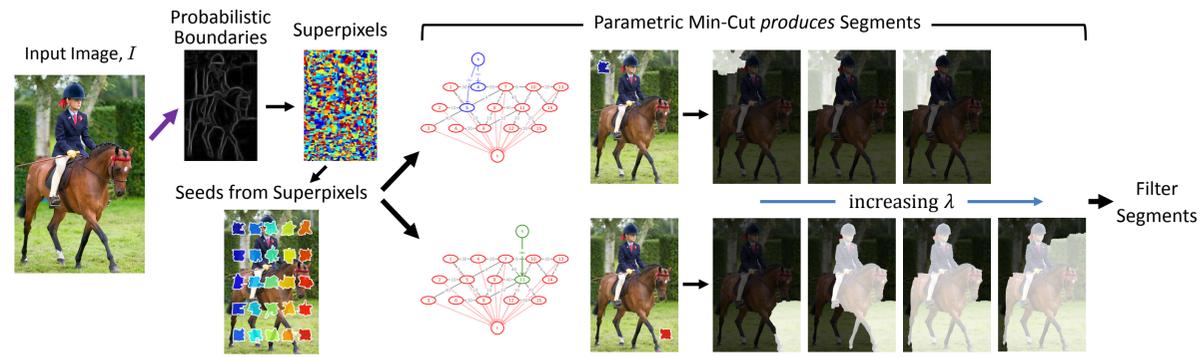


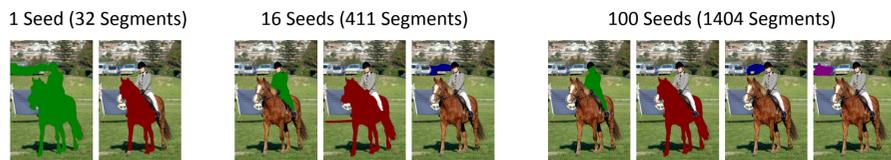
Goal: Generate unsupervised multiple *figure-ground segmentations*, in an order of magnitude faster than the state-of-art, without loss of accuracy.

Method Overview: Our algorithm has the following stages:



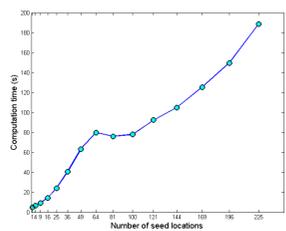
Fact 1

More seed locations gives higher recall



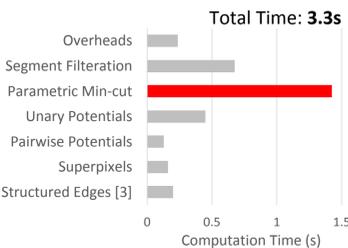
Fact 2

Obtaining segments from more seed locations is slow



Fact 3

PMC takes 45% of the total time after all other improvements (boundaries, superpixels, etc.)



Question: What information can be shared when minimizing N energy functions for parametric min-cut, if pairwise costs, V_{uv} remain same across the functions?

Seed: $D_{\lambda}^i(x_u) = \infty$ iff $x_u \in S_i$ and $x_u = 0$. **Condition:** $S_i \cap S_j = \emptyset$, for all i, j

$$S_i \rightarrow E_{\lambda}^i(\mathbf{X}) = \sum_{u \in \mathcal{V}} D_{\lambda}^i(x_u) + \sum_{(u,v) \in \mathcal{E}} V_{uv}(x_u, x_v)$$

$$S_j \rightarrow E_{\lambda}^j(\mathbf{X}) = \sum_{u \in \mathcal{V}} D_{\lambda}^j(x_u) + \sum_{(u,v) \in \mathcal{E}} V_{uv}(x_u, x_v)$$

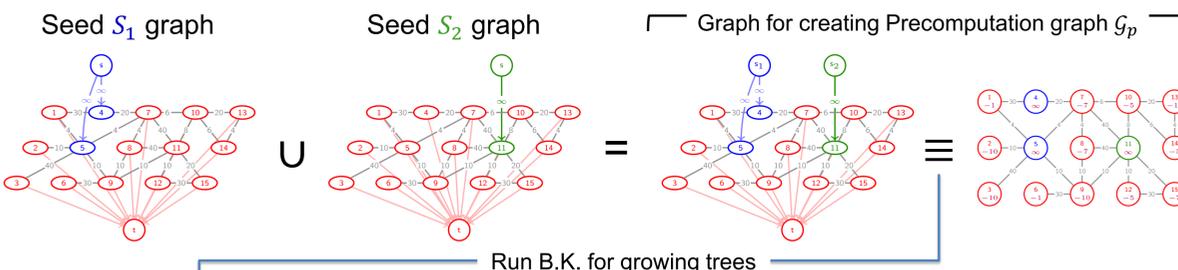
↑ Share information

Key Insight: Most of the edgelets between superpixels never get used in any parametric min-cut for any seed (about 43% edgelets).

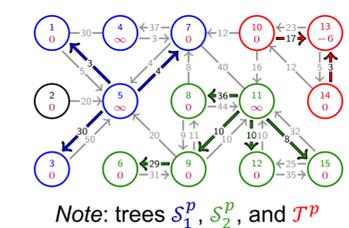


Novelty: Use a set of Boykov Kolmogorov [2] trees (one for each seed) to precompute information useful for all parametric min-cuts, in the case where pairwise costs do not change..

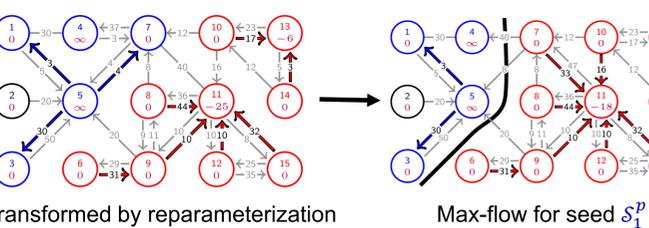
Step 1: Combine all seeds into one Precomputation Graph



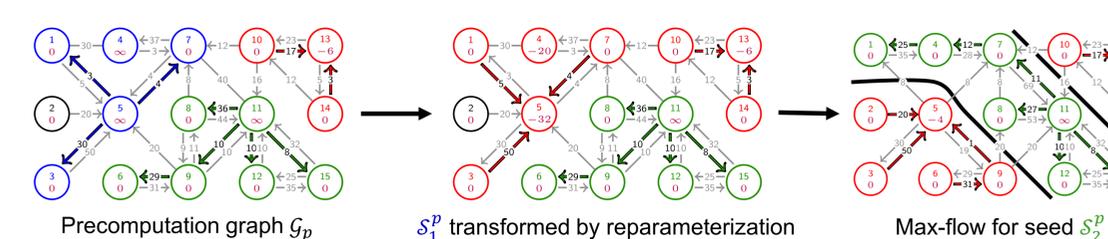
Step 2: G_p at end of Precomputation stage



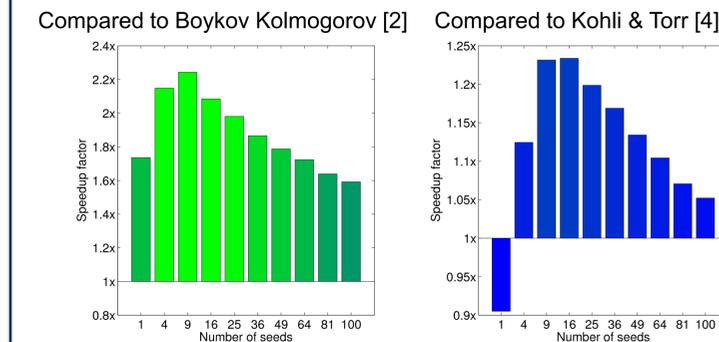
Step 3: Reparameterize and cut by Max-flow (this retains the optimal solution)



Repeat Step 3: Reparameterize and cut by Max-flow for all other seeds



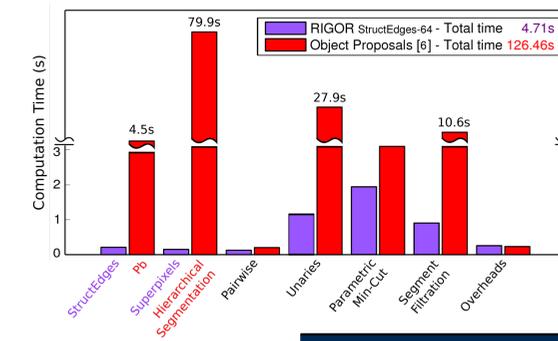
Results: Parametric min-cut timing comparison with increasing # of seeds.



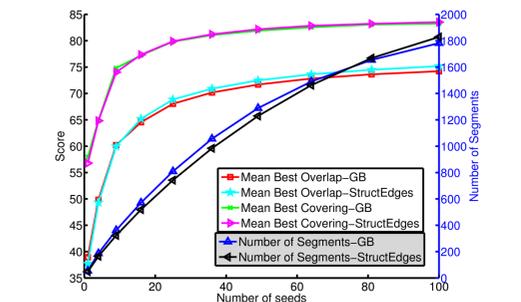
Parametric max-flow/min-cut timing comparison using different methods. StructEdges [3] was used in all tests.

	9 Seeds	25 Seeds	64 Seeds
RIGOR	282.1	688.5	1,846.2
[2]	632.8	1,363.8	3,181.3
[4]	347.5	825.4	2,038.8

Pipeline timing comparison to Object Proposals [6]



Algorithm performance with changing # of seeds



Performance:

Our algorithm performs slightly better and is an order of magnitude faster than CPMC [1]. It is ~25x times faster than Object Proposals [6] and ~100x faster than Shape Sharing [7].

Method	Mean Best Overlap	Mean Best Covering	Run Time (s)	# Segments
CPMC [1]	70.67	82.24	34.01	624.1
Object Proposals [6]	71.48	80.98	126.46	1544.1
Shape Sharing [7]	67.82	82.71	410.31	1115.4
Selective Search (fast) [8]	73.48	77.71	3.80	3574.0
RIGOR GB-25	68.04	79.83	4.62	808.7
RIGOR StructEdges-25	68.85	79.89	2.16	741.9
RIGOR GB-64	72.83	82.55	6.99	1490.3
RIGOR StructEdges-64	73.64	82.84	4.71	1462.8
RIGOR GB-100	74.22	83.25	9.26	1781.9
RIGOR StructEdges-100	75.19	83.52	6.84	1828.7

References:

[1] J. Carreira and C. Sminchisescu. *CPMC: Automatic Object Segmentation Using Constrained Parametric Min-Cuts*. PAMI, 2012.
 [2] Y. Boykov and V. Kolmogorov. *An experimental comparison of min-cut/maxflow algorithms for energy minimization in vision*. PAMI, 2004.
 [3] P. Dollar and C. Zitnick. *Structured forests for fast edge detection*. In ICCV, 2013.
 [4] P. Kohli and P. H. Torr. *Dynamic graph cuts for efficient inference in markov random fields*. PAMI, 2007.
 [5] D. S. Hochbaum. *The pseudoflow algorithm: A new algorithm for the maximum-flow problem*. Operations research, 2008.
 [6] I. Endres and D. Hoiem. *Category independent object proposals*. In ECCV, 2010.
 [7] J. Kim and K. Grauman. *Shape sharing for object segmentation*. In ECCV, 2012.
 [8] J. Uijlings, K. van de Sande, T. Gevers, and A. Smeulders. *Selective search for object recognition*. IJCV, 2013.

