

# Selective Removal of Extraneous Photographs

K. Armin Samii  
ksamii@ucsc.edu

Allison Carlisle  
acarlisl@ucsc.edu

Uliana Popov  
uliana@soe.ucsc.edu

James Davis  
davis@soe.ucsc.edu

## ABSTRACT

In this paper, we propose a method for selecting the most representative high-quality images from a set of user photographs. To avoid redundancy arising from many similar images, we find all sets of near-duplicates. We then rank images based on three technical qualities: exposure, blur, and contrast. Finally, we provide an ordering of the images which accounts for both quality and uniqueness. Our goal is not to rank images based on subjective aesthetic qualities, but instead to help a photographer filter out technically flawed photographs and focus on objectively high quality shots.

## Categories and Subject Descriptors

K.8.3 [Computing Milieux]: Personal Computing

## General Terms

Management

## Keywords

image processing, quality screening, exposure, meta data, color content

## 1. INTRODUCTION

The ease of using digital cameras allows for a large number of photographs to be taken at any given event. This increase in photograph quantity can lead to many unwanted low-quality photographs that need to be filtered by the user, which can be time-consuming and repetitive.

To automate this process, an aesthetics ranking technique could be implemented to find the appeal each image has, but this method would discard many images with potential to be fixed through retouching.

Implementing a technical quality ranking technique will keep the photographs suited for retouching, but does not take into account redundancy of similar images, and the

user would be left with a set which either includes many duplicates or does not span the entire event.

We thus propose a technique which ranks images based on their technical qualities with a bias for unique images.

To focus our research and determine which image qualities are relevant to technical image quality, we assume three general steps a photographer takes between shooting and using a picture, which we call the Photographer's Process:

1. Remove: Sort through the imported images and remove the ones least suited for retouching.
2. Retouch: Modify the raw files just selected to stylize and enhance them.
3. Retrieve: Select the retouched images most suitable to a given task.

This work focuses on the removal stage and contains both quality assessment and similarity ranking. We implement technical quality assessment using three modules: blur detection, exposure, and contrast. The blur detection uses three measures: the gradient magnitude of edges, the similarity to a predicted out-of-focus image, and the probability of a linear point-spread function existing. Exposure is measured by finding the balance in brightness throughout the image, both locally and globally. Contrast is measured based on the local changes in brightness throughout an image. We implement similarity ranking using global color similarity, foreground color similarity, and timestamps. These modules are combined to obtain a reordering of the set representing how suited they are for removal (the "Importance Order"). A user study and tests with publicly available data has shown 89% accuracy in our absolute ranking of photographs. The similarity ranking was evaluated separately and found to have 81% accuracy. Fig. 3 shows our accuracy in finding the subject and rating quality.

## 2. RELATED WORK

Previous works have classified images as professional vs. amateur with high accuracy[1][2][3], which is useful for search engines' retrieval of high quality images. Others have focused on personalized *aesthetic* rankings of photographs[4][5]. All of these works focus on finalized images, while we focus on images that are still raw and may be retouched afterwards.

Event classification has been explored in order to choose the best images from an entire event[6][7], while we choose the best images from each scene within each event: images

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

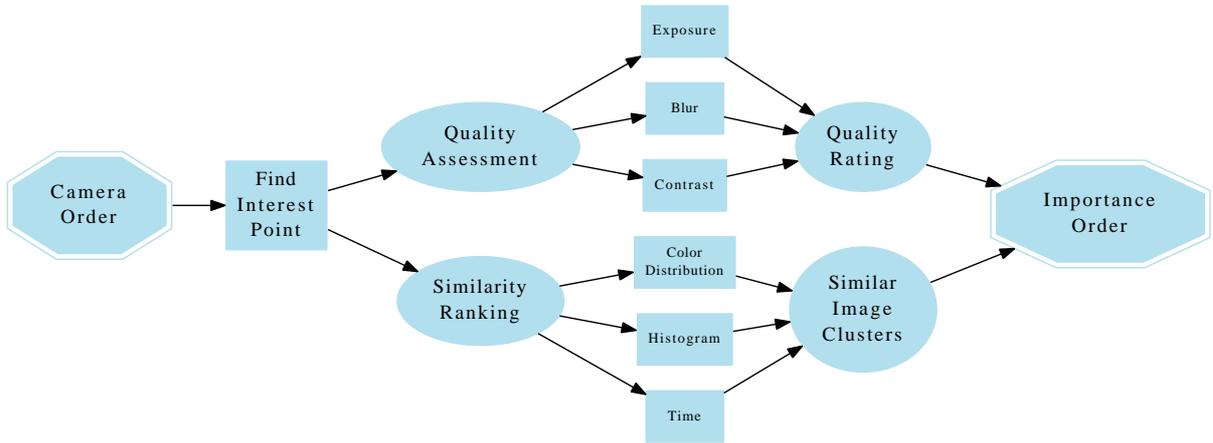


Figure 1: Interest points are used to find the foreground subject. Blur, contrast, and exposure assessment algorithms are used to calculate an image’s quality rating. Images are clustered into groups of similar images based on their histograms, color distribution, and timestamps. A reordering of the input results, allowing the user to remove images at the end of this ordering.

with the same subject content rather than the same event context.

Kormann, Dunker, and Paduschek[8] describe a method of automatically rating and ranking images based on image content and time-metadata, but do not report numerical results, only that their results are better than random.

### 3. QUANTIFYING IMAGE QUALITY

To find the Importance Order of the user’s photographs, we use both quality assessment and similarity clustering. Quality is determined by the rankings of the blur, exposure, and contrast modules. Similarity of images is determined by matching the color content and timestamp of every pair of images. Both modules make use of a foreground region detection module based on interest points. We combine these measures to get the final ordering (Fig. 1).

The quality modules (blur, exposure, and contrast assessment) provide a ranking between zero and nine, with larger numbers indicating higher quality. Because an image which is flawed in any of the three factors would be considered poor, we take the weighted sum-of-squares to penalize low scores more than we reward high scores. We assign weights to each module as follows: blur 50%, exposure 40%, and contrast 10%.

#### 3.1 Foreground detection

Several modules make use of foreground detection. We use a Harris corner detector[9] (utilizing NASA’s Vision Workbench[10]) to obtain interest points for each image. To extract a bounding box from these points, the most dense rectangle is calculated by maximizing the ratio of interest points to rectangle area. We assume this to be the primary foreground subject.

#### 3.2 Exposure

Exposure is a measure of how appropriate the lighting is in a given image. While easily measured in-camera using a light meter, we must evaluate the lighting condition based on pixel intensity.

To assess an image’s exposure quality, we use both global

and local calculations to compare the average grayscale value of the region to middle gray (halfway between black and white on a lightness scale)[11].

The global method ranks an image highly when the average grayscale value is near middle gray. The local method divides the image into 20x20 grid. Images are ranked highly when the bright and dark grid cells average near middle gray, with no cells being extremely light or dark. We also penalize an image if there is not enough variation between grid cells. These two measures are averaged to provide a final ranking.

#### 3.3 Blur Detection

There are three measures used to assess the amount of blur: edge width, global point-spread, and a comparison against a computed blurry image.

The first two rely on an edge-detection algorithm which is done in three steps. First, we calculate the differences between a pixel’s luminosity and the image’s average luminosity. Large differences are considered edge pixels. Mislabeled edge pixels are then removed by ensuring each one is connected to another edge pixel. Finally, edges are spaced out by ensuring there is a large radius of non-edge pixels around each set of edge pixels; if there is not, the edge pixel with the largest difference value from the first step is kept.

The first measure finds the width of each edge by finding a “line of maximum contrast” from the brightest to darkest pixel around each edge pixel. Sharper edges will have shorter line lengths.

The second measure looks at the orientation of the lines of maximum contrast. If there are significantly more lines oriented in a single direction than the average, we assume the point-spread to be in that direction.

Finally, if the previous two measures are inconclusive, we compute a predicted blurred model of the foreground using a Gaussian blur filter on the image. The closer an image is to this model, the blurrier it is.

These three values are combined to determine the amount of blur. Only one needs to have high confidence for the final result. The results are averaged if none have high confidence.

#### 3.4 Contrast

Contrast is the difference in visual properties that makes an object distinguishable from other objects and the background. We use luminance contrast and Root Mean Square (RMS) algorithm, which does not depend on the spatial frequency content or the spatial distribution of contrast in the image. It is defined as the standard deviation of the pixel intensities[12].

## 4. SIMILAR-IMAGE CLUSTERING

To find similar images we use three measures: time nearness, histogram similarity, and color distribution.

### 4.1 Timestamp

First, we use the timestamp to obtain a time similarity index between all pairs of images. Prior work focused on finding large timestamp gaps[13]. We want to allow nonconsecutive images to be grouped together, so we have derived a formula which finds temporal closeness between every pair of images. We find the similarity  $S_{i,j}$  between two images  $i$  and  $j$  on a 0-9 scale:

$$S_{i,j} = \frac{G_{i,j}}{A_{i,j}} \quad (1)$$

Where  $A_{i,j}$  is the average gap of the 8 images directly before  $i$  and after  $j$ , and  $G_{i,j}$  is the log of the time gap between the two images taken at time  $T_i$  and  $T_j$ :

$$G_{i,j} = \log(|T_i - T_j|) \quad (2)$$

The calculated similarity index  $S_{i,j}$  is used to weight the results of the next two steps.

### 4.2 Histogram Similarity

Directly comparing histograms will produce many false positives due to similar-colored images with different color distribution and content, and many false negatives due to changes in exposure affecting the histograms. Our method is robust against both of these.

The image is first divided into a 2x2 grid, with four pairs of values describing each grid cell: red, green, blue, and grayscale histograms, paired with the median color for each histogram. The grid protects against false positives. We then divide each histogram into 20 bins, each containing 5% of the range of values, plus 2% of the values in adjacent bins, which provides resistance against minor exposure changes. So, for example, the second bin will contain all pixels between 3% and 12%.

When comparing the histograms of two images, the bins are scaled so that all channels have the same median value. This provides resilience against larger exposure changes. A similarity ranking for each histogram is derived from the differences in the size of each scaled bin. Finally, each grid cell is weighted by the amount of the foreground that it contains. The result is the grid-weighted average of differences between each bin, with smaller differences indicating higher similarity:

$$\text{Difference}_{I,J} = \sum_{g=1}^4 W_g \sum_{c=1}^4 \sum_{b=1}^{20} |I_{b,c} - J_{b,c}| \quad (3)$$

Where  $g$  is the grid cell,  $W_g$  is the weight of the cell based on foreground overlap,  $c$  is the channel,  $b$  is the bin, and

$|I_{b,c} - J_{b,c}|$  is the difference between the number of pixels in bin  $b$ , channel  $c$  for images  $I$  and  $J$ .

### 4.3 Foreground Color Distribution

We degrade an image to remove features from the foreground while maintaining general color and shape information by applying a strong Gaussian blur filter to the foreground. When comparing two images, a smaller pixel-by-pixel difference indicates higher similarity. Because only the foreground is used, this method is robust against composition changes and subject movement between photographs.

### 4.4 Clustering

Clusters are then formed using a Quality-Threshold algorithm. We first put each image into its own group, then iteratively merge the two most similar groups until a minimum similarity threshold is reached. In this way, groups vote for other similar groups, rather than single image voting for other images, resulting in a noise-resistant clustering. This method successfully groups panoramic sets of images together, even when the first image's histogram does not match the last image's.

To determine the Importance Order, we iteratively choose the image with the highest rank, then penalize every other image in the same set.

## 5. RESULTS

We asked Amazon Mechanical Turk users to group together photographs within four sets of twenty images based on their own personal measure of "similarity." Our similarity index for each pair of images agrees with Turk users' decisions of whether a given pair of images is in the same group with 81% accuracy.

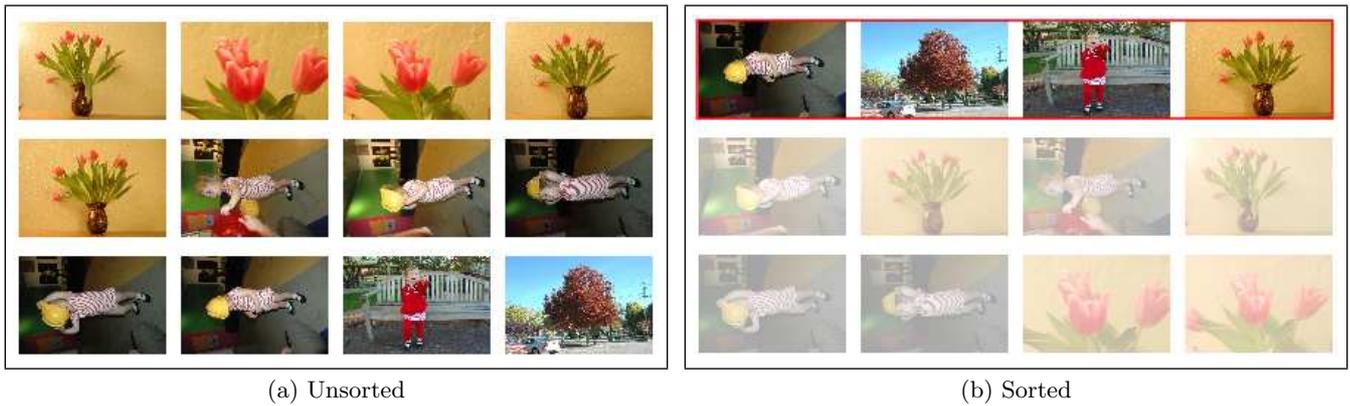
We also assessed blur and exposure individually using a binary classification. First, we use our own data to confirm an image is low quality. We gathered 70 out-of-focus images uploaded by Turk users and correctly classified 62 of them as blurry (89%). We then gathered 86 poorly exposed images (as determined by the in-camera light meter) and correctly classified 79 of them as having poor exposure (92%).

To assess technically high quality images, we use INRIA's publicly available "Original Image" 157-image dataset[14]. We correctly classified 139 as not blurry (89%) and 135 as well-exposed (86%). Averaging the four results gives 89% accuracy overall.

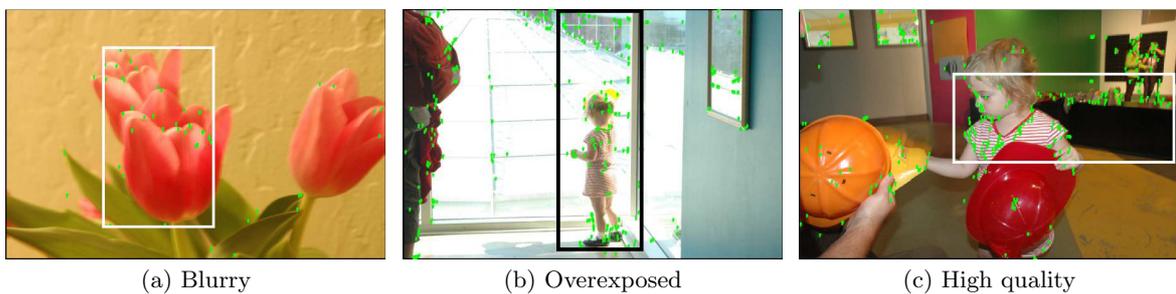
Fig. 2 shows an example of a user's input images being resorted by quality and uniqueness. The top four images in Fig. 2b represent the highest quality images in each similar-image set. The nontrivial case of zoomed-in flowers was recognized as the same foreground subject. The top images are consistent with Turk users' votes.

## 6. CONCLUSION AND FUTURE WORK

We propose a method of applying current research to automate another step of the Photographer's Process. By focusing on obtaining an ordering which is representative of all photographs taken, we obtain a diverse set of high quality images similar to what a user would have chosen manually. We derive a novel algorithm for analyzing exposure quality. We improve upon previous algorithms which find temporal gaps between images to obtain a metric for temporal nearness. The final Importance Order depends on both the



**Figure 2:** (a) shows four sets of similar photographs provided by the user. (b) shows the reordered set. We asked 10 Turk users to rank each image and we correctly chose the top image from each set.



**Figure 3:** Examples of (a)(b) low quality and (c) high-quality images, as determined by our modules. The dots are interest points found; the rectangle is the bounding box considered to be the subject.

quality ranking and the number of similar images which have appeared earlier in the ordering.

Our work focuses on a small portion of the Photographer’s Process. In the future, we would like to see the three steps combined to automate the entire Process with minimal human input. This would facilitate a photographer’s creativity by allowing more time spent on pure aesthetics and leaving technical details to the computer.

## 7. REFERENCES

- [1] Yan Ke, Xiaoou Tang, and Feng Jing, “The design of high-level features for photo quality assessment,” in *CVPR 2006*, June 2006, vol. 1.
- [2] R. Datta, D. Joshi, J. Li, and J. Wang, “Studying aesthetics in photographic images using a computational approach,” in *ECCV*, vol. 3953 of *Lec. Notes in Comp. Sci.* Springer Berlin / Heidelberg, 2006.
- [3] Y. Luo and X. Tang, “Photo and video quality evaluation: Focusing on the subject,” in *ECCV 2008*. Springer Berlin / Heidelberg, 2008.
- [4] X. Sun, H. Yao, R. Ji, and S. Liu, “Photo assessment based on computational visual attention model,” in *Proc. ACM Multimedia*. 2009, MM ’09, ACM.
- [5] C.H. Yeh, Y.C. Ho, B.A. Barsky, and M. Ouhyoung, “Personalized photograph ranking and selection system,” in *Proc. of the Int. Conf. on MM*. 2010, MM ’10, ACM.
- [6] A.C. Loui and A. Savakis, “Automated event clustering and quality screening of consumer pictures for digital albuming,” *IEEE Mult. Trans.*, Sept. 2003.
- [7] C. Jang, J. Lee, J. Lee, and H. Cho, “Smart management system for digital photographs using temporal and spatial features with exif metadata,” in *Digital Information Management*, 2007.
- [8] D. Kormann, P. Dunker, and R. Paduschek, “Automatic rating and selection of digital photographs,” in *Semantic Multimedia*. Springer Berlin / Heidelberg, 2009.
- [9] C. Harris and M. Stephens, “A combined corner and edge detector,” *Proc. Alvey Vision Conf.*, 1988.
- [10] M.D. Hancher, M.J. Broxton, and L.J. Edwards, *A User’s Guide to the NASA Vision Workbench.*, NASA Ames, Intelligent Syst. Div., 2006.
- [11] S. Barclay, “The motion picture image: from film to digital,” 2000.
- [12] E. Peli, “Contrast in complex images,” in *Journal of the Optical Society of America*. 1990.
- [13] J.C. Platt, M. Czerwinski, and B.A. Field, “Phototoc: automatic clustering for browsing personal photographs,” in *ICICS and PCM*, 2003.
- [14] H. Jegou, M. Douze, and C. Schmid, “Hamming embedding and weak geometric consistency for large scale image search,” in *Proc. of the Eur. Conf. on Comp. Vis.* Oct 2008, Springer.