

An Exposure Adjustment Algorithm Which Can Be Applied Concurrently with JPEG Image Coding

James R. Boyack and Ibrahim Hajj-ahmad
Polaroid Corporation
Cambridge, MA

Abstract

Various electronic imaging devices such as Electronic Still Cameras, scanners and printers support and implement the JPEG image-coding standard. Coders and decoders are widely available in both software and hardware. In many imaging systems, the visual quality of some images is lower than expected because of exposure errors. In cases such as ESC or scanners, where the images have to be compressed using the JPEG standard, we propose a new method to correct the exposure of the images while performing the image coding steps. The correction is done automatically using data in the luminance channel of the DC band. Analysis of a histogram created by plotting the number of over-threshold local variations in luminance against the luminance values around which these variations occur yields an estimate of an appropriate shift which can be applied to the DC luminance channel. Simulation results show that the proposed method improves exposure accuracy as compared to traditional methods.

Introduction

When a digital image is captured, by an electronic camera or a scanner, the exposure control is not always optimal. The object of a Scene Analysis lightness adjustment is to refine the effective exposure by means of a mathematical analysis of the image data alone in order to assure the highest possible quality when the image is displayed or printed. The lightness is set by determining some control metric specific to the scene in question and comparing it to a reference point which is a constant of the imaging system.

The simplest method of exposure control is just to take the overall luminance average of the scene and set it to coincide with the imaging system's reference grey value. This works well for scenes wherein the average luminance of the principle subject is the same as the overall average. And this is often the case. However, this method is ineffective for 'difficult' scenes which are back-lit or contain unbalanced proportions of light and dark areas in the background or have specular reflectances which can unduly affect the overall average upon which the simple exposure meter is based.

More sophisticated exposure control systems are typically just extensions of the simple method which use more complicated averaging schemes, usually based on the idea that if one could locate the principal subject in the scene, then the average of the subject luminance could become the control metric. Polaroid cameras, for example, use a center weighted luminance average, since the principal subject is often placed near the center of the picture. The highest weighting is applied a little below the geometrical center in order to reduce the influence of a bright sky, which might be in the background. Other methods segment the scene into a pattern of central and peripheral areas and determine their control metric from some logical combination of the luminance averages of these areas. These refinements, though they represent a considerable improvement, are, when presented with a 'difficult' scene, still subject to the errors pointed out above which are inherent in any method which depends upon large area luminance averaging.

Exposure Adjustment Algorithm Through Scene Analysis

Our eyes sense a scene as a luminance bitmap. But we do not perceive or understand it that way. Rather, we know what the image is about by reading the 'information elements' it contains - outlines of forms, edges, textures, shadings. The quality of an image of a scene depends critically upon how well and to what extent these information elements are represented in the displayed image. The optimal exposure, we hypothesize, will be the one which best preserves the critical information elements of the scene in easily readable form.

We have taken a step away from the large area luminance average and base our approach on the concept of the 'information elements' introduced above. These elements all have in common that they are represented by *changes* in luminance (which we shall call 'activity'). There are three parameters which define an activity element - the magnitude of the luminance change, a characteristic luminance value at which the change occurs, and the geometric distance over which it occurs.

The Activity Based Histogram

The algorithm which we have developed uses image data expressed as the logarithm of world luminance. In using this algorithm with JPEG, we use the DC band of the luminance as the input. The DC band is obtained by collecting the DC coefficients from each of the 8X8 DCT blocks. Each DC coefficient is 8 times the average of the 64 pixels in the corresponding block. This low resolution image is divided into 2X2 sectors. The magnitude of the sector activity is taken as

$$\text{Activity} = \text{Max} - \text{Min},$$

where *Max* and *Min* are the maximum and minimum luminance values of the four 8X8 averages in a sector. The characteristic luminance for that activity can be taken as

$$\text{Activity Luminance} = (\text{Max} + \text{Min})/2.$$

Since a luminance change which is not noticeable is not important, we count only activities whose magnitudes exceed some threshold value. Our best results have come using a threshold of about 1/3 stop.

This activity metric is a non-linear, omni-directional 'detail' finder which has some sensitivity to all of the different types of information elements mentioned above, on a half wave scale of eight (high resolution) pixels. We form a histogram by counting the number of over-threshold activities as a function of the *Activity Luminance*. This forms a luminance histogram of image detail and is the basis of our analysis to estimate the optimal exposure. Notice that large light or dark areas which contain no detail will not affect our result.

If the dynamic range of the detail luminance histogram is the same or smaller than that of the print or display medium, then it is only necessary to move the histogram on a logarithmic luminance scale such that it fits within the print window, the range of luminance values which can be reproduced by a print or display device. The print window for a display system with a film-like characteristic curve having a slope of 1.5 is about 4 stops, considerably smaller than the range of a typical Activity Luminance histogram, which can be up to 2 or 3 stops greater than this. Since there is generally a considerable overlap outside the print window, some of the detail information must be lost - clipped by the limitations of the print mechanism. The question of the optimal exposure shift now becomes one of how to position the histogram over the print window to get the best result, recognizing that the information represented by the ends of the histogram which extend beyond the print window will be lost. An example of an image histogram based on Activity Luminance is shown below.

Determination of the Exposure Correction

We have applied four methods to position the histogram over the print window. They are:

- MidShift: Set the midpoint of the histogram range to the midpoint of the print window.
- MeanShift: Determine the weighted mean of the *Activity Luminance*, using the activity counts as weights and set it to the midpoint of the print window.
- MaxShift: Shift the histogram on the log luminance scale such that the maximum possible number of counts are inside the print window.
- EqEndShift: Shift the histogram such that the same number of counts are excluded from the print window at either end of the histogram.

These four parameters differ in their sensitivity to the range, the shape and the symmetry of the histogram. The best results using only one of the above parameters comes from the EqEndShift. Somewhat better results can be obtained either by averaging all four, or by averaging the MidShift, the MeanShift, and the EqEndShift.

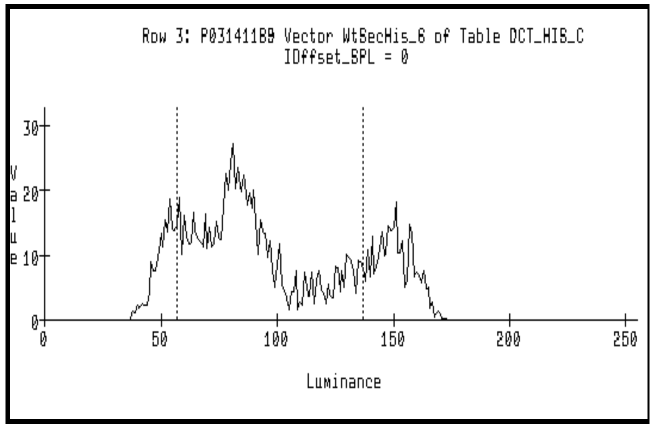

Sample Histogram: The Luminance scale is in logarithmic (base 2) units, 20 digits per stop with 100% reflectance at digit 144. The vertical dotted lines represent the print window. The histogram is placed on the luminance axis in the position of best exposure for this image.

Incorporation Into JPEG

JPEG (Joint Photographic Experts Group) is a widely used standard for color image compression. The main goal of image compression is to reduce the number of bits that are needed to describe an image at a specific level of quality or a given bit rate. Thus the time required for transmitting the image and the space needed to store the image are reduced. Many images, especially on the web, are coded using the JPEG coding standard where fast and efficient coders/ decoders are widely available in software and hardware. Also various imaging devices such as an ESC, scanners and printers support and implement the JPEG standard. The exposure adjustment algorithm that we have described in the previous sections can be used concurrently with the existing JPEG standard. The exposure adjustment can be done either as a part of the encoding or the decoding process depending on the applications. For example, for input devices such as an ESC, exposure adjustment could be done as a part of the encoding process, while for output devices such as printers it could be done as a part of the decoding. In this section, we will illustrate our method for ESC where the image will be enhanced by our exposure adjustment algorithm while it is being encoded and stored in the camera's memory.

A general block diagram for the encoder part of the JPEG standard is shown below. In this standard, there are three main components. The transform takes 8X8 DCT block of the image and maps the input image to a domain that is more suitable for coding. In the quantization block, the transform coefficients are represented with a fixed number of steps.

Finally, in the entropy coding block, any redundancy in the bit stream is removed by assigning short codewords for the frequently used steps and longer codewords for the other steps.

If we take a closer look at the Quantize block, we can see that the DC coefficients are separated from the AC coefficients where we perform a DPCM. The DC Coefficients of the DCT blocks of an image are coded differently from the AC coefficients. In order to reduce the redundancy among the DC coefficients, the difference of the current DC block is subtracted from the DC value of the previous block and the resulting differences are then sent to the entropy coding. Entropy coders such as Huffman or arithmetic coding are used to assign specific code words for each of the possible quantization steps or values.

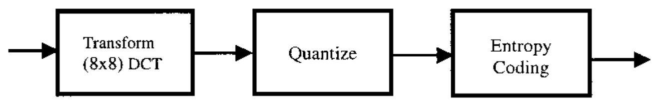

Block Diagram of JPEG Encoder

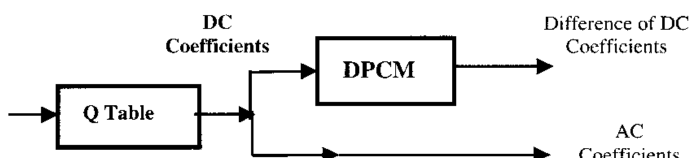

Quantization Step in JPEG

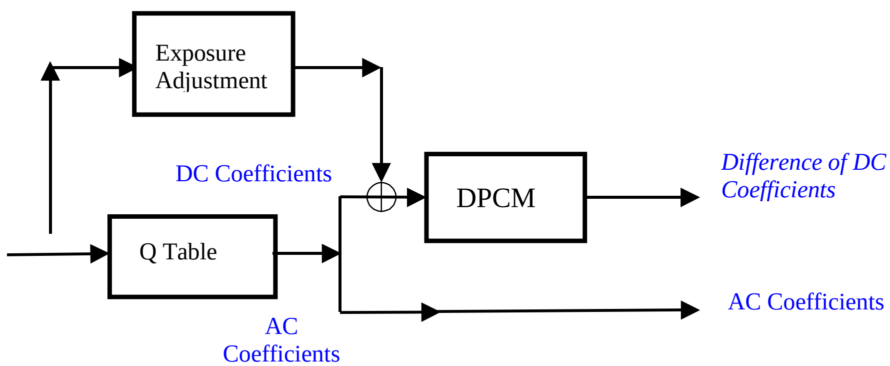

Exposure Adjustment with JPEG Encoding

The figure above shows how our exposure adjustment algorithm can be used with the JPEG encoder. Since our algorithm works only on the luminance channel, by automatically adjusting it with the appropriate correction it affects only the values of the DC coefficients of the luminance channel. Applying the Scene Analysis correction in the JPEG coding system is done by performing a global shift (adding the correction value) to all of the DC values. However, because the DC band in JPEG is coded using the DPCM, we only need only to adjust the value of the first block. This will adjust the DC values of all the remaining blocks because they are calculated recursively as the difference with the DC value of the previous block. Thus we obtain a JPEG coded image that has been adjusted for exposure error without affecting the bit rate.

Testing and Evaluation

The performance of the methodology outlined above has been tested using a library of digital images. This consists of a set of digital images which includes an oversampling of "difficult" scenes—snow scenes, beach and water scenes, high contrast back- and side-lit scenes, etc.

Each image was obtained by photographing a scene using Kodak VPS color negative film and scanning the processed negative with a high resolution Itek scanner. The scanner was calibrated to world XYZ by photographing a set of known color reference patches over a broad exposure range, scanning the resulting negatives, and correlating the digital data to the reference color values. Having this calibration enables printing of the digital scanned images on a similarly calibrated high quality color printer.

The digital image library contains about 600 images of a wide variety of scenes obtained in this way. Each one was subjected to psychovisual testing with a panel of 25 subjects who indicated their preferred exposure from a series of prints of the same scene made at eight equally spaced levels of lightness. This provided the 'Psychovisual Best' exposure data against which we can compare the estimates from our Scene Analysis exposure algorithm. The differences between the Scene Analysis estimate and the corresponding Psychovisual Best exposure were determined for 594 images. Success is measured by the standard deviation of these differences. Some results are summarized below. The first entry is the result for the Spectra camera center weighted luminance average. The last two entries are the result of averaging the exposure estimates from three, or all four of the Scene Analysis histogram evaluation methods.

Results of Scene Analysis Algorithms

Values are log Exposure, 20 digits per stop.

	Std. Dev	Max. Error
Spectra Center Wt.	13.7	53
MidShift	12.5	51
MeanShift	11.3	34
MaxShift	11.3	42
EqEndShift	10.3	35
Averages		
EqEndMeanMidAve	10.0	30
AllAve	10.0	31

Conclusions

Almost any method of adjusting exposure—even a good guess based upon prior experience and personal judgement -- can give quite acceptable results for a majority of scenes, since the most common scenes tend to contain a well distributed set of lightness values. Improvements in exposure control reduce the variability in the common, “easy” scenes and bring a greater range of “difficult” scenes into the acceptable and higher quality categories. Our Scene Analysis method reduces the maximum error in our library image set from over 2.5 stops for the Spectra center-weighted average to just 1.5 stops and makes a 30% reduction in the standard deviation of the exposure error. Also, this algorithm can be applied concurrently with the widely used image coding standard, JPEG. This allows us to leverage the existing fast software implementations and cheap and powerful hardware for the DCT. Further development of this method using the same principles should yield still better results.

References

1. J. Bullitt and J. Thornton, IQA: A One Step, Automatic, Image Processing System, Proc. IS&T 51th Conf., 1998.
2. J. Boyack, US Patent # 5,724,456, 1998.
3. W. Pennebaker and J. Mitchell, JPEG Still Image Compression Standard, Van Norstrand Rienhold, 1993.