Automatic Exposure Correction of Consumer Photographs

Lu Yuan and Jian Sun

Microsoft Research Asia

Abstract. We study the problem of automatically correcting the exposure of an input image. Generic auto-exposure correction methods usually fail in individual over-/under-exposed regions. Interactive corrections may fix this issue, but adjusting every photograph requires skill and time. This paper will automate the interactive correction technique by estimating the image specific S-shaped non-linear tone curve that best fits the input image. Our first contribution is a new Zone-based region-level optimal exposure evaluation, which would consider both the visibility of individual regions and relative contrast between regions. Then a detail-preserving S-curve adjustment is applied based on the optimal exposure to obtain the final output. We show that our approach enables better corrections comparing with popular image editing tools and other automatic methods.

1 Introduction

Exposure is one of the most important factors of determining the quality of a photograph. In over-exposed or under-exposed regions, details are lost, and colors are washed out. Despite that sophisticated metering techniques have been equipped on the cameras, taking well-exposed photos remains a challenge for normal users. There are several reasons: 1) the camera's metering (*e.g.*, spot, center-weighted, average, or multi-zone metering) is not perfect. If the metering points/areas are not targeting the subject or there are multiple subjects, the metering may fail. Fig. 1(a) is a failure case caused by the backlit; 2) the assumption that the mid-tone of the subject is gray is sometimes invalid due to the complex reflectance of the world (*e.g.*, a snow-white rabbit is often captured as an undesired grayish rabbit without exposure compensation); 3) in-camera post-processing capability is limited, especially for the low-end cameras.

To address this issue, some automatic methods like auto-level stretch [1] and histogram equalization [1] have been proposed to correct the exposure. For example, auto-level stretch linearly maps the brightness to the maximum tonal range (*e.g.*, [0, 255]). This method, however, only uses the statistics of the whole image, without considering each image region individually. For the backlit case in Fig. 1, auto-level stretch does not take effect (see Fig. 1 (b)) since the image histogram has reached the maximum tonal range (top-left of Fig. 1(a)). Histogram equalization [1] (and its variations [2]) better distributes the intensity values over the histogram. Unfortunately, it would produce unrealistic effects in photographs (see Fig. 1(c)).

If user assistance is allowed, the interactive correction method is more effective. For instance, most photo editing software allow the user to manually adjust a non-linear tone

A. Fitzgibbon et al. (Eds.): ECCV 2012, Part IV, LNCS 7575, pp. 771–785, 2012.

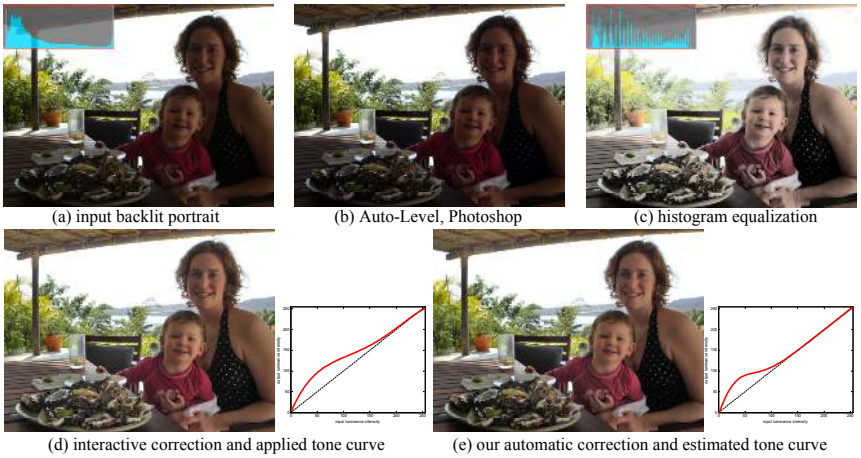

(a) input backlit portrait (b) Auto-Level, Photoshop (c) histogram equalization

(d) interactive correction and applied tone curve (e) our automatic correction and estimated tone curve

Fig. 1. A typical under-exposed photo. On the top-left of (a), we show the luminance histogram of the input image which has the maximum tonal range and peaks in shadows and highlights.

curve [3] (*e.g.,* S-curve) to correct the dark/mid-tone/bright regions separately. Fig. 1(d) is the assisted result by expert. But the best shape of the curve varies a lot from image to image. Touching up every single image is impractical for typical consumers.

In this paper, we present an *automatic* exposure correction method that can estimate the best image specific non-linear tone curve (the S-curve in our case) for a given image. Unlike [4], we need no training data. Note that it is a non-trivial task since the variation of input consumer photographs is so large. The key to the success of an automatic correction is to know what the best exposure should be for every image region.

To address this fundamental issue, we borrow the concept of "Zone" from the well-developed Zone System [5] in photography. The Zone system quantizes the whole exposure range as eleven discrete zones. We formulate the exposure correction as a zone estimation problem - we optimize a global objective function to estimate the desired zone in each image region by simultaneously considering two goals: maximizing the local details in each region, and preserving the relative contrast between regions.

After getting the estimated zone of every region, we propose a new non-linear curve-based correction algorithm called *detail-preserving S-curve adjustment*, to push each region to its desired zone, as much as possible. Compared with generic S-curve adjustment [6][7][8], our detail-preserving S-curve adjustment can maintain local details and avoid halo effects. Fig. 1(e) shows our estimated curve and final corrected result.

Like most automatic approaches, our approach does not address the user preference issue [9]. The "correct" exposure may be defined as the one that achieves the effect the photographer intended. However, our user studies show that an automatic correction still benefits most typical consumers - especially for their daily photos processing. We also show our new exposure optimization provides significant visual quality improvement over pervious work. Since our correction is simple and robust, it can be chosen as a better alternate in photo editing tools and a built-in camera component.

2 Related Work

Automatic exposure control is one of the most essential research issues for camera manufacturers. The majority of developed techniques are hardware-based. Representative work include HP "Adaptive Lighting" technology [10] , Nikon "D-Lighting" technology [11]. These methods compress the luminance range of images by a known tone mapping curve (*e.g.*, Log curve) and further avoid local contrast distortion by "Retinex" processing [12]. Specific hardware has been designed to perform per-pixel exposure control [13] or scene-based (*e.g.*, backlit, frontlit [14] or face [15]) exposure control. Some automatic techniques (*e.g.* [16]) are proposed to estimate the optimal exposure parameters (shutter speed and aperture) during taking photos.

There are numerous techniques about software-based exposure adjustment, including most popular global correction (*e.g.*, auto-level stretch, histogram equalization [1]) and local exposure correction [17][18]. However, these methods only use some heuristic histogram analysis to map per-pixel exposure to the desired one, without considering the spatial information of pixels (or regions). An interesting work [19] tries to enhancement image via frequency domain (*i.e.*, block DCT). But some fixed tone curves are used for each image and blocking artifacts occasionally occur in their results.

Some algorithms [8][20] only consider the exposure of the regions of interest (ROI) and assume it is most important to the whole image correction. Different from ours, they use a known and predefined tone curve but we will estimate the specific curve for every image. Some tone mapping algorithms [21] can also be used to estimate the key of scene and infer a tone curve to map its original exposure to the desired key. However, the key estimation is based on the global histogram analysis and is sometimes inaccurate. Exposure fusion [22] combines well-exposed regions together from an image sequence with bracketed exposures. In contrast, we only use a single image as the input.

Since the exposure correction is kind of *subjective*, recent methods [23][4][9] enhance the input image using training samples from internet or personalized photos. However, our exposure correction is not relied on the selection of training images and only focuses on the input image itself. Another issue worth mentioning is that our approach does not aim to restore completely saturated pixels like [24].

3 Automatic Exposure Correction Pipeline

Our exposure correction pipeline is depicted in Fig. 2 and divided to two main steps: exposure evaluation and S-curve adjustment. Both components are performed in the luminance channel. To avoid bias due to different camera metering systems, or user's manual settings, we would linearly normalize the input tonal range to [0, 1] at first.

The heart of our system is an optimization-based, region-level exposure evaluation (see Section 4). In the exposure evaluation, we apply a Zone-based exposure analysis to estimate the desired zone (*i.e.*, exposure) for each image region. We first segment the input image into individual regions (*i.e.*, super-pixels). In each region, we measure visible details, region size, and relative contrast between regions. Then we formulate the optimal zone estimation as a global optimization which takes into account all these factors. We also use the high level information (*e.g.*, face) to set the priority of the regions.

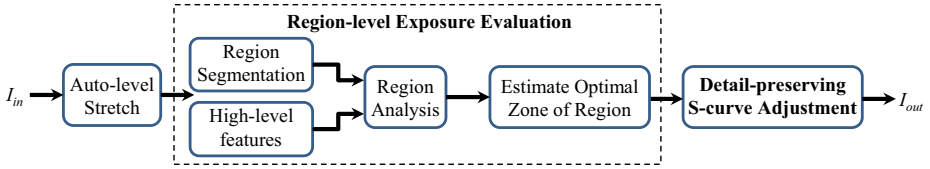

Fig. 2. Our automatic exposure correction pipeline

After the exposure evaluation, we estimate a best non-linear curve (S-curve) mapping for the entire image to push each region to its optimal zone. We further introduce a detail-preserving S-curve adjustment (see Section 5) instead of naïve S-curve mapping to preserve local details and suppress halo effects in the final result.

4 Region-Level Exposure Evaluation

The aim of our exposure evaluation is to infer the image specific tone curve for the consequent detail-preserving S-curve adjustment. To achieve this goal, we first need to know what is the "best" exposure of each region and how to estimate them all together.

4.1 Zone Region

To measure the exposure, we borrow the concept of "Zone" from Ansel Adams' Zone System [5], which is shown in Fig. 3(d). In Zone System, the entire luminance range [0, 1] is equally divided into 11 zones, ranging from 0 to X denoted by Roman numbers, with 0 representing black, V middle gray, and X pure white; these values are known as zones. In each zone, the mean intensity value is referred as its corresponding exposure. This concept was also used in recent HDR tone mapping applications [21][25] and realistic image composition [26].

We represent the image by a number of zone regions. We first decompose the image into a set of regions by graph-based segmentation [27]. Each region falls into one of the zones. Then, we merge the neighboring regions with the same zone value. To extract high-level information (*e.g.*, face/sky) for high priority of adjustment, we need to detect facial regions [28] and sky regions [29]. All connected regions belonging to face/sky regions are also merged. We call the final merged region as "zone region". Fig. 3(a-c) shows the procedure of the zone region extraction.

4.2 Optimal Zones Estimation

The optimal zone estimation can be formulated as a global optimization problem by considering two aspects: maximizing the visual details and preserving the original relative contrast between neighboring zone regions.

Measure of Visible Details. The amount of visible details in under-/over-exposed regions can be measured by the difference of the detected edges in these images which are generated by applying different gamma-curves on the input image I (the process is

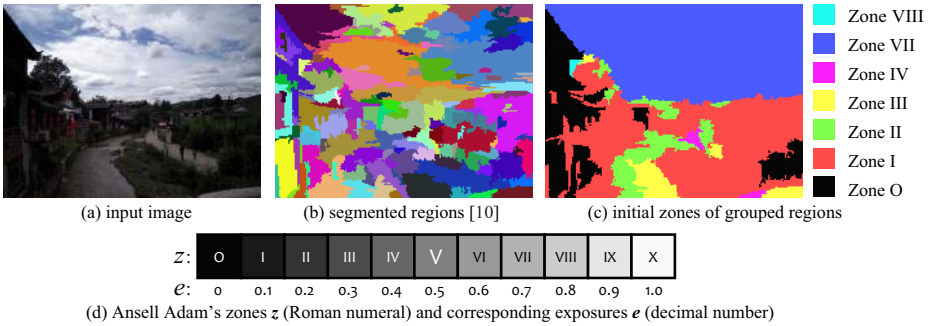

(a) input image (b) segmented regions [10] (c) initial zones of grouped regions

(d) Ansell Adam's zones z (Roman numeral) and corresponding exposures e (decimal number)

Fig. 3. Zone region extraction. In (c), different colors denote different zone regions, in which the Roman numbers denote corresponding zone values.

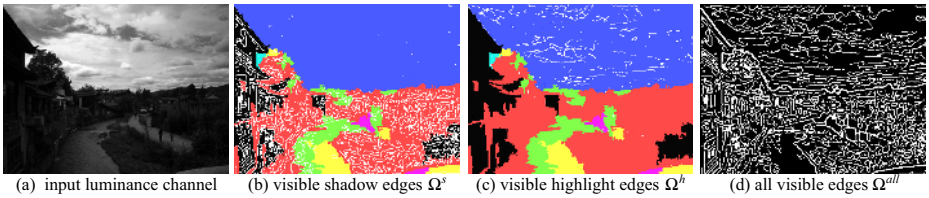

(a) input luminance channel (b) visible shadow edges Ω^s (c) visible highlight edges Ω^h (d) all visible edges Ω^{all}

Fig. 4. Measure of visible details. In (b-d), white lines are detected edges by Canny operator.

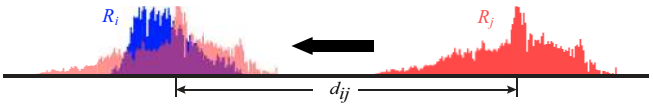

Fig. 5. Relative contrast d_{ij} is the histogram distance between two neighboring regions R_i, R_j

denoted as $I_{gamma} = I^{\gamma}$). It is based on an observation: in an under-exposed region, we can detect more/less visible edges after the gamma correction when the gamma γ is smaller/larger, and the edge number difference between two gamma-corrected images (one with small gamma, the other with large gamma) indicates the amount of recoverable details. The similar process can be applied to the over-exposed region as well.

In our implementation, we use two default gamma values $\gamma = 2.2$ and $\gamma^{-1} = 0.455$. We first detect edges on three images $I, I^{\gamma}, I^{\gamma^{-1}}$ by the same Canny operator [1] to obtain three edge sets: $\Omega_1, \Omega_{\gamma}, \Omega_{\gamma^{-1}}$. The visible details in the shadow region (zone value $<$ V) and the highlight region (zone value $>$ V) are measured by: $\Omega^s = \Omega_{\gamma^{-1}} - \Omega_{\gamma^{-1}} \bigcap \Omega_{\gamma}$ and $\Omega^h = \Omega_{\gamma} - \Omega_{\gamma^{-1}} \bigcap \Omega_{\gamma}$, shown in Fig. 4 (b)(c).

Note that the absolute differences Ω^s and Ω^h cannot be directly used since they vary from image to image. To obtain a comparable measure, we compute the relative visibility of details:

$$\nu^s = |\Omega^s| / |\Omega^{all}|, \qquad \nu^h = |\Omega^h| / |\Omega^{all}| \tag{1}$$

where $|\cdot|$ indicates the edge number in a set, and $\Omega^{all} = \Omega_1 \bigcup \Omega_\gamma \bigcup \Omega_{\gamma^{-1}}$ is the union of all three sets, shown in Fig. 4 (d).

Measure of Relative Contrast. We measure the relative contrast between zone regions using their intensity histogram distance. This distance is defined as the minimum shifting distance of two histograms to maximize their intersection (shown in Fig. 5). We use the term "relative contrast" for this distance. For example, when their histograms are too close, we say their relative contrast is small.

Zones Estimation as an Optimization. With the two measures defined, we formulate the best zone estimation as a graph-based labeling problem. Each zone region is regarded as a node and any two neighboring zone regions are connected by a link. The optimal labels $Z = \{z_i^*\}$ of nodes are the final desired zones. We define the Markov Random Field (MRF) energy function $E(Z)$ of the graph as:

$$Z^* = \arg\min_Z E(Z) = \arg\min_Z (\sum\nolimits_i E_i + \lambda \sum\nolimits_{i,j} E_{ij}),$$

where E_i is the data term of an individual region i, and E_{ij} is the pairwise term between two adjacent regions i and j. In our work, the data term and pairwise term are respectively specified by the form: $E_i = -log(P(i))$ and $E_{ij} = -log(P(i,j))$.

The likelihood $P(i)$ of a region i is measured by its visibility of details ν_i, the region size C_i (normalized by the whole image size), and the important region size θ_i (normalized by the whole image size). The important region is directly computed from the probability map of facial/sky detector. We take into account all the three factors:

$$P(i) = \begin{cases} \nu_i^s \times C_i \times \theta_i \times \rho\left(\hat{z}_i - z_i\right), & (z_i < \mathrm{V}) \\ \nu_i^h \times C_i \times \theta_i \times \rho\left(z_i - \hat{z}_i\right), & (z_i > \mathrm{V}) \end{cases}, \tag{2}$$

where z_i is the original zone, $\hat{z}_i$ is the new zone and $\rho(t) = 1/\left(1 + \exp(-t)\right)$ is a sigmoid function. The likelihood would encourage shadow/highlight regions to move to higher/lower zones. For mid-zones (zone V), it takes no effect.

The coherence $P(i,j)$ is defined by the change of relative contrast between two neighboring regions, from the original relative contrast d_{ij} (before the optimization) to the new relative contrast $\hat{d}_{ij}$ (after the optimization), which is denoted by

$$P(i,j) = C_j \times \mathcal{G}(\hat{d}_{ij} - d_{ij}), \tag{3}$$

where $\mathcal{G}(\cdot)$ is a zero-mean gaussian function with variance 0.15 and the weight C_j is used so that relatively smaller regions contribute less. The coherence would penalize the dramatic change of relative contrast.

To obtain the global optimum, we use a brute-force searching method to travel all combinations of zone candidates for all regions because the total number of zone regions is not very high after region merging. To automatically estimate the weight λ, we first calculate the sum of data terms and the sum of pairwise terms across all combinations of zone candidates. Then we set λ to the ratio of two summations. We found it works very well in our experiments and does not require any tuning.

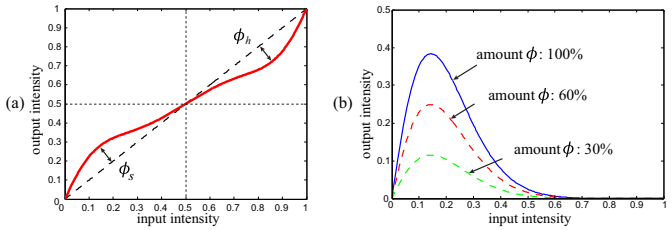

Fig. 6. (a) S-curve, ϕ_s, ϕ_h control the magnitude of S-curve adjustment in the shadow range and the highlight range respectively. (b) the curves of $f_\Delta(x)$ weighted by different amount ϕ.

5 Detail-Preserving S-Curve Adjustment

After getting the optimal zone for every region, we might have mapped the zone value (*i.e.*, exposure) of each region to its desired zone individually. However, this local mapping has the risk to produce exposure distortion in relatively small regions because these regions often contain insufficient information to estimate their optimal zones. To address this issue, we use a non-linear tone curve to globally map the brightness of every pixel to its desired exposure. We further preserve local contrast by fusion between the global curve mapping and an adaptive local detail enhancement.

S-Curve Adjustment. Most photographers often use an S-shaped non-linear curve (S-curve) to manually adjust the exposure in shadow/mid-tone/highlight areas. Fig. 6 (a) shows a typical (inverse) S-curve. This kind of S-curve can be simply parameterized by two parameters: shadow amount ϕ_s and highlight amount ϕ_h, which is denoted by:

$$f(x) = x + \phi_s \times f_\Delta(x) - \phi_h \times f_\Delta(1-x), \tag{4}$$

where x and $f(x)$ are the input and output pixel intensities. $f_\Delta(x)$ is the incremental function and empirically defined as: $f_\Delta(x) = \kappa_1 x \exp\left(-\kappa_2 x^{\kappa_3}\right)$, where κ_2 and κ_3 control the modified tone range of the shadows or highlights. We use the default parameters ($\kappa_1 = 5, \kappa_2 = 14, \kappa_3 = 1.6$) of $f_\Delta(x)$ to make the modified tonal range fall in $[0, 0.5]$. The effect of shadow/highlight amounts (ϕ_s, ϕ_h) is shown in Fig. 6 (b).

Inference of Correction Amounts. We infer the amounts (ϕ_s, ϕ_h) from the estimated optimal zone in every region. For the shadow regions, we want to set the amount ϕ_s so that the original zone value of each shadow region can be moved to its optimal zone value, as much as possible. The amount ϕ_h can be estimated in a similar way.

Suppose the original exposure and new exposure of a shadow region i are respectively e_i and $\hat{e}_i$. (The relationship between the exposure and its corresponding zone value is shown on Fig. 3(d)). The original exposure is calculated by the intensity mean: $e_i = \sum I/c_i$, where I is original intensity and c_i is the region size. After the S-curve adjustment (by Eqn. 4), the new exposure $\hat{e}_i = \sum f(I)/c_i = \sum (I + \phi_s \times f_\Delta(I))/c_i$. Thus, the shadow amount ϕ_s of this region should be: $\phi_s = (\hat{e}_i - e_i) \times c_i \times \sum f_\Delta(I)$. To consider all regions, we take the weighted average of the estimated shadow amounts of all regions. We use the percentage of region size as the weight.

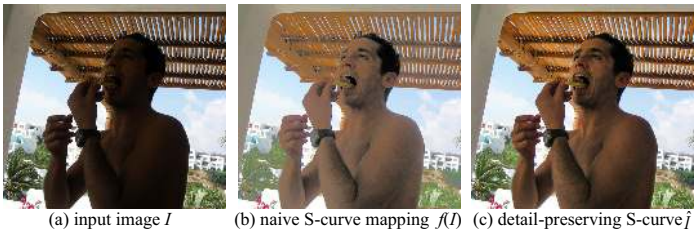

(a) input image I (b) naive S-curve mapping $f(I)$ (c) detail-preserving S-curve $\hat{I}$

Fig. 7. Comparison between direct S-curve mapping and detail-preserving S-curve adjustment

(a) input image (b) using Gaussian filter (c) using guided filter

Fig. 8. Comparisons of halo effects reduction between Gaussian filter and guided filter [30]

Detail-Preserving S-Curve Adjustment. If we directly apply the S-curve mapping (in Eqn. 4) to the input image, we may lose local details. Fig. 7(b) shows such a case, where the result looks too flat although dark areas are lightened. This undesired effect is due to: moving the intensities from shadows and highlights to the middle will compress the mid-tones. Since the S-curve is usually monotonic, the contrast between two neighboring pixels in the mid-tones could be reduced.

To address this issue, we propose a detail-preserving S-curve adjustment. Given an input image I, we adaptively fuse its S-curve result $f(I)$ with a local detail image ΔI. Note that ΔI is the difference between the input image I and its low-pass filtered version I_F: $\Delta I = I - I_F$. Here, we compute I_F by a fast edge-preserving low-pass filter, the so-called guided filter [30] to suppress halo effects. In Fig. 8, we show the result against a Gaussian filter. In our implementation, the radius is set to 4% of the short side of the image I. The final output image $\hat{I}$ is a weighted linear combination:

$$\hat{I} = f(I) + [2 \times f(I)(1 - f(I))] \times \Delta I, \tag{5}$$

where the second term on the right side adaptively compensates for the reduction of local details. The weight $f(I)(1 - f(I))$ reaches its maximum (when $f(I) = 0.5$) in the mid-tone range where there is notable loss in local details. In other words, we add more details back to the mid-tone than the shadow or highlight range. Specially in smooth regions, the output is mainly determined by the S-curve results. Such an adaptive adjustment mechanism can help us produce more natural-looking results (Fig. 7(c))

For a color image, we need to compensate the possible reduction of color saturation caused by the luminance adjustment, especially on shadows. To avoid this issue, we transform it to YIQ color space and then scale the corresponding I, Q chroma values by the adjustment of Y luminance values.

Efficient Implementation. For efficient computation, we enforce two extra constraints to largely reduce our search space of possible zone values: 1) Our adjustment uses the global S-curve which would map the same input pixel values to the same output. Thus we can consider the change of zone should be the same for the regions with the same original zone values; 2) Since our employed S-curve won't change values across the middle gray (0.5), we can consider that the change of every zone is not allowed across zone V. In addition, our exposure is evaluated on the down-scaled image with their long edge no more than 400 pixels. So our segmentation and face/sky detection can be very efficient. For an 16-megapixel RGB image, the whole evaluation and correction time is 0.3 second on Core2 Duo CPU 3.16GHz with single-thread, no SSE acceleration.

6 Experiments

6.1 Usability Study

Dataset: We perform our evaluation using a database of 4,000 images taken by our friends (including amateur and professional photographers) with direct camera output. These images varies on scenes, locations, lighting conditions and camera models (*e.g.*, DSLR, compact, mobile cameras). We ask five subjects to divide all images into three groups according to different extents of exposure problem. Three groups are "severely badly-exposed, definitely need correction" (Group A), "slightly badly-exposed photos, may require a little correction" (Group B), and "well-exposed, no more correction" (Group C). Finally, we obtain three different datasets respectively: "Group A" (975 images), "Group B" (1,356 images) and "Group C" (1,669 images) according to the majority agreement of five subjects. Fig. 9 (a) shows several examples.

Procedure: We will compare with automatic exposure corrections in several popular photo editing tools to manifest our method would become a better candidate. All of results are achieved by default parameters. We invite other 12 volunteers (7 males and 5 females) with balanced expertise in photography and camera use to perform pairwise comparison between our result and one of three other images: 1) input image, 2) result by Windows Live Photo Gallery's *Auto-adjust, exposure only* (http://download.live.com/photogallery), 3) result by Google Picasa's *Auto-contrast* (http://picasa.google.com/). For each pairwise comparison, the subject has three options: better, or worse, or no preference. Subjects are allowed to view each image pair back and forth for the comparison. To avoid the subjective bias, the group of images, the order of pairs, and the image order within each pair are randomized and unknown to each subject. This usability study is conducted in the same settings (room, light, and monitor).

Usability Study Results: The main user study results are summarized in Fig. 9 (b). Each color bar is the averaged percentage of the favored image over all 12 subjects (I-shape error bar denotes the standard deviation). From results on "All Groups" (without distinguishing the photos from different groups), we can see that the participants overwhelmingly select our result over the input (70.2% vs. 5.9%), Photo Gallery (60.5% vs. 29.6%), and Picasa (58.3% vs. 12.5%).

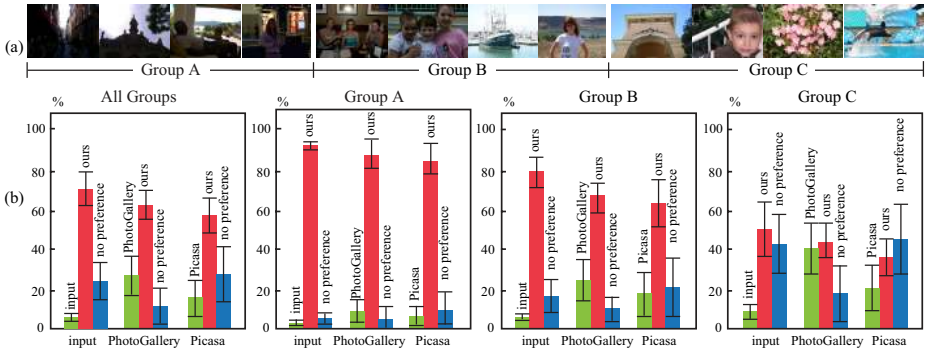

Fig. 9. Usability studies. (a) Examples from three groups: A (severely badly-exposed), B (slightly badly-exposed), C (well-exposed). (b) pairwise comparison of ours against the input, Photo Gallery, and Picasa, in all groups and three different groups respectively. Each color bar denotes the average percentage of favored image (with I-shape standard deviation bars).

Fig. 10. Examples randomly chosen from Group **A**. We can notice more details on foreground faces (a), foreground audiences (b) and street scene (c). **(Better View in Electronic Version)**.

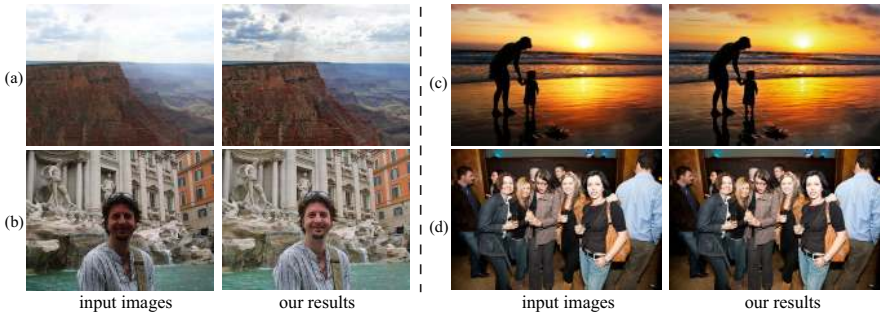

Fig. 11. Two examples randomly chosen from Group **B** (a-b) and two from Group **C** (c-d)

"Group A" results show that our approach works significantly better for severely badly-exposed photos. The participants show a strong bias in preference towards our correction when compared to input images (92.3% vs. 2.7%) and other automatic tools (87% vs. 8.5% against Photo Gallery, 84% vs. 6.4% against Picasa). The results from "Group B" indicate that slightly badly-exposed photos can benefit more from our correction than other methods as well. In "Group C", our approach also performs very well - for near 92% photos, our method does not make the result worse. It is quite nontrivial and very important for practical use, especially for batch-processing photos.

Fig. 9 (b) also graphically show two phenomenons on "Group C" compared with "Group A": 1) the margin between our result favored and no preference is smaller, and 2) all standard derivations are larger. They both indicate that the exposure correction itself is somewhat *subjective* especially for "not bad" photos. Subjects show different tastes for good photos correction, which has been discussed in [9][4], but most of these subjects consistently agree with our correction for relatively bad photos.

After the user study, we also ask all participants to articulate the criteria for their feedbacks. We conclude the main criteria: 1) the over-/under-exposed regions of interest should be well corrected; 2) well-exposed regions should not be over-corrected; and 3) the colors in corrected images should look natural. Other feedbacks include "the color of a few individual regions sometimes looks slightly unrealistic", "in some cases, the corrected results bring in some noise", and "I want some parameters tuning so that I can control the results.". Overall, most participants like our correction and want to use it for their daily photos processing.

Visual Quality Comparisons: Fig. 10 shows three examples from "Groups A". These photos show several common badly-exposed scenarios, such as outdoor backlit, dim-light indoor environment, which are very challenging for existing tools. As we can see, their corrections take no effect, but our method brings more visible details into badly-exposed areas while preserving the original appearance in well-exposed areas. Fig. 11(a)(b) show two examples from "Groups B", whose exposures look somewhat problematic. Our results look much more appealing, especially on important areas, *e.g.*, over-exposed sky (Fig. 11(a)) and under-exposed face (Fig. 11(b)). Fig. 11(c)(d) show two well-exposed examples from "Groups C". Our corrections seem to be imperceptible because the dark silhouette regions (Fig. 11(c)) have few detectable visible details and the black clothes (Fig. 11(d)) have lower priority than well-exposed faces, which would contribute little to the change of zone in our optimization.

6.2 Comparisons with Other Academic Methods

In consequent comparisons, our results are generated by the same parameters used in useability study. In Fig. 12(b)(c), we compare with two traditional histogram equalization algorithm [1][2] (by Matlab function *histeq*, *adapthisteq*). We can notice local contrast reduction and undesired halo effects in their correction results shown in Fig. 12(b)(c). However, our result shown in Fig. 12(e) looks more natural. We also compare our method with a well-known tone-mapping operator [21] (shown in Fig. 12(d)). Since their automatically estimated scene key is not accurate and tends to be higher than the actual key in this case, their result looks a little over-exposed.

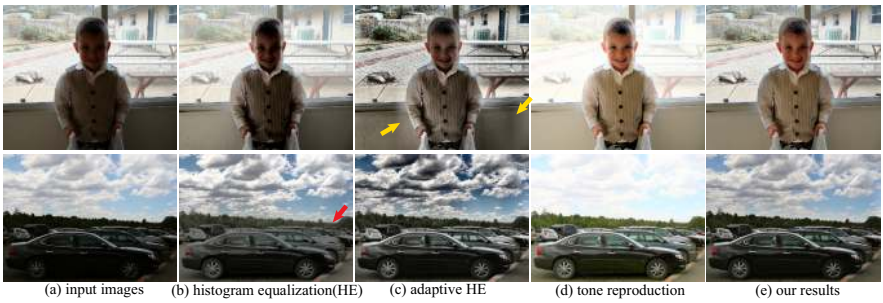

Fig. 12. Comparisons with histogram equalizations [1], adaptive histogram equalization [2] and tone reproduction [21]. The yellow/red arrows show unwanted halo effect/contrast reduction.

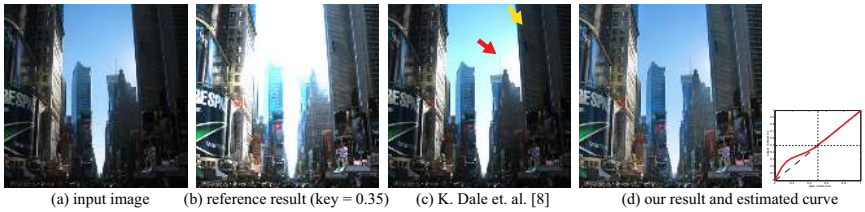

Fig. 13. Comparison with internet-based restoration [23]. Images (a-c) are taken from their paper. The reference result (b) is applied a fixed key. The yellow/red arrows show under-/over-exposed.

Fig. 14. Comparison with Exposure Fusion [22] on input image sequence (taken from their paper). Our algorithm only uses the single frame (b) as the input.

Fig. 15. Comparison with Exposure Fusion [22] on a single input image. We only use the input image (depicted in Green box) while Exposure Fusion uses the synthesized image sequence with different exposures from the input image. The red arrow shows unwanted artifacts.

Fig. 16. Comparisons with learning-based tonal adjustment [4]. Images (a,b,d) are taken from [4].

Fig. 17. Our failure case on noise amplification

In Fig. 13, we directly use the image and result from internet-based image restoration [23] for comparison. In this case, we can see our result has more visual details in local under-exposed areas than their provided result. Besides, their approach exaggerates over-exposed sky areas while our method can preserve their original appearance.

Exposure Fusion [22] is a fairly new concept that fuses all well-exposed regions together from a series of bracketed exposures. The good exposure is measured by some features: contrast, saturation and closeness to middle gray. Fig. 14 shows an example from their paper. We can see our result is visually approaching theirs, but our input is only a single frame from their input sequence. To perceive how well their algorithm works on a single input image, we make a modification of their method for comparison: (1) applying a series of global brightness adjustment (*e.g.*, multiplying luminance with 1/4, 1/2, 1, 2, 4) in Fig. 15(a); (2) applying a set of different gamma curves (*e.g.*, gamma values -3, -1.5, 1, 1.5, 3) in Fig. 15(b). Their results look either less vivid, or have lower global contrast than ours.

We show the comparison with learning-based adjustment [4] and assisted correction by expert in Fig. 16. As we can see, our result has more luminance details than their result on under-exposed areas and even much closer to the assisted result (from "Retoucher E" mentioned in [4]). Here, please ignore the difference in colors and focus on the luminance modification since the assisted adjustment includes both exposure correction and white balance. Without the need of training images, our approach obtain appealing results as well.

Fig. 17 (d) shows the limitation of our method. Since the correction does not consider the noise issue in our exposure evaluation, noise would become noticeable after we lighten dark areas. The issue may be addressed by suppressing the excessive noise amplification or applying denoising for these regions as preprocessing. We will further explore this issue in the future work.

7 Conclusions

We have presented an automatic method for the exposure correction of consumer photographs. The heart of this method is an optimization-based exposure evaluation and a detail-preserving curve adjustment algorithm. By simultaneously considering visible details in each region and relative contrast between regions, we are able to obtain appropriate exposure at the region level and produce natural-looking results.

References

1. Russ, J.C.: Image Processing HandBook, 3rd edn. CRC Press (1998)
2. Zuiderveld, K.: Contrast limited adaptive histograph equalization. In: Graphic Gems IV, pp. 474–485. Academic Press Professional, San Diego (1994)
3. CS5, A.P.: Adjust color and tonality with curves, Adobe Systems Inc. San Jose, CA (2010)
4. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustments with a database of input/output image pairs. In: CVPR (2011)
5. Ansel, A.: The Negative. The Ansel Adams Photography Series (1981)
6. Battiato, S., Bosco, A., Castorina, A., Messina, G.: Automatic image enhancement by content dependent exposure correction. Journal on Applied Signal Processing 12, 1849–1860 (2004)
7. Bhukhanwala, S.A., Ramabadran, T.V.: Automated global enhancement of digitized photographs. IEEE Trans. on. Consumer Electronics 40, 1–10 (1994)
8. dong Lee, K., Kim, S., Kim, S.D.: Dynamic range compression based on statistical analysis. In: ICIP (2009)
9. Kang, S.B., Kapoor, A., Lischinski, D.: Personalization of image enhancement. In: CVPR (2010)
10. Sobol, R.: Improving the retinex algorithm for rendering wide dynamic range photographs. Journal of Electronic Imaging 13(1) (2004)
11. Chesnokov, V.: Dynamic range compression preserving local image contrast. GB Patent 2417381 (2006)
12. Jobson, D.J., Rahman, Z., Woodell, G.A.: Properties and performance of a center/surround retinex. IEEE Trans. on Image Processing 6, 451–462 (1997)
13. Nayar, S.K., Branzoi, V.: Adaptive dynamic range imaging: Optical control of pixel exposures over space and time. In: ICCV (2003)
14. Shimizu, S., Kondo, T., Kohashi, T., Tsuruta, M., Komuro, T.: A new algorithm of exposure control based on fuzzy logic for video cameras. In: ICCE (1992)
15. Yang, M., Wu, Y., Crenshaw, J., Augustine, B., Mareachen, R.: Face detection for automatic exposure control in handheld camera. In: ICVS (2006)
16. Ilstrup, D., Manduchi, R.: One-Shot Optimal Exposure Control. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part I. LNCS, vol. 6311, pp. 200–213. Springer, Heidelberg (2010)
17. Brajovic, V.: Brightness perception, dynamic range and noise: a unifying model for adaptive image sensors. In: CVPR (2004)
18. Safonov, I.: Automatic correction of amateur photos damaged by backlighting. GraphiCon (2006)
19. Mukherjee, J., Mitra, S.K.: Enhancement of color images by scaling the dct coefficients. IEEE Trans. on Image Processing 17, 1783–1794 (2008)
20. Ovsiannikov, I.: Backlit subject detection in an image. US Patent 7813545 (2010)
21. Reinhard, E., Stark, M., Shirley, P., Ferwerda, J.: Photographic tone reproduction for digital images. SIGGRAPH (2002)

22. Mertens, T., Kautz, J., Reeth, F.V.: Exposure fusion. In: Pacific Conf. on Computer Graphics and Applications (2007)
23. Dale, K., Johnson, M.K., Sunkavalli, K., Matusik, W., Pfister, H.: Image restoration using online photo collections. In: ICCV (2009)
24. Guo, D., Cheng, Y., Zhuo, S., Sim, T.: Correcting over-exposure in photographs. In: Proc. IEEE CVPR, pp. 515–521 (2010)
25. Lischinski, D., Farbman, Z., Uyttendaele, M., Szeliski, R.: Interactive local adjustment of tonal values. ACM Trans. on Graph. 25(3), 646–653 (2006)
26. Xue, S., Agarwala, A., Dorsey, J., Rushmeier, H.: Understanding and improving the realism of image composites. ACM Trans. Graph. (2012)
27. Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient graph-based image segmentation. IJCV 59 (2004)
28. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: CVPR (2001)
29. Tao, L., Yuan, L., Sun, J.: Skyfinder: Attribute-based sky image search. ACM Trans. Graph. 28(3), 68:1–68:5 (2009)
30. He, K., Sun, J., Tang, X.: Guided Image Filtering. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010, Part I. LNCS, vol. 6311, pp. 1–14. Springer, Heidelberg (2010)