Intelligent Ad Resizing

Anthony P. Badali
abadali@comm.utoronto.ca

Parham Aarabi
parham@ecf.utoronto.ca

Ron Appel
ron@comm.utoronto.ca

Department of Electrical & Computer Engineering, University of Toronto, Toronto, ON, Canada

Categories and Subject Descriptors

H.4.0 [**Information Systems**]: General

General Terms

Algorithms

1. INTRODUCTION & BACKGROUND

Over the last two decades, the World Wide Web has grown drastically as a market, with image and text advertising as one of the leading revenue streams. In practice, a web publisher designates a rectangular region, and an advertiser designs an ad to fit that region. This burdens both the publisher, potentially having to waste ad-space, and the advertiser, having to generate different versions of the same ad to fit various dimensions. An ideal solution would accommodate any ad to fit any region, further allowing for the monetization of currently unusable ad-spaces within websites. This work investigates such a solution: the automatic generation of visual banners of any size, given a single prototype ad. Specifically, the contributions of this work are an optimization framework for determining resized ad parameters as well as a novel saliency map algorithm for estimating regions of importance within an image, which outperformed conventional methods by over 20% in a user study.

The automatic resizing of image-and-text ads is closely related to the general problem of automatic image resizing (AIR), an area that has received much attention in recent years. The resizing of an image consists of the construction of a saliency map of the image denoting the important areas therein, and a method of pixel reduction such as rectangular cropping [2][7], seam removal [5][1], and even object "cutting and pasting" [6]. The majority of prior resizing algorithms employ Itti and Koch (IK) to generate their saliency maps [3] based on a heuristic model of low-level human vision. Alternatively, Santella et. al. proposed a semi-automated cropping technique which used eye tracking data to improve saliency estimates [5]. In general, AIR objective functions maximize the saliency of the remaining regions [7].

2. AUTOMATIC POSTER RESIZING

Formally, this problem is defined as follows: given an initial poster (an image overlaid with text of a specified size and position), a new poster must be generated, with new dimensions and appropriate text size and position, while still maintaining the contextual information conveyed by the initial poster. The proposed approach consists of generating a saliency map, resizing the poster according to the generated map, and finally, placing the text in the least obtrusive location.

The generation of the saliency map consists of four main stages. Firstly, a representative set of pixels is randomly sampled from regions that are assumed a priori to be part of the background (such as peripheral pixels). This set is compared against the entire image, generating an initial estimate of which pixels are part of the background. Secondly, these samples are modeled as a simple (tri-modal) mixture of Gaussians, with the resulting means representing centers of importance in the image. Thirdly, an energy map is generated by taking into account aspects such as distance from the means and feature similarity with neighboring pixels. Fourthly, the final saliency map is generated by filtering out regions that do not surpass a threshold of minimal gradient in the initial image. This is a very high-level description; for more technical details, please refer to the full paper at http://ronappel1.com/papers/IAR.pdf. Figure 1 shows sample images and saliency maps generated by the Itti Koch method as well as the proposed method.

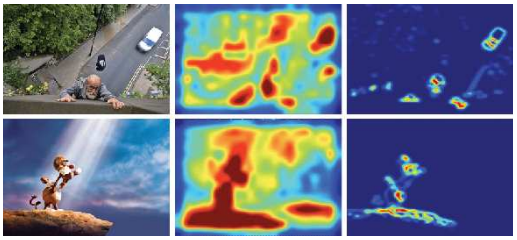

Figure 1: (Left) Original Image, (Center) IK Saliency Map, (Right) CS Energy Map

Once the saliency (importance) map is computed, a poster of any size can be generated, using the following resizing framework. The initial image is cropped (in either the vertical or horizontal direction) to the desired aspect ratio by searching for a window that retains the largest region of importance. Uniform scaling is then used to fit any exact dimension. Within this output window, the size and placement of the text is determined, this time, with the objective of covering the region of least importance. A block diagram of the overall procedure is given in Figure 2.

(a)

Method	Selection
APR-CS	69%
APR-G	60%
APR-IK	50%
Center	45%
Saliency Sum	45%
Seam Carving	29%

(b)

		Method Compared Against					
		APR-CS	APR-G	APR-IK	CEN	SSUM	SC
Winning %	APR-CS	-	58%	69%	67%	70%	80%
	APR-G	42%	-	60%	64%	60%	73%
	APR-IK	31%	40%	-	57%	56%	67%
	CEN	33%	36%	43%	-	51%	68%
	SSUM	30%	40%	44%	49%	-	65%
	SC	20%	27%	33%	32%	35%	-

Table 1: User Study Results: (a) Overall Selection Percentages, (b) Head-to-Head Selection Percentages

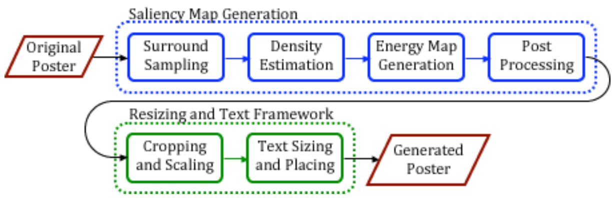

Figure 2: Block diagram of the overall procedure

3. RESULTS & DISCUSSION

Figure 3 shows the results of automatic poster resizing in the horizontal and vertical directions, and for various scaling factors. Overall, the technique yields satisfactory results in most cases, with resizing problems attributable to the low-level nature of the importance map generation. In all cases, the degree of resizing is limited by the image content, but despite such limitations, the results do show promise for automatic poster resizing applications.

In evaluating these results, a user study was conducted in a forced-choice format [4][5] where participants were shown a series of poster pairs and asked to select the better (in terms of image content, text position, and text size). All together, posters were resized according to six different methods.

Three methods were based on the proposed resizing and text-placing framework (differing only in saliency map generation method): APR-IK, using IK saliency as importance map; APR-G, using squared image gradient as importance map; and APR-CS, using the described CS energy-based importance map.

The remaining three methods were extensions of popular approaches for standard image resizing: CEN, simply cropping the image around its center; SSUM, cropping the image to the region of highest IK saliency; and SC, seam-carving the image using a gradient energy function. For these alternative methods, the text size and position were determined as a relative scaling of the original size and position.

The study consisted of 12 posters (based on typical advertisements), each resized to 6 different sizes, using the 6 different methods, for a total of 432 posters. 26 participants made approximately 2600 comparisons between pairs of posters. The results of the study are given in Table 1.

The APR methods decisively outperformed the alternatives, with APR-CS as the top performer by a significant margin, in both the absolute and head-to-head results. This suggests that the text relevance component alone improves resizing performance.

Figure 3: Sample Posters Resized Vertically and Horizontally

4. REFERENCES

[1] S. S. A. Avidan. Seam carving for content-aware image resizing. *ACM Trans. Graph.*, 26(3):10, 2007.

[2] L. Chen, X. Xie, X. Fan, W. Ma, H. Zhang, and H. Zhou. A visual attention model for adapting images on small displays. *ACM Multimedia*, 2003.

[3] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE PAMI*, 20(11):1254–1259, 1998.

[4] F. Liu and M. Gleicher. Automatic image retargeting with fisheye-view warping. In *ACM U.I.S.T. '05*, pages 153–162, New York, NY, USA, 2005.

[5] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, and M. Cohen. Gaze-based interaction for semi-automatic photo cropping. In *CHI '06*, pages 771–780, New York, NY, USA, 2006.

[6] V. Setlur, S. Takagi, R. Raskar, M. Gleicher, and B. Gooch. Automatic image retargeting. In *MUM '05: Proceedings of the 4th international conference on Mobile and ubiquitous multimedia*, pages 59–68, New York, NY, USA, 2005.

[7] B. Suh, H. Ling, B. B. Bederson, and D. W. Jacobs. Automatic thumbnail cropping and its effectiveness. In *UIST '03*, pages 95–104, New York, NY, USA, 2003.