

Resolution-independent superpixels based on convex constrained meshes without small angles

Jeremy Forsythe^{1,2}, Vitaliy Kurlin³, Andrew Fitzgibbon⁴

¹Vienna University of Technology, Favoritenstr. 9-11 / E186, A-1040 Vienna, Austria

²Department of Mathematical Sciences, Durham University, Durham DH1 3LE, UK

³Computer Science department, University of Liverpool, Liverpool L69 3BX, UK

⁴Microsoft Research, 21 Station Road, Cambridge CB1 2FB, UK

Abstract. The over-segmentation problem for images is studied in the new resolution-independent formulation when a large image is approximated by a small number of convex polygons with straight edges at subpixel precision. These polygonal superpixels are obtained by refining and extending subpixel edge segments to a full mesh of convex polygons without small angles and with approximation guarantees. Another novelty is the objective error difference between an original pixel-based image and the reconstructed image with a best constant color over each superpixel, which does not need human segmentations. The experiments on images from the Berkeley Segmentation Database show that new meshes are smaller and provide better approximations than the state-of-the-art.

1 Introduction: motivations, problem and contributions

1.1 Spatially Continuous Model for Over-segmentation of Images

Digital images are given by pixel values at discrete positions. Since images represent a spatially continuous world, the reconstruction problem should be solved in terms of functions defined over a *continuous image domain*, not over a discretization such as a regular grid. For example, grayscale values across a real image edge rarely drop from 255 (white) to 0 (black), but change gradually over 2-3 pixels, see details in [1, Fig. 1]. Hence a real edge between objects is often not along pixel boundaries and should be considered in the infinite family of line segments with any slope and endpoints having real coordinates. The first algorithm to output subpixel edges with theoretical guarantees is LSD [2].

The *over-segmentation problem* is to split an image into *superpixels* (larger than pixels and usually smaller than real objects) that have a nice shape and low variation of color. *Traditional superpixels* are formed by merging square-based pixels, e.g. by clustering. These superpixels often have irregular shapes with zigzag boundaries and holes inside. The *resolution-independent* approach [1] models a superpixel as a convex polygon with straight edges and vertices at subpixel resolution. Such a polygonal mesh can be rendered at any higher resolution by choosing a best color for each polygon in the *reconstructed image*.

38 A resulting mesh with constant colors over all polygons can be used to sub-
 39 stantially speed-up any higher level processing such as object detection or recog-
 40 nition. Fig. 1 shows that only 231 convex polygons are enough to approximate
 41 the original 512×512 image with a small reconstruction error from Definition 1.

42 1.2 Energy Minimization for Resolution-Independent Superpixels

43 A real image is modeled as a function I that is defined at any point of a con-
 44 tinuous image domain $\Omega \subset \mathbb{R}^2$ and takes values in \mathbb{R} (grayscale) or \mathbb{R}^3 (color
 45 images). We consider the function $I(\mathbf{x})$ taking the same color value at any point
 46 $\mathbf{x} \in \Omega$ within every square pixel B_p considered as a continuous subset of Ω . This
 47 function $I(\mathbf{x})$ defines a piecewise constant surface over the image domain Ω .




Fig. 1. Left: 512×512 input. **Middle:** 275 Voronoi superpixels have $\text{nRMS} \approx 10.2\%$.
Right: 246 superpixels based on a Convex Constrained Mesh have $\text{nRMS} \approx 4.48\%$.

48 **The reconstruction problem** is to find a latent image represented by a
 49 function $u(\mathbf{x})$ that minimizes the energy $E = \iint_{\Omega} \|I(\mathbf{x}) - u(\mathbf{x})\| d\mathbf{x} + R$, where R
 50 is a regularizer that penalizes degenerate solutions or reflects an image prior.

51 The energy E will be the reconstruction error from Definition 1. Usually $u(\mathbf{x})$
 52 is simpler than $I(\mathbf{x})$ in a certain sense. In our case $u(\mathbf{x})$ will have constant values
 53 over geometric polygons (superpixels) that are much larger than original pixels.
 54 The regularizer will forbid small angles, because narrow triangles may not cover
 55 even one pixel, while large angles (even equal to 180°) cause no difficulties.

56 So the reconstruction problem is to split a large image into a fixed number of
 57 polygons minimizing a difference between the original image function $I(\mathbf{x})$ over
 58 many pixels and the reconstructed image $u(\mathbf{x})$ over fewer convex polygons.

59 1.3 Contribution: Convex Constrained Mesh of Superpixels (CCM)

60 Here are the stages of the algorithm for resolution-independent superpixels.

61 1. The Line Segment Detector [2] finds line segments at subpixel resolution.

- 62 2. The LSD output is refined to resolve line intersections and small angles.
- 63 3. The resulting graph is extended to a triangulation without small angles.
- 64 4. Triangles are merged in convex polygons that also have no small angles.
- 65 5. The reconstructed image is obtained by finding the best constant color of any
- 66 convex superpixel after minimizing the approximation error in Definition 1.

67 The input of the LSD and CCM algorithms above is a grayscale image. The
 68 Convex Constrained Mesh (CCM) built at Stage 4 is introduced in Definition 2
 69 and has guarantees in Theorem 5 in terms of the following parameters.

- 70 • `Min_Angle` is the minimum angle between adjacent edges in a final mesh.
- 71 • `Min_Distance` is an approximation tolerance of LSD segments by CCM edges.

72 The default values are 3 pixels and 30° motivated by a similar angle bound
 73 in Shewchuk triangulations used at Stage 3. Here are the main contributions.

- 74 • The new concepts of the *reconstruction error* (a new quality measure for
- 75 resolution-independent superpixels not relying on ground truth segmentations)
- 76 and a *Convex Constrained Mesh* (CCM) are introduced in Definitions 1–2.
- 77 • The *LSD refinement* (Algorithm 3): disorganized line segments are converted
- 78 into a planar graph well approximating the original LSD with guarantees.
- 79 • *Shewchuk’s Triangle extension* (Algorithm 4): a triangulation is upgraded to a
- 80 Convex Constrained Mesh without small angles as guaranteed by Theorem 5.
- 81 • The *experiments on BSD* [3] in section 4 show that CCM have smaller sizes
- 82 and reconstruction errors than other resolution-independent superpixels, also
- 83 achieving similar benchmark results in comparison with traditional superpixels.

84 2 Pixel-based and Resolution-Independent Superpixels

85 A pixel-based image is represented by a lattice L whose nodes are in a 1–1
 86 correspondence with all pixels, while all edges of L represent adjacency relations
 87 between pixels. Usually each pixel is connected to its closest 4 or 8 neighbors.

88 The seminal *Normalized Cuts* algorithm by Shi and Malik [4] finds an optimal
 89 partition of L into connected components, which minimizes an energy taking into
 90 account all nodes of L . The algorithm by Felzenszwalb and Huttenlocher [5] was
 91 faster, but sometimes produced superpixels of irregular sizes and shapes as found
 92 by Levinstein et al. [6]. The *Lattice Cut* algorithm by Moore et al. [7] guarantees
 93 that the final mesh of superpixels is regular like the original grid of pixels. The
 94 best quality in this category is achieved by the Entropy Rate Superpixels (ERS)
 95 of Lie et al. [8] minimizing the entropy rate of a random walk on a graph.

96 The *Simple Linear Iterative Clustering* (SLIC) algorithm by Achanta et al. [9]
 97 forms superpixels by k -means clustering in a 5-dimensional space using 3 colors
 98 and 2 coordinates per pixel. Because the search is restricted to a neighborhood
 99 of a given size, the complexity is $O(kmn)$, where n and m are the numbers of
 100 pixels and iterations. This gives an average time of 0.2s per BSD500 image.

SEEDS (Superpixels Extracted via Energy-Driven Sampling) by Van den Bergh et al. [10] seems the first superpixel algorithm to use a *coarse-to-fine optimization*. The colors of all pixels within each fixed superpixel are put in bins, usually 5 bins for each color channel. Each superpixel has the associated sum of deviations of all bins from an average bin within the superpixel. This sum is maximal for a superpixel whose pixels have colors in one bin. SEEDS iteratively maximizes the sum of deviations by shrinking or expanding superpixels.

Almost all past superpixels have *no geometric or topological constraints*, only in a soft form of a regularizer [11]. If a final cluster of pixels in SLIC is disconnected or contains holes, post-processing is needed. TopoCut [12] by Chen et al. has a hard topological constraint in a related problem of image segmentation.

The *key limitation of pixel-based superpixels* is the fixed resolution of an original pixel grid. Resolution-independent superpixels are the next step in approximating images by polygons whose vertices have any subpixel precision.

The *only past resolution-independent superpixels* by Duan and Lafarge [13] and new CCM superpixels use constrained edges from the LSD algorithm of Grompone von Gioi et al. [2], which outputs thin rectangles such that the color substantially changes at their long middle lines, see Fig. 3. The parameters are a tolerance τ for angles between gradients and a threshold ε for false alarms.

Voronoi superpixels [13] are obtained by splitting an image into Voronoi faces whose centers are chosen along LSD edges. The natural input would be a set of centers, however the algorithm first runs LSD [2] and then chooses centers on both sides of LSD edges. So the edges were soft constraints without proved guarantees yet. By Theorem 5 all given edges are a hard constraint for CCMs.

A *Shewchuk triangulation* is produced by the state-of-the-art Triangle software [14] that guarantees a lower bound (as large as 28°) for all angles. A Convex Constrained Mesh introduced in Definition 2 extends a Shewchuk triangulation to a mesh of convex polygons that also have no small angles by construction.

3 A Convex Constrained Mesh (CCM) with Guarantees

A superpixel in Definition 1 can be a union of square pixels or any polygon.

Definition 1 Let an image I have n pixels, each pixel be the 1×1 square B_p and have $\text{Intensity}(p) \in [0, 255]$. Let I be split in superpixels F_j (polygons or unions of pixels) with $\text{Color}(F_j) \in [0, 255]$, $j = 1, \dots, s$. The Reconstruction Error is

$$RE = \min \sum_{\text{pixels } p} \left(\text{Intensity}(p) - \sum_{j=1}^s \text{Area}(B_p \cap F_j) \text{Color}(F_j) \right)^2, \quad (1a)$$

where the minimum is over all $\text{Color}(F_j)$, $j = 1, \dots, s$. The internal sum in RE is small, because each square B_p non-trivially intersects only few superpixels F_j , so the intersection $\text{Area}(B_p \cap F_j)$ is almost always 0 (when B_p is outside F_j)

or 1 (when F_j covers B_p). For a fixed splitting $I = \cup_{j=1}^s F_j$, the function RE quadratically depends on $\text{Color}(F_j)$, which are found from a linear system.

$$\text{The normalized Root Mean Square is } nRMS = \sqrt{\frac{RE}{n}} \cdot \frac{100\%}{255}. \quad (1b)$$

131 The reconstructed image is the superpixel mesh with all optimal $\text{Color}(F_j)$ min-
132 imizing $nRMS$. This colored mesh can be rendered at any resolution, see Fig. 2.

133 In Definition 1 if a superpixel F_j is a union of square pixels, then $\text{Area}(B_p \cap F_j)$
134 is always 0 or 1, so the optimal $\text{Color}(F_j)$ is the mean color of all pixels in F_j .



Fig. 2. Left: 589 Voronoi superpixels (mesh and reconstruction) have $nRMS \approx 9.22\%$.
Right: 416 CCM superpixels (red mesh and reconstruction) have $nRMS \approx 6.32\%$

135 Another important motivation for the new CCM superpixels is in Fig. 2,
136 where the reconstructed image from Definition 1 in the second picture is consid-
137 ered as the input for any higher level processing. Since boundaries of a Voronoi
138 mesh may not well approximate constrained edges, the reconstructed image may
139 miss long thin structures, such as legs of a camera tripod in Fig. 2.

140 **Definition 2** Let G be a planar straight line graph with angles at least $\varphi \leq 60^\circ$.
141 A Convex Constrained Mesh $\text{CCM}(G)$ is a piecewise linear complex such that

142 (2a) $\text{CCM}(G)$ has convex polygons with angles $\geq \text{Min_Angle} = \arcsin\left(\frac{1}{\sqrt{2}} \sin \frac{\varphi}{2}\right)$;
143

144 (2b) the graph G is covered by the edges of the Convex Constrained Mesh $\text{CCM}(G)$.

145 Any Shewchuk triangulation is an example of a Convex Constrained Mesh.
146 However, Definition 2 allows general meshes of any convex polygons without
147 small angles. We build CCM by converting the LSD output in Algorithm 3 into
148 a planar graph G without self-intersections and then by extending G into a
149 polygonal mesh without small angles. All steps below are needed to satisfy main
150 Theorem 5. Subsection 4.1 confirms that CCMs are smaller than past meshes.

151 **Algorithm 3** We convert disorganised line segments with self-intersections from
152 the LSD output into a straight line graph as follows, see details in [15].

153 (3.1) When a segment almost meets another segment (within the offset parameter
154 $\text{Min_Distance} = 3$ pixels), we extend the first one to a proper intersection .

155 (3.2) When two segments almost meet (endpoints within Min_Distance), we ex-
156 tend both to the intersection to avoid small angles/triangles in Algorithm 4.

157 (3.3) When segments meet, we insert their intersection as a vertex in the graph.

158 **Algorithm 4** We extend a graph G from Algorithm 3, see details in [15].

159 (4.1) The Triangle [14] extends the constrained edges of the graph G to a trian-
160 gulation that has more edges, no angles smaller than $\text{Min_Angle} = 30^\circ$.

161 (4.2) We merge adjacent faces along their common edge e if the resulting face is
162 still convex. If two new angles at the endpoints of e are almost convex, we try to
163 perturb them within Min_Distance to guarantee convexity and no small angles.

164 (4.3) We collapse unconstrained edges if all constrained edges remain fixed.

165 The steps above guarantee no small angles in CCM. Theorem 5 is proved in [15].

166 **Theorem 5** Let line segments S_1, \dots, S_k have m intersections. Algorithm 3
167 builds a CCM in time $O((k + m) \log(k + m))$ so that

168 (5a) any internal angle in a CCM face is not smaller than Min_Angle ;

169 (5b) the union $\cup_i S_i$ is covered by the Min_Distance -offset of the CCM's edges.

170 4 Experimental Comparisons and Conclusions

171 The sizes and reconstruction errors of the CCM and Voronoi superpixels are
172 compared in subsections 4.1 and 4.2. Then two more superpixel algorithms SLIC
173 [9] and SEEDS [10] are also included into BSD benchmarks in subsection 4.3.

174 4.1 Sizes of CCMs, Shewchuk's Triangulations and Voronoi meshes

175 The first picture in Fig. 3 is the original LSD output. The second picture shows
176 the graph G obtained by the LSD refinement in Algorithm 3. The refined LSD
177 output has more edges than the original LSD, because we include boundary
178 edges of images and also intersection points, which become vertices of graphs.

179 We use $\phi = 30^\circ$ for the LSD refinement, which leads to $\text{Min_Angle} \approx 10.5^\circ$ in
180 Shewchuk's Triangle [14]. We compare Shewchuk triangulations on the original
181 LSD output and CCM on the refined LSD output in Fig. 3, where the 3rd
182 picture shows a zoomed-in green box with many tiny triangles. The final picture
183 in Fig. 3 contains only few faces after merge operations in Algorithm 4. The
184 ratio of Shewchuk triangles to the number of faces in CCMs across BSD is 7.6.

185 The first step for Voronoi superpixels [13] is to post-process the LSD output
186 when close and near parallel lines are removed, because the target application
187 was satellite images of urban scenes with many straight edges of buildings. Then
188 long thin structures such as legs of a camera tripod in Fig. 3 are represented
189 only by one edge and may not be recognized in any further processing.

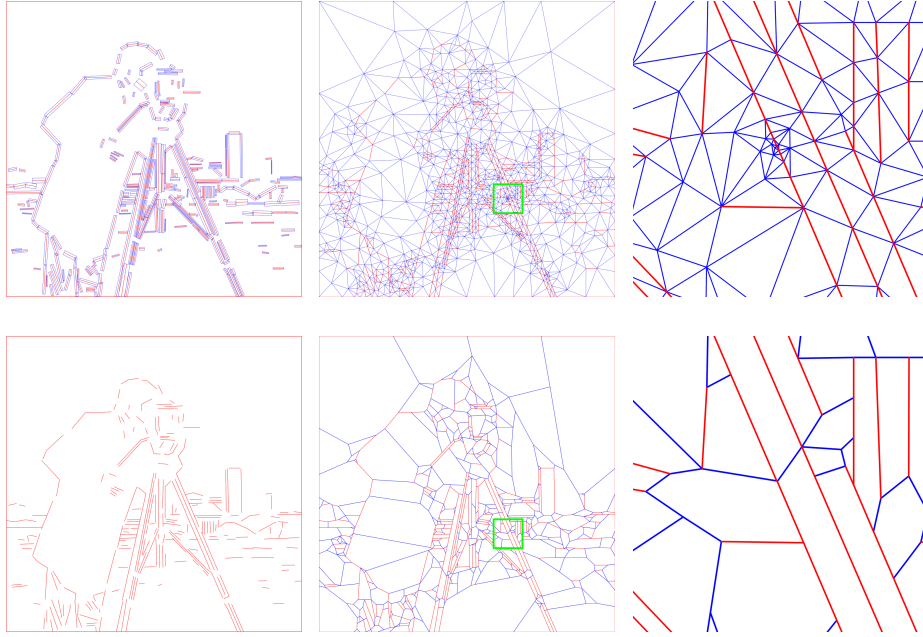


Fig. 3. **Top left:** 259 LSD red middle segments in blue rectangles before the refinement in Algorithm 3. **Bottom left:** the refined LSD output (a graph G) with 294 edges. **Top middle:** Shewchuk triangulation $T(G)$ with 2260 triangles. **Bottom middle:** the Convex Constrained Mesh $CCM(G)$ with 416 faces. **Top right:** zoomed in green box with tiny triangles. **Bottom right:** zoomed in green box, all tiny triangles are merged.

190 That is why the LSD refinement in section 3 follows another approach and
 191 offers guarantees leading to Theorem 5. Table 1 displays the average ratios of
 192 face numbers over BSD images. Even when the parameter Eps.Radius of Voronoi
 193 superpixels is increased to 12, these ratios converge to a factor of about 3.25.

194 4.2 Approximation Quality of the CCM and Past Superpixels

195 Since the aim of superpixels is to approximate a large image by a reconstructed
 196 image based on a smaller superpixel mesh, the important quality is the standard
 197 statistical error $nRMS$ over all pixels, which is introduced in Definition 1.

Table 1. Ratios of the face numbers for CCM and Voronoi meshes on the same LSD edges, averaged across BSD images [3]. The parameter Eps.Radius is in pixels.

Eps.Radius of a superpixel	4	5	6	7	8	9	10	11	12
Mean $\frac{\text{Voronoi superpixels [13]}}{\text{number of faces in CCM}}$	8.91	6.21	4.86	4.03	3.96	3.43	3.27	3.27	3.26

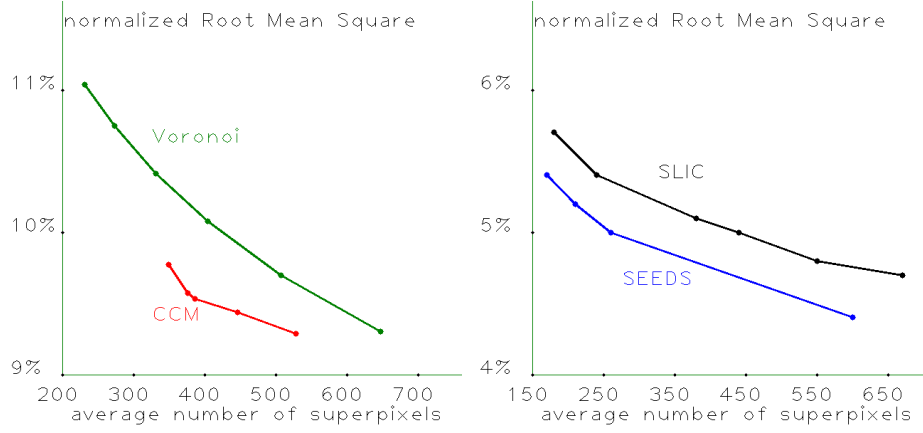


Fig. 4. The normalized Root Mean Squares in percents for Voronoi and CCM superpixels (on the left), SLIC and SEEDS (on the right) averaged over BSD500 images.




Fig. 5. Left: 791 Voronoi superpixels (mesh and reconstruction) with $nRMS \approx 8.45\%$. Right: 791 CCM superpixels (red mesh and reconstruction) with $nRMS \approx 7.22\%$.

Fig. 4 shows that the reconstructed images of CCM superpixels better approximate original images than Voronoi superpixels. Some convex polygons of CCMs are much larger than Voronoi superpixels, simply because the corresponding regions in images indeed have almost the same intensity, e.g. the sky. Hence taking the best constant color over each superpixel is reasonable.

Voronoi superpixels have similar sizes, because extra centers are added to empty regions using other non-LSD edges. Despite CCMs being obtained from only LSD edges without using colors, the reconstructions have smaller errors in comparison with Voronoi meshes containing more superpixels in Fig. 5.

Fig. 4 confirms smaller approximation errors of CCM superpixels across all BSD500 images, where we used the same LSD parameters for CCM and Voronoi superpixels. For all superpixels, we computed optimal colors minimizing the reconstruction error and measured $nRMS$ in percents, see Definition 1.

Each BSD experiment outputs 500 pairs (number of faces, $nRMS$). We average each coordinate of these pairs and output a single dot per experiment. The first red dot at (377.1, 9.626%) in Fig. 4 means that CCMs have 377 faces and an

approximation error of 9.6% on average. For a fixed image, the LSD algorithm outputs roughly the same number of edges for all reasonable parameters τ, ε .

So smaller CCMs seem impossible, because all LSD edges are hard constraints, while all faces should be convex. To get larger CCMs, we stop merging faces in Algorithm 4 after getting a certain number of convex faces. The five experiments on Voronoi superpixels with $\text{Eps_Radius} = 7, 8, 9, 10, 11$ produced 5 dots along a decreasing curve. Fig. 4 implies that Voronoi meshes require more superpixels (507.3 on average) to achieve the similar $nRMS = 9.696\%$.

4.3 Standard Benchmarks for CCM and Past Superpixels

The benchmarks BR and CUE are designed for pixel-based superpixels and use human segmentations from BSD [3], see details in [15]. We discretize CCM and Voronoi superpixels by drawing lines in OpenCV to detect boundary pixels. We put all pixels into one superpixel if their centers are in the same polygon.

It is unfair to compare discretized resolution-independent superpixels and pixel-based superpixels on benchmarks designed for the latter superpixels. CCM achieves smaller undersegmentation errors than SEEDS/SLIC and most importantly beats Voronoi superpixels on the objective $nRMS$ as well as on BR.

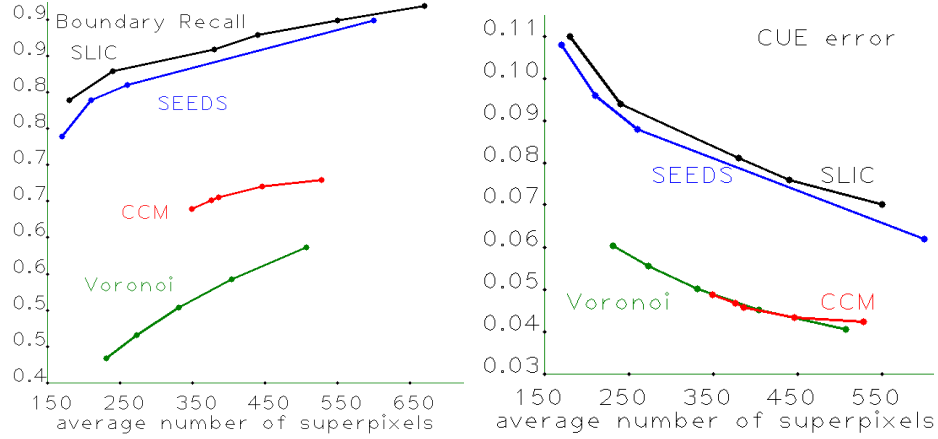


Fig. 6. Left: Boundary Recall (BR). **Right:** Corrected Undersegmentation Error.

Pixel-based superpixels SLIC and SEEDS achieve better results on $nRMS$ and Boundary Recall (BR) in Fig. 6, because their superpixels can have irregular boundaries (of only horizontal and vertical edges). However, humans are more likely to sketch straight edges than boundaries consisting of short zigzags.

So irregular pixel-based superpixels are often split by straight ground truth boundaries. Resolution-independent superpixels are convex polygons with straight edges and are expected to have smaller undersegmentation errors in Fig. 6.

Since only a Windows demo is available for Voronoi superpixels [13], we couldn't directly compare the running times of resolution-independent superpixels. We worked on a different platform and confirm that the running time for the CCM on a laptop with 8G RAM is about 0.15s across BSD500 images.

The key contribution is the new concept of a Convex Constrained Mesh (CCM), which extends any constrained line segments to a mesh of convex polygons without small angles. The paper focused on the quality of CCM superpixels, which seem ideal for detecting long thin structures in urban scenes, see Fig. 2.

- Theorem 5 guarantees the approximation quality and no small angles in CCMs, which also have smaller sizes on the same input in comparison with [14], [13].

- The CCM outperforms the only past algorithm [13] for resolution-independent superpixels on BR (Boundary Recall) and the new error $nRMS$ in Fig. 4, and even outperforms pixel-based superpixels on the CUE benchmark in Fig. 6.

The first author was supported by the project FWF P24600-N23 at TU Wien.

References

1. Viola, F., Fitzgibbon, A., Cipolla, R.: A unifying resolution-independent formulation for early vision. In: Proceedings of CVPR. (2012) 494–501
2. Grompone von Gioi, R., Jakubowicz, J., Morel, J.M., Randall, G.: Lsd: a line segment detector. Image Processing On Line **2** (2012) 35–55
3. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. Transactions PAMI **33** (2011) 898–916
4. Shi, J., Malik, J.: Normalized cuts and image segmentation. Transactions PAMI **22** (2000) 888–905
5. Felzenszwalb, P., Huttenlocher, D.: Efficient graph-based image segmentation. Int J Computer Vision **59** (2004) 167–181
6. Levinstein, A., Stere, A., Kutulakos, K., Fleet, D., Siddiqi, K.: Turbopixels: fast superpixels using geometric flows. Transactions PAMI **31** (2009) 2290–2297
7. Moore, A., Prince, S., Warrell, J.: Lattice cut – constructing superpixels using layer constraints. In: Proceedings of CVPR. (2010) 2117–2124
8. Liu, M.Y., Tuzel, O., Ramalingam, S., Chellappa, R.: Entropy rate superpixel segmentation. In: Proceedings of CVPR. (2011) 2097 – 2104
9. Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., Süsstrunk, S.: Slic superpixels compared to state-of-the-art superpixel methods. T-PAMI **34** (2012)
10. Van de Bergh, M., Boix, X., Roig, G., Van Gool, L.: Seeds: superpixels extracted via energy-driven sampling. Int J Computer Vision **111** (2015) 298–314
11. Veksler, O., Boykov, Y., Mehrani, P.: Superpixels and supervoxels in an energy optimization framework. In: Proceedings of ECCV. (2010) 211–224
12. Chen, C., Freedman, D., Lampert, C.: Enforcing topological constraints in random field image segmentation. In: Proceedings of CVPR. (2011) 2089–2096
13. Duan, L., Lafarge, F.: Image partitioning into convex polygons. In: Proceedings of CVPR (Computer Vision and Pattern Recognition). (2015) 3119–3127
14. Shewchuk, J.R.: Delaunay refinement algorithms for triangular mesh generation. Computational Geometry: Theory and Applications **22** (2002) 21–74
15. Forsythe, J., Kurlin, V., Fitzgibbon, A.: Resolution-independent superpixels based on convex constrained meshes (full version) (2016) <http://kurlin.org>.