

Salient Object Detection and Segmentation

Journal:	<i>Transactions on Pattern Analysis and Machine Intelligence</i>
Manuscript ID:	TPAMI-2013-06-0454.R1
Manuscript Type:	Regular
Keywords:	I.4.8 Scene Analysis < I.4 Image Processing and Computer Vision < I Computing Methodologies, I.4.6 Segmentation < I.4 Image Processing and Computer Vision < I Computing Methodologies, I.2.10 Vision and Scene Understanding < I.2 Artificial Intelligence < I Computing Methodologies

SCHOLARONE™
Manuscripts

Summary of Changes

Paper title: Salient Object Detection and Segmentation

Authors: M. Cheng, N. J. Mitra, X. Huang, P.H.S. Torr, and S.M. Hu

Original log number: TPAMI-2013-06-0454, TPAMI-2011-10-0753.R1

We thank the editor and the reviewers for their comments and valuable suggestions. We have revised our paper according to the suggestions and summarize below the changes made in this revision.

Response to Editor Comments

We thank the editor for supporting our paper. We have revised the paper according to Reviewer 1 and Reviewer 2's comments, to discuss more clearly the new contributions presented in this TPAMI manuscript. Please see our response to reviewer comments below.

Response to Reviewer 1 comments

Comment 1. "They (i.e. changes made to methods such as spatial prior, border regions) seem a bit hacky, and rely on a few parameters fitted to the author's dataset. However, the extended quantitative evaluation, new benchmark dataset, and additional technical details the authors provide for the proposed methods -- all are positive additions."

Response: We thank the reviewer for acknowledging the positive additions we made to this TPAMI manuscript. Regarding the several parameters used in the spatial prior and border region steps, we would like to clarify that these algorithmic steps are not parameter sensitive. A default set of parameters can be used for processing all images. To demonstrate this, we directly ran our binary source code (made publicly available in 2011 with the submission of TPAMI-2011-10-0753) with default parameters to test on the 'Extended Complex Scene Saliency Dataset'¹ provided by The Chinese University of Hong Kong. We found our results to outperform their CVPR2013 results (Hierarchical Saliency Detection, Qiong Yan, Li Xu, Jianping Shi, Jiaya Jia, CVPR 2013). The statistics are shown below. Although there are many methods that have been developed since the initial version of this work, to the best of our knowledge, our salient object segmentation results are still the best results reported on the most widely used benchmark. We have added related discussion to the manuscript.

¹ <http://www.cse.cuhk.edu.hk/leoija/projects/hsaliency/dataset.html>

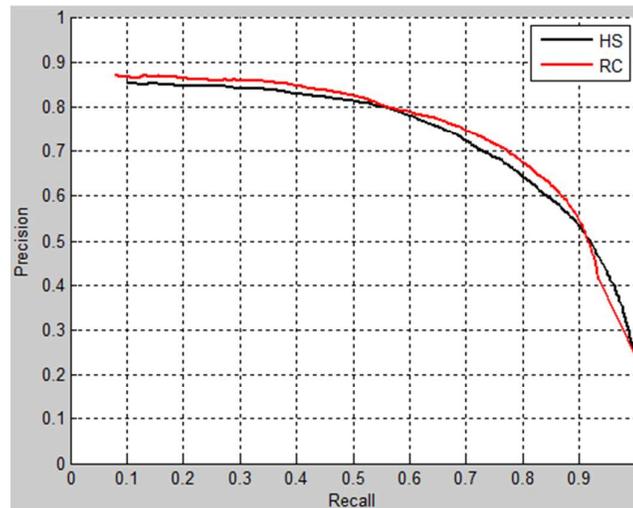


Figure 1: Precision recall curve of our method (RC) and the HC method (Hierarchical Saliency Detection, CVPR 2013).

Comment 2. Is there some hard assumption on how many salient components exist in the image? Moreover, in all the segmentation examples it seems that the method identifies exactly one salient component? Is that enforced explicitly? If so -- how? If not -- what prevents the results from containing two or more spatially disconnected segments?

Response: As described in the introduction, we focus on bottom-up data driven salient object detection using image contrast under the assumption that a salient object exists in an image (similar to [1]). If multiple objects occur in an image, the saliency maps could still highlight them (see Fig. 17). In our current SaliencyCut implementation, in case of ambiguity, we take the largest connected region as the candidate for the most salient object. One could change this to the top k connected regions.

Comment 3. Section 5.1 got me confused by the several thresholds. ... Please clarify.

Response: There is only one threshold in Section 5.1. We have rephrased the text to remove the confusion.

Comment 4. How the annotations of the 10K images were carried out?

Response: The ground truth binary segmentation for each of the 10K images is marked by *a single user*. These images are derived from MSRA benchmarks, where a salient object exists in an image. These images are carefully selected from a much larger image pool so that each image contains a clear, unambiguous object of interest (such consistency is confirmed by 3-9 users in terms of bounding boxes). During the pixel accuracy annotation, these bounding boxes are used to guide the segmentation, thus the ambiguous of salient region is unnoticeable.

Comment 5. Related work, 'Unsupervised Joint Object Discovery and Segmentation in Internet Images, CVPR13' is worth mentioning.

Response: Using our saliency measure they showed that saliency-based segmentation produces state-of-the-art results on several co-segmentation benchmarks (i.e. without using co-segmentation) – this is very encouraging. We have added the relative discussions and highlighted it in our project page.

Comment 6. Comments on exposition.

Response: Thanks for the comments. We have incorporated them.

Response to Reviewer 2 comments

Comment 1. The method is simple and provides very good results for salient object segmentation, and therefore very useful for many applications. ..., the paper is cited a lot and the software has many downloads, indicating that CVPR was a very good stage ... I think that the additions presented in this version ... are too minor to justify a PAMI publication.... These three changes do not significantly improve the model's performance.

Response: The additions, although quite simple, did improve the models performance significantly. In most part of the precision recall curve, the precision in the PAMI version is about 5% better than the CVPR version. This is a very significant improvement considering the precision in CVPR version is about 90%. Many saliency detection papers compared with our CVPR paper and claimed their new method gets better results (typically with more computational cost). Our results, initially submitted to PAMI at 2011, outperforms most recent papers like at CVPR 2012, ICCV 2012, CVPR 2013 salient object detection methods. See also response to the first question of review 1.

Comment 2. Other minor comments.

Response: Thanks for the detailed comments. We revised the paper accordingly.

Salient Object Detection and Segmentation

Ming-Ming Cheng, Niloy J. Mitra, Xiaolei Huang, Philip H. S. Torr, and Shi-Min Hu

Abstract—Automatic estimation of salient object regions across images, without any prior assumption or knowledge of the contents of the corresponding scenes, enhances many computer vision and computer graphics applications. We introduce a regional contrast based salient object detection algorithm, which simultaneously evaluates global contrast differences and spatial weighted coherence scores. The proposed algorithm is simple, efficient, naturally multi-scale, and produces full-resolution, high-quality saliency maps. These saliency maps are further used to initialize a novel iterative version of GrabCut, namely SaliencyCut, for high quality unsupervised salient object segmentation. We extensively evaluated our algorithm using traditional salient object detection datasets, as well as a more challenging Internet image dataset. Our experimental results demonstrate that our algorithm consistently outperforms 15 existing salient object detection and segmentation methods, yielding higher precision and better recall rates. We also show that our algorithm can be used to efficiently extract salient object masks from Internet images, enabling effective sketch-based image retrieval (SBIR) via simple shape comparisons. Despite such noisy internet images, where the saliency regions are ambiguous, our saliency guided image retrieval achieves a superior retrieval rate compared with state-of-the-art SBIR methods, and additionally provides important target object region information.

Index Terms—Salient object detection, visual attention, saliency map, unsupervised segmentation, image retrieval.

1 INTRODUCTION

WE, as humans, are experts at quickly and accurately identifying the most visually noticeable foreground object in the scene, known as salient objects, and adaptively focus our attention on such perceived important regions. In contrast, computationally identifying such salient object regions [1], [2], that match the human annotators' behaviour when they have been asked to pick a salient object in an image, is very challenging. Being able to automatically, efficiently, and accurately estimate salient object regions, however, is highly desirable given the immediate ability to characterise the spatial support for feature extraction, isolate the object from potentially confusing background, and preferentially allocate finite computational resources for subsequent image processing.

While essentially solving a segmentation problem, salient object detection approaches segment only the salient foreground object from the background, rather than partition an image into regions of coherent properties as in general segmentation algorithms [2]. Salient object detection models also differ from eye fixation prediction models that predict a few fixation points in an image rather than uniformly highlighting the entire salient object region [2]. In practice, salient object detection methods are commonly used as a first step of many graphics/vision applications including object-of-interest image segmentation [3], object recognition [4], adaptive compression of images [5], content-aware image editing [6], [7], image retrieval [8]–[10], etc.

- M.M. Cheng and P.H.S. Torr are with the University of Oxford. The major part of this work was finished in Tsinghua University. N.J. Mitra is with UCL/KAUST. X. Huang is with Lehigh University. Shi-Min Hu (corresponding author) is with TNList, Tsinghua University.



Fig. 1. Given an input image (top), a global contrast analysis is used to compute high resolution saliency map (middle), which can be used to produce unsupervised segmentation mask (bottom) for object interest.

Although extraction of salient objects in a scene is related to accurate image segmentation and object retrieval, interestingly reliable saliency estimation is often feasible at the image-level *without* any actual scene understanding. This is feasible since often saliency, as widely believed, is bottom-up. Such a hypothesis is favored by an evolutionary incentive to parallel process large volumes of low-level image cues, without requiring any computationally expensive global coupling. Hence, saliency often originates from visual uniqueness, unpredictability, rarity, or surprise, and is often attributed to variations in image attributes like color, gradient, edges, and boundaries.

Visual saliency, being tightly related to our perception and processing of visual stimuli, is investigated across many disciplines including cognitive psychology [11], [12], neurobiology [13], [14], and computer vision [15]–

[17]. Based on our observed reaction times and estimated signal transmission times along biological pathways, human attention theories hypothesize that the human vision system processes only parts of an image in detail, while leaving the rest nearly unprocessed. Early work by Treisman and Gelade [18], Koch and Ullman [19], and subsequent attention theories proposed by Itti, Wolfe and others, suggest two stages of visual attention: a fast, pre-attentive, bottom-up, data driven saliency extraction; followed by a slower, task dependent, top-down, goal driven saliency extraction.

We focus on bottom-up data driven salient object detection using image contrast (see Fig. 1)¹ under the assumption that a salient object exists in an image [1]. The proposed method is simple, fast, and produces high quality results on benchmark datasets. Motivated by the popular belief that human cortical cells may be *hard wired* to preferentially respond to high contrast stimulus in their receptive fields [35], we propose contrast analysis for extracting high-resolution, full-field saliency maps based on the following considerations:

- A global contrast based method, which separates a large-scale object from its surroundings, is desirable over local contrast based methods producing high saliency values at or near object boundaries. Global considerations enable assignment of comparable saliency values across similar image regions, and can uniformly highlight entire objects.
- Saliency of a region mainly depends on its contrast with respect to its nearby regions, while contrasts to distant regions are less significant (see also [36]).
- In man-made photographs, object are often concentrated towards the inner regions of the images, away from image boundaries (see [33]).
- Saliency maps should be fast, accurate, have low memory footprints, and easy to generate to allow processing of large image collections, and facilitate efficient image classification and retrieval.

We propose a *histogram-based contrast method* (HC) to measure saliency. HC-maps assign pixel-wise saliency values based simply on color separation from all other image pixels to produce full resolution saliency maps. We use a histogram-based approach for efficient processing, while employing a smoothing procedure to reduce quantization artifacts.

As an improvement over HC-maps, we incorporate spatial relations to produce *region-based contrast* (RC) maps where we first segment the input image into regions, and then assign saliency values to them. The saliency value of a region is then calculated using a global contrast score, measured by the region's contrast and spatial distances to other regions in the image. Note that this approach better acknowledges the relation between image segmentation and saliency determination.

Segmenting regions of interest in still images is of great practical importance in many computer vision

and graphics applications. Researchers have devoted significant efforts to minimize user interaction during this process. GrabCut [37], which iteratively optimizes the energy function and considers both texture and edge information, has successfully simplified the user interaction to simply dragging a rectangle around the desired object. We propose SaliencyCut, an improved iterative version of GrabCut, and combine it with our saliency detection method to achieve superior performance compared to state-of-the-art unsupervised salient object extraction methods.

In order to evaluate the proposed algorithms and compare with state-of-the-art alternatives, we build a database with 10,000 pixel-accurate human-labeled ground truth images (see also Sec. 6.1.1), which is an order of magnitude bigger than previous largest public available dataset of its kind [31]. We have extensively evaluated our methods on this dataset, and compared our methods with 15 state-of-the-art saliency methods as well as with manually created ground truth annotations². The experiments show significant improvements over previous methods both in terms of precision and recall rates. Overall, compared with HC-maps, RC-maps produce better precision and recall rates, but at the cost of increased computational overhead. In our extensive empirical evaluations, we observe that the unsupervised segmentation results produced by our SaliencyCut method are, in most cases, are comparable to the manually annotated ground truths. We also demonstrate applications of the extracted saliency maps to segmentation and sketch-based image retrieval.

2 RELATED WORK

Our work belongs to the active research field of visual attention modeling, for which a comprehensive discussion is beyond the scope of this paper. We refer readers to recent survey papers for a detailed discussion of 65 models [38], as well as quantitative analysis of different methods in the two major research directions: fixation prediction [39], [40] and salient object detection [2].

We focus on relevant literature targeting pre-attentive bottom-up saliency region detection, which are biologically motivated, or purely computational, or involve both aspects. Such methods utilize low-level processing to determine the contrast of image regions to their surroundings, and use feature attributes such as intensity, color, and edges [31]. We broadly classify the algorithms into local and global schemes. Note that the classification is not strict as some of the research efforts can be listed under both categories.

Local contrast based methods investigate the rarity of image regions with respect to (small) local neighborhoods. Based on the highly influential biologically inspired *early representation* model introduced by Koch and

2. The C++ source code, benchmark dataset, corresponding results, links to competing methods and highlighted applications, are available via our project page: <http://mmcheng.net/SalObj/>.

1. A preliminary version of this work appeared at CVPR [34].

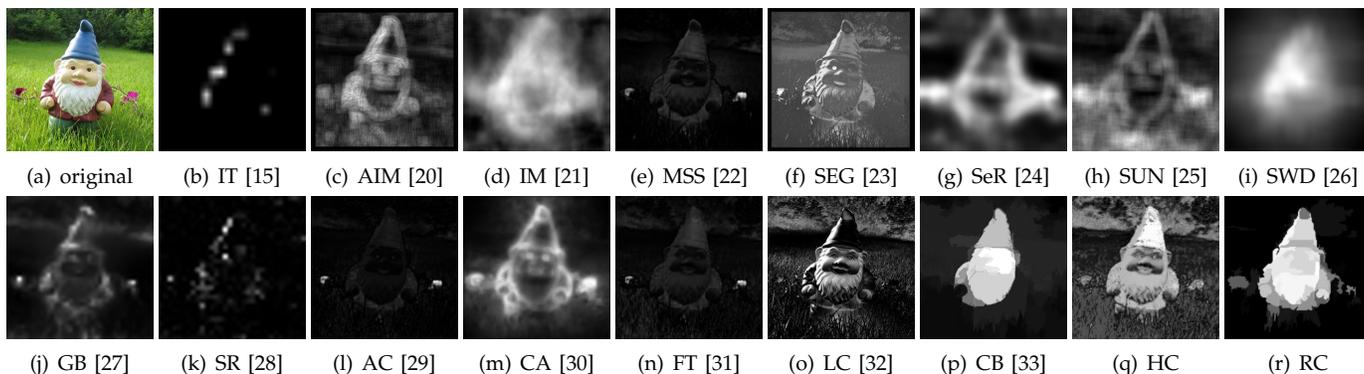


Fig. 2. Saliency maps computed by different state-of-the-art methods (b-p), and with our proposed HC (q) and RC methods (r). Most results highlight edges, or are of low resolution. See also Fig. 9 and our project webpage.

Ullman [19], Itti et al. [15] define image saliency using central-surrounded differences across multi-scale image features. Ma and Zhang [41] propose an alternate local contrast analysis for generating saliency maps, which is then extended using a fuzzy growth model. Harel et al. [42] propose a bottom-up visual saliency model to normalize the feature maps of Itti et al. to highlight conspicuous parts and permit combination with other importance maps. The model is simple, biologically plausible, and easy to parallelize. Liu et al. [1] find multi-scale contrast by linearly combining contrast in a Gaussian image pyramid. More recently, Goferman et al. [30] simultaneously model local low-level clues, global considerations, visual organization rules, and high-level features to highlight salient objects along with their contexts. Such methods using local contrast tend to produce higher saliency values near edges instead of uniformly highlighting salient objects (see Fig. 2). Note that Reinagel et al. [36] observe that humans tend to focus attention in image regions with high spatial contrast and local variance in pixel correlation.

Global contrast based methods evaluate saliency of an image region using its contrast with respect to the entire image. Zhai and Shah [32] define pixel-level saliency based on a pixel’s contrast to all other pixels. However, for efficiency they use only luminance information, thus ignoring distinctiveness clues in other channels. Achanta et al. [31] propose a frequency tuned method that directly defines pixel saliency using a pixel’s color difference from the average image color. The elegant approach, however, only considers first order average color, which can be insufficient to analyze complex variations common in natural images. In Figures 9 and 11, we show qualitative and quantitative weaknesses of such approaches. Furthermore, these methods ignore spatial relationships across image parts, which can be critical for reliable and coherent saliency detection (see Sec. 6).

Saliency maps are widely employed for unsupervised object segmentation: Ma and Zhang [41] find rectangular salient regions by fuzzy region growing on their saliency maps. Ko and Nam [43] select salient regions using a support vector machine trained on image segment

features, and then cluster these regions to extract salient objects. Han et al. [44] model color, texture, and edge features in a Markov random field framework to grow salient object regions from seed values in the saliency maps. More recently, Achanta et al. [31] average saliency values within image segments produced by mean-shift segmentation, and then find salient objects via adaptive thresholding. We propose a different approach that extends GrabCut [37] method and automatically initialize it using our saliency detection methods. Experiments on our 10,000 images dataset (see Sec. 6.1.1) demonstrate the significant advantages of our method compared to other state-of-the-art methods.

Subsequent to our preliminary results [34], Jiang et al. [33] propose a comparable method also making use of region level contrast to model image saliency. In the segmentation step, their method also expands and shrinks the initial trimap and iteratively applies graphcut and histogram appearance model. Since GrabCut is an iterative process of using graphcut and GMM appearance mode, the two segmentation methods share a strong similarity. Compared to the CB method [33], experimental results show that our RC salient object region detection and segmentation is more accurate (Fig. 11(a)(c)), 20× faster (Fig. 7), and more robust to center-bias (Fig. 11(b)).

3 HISTOGRAM BASED CONTRAST

Our biological vision system is highly sensitive to contrast in visual signal. Based on this observation, we propose a histogram-based contrast (HC) method to define saliency values for image pixels using color statistics of the input image. Specifically, the saliency of a pixel is defined using its color contrast to all other pixels in the image, i.e., the saliency value of a pixel I_k in image I is,

$$S(I_k) = \sum_{\forall I_i \in I} D(I_k, I_i), \quad (1)$$

where $D(I_k, I_i)$ is the color distance metric between pixels I_k and I_i in the $L^*a^*b^*$ space for perceptual accuracy. (1) can be expanded by pixel order as,

$$S(I_k) = D(I_k, I_1) + D(I_k, I_2) + \dots + D(I_k, I_N), \quad (2)$$

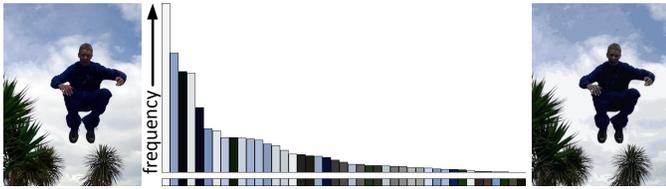


Fig. 3. Given an input image (left), we compute its color histogram (middle). Corresponding histogram bin colors are shown in the lower bar. The quantized image (right) uses only 43 histogram bin colors and still retains sufficient visual quality for saliency detection.

where N is the number of pixels in image I . It is easy to see that pixels with the same color have the same saliency under this definition, since the measure is oblivious to spatial relations. Thus, rearranging (2) such that the terms with the same color value c_j are grouped together, we get saliency value for each color as,

$$S(I_k) = S(c_l) = \sum_{j=1}^n f_j D(c_l, c_j), \quad (3)$$

where c_l is the color value of pixel I_k , n is the number of distinct pixel colors, and f_j is the probability of pixel color c_j in image I .

3.1 Histogram based speed up

Naively evaluating the saliency value for each image pixel using (1) takes $O(N^2)$ time, which is computationally too expensive even for medium sized images. The equivalent representation in (3), however, takes $O(N) + O(n^2)$ time, implying that computational efficiency can be improved to $O(N)$ if $O(n^2) \leq O(N)$. Thus, the key to speed up is to reduce the number of pixel colors in the image. However, the true-color space contains 256^3 possible colors, which is typically larger than the number of image pixels.

Zhai and Shah [32] reduce the number of colors, n , by only using luminance. In this way, $n^2 = 256^2$ (typically $256^2 \ll N$). The method, however, ignores distinctiveness of color information. In this work, we use the full color space instead of luminance only. To reduce the number of colors needed to consider, we first quantize each color channel to have 12 different values, which reduces the number of colors to $12^3 = 1728$. Considering that color in a natural image typically covers only a small portion of the full color space, we further reduce the number of colors by ignoring less frequently occurring colors. By choosing more frequently occurring colors and ensuring these colors cover the colors of more than 95% of the image pixels, we typically are left with around $n = 85$ colors (see Sec. 6 for experimental details). The colors of the remaining pixels, which comprise fewer than 5% of the image pixels, are replaced by the closest colors in the histogram. A typical example of such quantization is shown in Fig. 3. Note that due to efficiency considerations we select the simple histogram based quantization instead of optimizing for an image specific color palette.

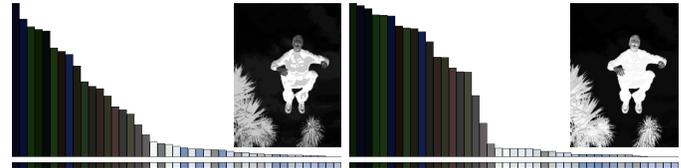


Fig. 4. Saliency of each color before (left) and after (right) color space smoothing. Corresponding saliency maps are shown in the respective insets.

3.2 Color space smoothing

Although we can efficiently compute color contrast by building a compact color histogram using color quantization and choosing more frequent colors, the quantization itself may introduce artifacts. Some similar colors may be quantized to different values. In order to reduce noisy saliency results caused by such randomness, we use a smoothing procedure to refine the saliency value for each color. We replace the saliency value of each color by the weighted average of the saliency values of similar colors. This is actually a smoothing process in the color feature space. We choose $m = n/4$ nearest colors to refine the saliency value of color c by,

$$S'(c) = \frac{1}{(m-1)T} \sum_{i=1}^m (T - D(c, c_i)) S(c_i) \quad (4)$$

where $T = \sum_{i=1}^m D(c, c_i)$ is the sum of distances between color c and its m nearest neighbors c_i , and the normalization factor comes from $\sum_{i=1}^m (T - D(c, c_i)) = (m-1)T$. Note that we use a linearly-varying smoothing weight $(T - D(c, c_i))$ to assign larger weights to colors closer to c in the color feature space. In our experiments, we found that such linearly-varying weights are better than Gaussian weights, which fall off too sharply. Fig. 4 shows the typical effect of color space smoothing with the corresponding histograms sorted by decreasing saliency values. Note that similar histogram bins are closer to each other after such smoothing, indicating that similar colors is more likely to be assigned similar saliency values, thus reducing quantization artifacts (see Fig. 11).

3.3 Implementation details

To quantize the color space into 12^3 different colors, we uniformly divide each color channel into 12 levels. While the quantization of colors is performed in the RGB color space, we measure color differences in the $L^*a^*b^*$ color space given its perceptual accuracy. We do not, however, perform quantization directly in the $L^*a^*b^*$ color space since not all colors in the range $L^* \in [0, 100]$, and $a^*, b^* \in [-127, 127]$ necessarily correspond to real colors. Experimentally we observed worse quantization artifacts using direct $L^*a^*b^*$ color space quantization. Best results were obtained by quantization in the RGB space while measuring distance in the $L^*a^*b^*$ color space, as opposed to performing both quantization and distance calculation in a single color space, either RGB or $L^*a^*b^*$.

4 REGION BASED CONTRAST

Humans pay more attention to image regions with high contrast to their surroundings [45]. Besides contrast, spatial relationships are important in human attention. High contrast to ones surrounding regions is usually stronger evidence for saliency of a region than comparable contrast to far-away regions. Since directly introducing spatial relationships when computing pixel-level contrast is computationally expensive, we introduce a contrast analysis method, *region contrast* (RC), so as to integrate spatial relationships into region-level contrast computation. In RC, we first segment the input image into regions, then compute color contrast at the region level, and finally define saliency for each region as the weighted sum of the region's contrasts to all other regions in the image. The weights are set according to the spatial distances with farther regions being assigned smaller weights.

4.1 Region contrast by histogram comparison

We first segment the input image into regions using a graph-based image segmentation method [46]. Then we build the color histogram for each region as in Sec. 3. For a region r_k , we compute its saliency value by measuring its color contrast to all other regions in the image,

$$S(r_k) = \sum_{r_i \neq r_k} w(r_i) D_r(r_k, r_i), \quad (5)$$

where $w(r_i)$ is the weight of region r_i and $D_r(\cdot, \cdot)$ is the color distance metric between the two regions. We weight the distances by the number of pixels in r_i as $w(r_i)$ to emphasize color contrast to bigger regions. The color distance between two regions r_1 and r_2 is,

$$D_r(r_1, r_2) = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} f(c_{1,i}) f(c_{2,j}) D(c_{1,i}, c_{2,j}) \quad (6)$$

where $f(c_{k,i})$ is the probability of the i -th color $c_{k,i}$ among all n_k colors in the k -th region r_k , $k = \{1, 2\}$. Note that we use the probability of a color in the probability density function (i.e., normalized color histogram) of the region as the weight for this color to further emphasize the color differences between dominant colors.

Storing and calculating the regular matrix format histogram for each region is inefficient since each region typically contains a small number of colors in the color histogram of the whole image. Instead, we use a sparse histogram representation for efficient computation.

4.2 Spatially weighted region contrast

We further incorporate spatial information by introducing a spatial weighting term in (5) to increase the effects of closer regions and decrease the effects of farther regions. Specifically, for any region r_k , the spatially weighted region contrast based saliency is:

$$S(r_k) = w_s(r_k) \sum_{r_i \neq r_k} e^{-\frac{D_s(r_k, r_i)}{\sigma_s^2}} w(r_i) D_r(r_k, r_i) \quad (7)$$

where $D_s(r_k, r_i)$ is the spatial distance between regions r_k and r_i , σ_s controls the strength of spatial distance weighting, $w(r_i)$ is the weight of region r_i defined by the number of pixels in r_i , and $w_s(r_k)$ is a spatial prior weighting term similar to center bias (CB [33]). We use $w_s(r_k) = \exp(-9d_k^2)$, where d_k is the average distance between pixels in region r_k and the center of the image, with pixel coordinates normalized to $[0, 1]$. Thus, $w_s(r_k)$ gives a high value if region r_k is close to the center of the image and it gives a low value if the region is a border region away from the center. For σ_s , larger values of σ_s reduce the effect of spatial weighting so that contrast to farther regions would contribute more to the saliency of the current region. The spatial distance between two regions is defined as the Euclidean distance between their centroids. In our implementation, we use $\sigma_s^2 = 0.4$ with pixel coordinates normalized to $[0, 1]$.

4.3 Further improvement of RC saliency maps

We further refine our RC saliency maps in two steps. First, we use the spatial prior to explicitly estimate the non-salient (background) region. Second, we apply the color space smoothing as described in Sec. 3.2.

We observe that regions with long borders overlapping with image borders are typically non-salient background regions, which we call border regions. We incorporate them as another spatial prior ($w_s(\cdot)$ in (7)) to detect non-salient regions. In our implementation, we normalize the number of pixels located in the 15 pixel-wide image-border area by the region size, and consider regions with this value higher than a threshold to be border regions. In practice, this hard constraint improves both the saliency maps as well as the convergence speed of SaliencyCut (Sec. 5) by improving the initial condition. Our border region estimation aims at high precision, rather than high recall. A strict fixed threshold, which on average corresponds to 2% miss alarm rate in our dataset, is chosen to detect border regions.

In order to uniformly highlight the entire saliency region of the image, we get the average saliency of each color in the color histogram and adopt the color space smoothing (Sec. 3.2) to improve our RC saliency map. After smoothing, some border region pixels may get non-zero saliency values. We reset the saliency of border region to zero and re-estimate the saliency of each region as the average saliency value of its corresponding pixels. Since initial RC maps are typically more uniformly highlighted compared to HC saliency maps without color space smoothing, we typically choose smaller number of nearest colors ($m = n/10$ in this part). Fig. 5(f) demonstrates such an example. The jumping man region is more uniformly highlighted compared to Fig. 5(e).

5 SALIENCYCUT: AUTOMATIC SALIENT REGION EXTRACTION

In a highly influential work, GrabCut [37] made critical changes to the graphcut formulation to allow processing

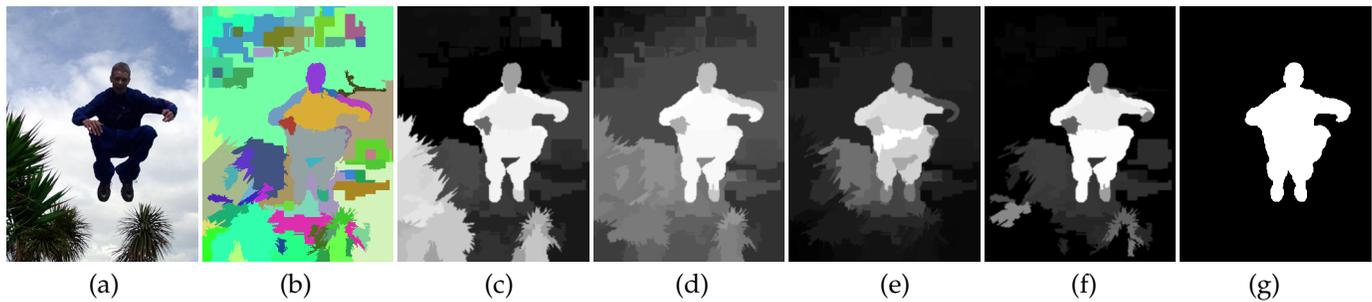


Fig. 5. Region based contrast computation: (a) input image, (b) image regions generated by Felzenszwalb and Huttenlocher's segmentation method [46], (c) region contrast without distance weighting and spatial prior ((5)), (d) region contrast with distance weighting, (e) region contrast further considering spatial prior ((7)), (f) region contrast after improvement by border region estimation and color space smoothing, (g) using our SaliencyCut (Sec. 5), we get a high quality cut that is comparable to human labeled ground truth.

of noisy initialization. This enabled users to roughly annotate (e.g., using a rectangle) a region of interest, and then use GrabCut to extract a precise image mask. Using our estimated saliency masks, we remove even the need for user annotated rectangular regions. In this section, we introduce *SaliencyCut*, which uses the computed saliency map to assist in automatic salient object segmentation. This immediately enables automatic analysis of large internet image repositories. Specifically, we make two enhancements to GrabCut [37]: “iterative refine” and “adaptive fitting”, which together handle considerably more noisy initializations. Thanks to the robustness of the new approach, we are able to automatically initialize the segmentation according to the detected saliency map.

5.1 Algorithm initialization

Instead of manually selecting a rectangular region to initialize the process, as in classical GrabCut, we automatically initialize using a segmentation obtained by binarizing the saliency map using a fixed threshold T_b . Similar to GrabCut, we use incomplete trimap for the initialization. For image pixels with saliency value bigger than T_b , the largest connected region is considered as initial candidate region of the most dominate salient object. This candidate region is labeled as unknown part of the trimap, while other regions are labeled as background. Note that we do not initialize any hard foreground labeling. These unknown regions are initially used to train foreground color models thus helps the algorithm to identify the foreground pixels.

Since the initial background regions are retained while other regions may be changed during the GrabCut optimization, we give preference to confident background labels in the trimaps. Thus we initialize the GrabCut algorithm using threshold given high recall of potential foreground region and let the iterative optimization process to increase its precision. In our experiments, the threshold is chosen empirically to be the threshold that gives 95% recall rate in our fixed thresholding experiments (see Sec. 6.2). When initialized using RC saliency maps, we use $T_b = 70$ with saliency values normalized to $[0, 255]$.

5.2 Segmentation by iterative fitting

Once initialized, we iteratively run GrabCut [37] to improve the SaliencyCut result (maximum of 4 iterations in our experiments). After each iteration, we use dilation and erosion operations on the current segmentation result to get a new trimap for the next GrabCut iteration. As shown in Fig. 6(c, d), the region outside the dilated region is set to background, the region inside the eroded region is set to foreground, and the remaining areas are set to unknown in the trimap. GrabCut, which by itself is an iterative process using Gaussian mixture models and graph-cut [47], helps to refine salient object regions at each step.

Different from one-pass GrabCut or the even simpler graph cut based saliency segmentation [48], the new scheme in SaliencyCut *iteratively refines* the initial salient regions. Such an iterative design is important to handle noisy initializations supplied by the saliency detection algorithm rather than human annotations. In case of incorrect initialization as shown in flower example in Fig. 6 (b), the initial background region incorrectly contains foreground object(s). Although we can still get a segmentation result containing many parts of the flower using GrabCut, the remaining flowers in the initial background region would never be correctly extracted using GrabCut since the background gets a hard labeling. One may consider relaxing the hard constrain of GrabCut to solve this problem. However, experimental results show this would make the method not stable, often producing results containing all foreground or all background.

We iteratively refine the initial segmentation and adaptively change the initial condition to fit with newly segmented salient region. The *adaptive fitting* is based on an important observation: regions closer to an initial salient object region are more likely to be part of that salient object than far-away regions. Thus, our new initialization enables GrabCut to include nearby salient regions, and exclude non-salient regions according to color feature dissimilarity. After each GrabCut iteration, SaliencyCut incorporates the constraints given by the newly obtained trimap, and train a better appearance model according to previous results.

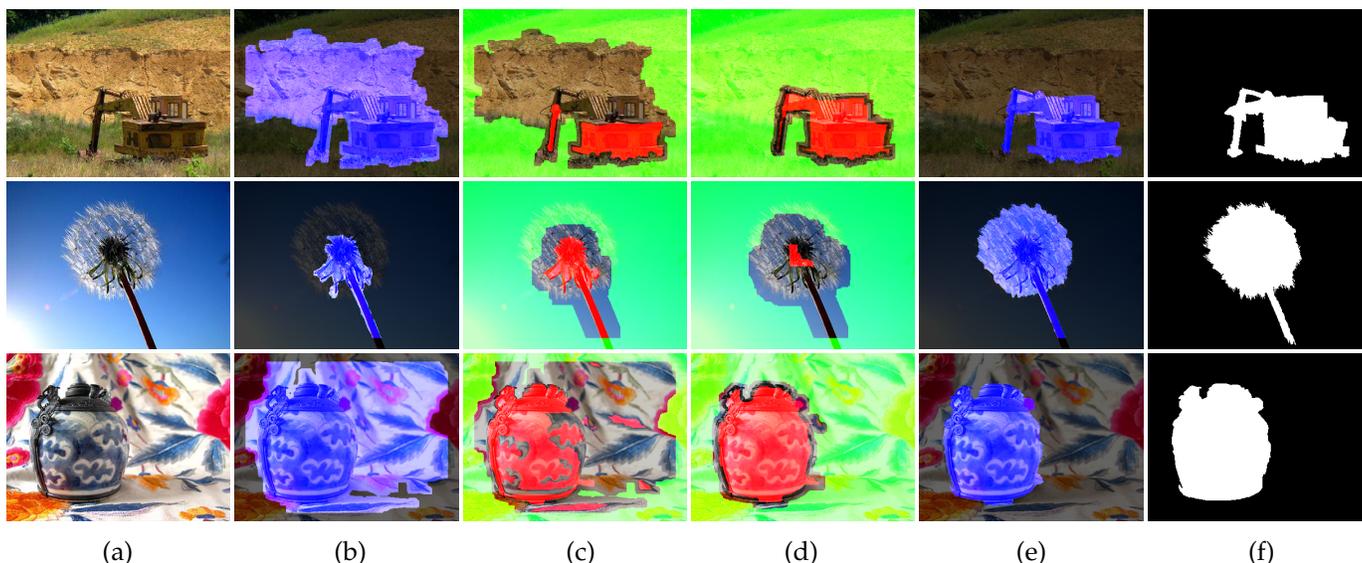


Fig. 6. Demonstration of SaliencyCut: (a) original image, (b) initial segmentation by fixed thresholding the saliency map, (c) trimap after first iteration, (d) trimap after second iteration, (e) final segmentation, and (f) manually labeled ground truth. In the segmented images (e), blue is foreground, gray is background, while in the trimaps (b–d), foreground is red, background is green, and unknown regions are left unchanged.

Fig. 6 shows three examples. In the flower example, SaliencyCut successfully expanded the initial salient regions (obtained directly from the saliency map) and converged to an accurate segmentation result. In the excavator and teapot examples, unwanted regions are correctly excluded during GrabCut iterations. The intermediate steps show how SaliencyCut successfully extracted the object regions of interest in these challenging examples. A comprehensive quantitative evaluation of different saliency segmentation methods is presented in Sec. 6.3.

6 EXPERIMENTAL COMPARISONS

In this work, we extensively evaluated our saliency detection method on three different types of benchmark datasets, and compared it against 15 alternate methods — SR [28], IT [15], IM [21], SUN [25], AC [29], SeR [24], AIM [20], GB [27], LC [32], CA [30], FT [31], SWD [26], SEG [23], MSS [22], LP [49] and CB [33], respectively. Following [31], we selected these methods according to: number of citations (IT, SR, SUN, AIM and FT), recency (SeR, MSS, SEG, IM, CA and SWD), variety (IT is biologically-motivated, LC is purely computational, GB and LP are hybrid, SR works in the frequency domain, AC and FT output full resolution saliency maps), and being related to our approach (LC and CB).

Fig. 7 compares the average time taken by each method on a Dual Core 2.6 GHz machine with 2GB RAM. Our algorithms, HC and RC, are implemented in C++. For the other methods namely IT, AIM, IM, MSS, SEG, SeR, SUN, GB, SR, AC, CA, FT and CB, we used the authors’ implementations, while for LC, we implemented the algorithm in C++ since we failed to obtain the authors’ implementation. For typical natural images, our HC method runs in $O(N)$, which is sufficient

for real-time applications. In contrast, our RC variant is slower as it requires image segmentation [46], but produces superior quality saliency maps.

The true effectiveness of a saliency detection method depends on the applications [31]. We evaluated our method on three core computer vision and graphics applications, namely salient region segmentation by fixed thresholding, object of interest image segmentation, and sketch based image retrieval (SBIR).

6.1 Benchmark datasets for saliency detection

6.1.1 Images with unambiguous salient objects

Similar to existing salient object region detection methods [1], [22], [31], [33], we first evaluate our methods on images with unambiguous salient object. The largest dataset of this kind is provided by Liu et al. [1]. This dataset contains 20,000+ images (mostly at 400×300 resolution), with bounding box labeling by 3-9 users. These images are selected from an initial set of 130,099 images,

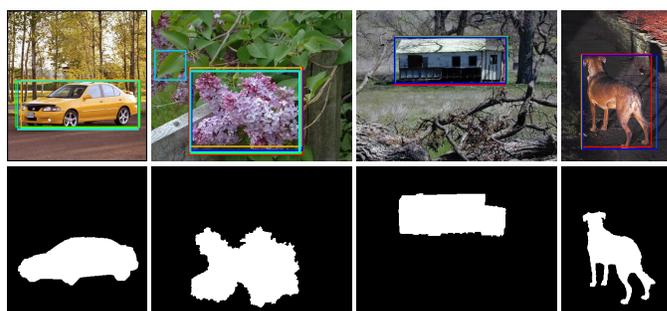


Fig. 8. Ground truth examples: (first row) original images with ground truth rectangles from [1], (second row) our ground truth, which have more precisely marked important regions at pixel level accuracy.

Method	IT [15]	AIM [20]	IM [21]	MSS [22]	SEG [23]	SeR [24]	SUN [25]	SWD [26]	CB [33]
Time (s)	0.246	4.288	0.991	0.106	4.921	1.019	1.116	0.100	5.568
Code Type	Matlab	Matlab	Matlab	C++	Matlab	Matlab	Matlab	Matlab	Matlab & C
Method	GB [27]	SR [28]	FT [31]	AC [29]	CA [30]	LC [32]	HC	RC	
Time (s)	1.614	0.064	0.102	0.109	53.1	0.018	0.019	0.254	
Code Type	Matlab	Matlab	C++	Matlab	Matlab	C++	C++	C++	

Fig. 7. Average time taken to compute a saliency map for images in the THUS10000 database (most have resolution 400×300). We use parallel computing environment for all Matlab functions for efficient computation.

such that each image contains a clear, unambiguous object of interest. Since objects can still be recognized at low resolution, the dataset has limited scale and location variations of salient objects, i.e., implicitly the images have scale and location priors (Flickr like).

Although an invaluable recourse to evaluate saliency detection algorithms, the database with the marked bounding boxes, however, is often too coarse for fine grained evaluation as observed by Wang and Li [50], and Achanta et al. [31]. In order to do more extensive and accurate evaluation, we randomly selected 10,000 images with consistent bounding box labeling in database provided by Liu et al. [1] and the consistent measure is the same as choosing image dataset B in their paper. As shown in Fig. 8, we accurately marked pixels in salient object regions. We call this dataset THUS10000 because it contains 10,000 images with *pixel-level* saliency labeling (publicly available on our project page). Our dataset is 10 times bigger than what was previously the largest public available database of its kind [31]. In our experiments, we find that saliency detection methods using pixel level contrast (FT, HC, LC, MSS) do not scale well on this larger benchmark (see Fig. 11(a)), suggesting the importance of region level analysis.

6.1.2 Randomly-selected internet images

While state-of-the-art methods consistently produce excellent results when evaluated using the traditional benchmark dataset [31] (see Fig. 11(c)), ordinary users often report less satisfactory experiences when using their own images. This encourage us to think about two questions: ‘How would these methods deal with random internet images?’ and ‘When can we trust the results of these methods?’ To better explore these issues, we evaluated salient object segmentation methods on a dataset with randomly-selected internet images [51]. This benchmark dataset, namely THUR15000 [51], contains about 3000 images downloaded from Flickr for each of the 5 keywords: “butterfly,” “coffee mug,” “dog jump,” “giraffe,” and “plane.” Salient regions in THUR15000 images are marked at pixel accuracy. Note that not every image in the THUR15000 dataset contains a salient region label, as some images do not have any salient object region. Besides saliency detection, this dataset can also be used to evaluate the performance of sketch based image retrieval (SBIR).

6.1.3 Human fixation dataset

While our algorithm targets salient object detection, it is also interesting to test its performance on human fixation prediction benchmarks. We use the most widely adopted human fixation benchmark [49] for such evaluation.

6.2 Segmentation by fixed thresholding

The simplest way to get a binary segmentation of salient objects is to threshold the saliency map with a threshold $T_f \in [0, 255]$. To reliably compare how well various saliency detection methods highlight salient regions in images, we vary the threshold T_f from 0 to 255. Fig. 11(a) shows the resulting precision vs. recall curves. Typical qualitative comparison of saliency maps obtained by the various methods are presented in Fig. 2 and Fig. 9.

Unlike most other methods, both the CB method and our RC method use the center location prior of the man-made photographs. However, for a fair comparison, Fig. 11(b) shows comparisons while disabling such a location prior. Specifically, RC1 shows the effect disabling the center location weighting ((7)) of RC method, while RC2 shows the effect of further disabling border region estimation (Sec. 4.3). Other methods in Fig. 11(b) also improve when we use the same segmentation, as used in RC, to average saliency values within each segment and re-normalize to $[0, 255]$ by uniform scaling. Note that many of these methods aim to predict human eye movements rather than perform salient object segmentation, as is our focus.

The precision and recall curves clearly show that our RC method outperforms the other methods. We observe a significant loss in precision Fig. 11(b) for the CB method (which has best performance in the benchmark paper [2]) indicating that the method heavily relies on location prior. The extremities of the precision vs. recall curve are interesting: At maximum recall where $T_f = 0$, all pixels are retained as positives, i.e., considered to be

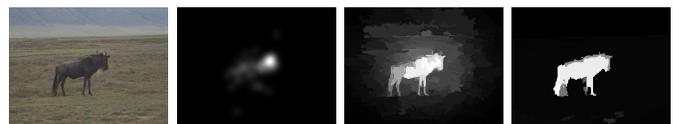
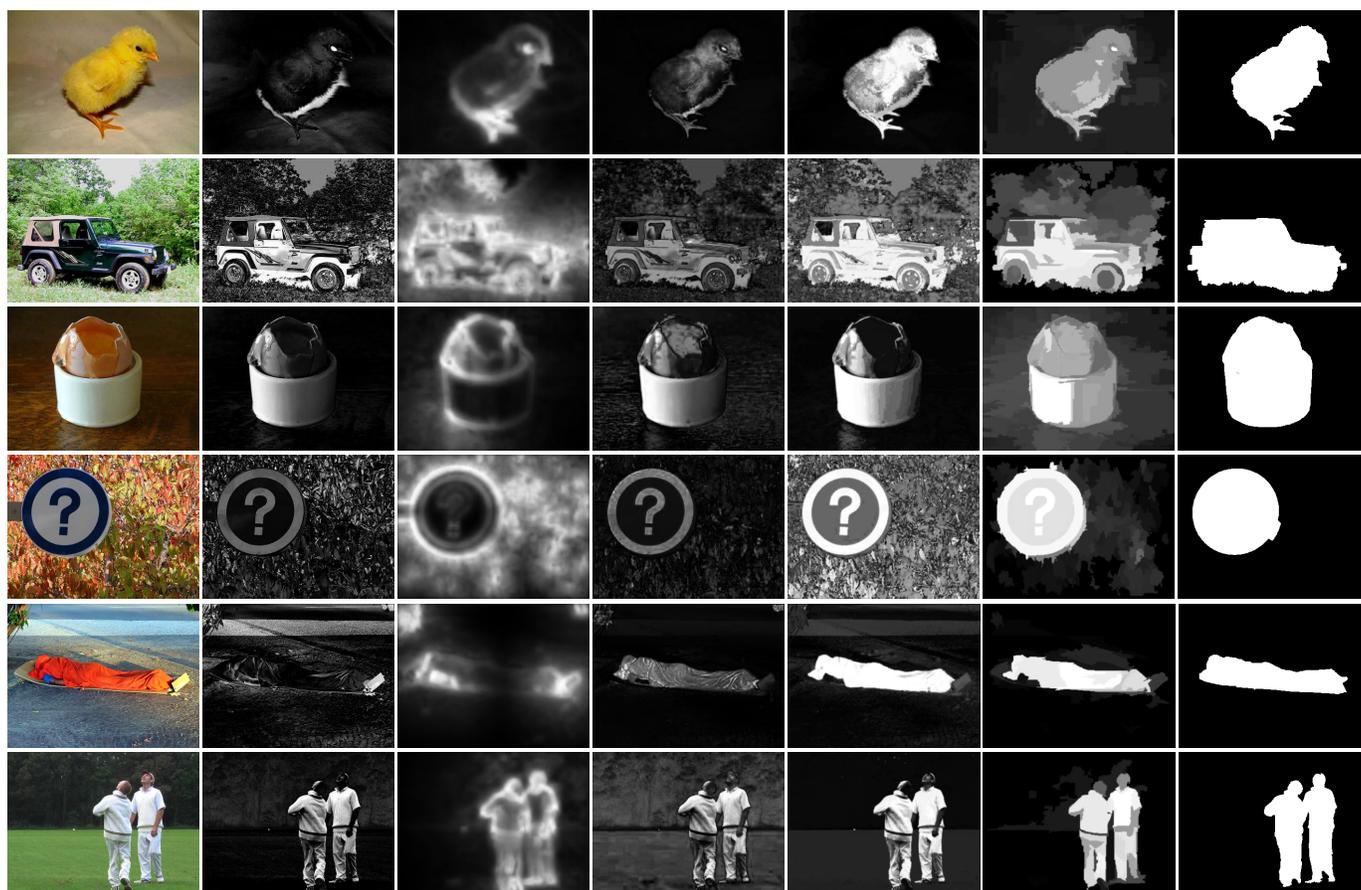


Fig. 10. From left to right, we show source image, ground truth eye fixation map by human observer, our RC result with the term encouraging similar appearance region receive similar saliency (Sec. 4.3) disabled, and result by our full RC method.



(a) original (b) LC (c) CA (d) FT (e) HC-maps (f) RC-maps (g) RCC

Fig. 9. Visual comparison of saliency maps. (a) original images, saliency maps produced using (b) Zhai and Shah [32], (c) Goferman et al. [30], (d) Achanta et al. [31], (e) our HC and (f) RC methods, and (g) SaliencyCut. Our methods generate uniformly highlighted salient regions (see our project webpage for all results on the full benchmark dataset).

foreground, so all the methods have the same precision and recall values; precision 0.22 and recall 1.0 at this point indicate that, on average, there are 22% image pixels belonging to the ground truth salient regions. At the other end, the minimum recall values of our RC method are higher than those of the other methods, because the saliency maps computed by our RC method are smoother and contain more pixels with the saliency value 255. Our HC method also has better precision and recall compared to methods with similar computational efficiency (SR, FT, and LC). After comparison of a large number of saliency detection models, Borji et al. [2] proposed a combined model and show that integration of the few best models (with the initial version of our method as one of them) outperforms all models. We believe that the combined model of [2] will further benefit from performance improvement due to our method.

As also observed in the survey papers [2], [38]–[40], center-bias naturally exists in human captured photos. Judd et al. [49] further found that a simple Gaussian blob performs better than many saliency detection methods when evaluated in famous eye fixation dataset. We experimentally find that such simple Gaussian blob, represented by ‘Gau’ in Fig. 11(a)(c), also performs better

Method	MSS [22]	CA [30]	LP [49]	Ours
ROC Area	0.683	0.844	0.849	0.830

Fig. 12. Comparison on human fixation dataset [49].

than many existing models for saliency region detection task. However, in the absence of explicit information, we prefer not to use such a strong prior that can potentially produce biased results, e.g., in automated imaging systems. When disabling the center bias term, our method still produces better results than other alternatives Fig. 11(b).

In the context of fixation prediction, the CA [30] and LP [49] methods report the best performance. Although it avoids the heavy learning for combining multi-saliency models and object detectors, the CA method still needs about 1 min to calculate a saliency map even for small images. Fig. 12 and Fig. 7 shows that our method, although initially designed for saliency region detection, has only slightly lower performance to state-of-the-art methods for predicting human fixation points, while being 200+ times more efficient. Readers can refer to [30], [40], [49] for more comparisons. Notice that the good performance of our RC method for predicting eye

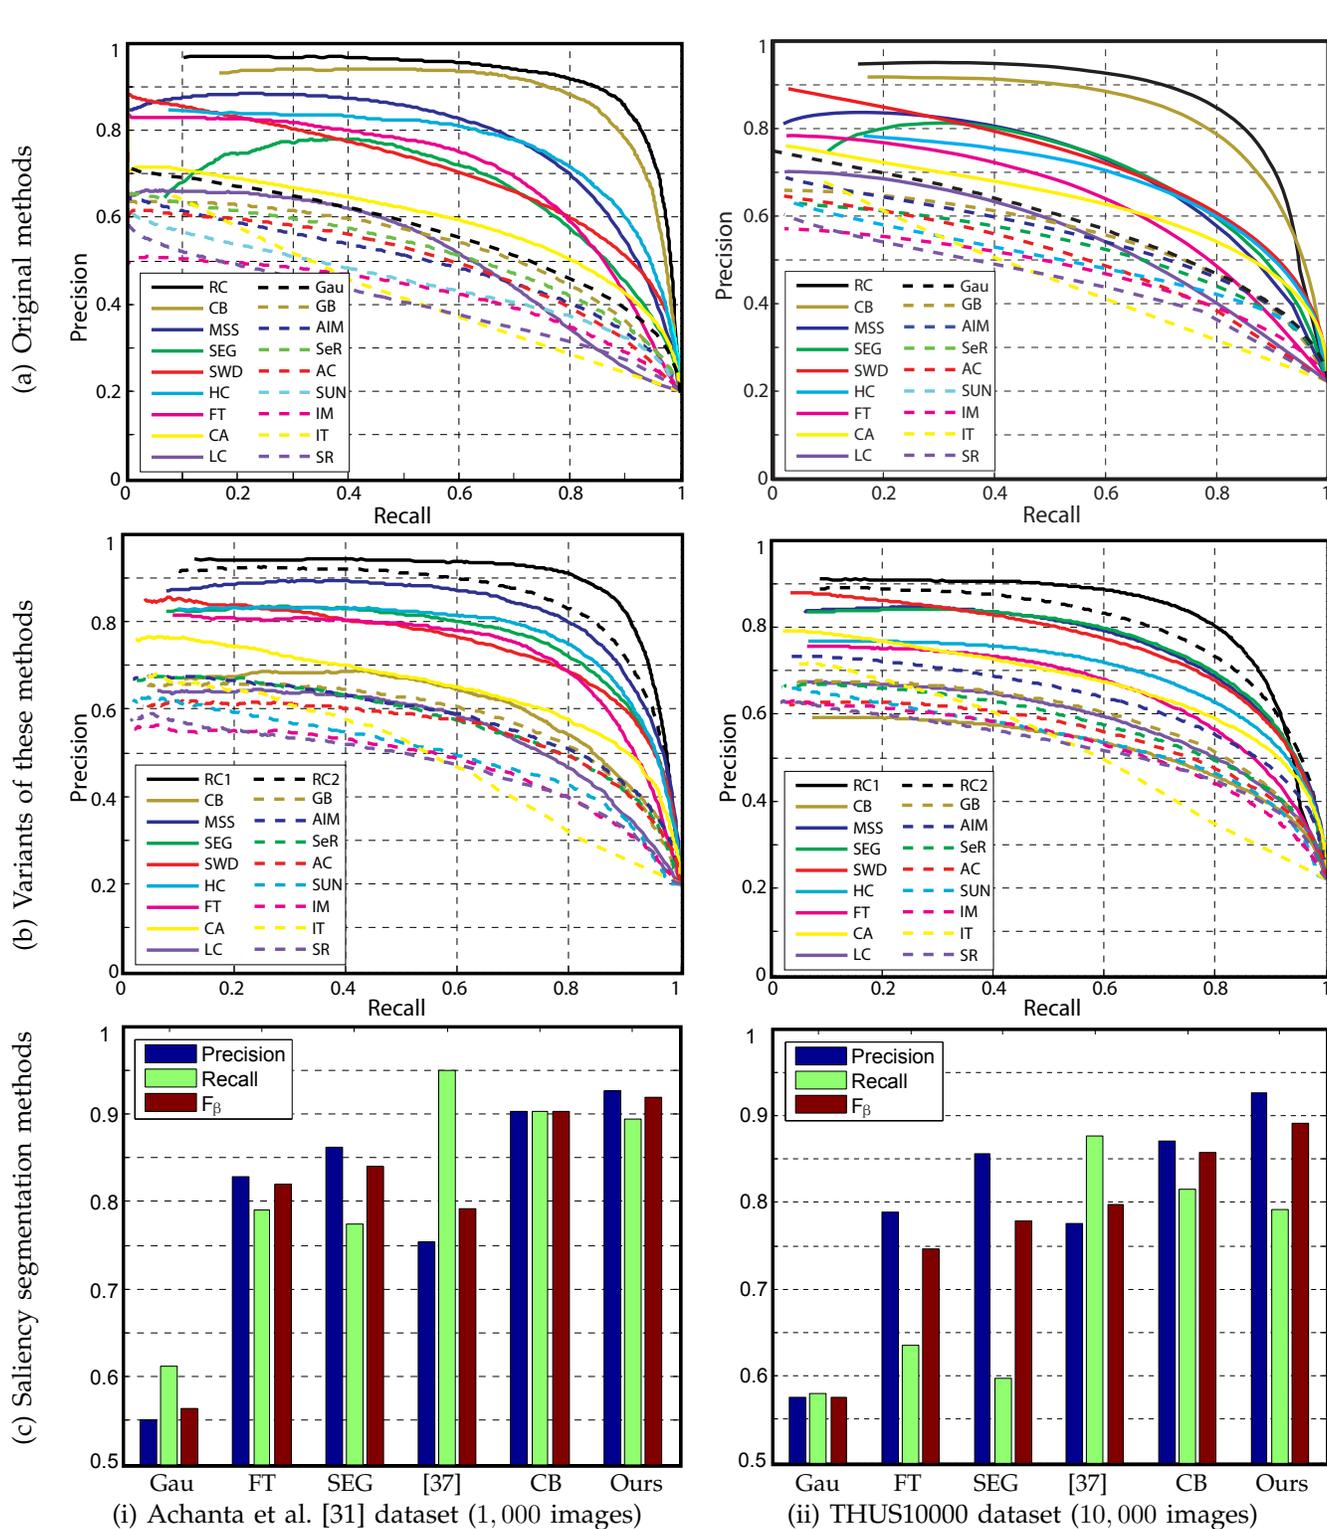


Fig. 11. Statistical comparison results of (a) different saliency detection methods, (b) their variants, and (c) object of interest region segmentation methods, using largest public available dataset (i) and (ii) our THUS10000 dataset.

fixation points shown in Fig. 12 is achieved by disabling the term encouraging similar appearance region receive similar saliency value, thus improves human fixation point prediction as demonstrated in Fig. 10. Although disabling the process explained in Sec. 4.3 improves eye fixation prediction performance, we argue that uniformly highlighting the entire object region is better in

many applications, including content aware image resizing [6], non-photorealistic rendering [34], adaptive image compression [5], and image mosaic [30]. Thus, although their own method [30] achieves best performance on eye fixation dataset [49], Margolin et al. [52] still choose to integrate our RC saliency maps to achieve better effects for various of image manipulation applications.

6.3 Object of interest image segmentation

To objectively evaluate our new SaliencyCut method using our RC-map as initialization, we compare our results with results obtained by other state-of-the-art methods for object of interest segmentation, i.e., FT [31], SEG [23], GrabCut [37] (initialized using 5 pixel wide image boundary), and CB [33] (best parameters are selected for these methods). Average precision, recall, and F -Measure are compared against the entire ground-truth database [31], with the F -Measure defined as:

$$F_{\beta} = \frac{(1 + \beta^2)Precision \times Recall}{\beta^2 \times Precision + Recall}. \quad (8)$$

We use $\beta^2 = 0.3$ as in Achanta et al. [31] to weigh precision more than recall. As can be seen from the comparison (see Fig. 11(c)), SaliencyCut using our RC saliency maps significantly outperforms other methods. As discussed by Liu et al. [1], recall rate is not as important as precision for attention detection. For example, a 100% recall rate can be achieved by simply selecting the whole image. Our approach reduced 57.2%, 50.9%, 46.5%, and 23.7% overall error rates on F-measure, compared with FT [31], SEG [23], GrabCut [37], and CB [33], respectively when evaluated using large accurate dataset (THUS10000). Besides producing higher F-Measure and robustness to location prior, our SaliencyCut is about 60 times faster (see Fig. 14) compared to CB [33]. Although several new methods [53], [54] have been developed since the initial version of this work [34], to the best of our knowledge, our salient object segmentation results are still the best results reported on the most widely used benchmark [31].

Although producing quite promising results for simple images as evaluated in Fig. 11, evaluation results for randomly-selected internet images Fig. 13 shows that there is still a need to develop more robust methods. For

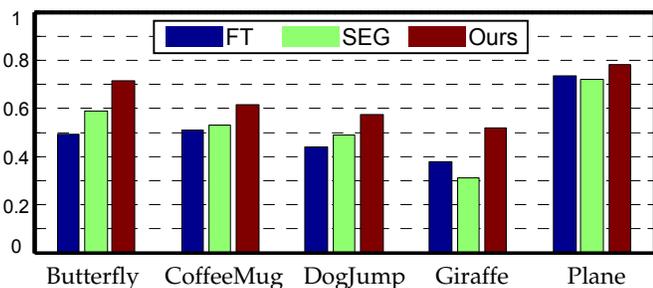


Fig. 13. Comparison of average F_{β} for different saliency segmentation methods: FT [31], SEG [23], and ours, on THUR15000 dataset [51].

Method	FT [31]	SEG [23]	CB [33]	Ours
Time (s)	0.247	7.48	36.5	0.621
Code	Matlab	Matlab & C	Matlab & C	C++

Fig. 14. Comparison of average time taken for different saliency segmentation methods. Segmentation results for THUS10000 dataset are available via our project page.

both datasets, our SaliencyCut’s performance is the best.

6.4 Sketch based image retrieval

Outline sketches are typically easier and faster for users to generate than a complete color description of the desired image. Sketch based image retrieval (SBIR) techniques become vital for users to leverage the increasing volumes of available image database. A large majority of potential users fail to precisely express fine details in their drawings. Thus most SBIR systems, which employ global descriptors, are unsatisfactory as they are unreliable under affine variations. To overcome such drawbacks, Eitz et al. [56], [57] use local descriptors to achieve state-of-the-art retrieval performance. The success of their methods is mainly attributed to translation invariance of local descriptors while using large local feature size (in the order of 20 – 25% of the image’s diagonal) to still retain large scale characteristics. However, for such large window sizes, there is simply not much space left for translating the sketch, thus limiting the translation invariance. SBIR still suffers from relatively low accuracy thus restricting its commercial potential.

Matching object shapes with clean background, however, is a relatively mature field. Even for the very challenging MPEG-7 dataset, state-of-the-art methods can achieve 91.61% retrieval rates [58]. Classical shape methods such as Shape Contexts (SC) [55] and Chamfer Matching [59] are mostly successful when dealing with limited background clutter. Selecting clean object outlines without influence from irrelevant image edges has great potential to improve current SBIR systems. Based on the observation that good results cannot be achieved without selection of segments, Bai et al. [60] use a shape band model to coarsely select candidate of edge segments while using Shape Context distance to decide the optimal matching. However, the shape band model requires user sketch for further detection thus does not allow preprocessing. It needs a few minutes to process a single image making it unsuited for real-time image retrieval applications.

Our SaliencyCut algorithm provides another possibility for automatically finding the outlines of object of interest on large scale image datasets. After such preprocessing, it becomes possible to make use of proven shape matching algorithms. We simply rank the images by SC [55] distance between their salient region outlines and user input sketches and compare with a state-of-the-art SBIR method using SHoG [56].

Experiments indicate that although our SaliencyCut method may produce less optimal results for noisy internet images, the shape matching method is very efficient in selecting those well segmented results. A quantitative evaluation in our THUR15000 dataset is shown in Fig. 16. One can see that our retrieval method is more effective than SHoG in terms of selecting user-desired candidates. Sample qualitative results are shown in Fig. 15. Compared with SHoG, our method gives results that



Fig. 15. Sketch based image comparison: first row shows images downloaded from Flickr using keyword ‘giraffe’, second row shows our retrieval results obtained by comparing user input sketch with SaliencyCut result using shape context measure [55]; third row shows corresponding sketch based retrieval results using SHoG [56].

TPR (%)	Butterfly		Coffee mug		Dog jump		Giraffe		Plane		Average	
	T50	T100	T50	T100	T50	T100	T50	T100	T50	T100	T50	T100
Flickr	28	28	58	51	56	55	30	25	44	48	41.4	43.2
Ours	58	52	88	93	86	90	72	66	88	90	78.2	78.4
SHoG [56]	36	40	82	78	74	73	18	18	90	91	60.0	60.0

Fig. 16. True positive ratios (TPR) among top 50 and 100 retrieval results. Results for SHOG are supplied by original authors. An image is considered as true positive if it contains a target object specified by the keywords.

are more relevant to user input sketches. Moreover, our method produces the precise boundary of the desired object, which makes it possible to reuse these segmented image components in many applications. Note that the extracted salient region features are complementary to other features like color, texture, and local edges.

7 DISCUSSION

We presented global contrast based saliency computation methods, namely Histogram based Contrast (HC) and spatial information-enhanced Region based Contrast (RC). While the HC method is fast and generates results with fine details, the RC method generates spatially coherent high quality saliency maps at the cost of reduced computational efficiency. Based on the proposed saliency detection method, we introduced a novel unsupervised segmentation algorithm, namely SaliencyCut, to automatically segment the most salient object in an image, without requiring expensive training data. The proposed methods were evaluated on several large scale publicly available benchmarks. The experimental results show that our methods consistently outperform other state-of-the-art methods in terms of precision and recall, while still being simple, fast, and efficient. For noisy internet images, although our saliency detection and segmentation methods cannot guarantee robust performance on individual images, their efficiency and simplicity makes it possible to automatically process a large number of images, which can then be further filtered for reliability and accuracy.

Limitations. Our methods aim at finding the most salient object in an image. It might produce sub-optimal results for images with multiple objects (e.g., Fig. 17), especially if the objects occlude each other (e.g., PASCAL VOC images [61]), for which even specialized object detectors fail to reliably generate good results for most



Fig. 17. Saliency object detection results for benchmark images [65] with multiple objects.

object classes. Rather than making hard decisions early on, proposing some candidate object regions [62]–[64] can be useful for those applications requiring high detection rate, e.g., object detection in cluttered scenes.

Future work. We further discuss possible applications and extensions by highlighting a few of the many exciting works using the preliminary version of our work:

- The ability to generate high quality saliency maps is essential for many applications including content-aware image manipulation [66], [67], non-photorealistic rendering [68], [69], image scene analysis [70]–[72] adaptive compression [73], forgery detection [74], [75], etc.
- Unsupervised segmentation of the entire salient object, without extensive training data annotation, naturally benefits applications like auto-cropping [52], scene classification [76], semantic image editing [77], [78], and data-driven image synthesis [79], [80].
- A tool to retrieve internet images and get precise object of interest regions is powerful to explore this big data for image composition [81], semantic colorization [82], information discovery [83], [84], image retrieval [85]–[87], etc.
- The proposed saliency measure has already been used to produce state-of-the-art results on cosegmentation benchmarks without using cosegmentation [88], or simultaneously analyze multiple images for better salient object extraction [51], [89].

ACKNOWLEDGMENTS

We would like to thank Prof. Z. Tu (for the suggestion of region level analysis), Prof. C. Rother, Prof. A. Zisserman, and the anonymous reviewers for their many useful feedbacks. This research was supported by the 973 Program (2011CB302205), the 863 Program (2009AA01Z327), the Key Project of S&T (2011ZX01042-001-002), NSFC (U0735001), EPSRC (EP/I001107/1), and ERC Starting Grant SmartGeometry.

REFERENCES

- [1] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, T. X., and S. H.Y., "Learning to detect a salient object," *IEEE TPAMI*, no. 2, 2011.
- [2] A. Borji, D. N. Sihite, and L. Itti, "Salient object detection: A benchmark," in *ECCV*, 2012.
- [3] M. Donoser, M. Urschler, M. Hirzer, and H. Bischof, "Saliency driven total variation segmentation," in *IEEE ICCV*, 2009, pp. 817–824.
- [4] U. Rutishauser, D. Walther, C. Koch, and P. Perona, "Is bottom-up attention useful for object recognition?" in *IEEE CVPR*, 2004.
- [5] C. Christopoulos, A. Skodras, and T. Ebrahimi, "The JPEG2000 still image coding system: an overview," *IEEE Trans. Consumer Elec.*, vol. 46, no. 4, pp. 1103–1127, 2002.
- [6] G.-X. Zhang, M.-M. Cheng, S.-M. Hu, and R. R. Martin, "A shape-preserving approach to image resizing," *Comput. Graph. Forum*, vol. 28, no. 7, pp. 1897–1906, 2009.
- [7] M.-M. Cheng, F.-L. Zhang, N. J. Mitra, X. Huang, and S.-M. Hu, "Repfinder: Finding approximately repeated scene elements for image editing," *ACM TOG*, vol. 29, no. 4, pp. 83:1–8, 2010.
- [8] T. Chen, M.-M. Cheng, P. Tan, A. Shamir, and S.-M. Hu, "Sketch2photo: Internet image montage," *ACM TOG*, vol. 28, no. 5, pp. 124:1–10, 2009.
- [9] S.-M. Hu, T. Chen, K. Xu, M.-M. Cheng, and R. R. Martin, "Internet visual media processing: a survey with graphics and vision applications," *The Visual Computer*, pp. 1–13, 2013.
- [10] Y. Gao, M. Wang, H. Luan, J. Shen, S. Yan, and D. Tao, "Tag-based social image search with visual-text joint hypergraph learning," in *ACM MM*, 2011, pp. 1517–1520.
- [11] H. Teuber, "Physiological psychology," *Annual Review of Psychology*, vol. 6, no. 1, pp. 267–296, 1955.
- [12] J. M. Wolfe and T. S. Horowitz, "What attributes guide the deployment of visual attention and how do they do it?" *Nature Reviews Neuroscience*, pp. 5:1–7, 2004.
- [13] R. Desimone and J. Duncan, "Neural mechanisms of selective visual attention," *Annual review of neuroscience*, vol. 18, no. 1, pp. 193–222, 1995.
- [14] S. K. Mannan, C. Kennard, and M. Husain, "The role of visual salience in directing eye movements in visual object agnosia," *Current biology*, vol. 19, no. 6, pp. 247–248, 2009.
- [15] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE TPAMI*, vol. 20, no. 11, pp. 1254–1259, 1998.
- [16] M.-M. Cheng, J. Warrell, W.-Y. Lin, S. Zheng, V. Vineet, and N. Crook, "Efficient salient region detection with soft image abstraction," in *IEEE ICCV*, 2013, pp. 1529–1536.
- [17] X. Hou, J. Harel, and C. Koch, "Image signature: Highlighting sparse salient regions," *IEEE TPAMI*, vol. 34, no. 1, 2012.
- [18] A. M. Triesman and G. Gelade, "A feature-integration theory of attention," *Cognitive Psychology*, vol. 12, pp. 97–136, 1980.
- [19] C. Koch and S. Ullman, "Shifts in selective visual attention: towards the underlying neural circuitry," *Human Neurobiology*, vol. 4, pp. 219–227, 1985.
- [20] N. Bruce and J. Tsotsos, "Saliency, attention, and visual search: An information theoretic approach," *Journal of Vision*, vol. 9, 2009.
- [21] N. Murray, M. Vanrell, X. Otazu, and C. A. Parraga, "Saliency estimation using a non-parametric low-level vision model," in *IEEE CVPR*, 2011, pp. 433–440.
- [22] R. Achanta and S. Süsstrunk, "Saliency detection using maximum symmetric surround," in *IEEE ICIP*, 2010, pp. 2653–2656.
- [23] E. Rahtu, J. Kannala, M. Salo, and J. Heikkilä, "Segmenting salient objects from images and videos," in *ECCV*, 2010.
- [24] H. Seo and P. Milanfar, "Static and space-time visual saliency detection by self-resemblance," *Journal of vision*, vol. 9, 2009.
- [25] L. Zhang, M. Tong, T. Marks, H. Shan, and G. Cottrell, "SUN: A bayesian framework for saliency using natural statistics," *Journal of Vision*, vol. 8, no. 7, pp. 32:1–20, 2008.
- [26] L. Duan, C. Wu, J. Miao, L. Qing, and Y. Fu, "Visual saliency detection by spatially weighted dissimilarity," in *IEEE CVPR*, 2011, pp. 473–480.
- [27] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," in *NIPS*, 2007, pp. 545–552.
- [28] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *IEEE CVPR*, 2007, pp. 1–8.
- [29] R. Achanta, F. Estrada, P. Wils, and S. Süsstrunk, "Salient region detection and segmentation," in *IEEE ICVS*, 2008.
- [30] S. Goferman, L. Zelnik-Manor, and A. Tal, "Context-aware saliency detection," *IEEE TPAMI*, vol. 34, no. 10, 2012.
- [31] R. Achanta, S. Hemami, F. Estrada, and S. Süsstrunk, "Frequency-tuned salient region detection," in *IEEE CVPR*, 2009.
- [32] Y. Zhai and M. Shah, "Visual attention detection in video sequences using spatiotemporal cues," in *ACM Multimedia*, 2006, pp. 815–824.
- [33] H. Jiang, J. Wang, Z. Yuan, T. Liu, N. Zheng, and S. Li, "Automatic salient object segmentation based on context and shape prior," in *BMVC*, 2011, pp. 1–12.
- [34] M.-M. Cheng, G.-X. Zhang, N. J. Mitra, X. Huang, and S.-M. Hu, "Global contrast based salient region detection," in *IEEE CVPR*, 2011, pp. 409–416.
- [35] J. Reynolds and R. Desimone, "Interacting roles of attention and visual salience in v4," *Neuron*, vol. 37, no. 5, 2003.
- [36] P. Reinagel, A. Zador *et al.*, "Natural scene statistics at the centre of gaze," *Network: Computation in Neural Systems*, vol. 10, no. 4, pp. 341–350, 1999.
- [37] C. Rother, V. Kolmogorov, and A. Blake, "GrabCut"– Interactive foreground extraction using iterated graph cuts," *ACM TOG*, vol. 23, no. 3, pp. 309–314, 2004.
- [38] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE TPAMI*, 2012.
- [39] A. Borji, D. Sihite, and L. Itti, "Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study," *IEEE TIP*, 2012.
- [40] T. Judd, F. Durand, and A. Torralba, "A benchmark of computational models of saliency to predict human fixations," MIT tech report, Tech. Rep., 2012.
- [41] Y.-F. Ma and H.-J. Zhang, "Contrast-based image attention analysis by using fuzzy growing," in *ACM Multimedia*, 2003.
- [42] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," in *NIPS*, 2006, pp. 545–552.
- [43] B. Ko and J. Nam, "Object-of-interest image segmentation based on human attention and semantic region clustering," *J. Opt. Soc. Am.*, vol. 23, no. 10, pp. 2462–2470, 2006.
- [44] J. Han, K. Ngan, M. Li, and H. Zhang, "Unsupervised extraction of visual attention objects in color images," *IEEE TCSV*, vol. 16, no. 1, pp. 141–145, 2006.
- [45] W. Eihhauser and P. König, "Does luminance-contrast contribute to a saliency map for overt visual attention?" *European Journal of Neuroscience*, vol. 17, pp. 1089–1097, 2003.
- [46] P. Felzenszwalb and D. Huttenlocher, "Efficient graph-based image segmentation," *IJCV*, vol. 59, no. 2, pp. 167–181, 2004.
- [47] Y. Boykov, O. Veksler, and R. Zabih, "Fast approximate energy minimization via graph cuts," *IEEE TPAMI*, vol. 23, no. 11, pp. 1222–1239, 2001.
- [48] P. Mehrani and O. Veksler, "Saliency segmentation based on learning and graph cut refinement," in *BMVC*, 2010.
- [49] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *IEEE ICCV*, 2009.
- [50] Z. Wang and B. Li, "A two-stage approach to saliency detection in images," in *IEEE ICASSP*, 2008, pp. 965–968.
- [51] M.-M. Cheng, N. J. Mitra, X. Huang, and S.-M. Hu, "SalientShape: Group Saliency in Image Collections," *The Visual Computer*, 2013.
- [52] R. Margolin, L. Zelnik-Manor, and A. Tal, "Saliency for image manipulation," *The Visual Computer*, pp. 1–12, 2012.
- [53] Q. Yan, L. Xu, J. Shi, and J. Jia, "Hierarchical saliency detection," *CVPR*, 2013.
- [54] M. Tang, L. Gorelick, O. Veksler, and Y. Boykov, "Grabcut in one cut," in *IEEE ICCV*, 2013.

- [55] S. Belongie, J. Malik, and J. Puzicha, "Shape matching and object recognition using shape contexts," *IEEE TPAMI*, vol. 24, no. 4, pp. 509–522, 2002.
- [56] M. Eitz, K. Hildebrand, T. Boubekeur, and M. Alexa, "Sketch-based image retrieval: benchmark and bag-of-features descriptors," *IEEE TVCG*, 2011, preprint.
- [57] —, "An evaluation of descriptors for large-scale image retrieval from sketched feature lines," *Computers & Graphics*, vol. 34, no. 5, pp. 482–498, 2010.
- [58] X. Bai, X. Yang, L. Latecki, W. Liu, and Z. Tu, "Learning context-sensitive shape similarity by graph transduction," *IEEE TPAMI*, pp. 861–874, 2010.
- [59] A. Thayananthan, B. Stenger, P. Torr, and R. Cipolla, "Shape context and chamfer matching in cluttered scenes," in *IEEE CVPR*, vol. 1, 2003, pp. 127–133.
- [60] X. Bai, Q. N. Li, L. J. Latecki, W. Y. Liu, and Z. W. Tu, "Shape band: A deformable object detection approach," in *IEEE CVPR*, 2009, pp. 1335–1342.
- [61] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The PASCAL visual object classes (voc) challenge," *IJCV*, vol. 88, no. 2, pp. 303–338, 2010.
- [62] B. Alexe, T. Deselaers, and V. Ferrari, "Measuring the objectness of image windows," *IEEE TPAMI*, vol. 34, no. 11, 2012.
- [63] I. Endres and D. Hoiem, "Category-independent object proposals with diverse ranking," *IEEE TPAMI*, 2013.
- [64] M.-M. Cheng, Z. Zhang, and P. Torr, "BING: Binarized normed gradients for objectness estimation at 300fps," in *IEEE CVPR*, 2014.
- [65] S. Alpert, M. Galun, R. Basri, and A. Brandt, "image segmentation by probabilistic bottom-up aggregation and cue integration," in *IEEE CVPR*, "2007".
- [66] D. Panozzo, O. Weber, and O. Sorokin, "Robust image retargeting via axis-aligned deformation," vol. 31, no. 21, pp. 229–236, 2012.
- [67] L. Mai, H. Le, Y. Niu, and F. Liu, "Rule of thirds detection from photograph," in *IEEE ISM*, 2011, pp. 91–96.
- [68] P. L. Rosin and Y.-K. Lai, "Artistic minimal rendering with lines and blocks," *Graphical Models*, vol. 75, no. 4, pp. 208–229, 2013.
- [69] J. Han, E. J. Pauwels, and P. De Zeeuw, "Fast saliency-aware multi-modality image fusion," *Neurocomputing*, pp. 70–80, 2013.
- [70] A. Borji, D. N. Sihite, and L. Itti, "What stands out in a scene? a study of human explicit saliency judgment," *Vision research*, vol. 91, pp. 62–77, 2013.
- [71] C. Zhang, X. Li, X. Ruan, Y. Zhao, and M.-H. Yang, "Discriminative generative contour detection," in *BMVC*, 2013.
- [72] S. Zheng, M.-M. Cheng, J. Warrell, P. Sturgess, V. Vineet, C. Rother, and P. Torr, "Dense semantic image segmentation with objects and attributes," in *IEEE CVPR*, 2014.
- [73] C. Deng, W. Lin, and J. Cai, "Content-based image compression for arbitrary-resolution display devices," *IEEE TMM*, 2012.
- [74] H. Fu and X. Cao, "Forgery authentication in extreme wide-angle lens using distortion cue and fake saliency map," *IEEE Transactions on Information Forensics and Security*, no. 4, 2012.
- [75] L.-Q. Ma, K. Xu, T.-T. Wong, B.-Y. Jiang, and S.-M. Hu, "Change blindness images," *IEEE TVCG*, pp. 1808–1819, 2013.
- [76] Z. Ji, J. Wang, Y. Su, Z. Song, and S. Xing, "Balance between object and background: Object-enhanced features for scene image classification," *Neurocomputing*, vol. 120, pp. 15–23, 2013.
- [77] X. Wang, J. Jia, and L. Cai, "Affective image adjustment with a single word," *The Visual Computer*, pp. 1121–1133, 2013.
- [78] M.-M. Cheng, S. Zheng, W.-Y. Lin, J. Warrell, V. Vineet, P. Sturgess, N. Crook, N. Mitra, and P. Torr, "ImageSpirit: Verbal guided image parsing," *arXiv:1310.4389*, 2013.
- [79] C. Goldberg, T. Chen, F.-L. Zhang, A. Shamir, and S.-M. Hu, "Data-driven object manipulation in images," pp. 265–274, 2012.
- [80] R.-F. Tong, Y. Zhang, and K.-L. Cheng, "Stereopasting: Interactive composition in stereoscopic images," *IEEE TVCG*, 2013.
- [81] H. Huang, L. Zhang, and H.-C. Zhang, "Arcimboldo-like collage using internet images," *ACM TOG*, vol. 30, 2011.
- [82] Y. S. Chia, S. Zhuo, R. K. Gupta, Y.-W. Tai, S.-Y. Cho, P. Tan, and S. Lin, "Semantic colorization with internet images," *ACM TOG*, vol. 30, no. 6, pp. 156:1–156:8, 2011.
- [83] H. Liu, L. Zhang, and H. Huang, "Web-image driven best views of 3d shapes," *The Visual Computer*, pp. 1–9, 2012.
- [84] J.-Y. Zhu, J. Wu, Y. Wei, E. Chang, and Z. Tu, "Unsupervised object class discovery via saliency-guided multiple class learning," in *IEEE CVPR*, 2012, pp. 3218–3225.
- [85] J. He, J. Feng, X. Liu, T. Cheng, T.-H. Lin, H. Chung, and S.-F. Chang, "Mobile product search with bag of hash bits and boundary reranking," in *IEEE CVPR*, 2012, pp. 3005–3012.
- [86] T. Chen, P. Tan, L.-Q. Ma, M.-M. Cheng, A. Shamir, and S.-M. Hu, "Poseshop: Human image database construction and personalized content synthesis," *IEEE TVCG*, vol. 19, no. 5, pp. 824–837, 2013.
- [87] Y. Gao, M. Wang, D. Tao, R. Ji, and Q. Dai, "3-d object retrieval and recognition with hypergraph analysis," *IEEE TIP*, vol. 21, no. 9, pp. 4290–4303, 2012.
- [88] M. Rubinstein, A. Joulin, J. Kopf, and C. Liu, "Unsupervised joint object discovery and segmentation in internet images," in *IEEE CVPR*, 2013, pp. 1939–1946.
- [89] H. Fu, X. Cao, and Z. Tu, "Cluster-based co-saliency detection," *IEEE TIP*, 2013.



Ming-Ming Cheng received his PhD degree from Tsinghua University in 2012. He is currently a research fellow in The University of Oxford, working with Prof. Philip Torr. His research interests include computer graphics, computer vision, image processing, and image retrieval. He has received the Google PhD fellowship award, the IBM PhD fellowship award, and the "New PhD Researcher Award" from Chinese Ministry of Education. He reviews papers regularly for IEEE TPAMI, ACM SIGGRAPH, IEEE TIP, etc.



Niloy J. Mitra received his PhD degree from Stanford University in 2006. He is currently a reader (associate professor) in the Computer Science Department at the University College London (UCL). His research interests span a range of topics including shape analysis, geometry processing, shape manipulation, image analysis, and recreational art. He has received ACM SIGGRAPH 'Significant New Researcher' award. He is on the editorial boards of ACM Transactions on Graphics, Computer Graphics Forum, Computer & Graphics, and The Visual Computer.



Xiaolei Huang received her PhD degrees from Rutgers University in 2006. She is currently an assistant professor in Lehigh University. Her research interests are in computer vision, computer graphics, and multimedia retrieval. Dr. Huang serves on the program committees of several international conferences on computer vision and computer graphics, and she reviews papers regularly for journals including the IEEE TPAMI, and the IEEE TIP. She is a member of the IEEE.



Philip H. S. Torr received the PhD degree from Oxford University. After working for another three years at Oxford, he worked for six years as a research scientist for Microsoft Research, first in Redmond, then in Cambridge, founding the vision side of the Machine Learning and Perception Group. He is now a professor at Oxford University. He has won awards from several top vision conferences, including ICCV, CVPR, ECCV, NIPS and BMVC. He is a senior member of the IEEE, Royal Society Wolfson Research Merit Award holder, and program co-chair of ICCV 2013.



Shi-Min Hu received the PhD degree from Zhejiang University in 1996. He is currently a chair professor of computer science in the Dept. of Computer Science and Technology, Tsinghua University, Beijing. His research interests include digital geometry processing, video processing, rendering, computer animation, and computer-aided geometric design. He is an associate editor-in-chief of The Visual Computer, and on the editorial board of IEEE TVCG, Computer Aided Design.

Global Contrast based Salient Region Detection

Ming-Ming Cheng¹ Guo-Xin Zhang¹ Niloy J. Mitra² Xiaolei Huang³ Shi-Min Hu¹
¹ TNList, Tsinghua University ² KAUST ³ Lehigh University

chengmingvictor@gmail.com

Abstract

Reliable estimation of visual saliency allows appropriate processing of images without prior knowledge of their contents, and thus remains an important step in many computer vision tasks including image segmentation, object recognition, and adaptive compression. We propose a regional contrast based saliency extraction algorithm, which simultaneously evaluates global contrast differences and spatial coherence. The proposed algorithm is simple, efficient, and yields full resolution saliency maps. Our algorithm consistently outperformed existing saliency detection methods, yielding higher precision and better recall rates, when evaluated using one of the largest publicly available data sets. We also demonstrate how the extracted saliency map can be used to create high quality segmentation masks for subsequent image processing.

1. Introduction

Humans routinely and effortlessly judge the importance of image regions, and focus attention on important parts. Computationally detecting such salient image regions remains a significant goal, as it allows preferential allocation of computational resources in subsequent image analysis and synthesis. Extracted saliency maps are widely used in many computer vision applications including object-of-interest image segmentation [13, 18], object recognition [25], adaptive compression of images [6], content-aware image editing [28, 33, 30, 9], and image retrieval [4].

Saliency originates from visual uniqueness, unpredictability, rarity, or surprise, and is often attributed to variations in image attributes like color, gradient, edges, and boundaries. Visual saliency, being closely related to how we perceive and process visual stimuli, is investigated by multiple disciplines including cognitive psychology [26, 29], neurobiology [8, 22], and computer vision [17, 2]. Theories of human attention hypothesize that the human vision system only processes parts of an image in detail, while leaving others nearly unprocessed. Early work by Treisman and Gelade [27], Koch and Ullman [19], and subse-



Figure 1. Given input images (top), a global contrast analysis is used to compute high resolution saliency maps (middle), which can be used to produce masks (bottom) around regions of interest.

quent attention theories proposed by Itti, Wolfe and others, suggest two stages of visual attention: fast, pre-attentive, bottom-up, data driven saliency extraction; and slower, task dependent, top-down, goal driven saliency extraction.

We focus on bottom-up data driven saliency detection using image contrast. It is widely believed that human cortical cells may be *hard wired* to preferentially respond to high contrast stimulus in their receptive fields [23]. We propose contrast analysis for extracting high-resolution, full-field saliency maps based on the following observations:

- A global contrast based method, which separates a large-scale object from its surroundings, is preferred over local contrast based methods producing high saliency values at or near object edges.
- Global considerations enable assignment of comparable saliency values to similar image regions, and can uniformly highlight entire objects.
- Saliency of a region depends mainly on its contrast to the nearby regions, while contrasts to distant regions are less significant.
- Saliency maps should be fast and easy to generate to allow processing of large image collections, and facilitate efficient image classification and retrieval.

We propose a *histogram-based contrast method* (HC) to measure saliency. HC-maps assign pixel-wise saliency val-



(a) original (b) IT[17] (c) MZ[21] (d) GB[14] (e) SR[15] (f) AC[1] (g) CA[12] (h) FT[2] (i) LC[32] (j) HC (k) RC
 Figure 2. Saliency maps computed by different state-of-the-art methods (b-i), and with our proposed HC (j) and RC methods (k). Most results highlight edges, or are of low resolution. See also Figure 6 (and our project webpage).

ues based simply on color separation from all other image pixels to produce full resolution saliency maps. We use a histogram-based approach for efficient processing, while employing a smoothing procedure to control quantization artifacts. Note that our algorithm is targeted towards natural scenes, and maybe suboptimal for extracting saliency of highly textured scenes (see Figure 12).

As an improvement over HC-maps, we incorporate spatial relations to produce *region-based contrast* (RC) maps where we first segment the input image into regions, and then assign saliency values to them. The saliency value of a region is now calculated using a global contrast score, measured by the region’s contrast and spatial distances to other regions in the image.

We have extensively evaluated our methods on publicly available benchmark data sets, and compared our methods with (eight) state-of-the-art saliency methods [17, 21, 32, 14, 15, 1, 2, 12] as well as with manually produced ground truth annotations¹. The experiments show significant improvements over previous methods both in precision and recall rates. Overall, compared with HC-maps, RC-maps produce better precision and recall rates, but at the cost of increased computations. Encouragingly, we observe that the saliency cuts extracted using our saliency maps are, in most cases, comparable to manual annotations. We also present application of the extracted saliency maps to segmentation, context aware resizing, and non-photo realistic rendering.

2. Related Work

We focus on relevant literature targeting pre-attentive bottom-up saliency detection, which may be biologically motivated, or purely computational, or involve both aspects. Such methods utilize low-level processing to determine the contrast of image regions to their surroundings, using feature attributes such as intensity, color, and edges [2]. We broadly classify the algorithms into local and global schemes.

Local contrast based methods investigate the rarity of image regions with respect to (small) local neighborhoods. Based on the highly influential biologically inspired *early representation* model introduced by Koch and Ullman [19], Itti et al. [17] define image saliency using central-surrounded differences across multi-scale image features.

¹Results for 1000 images and prototype software are available at the project webpage: <http://cg.cs.tsinghua.edu.cn/people/%7Ecmm/saliency/>

Ma and Zhang [21] propose an alternative local contrast analysis for generating saliency maps, which is then extended using a fuzzy growth model. Harel et al. [14] normalize the feature maps of Itti et al., to highlight conspicuous parts and permit combination with other importance maps. Liu et al. [20] find multi-scale contrast by linearly combining contrast in a Gaussian image pyramid. More recently, Goferman et al. [12] simultaneously model local low-level clues, global considerations, visual organization rules, and high-level features to highlight salient objects along with their contexts. Such methods using local contrast tend to produce higher saliency values near edges instead of uniformly highlighting salient objects (see Figure 2).

Global contrast based methods evaluate saliency of an image region using its contrast with respect to the entire image. Zhai and Shah [32] define pixel-level saliency based on a pixel’s contrast to all other pixels. However, for efficiency they use only luminance information, thus ignoring distinctiveness clues in other channels. Achanta et al. [2] propose a frequency tuned method that directly defines pixel saliency using a pixel’s color difference from the average image color. The elegant approach, however, only considers first order average color, which can be insufficient to analyze complex variations common in natural images. In Figures 6 and 7, we show qualitative and quantitative weaknesses of such approaches. Furthermore, these methods ignore spatial relationships across image parts, which can be critical for reliable and coherent saliency detection (see Section 5).

3. Histogram Based Contrast

Based on the observation from biological vision that the vision system is sensitive to contrast in visual signal, we propose a histogram-based contrast (HC) method to define saliency values for image pixels using color statistics of the input image. Specifically, the saliency of a pixel is defined using its color contrast to all other pixels in the image, i.e., the saliency value of a pixel I_k in image I is defined as,

$$S(I_k) = \sum_{\forall I_i \in I} D(I_k, I_i), \quad (1)$$

where $D(I_k, I_i)$ is the color distance metric between pixels I_k and I_i in the $L^*a^*b^*$ space (see also [32]). Equation 1 can be expanded by pixel order to have the following form,

$$S(I_k) = D(I_k, I_1) + D(I_k, I_2) + \dots + D(I_k, I_N), \quad (2)$$

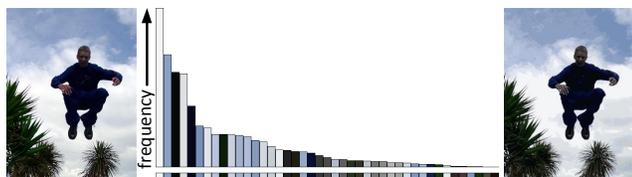


Figure 3. Given an input image (left), we compute its color histogram (middle). Corresponding histogram bin colors are shown in the lower bar. The quantized image (right) uses only 43 histogram bin colors and still retains sufficient visual quality for saliency detection.

where N is the number of pixels in image I . It is easy to see that pixels with the same color value have the same saliency value under this definition, since the measure is oblivious to spatial relations. Hence, rearranging Equation 2 such that the terms with the same color value c_j are grouped together, we get saliency value for each color as,

$$S(I_k) = S(c_l) = \sum_{j=1}^n f_j D(c_l, c_j), \quad (3)$$

where c_l is the color value of pixel I_k , n is the number of distinct pixel colors, and f_j is the probability of pixel color c_j in image I . Note that in order to prevent salient region color statistics from being corrupted by similar colors from other regions, one can develop a similar scheme using varying window masks. However, given the strict efficiency requirement, we take the simple global approach.

Histogram based speed up. Naively evaluating the saliency value for each image pixel using Equation 1 takes $O(N^2)$ time, which is computationally too expensive even for medium sized images. The equivalent representation in Equation 3, however, takes $O(N) + O(n^2)$ time, implying that computational efficiency can be improved to $O(N)$ if $O(n^2) \leq O(N)$. Thus, the key to speed up is to reduce the number of pixel colors in the image. However, the true-color space contains 256^3 possible colors, which is typically larger than the number of image pixels.

Zhai and Shah [32] reduce the number of colors, n , by only using luminance. In this way, $n^2 = 256^2$ (typically $256^2 \ll N$). However, their method has the disadvantage that the distinctiveness of color information is ignored. In this work, we use the full color space instead of luminance only. To reduce the number of colors needed to consider, we first quantize each color channel to have 12 different values, which reduces the number of colors to $12^3 = 1728$. Considering that color in a natural image typically covers only a small portion of the full color space, we further reduce the number of colors by ignoring less frequently occurring colors. By choosing more frequently occurring colors and ensuring these colors cover the colors of more than 95% of the image pixels, we typically are left with around $n = 85$ colors (see Section 5 for experimental details). The colors

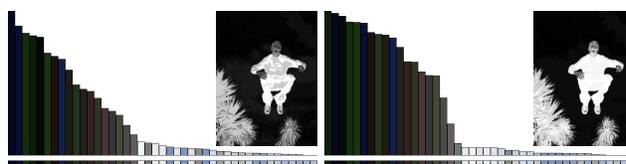


Figure 4. Saliency of each color, normalized to the range $[0, 1]$, before (left) and after (right) color space smoothing. Corresponding saliency maps are shown in the respective insets.

of the remaining pixels, which comprise fewer than 5% of the image pixels, are replaced by the closest colors in the histogram. A typical example of such quantization is shown in Figure 3. Note that again due to efficiency requirements we select the simple histogram based quantization instead of optimizing for an image specific color palette.

Color space smoothing. Although we can efficiently compute color contrast by building a compact color histogram using color quantization and choosing more frequent colors, the quantization itself may introduce artifacts. Some similar colors may be quantized to different values. In order to reduce noisy saliency results caused by such randomness, we use a smoothing procedure to refine the saliency value for each color. We replace the saliency value of each color by the weighted average of the saliency values of similar colors (measured by $L^*a^*b^*$ distance). This is actually a smoothing process in the color feature space. Typically we choose $m = n/4$ nearest colors to refine the saliency value of color c by,

$$S'(c) = \frac{1}{(m-1)T} \sum_{i=1}^m (T - D(c, c_i)) S(c_i) \quad (4)$$

where $T = \sum_{i=1}^m D(c, c_i)$ is the sum of distances between color c and its m nearest neighbors c_i , and the normalization factor comes from $\sum_{i=1}^m (T - D(c, c_i)) = (m-1)T$. Note that we use a linearly-varying smoothing weight $(T - D(c, c_i))$ to assign larger weights to colors closer to c in the color feature space. In our experiments, we found that such linearly-varying weights are better than Gaussian weights, which fall off too sharply. Figure 4 shows the typical effect of color space smoothing with the corresponding histograms sorted by decreasing saliency values. Note that similar histogram bins are closer to each other after such smoothing, indicating that similar colors have higher likelihood of being assigned similar saliency values, thus reducing quantization artifacts (see Figure 7).

Implementation details. To quantize the color space into 12^3 different colors, we uniformly divide each color channel into 12 different levels. While the quantization of colors is performed in the RGB color space, we measure color differences in the $L^*a^*b^*$ color space because of its perceptual accuracy. However, we do not perform quantization directly



Figure 5. Image regions generated by Felzenszwalb and Huttenlocher’s segmentation method [11] (left), region contrast based segmentation with (left-middle) and without (right-middle) distance weighting. Incorporating the spatial context, we get a high quality saliency cut (right) comparable to human labeled ground truth.

in the $L^*a^*b^*$ color space since not all colors in the range $L^* \in [0, 100]$, and $a^*, b^* \in [-127, 127]$ necessarily correspond to real colors. Experimentally we observed worse quantization artifacts using direct $L^*a^*b^*$ color space quantization. Best results were obtained by quantization in the RGB space while measuring distance in the $L^*a^*b^*$ color space, as opposed to performing both quantization and distance calculation in a single color space, either RGB or $L^*a^*b^*$.

4. Region Based Contrast

Humans pay more attention to those image regions that contrast strongly with their surroundings [10]. Besides contrast, spatial relationships play an important role in human attention. High contrast to its surrounding regions is usually stronger evidence for saliency of a region than high contrast to far-away regions. Since directly introducing spatial relationships when computing pixel-level contrast is computationally expensive, we introduce a contrast analysis method, *region contrast* (RC), so as to integrate spatial relationships into region-level contrast computation. In RC, we first segment the input image into regions, then compute color contrast at the region level, and define the saliency for each region as the weighted sum of the region’s contrasts to all other regions in the image. The weights are set according to the spatial distances with farther regions being assigned smaller weights.

Region contrast by sparse histogram comparison. We first segment the input image into regions using a graph-based image segmentation method [11]. Then we build the color histogram for each region as in Section 3. For a region r_k , we compute its saliency value by measuring its color contrast to all other regions in the image,

$$S(r_k) = \sum_{r_k \neq r_i} w(r_i) D_r(r_k, r_i), \quad (5)$$

where $w(r_i)$ is the weight of region r_i and $D_r(\cdot, \cdot)$ is the color distance metric between the two regions. Here we

use the number of pixels in r_i as $w(r_i)$ to emphasize color contrast to bigger regions. The color distance between two regions r_1 and r_2 is defined as,

$$D_r(r_1, r_2) = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} f(c_{1,i}) f(c_{2,j}) D(c_{1,i}, c_{2,j}) \quad (6)$$

where $f(c_{k,i})$ is the probability of the i -th color $c_{k,i}$ among all n_k colors in the k -th region r_k , $k = \{1, 2\}$. Note that we use the probability of a color in the probability density function (i.e. normalized color histogram) of the region as the weight for this color to emphasize more the color differences between dominant colors.

Storing and calculating the regular matrix format histogram for each region is inefficient since each region typically contains a small number of colors in the color histogram of the whole image. Instead, we use a sparse histogram representation for efficient storage and computation.

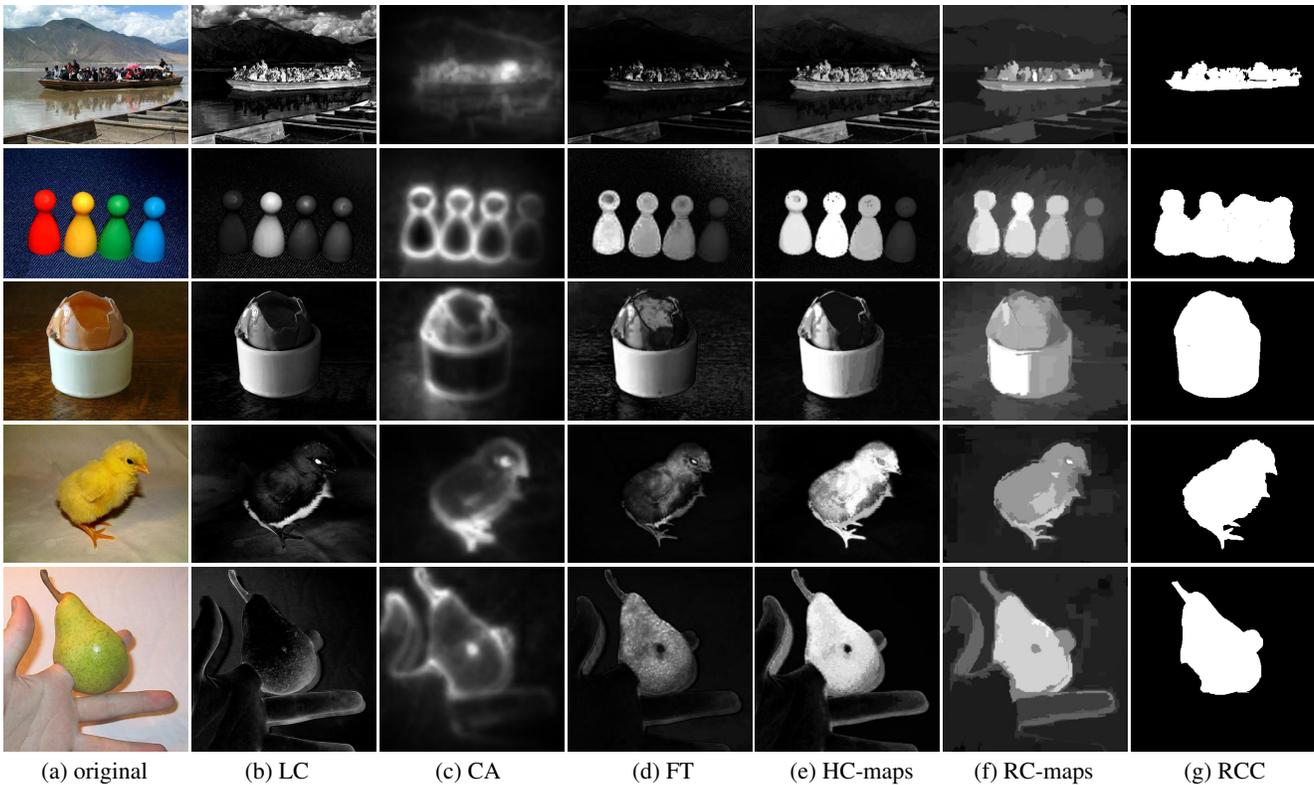
Spatially weighted region contrast. We further incorporate spatial information by introducing a spatial weighting term in Equation 5 to increase the effects of closer regions and decrease the effects of farther regions. Specifically, for any region r_k , the spatially weighted region contrast based saliency is defined as:

$$S(r_k) = \sum_{r_k \neq r_i} \exp(-D_s(r_k, r_i)/\sigma_s^2) w(r_i) D_r(r_k, r_i) \quad (7)$$

where $D_s(r_k, r_i)$ is the spatial distance between regions r_k and r_i , and σ_s controls the strength of spatial weighting. Larger values of σ_s reduce the effect of spatial weighting so that contrast to farther regions would contribute more to the saliency of the current region. The spatial distance between two regions is defined as the Euclidean distance between their centroids. In our implementation, we use $\sigma_s^2 = 0.4$ with pixel coordinates normalized to $[0, 1]$.

5. Experimental Comparisons

We have evaluated the results of our approach on the publicly available database provided by Achanta et al. [2]. To the best of our knowledge, the database is the largest of its kind, and has ground truth in the form of accurate human-marked labels for salient regions. We compared the proposed global contrast based methods with 8 state-of-the-art saliency detection methods. Following [2], we selected these methods according to: number of citations (IT[17] and SR[15]), recency (GB[14], SR, AC[1], FT[2] and CA[12]), variety (IT is biologically-motivated, MZ[21] is purely computational, GB is hybrid, SR works in the frequency domain, AC and FT output full resolution saliency maps), and being related to our approach (LC[32]).



(a) original (b) LC (c) CA (d) FT (e) HC-maps (f) RC-maps (g) RCC
 Figure 6. Visual comparison of saliency maps. (a) original images, saliency maps produced using (b) Zhai and Shah [32], (c) Goferman et al. [12], (d) Achanta et al. [2], (e) our HC and (f) RC methods, and (g) RC-based saliency cut results. Our methods generate uniformly highlighted salient regions. (See our project webpage for all results on the full benchmark dataset.)

We used our methods and the others to compute saliency maps for all the 1000 images in the database. Table 1 compares the average time taken by each method. Our algorithms, HC and RC, are implemented in C++. For the other methods namely IT, GB, SR, FT and CA, we used the authors’ implementations, while for LC, we implemented the algorithm in C++ since we could not find the authors’ implementation. For typical natural images, our HC method needs $O(N)$ computation time and is sufficiently efficient for real-time applications. In contrast, our RC variant is slower as it requires image segmentation [11], but produces superior quality saliency maps.

In order to comprehensively evaluate the accuracy of our methods for salient object segmentation, we performed two experiments using different objective comparison measures. In the first experiment, to segment salient objects and calculate precision and recall curves [15], we binarized the saliency map using every possible fixed threshold, similar to the fixed thresholding experiment in [2]. In the second experiment, we segment salient objects by iteratively applying GrabCut [24] initialized using thresholded saliency maps, as we will describe later. We also use the obtained saliency maps as importance weighting for content aware image resizing and non-photo realistic rendering.

Segmentation by fixed thresholding. The simplest way to get a binary segmentation of salient objects is to threshold the saliency map with a threshold $T_f \in [0, 255]$. To reliably compare how well various saliency detection methods highlight salient regions in images, we vary the threshold T_f from 0 to 255. Figure 7 shows the resulting precision vs. recall curves. We also present the benefits of adding the color space smoothing and spatial weighting schemes, along with objective comparison with other saliency extraction methods. Visual comparison of saliency maps obtained by the various methods can be seen in Figures 2 and 6.

The precision and recall curves clearly show that our methods outperform the other eight methods. The extremities of the precision vs. recall curve are interesting: At maximum recall where $T_f = 0$, all pixels are retained as positives, i.e., considered to be foreground, so all the methods have the same precision and recall values; precision 0.2 and recall 1.0 at this point indicate that, on average, there are 20% image pixels belonging to the ground truth salient regions. At the other end, the minimum recall values of our methods are higher than those of the other methods, because the saliency maps computed by our methods are smoother and contain more pixels with the saliency value 255.

Saliency cut. We now consider the use of the computed saliency map to assist in salient object segmentation.

Method	IT[17]	MZ[21]	GB[14]	SR[15]	FT[2]	AC[1]	CA[12]	LC[32]	HC	RC
Time(s)	0.611	0.070	1.614	0.064	0.016	0.109	53.1	0.018	0.019	0.253
Code	Matlab	C++	Matlab	Matlab	C++	C++	Matlab	C++	C++	C++

Table 1. Average time taken to compute a saliency map for images in the database by Achanta et al. [2]. Most images in the database (see our project webpage) have resolution 400×300 . Algorithms were tested using a Dual Core 2.6 GHz machine with 2GB RAM.

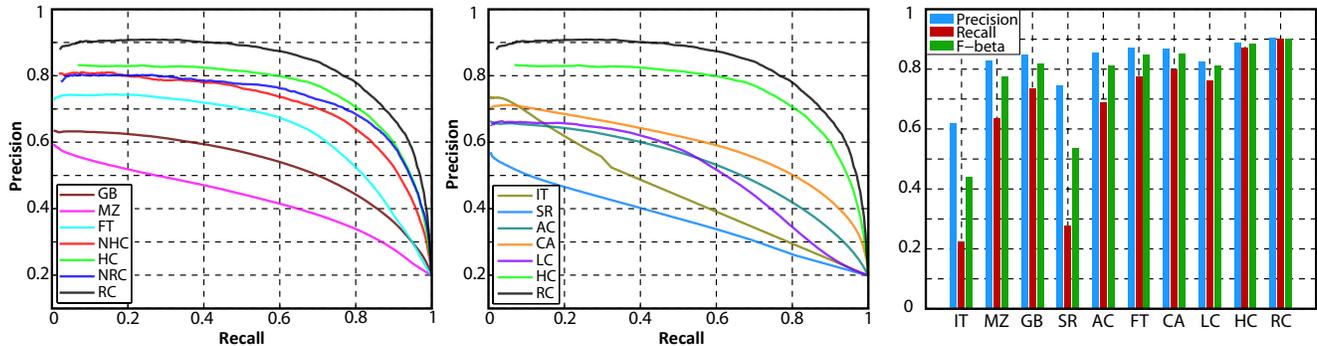


Figure 7. Precision-recall curves for naive thresholding of saliency maps using 1000 publicly available benchmark images. (Left, middle) Different options of our method compared with GB[14], MZ[21], FT[2], IT[17], SR[15], AC[1], CA[12], and LC[32]. NHC denotes a naive version of our HC method with color space smoothing disabled, and NRC denotes our RC method with spatial related weighting disabled. (Right) Precision-recall bars for our saliency cut algorithm using different saliency maps as initialization. Our method RC shows high precision, recall, and F_β values over the 1000-image database. (Please refer to our project webpage for the respective result images.)

Saliency maps have been previously employed for unsupervised object segmentation: Ma and Zhang [21] find rectangular salient regions by fuzzy region growing on their saliency maps. Ko and Nam [18] select salient regions using a support vector machine trained on image segment features, and then cluster these regions to extract salient objects. Han et al. [13] model color, texture, and edge features in a Markov random field framework to grow salient object regions from seed values in the saliency maps. More recently, Achanta et al. [2] average saliency values within image segments produced by mean-shift segmentation, and then find salient objects by identifying image segments that have average saliency above a threshold that is set to be twice the mean saliency value of the entire image.

In our approach, we iteratively apply GrabCut [24] to refine the segmentation result initially obtained by thresholding the saliency map (see Figure 8). Instead of manually selecting a rectangular region to initialize the process, as in



Figure 8. Saliency Cut. (Left to right) Initial segmentation, trimap after first iteration, trimap after second iteration, final segmentation, and manually labeled ground truth. In the segmented images, blue is foreground, gray is background, while in the trimaps, foreground is red, background is green, and unknown regions are left unchanged.

classical GrabCut, we automatically initialize GrabCut using a segmentation obtained by binarizing the saliency map using a fixed threshold; the threshold is chosen empirically to be the threshold that gives 95% recall rate in our fixed thresholding experiments.

Once initialized, we iteratively run GrabCut to improve the saliency cut result. (At most 4 iterations are run in our experiments.) After each iteration, we use dilation and erosion operations on the current segmentation result to get a new trimap for the next GrabCut iteration. As shown in Figure 8, the region outside the dilated region is set to background, the region inside the eroded region is set to foreground, and the remaining areas are set to unknown in the trimap. GrabCut, which by itself is an iterative process using Gaussian mixture models and graph cut, helps to refine salient object regions at each step. Regions closer to an initial salient object region are more likely to be part of that salient object than far-away regions. Thus, our new initialization enables GrabCut to include nearby salient regions, and exclude non-salient regions according to color feature dissimilarity. In the implementation, we set a narrow image-border region (15 pixels wide) to be always in the background in order to avoid slow convergence in the border region.

Figure 8 shows two examples of our saliency cut algorithm. In the flag example, unwanted regions are correctly excluded during GrabCut iterations. In the flower example, our saliency cut method successfully expanded the initial salient regions (obtained directly from the saliency map) and converged to an accurate segmentation result.

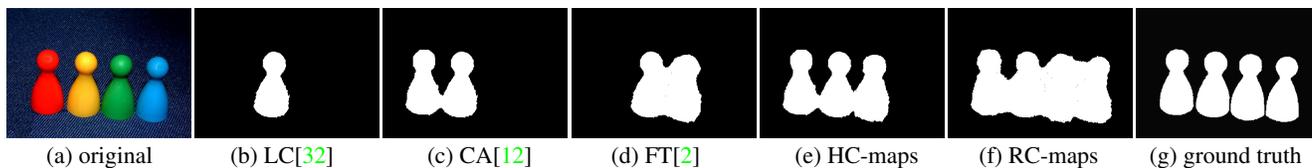


Figure 9. Saliency cut using different saliency maps for initialization. The respective saliency maps are shown in Figure 6.

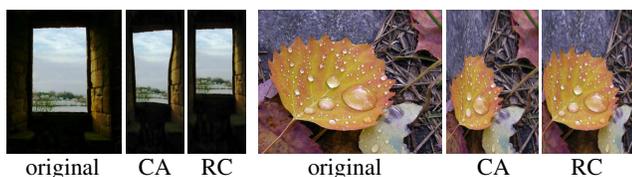


Figure 10. Comparison of content aware image resizing [33] results using CA[12] saliency maps and our RC saliency maps.

To objectively evaluate our new saliency cut method using our RC-map as initialization, we compare our results with results obtained by coupling iterative GrabCut with initializations from saliency maps computed by other methods. For consistency, we binarize each such saliency map using a threshold that gives 95% recall rate in the corresponding fixed thresholding experiment (see Figure 7). A visual comparison of the results is shown in Figure 9. Average precision, recall, and F -Measure are compared over the entire ground-truth database [2], with the F -Measure defined as:

$$F_{\beta} = \frac{(1 + \beta^2) Precision \times Recall}{\beta^2 \times Precision + Recall}. \quad (8)$$

We use $\beta^2 = 0.3$ as in Achanta et al. [2] to weigh precision more than recall. As can be seen from the comparison (see Figures 7-right and 9), saliency cut using our RC and HC saliency maps significantly outperform other methods. Compared with the state-of-the-art results on this database by Achanta et al. (precision = 75%, recall = 83%), we achieved better accuracy (precision = 90%, recall = 90%). (Our demo software is available at our project webpage.)

Content aware image resizing. In image re-targeting, saliency maps are usually used to specify relative importance across image parts (see also [3]). We experimented with using our saliency maps in the image resizing method proposed by Zhang et al. [33]², which distributes distortion energy to relatively non-salient regions of an image while preserving both global and local image features. Figure 10 compares the resizing results using our RC-maps with the results using CA[12] saliency maps. Our RC saliency maps help produce better resizing results since the saliency values in salient object regions are piece-wise smooth, which is important for energy based resizing methods. CA saliency maps, having higher saliency values at object boundaries, are less suitable for applications like resizing, which require entire salient objects to be uniformly highlighted.

²The authors' original implementation is publicly available.



Figure 11. (Middle, right) FT[2] and RC saliency maps are used respectively for stylized rendering [16] of an input image (left). Our method produces a better saliency map, see insets, resulting in improved preservation of details, e.g., around the head and the fence regions.

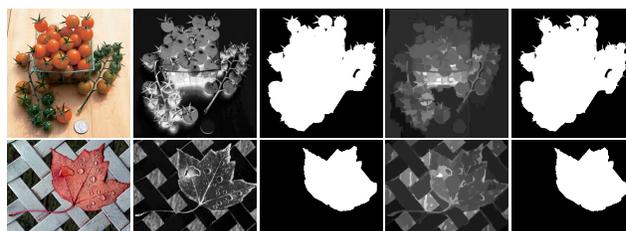


Figure 12. Challenging examples for our histogram based methods involve non-salient regions with similar colors as the salient parts (top), or an image with textured background (bottom). (Left to right) Input image, HC-map, HC saliency cut, RC-map, RC saliency cut.

Non-photorealistic rendering. Artists often abstract images and highlight meaningful parts of an image while masking out unimportant regions [31]. Inspired by this observation, a number of non-photorealistic rendering (NPR) efforts use saliency maps to generate interesting effects [7]. We experimentally compared our work with the most related, state-of-the-art saliency detection algorithm [2] in the context of a recent NPR technique [16] (see Figure 11). Our RC-maps give better saliency masks, which help the NPR method to better preserve details in important image parts and region boundaries, while smoothing out others.

6. Conclusions and Future Work

We presented global contrast based saliency computation methods, namely Histogram based Contrast (HC) and spatial information-enhanced Region based Contrast (RC). While the HC method is fast and generates results with fine details, the RC method generates spatially coherent high quality saliency maps at the cost of reduced computational efficiency. We evaluated our methods on the largest publicly available data set and compared our scheme with eight other state-of-the-art methods. Experiments indicate the

proposed schemes to be superior in terms of both precision and recall, while still being simple and efficient.

In the future, we plan to investigate efficient algorithms that incorporate spatial relationships in saliency computation while preserving fine details in the resulting saliency maps. Also, it is desirable to develop saliency detection algorithms to handle cluttered and textured background, which could introduce artifacts to our global histogram based approach. Finally, it may be beneficial to incorporate high level factors like human faces, symmetry into saliency maps. We believe the proposed saliency maps can be used for efficient object detection [13], reliable image classification, robust image scene analysis [5], leading to improved image retrieval [4].

Acknowledgements. This research was supported by the 973 Program (2011CB302205), the 863 Program (2009AA01Z327), the Key Project of S&T (2011ZX01042-001-002), and NSFC (U0735001). Ming-Ming Cheng was funded by Google PhD fellowship, IBM PhD fellowship, and New PhD Researcher Award (Ministry of Edu., CN).

References

- [1] R. Achanta, F. Estrada, P. Wils, and S. Süsstrunk. Saliency region detection and segmentation. In *ICVS*, pages 66–75. Springer, 2008. 410, 412, 414
- [2] R. Achanta, S. Hemami, F. Estrada, and S. Süsstrunk. Frequency-tuned salient region detection. In *CVPR*, pages 1597–1604, 2009. 409, 410, 412, 413, 414, 415
- [3] R. Achanta and S. Süsstrunk. Saliency Detection for Content-aware Image Resizing. In *ICIP*, 2009. 415
- [4] T. Chen, M.-M. Cheng, P. Tan, A. Shamir, and S.-M. Hu. Sketch2photo: Internet image montage. *ACM Trans. Graph.*, 28(5):124:1–10, 2009. 409, 416
- [5] M.-M. Cheng, F.-L. Zhang, N. J. Mitra, X. Huang, and S.-M. Hu. Repfinder: Finding approximately repeated scene elements for image editing. *ACM Trans. Graph.*, 29(4):83:1–8, 2010. 416
- [6] C. Christopoulos, A. Skodras, and T. Ebrahimi. The JPEG2000 still image coding system: an overview. *IEEE Trans. Consumer Elec.*, 46(4):1103–1127, 2002. 409
- [7] D. DeCarlo and A. Santella. Stylization and abstraction of photographs. *ACM Trans. Graph.*, 21(3):769–776, 2002. 415
- [8] R. Desimone and J. Duncan. Neural mechanisms of selective visual attention. *Annual review of neuroscience*, 18(1):193–222, 1995. 409
- [9] M. Ding and R.-F. Tong. Content-aware copying and pasting in images. *The Visual Computer*, 26:721–729, 2010. 409
- [10] W. Eihhauser and P. Konig. Does luminance-contrast contribute to a saliency map for overt visual attention? *European Journal of Neuroscience*, 17:1089–1097, 2003. 412
- [11] P. Felzenszwalb and D. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2):167–181, 2004. 412, 413
- [12] S. Goferman, L. Zelnik-Manor, and A. Tal. Context-aware saliency detection. In *CVPR*, 2010. 410, 412, 413, 414, 415
- [13] J. Han, K. Ngan, M. Li, and H. Zhang. Unsupervised extraction of visual attention objects in color images. *IEEE TCSV*, 16(1):141–145, 2006. 409, 414, 416
- [14] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In *NIPS*, pages 545–552, 2006. 410, 412, 414
- [15] X. Hou and L. Zhang. Saliency detection: A spectral residual approach. In *CVPR*, pages 1–8, 2007. 410, 412, 413, 414
- [16] H. Huang, L. Zhang, and T.-N. Fu. Video painting via motion layer manipulation. *Comput. Graph. Forum*, 29(7):2055–2064, 2010. 415
- [17] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI*, 20(11):1254–1259, 1998. 409, 410, 412, 414
- [18] B. Ko and J. Nam. Object-of-interest image segmentation based on human attention and semantic region clustering. *J Opt Soc Am*, 23(10):2462, 2006. 409, 414
- [19] C. Koch and S. Ullman. Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, 4:219–227, 1985. 409, 410
- [20] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, T. X., and S. H.Y. Learning to detect a salient object. *IEEE TPAMI*, 33(2):353–367, 2011. 410
- [21] Y.-F. Ma and H.-J. Zhang. Contrast-based image attention analysis by using fuzzy growing. In *ACM Multimedia*, pages 374–381, 2003. 410, 412, 414
- [22] S. K. Mannan, C. Kennard, and M. Husain. The role of visual salience in directing eye movements in visual object agnosia. *Current biology*, 19(6):247–248, 2009. 409
- [23] J. Reynolds and R. Desimone. Interacting roles of attention and visual salience in v4. *Neuron*, 37(5):853–863, 2003. 409
- [24] C. Rother, V. Kolmogorov, and A. Blake. “Grabcut”—Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, 2004. 413, 414
- [25] U. Rutishauser, D. Walther, C. Koch, and P. Perona. Is bottom-up attention useful for object recognition? In *CVPR*, pages 37–44, 2004. 409
- [26] H. Teuber. Physiological psychology. *Annual Review of Psychology*, 6(1):267–296, 1955. 409
- [27] A. M. Triesman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, 1980. 409
- [28] Y.-S. Wang, C.-L. Tai, O. Sorkine, and T.-Y. Lee. Optimized scale-and-stretch for image resizing. *ACM Trans. Graph.*, 27(5):118:1–8, 2008. 409
- [29] J. M. Wolfe and T. S. Horowitz. What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, pages 5:1–7, 2004. 409
- [30] H. Wu, Y.-S. Wang, K.-C. Feng, T.-T. Wong, T.-Y. Lee, and P.-A. Heng. Resizing by symmetry-summarization. *ACM Trans. Graph.*, 29(6):159:1–9, 2010. 409
- [31] S. Zeki. *Inner vision: An exploration of art and the brain*. Oxford University Press, 1999. 415
- [32] Y. Zhai and M. Shah. Visual attention detection in video sequences using spatiotemporal cues. In *ACM Multimedia*, pages 815–824, 2006. 410, 411, 412, 413, 414, 415
- [33] G.-X. Zhang, M.-M. Cheng, S.-M. Hu, and R. R. Martin. A shape-preserving approach to image resizing. *Comput. Graph. Forum*, 28(7):1897–1906, 2009. 409, 415