A framework for visual saliency detection with applications to image thumbnailing

Luca Marchesotti, Claudio Cifarelli and Gabriela Csurka Textual and Visual Pattern Analysis (TVPA) Xerox Research Centre Europe (XRCE), France {Luca.Marchesotti, Claudio.Cifarelli, Gabriela.Csurka}@xerox.com

Abstract

We propose a novel framework for visual saliency detection based on a simple principle: images sharing their global visual appearances are likely to share similar salience. Assuming that an annotated image database is available, we first retrieve the most similar images to the target image; secondly, we build a simple classifier and we use it to generate saliency maps. Finally, we refine the maps and we extract thumbnails. We show that in spite of its simplicity, our framework outperforms state-of-the-art approaches. Another advantage is its ability to deal with visual pop-up and application/task-driven saliency, if appropriately annotated images are available.

1. Introduction

Image thumbnailing consists in the identification of one or more regions of interest in an input image: salient parts are aggregated in foreground regions, whereas redundant and non informative pixels become part of the background. The range of applications where thumbnailing can be employed is broad. It includes traditional problems like image compression, visualization, summarization and more recent applications like variable data printing [24], assisted content creation [3], etc.

Thumbnailing and more generally visual saliency detection are intrinsically challenging problems. In fact, despite the many theories recently formulated [15, 17], it is still not completely clear how the human visual attention processes work. However, all theories seem to agree upon the fact that : subjects selectively direct attention to objects in a scene using both bottom-up, image-based saliency cues and top-down, task-dependent cues.

Bottom-up saliency can be considered task-independent. In fact, if a stimulus is sufficiently salient, it will pop-up from a scene as in Figure 1, image on the left. In this case, saliency is fairly unambiguous. However, in a more clut-



Figure 1. Left, sample image from MRSA dataset [19], saliency is fairly unambiguous. Right, sample from PASCAL dataset [10], in this case, saliency is more complex, ambiguous and application dependent.

tered image (e.g. see Figure 1 on the right), multiple objects make the saliency analysis subjective and more ambiguous. This ambiguity might be reduced if we take into account the final intent of the thumbnailing operation (e.g. a thumbnail containing the person might be salient for albuming applications, whereas the dog might be more relevant for highlighting the query results in image retrieval applications).

We designed a generic framework able to deal with visual pop-up and task-driven saliency, respectively using images annotated with salient thumbnails [19] or salient thumbnails and semantic labels [10].

The framework was built upon a simple idea: images sharing global visual appearance are likely to share similar salient regions. Following this principle, we approach thumbnailing as a learning by example problem, and we show that the visual similarity is advantageous to detect saliency and to build thumbnails. Finally, we show that in spite of its simplicity, our approach outperforms state-of-the art saliency detection methods.

The rest of the paper is organized as follows. Firstly, in section 2 we study the literature around visual saliency and image thumbnailing. Then, we describe in detail in section 3 our framework. Section 4 presents the experimental validations and finally in section 5 we discuss the advantages of the methods and the future challenges.

2. State-of-the-art

Most models for visual saliency detection and thumbnail extraction were inspired by the human visual system and can be grouped in bottom-up, top-down and hybrid approaches.

Bottom-up. Methods falling in this category are stimulus-driven. Then, the idea is to seek for the so-called "visual pop-out" saliency. In fact, human attention is interpreted by some, as a cognitive process that selectively concentrate on the most unusual aspects of an environment while ignoring more common things. To model this behavior, various approaches were proposed such as center-surround operation [16] or graph based activation maps [12]. Gao et al [11] reformulated the "center-surround" hypothesis in a decision theoretic framework. Hou and Zhang in [13] proposed a method based on residual of images in the spectral domain.

Top-down. Top-down visual attention processes are considered driven by voluntary control, and related to the observer's goal when analyzing a scene [32]. Object detection can be seen as a particular case of top-down saliency detection, where the predefined task is given by the object class to be detected [21].

Hybrid. Most of the saliency detection methods are hybrid models leveraging the combinations of the bottom-up and top-down approaches [15, 4, 31, 33]. In general, they are structured in two levels, a top-down layer filters out noisy regions in saliency maps created by the bottom-up layer. In most part of the cases, the top-down component is actually a human face detector [15, 31]. However, Chen et al [4] combined a face and text detector finding optimal solutions efficiently through a branch and bound algorithm. Instead, Wang and Li combines spectral residual for bottom-up analysis with features capturing similarity and continuity based on Gestalt principles [33].

Recent approaches suggest that saliency can be learned from manually labeled examples. Liu et al. in [19] formulate salient object detection as an image segmentation problem, where they separate the salient object from the image background.

Other approaches targeting thumbnailing applications can be found in the state-of-the-art, however they use standard bottom-up saliency maps or they discuss the advantages of auto-crop over simple image resizing methods through user preference experiments [31, 9, 27]. An alternative to thumbnailing has been recently studied in computer graphics [28, 1, 29] and it consists in intelligent rescaling and re-targeting of several relevant regions.

3. The Framework

We assume, the existence of an annotated image database representing a wide variety of subjects, where in



Figure 2. On-line saliency detection. Given an image to thumbnail, we retrieve the K most similar images (see section 3.2) and we train a simple classifier to detect salient (foreground) and nonsalient (background) regions. Saliency maps are generated and thumbnails extracted (see section 3.4).

each image we have salient and non-salient regions manually annotated. Our framework operates in two different phases:

- 1. **Off-line database indexation**. For each image in the database we extract local patches and associated low level descriptors. In the low level feature space, we build a visual vocabulary. We then compute high level image representations for salient and non salient regions. Finally, for each image in the dataset, we store a signature based on the high level representations (see details in section 3.1).
- 2. On-line saliency detection and thumbnail. Given a new image, we apply the steps sketched in Figure 2: (1) we retrieve the K most similar images from the indexed database (see section 3.2), (2) we compute a salient (foreground) and non-salient (background) model, (3) we classify each image patch as salient/non-salient, (4) we propagate the result of classification to pixels generating a saliency map, (5) we refine the map and we extract the thumbnail.

3.1. High Level Visual Features

In the image classification literature, the traditional approach to transform low-level features into high-level representations is the bag-of-visual-words (BOV) [30, 6]. However, Perronnin et Dance [23] showed that Fisher Kernel outperforms bag-of-words in image categorization scenarios. Additionally, Fisher representation was successfully used in image retrieval [5] and semantic segmentation [7]. For this reason, we employed Fisher vectors as high level image descriptors. Given a generative model p with parameters λ , Fisher kernel [23] consists in deriving a fixed-length representation of variable length sample sets $X = \{x_t, t = 1...T\}$ using the following gradient vector: $\nabla \log p(X|\lambda)$. This Fisher representation can be interpreted as the direction in which the parameters λ of the generative model should be modified to best fit the data set X.

If we further assume independence between the samples and using the linearity of the gradient we can write:

$$\nabla \log p(X|\lambda) = \sum_{t=1}^{T} \nabla \log p(x_t|\lambda)$$
(1)

In our case, we employ a Gaussian mixture model (GMM) to build a visual vocabulary in some low level feature space where each Gaussian corresponds to a visual word. Let $\lambda = \{w_i, \mu_i, \sigma_i, i = 1...N\}$ denote the set of parameters of the GMM, where N is the number of Gaussians and w_i , μ_i and σ_i are respectively the weight, the mean vector and the variance vector (representing the diagonal covariance matrix Σ_i of the *i*th Gaussian). The GMM vocabulary is trained using maximum likelihood estimation (MLE) considering all or a random subset of the low level descriptors extracted from the training set.

Given a new low level descriptor x_t , the probability that it was generated by the GMM is $p(x_t|\lambda) = \sum_{i=1}^{N} w_i p_i(x_t|\lambda)$, where:

$$p_i(x_t|\lambda) = \frac{\exp\left\{-\frac{1}{2}(x_t - \mu_i)'\Sigma_i^{-1}(x_t - \mu_i)\right\}}{(2\pi)^{D/2}|\Sigma_i|^{1/2}}$$

The partial derivatives of $\log p(x_t|\lambda)$ according to the GMM parameters can be computed by the following formulas [23]:

$$\frac{\partial \log p(x_t|\lambda)}{\partial \mu_i^d} = \gamma_i(x_t) \left[\frac{x_t^d - \mu_i^d}{(\sigma_i^d)^2} \right], \qquad (2)$$

$$\frac{\partial \log p(x_t|\lambda)}{\partial \sigma_i^d} = \gamma_i(x_t) \left[\frac{(x_t^d - \mu_i^d)^2}{(\sigma_i^d)^3} - \frac{1}{\sigma_i^d} \right] .$$
(3)

where the superscript d denotes the d-th dimension of a vector and $\gamma_i(x_t)$ is the occupancy probability given by:

$$\gamma_i(x_t) = \frac{w_i p_i(x_t)}{\sum_{j=1}^N w_j p_j(x_t)}.$$

Here, we use only the gradient with respect to the mean and standard deviation¹.

The Fisher gradient vector $\mathbf{f}_t = \nabla \log p(x_t|\lambda)$ of the descriptor x_t is by definition the concatenation of the partial derivatives shown in equations (2) and (3) leading to a

2xDxN dimensional vector, where D is the dimension of the low level feature space. While the Fisher representation is in general high dimensional, it can be made relatively sparse, as only a small number of components (relevant Gaussians) have non-negligible values.

From (1) the Fisher Vector of the set of descriptors $X = \{x_t, t = 1...T\}$ is the sum of individual Fisher vectors:

$$\nabla \log p(X|\lambda) = \mathbf{f}_X = \sum_{t=1}^T \mathbf{f}_t \tag{4}$$

3.2. Image Indexation and Retrieval

For each image in the database, we extract a set of local image patches and we label each of them according to their overlap² with the manually annotated salient regions. This leads to two sets of annotated descriptors X^+ and X^- (salient and non-salient). Using equation (4), we compute the two corresponding Fisher vectors f_{X^+} and f_{X^-} and we use them as a pair of signatures for each image.

Given a new image to thumbnail, we retrieve the K most similar images as follows. First, we extract a set of local image patches with their low level descriptors $Y = \{y_1, y_2, ..., y_M\}$. We use the visual vocabulary (GMM) trained off-line for image indexation and we compute a Fisher vector $\mathbf{f}_{\mathbf{Y}}$ using equation (4). To compute the similarities between two images, we use the following normalized L_1 similarity measure:

$$sim(X,Y) = -||\mathbf{\hat{f}}_X - \mathbf{\hat{f}}_Y||_1 = -\sum_i |\hat{f}_X^i - \hat{f}_Y^i| \quad (5)$$

where $\hat{\mathbf{f}}$ is the vector \mathbf{f} normalized to norm L_1 equal 1 $(||\hat{\mathbf{f}}||_1 = 1)$ and $\mathbf{f}_X = \mathbf{f}_{X^+} + \mathbf{f}_{X^-}$ represents the global set of descriptors (salient and non salient) of image X.

3.3. Saliency Detection

For each image to thumbnail, in analogy with [34], we retrieved the K most similar images as described in section 3.2.

Nearest neighborhood images were used in the past for several related problems such as scene classification [8], similar images alignment [18] and object detection [26]. We have seen that for each retrieved image X_j we have two available Fisher vectors: $\mathbf{f}_{X_j^+}$ and $\mathbf{f}_{X_j^-}$ corresponding to the salient and to the non salient regions. We sum all Fisher vectors associated to the K similar images for salient and non-salient regions:

$$\mathbf{f}_{FG} = \sum_{j=1}^{K} \mathbf{f}_{X_j^+} \quad \text{and} \quad \mathbf{f}_{BG} = \sum_{j=1}^{K} \mathbf{f}_{X_j^-} \tag{6}$$

¹According to [23], adding the gradient with respect to the mixture weights does not add significant information.

 $^{^{2}}$ The patch is labeled as salient if its overlap with the salient region is above 50% of its area and non relevant otherwise.

and we call them (abusively) foreground and background Fisher models.

A patch x_i then is considered salient, if its normalized L_1 distance to the foreground Fisher model is smaller than to the background Fisher model:

$$||\hat{\mathbf{f}}_{x_i} - \hat{\mathbf{f}}_{FG}||_1 < ||\hat{\mathbf{f}}_{x_i} - \hat{\mathbf{f}}_{BG}||_1$$
 (7)

However, this classifier is too dependent on a single local patch which makes it locally unstable. Therefore, in order to increase the model's robustness, instead of considering a single patch we sum Fisher vectors of patches over a neighborhood \mathcal{N} :

$$\mathbf{f}_{\mathcal{N}} = \sum_{x_i \in \mathcal{N}} \mathbf{f}_i \tag{8}$$

Furthermore, we replace the binary classifier with nonbinary score which is a simple function of the normalized L_1 distances:

$$s(\mathcal{N}) = ||\hat{\mathbf{f}}_{\mathcal{N}} - \hat{\mathbf{f}}_{FG}||_1 - ||\hat{\mathbf{f}}_{\mathcal{N}} - \hat{\mathbf{f}}_{BG}||_1$$
(9)

The proposed method is based on a simple distance metric and can be replaced by more complex patch classifiers [7]. However, the method proposed here has two main advantages. It is simple and computationally efficient. Moreover, it requires neither pre-trained patch class models, nor the pre-processing of the retrieved images to extract patch descriptors, but it uses directly the pair of image signatures.

Finally, to build a "saliency map" **S** we can consider that each pixel in the neighborhood region \mathcal{N} takes the value $s_{\mathcal{N}} = s(\mathcal{N})$. However, this is not a good strategy especially if we consider overlapping regions. Instead, we assign the value $s_{\mathcal{N}}$ to the center pixel of each region \mathcal{N} and then we either interpolate the values between these centers or we use a Gaussian propagation of these values. The latter can be done by averaging over all Gaussian weighted scores:

$$s(p) = \frac{\sum_{\mathcal{N}} s_{\mathcal{N}} w_{\mathcal{N}}(p)}{\sum_{\mathcal{N}} w_{\mathcal{N}}(p)}$$
(10)

where w_N is the value in pixel p of the Gaussian centered in the geometrical center of each the region N. In our experiments we used a diagonal isotropic covariance matrix with values $(0.6 * R)^2$, RxR being the size of N.

3.4. Map Refinement and Thumbnail Extraction

The aim of this step is to build one or more thumbnails from the saliency map **S**. A straightforward option is to binarize **S** with an appropriate threshold th_{bin} leading to a the binarized saliency map **S**_B. Note that by increasing or decreasing the threshold, we can give more importance respectively to precision or to recall. The drawback of this simple approach is that it does not take into account the contours of the salient object. However, we can overcome this drawback using a segmentation method inspired by [25]. The main idea is to use the saliency map S_B to initialize the Graph-Cut algorithm, then iterate between energy minimization based region labeling and foreground and background GMM updates. First, we choose two thresholds (one positive th_+ and one negative th_-) that separate the saliency map S into 3 different regions: pixels u labeled as salient ($S(u) > th_+$), pixels v labeled as non-salient ($S(v) < th_-$) and unknown (the others). Two Gaussian Mixture Models (GMMs) Ω_1 and Ω_0 are created, one using RGB values of salient (foreground) pixels and one using RGB values of non salient (background) pixels. Then the following energy is minimized:

$$E(L) = \sum_{u \in \mathcal{P}} D_u(u) + \sum_{(u,v) \in \mathcal{C}} V_{u,v}(u,v)$$
(11)

where the data penalty function $D_u(u) = -\log p(u|l_u, \Omega_{l_u})$ is the negative log likelihood that the pixel u belongs to Ω_{l_u} , with $l_u \in 0, 1$ and the contrast term:

$$V_{u,v}(u,v) = \gamma \sum_{(u,v)\in\mathcal{C}} \delta_{l_u,l_v} \exp\left(-\frac{\|u-v\|^2}{2*\beta}\right) \quad (12)$$

with $\delta_{l_u,l_v} = 1$ if $l_u = l_v$, C representing 4-way cliques, and $\beta = \mathbf{E}(||u - v||^2)$ (see further details in [25]).

The energy (11) is minimized using the min-cut/maxflow algorithms proposed in [2] leading to a binary annotation of the image. Using the new labels, we update the two GMM parameters and similarly to [25] iterate between energy minimization (11) and GMM updates until no modifications are made to the binary labels. This binary map can be considered as a new saliency map, denoted by S_G .

Note, that the above minimization methods works well if we have a relatively good initialization, but might fail otherwise. Therefore, we want to keep the refined map S_G only when the risk to deteriorate S_B is low. Seen that (a) we cannot directly estimate this risk and (b) Graph Cut does not reuse information about saliency during iterations, we introduce a simple decision mechanism to choose a posteriori between S_B and S_G : if the overlap between the two maps is above a certain threshold th_d (i.e. we did not diverge too much from initialization), we choose S_G otherwise keep S_B :

$$\mathbf{S}^* = \begin{cases} \mathbf{S}_G & \text{if } \frac{\mathbf{S}_B \cap \mathbf{S}_G}{\mathbf{S}_B \cup \mathbf{S}_G} > th_d \\ \mathbf{S}_B & \text{otherwise} \end{cases}$$

with $0 < th_d < 1$ (set to $th_d = 0.1$ in our experiments).

Finally, different strategies can be designed to extract a thumbnail from the binary map: we can select the biggest, most centered salient region as thumbnail or, alternatively, all the detected salient regions and re-target them into a single thumbnail as proposed in [28].

4. Experimental Results

The main objectives of the experiments were (1) to show that image similarity is advantageous for detecting saliency, (2) to prove that our approach can compete with state-of-the-art methods and (3) to test the behavior of the framework in more challenging scenarios where saliency is "target/application-dependent".

4.1. The datasets

We used two state-of-the-art datasets with thumbnails manually annotated by one or several users:

- MRSA. The dataset described in [19] is composed of two parts: Part A (15000 natural images) and part B (5000 images) with good variety of different subjects. Only Part B was used, as it is the only one currently made available by the authors. The annotation are highly consistent with generally small variance over 9 thumbnails. This dataset was used to test the performances of the framework in the case of "visual pop-up" saliency (i.e unambiguous and objective thumbnails). In fact, MRSA is composed by images containing a wide range of different objects, nevertheless in most cases there is a single object per image that "pops-out" from a relatively simple background. Ground truth data consist of 9 different thumbnails annotated for each image by 9 different users (see Figure 1 on the left). The users manually selected a thumbnail containing the region of interest, which is typically represented by a full object or, in some cases by a subpart of the object (e.g. head of the dog). In average, the MRSA thumbnails represent the 35% of the total area of the image and the distance of their center of mass from the center of the image is within 42 pixels.
- PASCAL VOC. To quantify the performances in the case of "targeted" visual saliency, we opted for PAS-CAL Visual Object Classes Challenge data [10]. Indeed, no other dataset is currently available for this kind of analysis. Moreover, PASCAL has some features which fit well our purpose: (1) Images are annotated (see Figure 1 on the right) with multiple thumbnail regions as well as thumbnail labels for 20 object classes (person, bird, cat, cow, dog, horse, sheep, aeroplane, bicycle, boat, bus, car, motorbike, train, bottle, chair, dining table, potted plant, sofa, tv monitor). (2) The images are much more challenging with respect to visual complexity; they contain multiple, ambigous, often small objects and very cluttered backgrounds. In our experiments, we used the VOC 2008 Trainval dataset to build our indexed database and the VOC 2007 Test Dataset to perform the tests. The main reason for this choice was to have independent train-

ing and test set and hence to avoid any bias in the performances which might be caused by near duplicate images (especially as our method is based on nearest neighbor search).

4.2. Experimental Set-up

In all the experiments, we used two types of low level descriptors: SIFT-like gradient orientation histograms [20] and simple local color statistics (RGB mean and variations). They were computed on local patches extracted on multiscale grid. The dimensionality of the original features were subsequently reduced to 50 by PCA. A GMM composed by 32 Gaussian was computed in both PCA projected feature spaces leading to two visual vocabularies. Hence, for each patch, we computed two separate Fisher vectors: one for SIFT features and the second for color features. Finally, the two Fisher Vectors were normalized and concatenated. The saliency maps S were created by setting the dimensions of the neighborhood \mathcal{N} to 50x50 pixels (close to the average patch size).

To evaluate the performance we followed the startegy proposed in [19], and we computed precision, recall, F_{α} with $\alpha = 0.5$ and BDE (Bounding Box Displacement Error) [22].

In the case of the MSRC database, we used a leave-oneout startegy and then we averaged the performances over the whole dataset. In PASCAL, we used independent training and test data and the average was evaluated per category.

4.3. Detecting saliency through visual similarity

First, we tested the performances of the saliency detector by varying the parameter K (number of most similar images considered). To quantify the advantage of visual similarity in cases of visual pop-up saliency, we performed the following test: instead of taking into account the K most similar images, we selected in the indexed database, K random images.

These results are plotted in Figure 3, the upper curve shows F_{α} with K nearest neighbor images and the bottom curve F_{α} with K random images. First, we notice that for all tested values of K, we improve the performance using visual similarity. As expected, we increase in both cases the performance by increasing K. While we have large improvement for low Ks, after K = 20 the performance increases less. Beyond K = 50 the improvement becomes negligible. Moreover, we know that even if both curves continue to slightly increase after 100, they begin to decrease with large Ks arriving both to $F_{\alpha} = 0.64$ for K = 4999(i.e. all dataset except the test image). As we didn't tested K bigger that 100, we do not have the ideal K (max of the curve). Nevertheless, this is not a very significant value as it will vary from database to database. In fact, what is more important is to find a good compromise between accuracy



Figure 3. Performances of the saliency detector by varying the parameter K, using nearest neighbor images and random images.



Figure 4. Thumbnail extraction strategies compared: Pixel Connectivity only, Pixels Connectivity + Graph-Cut.

and cost. To this end, we believe that K = 50 is a good choice and we used it in the rest of the experiments.

Finally, we evaluated the refinement of the saliency maps through Graph-Cut. In figure 4, we plot F_{α} obtained by varying the binarization threshold³ th_{bin} from 0.1 to 1 with step 0.1. As expected, Graph-Cut combined with pixel connectivity provides better results for all values of th_{bin} with a global maximum for $th_{bin} = 0.6$. We used this parameterization in the rest of the tests.

In Figure 5 and 6 we show some qualitative results obtained on the MRSA dataset.



Figure 5. Some qualitive results collected in MRSA dataset obtained using Graph-Cut refinement.



Figure 6. Some qualitive results collected in MRSA dataset where the decision mechanism rejected Graph-Cut refinement.

4.4. Comparison with State-of-the-art methods

We compared our framework with three state-of-the-art methods designed for saliency and thumbnail detection:

- **ITTI**: a classical approach based on Itti's method [14], that leverages a neuromorphic model simulating which elements are likely to attract visual attention⁴,
- **SR**: a more recent method described in [13] based on the analysis of the residual of an image in the spectral domain⁵.
- **CRF**: a learning method proposed in [19], based on Conditional Random Field.

As the first two methods evaluate only saliency maps, we extracted thumbnails to make the comparison possible. In particular, we applied the thumbnail strategy described in section 3.4 and we chose the most performing parameter configurations for our method (see section 4.3). For the third method, we directly reported the results given by [19] on the same dataset.

³As the range of the scores vary in the images, we first normalize the value to be between 0 and 1. This allows to select the same threshold for different images.

⁴We used the Matlab implementation available at http://www.saliencytoolbox.net/

⁵Matlab implementation available at http://bcmi.sjtu.edu.cn/ houxiaodi



Figure 7. Above, comparison of our method with (SR), (ITTI) and (CRF) using precision, recall, F_{α} measures. Below, comparison of the same methods using BDE (Bounding Box Displacement Error).

As Figure 7 (above) shows, our method outperforms **ITTI**) and (**SR**) with a consistent margin on Precision, Recall and F_{α} -measure, and it gives slightly better results than (**CRF**). Comparisons using BDE (see Figure 7, below) show similar behavior.

4.5. Target-driven Visual Saliency

The experiments were performed on each target class c as follows: 1. Among the training samples containing at least one object of the class c, we retrieved the K most similar images (K=50). 2. When we build the foreground/background model, we consider as salient only the bounding boxes labeled with the class c. 3. We average the results taking into account the labels available for the test images. This is motivated by the fact that in many applicative scenarios such as variable data printing or image retrieval, we only need to localize the most representative thumbnail for the targeted class . 4. We apply the same strategy used for processing the MRSA dataset to extract the thumbnails.

Figure 8 shows the quantitative results of the experiments performed on the PASCAL dataset with and without Graph-Cut. Instead, in Figure 9 we display examples of qualitative results (with Graph-Cut). Not surprisingly, the performances are in general lower in comparison to MRSA. In fact, PASCAL is a more challenging database character-



Figure 8. Performances (F_{α}) in the PASCAL dataset [10] with and without Graph-Cut).



Figure 9. Some qualitive results collected in PASCAL dataset.

ized by high visual complexity⁶. Also, Graph-Cut improves only a few categories because we have in most cases cluttered background and occluded objects.

In this experiment, we cannot directly compare the results with other methods as there is no available benchmark data or quantitative results reported in the literature for this specific problem. Object detection could be considered as the closest scenario, however the comparison would be unfair because here we solve a different problem. Indeed, while object detection aims at detecting and localizing every object individually, target-driven thumbnailing is less restrictive: it consists in higlighting region(s) containing one or multiple instance of the targeted objects.

5. Conclusions

We proposed a framework for image thumbnailing based on visual similarity. The underlying assumption was that images sharing their global visual appearance are likely to share similar salience. Through exhaustive experimental setup we showed that in spite of its simplicity, the framework achieves satisfactory results with respect to other state-of-the art methods. In addition we showed that the

⁶This is also confirmed by the very low detection and segmentation results (about 25% accuracy) reported in the literature on this dataset.

framework can be used in various thumbnailing scenarios, based on "visual pop-ups" and "target oriented" salency. At present, we took into account only natural images, however the framework is generic enough to be applicable to other type of images such as medical or document images providing that appropriate training data is available.

Acknowledgments

This work was supported by the Agence Nationale de la Recherche through the OMNIA project (ANR-07-MDCO-009-02).

References

- [1] S. Avidan and S. Avidan. Seam carving for content-aware image resizing. In *SIGGRAPH*, 2007.
- [2] Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *PAMI*, 26, 2004.
- [3] M. Bressan, G. Csurka, Y. Hoppenot, and J. Renders. Travel blog assistant system. In *Proceedings of the International Conference on Computer Vision Theory and Applications.*, 2008.
- [4] L.-Q. Chen, X. Xie, X. Fan, W.-Y. Ma, H.-J. Zhang, and H.-Q. Zhou. A visual attention model for adapting images on small displays. *ACM Multimedia Systems Journal*, 9(4), 2003.
- [5] S. Clinchant, J.-M. Renders, and G. Csurka. Trans-media pseudo-relevance feedback methods in multimedia retrieval. In Advances in Multilingual and Multimodal Information Retrieval, volume LNCS 5152, 2008.
- [6] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray. Visual categorization with bags of keypoints. In ECCV Workshop on Statistical Learning for Computer Vision, 2004.
- [7] G. Csurka and F. Perronnin. A simple high performance approach to semantic segmentation. In *BMVC*, 2008.
- [8] S. K. Divvala, A. A. Efros, and M. Hebert. Can similar scenes help surface layout estimation?. In *IEEE Workshop* on Internet Vision, at CVPR'08., 2008.
- [9] B. Erol, K. Berkner, and S. Joshi. Multimedia thumbnails for documents. In *Proceedings of ACM Multimedia*, 2006.
- [10] M. Everingham, A. Zisserman, C. Williams, and L. V. Gool. The pascal visual object classes challenge. http://pascallin.ecs.soton.ac.uk/challenges/VOC/.
- [11] D. Gao, V. Mahadevan, and N. Vasconcelos. The discriminant center-surround hypothesis for bottom-up saliency. In *NIPS*, 2007.
- [12] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In Advances in Neural Information Processing Systems 19, pages 545–552. MIT Press, 2007.
- [13] X. Hou and L. Zhang. Saliency detection: A spectral. residual approach. In CVPR, 2007.
- [14] L. Itti and C. Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40, 2000.

- [15] L. Itti and C. Koch. Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2(3), 2001.
- [16] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *PAMI*, 20(11), 1998.
- [17] S. W. Kienzle, F. A. Wichmann, B. Scholkopf, and M. O. Franz. A nonparametric approach to bottom-up visual saliency. In *NIPS*, 2006.
- [18] C. Liu, J. Yuen, A. Torralba, J. Sivic, and W. T. Freeman. Sift flow: dense correspondence across different scenes. In *European Conference on Computer Vision (ECCV) 2008.*, 2008.
- [19] T. Liu, J. Sun, N. Zheng, X. Tang, and H. Shum. Learning to detect a salient object. In CVPR, 2007.
- [20] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [21] J. Luo. Subject content-based intelligent cropping of digital photos. In *IEEE International Conference on Multimedia* and Expo, 2007.
- [22] D. R. Martin, C. C. Fowkles, and J. Malik. Learning to detect natural image boundaries using local brightness, color and texture cues. *PAMI*, 26(5), 2004.
- [23] F. Perronnin and C. Dance. Fisher kernels on visual vocabularies for image categorization. In CVPR, 2007.
- [24] F. Romano. An investigation into printing industry trends., 2004.
- [25] C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *SIGRAPH*, 2004.
- [26] B. C. Russell, A. Torralba, C. Liu, R. Fergus, and W. T. Freeman. Object recognition by scene alignment. In Advances in Neural Information Processing Systems., 2007.
- [27] R. Samadani, T. Mauer, D. Berfanger, J. Clark, and B. Bausk. Representative image thumbnails: Automatic and manual. Technical Report HPL-2008-6, HP Laboratories Palo Alto, January 2008.
- [28] V. Setlur, S. Takagi, R. Raskar, M. Gleicher, and B. Gooch. Automatic image retargeting. In *Mobile and Ubiquitous Multimedia (MUM)*, 2005.
- [29] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani. Summarizing visual data using bidirectional similarity. In CVPR, 2008.
- [30] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *ICCV*, volume 2, pages 1470–1477, 2003.
- [31] B. Suh, H. Ling, B. Bederson, and D. Jacobs. Thumbnail cropping and its effectiveness. In ACM User Interface Software and Technology, 2003.
- [32] V. Sundstedt, A. Chalmers, K. Cater, and K. Debattista. Topdown visual attention for efficient rendering of task related scenes. In *In Vision, Modeling and Visualization*, pages 209– 216, 2004.
- [33] Z. Wang and B. Li. A two-stage approach to saliency detection in images. In *ICASSP*, 2008.
- [34] H. Zhang, A. C. Berg, M. Maire, and J. Malik. Svm-knn: Discriminative nearest neighbor classification for visual category recognition. In *ICCV*, 2006.