

Digital Processing of Scanned Negatives

Qian Lin and Daniel Tretter
Hewlett-Packard Laboratories
Palo Alto, CA, USA

ABSTRACT

One source of high quality digital image data is scanned photographic negatives, which can be processed to produce high quality color images. The scanned data must be inverted and processed to adjust for the film and scene characteristics. This paper proposes an approach to processing scanned negatives with unknown film type to produce output color images suitable for viewing on a computer monitor or for printing.

Our processing pipeline contains an adaptive stage that automatically adjusts the white and black point according to the image characteristics. Other stages invert the scanned data and adjust the midtone values. Finally, a postprocessing stage is used to detect dark and backlit scenes, which are then brightened. The pipeline has been tested on several hundred scanned negatives using two different film scanners.

Introduction

Digital processing and archiving of photographic images has several advantages over traditional analog means. Digital processing allows more flexibility in adjusting the image to produce higher image quality or to better match individual user preferences. Storing the images in digital form allows them to be cataloged in an orderly fashion for efficient searching and retrieval using standard database tools. With the development of consumer grade photographic printers, digital storage also permits the user to make multiple prints easily and inexpensively without going through a photofinisher. One way to convert existing photographic images from analog to digital form is to scan an analog negative or reflective print using a digital scanner.

Prints are generally easier to digitize because the colors have already been adjusted for viewing and the originals are large enough to give a high resolution digital image with a modest scanning resolution. Unfortunately, prints have considerably less dynamic range than the original negatives, so they do not reproduce specular and semispecular highlights or shadow regions well [Jam66]. The resulting digital images inherit any shortcomings of the print as well as any artifacts in the negative, so higher quality images can often be produced by scanning the negative directly.

Figure 1 illustrates a typical 35mm film scanner suitable for scanning negatives. A negative strip is inserted into the

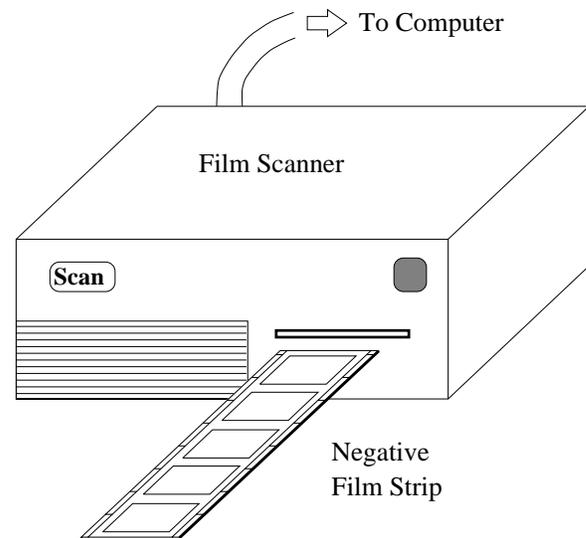


Figure 1: A film scanner can produce a digitized image from a film negative.

scanner, usually after manually placing it in a special film holder. The recently introduced APS (Advanced Photo System) film format does not require the user to handle the negatives directly, so a scanner for this format would likely be simpler to use [Tui96]. A negative frame is scanned and converted to a digital image, which is processed and transferred to a host computer. The processing can either take place on the scanner or as a software application on the computer. To produce a viewable image from a scanned negative, the data must be inverted and processed to adjust for the film and scene characteristics. In this report we describe a proposed processing pipeline to automatically convert scanned negative images to output images suitable for viewing on a computer monitor or for printing.

Both film characteristics and scene conditions can vary widely from frame to frame, so any automatic processing algorithm must adapt to each scan based on image statistics. Human intervention can also be used to adapt algorithm parameters, but for consumer or high volume applications, we want to minimize human intervention. Several researchers have proposed techniques to adjust image colors based on the image data.

Evans uses a “grayworld” assumption, where all the pixels in an image are assumed to average to gray[Eva51]. This assumption can be too restrictive, however, and it can cause undesirable color errors in some processed images. More recently, a number of authors have proposed schemes for estimating a scene illuminant from the image data[Lee86, Fun95, Fin96]. These methods are often computationally complex, and they generally assume the data is in some known linear color space. The algorithms thus correct for scene illumination, but film characterization and correction would need to be done separately if one of these methods is to be used when processing a scanned negative. Funt, Cardei, and Barnard propose a neural network based approach to color constancy that does not have these built-in constraints[FCB96]. This approach might work well for processing scanned negatives, although as far as we know it has not yet been tested on this task.

Traditional analog processing of negatives includes characterizing and adjusting for the film type and scene data using both global statistics for the entire film strip and local statistics for the desired frame[Tui96]. Manual adjustments may be necessary for some images. Tuijn proposes a similar procedure for processing digital scans of negatives[Tui96]. He starts with predefined characteristic curves for negative film and modifies the curves to better reflect the individual strip being scanned. He then uses an algorithm called TFS (Total Film Scanning) to produce a virtual point on the negative which is mapped to a neutral reference point. He reports that this system performs well on more than 99 percent of the images scanned.

Our approach uses the red, green, and blue histograms for a scanned image to adapt to changing film and scene characteristics. The histograms are used to compute a white point and black point for each image scanned. Fixed lookup tables are then used to adjust the color balance in the image midtones. Finally, a postprocessing step is included which detects dark or backlit images and brightens them. Section 2 discusses our processing pipeline in more detail. Experimental results are given in section 3, and section 4 gives general conclusions.

Processing Pipeline

Our negative processing pipeline consists of four stages; one adapts to the image data, while the other three are fixed lookup tables. We also implement a separate postprocessing step that adjusts for dark and backlit images. Figure 2 illustrates the four processing steps in the pipeline. The graphs all assume 8-bit per color plane scan data; axis numbers need to be scaled for other bitdepths. The first stage is a simple inversion step, converting the scanned image from a negative to a positive. The input/output characteristic shown in Figure 2 is applied to all three scanned colors.

The second pipeline stage remaps the white and black point of the image based on the red, green, and blue his-

tograms. This mapping function is therefore adaptive, adjusting to the measured image statistics. The operation remaps the image data as illustrated in Figure 2. Using the red, green, and blue histograms, a bounding box is constructed containing the scanned data values. The data is then remapped to fill the entire RGB cube. In its simplest form, this amounts to setting the maximum red, green, and blue (R , G , and B) values to the white point and the minimum values to the black point.

$$\begin{aligned} \{R_{MAX}, G_{MAX}, B_{MAX}\} &\Rightarrow \text{White} \\ \{R_{MIN}, G_{MIN}, B_{MIN}\} &\Rightarrow \text{Black} \end{aligned}$$

This remapping of the white and black points helps adjust for both scene illumination and mask density of the negative. The procedure assumes that the minimum and maximum scene reflectances are indicative of white and black points. It also assumes that only reflected light is in the scene. Any light sources or fluorescent objects can result in a color cast to the image or in an image that is too dark. Finally, if the absolute maximum and minimum measured values are used for white and black points, any noise in the image, including that due to dust or scratches on the negative film, can cause errors in the remapping.

We alleviate some of these potential problems by implementing a more robust form of the white and black point adjustment. White and black point estimation in the presence of noise is improved by using red, green, and blue values that are somewhat below the absolute maximum and above the absolute minimum. For example, we can map the 95th percentile R , G , and B values to white and the 5th percentile values to black. This eliminates outliers due to noise.

The white point computed for an image can result in an undesirable color cast or underexposure if the scene contains light sources or fluorescent objects. Color casts can be reduced by identifying when the computed white point lies outside of some expected range of values. Although the white point can vary considerably from image to image, a large number of processed images can be used to compute a distribution of likely white point values. If the computed white point for a particular image lies far enough outside this distribution, we adjust the value to pull it in closer to the expected range before remapping the data.

If the white and black points are remapped so the bounding box expands to fill the entire RGB cube, detail can be lost in the highlight and shadow regions of the image. The highlights will be pushed into saturation, while shadow regions will have remapped pixel values very near zero. In order to prevent this phenomenon, which is sometimes referred to as “blowing out” the highlights, our remapping function contains soft shoulders. The bounding box is expanded to fill most of the RGB cube, but a small set of overload values is left around all sides. These values are used to represent the highlights and shadows that would otherwise be clipped. A one dimensional representation of this technique is given in

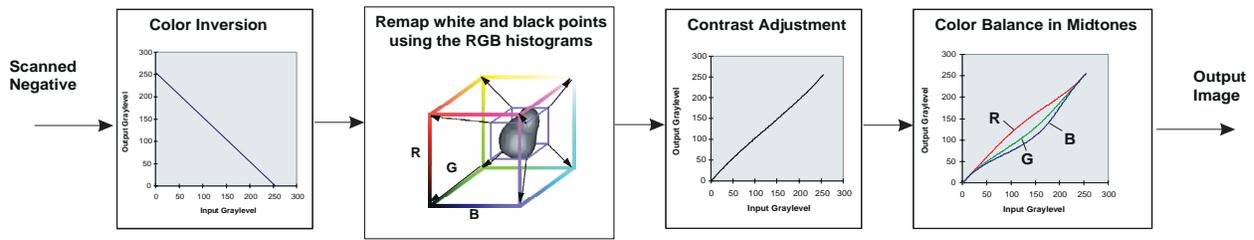


Figure 2: Our processing pipeline contains four stages, one of which adapts to the image data.

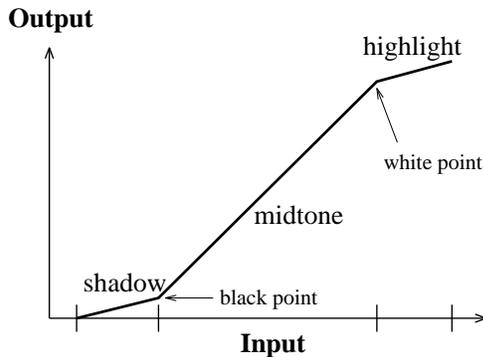


Figure 3: We use soft shoulders on our white/black point remapping to avoid clipping in the highlight and shadow regions.

Figure 3.

The white and black point mapping is followed by two stages that adjust the image midtone characteristics, as shown in Figure 2. The contrast adjustment step applies a nonlinear remapping function to all three color planes. This function is designed to adjust for some of the nonlinear characteristics of the negative. We use a parameterized set of inverse sigmoidal functions for contrast adjustment. If the input range x between shadow and highlight regions is normalized to $[0, 1]$, the function takes the form:

$$y = \begin{cases} 0.5 (2x)^{\gamma_c} & x \leq 0.5 \\ 1 - 0.5 (2 - 2x)^{\gamma_c} & x > 0.5 \end{cases}$$

The parameter γ_c is determined experimentally.

Finally, the color balance is adjusted in the midtone regions using a procedure suggested by Michael Stokes. The three preceding processing stages can all be incorporated into a set of three one dimensional lookup tables, one for each color plane. The curves corresponding to these lookup tables will be rotated clockwise by 45 degrees, added to a midtone adjustment curve, and rotated back 45 degrees counterclockwise. This procedure is illustrated for a single color plane in Figure 4. Our midtone adjustment curves are computed as

$$y = \beta \sin^2(\pi x),$$

where the input range is again normalized to $[0, 1]$, and the parameter β is determined experimentally for each color plane.

The midtone adjustment will balance the colors in the midtone regions to remove color casts along the neutral axis.

The entire pipeline of Figure 2 can be computed as a single image dependent input/output lookup table for each color plane.

Experimental Results

We used commercial film scanners to obtain our test image data. In order to prevent the scanners from applying their own correction algorithms on the data, the scanners were set to positive mode for scanning. The data we use in our experiments have therefore been processed by the scanner as if the originals were transparencies. Our experimental images are not raw data samples from the scanner, but have probably undergone a color space transformation and gamma correction. Due to the structure of our pipeline, however, we expect that results obtained during testing are similar to those we would get using raw data samples. The various pipeline parameters, particularly the β parameters and γ_c , could be adjusted appropriately to yield similar output image characteristics.

The algorithm was applied to a large collection of photos. The photos, taken by around 50 employees of HP, have a wide variety of subjects, including people, animals, outdoor scenes, night life, etc. The photos were scanned from Kodak Professional RFS 2035 Plus Film Scanner. There are a total of 443 photos. We also obtained the regular photo prints from a traditional photofinisher, as well as digital images processed by the Kodak scanner when we set the scanner in the automatic mode.

This is a very challenging collection of images, because of the wide variety of cameras used, the different illumination conditions, AND the deliberate selection of film types. The films include

- 3M 100
- 3M 200
- Agfa HDC 200
- Fuji Super HG 100
- Fuji Super HGII 100
- Fuji Super HGII 200
- Fuji Super G+ 100
- Fuji Super G+ 400
- Kodak Ektachrome 200
- Kodak Ektra 200

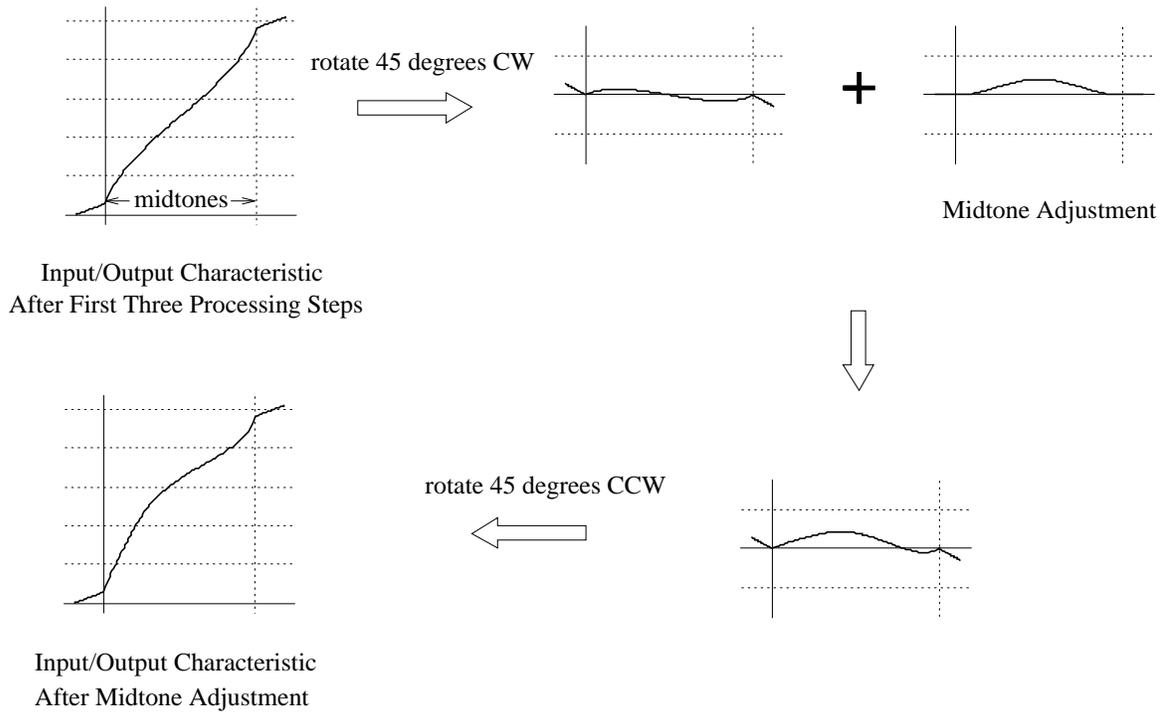


Figure 4: The midtone color adjustments are computed by rotating the input/output characteristic to horizontal, adding an adjustment function, and rotating back.

- Kodak Gold 100
- Kodak Gold 200
- Kodak Gold 400
- Kodak Gold Ultra 400
- Kodak Plus 100
- Kodak Royal Gold 100
- Kodak Royal Gold 200
- Kodak Royal Gold 1000
- MotoPhoto (Agfa) 100
- MotoPhoto (Agfa) 200
- Ritz 200

We applied the processing pipeline to the negative images. Figure 5 shows the Chromaticity coordinates a^* and b^* in the $L^*a^*b^*$ space of the white points for 329 images. The white points are calculated after the inversion process, hence the white points are located in the cyan region in the $L^*a^*b^*$ space.

Figure 6 shows the brightness coordinates L^* of the white points for 329 images.

From Figure 5 and Figure 6, it is clear that there is a large variation in the exposure, illumination, and film type from image to image. As a result, the required adjustments vary considerably across images. The location of the white point for a particular image is related to the illumination and film type. However, it appears that there is no clear separation that would allow one to identify the illumination or film type from the white point location. On the other hand, the distribution of the white points allows us to make the white point

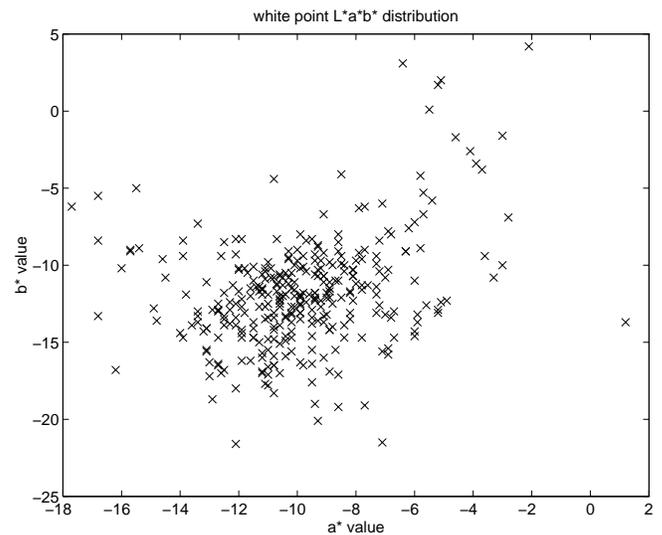


Figure 5: Chromaticity coordinates of the estimated white points

identification process more robust. This can be accomplished, for example, by checking the color of each image point and throwing out those that clearly are not potential white points in the histogram building stage. In one of the photos taken in a nightclub with neon lights in the background, for example, the bright red neon lights biased the white point toward the red color. Hence the image appeared greenish. With the more

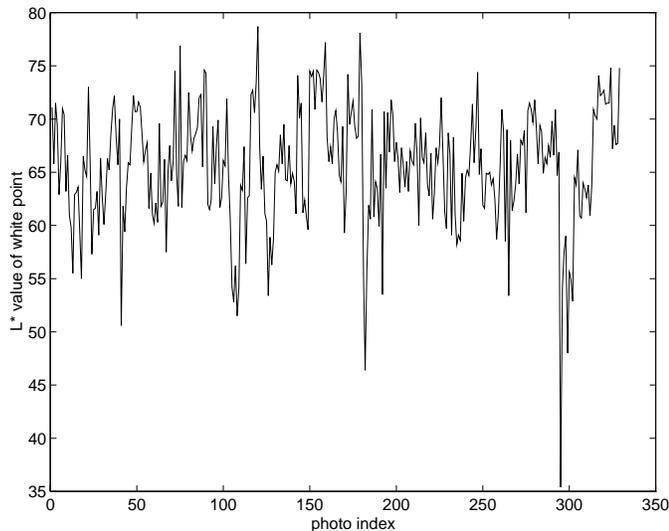


Figure 6: Brightness coordinates of the estimated white points

robust estimation, the image became much more natural.

Compared with the automatic processing implemented in the Kodak scanner, our processing pipeline produces image of good color balance and good contrast. Images processed by the Kodak algorithm have very strong red cast. Our algorithm is also very fast since it is only a one dimensional table look-up, while the table can be built with a much smaller subsampled image. The traditional prints, which use operator assistance with film type and possibly manual color/exposure adjustments, have good color balance and contrast, but often they have poor bright details. This is particularly obvious in some of the wedding photos, where details in the wedding gown failed to come up.

The exposure control in our algorithm is good for most images. However, a few images that were backlit came out too dark. The Kodak processing suffered the same problem, while the traditional prints were quite good. A separate post-processing step follows the pipeline to identify and brighten backlit and dark images.

Backlit images most often occur when the sky is in the background, and the foreground is not well lit. This can occur, for example, when an indoor photo is taken with a window to the outside somewhere in the scene or when an outdoor photo is taken with the foreground (often people) in shadow and the background a bright sunlit scene. The black and white point mapping can result in a dark image when light sources, such as lit candles or incandescent bulbs, are visible in the scene. Our postprocessing algorithm uses heuristics to identify these situations, which are then brightened using a nonlinear power mapping.

We use a set of six scalar image statistics, η_1, \dots, η_6 , to detect backlit and dark images. The statistics used by our

heuristic algorithm are:

- η_1 = # image pixels with $R > 200$
- η_2 = # image pixels with R, G , and $B > 200$
- η_3 = # image pixels with $R < 128$
- η_4 = # image pixels with $R < 128, G < 0.8R$, and $B < G$
- η_5 = percentage of pixels along the edges of the image with R, G , and $B > 200$
- η_6 = percentage of pixels along the edges of the image with R, G , and $B < 100$

If η_1 is less than a set threshold, the image is classified as too dark. These images are brightened using a power function with the power factor determined from image histograms such that the average output median is mapped to half scale.

The heuristic used to determine a backlit image is somewhat more complicated. We assume a backlit image has a bright background and a dark foreground, both of which extend to the edges of the image. If $1.2\eta_2 > \eta_1$ and $10\eta_4 > \eta_3$, the image may be backlit. These conditions establish that the bright image regions are near neutral in color (like sky) and a significant percentage of the darker regions have a reddish tint (like skin tones). In addition to these conditions, we also require $\eta_5 > 15\%$ and $\eta_6 > 15\%$. This requirement guarantees that the bright and dark regions each extend to the image edges. If all of the above conditions are met, the image is classified as backlit. Backlit images are brightened using a standard gamma correction curve for a fixed preset gamma.

We also applied the processing pipeline to some negative images scanned from Nikon Super CoolScan LS-1000. We found that the same pipeline can be used with minor adjustment of contrast and color balance parameters.

Conclusions

We presented a fully automatic processing pipeline for processing scanned negative images. The algorithm first builds up a histogram based on subsampled image data, then estimates white/black points of the image, and performed one dimensional look-up table operation to adjust for color balance, contrast, as well as black/white point mapping. A post processing step then follows to identify backlit images and performs further exposure adjustments. Without any knowledge of the film type, the processed images has better color balance than the corresponding ones processed by the Kodak automatic algorithm, and comparable to the traditional prints. Exposure control in our algorithm is comparable to the Kodak algorithm, but slightly falls short of the traditional prints. Further research is required to fully take advantage of the flexibility of digital processing to optimize image quality.

Acknowledgements

The authors would like to thank Michael Stokes for suggesting the procedure we use for midtone color balancing as well as for other valuable suggestions and comments. We would also like to thank Malcolm Rix for providing numerous test images and comments on negative image characteristics.

References

- [Eva51] R. M. Evans. Method for correcting photographic color prints, 1951. US Patent 2 571 697.
- [FCB96] Brian Funt, Vlad Cardei, and Kobus Barnard. Learning color constancy. In *The Fourth Color Imaging Conference: Color Science, Systems and Applications*, Scottsdale, AZ, 1996. IS&T/SID.
- [Fin96] G. D. Finlayson. Color in perspective. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(10):1034–1038, October 1996.
- [Fun95] Brian Funt. Linear models and computational color constancy. In *Proceedings of the 1995 Color Imaging Conference: Color Science, Systems and Applications*, pages 26–29, Scottsdale, AZ, 1995. IS&T/SID.
- [Jam66] T. H. James, editor. *The Theory of the Photographic Process, Third Edition*. The Macmillan Company, New York, New York, 1966.
- [Lee86] Hsien-Che Lee. Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, 3(10):1694–1699, October 1986.
- [Tui96] Chris Tuijn. Scanning color negatives. In *The Fourth Color Imaging Conference: Color Science, Systems and Applications*, pages 33–38, Scottsdale, AZ, 1996. IS&T/SID.