

Saliency Retargeting: An Approach to Enhance Image Aesthetics

Lai-Kuan Wong Kok-Lim Low

School of Computing, National University of Singapore
{lkwong, lowkl}@comp.nus.edu.sg

Abstract

A photograph that has visually dominant subjects in general induces stronger aesthetic interest. Inspired by this, we have developed a new approach to enhance image aesthetics through saliency retargeting. Our method alters low-level image features of the objects in the photograph such that their computed saliency measurements in the modified image become consistent with the intended order of their visual importance. The goal of our approach is to produce an image that can redirect the viewers' attention to the most important objects in the image, and thus making these objects the main subjects. Since many modified images can satisfy the same specified order of visual importance, we trained an aesthetics score prediction model to pick the one with the best aesthetics. Results from our user experiments support the effectiveness of our approach.

1. Introduction and Motivation

A photograph that has visually dominant subjects in general induces stronger aesthetic interest. In the study of photography and aesthetics, Wollen [19] revealed that photographers deliberately avoid uniform sharpness of focus and illumination as an approach to achieve higher image aesthetics. This approach is based on the basis that our eyes are attracted to salient elements that are acutely sharp, bright or colorful in images. Figure 1 shows examples of how professional photographers utilize contrast in sharpness, lighting, and color to bring out the visual dominance of subjects so that the viewer is directed to where the photographers intended. However, due to the limited selective focusing capability of the imaging device, unfavorable lighting, and the infeasibility to exclude distracting background, it may be difficult, especially for casual photographers, to make the main subjects visually salient. This results in prolonged searching for the subjects in a photograph, leading to reduced satisfaction of viewing and thus, poorer aesthetics experience.

In this paper, we propose a new approach to enhance image aesthetics through *saliency retargeting*. The key idea of saliency retargeting is to alter image features of the objects in the photograph such that their computed



Figure 1: Visual dominance of the subject can be achieved using (a) acutely sharp focus, (b) lighting contrast, and (c) color contrast. Images courtesy of Roie Galitz [21].

saliency measurements in the modified image become consistent with the user-intended order of their visual importance. Our method generates many such modified images that satisfy the specified order of importance, and uses an improved aesthetics score prediction model to pick the one with the best aesthetics. The goal is to produce a maximally-aesthetic version of the input image that can redirect the viewers' attention to the most important objects in the image, and thus making these objects the main subjects. This is useful for enhancing photographs that do not have any obvious main subjects, or for photographs that one wishes to swap the role of the main subject with some other objects.

Figure 2 shows a simple result from our method (for accurate viewing of the images in this paper, the reader is advised to calibrate the display device for the sRGB color space). Figure 2(a) is the original image, in which the intended subject (the fish) does not stand out due to the distracting background, particularly the corals at the bottom of the image. This observation is supported by the saliency map in Figure 2(d), where a higher intensity corresponds to higher saliency measurements. Figure 2(f) is a segmentation map that is provided to our algorithm to indicate the boundaries of the objects and also for specifying the intended order of importance of these objects, where Object A (fish) should have higher importance than Object B (background). With the inputs, our algorithm then automatically applies non-global changes to three low-level image features for each segment in the input image to produce the result image in Figure 2(c), using our approach of saliency retargeting. The shift of saliency to the intended subject is evident in the resulting saliency map in Figure 2(g). In the resulting image, the saliency of the background has been




Figure 2: (a), (d) Original image and its saliency map. (b), (e) Globally-enhanced image and its saliency map. (c), (g) Image enhanced by saliency retargeting and its saliency map. (f) Object segments, where Objects A and B are in decreasing order of importance.

suppressed, making it less distracting, and the fish has become more salient, making it the most dominant subject. Image in Figure 2(b) is globally enhanced using level operation, saturation increase and unsharp mask. Its saliency map in Figure 2(e) is almost identical to the saliency map of the original image. It is generally difficult to use global image enhancements to alter the relative saliency among objects. A user study has shown that 84.4% of 32 viewers chose the saliency retargeted image over the original image as the more aesthetically pleasing image.

Our method also supports the change of visual importance of sub-parts of objects. For example, the features of the face in Figure 3(a) look flat due to unfavorable lighting. Based on the given order of importance of the regions marked by the mask in 3(c), our algorithm retargets the saliency of image 3(a) to produce image 3(b), making the facial features more dominant.

The primary application of our new approach is to enable casual photographers and novice photo-editing users to more easily improve their photographs. Casual photographers often take photographs that do not have clear subjects or have the wrong objects being the subjects. They are generally aware that the photographs are not good, but they often do not know what is wrong and do not know how to enhance the photographs through digital photo editing. However, they know very well which objects are the intended subjects in the photographs, and this information is all that is required by our method to enhance the photographs. For more advanced users, they can use this approach for advanced image editing that involves editing of sub-parts of objects in the image. This may result in many object segments being chosen and it is non-trivial for users to find the right combination of

image modifications of the object segments to achieve the desired saliency. Our approach becomes even more useful for batch image enhancements.

2. Related Work

Many existing techniques for automatic image enhancement, including contrast enhancement [13, 17], color correction and balancing [3, 11, 14], and edge sharpening [7, 12], alter only the global features or local features based on neighborhood information, without considering the content or subjects of the image. In addition, these methods are targeted to deal with image degradations rather than to enhance image aesthetics.

Inspired by the emergence of computational visual attention models [2, 4, 6] that simulate the human visual system to identify regions of interests (ROIs) in an image, Su et al. [16] first proposed the idea of altering saliency of an image by reducing the background saliency to redirect attention to the main subject. Their method utilizes texture power maps to de-emphasize texture variations to decrease the saliency of distracting regions. This method preserves key features, but while adding white noise maintains overall graininess, the resulting images appear too noisy and do not seem to be more aesthetically pleasing than their originals.

A few attempts have been made to improve image aesthetics based on photographic composition rules. Barnajee [1] pioneered this area of research by developing an optical-digital framework for in-camera automation of photographic composition rules. Out-of-focus blur information in a supplementary image is used to determine the ROIs of the image and to select the composition rules (rule of thirds, background blurring, and mitigation of merger) to re-compose the image. However, this approach has limited practicality as it only works for simple images with low depth of field. In another attempt, Kao et al. [8] used a visual attention model [6] and a composition scoring system to perform automatic re-composition using a limited set of composition rules that consists of only rotational correction and cropping. A more promising work [9] employs a compound operator of crop-and-retarget to change the relative position of salient regions in the image, thus modifying the composition aesthetics of




Figure 3: (a) Original image. (b) Enhanced image based on the mask in (c) with the order of importance given to the segments.

the image. A maximally-aesthetic version of the input image is produced using particle swarm optimization and the aesthetics of an image is evaluated based on a selected set of composition rules. While this approach can produce good results, it works well for images with subjects that are sufficiently salient but not otherwise. To bring out the dominance of the subjects, merely changing the relative position may not be enough to redirect viewers’ attention to the important subjects. Alteration to low-level features of the subjects is needed and we believe our saliency retargeting approach can complement existing re-composition methods to further enhance image aesthetics.

3. Our Approach

For an image to be enhanced using our approach, the user first has to create a segmentation map to partition the image into object segments. Fully automatic segmentation is not appropriate in this case because the input images may have non-salient main subjects, causing automatic segmentation to have high failure rates in identifying the object segments. Thus, we provide an interactive interface that incorporates GrabCut [15] for user to make trivial selection to partition the image into object segments. Next, the user specifies the *order of importance* for the object segments, where internally the object segments will be assigned unique *importance values* that accord with the order of importance.

With the above inputs, our method then modifies the input image to produce an enhanced image that satisfies the specified order of importance and has increased aesthetics. Our method consists of two major components—a *saliency retargeting process* that modifies an image and an *aesthetics score prediction model* that measures image aesthetics.

The core of the *saliency retargeting process* is an optimization algorithm that finds a set of image modifications to produce an image such that the saliency measurements of the object segments “match” their target importance values. In this work, only three low-level features in each object segment are modified, namely, luminance, color saturation, and sharpness. The reason for the choice of these features is that they correspond directly to the main features used in the widely-used Itti-Koch saliency model [6], which we use to compute the image saliency. Figure 4 shows the effects of the image modifications on the conspicuity maps that collectively form the saliency map in Figure 2. This example shows that our saliency retargeting algorithm has successfully altered the intensity, color contrast and orientation conspicuity maps through local modifications to the luminance, color saturation and sharpness respectively for each object segment, leading to the increase in dominance of the fish.

To enable aesthetics maximization, we trained an *aesthetics score prediction model* to measure image




Figure 4: Effects of the image modifications on the conspicuity maps. (a)-(d) Original image and its thresholded conspicuity maps. (e)-(h) Enhanced image produced by our saliency retargeting algorithm and its thresholded conspicuity maps.

aesthetics. Our aesthetics maximization algorithm first generates a set of vectors where each contains a different sequence of importance values that accords with the user-specified order of importance for the object segments. Each vector of importance values is then passed to our saliency retargeting algorithm to generate an enhanced image in which the average saliency values of the object segments are consistent with the respective importance values. Finally, our aesthetics score prediction model computes an aesthetics score for each enhanced image in this resulting image set and return the maximally-aesthetics version as the result. Figure 5(c) and 5(d) shows two of the enhanced images produced by our saliency retargeting algorithm for different vectors of important values generated based on the user-specified importance map in Figure 5(b). Both enhanced images have higher aesthetics score than the original image in Figure 5(a). Image 5(d), the image with the highest aesthetics score is returned as the result enhanced image. Algorithmic details of our approach are described in the following subsections.

3.1. Saliency Retargeting

Given an image I and a set of N object segments with target importance value T_i for each object segment i , we aim to enhance the aesthetics of the image by applying a set of low-level image modifications x to the input image I to produce an output image with saliency value M_i that “matches” the target importance value T_i for each object segment i . We formulate the saliency

retargeting process as a *constrained optimization problem* as follows:

$$\min_x f(x) = \sum_{i=1}^N |q(T_i) - q(M_i)| \quad (1)$$

where

- object segment $i=1$ is the most important object segment while $i=N$ is the least important object segment,
- $x = \{v_i, s_i, \sigma_i \mid i = 1, 2, \dots, N\}$ and $|x| = 3N$,
- v_i is the increase of average luminance in object segment i ,
- s_i is the increase of average color saturation in object segment i ,
- σ_i is the increase of sharpness in object segment i (positive σ represents image sharpening, and negative σ image blurring), and
- $q(\cdot)$ is a normalization function defined as

$$q(X_i) = \frac{X_i}{\sum_{j=1}^N X_j}. \quad (2)$$

Each region saliency M_i is the average saliency value within object segment i in the output image. A higher saliency value corresponds to higher saliency.

To avoid intensity and color saturation reversal, the image modifications $\{v_1, s_1\}$ for the most important object segment 1 are set to be positive. In addition, the constraint on the sharpness parameter σ_i is based on the assumption that the more important object segment should be sharper than or as sharp as the less important object. Therefore, the optimization is subject to the set of *inequality constraints*

$$\begin{aligned} v_1 &\geq 0, \quad s_1 \geq 0, \text{ and} \\ \sigma_i &\geq \sigma_j \text{ iff } i \leq j. \end{aligned} \quad (3)$$

In addition, we define a set of *bound constraints* to ensure that the image is not over-modified resulting in image unnaturalness:

$$-v_B \leq v_i \leq v_B, \quad -s_B \leq s_i \leq s_B, \quad -\sigma_B \leq \sigma_i \leq \sigma_B. \quad (4)$$

3.1.1 Implementation

We employed Sequential Quadratic Programming (SQP) [5] to solve our saliency retargeting problem. Given the optimization problem defined in Eq. (1), SQP tries to solve the Lagrangian function

$$L(x, \lambda) = f(x) + \sum_{j=1}^m \lambda_j \cdot G_j(x), \quad (5)$$

where $G_j(x)$ is the combined set of inequality and bound constraints defined in Eq. (3) and (4).

We set the initial values of the optimization parameters $v_i^{(0)}$, $s_i^{(0)}$, and $\sigma_i^{(0)}$ to correlate with the difference between the saliency of the input image and its corresponding importance value:

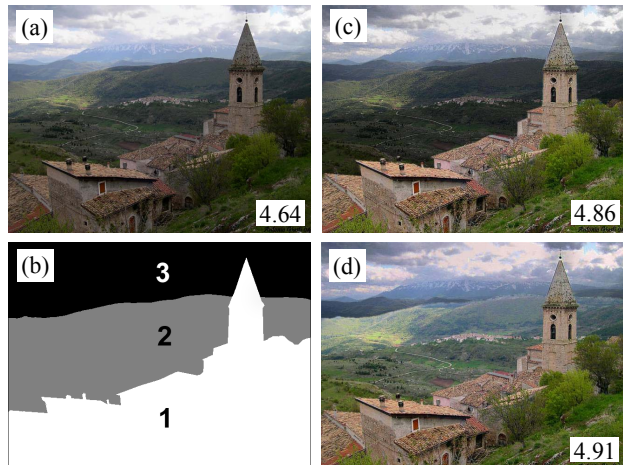


Figure 5: Example enhanced images (aesthetics score is at the bottom right of each image). (a) Original image. (b) Segmentation map with order of importance. (c)-(d) Two result images produced by our saliency retargeting algorithm for different sets of important values.

$$v_i^{(0)}, s_i^{(0)}, \sigma_i^{(0)} = \begin{cases} 0.1 & q(T_i) \leq q(M_i) \\ -0.1 & q(T_i) > q(M_i) \end{cases} \quad (6)$$

For bounds setting, we have found that $v_B = 0.15$, $s_B = 0.08$ and $\sigma_B = 0.2$ produce reasonably good and natural results for most images.

3.1.2 Image Modifications

The set of optimized parameters $x = \{v_i, s_i, \sigma_i\}$ is used to modify the input image. The modification is performed in the HSV color space. To obtain more natural images, it is important that we consider the contrast sensitivity of human visual perception [10] and maintain the contrast within an image segment when adjusting its luminance. Therefore, we apply non-linear contrast-preserving changes to the V channel based on the perception theory that our eyes perceive higher contrast at the lower range of luminance compared to the higher range of luminance. For every pixel in object segment i , the new luminance value is computed as

$$V' = x + 6v_i(V^2 - V). \quad (7)$$

This formula ensures that the sum of luminance changes is the same as adding the value v_i directly to the V channel of each pixel in the object segment. Modification to the color saturation is performed by adding the parameter s_i directly to the S channel of the respective object segments. The sharpness effect is applied to the V channel after the luminance adjustment. The parameter σ_i is multiplied by a user-specified parameter γ such that $|\gamma\sigma_i|$ represents the sigma of the unsharp mask filter for sharpening or the standard deviation of the

Gaussian filter for blurring. This γ parameter allows users to set the magnitude of the sharpness transformation to achieve their desired depth of field. In our experiments, we used $\gamma = 3$ for all images.

3.2. Aesthetics Maximization

For aesthetics maximization, we first generate a set of vectors, each containing a different sequence of importance values that satisfies the user-specified order of importance of the object segments. Each vector is then used by our saliency retargeting algorithm to produce an enhanced image in which the region saliency values of the object segments match the importance values in the vector. Our aesthetics score prediction model will then compute an aesthetics score for every enhanced image and the image with the highest score will be selected as the result image.

We set the initial vector of importance values as a geometric sequence where

$$T_i^{(0)} \leftarrow p^{N-i}, \quad (8)$$

for object segment $i = 1, 2, \dots, N$. We then iteratively update this vector of importance values exponentially, i.e.

$$T_i^{(k+1)} \leftarrow (T_i^{(k)})^\phi, \quad (9)$$

to generate the set of vectors containing importance values. Each vector $(T_1^{(k)}, T_2^{(k)}, \dots, T_N^{(k)})$ is passed to the saliency retargeting algorithm to generate an enhanced image. Using the geometric sequence in Eq. (8) to generate the initial importance values signifies that the initial relative importance from one object segment to the next object segment is constant. The vector updating process in Eq. (9) gradually widens the relative difference in importance values between the object segments. Values of p and ϕ that are too small lead to generation of redundant enhanced images that are very similar, thus increasing computational time unnecessarily. On the other hand, values of p and ϕ that are too large may potentially skip important solutions. From our experiments, we have found the values $p = \sqrt{2}$ and $\phi = 1.2$ to be appropriate. To further reduce the search space, we limit the normalized importance value $q(T_i)$ to a maximum value of 0.85. We observed that any object segment that carries an importance value more than 0.85 tends to produce over-dominance effect that often leads to reduced aesthetics.

The reason for updating the importance value vector as described in Eq. (9) is that we want to generate a set of enhanced images where the main subject becomes progressively more dominant. Our studies with image aesthetics have shown that the main subject’s degree of dominance is one of the most significant features that affect aesthetics measurements.

3.2.1 Aesthetics Score Prediction

Instead of using a scoring system that gives a score based only on features used in image modifications

[8, 9], we trained a generic aesthetics score prediction model that is based on a more complete set of aesthetics features. We believe this approach will reduce bias and also increase the accuracy of aesthetics evaluation. In addition, it also allows our method to be extended to support any image modifications including re-composition.

We extend Wong and Low’s saliency-enhanced aesthetics classification approach [18] to develop an aesthetics score prediction model. In addition to a set of global image features, their classification model uses a set of salient features that characterize the subject and the subject-background relationship. It is therefore pertinent for our context. Adopting their automatic subject extraction algorithm, we extract the main subject(s) from a set of 2682 training images from the photo.net online photo community database (<http://www.photo.net>) for feature computation. We use a set of 34 features for training our model that consists of a subset of 29 features from Wong and Low’s feature set (excluding their 15 wavelet features) and additional five new features—rule-of-thirds [9], visual balance [9], average saliency of first subject, average saliency of second subject, and difference between average saliency of first subject and background. This feature set can be categorized into two types—low-level features such as contrast, exposure, saturation, sharpness, and saliency and photographic rules such as the rule-of-thirds and visual balance. We performed 5th-degree polynomial regression on this feature set to infer an aesthetics score in the range of 1 to 7 to train our prediction model. The regression produces a residual sum-of-squares $R_{\text{res}}^2 = 0.33$ and correlation coefficient value of 0.58. Despite the correlation is only moderately high, results from our user evaluation in Section 5.2 shows that our model is effective for comparing images before and after the enhancements.

4. Results

We have tested our algorithm on a set of test images selected from our personal photo collections, from photo.net, and from the BSDS segmentation benchmark image set [20]. Based on our own judgment, we appropriately segmented each test image into 2 to 6 object segments, and ranked the importance of the object segments. It typically takes 30 to 90 seconds to segment an image. For 80% of the test images, our algorithm is able to converge to solutions that produce retargeted images with improved aesthetics scores. Interestingly, for the 20% of the test images that result in retargeted images that do not have improved aesthetics scores, our human subject experiment shows that these retargeted images are not preferred over the originals by human viewers.

Figure 6 and 7 show some of the results produced by our algorithm. For the example in Figure 6, the average

saliency values of the object segments in the input images do not correlate with the desired order of importance. In the result, our algorithm has successfully altered the input image to achieve the targeted order of importance. We can observe that the enhanced image is produced by local modification of the low-level features in each object segment. The computed aesthetics score of this test image increases from 4.59 to 4.94 after the enhancement. More results are presented in Figure 7.

5. User Evaluation

For an objective evaluation of our approach, we conducted two human subject experiments. In Experiment 1, we let each test subject view either a set of 16 original images or a set of 16 images enhanced from the original images, and track his/her gaze while he/she is looking at each image. Each image is displayed for 3 seconds. The test subject does not know which set of images he/she is looking at. The goal of Experiment 1 is to evaluate the effectiveness of our method in modifying the actual visual

saliency in the images. In Experiment 2, we let each test subject compare 40 pairs of images. Each time, two images are shown side-by-side on the screen, where one is the original and the other has been enhanced by our algorithm, and the test subject is required to choose one that he/she prefers. The screen positions (left or right) of the original image and its corresponding enhanced one are chosen at random and so the test subject does not know for sure which is the enhanced image. The purpose of Experiment 2 is to study the effectiveness of our method in enhancing the aesthetics of the images. The images we used in the user experiments are a subset of the test images mentioned in Section 4. We recruited 28 subjects for Experiment 1 and 32 for Experiment 2. The following subsections describe the results of the experiments.

5.1. Results of Experiment 1

Table 1 shows the correlation between the number of gaze fixations on each object segment and the desired importance value of the segment. One can see a significant increase in the correlation for the enhanced



Figure 6: An example result. (a) Object segments and the desired importance order. (b) The input image. (c) The average saliency of each object segment in the input image. (d) The result image. (e) The average saliency of each object segment in the result image.

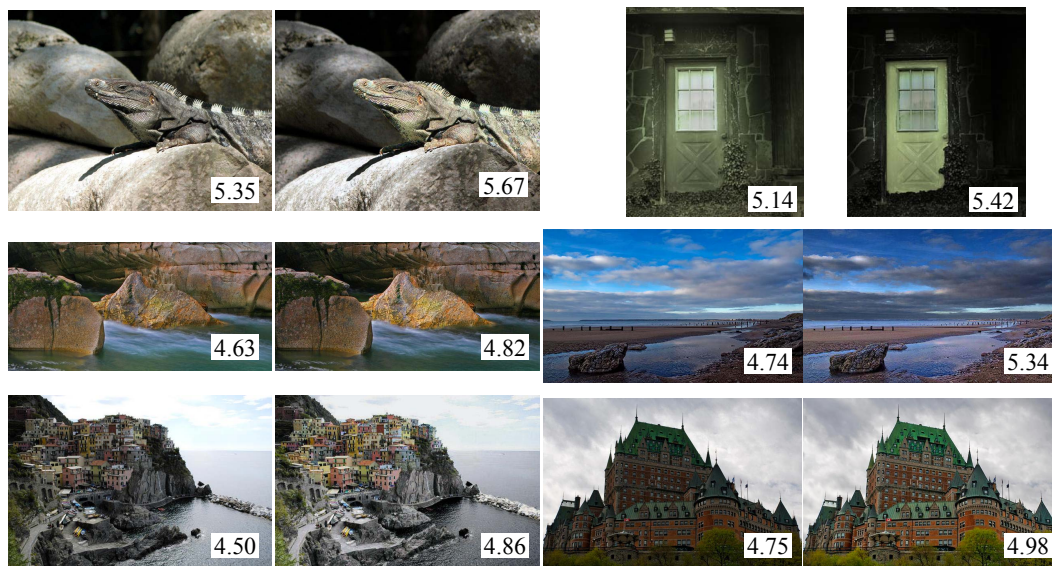


Figure 7: Example results. For each pair of images, the one on the left is the original and the one on the right is the enhanced. At the bottom right is the computed aesthetics score for each image.

images. It shows that our algorithm is quite effective in retargeting the actual saliency in the images. Figure 8 compares the scan paths on two original and their enhanced images. In both examples, the scan paths on the original images are diverse and less time is spent on the main subjects. On the other hand, on the enhanced images, the viewers' attention is concentrated around the main subjects, indicating that the viewers can more easily identify the main subjects in the image, which potentially leads to better aesthetics experience.

Table 1: Correlation between the number of gaze fixations and the desired importance value.

	Correlation coefficient	<i>p</i> -value
Original images	0.284	0.0757
Enhanced images	0.5464	0.0003

5.2. Results of Experiment 2

Results from this experiment clearly indicate that test subjects have a preference for the enhanced images. The chart in Figure 9(a) shows that for every image pair used in the experiment, at least 40% of the test subjects chose the enhanced image. For 87.5% of the image pairs, majority of the test subjects chose the saliency retargeted images. If we treat images with 40%–60% preference votes as noise, where there is no significant preference for the original or enhanced image, we can conclude that our algorithm is unlikely to produce images less aesthetically pleasing than the input images.

Next, we study the relationship between image aesthetics and the total saliency change. Aesthetics is represented by the ground truth of the number of preference votes an enhanced image gets. The total saliency change is computed by summing the saliency differences between the same object segments in the original image and in its corresponding enhanced image. The chart in Figure 9(b) shows two peaks, indicating there is high saliency change for two ranges of enhanced images. The total saliency change peaks at the enhanced images with the highest preference. The second peak falls within the range of 50%–59.9%, which can be considered noise. Studying the images within this range suggests that high saliency change in an image may potentially result in unnaturalness in the result, making it less likeable. In addition, computing the correlation between the number of preference votes an enhanced image gets and its total saliency change for images with more than 60% preference votes gives a coefficient of 0.42 (with *p*-value of 0.035). This suggests there is a significant positive correlation between image aesthetics and saliency change.

To study the performance of our aesthetics score




Figure 8: Scan paths on images. For each pair of images, the left one is the original and the right one is the result.

prediction model, we evaluate the accuracy of our model in predicting the relative score between an original image and its corresponding enhanced images. For images with more than 60% preference votes, we find that our algorithm is able to predict an increase in aesthetics score with 88% accuracy. We notice that majority of the enhanced images that are given reduced scores by our aesthetics model fall under the category of images where consensus among human test subjects is low. This result is indeed encouraging, considering the subjectivity in evaluating aesthetics is unavoidable.

6. Limitations and Future Work

We are currently using fixed bounds for the amount of changes to the three low-level image features. We have

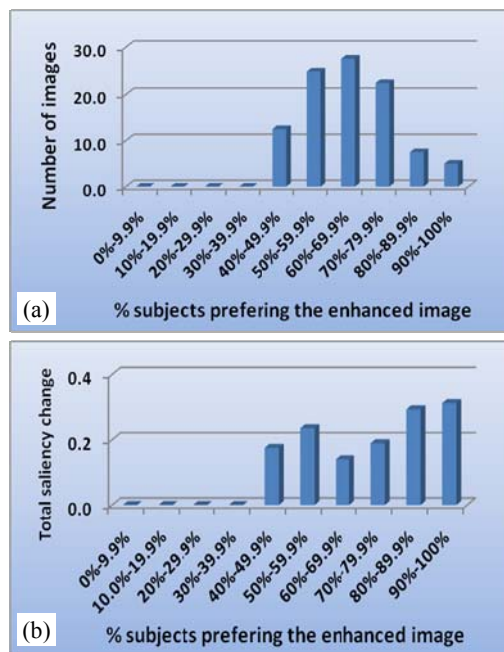


Figure 9: Results from Experiment 2.

observed that for some cases, these bounds are too conservative, and do not allow sufficient change of the image to satisfy the target saliency. On the other hand, the bounds appear to be too loose for some cases, and it results in overly-enhanced images that look unnatural. Ideally, the image modifications should use image “naturalness” to limit the amount of changes; however, accurately evaluating the “naturalness” of images is still an open problem.

In general, the composition within a photograph is still one of the most important elements that determine the aesthetics of the image. Often, photographs are so poorly composed that no amount of saliency retargeting can make them more acceptable than before. Vice versa, in cases where main subjects are not salient, re-composition alone cannot help to improve the image aesthetics much. Therefore, saliency retargeting and re-composition method such as an automatic cropping and image retargeting tool [9] serves as good complement for each other. These two image editing approaches could be integrated to produce a more complete content-aware image enhancement tool.

7. Contributions and Conclusion

The main contribution of our work is the novel idea of using saliency retargeting as a means for image aesthetics enhancement, and a simple practical algorithm for the approach. The saliency retargeting is performed by modifying only the low-level image features that correspond directly to the features used in the saliency computation. Very importantly, the relationship between image aesthetics and image saliency, the goal of saliency retargeting, and the image features used for the saliency evaluation together provide a clear guidance to how the image can be modified to enhance its aesthetics. Results from our user experiments have supported the effectiveness of our approach for enhancing image aesthetics.

8. Acknowledgements

Vlad Hosu contributed to the initial development of the saliency retargeting framework. This work is supported by the Singapore MOE Academic Research Fund (Project Number: T1 251RES0804).

9. References

[1] Banerjee, S. & Evans, B.L. (2007) In-Camera Automation of Photographic Composition Rules. *IEEE Transactions on Image Processing*, 16(7), 1807-1820.
 [2] Achanta, R. & Ssstrunk, S. (2009) Saliency Detection for Content-aware Image Resizing. *IEEE International Conference on Image Processing*, pp. 1005-1008.

[3] Barnard, K., Cardei, V. & Funt, B. (2002) Comparison of Computational Color Constancy Algorithms-Part 1: Methodology and Experiments with Synthesized Sata. *IEEE Transactions on Image Processing*, 11(9), 972-983.
 [4] Harel, J., Koch, C. & Perona, P. (2006) Graph-Based Visual Saliency, In *Advances in Neural Information Processing*, pp. 545-552.
 [5] Gill, P., Murray, W., and Wright, M. (1981) *Practical optimization*. New York:Academic.
 [6] Itti, C. Koch, E. Niebur, A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), pp. 1254-1259.
 [7] Kashyap, R.L. (1994) A robust variable length nonlinear filter for edge enhancement and noise smoothing In *Proc. of Int. Conf. on Signal Processing*, pp. 505-510.
 [8] Kao, H.C., Ma, W.C., and Ming, Q. (2008) Esthetics-based Quantitative Analysis of Photo Composition. In *Proc. of Pacific Graphics Conference*, Tokyo, Japan.
 [9] Liu, L.,G., Chen, R.J., Wolf, L., &Cohen-Or, D. (2010) Optimizing Photo Composition, In *Proc. of Eurographics 2010*, Norrkping, Sweden.
 [10] Manos, J.L., and Sakrison, D. J. (1974), The Effects of a Visual Fidelity Criterion on the Encoding of Images, *IEEE Transactions on Information Theory*, 20(4), 525-535.
 [11] Moroney, N. (2000) Local color correction using non-linear masking. In *Proc. IS&T/SID Eight Color Imaging Conference*, pp. 108-111.
 [12] Polesel, A., Ramponi, G., & Mathews, V.J. (2000) Image enhancement via adaptive unsharp-masking. *IEEE Transactions on Image Processing*, 9(3), 505-510.
 [13] Rahman, Z., Jobson, D., & Woodell, G. (2004) Retinex processing for automatic image enhancement. *Journal of Electronic Imaging*, 13(1), 100-110.
 [14] Rizzi, A., Gatta, C., & Marini, D. (2003) A new algorithm for unsupervised global and local color correction. *Pattern Recognition Letters*, 24, 1663-1677.
 [15] Rother, C., Kolmogorov, V., & Blake, A. (2004) GrabCut: Interactive foreground extraction using iterated graph cuts. *ACM Transaction on Graphics*, vol. 23(3), 309-314.
 [16] Su, S., Durand, F., & Agrawala, M. (2004) An Inverted Saliency Model for Display Enhancement. In *Proc. of MIT Student Oxygen Workshop*, Ashland, MA.
 [17] Tomasi, C., & Manduchi, R. (1998), Bilateral Filtering for gray and color images, In *Proc. of IEEE Conference on Computer Vision*, pp. 836-846.
 [18] Wong, L.K., & Low, K.L. (2009) Saliency-enhanced image aesthetics class prediction. In *Proc. of IEEE Conference of Image Processing*, pp. 997-1000.
 [19] Wollen, P. (1978), *Photography and Aesthetics*. Oxford Journal of Arts, 4(19), 9-28.
 [20] Berkeley Segmentation Dataset Images (BSDS) <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/bsds/BSDS300/html/dataset/images.html>
 [21] Galitz, R., Composition: Frame Construction 101 <http://www.galitz.co.il/en/articles/composition.shtml>