

Dynamic Spectral Residual Superpixels

Jianchao Zhang^{1*}, Angelica I. Aviles-Rivero^{2*}, Daniel Heydecker^{2*},
Xiaosheng Zhuang¹, Raymond Chan¹ and Carola-Bibiane Schönlieb²

¹ City University of Hong Kong (CityU), Hong Kong.

² DPMMS and DAMTP, University of Cambridge, UK.

jzhang537-c@my.cityu.edu.hk, {ai323,dh489,cbs31}@cam.ac.uk, {xzhuang7,rchan.sci}@cityu.edu.hk




Figure 1. Input-Output examples of our proposed approach for three images with different characteristics such as diverse structure and colours. Our approach seeks to adhere better to the boundaries through a structure measurement whilst connecting meaningful regions, for example see the eyes and mouth at the three outputs.

Abstract

We consider the problem of segmenting an image into superpixels in the context of k -means clustering, in which we wish to decompose an image into local, homogeneous regions corresponding to the underlying objects. Our novel approach builds upon the widely used Simple Linear Iterative Clustering (SLIC), and incorporate a measure of objects' structure based on the spectral residual of an image. Based on this combination, we propose a modified initialisation scheme and search metric, which helps keeps fine-details. This combination leads to better adherence to object boundaries, while preventing unnecessary segmentation of large, uniform areas, while remaining computationally tractable in comparison to other methods. We demonstrate through numerical and visual experiments that our approach outperforms the state-of-the-art techniques.

*Equal Contribution.

1. Introduction

Image segmentation has been a widely explored task in computer vision yet a still open problem. In particular, superpixels segmentation has become a pre-processing tool for several applications including classification [23, 25], optical flow [18, 15], colour transfer [8], depth estimation [14, 4] and tracking [34, 37]. The central idea of superpixels is to split a given image in multiple clusters, which reflect semantically meaningful regions.

There are several advantages of using superpixel representation instead of working at pixel-wise level. Firstly, an application becomes computationally and representationally efficient as the number of primitives are significantly reduced. Secondly, the natural redundancy in an image is exploited, and therefore, features can be extracted on representative regions whilst reducing noise and increasing discriminative information [22, 1, 28].

Since the pioneering work of Ren and Malik [22], the community has devoted to develop different algorithmic approaches to improve over [22]. These approaches can be

roughly divided in: graph-based methods e.g. [22, 10], path-based approaches e.g. [29], density-based models e.g. [31], contour models e.g. [11] and clustering methods e.g. [1, 17].

Out of all of the approaches reported in the literature, the Simple Linear Iterative Clustering (SLIC) [1] is perhaps the most popular method that offers a good performance whilst demanding low computational cost, by building on Lloyd’s algorithm [16] for k -means. The central idea of SLIC is to perform the superpixels partition based on an iterative scheme that search for similarities between points ensuring at each step that we assign points to the nearest cluster from the previous step.

The ability of SLIC to obtain a good segmentation with low computational cost comes from the observation that, by using a similarity metric, one can greatly reduce the number of distance calculations required. However, SLIC is also limited by its own construction, and in particular, by its search range, and one can thus observe two major limitations. Firstly, SLIC tends to segment large uniform regions in an image with more superpixels than are intuitively necessary, which limits resolution in other parts of the image. Secondly, in structure-rich parts of the image, the final superpixel size is much smaller than the search radius of SLIC, which leads to many unnecessary distance computation. Finally, since we expect structure-rich parts of the image to have a higher density of superpixels, it may be efficient to perform the initial seeding of cluster centres in anticipation of this inhomogeneity.

In this work, we propose a new algorithmic approach, exhibited in Fig. 1, that improves upon the SLIC approach, motivated by the drawbacks discussed above. We show that our approach outperforms SLIC and several works on the body of literature. Our main contributions are as follows.

- We propose a new superpixel approach, which incorporates the *saliency function* $\mathcal{S}(x)$ of Hou et al. [9] as a proxy for object density. This leads to the following advantages.
 - By incorporating the saliency $\mathcal{S}(x)$ into the distance computation, we can prevent unnecessary over-segmentation of large, uniform regions, such as the sky in the first example of Fig. 2, and allowing greater focus on structure-rich parts of the image.
 - We propose a new seeding strategy, based on the inhomogeneity described by \mathcal{S} . This allows for greater resolution changes at fewer iterations by focusing on relevant structures, and hence keeping fine-details of the structures in the final segmentation.
- We extensively evaluate our approach with a large range of numerical and visual experiments.

- We demonstrate that our two major contributions mitigates the major drawbacks of the state-of-the-art techniques, by reporting the lowest undersegmentation error and highest boundary recall.

2. Related Work

In this section, we review the body of literature in turn. We then highlight the advantages of clustering based methods, and their current drawbacks that motivate our new algorithmic approach.

There have been different attempts in the literature to improve superpixels segmentation. A set of approaches tackle the problem using graph representation of the image and the partition is based on the similarity of the nodes, e.g. colour, including [22, 24, 26, 6, 19, 35]. However, although promising results are reported, the computational time is often very high. Another perspective has been followed by local mode-seeking algorithms including the well-known Quick Shift, which partition is based on an approximation of kernelised mean-shift [31]. However, there is not control on the number of superpixels or compactness.

Another set of approaches addressed the superpixel partition problem as the task of finding the shortest path between seeds, for instance using the the well-known Dijkstra algorithm, as reported in [29, 7]. However, this type of approach is usually unable to control the compactness. We briefly mention other methods for superpixels segmentation. A body of work has proposed algorithms for image segmentation based on *gradient ascent* and other geometric methods [5, 32, 11]. For an extensive review of the literature, we refer the reader to [28].

In particular, in this work we focus on, probably, the most popular superpixel category, which is clustering based approaches. The basis of this perspective builds on Lloyd’s algorithm [16] for k -means clustering. The main idea of this algorithm is to partition a set of observations into k clusters, in which each observation is assigned to the cluster with the nearest mean, and produces excellent results at the cost of high computational intensity. Within this category, one can find a top reference approach called Simple Linear Iterative Clustering (SLIC). SLIC was proposed by Achanta et al. [1], in which authors propose a local version of the Lloyd’s algorithm, which is computationally much simpler while keeping excellent segmentation quality.

Following this philosophy, different algorithmic approaches have been proposed including [33, 21, 20, 13]. Most recently, in [2] authors proposed an improved version of SLIC, in which they proposed to compute a polygonal partition to adapt better to the geometry of the objects in the image. Maierhofer et al [17] propose a *dynamic refinement* of this method, called dSLIC, which seeds the initial cluster points inhomogeneously and allows the search radius to

vary across clusters, both according to a measure of local object density. This allows better capturing of fine details in structure-rich regions, and further reduces computational complexity by eliminating unnecessary searches.

Let us also mention the closely related problem of *salient object detection*. In this problem, one has the simpler goal of identifying which regions in an image contain *salient* or novel information, and which contain only patterns and structures repeated throughout the image. This problem shares some similarities with the problem of image segmentation; for instance, one might hope that the salient objects are identified as superpixel regions. A hugely successful method in this problem, based on Fourier analysis, was proposed by Hou et al. [9], which inspires our current approach. More recent works include techniques based on graphs [35] or machine-learning [12, 30].

3. Proposed Approach

In this section, we describe in detail our superpixel approach. Firstly, we formalise the definition of superpixels in terms of a clustering task. We then introduce the details of both our new measure of structure function and our initialisation strategy.

We view an input image, of width A and height B , as a map $I : \mathcal{X} \rightarrow \Omega$, where $\mathcal{X} = [A] \times [B]$ is a rectangular domain, and $\Omega \subset \mathbb{R}^3$ appropriate colour domain. We also define a metric d on $\mathcal{X} \times \Omega$, representing the similarity of points in space and with different colour values, and a *feature map* \mathcal{F} , which takes a subset $S \subset \mathcal{X}$ and returns a pair in $\mathcal{X} \times \Omega$. *k-means clustering* now seeks a partition of \mathcal{X} into path-connected sets $\{S_i\}_{i=1}^n$ such that, for each i , S_i is exactly the set of points where the infimum $\inf_j d((x, I(x)), \mathcal{F}(S_j))$ is attained at $j = i$.

3.1. Object Density Measure via Spectral Residual

The key strength of dSLIC [17] over SLIC is the recognition that objects in an image are not distributed uniformly, and that image segmentation can exploit this to improve segmentation results and computational efficiency. Our approach is to exploit this same principle further, and use the strength of the *Spectral Residual* approach proposed by [9] as a better measure of object detection.

We briefly review the Fourier analysis leading to the definition of the spectral residual in [9]. For an image I , we write $\mathfrak{F}I$ for the Fourier transform, which is a matrix of the same dimensions as I , and whose arguments we will write as two-dimensional frequencies f . The log-spectrum of an image I is then given by

$$\mathcal{L}(f) = \log(\Re(\mathfrak{F}I)(f)) \quad (1)$$

where \Re denotes the real part; we also write $\mathcal{P}(f) = \Im(\mathfrak{F}I)(f)$ for the imaginary part, or phase spectrum. The

key insight of [9] is that much of the information contained within \mathcal{L} is redundant, because \mathcal{L} is, to a good approximation, locally linear. These features are then encapsulated in the local average $\mathcal{A}(f) = (h_n \star \mathcal{L})(f)$, where h_n is the matrix consisting entirely of $\frac{1}{|\mathcal{X}|}$, and the *residual* log-spectrum, corresponding to the salient features, is given by

$$\mathcal{R}(f) = \mathcal{L}(f) - \mathcal{A}(f) = \mathcal{L}(f) - (h_n \star \mathcal{L})(f). \quad (2)$$

The final saliency map, which we take as our measure of object density, is then given by recombining with the phase spectrum, inverting the Fourier transform and adjusting the resulting function. Therefore, our proposed function reads:

$$\mathcal{SR}(x) = g_\sigma \star \left(\mathfrak{F}^{-1} \left\{ e^{\mathcal{R}(f) + \mathcal{P}(f)} \right\} \right)^2(x), \quad (3)$$

where the squaring ensures that the quantity considered is nonnegative, and the convolution with a Gaussian kernel g_σ ensures that the final result is smooth. For our purposes, we found that in practice $\sigma = 20$ is an excellent value. We then set a rescaling step for the search radius using (3) as $\mathcal{G}(x) := \exp(\mathcal{SR}(x) - \overline{\mathcal{SR}})$, where $\overline{\mathcal{SR}}$ denotes the mean of the structure function on the image grid. We then propose to have the distance computations depending on our function, which reads:

```

1 #pseudocode for computing distances
2 input m (compactness), number of superpixels
3 while e_r ≤ outset
4   for x ∈ X, 1 ≤ i ≤ k:
5     if |x - (F(S_i^{(t)}))_1| ≤ 2SG(F(S_i^{(t)})):
6       #distance computation
7       Compute: d((x, I(x)), F(S_i^{(t)}))
8     else:
9       d((x, I(x)), F(S_i^{(t)})) = ∞
10    S_i^{t+1} = min_{1 ≤ j ≤ n} d((x, I(x)), F(S_j^{(t)}))
11 Compute Residual Error e_r
12 Increase t=t+1

```

By incorporating our proposed function, which we use as a measure of object density, into the distance computation one obtains two major advantages. Firstly, by doing a dynamic adjustment of the search range based on our function \mathcal{G} , one can connect uniform regions, and so avoid segmenting the images into unnecessary small superpixels. This effect is illustrated in Fig. 2, for example, at the second column where our approach was able to keep the sky in a same region, and the yellow car. Secondly, our approach focuses on segmenting fine details by capturing relevant structures; this is visible in the owl's eyes and head in Fig. 2. Furthermore, we remark that our approach does not introduce additional parameters to fine-tuning.

We now turn to explain our second major modification, which concerns the seeding initialisation.



Figure 2. Illustration output of our approach against SLIC. SLIC tends to over-segment uniform areas with more superpixels than necessary, such as the sky in the first image, and fails to preserve fine-structures, such as the owl’s eyes and the roller coaster at the zoom-in views. The numerical results in parenthesis denote (Undersegmentation Error, Boundary Recall). The lower the undersegmentation error the better whilst the higher the boundary recall the better.




Figure 3. Illustration of our initialisation strategy, which incorporates the object density measure \mathcal{G} , for two initial seeds. From left to right, first seed and second one.

3.2. Seeding Initialisation: A New Strategy

In this section, we describe our new seeding initialisation. Our main motivation is that we can use the object density measure \mathcal{G} defined above to help seed clusters in object-rich parts of the image, which we expect to contain more distinct regions. In this way, we obtain greater resolution at fewer iterations, and improve the focus on relevant and interesting regions.

We remark to the reader that SLIC initialisation is based on sampling pixels at the image grid. For comparison purposes, we start by defining the SLIC initialisation, which reads as follows.

```

1 #pseudocode for seeding initialisation SLIC
2 Set: Initialise cluster centers as
3  $C_i = [x_i, y_i, l_i, a_i, b_i]$  by sampling at regular grid
4 step:  $S = \sqrt{N/k}$ 
5 #where N is the size of the image
6 Move cluster centers to the lowest gradient
7 position in a  $3 \times 3$  neighborhood

```

Our proposed approach, which incorporates \mathcal{G} into this seeding, can be described informally as follows. We first set as an initial point the pixel with the lowest value in \mathcal{G} , and then we increase the values near to the initial point such that its neighbours are unlikely to be selected as another initial point. In this way, we guarantee that the distance between seed points is comparable to the search range, which will help reduce redundant searches. This process is illustrated in Fig. 3 for two initial seeds.

The hedging described above is carried out in two stages as follows.

- Points adjacent to the initial point are made unselectable, by setting the value of \mathcal{G} at these points excessively large.
- Points in the proximity of the initial point are made less likely, but not impossible, to select, again by altering \mathcal{G} . The influence range and to what extent are under consideration.

The advantage of these changes is that the density of area is limited twice compared with original method. The overall procedure of our method, which suitably sets the initialisation points according to our structure measure \mathcal{G} , is described formally as follows.

```

1 #pseudocode for seeding initialisation OURS
2 Set  $r = \max(\mathcal{G})/\min(\mathcal{G})$ 
3 While Enough Seeds:
4   Set  $\text{range} = \sqrt{\text{NumOfPixels}/\text{NumOfSuperpixels}}$ 

```

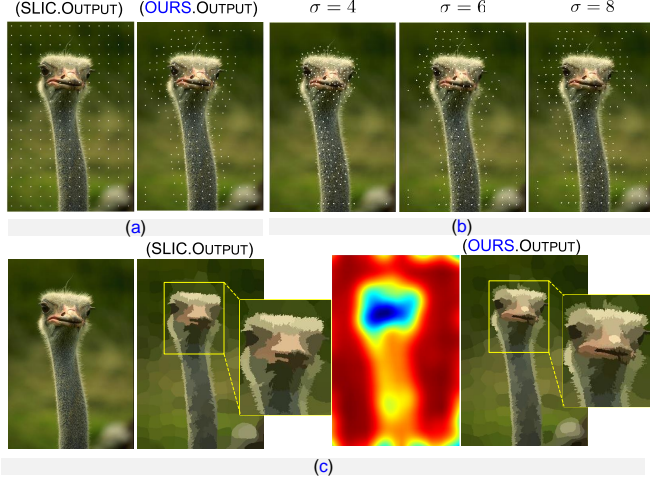



Figure 4. Visual comparison of the seeding initialisation of SLIC vs ours. (a) One can see that we seek to focus on relevant areas (i.e. other than background). (b) The effect of τ in our seeding strategy. (c) Visual illustration of the positive repercussion of our technique over SLIC. Our approach seeks to focus on structure-rich parts (see the face) whilst avoid over-segmenting uniform areas such as the background.

```

5  for each Superpixel center  $C_j$ :
6      Initialise  $S_i$  with the lowest value
7      in  $\mathcal{G}$ 
8      Set the adjacent neighbours of  $S_i$ 
9       $\mathcal{G} = \infty$ 
10     for all points  $S_j$  with  $d(S_i, S_j) < range$ 
11          $\mathcal{G}(S_j) = \mathcal{G}(S_j) * \sqrt{\tau}$ 
12     Smooth region with  $g_\tau$ ,  $\tau = 13/2$ 

```

An output example is displayed in Fig. 4. Subfigure (a) shows a initialisation comparison between our approach and SLIC, and we see that our approach gives more importance to the ostrich than the background. In subfigure (b), we evaluate possible choices for τ , and display outputs for $\tau = 4, 6, 8$. In practise, we found that the $\tau = 13/2$ works for a range of images. Finally, subfigure (c) shows a visual illustration of the advantage of our technique. By integrating our density function into the distance computation one can avoid unnecessary over-segmentation of large uniform regions (i.e. the background) and aim the attention at structure-rich parts (i.e. the ostrich). A closer look at the zoom-in regions one can see that our technique was able to preserve fine details in the face whilst SLIC lost relevant areas such as the left eye.

4. Experimental Results

In this section, we describe in detail the experiments that we ran to evaluate our approach.

4.1. Evaluation Protocol

Dataset Description. We evaluate our proposed approach on a publicly available dataset, the Berkeley Segmentation Dataset [3], which provides ground truth of the images for quantitative analysis. Moreover, to further support our technique, we included two interesting cases to study where fine-details are particularly difficult to retrieve: 1) out-of-the-focus images using the dataset of that [27] and 2) wide angle images using the dataset from [36]. All the measurements and comparisons in this section are taken from these datasets.

Comparison Methodology. We compare our approach to the SOTA methods on superpixels. For this, we design a two-part evaluation scheme. For the first part, we compared our approach against SLIC [1]; this comparison therefore demonstrate that our carefully design solution achieves better performance than the top reference in clustering-based methods. For the second part, we compared to state-of-the-art techniques: QS [31], TP [11], TPS [29], LRW [24], SNIC [2] and dSLIC [17]. We compare our approach qualitatively by visual comparisons and quantitatively by computing three metrics: Under-segmentation Error (UE), Boundary Recall (BR) and Boundary Precision (BP), which definitions are described next.

We assume that we are given an image, along with a *ground truth* $\Gamma = \{g_i\}_{i=1}^M$, representing the true regions of the image.

Boundary Recall measures the proportion of the boundary of the true regions in the ground truth which are close to a boundary in the segmentation. To quantify the notion of being close to a boundary, we recall the following definition.

Given a subset E of the edge set, we define the distance $d(e, E) = \inf\{|e - f| : f \in E\}$, where $|\cdot|$ denotes the l_2 norm of the difference, measured in pixels. We then define the *Boundary Recall* as

Definition 1 (Boundary Recall) Given a ground truth $\Gamma = \{g_i\}_{i=1}^M$ and a segmentation $\mathcal{S} = \{s_j\}_{j=1}^k$, we write $\partial\Gamma$ for the union of the edge boundaries ∂g_i , and similarly write $\partial\mathcal{S}$ for $\{s_j\}_{j=1}^k$. We define the boundary recall by

$$\partial(\Gamma, \mathcal{S}) = \frac{\#\{e \in \partial\Gamma : d(e, \partial\mathcal{S}) \leq 2\}}{\#\partial\Gamma}.$$

In words, the boundary recall is the proportion of true edges which are close to a superpixel edge.

Undersegmentation Error. Intuitively, this measures the size of all superpixels which *spill* across boundaries of the ground-truth.



Figure 5. From left to right. Quantitative comparison of our approach vs SLIC using three metrics: UE, BR and BP. Our approach reported the best scores metric-wise. Four visual outputs comparisons of SLIC vs our approach. In a closer look, one can see that our approach achieves better connection of structures and keeps fine-details. For example, see (A), (C) and (D) the faces and (B) the hand. The numerical results in parenthesis denote (Undersegmentation Error, Boundary Recall)

Definition 2 For a ground truth $\Gamma = \{g_i\}_{i=1}^M$, we fix thresholds $B_i, i = 1, \dots, M$. Given segmentation $\mathcal{S} = \{s_i\}_{i=1}^k$ of the image, the under-segmentation error is given by

$$U = \frac{1}{N} \left[\sum_{i=1}^M \left(\sum_{\#s_j \cap g_i \geq B_i} \#s_j \right) - N \right]$$

Observe that, since \mathcal{S} is a partition of the image, we can rewrite

$$U = \frac{1}{N} \sum_{i=1}^M \left[\left(\sum_{\#s_j \cap g_i \geq B_i} \#s_j \right) - \#g_i \right]$$

Hence, the undersegmentation error U measures how wasteful the coverings of the true regions g_i by the superpixel regions s_j are. We usually take B_i to be a fixed proportion of $\#g_i$

Parameter Selection. For all compared approaches QS [31], TP [11], TPS [29], SLIC [1], LRW [24], SNIC [2] and dSLIC [17], we set the parameters as suggested in the corresponding work. We also used the codes released from each corresponding author. For our approach, we set the $m = 10$ since it offers a good trade-off between shape uniformity and boundary adherence (see Supplementary Material Section 2 for further description on m). We performed the evaluation using up to a range of number of superpixels up to 600.

The experiments reported in this section were under the same conditions in a Matlab CPU-based implementation. We used an Intel Core i7 with 16GB RAM.

4.2. Results and Discussion

We divide this section in two parts, following the comparison methodology scheme presented in previous sections.

► **Is our Approach better than SLIC?** As SLIC approach remains a top reference, and is the basis of our approach, we start by evaluating our approach against it. Results are displayed in Fig. 5. In a closer look, at the right side, of this figure, one can see that our approach yields to a better segmentation of the structures, keeping fine details of the objects. Moreover, it avoids unnecessary oversegmentation on uniform areas. These positive properties of our approach can be seen, for example, in (B) the proper recovery of the hand; in (C) the hair, eyebrows and the lines patterns in the jumper that are correctly clustered; in (D) where our approach successfully capture the eyes and moustache, and in (A) with better preservation of the face structure including the nose and lips.

To further support of our visual results, we ran a quantitative analysis based on three metrics UE, BR and BP. The results are displayed at the left side of Fig. 5. The top part shows a comparison in terms of UE, where the results reflect conformity to the true boundaries. We can observe that our approach achieves the lowest UE for all superpixels counts. The same positive effect was found in terms of precision versus recall, in which our approach displayed the best performance. Overall, our proposed technique outperforms SLIC at all superpixels counts and reduces the error rate by 20%. This improvement is translated to our ap-




Figure 6. Superpixel outputs comparisons of our approach vs different methods from the body of literature: QS [31], SLIC [1] TP [11], TPS [29], LRW [24], SNIC [2] and dSLIC [17]. A closer inspection, one can see that our approach offers better superpixels segmentation. For example, (A), (B), and (C) the eyes; (D) the ostrich’s boundary and (E) the eyes and basket.

proach to be the best in terms of producing superpixels that respect the object boundaries.

▷ **Is our approach better than other Superpixel approaches?** As the second part of our evaluation, we compare our approach against SOTA models: QS [31], TP [11], TPS [29], LRW [24], SNIC [2] and dSLIC [17]. We selected for our comparison approaches coming from differ-

ent perspectives: graph-based, path-based, density-based and clustering-based approaches. Results are displayed in Figs. 6, 7 and 8.

We first present a visual comparison of a selection of images from the Berkeley dataset in Figs. 6 and 7. By visual inspection, one can see that QS and TPS are the ones that perform more poorly than the other compared approaches. They fail to provide good boundaries of the structures and




Figure 7. Superpixel segmentation outputs of our approach vs QS [31], SLIC [1] TP [11], TPS [29], LRW [24], SNIC [2] and dSLIC [17]. Visual assessment shows that the proposed algorithm performs better than the compared approaches. Examples are: (F) the leaves; (G) and (I) the face; (H) the house boundaries and (J) the hand.

they do not preserve relevant details. Examples can be seen in (A), (B), (E) and (G) the eyes; (C), (E) and (F) preservation of fine details. LRW offer a better edge adherence to the objects than QS and TPS but also fails to preserve relevant objects, for example (G) the eyes and moustache.

In contrast to those approaches, SLIC and dSLIC performs better than QS, TPS and LRW. One can observe that SLIC and dSLIC readily compete in terms of having bet-

ter boundary adherence to the structures and grouping correctly majority of the objects. However, in particular, SLIC still produces outputs with more superpixels than necessary in homogeneous parts of the structures; see, for instance, in (A) the fish eye, (B) the nose and (J) the hand. Although dSLIC performs slightly better than SLIC, it still fails to capture fine details.

Among those approaches, SNIC approach display more

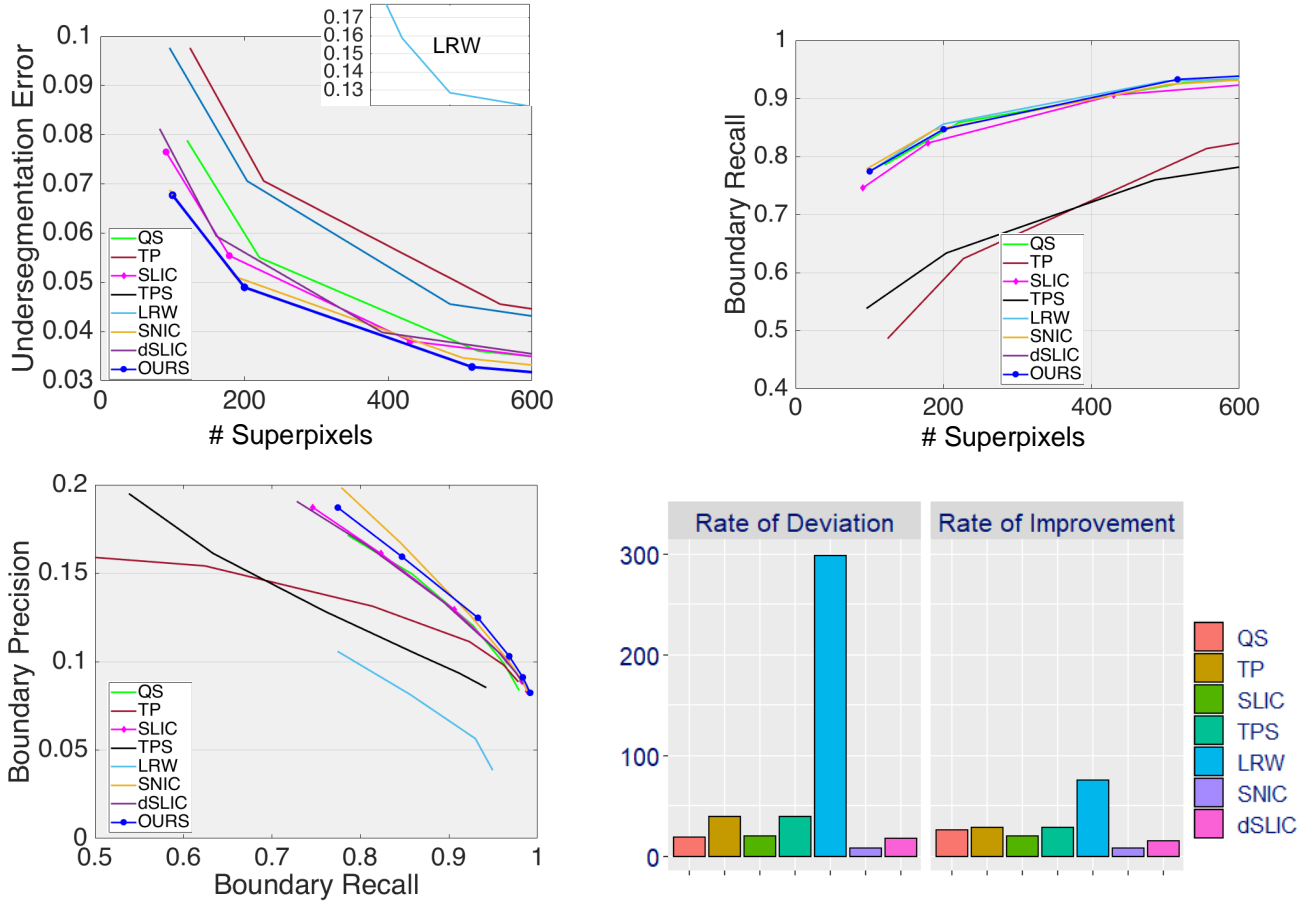


Figure 8. Metric-wise comparison of our approach vs SOTA techniques using UE, BR and BP. In a closer look, we can see that our approach, overall, offers the lowest UE and the highest BR. Finally, the good boundary adherence to the true edges is reflected in the last plot, in which our approach overall gets the best trade-off between those metrics. These results are further supported by the rate of deviation and rate of improvement.

robustness in terms of grouping structures correctly than the previous approaches. However, like SLIC it also tends to generate more superpixels than needed in uniform regions so that the final outputs do not capture fine details. Examples are (G), (E) and (C) face details; (I) the eyes and head and (F) the leaves.

These major drawbacks are mitigated by our model. Our algorithmic approach shows the best boundary adherence and regularity. This is visible in the leaves in image (F), in which our approach is able to better capture the structure, in (I) on the lips where our approach is able to capture the correct geometry, and (A) the fish eyes, where our approach is the only one that correctly segments the inner part. These positive properties of our approach are prevalent in all images. More examples include preservation of the geometry such as in (J) the hand and (G) face, in which our model is the only one able to correctly segment these fine details.

To further support our visual evaluation, we show a metric-wise comparison in Fig. 8. We start by evaluating

the approaches in terms of UE, which is displayed at the left side of this figure. Close observation shows that QS, TPS and TP perform poorly, and in particular, LRW that reported the highest Undersegmentation Error. dSLIC and SLIC show quite low UE, and SNIC ranks the second best. Overall, our approach shows substantial improvement over the compared approaches reporting the lowest UE for all superpixels counts. A similar effect is exhibited in terms of Boundary Recall. TP and TPS perform poorly while the other compared approaches reported better BR. Our approach readily competes with the other compared schemes and the overall BR of our approach was reported to be the best. The same effect is observed in terms of BP-vs-BR, which reflects that our approach overall adheres better to the truth boundaries. To have a better sense of the improvement, we finally report the rate of deviation and rate of improvement of all compared techniques with respect to ours, which results are displayed at the bottom left side of Fig. 8. With respect to the rate of deviation our technique displays




Figure 9. Visual comparison of our approach and SLIC using out-of-the-focus sport images. Zoom-in views display relevant areas, in which our approach successfully retrieve fine detail and blurry structures.

a clear difference ranging from 299% to 8% whilst in terms of rate of improvement a range of 75% to 10%. In particular, the worst performance was reported by LRW in which we improved substantially with 299% and 75% in the rate of deviation and improvement. Our closest competitor is SNIC with which we reduce the error rate by more than 10%. This improvement comes at a negligible cost time which can be seen in Figure 11.

▷ **How is the Computational Performance?** Finally, we evaluate our approach vs the SOTA models in terms of CPU performance in seconds. Results are displayed at Figure 11, using the average time across all images and over the range of [80, 2500] superpixels. From this plot, we observe that TP, TPS and LRW require high computational

time whilst QS and dSLIC slightly improve in this regard. Finally, SLIC, SNIC and OURS provide more feasible run-times that are appropriate for a pre-processing task. We remark that our negligible computational load ($\sim 5\%$) is justified by the substantially improved results over SLIC, SNIC with the same number of superpixels. Overall, the computational demand of our technique is of practical interest for several domains, where decreasing the computational load by reducing the number of primitives in the image domain is needed, including remote sensing, medical applications and video analysis.

▷ **Case Study: Out-of-the Focuss and Wide Field of View Images** To further support our technique, we included two interesting yet challenging cases. The first one




Figure 10. Visual assessment of our technique vs SLIC using wide angle images, in which the objects are distorted and the segmentation task becomes more challenging. Red arrows show interesting areas, in which our technique improves over SLIC.

is superpixels with out-of-the focus images. Whilst the second case is when using wide field of view images.

We start evaluating our approach using out-of-focus images and the visual outputs of our technique and SLIC are displayed in Figures 9 and 12. The study of this case is rele-

vant to assess the effectiveness of our technique in retrieving fine details and rich-structures with blurry and non blurry objects in the image. A common scenario of this case is in sport events due to the fast actions in the scene. In Figure 9, we displayed a set of outputs and in a closer look, one can




Figure 11. CPU time averaged comparison of our approach vs the body of literature. One can see that our approach improvement comes at a negligible cost in runtime in comparison with the fastest approaches SLIC and SNIC.

see that our technique was able to preserve relevant structures. We provide zoom-in snapshots of interesting cases. For example see the face details in all image, in which our technique segmented correctly the eyes, mouth and nose. Another example is the letters in the background from the first, second and third images. We also included a set of different images with different scenarios in Figure 12, in which previous positive effects are also observed. For example at the first image the left eye is correctly retrieved with our technique whilst SLIC failed to segmented correctly. We observed that blurry parts are also well-segmented by our technique, and in particular, also in keeping fine details. Overall, the outputs displayed in Figures 9 and 12 suggest that our proposed method produces visually more pleasant results than SLIC.

Our second case of study is when using wide angle images. The challenging when segmenting is due to the structure of the objects become highly distorted with respect to the distance between the camera point and the object position in the scene. This makes image segmentation harder than rectilinear images. We then investigate if our technique is still efficient in terms of keeping the object shape and fine details. With this purpose, we ran our technique and SLIC over a set of image samples from [36], and the results are reported in Fig. 10. A closer look at the outputs, one can observe that our technique fulfill its purpose of avoiding over-segmenting homogeneous parts such as the sky and focus on the rich parts for example buildings or people. Visual comparison of this effect can be seen as pointed out by the red arrows. Overall, our technique shows potentials for the two investigated cases by generating better output segmentations in terms of retrieving fine details and preserving the

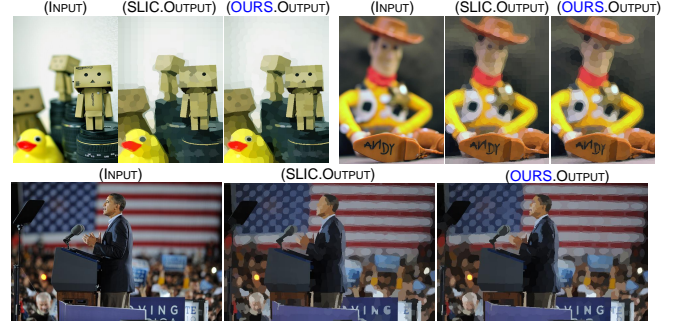


Figure 12. Visual comparison against SLIC using out-of-the-focus images. The selected frames show challenges cases when objects are highly blurry. Overall, our technique gives a better superpixel segmentation by preserving objects structures and fine details.

objects shape.

5. Conclusion

In this work, we proposed a new superpixel approach that builds upon SLIC technique. Our approach incorporates the notion of spectral residual as a proxy for object density and a novel seeding strategy. We demonstrated that our approach seeds clusters advantageously and modifies the local search radius. This leads to better segmentation, with a comparable computational load, to other state-of-the-art algorithms. These modifications leads to major advantages in terms of avoiding the over-segmentation of uniform areas in terms of fine-detail preservation, with a comparable computational load to other state-of-the-art algorithms.

We offer the following two remarks on how future work may reduce the computational load of our technique. First, we observe that our slightly higher computational load is justified by the substantially improved results over SLIC and SNIC with the same number of superpixels. This suggests that, by reducing the number of superpixels in our approach, we could achieve better segmentation than SLIC and SNIC while also reducing computational load. Secondly, we remark that the full saliency measure $\mathcal{SR}(x)$ contains more information that is strictly necessary for our technique, since it is only used as a *rough* measure of object density, while also adding substantially to the computational load. We therefore suggest that a density measure \mathcal{G} could instead be constructed from a *down-sampled* version of the image so that the computational load could be significantly reduced, while maintaining excellent segmentation. For instance, down-sampling by a factor of 2 will reduce the computational load of computing \mathcal{SR} by a factor of 4 and, provided that the superpixel resolution is much lower than the the resolution of the true image, have minimal effect on the initial seeding. We also note that the map \mathcal{SR} used here is certainly not the only possible technique for determin-

ing object density; future research may consider using other techniques, possibly based on machine learning, to replace our object density measure \mathcal{G} . Furthermore, future work will include performance evaluation over a range of applications, where superpixels makes significant positive effect in decreasing the computational cost. For example, in remote sensing, video analysis and medical image analysis.

Acknowledgments

AIAR gratefully acknowledges support from CMIH (EP/N014588/1 and EP/T017961/1) and CCIMI, University of Cambridge. DH is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L016516/1 for the University of Cambridge Centre for Doctoral Training, the Cambridge Centre for Analysis. XZ was partially supported by the Research Grants Council of Hong Kong (Project No. CityU 11301419) and City University of Hong Kong (Project No. 7005497). RC's research is supported by HKRGC Grants No. CUHK 14306316 and CUHK 14301718, CityU Grant 9380101, CRF Grant C1007-15G, AoE/M-05/12. CBS acknowledges support from the Leverhulme Trust project on 'Breaking the non-convexity barrier', the Philip Leverhulme Prize, the Royal Society Wolfson Fellowship, the EPSRC grants EP/S026045/1 and EP/T003553/1, EP/N014588/1, EP/T017961/1, the Wellcome Innovator Award RG98755, European Union Horizon 2020 research and innovation programmes under the Marie Skłodowska-Curie grant agreement No. 777826 NoMADS and No. 691070 CHiPS, the CCIMI and the Alan Turing Institute.

References

- [1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 34(11):2274–2282, 2012.
- [2] R. Achanta and S. Süsstrunk. Superpixels and polygons using simple non-iterative clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4651–4660, 2017.
- [3] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 898–916, 2010.
- [4] J. Chen, J. Hou, Y. Ni, and L.-P. Chau. Accurate light field depth estimation with superpixel regularization over partially occluded regions. *IEEE Transactions on Image Processing (TIP)*, pages 4889–4900, 2018.
- [5] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (5):603–619, 2002.
- [6] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 2004.
- [7] H. Fu, X. Cao, D. Tang, Y. Han, and D. Xu. Regularity preserved superpixels and supervoxels. *IEEE Transactions on Multimedia*, 16(4):1165–1175, 2014.
- [8] R. Giraud, V.-T. Ta, and N. Papadakis. Superpixel-based color transfer. In *IEEE International Conference on Image Processing (ICIP)*, pages 700–704, 2017.
- [9] X. Hou and L. Zhang. Saliency detection: A spectral residual approach. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. Ieee, 2007.
- [10] A. Humayun, F. Li, and J. M. Rehg. The middle child problem: Revisiting parametric min-cut and seeds for object proposals. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1600–1608, 2015.
- [11] A. Levinstein, A. Stere, K. N. Kutulakos, D. J. Fleet, S. J. Dickinson, and K. Siddiqi. Turbopixels: Fast superpixels using geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2009.
- [12] G. Li and Y. Yu. Deep contrast learning for salient object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 478–487, 2016.
- [13] Z. Li and J. Chen. Superpixel segmentation using linear spectral clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1356–1363, 2015.
- [14] F. Liu, C. Shen, and G. Lin. Deep convolutional neural fields for depth estimation from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5162–5170, 2015.
- [15] P. Liu, M. Lyu, I. King, and J. Xu. Selfflow: Self-supervised learning of optical flow. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4571–4580, 2019.
- [16] S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, pages 129–137, 1982.
- [17] G. Maierhofer, D. Heydecker, A. I. Aviles-Rivero, S. M. Alsaleh, and C.-B. Schonlieb. Peekaboo-where are the objects? structure adjusting superpixels. In *IEEE International Conference on Image Processing (ICIP)*, pages 3693–3697, 2018.
- [18] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3061–3070, 2015.
- [19] A. P. Moore, S. J. Prince, J. Warrell, U. Mohammed, and G. Jones. Superpixel lattices. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. Cite-seer, 2008.
- [20] P. Neubert and P. Protzel. Compact watershed and preemptive slic: On improving trade-offs of superpixel segmentation algorithms. In *International Conference on Pattern Recognition*, pages 996–1001, 2014.
- [21] J. Papon, A. Abramov, M. Schoeler, and F. Worgotter. Voxel cloud connectivity segmentation-supervoxels for point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2027–2034, 2013.

- [22] X. Ren and J. Malik. Learning a classification model for segmentation. In *IEEE International Conference on Computer Vision*, page 10, 2003.
- [23] P. Sellars, A. Aviles-Rivero, and C.-B. Schönlieb. Superpixel contracted graph-based learning for hyperspectral image classification. *arXiv preprint arXiv:1903.06548*, 2019.
- [24] J. Shen, Y. Du, W. Wang, and X. Li. Lazy random walks for superpixel segmentation. *IEEE Transactions on Image Processing*, pages 1451–1462, 2014.
- [25] C. Shi and C.-M. Pun. Multiscale superpixel-based hyperspectral image classification using recurrent neural networks with stacked autoencoders. *IEEE Transactions on Multimedia (TMM)*, 2019.
- [26] J. Shi and J. Malik. Normalized cuts and image segmentation. *Departmental Papers (CIS)*, page 107, 2000.
- [27] J. Shi, L. Xu, and J. Jia. Discriminative blur detection features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2965–2972, 2014.
- [28] D. Stutz, A. Hermans, and B. Leibe. Superpixels: An evaluation of the state-of-the-art. *Computer Vision and Image Understanding*, 166:1–27, 2018.
- [29] D. Tang, H. Fu, and X. Cao. Topology preserved regular superpixel. In *IEEE International Conference on Multimedia and Expo*, pages 765–768, 2012.
- [30] N. Tong, H. Lu, X. Ruan, and M.-H. Yang. Salient object detection via bootstrap learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1884–1892, 2015.
- [31] A. Vedaldi and S. Soatto. Quick shift and kernel methods for mode seeking. In *European conference on computer vision*, pages 705–718. Springer, 2008.
- [32] L. Vincent and P. Soille. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (6):583–598, 1991.
- [33] J. Wang and X. Wang. Vcells: Simple and efficient superpixels using edge-weighted centroidal voronoi tessellations. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2012.
- [34] S. Wang, H. Lu, F. Yang, and M.-H. Yang. Superpixel tracking. In *International Conference on Computer Vision (CVPR)*, pages 1323–1330, 2011.
- [35] C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3166–3173, 2013.
- [36] K. Yang, X. Hu, L. M. Bergasa, E. Romera, and K. Wang. Pass: Panoramic annular semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [37] D. Yeo, J. Son, B. Han, and J. Hee Han. Superpixel-based tracking-by-segmentation using markov chains. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1812–1821, 2017.