

HERS Superpixels: Deep Affinity Learning for Hierarchical Entropy Rate Segmentation

Hankui Peng¹, Angelica I. Aviles-Rivero¹, and Carola-Bibiane Schönlieb¹

¹DAMTP, University of Cambridge.
{hp467, ai323, cbs31}@cam.ac.uk

Abstract

*Superpixels serve as a powerful preprocessing tool in numerous computer vision tasks. By using superpixel representation, the number of image primitives can be largely reduced by orders of magnitudes. With the rise of deep learning in recent years, a few works have attempted to feed deeply learned features / graphs into existing classical superpixel techniques. However, none of them are able to produce superpixels in near real-time, which is crucial to the applicability of superpixels in practice. In this work, we propose a two-stage graph-based framework for superpixel segmentation. In the first stage, we introduce an efficient Deep Affinity Learning (DAL) network that learns pairwise pixel affinities by aggregating multi-scale information. In the second stage, we propose a highly efficient superpixel method called Hierarchical Entropy Rate Segmentation (HERS). Using the learned affinities from the first stage, HERS builds a hierarchical tree structure that can produce any number of highly adaptive superpixels instantaneously. We demonstrate, through visual and numerical experiments, the effectiveness and efficiency of our method compared to various state-of-the-art superpixel methods.*¹

1. Introduction

Superpixel segmentation is the task of partitioning an image into meaningful regions, within which the pixels share similar qualities such as colour, texture, or other low-level features. It is a powerful preprocessing tool for various computer vision tasks. For example, image classification [6, 27], optical flow [23, 18], object tracking [35, 38] and semantic segmentation [14, 42]. Superpixels offer more computationally digestible input data representations than

the conventional pixel-level format. By using superpixels, one can substantially reduce the number of image primitives, whilst highlighting the discriminative information [1].

The aforementioned advantages have encouraged the fast advancement of superpixel segmentation techniques (e.g. [15, 1, 17, 16, 21, 34]) since the seminal work of Ren and Malik [26]. For a superpixel segmentation method to be useful, it should be computationally efficient and preserve the structure of the objects well (i.e. adhere to the object boundaries). With the advent of deep learning, it becomes possible to explore more flexible representations in various superpixel segmentation techniques. However, to the best of our knowledge, there are only a few attempts to employ deep networks for superpixels (e.g. [32, 12, 39]).

There are several reasons for this limited adoption. Firstly, the conventional convolution operation in a neural network is designed to work efficiently with a regular image grid, whereas superpixel segmentations naturally give rise to irregular grids. Secondly, there does not exist ground truth superpixels in an image, but rather a delineation of the object structure. Thirdly, several existing superpixel techniques (e.g. [1, 17]) are non-differentiable, due to the nearest neighbour assignment used in the pixel-superpixel association. This imposes a challenge in the network training process which would be end-to-end trainable otherwise.

So far, there are only a few works that have addressed these challenges in integrating deep networks within superpixel methods. One possible approach is to redesign a network architecture to allow the computation of irregular superpixel grids, e.g. [9, 30, 39]. This is still challenging, particularly, if one wishes to integrate subsequent tasks within a single learning process. An alternative approach is to decouple the superpixel segmentation process from the deep network training process. For example, one can use a network to extract pixel-level features and then feed them into a superpixel segmentation method [12]. Most of these existing attempts are based on building upon a modified ver-

¹The code is available at: <https://github.com/hankuipeng/DAL-HERS>

sion of SLIC [1], which is based on Lloyd’s algorithm for k -means clustering, for its efficiency and simplicity. However, SLIC has a few limitations, such as the over-partition in homogeneous regions of an image and high computational cost in texture-rich regions.

As an alternative to clustering-based techniques, a number of researchers have demonstrated the potential of graph-based superpixel techniques [26, 32]. Under this framework, the superpixel segmentation task is translated into a graph partitioning one. The key challenge in this framework lies in obtaining a good partition of the graph. In addition, the problem of learning pixel-wise affinities from a standard deep network that can translate well into edge weights of a graph is not trivial. With that said, the graph-based framework provides several advantages such as computational tractability, natural representation of an image, and desirable tree structure.



Figure 1: HERS produces superpixels that are highly adaptive to the homogeneous regions of an image. **Left:** the original image. **Right:** segmentation by DAL-HERS.

In this work, we present a two-stage graph-based superpixel segmentation framework, where an output example is displayed in Figure 1. In the first stage, we propose a deep network structure, Deep Affinity Learning (DAL), to learn boundary-aware pixel-wise affinities. In the second stage, we propose an efficient graph partitioning method called Hierarchical Entropy Rate Segmentation (HERS). HERS produces highly adaptive superpixels that adhere well to the object boundaries by maximising the entropy rate of the graph. To summarise, our main contributions are:

- We propose DAL network that learns and aggregates multi-scale information, produces boundary-aware affinities, and can be trained efficiently.
- We propose an efficient superpixel segmentation algorithm called HERS that produces highly adaptive superpixels. It builds a hierarchical tree structure that allows instant generation of any number of superpixels.
- We evaluate our proposal on superpixel benchmark datasets that contain a range of indoor and outdoor scenes. Extensive experimental results demonstrate its effectiveness (numerically and visually) and efficiency against various state-of-the-art classical and deep learning based superpixel methods.

2. Related Work

Superpixel segmentation methods have been extensively studied in the literature, with the majority of existing techniques designed from the classical perspective. Most recently, a few works have reported the use of deep networks for superpixels. We review the existing techniques in turn.

Classic techniques for superpixels. Since the pioneering work of [26], several techniques have been proposed which can be roughly divided into patch-based models [8, 31], watershed techniques [10, 4, 20], clustering-based approaches [1, 16, 19] and graph-based techniques [26, 7, 17]. The last two categories are the most widely applied family of techniques, which we will cover in the rest of this section.

One set of techniques have been proposed based on clustering principles for superpixels. The most popular technique in this category is the Simple Linear Iterative Clustering (SLIC) [1]. This technique partitions a given image using a local version of the k -means algorithm. Whilst this technique offers simplicity and efficiency, it has several drawbacks. For example, it unnecessarily partitions uniform areas and computes unnecessary distances in dense areas. These issues motivated several improvements for k -means based techniques including reducing the number of distance calculations, improving the seeding initialisation and improving the feature representation [19, 16, 2, 21, 41].

Another set of techniques treat superpixel segmentation as a graph partitioning problem. The seminal paper of [26] uses the Normalised Cuts algorithm, where superpixels are subgraphs that are obtained as a result of the partitioning based on edge similarities. This technique produced fair results, with boundary adherence being the main issue. Another graph based technique was presented in [7], where the main criterion for graph partitioning is based on the evidence for a boundary between two regions. One of the most competitive graph-based techniques is Entropy Rate Superpixel Segmentation (ERS) [17]. It considers edge weights as transition probabilities of a random walk, and selects edges to form a number of subgraphs / superpixels that maximises the overall entropy rate of the partition. The problem is solved through a lazy greedy algorithm [25], which is not very computationally efficient.

Deep learning techniques for superpixels. Given the impressive results achieved by deep learning in recent years, a few recent works have introduced the idea of deeply learned features or edge affinities, in place of handcrafted ones, in the context of superpixel segmentation. In [12], the authors proposed the SSN model that uses a deep network to extract pixel-level features, and then use the learned features as input to a soft version of SLIC [1]. The key idea is to enforce soft-associations between pixels and superpixels to avoid the non-differentiability of SLIC. Although SSN extends the capability of SLIC through learned features, it

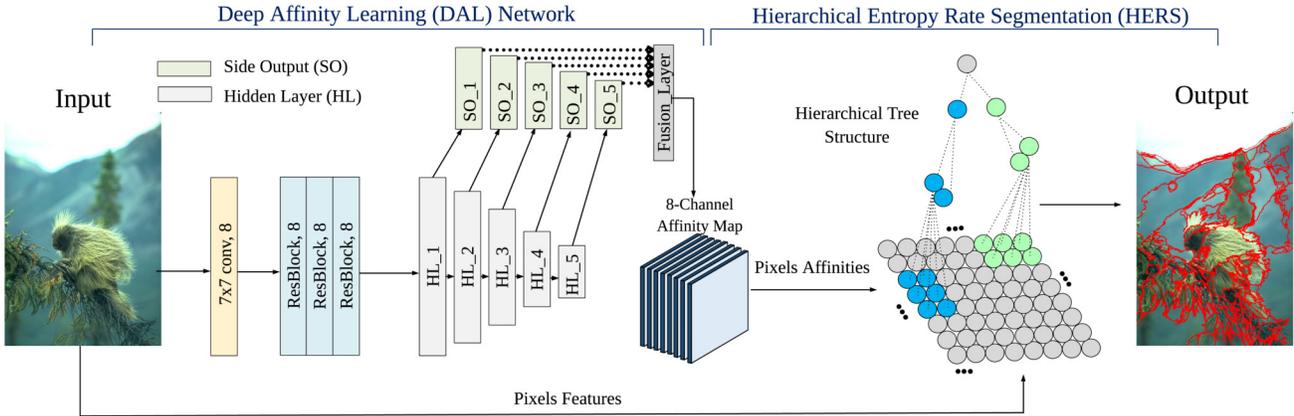


Figure 2: Illustration of our proposed superpixel framework. The left side of the scheme shows the flow of our proposed DAL network, which learns an 8-channel pixel affinity map. The right side displays our proposed HERS algorithm, which constructs a hierarchical tree structure that allows any number of highly adaptive superpixels to be generated instantaneously.

produces superpixels which mimics that of SLIC, which are not very adaptive in sizes. The authors of [32] introduced a graph-based model SEAL, where the key idea is to learn the pixel affinities for superpixel segmentation. This work is built upon ERS [17] along with a segmentation-aware loss. Although it produces very impressive results quantitatively, there are a number of drawbacks in its network training process and the resulting superpixels are not very regular or adaptive in size. Most recently, [39] uses fully convolutional network in an Encoder-Decoder structure to learn the association between pixels and superpixels on a regular image grid. As a result, it cannot produce the exact user-specified number of superpixels. Additionally, this method also requires a post-processing technique to remove any unwanted superpixels that are tiny in sizes.

Distinctions in our approach. Our technique is close in philosophy to that of [32] and [17]. Both our work and [32] seek to learn pixel-wise affinities, however our proposal has several major advantages. Firstly, the work of [32] is limited by its own design – as it only allows the learning of pairwise pixel affinities for one direction at a time. Therefore, the network needs to be applied twice to obtain both horizontal and vertical pixel affinities, which results only in a 4-connected affinity map. By contrast, our work learns a richer 8-connected affinity map within one training process. Secondly, the loss design of [32] requires applying a superpixel segmentation method for each training epoch, which is computationally costly. Additionally, it requires incorporating external edge information, which does not necessarily generalise well on the training set. In comparison, our work efficiently enforces boundary adherence by learning directly from the network without the need for super-

pixel segmentation or incorporation of external edge information throughout training. Finally, we introduce an efficient graph-based technique called Hierarchical Entropy Rate Segmentation (HERS). It can produce highly adaptive superpixels instantaneously without the need for any parameter. Notably, we offer a computational complexity of $\mathcal{O}(|\mathcal{V}|)$, as opposed to $\mathcal{O}(|\mathcal{V}|^2 \log(|\mathcal{V}|))$ [17].

3. The Proposed Methodology

In this section, we present the two core parts of our proposed graph-based superpixel segmentation framework: i) our proposed **Deep Affinity Learning (DAL)** neural network architecture for obtaining deeply learned pixel affinities, and ii) our proposed **Hierarchical Entropy Rate Segmentation (HERS)** algorithm. An overview of our proposal is displayed in Figure 2.

3.1. Deep Affinity Learning

In the first stage of our proposed framework, our goal is to design an effective and efficient network training scheme that produces an 8-connected affinity map $A \in \mathbb{R}^{8 \times H \times W}$ for any input image of height H and width W . The total number of pixels in an image is given by $N = W \times H$. For each pixel p_i , we compute the affinities between p_i and a maximum of 8 surrounding pixels (horizontally and diagonally) that lie in its closest neighbourhood \mathcal{N}_i .

Network design. Our proposed Deep Affinity Learning network, DAL for short, consists of two parts. In the first part, a 7×7 convolutional kernel is used in the first layer to capture both the horizontal and vertical changes in an image. This is followed by 3 standard residual blocks (ResBlock) [11], each of which has a kernel size of 3 and there

are 8 input and output channels. Using 3 ResBlocks would expedite the decrease of the training loss without inducing too much additional computational burden and extra model parameters. Using 8 input and output channels allows us to obtain an intermediate 8-channel affinity map that captures pairwise pixel proximity towards all possible directions.

In the second part, we seek to enforce the preservation of fine details and scene structures. To achieve this, we integrate the HED network structure from [37]. That is, we incorporate five continuous blocks of convolutional layers with increasing receptive field sizes (64, 128, 256, 512, 512) to capture neighbourhood information at varying scales. Each of these five varying scale blocks produces a side output of 8 channels through an additional convolutional layer. Finally, a weighted fusion layer is used to automatically learn how to combine the side outputs from multiple scales. Concretely, the weighted fusion layer is a convolutional layer with 40 (5×8) input channels and 8 output channels to produce the final 8-channel affinity map A .

Loss function. Each image comes with a ground truth segmentation mask $M \in \mathbb{R}^{H \times W}$, where $m_{ij} \in \{1, \dots, C\}$ denotes the pixel-specific class label, and C is the total number of classes in the image. By comparing the label of each pixel \mathbf{p}_i with the labels of its maximally 8 surrounding pixels, we can transform M into an 8-channel binary segmentation map $T \in \{0, 1\}^{8 \times H \times W}$, indicating whether two pairwise pixels belong to the same class label or not.

As such, the training loss can be quantified separately for both boundary pixels and non-boundary pixels according to T . The most common way of measuring the loss between the learned affinity map A and the ground truth map $T \in \{0, 1\}^{8 \times H \times W}$ is through the following Binary Cross Entropy (BCE) loss:

$$\begin{aligned} \mathcal{L}_{BCE}(A, T) = & -\frac{1}{8N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} (1 - t_{ij}) \log(1 - a_{ij}) \\ & - \frac{1}{8N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} t_{ij} \log(a_{ij}), \end{aligned} \quad (1)$$

where $t_{ij} \in \{0, 1\}$ denotes the ground truth relationship between pixel \mathbf{p}_i and \mathbf{p}_j , and $a_{ij} \in [0, 1]$ denotes the learned affinity between pixel \mathbf{p}_i and \mathbf{p}_j .

It can be seen that (1) encourages the learned affinities to be zeros for boundary pixels, and ones for non-boundary pixels. However, the ground truth segmentations are often given for object detection instead of superpixel segmentation. Therefore, there is potentially no supervision information available in the heterogeneous regions within an object. Motivated by this, we propose the following modified version of (1):

$$\begin{aligned} \mathcal{L}_{DAL}(A, T) = & -\frac{1}{8N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} (1 - t_{ij}) \cdot \log(1 - a_{ij}) \\ & - \frac{1}{8N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} t_{ij} \cdot |g_{ij} - a_{ij}|, \end{aligned} \quad (2)$$

where g_{ij} denotes the pre-computed pairwise pixel affinity between \mathbf{p}_i and \mathbf{p}_j . We use the Gaussian similarity to compute the pairwise pixel affinity g_{ij} as: $g_{ij} = -\exp\{d(\mathbf{p}_i, \mathbf{p}_j)/2\sigma^2\}$ [5], in which σ is the bandwidth parameter and $d(\mathbf{p}_i, \mathbf{p}_j)$ denotes the distance between pixel \mathbf{p}_i and \mathbf{p}_j . It can be computed as the ℓ_2 distance between the RGB pixel features. As such, the loss in (2) would encourage the learning of boundary information just as in the BCE loss, whilst learning additional local pixel affinity information for the non-boundary pixels guided by Gaussian similarity. Note that this is very different to the BCE loss, in which all non-boundary pairwise pixel affinities are treated as 1s. Since the majority of pixels in an image are non-boundary pixels, training the network using (2) enjoys a substantial advantage over the use of the BCE loss.

3.2. Hierarchical Entropy Rate Segmentation

After obtaining an 8-channel affinity map $A \in \mathbb{R}^{8 \times H \times W}$ from our trained network, the second part of our proposed framework is to use the extracted rich information to generate superpixels. To achieve this, we introduce Hierarchical Entropy Rate Segmentation (HERS) algorithm next.

The idea is to represent the extracted information as an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} denotes the set of nodes on the graph that correspond to the image pixels. The adjacent nodes / pixels are connected by a set of edges \mathcal{E} , whose weights reflect their pairwise similarities. We wish to select a subset of edges $E \subset \mathcal{E}$ from the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ such that the entropy rate of the segmentation is maximised. It has been shown in [17] that higher entropy rate corresponds to more compact and homogeneous clusters. The entropy rate is used to describe the uncertainty of a stochastic process. By modelling a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as a first-order Markov process, we can obtain the entropy rate of the graph as a conditional entropy.

In this framework, the transition probability between a pair of nodes v_i and v_j can be expressed as $p_{ij} = w_{ij}/w_i$, in which $w_i = \sum_{k: e_{ik} \in \mathcal{E}} w_{ik}$ denotes the sum of edge weights connecting to node v_i . Here w_{ij} denotes the pairwise similarity between v_i and v_j , which can be readily obtained from the affinity map A as $(a_{ij} + a_{ji})/2$. A stationary distribution of the Markov process can be given as follows:

$$\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_{|\mathcal{V}|}] = \left[\frac{w_1}{w_T}, \frac{w_2}{w_T}, \dots, \frac{w_{|\mathcal{V}|}}{w_T} \right]^T. \quad (3)$$

Then the entropy rate of the graph can be expressed as

$$\mathcal{H}(\mathcal{E}) = - \sum_i \frac{w_i}{w_T} \sum_j \frac{w_{ij}}{w_i} \log \frac{w_{ij}}{w_i}, \quad (4)$$

where $w_T = \sum_{i=1}^{|\mathcal{V}|} w_i$ denotes the sum of weights across all nodes. Our goal is to sequentially select edges, guided by the entropy rate criterion in (4), in order to *form balanced superpixels that are adaptive in sizes*. To this end, we propose to use Borůvka’s algorithm to optimise solely the entropy rate of the graph, whilst obtaining balanced and adaptive superpixels.

Borůvka’s algorithm is one of the oldest algorithms designed to identify the minimum spanning tree (MST) of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. It simultaneously goes through the following steps for all currently formed trees in the graph: i) identify the best edge for each tree, which consists of one node or a set of nodes that are already connected by existing edges, and ii) add those identified edges to the forest until one minimum spanning tree is formed connecting all nodes in the graph. Borůvka’s algorithm enjoys a number of desirable properties. Firstly, instead of having to sort all the edges globally as in the lazy greedy algorithm adopted by [17], Borůvka’s algorithm finds the best edges for all trees simultaneously. As a result, we can offer a computational complexity of $\mathcal{O}(|\mathcal{V}|)$, as opposed to $\mathcal{O}(|\mathcal{V}|^2 \log(|\mathcal{V}|))$ as in [17]. In addition, the parallel edge selection process in Borůvka’s algorithm allows for the generation of very balanced trees. Finally, by recording the orders in which the edges are added, a hierarchical structure can be formed which allows any number of trees to be generated instantaneously.

The idea of using Borůvka’s algorithm for the purpose of superpixel segmentation has been previously utilised in [36]. However, [36] considers edge weights as dissimilarities between pairs of nodes, whereas we associate edge weights with their potential contributions to the overall entropy rate of the graph. Concretely, the edge weight can be interpreted in terms of its contribution to the overall entropy rate of the graph, if it were to be added to the current set of selected edges E . Using the learned affinities $a_{ij} \in A$ for the edge weight w_{ij} between node v_i and v_j , the edge weight w_{ij} can be calculated as:

$$\begin{aligned} w_{ij} &= \mathcal{H}(E \cup \{e_{ij}\}) - \mathcal{H}(E) \\ &= (w_{ij} + \mu_i) \log(w_{ij} + \mu_i) + (w_{ij} + \mu_j) \log(w_{ij} + \mu_j) \\ &\quad - \mu_i \log \mu_i - \mu_j \log \mu_j - 2w_{ij} \log w_{ij}. \end{aligned}$$

We refer to our proposal of sequential edge selection guided by entropy rate via Borůvka’s algorithm as **Hierarchical Entropy Rate Segmentation (HERS)**. The algorithmic form of HERS is presented in Algorithm 1. Note that the number of superpixels K is not necessarily required as an input to the algorithm. Instead, a segmentation with

Algorithm 1: HERS Superpixels

```

Input: Graph:  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ 
1 Output: A minimum spanning tree  $T$ ; A set of
   selected edges  $E \subset \mathcal{E}$ 
2 Constructed tree:  $T = \emptyset$ ; Selected edges:  $E = \emptyset$ ;
   Unselected edges:  $U = \mathcal{E}$ 
3 function Borůvka( $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ )
4   do
5     - Identify the currently formed trees
        $T_1, T_2, \dots, T_k$  where  $k \leq |\mathcal{V}|$ 
6     - Find the best outgoing edge for each tree
        $e^*(T_i) = \arg \max_{e_{ij} \in U_i} \mathcal{H}(E \cup \{e_{ij}\}) - \mathcal{H}(E)$ ,
       where  $U_i$  denotes the unselected outgoing
       edges for tree  $T_i$ 
7     - Sort the best outgoing edges in descending
       order:  $e^*(T_1), e^*(T_2), \dots, e^*(T_k)$ 
8     for  $\ell = 1, \dots, k$  do
9       | Add the  $\ell$ -th outgoing edge to  $E$ ;
10      |  $K \leftarrow K - 1$ ;
11     end
12   while  $|T| > 1$ ;
13 end

```

any number of superpixels can be obtained instantaneously from the hierarchy that is constructed via Algorithm 1.

4. Experimental Results

In this section, we describe in detail the range of experiments that we conducted to demonstrate the effectiveness and efficiency of our proposed methodology.

Dataset Description & Evaluation Protocol. We evaluate the performance of our technique using two benchmark datasets for superpixels. We use The Berkeley Segmentation Dataset 500 (BSDS500) [3]. This dataset contains 500 images with provided ground truth segmentations. It provides a wide variety of outdoor scenes with different complex structures. We also use the NYU Depth Dataset V2 (NYUv2) [29], which consists of 1449 images with provided ground truth segmentations. The NYUv2 dataset provides a range of different indoor scenes.

We first justify the model design and support the advantage of our technique through a set of ablation studies. We then compare our technique with the following state-of-the-art techniques: i) classic techniques: ERS [17], SH [36], SLIC [1], SNIC [2], SEEDS [33], ETPS [40]; and ii) deep learning techniques: SSN [12], SEAL-ERS [32] and SP-FCN [39]. The most commonly used performance measures for evaluating superpixel segmentation algorithms include: Under-segmentation Error (UE) [33], Achievable Segmentation Accuracy (ASA) [17] and Boundary Recall

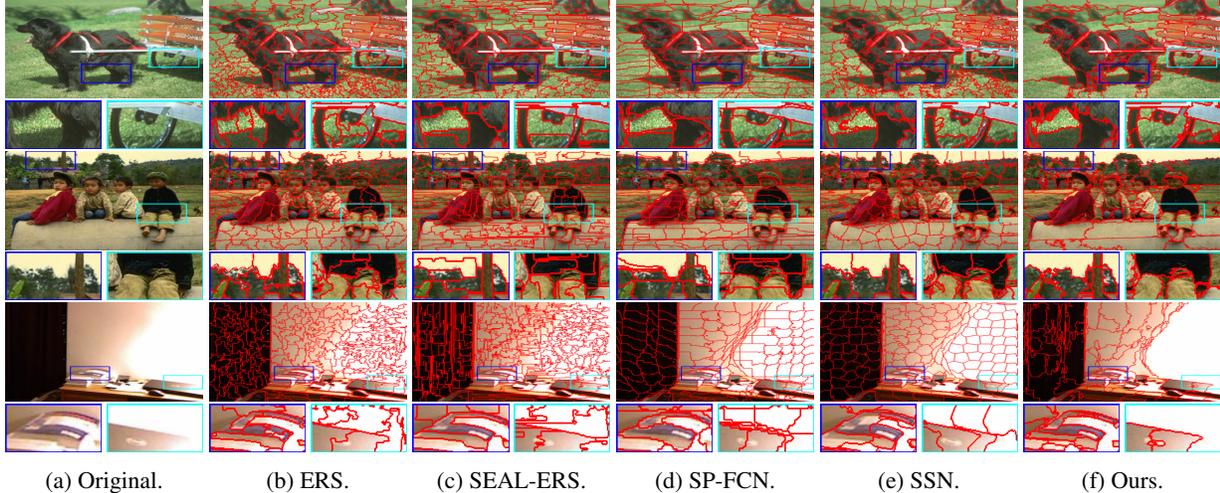


Figure 3: Visual comparison of our technique against a number of state-of-the-art algorithms using 200 superpixels. The selected images show variations in scenes from BSDS500 and NYUv2 datasets.

(BR) [22]. In particular, UE can be directly obtained from ASA as 1-ASA. Thus, we report the ASA and BR measures in our experiments. Additionally, we also report the Explained Variation (EV) [24] that quantifies the variance within an image that is captured by the superpixels without relying on any ground truth labels. Explicit definitions of the metrics can be found in the supplementary material.

Implementation & Training Details. We learn an 8-channel affinity map by training our proposed DAL network on the BSDS500 training set. We implement the network in PyTorch using the Adam [13] optimiser with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The DAL network is trained for 5k epochs, in which the input images are cropped to have size 200 by 200. The initial learning rate is set to $1e - 4$, and is reduced by a factor of 10 after 3k epochs. For all the competing methods, we use the code provided by the corresponding authors.

4.1. Ablation Study

Our proposed methodology contains two main building blocks including the deeply learned affinities from the DAL network and the proposed HERS algorithm. In this section, we evaluate the contribution from each of these components. The comparison across different variants are reported in Figure 4 in terms of EV, BR, and ASA scores.

Benefit of the proposed DAL network. The benefit of using learned affinities from the DAL network as opposed to handcrafted ones can be observed by comparing the performance of ERS (green) with that of DAL-ERS (pink) in Figure 4. ERS computes the pixel-wise affinities using the Gaussian kernel in which the RGB pixel values are used as features. Whereas DAL-ERS directly uses the learned affinities from the DAL network as input to the ERS algorithm. It is clear that DAL-ERS (pink) largely outperforms the baseline ERS algorithm (green) in all three measures.

Benefit of the proposed HERS algorithm. Next, we inspect the contribution of our proposed superpixel segmentation algorithm HERS. It is worth pointing out that HERS does not require any balancing term in the objective function and is therefore parameter-free. Whereas the balancing term in ERS plays a crucial role in avoiding extremely unbalanced superpixels.

By comparing DAL-HERS (cyan) with DAL-ERS (pink), we can see that the former enjoys much higher EV and BR scores at the cost of a lower ASA score. This trade-off can be explained by the fact that the superpixels produced by HERS are highly adaptive, i.e. *it keeps large homogeneous regions of an image intact whilst over-segmenting the texture-rich regions*. Adaptive superpixels are arguably more desirable than superpixels whose sizes are agnostic to the semantics of the image. However such adaptive behaviour has certain implications on the performance measures. Since the superpixels produced by HERS adhere strongly to the object boundaries and preserves the homogeneous regions, it is to be expected that it enjoys very good EV and BR scores. However, having adaptive superpixels also means that a small “leakage” in the boundary adherence would incur a big under-segmentation error, which negatively correlates with the ASA score.

Benefit of external edge information. For the purpose of achieving a higher ASA score, we could compromise a small amount of the performance gain in terms of EV and BR scores. The ASA score can be improved by strengthening the adherence to object boundaries given by the ground truth segmentations. Therefore, we additionally divide the edge weight by HED edge probabilities, which corresponds to DAL-HERS (HED) in Figure 4. As a result, we observe a notable increase in the ASA score of DAL-HERS (HED) as compared to that of DAL-HERS. The gain in the ASA

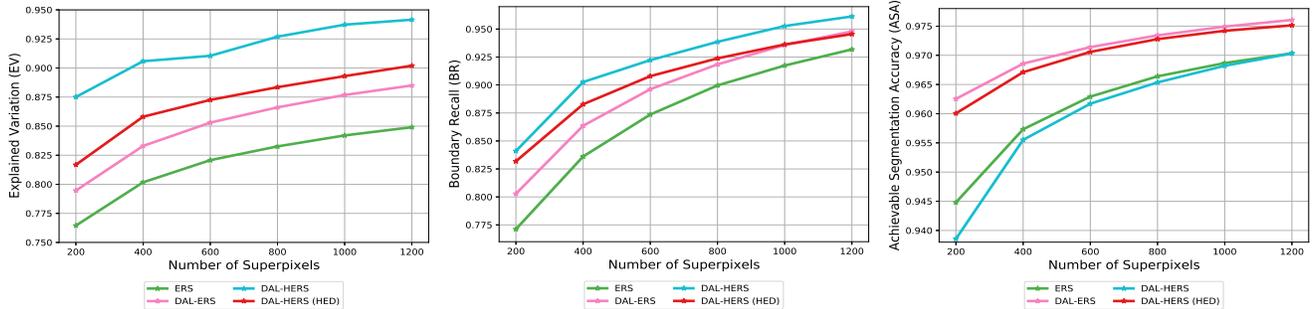


Figure 4: **Ablation study on the BSDS500 test set.** We compare the components in our proposed framework, which includes both the learned affinities from the DAL network and the HERS algorithm.

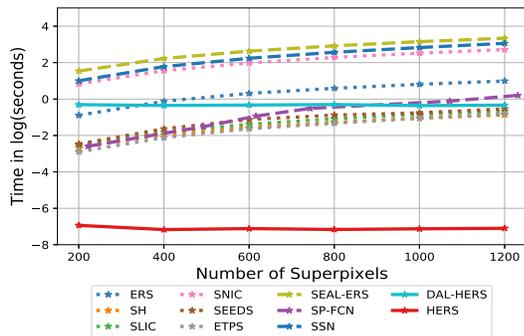


Figure 5: Cumulative runtime in \log_e (seconds) on the whole BSDS500 test dataset.

score comes at a small cost of the EV and BR scores. This is again to be expected, because highlighting more boundary pixels would inevitably increase the false positives of boundary pixels thus decrease the BR and EV scores.

Quantitative performance. Figure 6 shows the quantitative comparison across all methods on both the BSDS500 and the NYUv2 test set. We include both DAL-HERS and DAL-HERS (HED) in the comparison, and refer to them as Ours and Ours (HED) for short. It can be seen that Ours outperforms all other competing methods in terms of EV and BR scores on both datasets. This means that the superpixels generated by our method fully captures the semantically homogenous regions of images, and at the same time strongly adheres to the object boundaries (see Figure 3). Due to the highly adaptive nature of our superpixels, it does not lead to a high ASA score by design. The ASA score measures the overlap between the computed superpixels with the ground truth, where the ground truth labels are provided for object detection or semantic segmentation. As such, they do not delineate the fine details that our method captures so well, which results in the relatively low ASA score. To mitigate this, we could further incorporate HED edge information (as is explained in the ablation study). As a result, Ours (HED) achieves competitive performance against the majority of the superpixel methods that are being compared to.

Qualitative performance. In addition to the quantita-

tive advantages exhibited by our proposed schemes, additional advantages can be demonstrated via visual comparisons of the segmentation results. In Figure 3, we observe several advantages of our method over the compared ones. Firstly, our technique prevents the over-segmentation of homogeneous regions in the scene. Clear examples of this effect can be seen in the grass on the first row, the sky on the second row, and the walls on the third and fourth rows. In all images, our method preserves fine details on the objects by focusing on rich-structure parts rather than uniform regions. Secondly, our technique displays the best boundary adherence. That is, our technique is able to better capture the object structures well. Examples of this property can be found by inspecting the zoomed-in views of the book and chair on the third and fourth rows. We provide further visual comparisons in the supplementary material.

4.2. Comparison to State-of-the-Art Techniques

Computational efficiency. Another highly desirable property of any superpixel technique, as a stand-alone pre-processing tool, is the computational efficiency. We report the runtime (in log seconds) of our proposed DAL-HERS and HERS against other state-of-the-art methods in Figure 5. To highlight the advantage of HERS, we report the cumulative runtime of each method for various numbers of superpixels. For example, the time comparison at 400 on the x -axis reports the time that a method requires to produce both 200 and 400 superpixels.

It is notable that the runtimes of both DAL-HERS and HERS are constant across various numbers of superpixels. This is because HERS constructs a hierarchical tree structure from which any number of superpixels can be extracted instantaneously, whereas all other methods have to run several times for multiple numbers of superpixels. When comparing HERS to other classical methods, it can be seen that HERS has a noticeable computational advantage against other methods. When comparing DAL-HERS to other deep methods, it is clear that it still enjoys competitive performance. A main strength of our technique is the *constant*

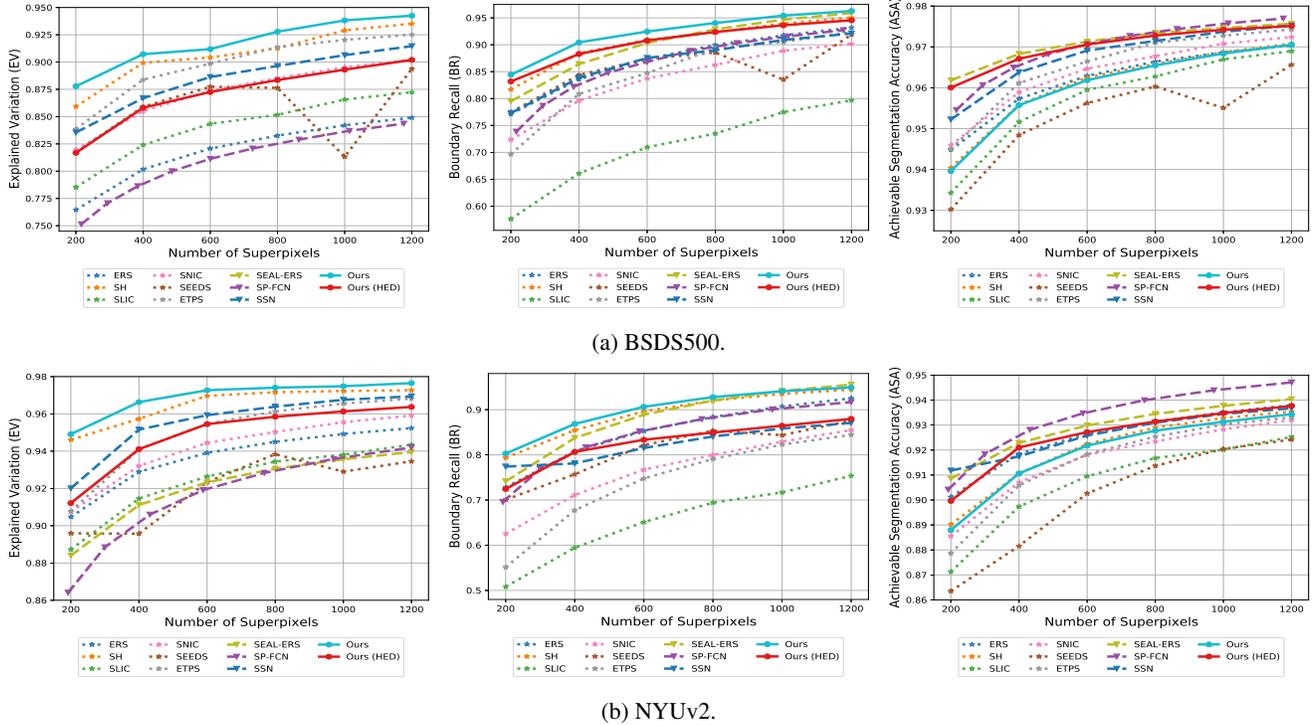


Figure 6: Comparison of our approach (solid lines) against the state-of-the-art of classical approaches (dotted lines) and deep learning based approaches (dashed lines) for superpixels. The performance of various methods are evaluated using EV, BR and ASA scores over different numbers of superpixels.

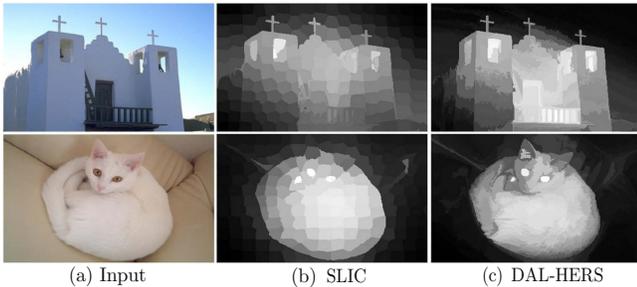


Figure 7: Visual comparison of saliency detection using SLIC and DAL-HERS with 200 superpixels on selected images from the ECSSD dataset [28].

time required to produce *any number of superpixels especially as the numbers of superpixels increase.*

Superpixels for saliency detection. To further support the advantage of DAL-HERS, we report results for the downstream task of saliency detection for a selection of images from the ECSSD dataset [28]. we used the Saliency Optimisation (SO) [43] technique as backbone, which originally uses SLIC to preprocess an image [1]. We replace SLIC with DAL-HERS and present a visual comparison in Figure 7. One can see that DAL-HERS produces smoother outputs whilst preserving better edges than SLIC. We provide a detailed explanation with metric and visual compari-

son in the supplementary material.

5. Conclusion & Future Work

In this paper, we present a graph-based framework consisting of a network for obtaining deeply learned affinities, and an efficient superpixel segmentation method for producing adaptive superpixels. Through experimental results, we show that our technique compares favourably against state-of-the-art methods. Moreover, unlike existing methods demanding linear time with respect to various numbers of user-specified superpixels, our technique exhibits constant time which makes it appealing to be used as a preprocessing tool. For future work, our method could be tailored to various computer vision tasks such as semantic segmentation, saliency detection, and stereo matching. In addition, it would be interesting to combine our deeply learned affinities with other graph-based superpixel methods.

Acknowledgements

HP acknowledges the financial support by Aviva Plc. AIAR acknowledges the CMIH and CCIMI, University of Cambridge. CBS acknowledges the Philip Leverhulme, the EPSRC grants EP/S026045/1 and EP/T003553/1, EP/N014588/1, EP/T017961/1, NoMADS 777826, the CCIMI, University of Cambridge.

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012.
- [2] Radhakrishna Achanta and Sabine Süsstrunk. Superpixels and polygons using simple non-iterative clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4651–4660, 2017.
- [3] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):898–916, 2010.
- [4] Wanda Benesova and Michal Kottman. Fast superpixel segmentation using morphological processing. In *Conference on Machine Vision and Machine Learning*, pages 67–1, 2014.
- [5] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 359–374. Springer, 2001.
- [6] Leyuan Fang, Shutao Li, Xudong Kang, and Jon Atli Benediktsson. Spectral–spatial classification of hyperspectral images with a superpixel-based discriminative sparse model. *IEEE Transactions on Geoscience and Remote Sensing*, 53(8):4186–4201, 2015.
- [7] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2):167–181, 2004.
- [8] Huazhu Fu, Xiaochun Cao, Dai Tang, Yahong Han, and Dong Xu. Regularity preserved superpixels and supervoxels. *IEEE Transactions on Multimedia*, 16(4):1165–1175, 2014.
- [9] Raghudeep Gadde, Varun Jampani, Martin Kiefel, Daniel Kappler, and Peter V Gehler. Superpixel convolutional networks using bilateral inceptions. In *Proceedings of the European Conference on Computer Vision*, pages 597–613. Springer, 2016.
- [10] Josep M Gonfaus, Xavier Boix, Joost Van de Weijer, Andrew D Bagdanov, Joan Serrat, and Jordi Gonzalez. Harmony potentials for joint classification and segmentation. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3280–3287, 2010.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [12] Varun Jampani, Deqing Sun, Ming-Yu Liu, Ming-Hsuan Yang, and Jan Kautz. Superpixel sampling networks. In *Proceedings of the European Conference on Computer Vision*, pages 352–368, 2018.
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [14] Suha Kwak, Seunghoon Hong, and Bohyung Han. Weakly supervised semantic segmentation using superpixel pooling network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 1, page 2, 2017.
- [15] Alex Levinstein, Adrian Stere, Kiriakos N Kutulakos, David J Fleet, Sven J Dickinson, and Kaleem Siddiqi. TurboPixels: Fast superpixels using geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2290–2297, 2009.
- [16] Zhengqin Li and Jiansheng Chen. Superpixel segmentation using linear spectral clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1356–1363, 2015.
- [17] Ming-Yu Liu, Oncel Tuzel, Srikumar Ramalingam, and Rama Chellappa. Entropy rate superpixel segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2097–2104. IEEE, 2011.
- [18] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4571–4580, 2019.
- [19] Yong-Jin Liu, Cheng-Chi Yu, Min-Jing Yu, and Ying He. Manifold SLIC: A fast method to compute content-sensitive superpixels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 651–659, 2016.
- [20] Vaia Machairas, Matthieu Faessel, David Cárdenas-Peña, Théodore Chabardes, Thomas Walter, and Etienne Decencière. Waterpixels. *IEEE Transactions on Image Processing*, 24(11):3707–3716, 2015.
- [21] Georg Maierhofer, Daniel Heydecker, Angelica I Aviles-Rivero, Samar M Alsaleh, and Carola-Bibiane Schönlieb. Peekaboo - Where are the objects? Structure adjusting superpixels. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3693–3697. IEEE, 2018.
- [22] David R Martin, Charless C Fowlkes, and Jitendra Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(5):530–549, 2004.
- [23] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015.
- [24] Alastair P Moore, Simon JD Prince, Jonathan Warrell, Umar Mohammed, and Graham Jones. Superpixel lattices. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [25] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming*, 14(1):265–294, 1978.
- [26] Xiaofeng Ren and Jitendra Malik. Learning a classification model for segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10–17, 2003.
- [27] Philip Sellars, Angelica I Aviles-Rivero, and Carola-Bibiane Schönlieb. Superpixel contracted graph-based learning for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 58(6):4180–4193, 2020.
- [28] Jianping Shi, Qiong Yan, Li Xu, and Jiaya Jia. Hierarchical image saliency detection on extended CSSD. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence*, 38(4):717–729, 2015.
- [29] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *European Conference on Computer Vision*, pages 746–760. Springer, 2012.
- [30] Teppei Suzuki, Shuichi Akizuki, Naoki Kato, and Yoshimitsu Aoki. Superpixel convolution for segmentation. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3249–3253. IEEE, 2018.
- [31] Dai Tang, Huazhu Fu, and Xiaochun Cao. Topology preserved regular superpixel. In *2012 IEEE International Conference on Multimedia and Expo*, pages 765–768. IEEE, 2012.
- [32] Wei-Chih Tu, Ming-Yu Liu, Varun Jampani, Deqing Sun, Shao-Yi Chien, Ming-Hsuan Yang, and Jan Kautz. Learning superpixels with segmentation-aware affinity loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 568–576, 2018.
- [33] Michael Van den Bergh, Xavier Boix, Gemma Roig, Benjamin de Capitani, and Luc Van Gool. SEEDS: Superpixels extracted via energy-driven sampling. In *Proceedings of the European Conference on Computer Vision*, pages 13–26. Springer, 2012.
- [34] Andrea Vedaldi and Stefano Soatto. Quick shift and kernel methods for mode seeking. In *European Conference on Computer Vision*, pages 705–718. Springer, 2008.
- [35] Shu Wang, Huchuan Lu, Fan Yang, and Ming-Hsuan Yang. Superpixel tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1323–1330, 2011.
- [36] Xing Wei, Qingxiong Yang, Yihong Gong, Narendra Ahuja, and Ming-Hsuan Yang. Superpixel hierarchy. *IEEE Transactions on Image Processing*, 27(10):4838–4849, 2018.
- [37] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1395–1403, 2015.
- [38] Fan Yang, Huchuan Lu, and Ming-Hsuan Yang. Robust superpixel tracking. *IEEE Transactions on Image Processing*, 23(4):1639–1651, 2014.
- [39] Fengting Yang, Qian Sun, Hailin Jin, and Zihan Zhou. Superpixel segmentation with fully convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 13964–13973, 2020.
- [40] Jian Yao, Marko Boben, Sanja Fidler, and Raquel Urtasun. Real-time coarse-to-fine topologically preserving segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2947–2955, 2015.
- [41] Yuhang Zhang, Richard Hartley, John Mashford, and Stewart Burn. Superpixels via pseudo-boolean optimization. In *2011 International Conference on Computer Vision*, pages 1387–1394. IEEE, 2011.
- [42] Wei Zhao, Yi Fu, Xiaosong Wei, and Hai Wang. An improved image semantic segmentation method based on superpixels and conditional random fields. *Applied Sciences*, 8(5):837, 2018.
- [43] Wangjiang Zhu, Shuang Liang, Yichen Wei, and Jian Sun. Saliency optimization from robust background detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2814–2821, 2014.