

Improving Scans of Black and White Photographs by Recovering the Print Maker's Artistic Intent

Daniel M. German*

dmg@uvic.ca
Dept. of Computer Science
University of Victoria, Canada

Jaume Rigau

jaume.rigau@udg.edu
Graphics and Imaging Laboratory
University of Girona, Spain

Abstract

In this paper we propose a method that reverse engineers the aesthetic decisions made by a print maker to produce a print from a negative, namely cropping, contrast selection, and dodging-and-burning. It then re-applies this process to the electronic negative in order to achieve an electronic version of such print with better tonal range and detail than one produced by scanning the print. We then extend this method to restore a print by combining scans of different versions of the same image.

1. Introduction

Individuals and institutions, such as the United States Library of Congress¹ are scanning their inventories of photographs, both negatives and prints. The most valuable negatives of a collection have usually been printed. It is not uncommon for a scan from a negative to look different than the scan of a print. This is often because prints (particularly those created in chemical darkrooms) were hand crafted, done by print makers (darkroom operators) who applied their skills and knowledge to try to create the best print possible from a negative (according her artistic vision). Negative scanning is preferred to print scanning because more detail can be recovered from the negative than from the print. Negatives (translucent media) have significantly higher dynamic range than photographic papers (reflective media); furthermore prints are a copy of the negative, and in the process some information is lost. Even with advanced scanners that are capable of recovering most of the information available from prints and from negatives, the best possible scan from a negative would usually contain more information than the best possible scan from one of its prints. And because there is significantly more dynamic range in the scan of the negative, its contrast will appear lower than the one of the print when they are visually compared. A human operator of image editing software is likely capable of turning a negative scan into a version that is very close to the look of the print, but it is time consuming and requires skill. Automatic scanning offers considerable

savings in time and money. It is desirable to automatically recover and reuse operations previously done by the print maker in the darkroom to reapply them to the negative to obtain a version that is as close as possible to the negative, but with the richness in tone and detail that the negative provides.

1.1. Print Making

The aphorism “the negative is the score, and the print is the performance” is credited to Ansel Adams [1, 2]. Photographs are rarely recognized via their negatives; what we recognize are their prints. Adams wrote “we start with the negative [...] and proceed through a series of ‘work’ prints to our ultimate objective, the fine print” [3]. It is sometimes desirable to digitally reconstruct, from a scan of its negative, a particular “fine print” as it was envisioned by its creator.

Often overlooked, the artistic decisions made by the print maker are what frequently set apart such fine prints. Print making is a craft. It requires a good understanding of the technical aspects of the printing process, an artistic eye that can apply such techniques to a negative to bring the best of it, and plenty of practice. Photographic printing merges art and craft to express an idea or a feeling that has been captured in a negative [2]. A print maker is responsible of making four important aesthetic decisions to create a specific print. These are [2]:

Cropping Many photographers object to having their negatives cropped and expect their photographs to always be printed with the edge of the negative (most notably Henri Cartier-Bresson). When a negative is cropped, is usually for one of these reasons: a) restrictions of the paper to use (for example, photo papers are rarely found in sizes

*Corresponding author

¹See Digital Collections & Services of the Library of Congress at <http://www.loc.gov/library/libarch-digital.html>.

proportional to 35mm negatives); b) for aesthetic reasons (to balance the image, to emphasize an element in it, or to exclude areas); c) to correct framing, such as a tilted negative; and d) to change the message of a photograph (by excluding information that might change the opinion of the viewer). Cropping is very subjective.

Brightness Determine the shadows (blacks), highlights (whites), and midtones (mid grays) of the print. Automatic printing and scanning often results in images that look too dark or too light. The perfect brightness of a print is highly subjective, and depends on both the contents of the image (e.g., snow versus a black subject) and the desired feeling that it should convey.

Contrast Determine the tonality of the print. The contrast of a print changes its mood and there are no rules with respect to what contrast brings the best from a print. Many photographers believe that every print should include almost any tone of gray from pure black to pure white (such as those using the Zone System developed by Adams [3]), while others take a totally opposite position and produce images with only pure white and pure black. The contrast of a print is usually determined by the type of paper used and how it is developed (the paper’s characteristic contrast curve).

Dodge-and-burn A print from a negative might require that some of its regions be lighter or darker. This is achieved by selectively under exposing (dodging) or over exposing (burning) such regions. A print that has not been dodged-and-burned is said to be a straight print. Determining the regions that should be dodged or burned is also a very subjective decision.

Scanning a negative and applying automatic midpoint and contrast adjustment produces the equivalent of a straight print, similar to those made by an automatic photo minilab: the scanning software has no knowledge of the contents of the negative and simply finds the white and black points, and applies a pre-determined contrast function. These scans will look different from prints that have been manually cropped or dodged-and-burned.

In this paper we propose a method that transforms a high-resolution scan of a negative into an image that matches the cropping, midtones, contrast and dodge-and-burning of one its prints—which we call the base print. The base print can be of significantly lower resolution and with loss of discernible detail in its brightest and lightest areas. We call the resulting image the match print of the negative and such base print. We extend this method to make it possible to combine two or more images in such a way that their sum is better than either one. We then propose several metrics to assess the effectiveness of the method presented.

1.2. Related Work

Few algorithms have been designed to automatically crop an image based upon its content but with little regard to their

aesthetic value [4, 5, 6]. Geigel and Musgrave [7] modeled the film development process and their goal was to create digital images that had the “look” of film. Reinhard et al. [8] proposed a method to automatically tone and dodge-and-burn a HDR image to a LDR; their method is based on the method used by Adams (the Zone System [3]) and attempts to provide a faithful reproduction of the original. Bae et al. [9] proposed a method to transfer the tonality of one image to another but it does not take into consideration dodge-and-burn. Automatic registration of different versions of the same image (even if they are of different resolutions) can be done using methods commonly used in the creation of panoramic images (see Szeliski [10] for an overview).

2. Model

Let neg be a negative image and $print$ a print originated from neg . The process of producing $print$ can be modelled as follows:

$$print \equiv paperF(projNeg, t)$$

$$projNeg \equiv scalecrop(proj(neg))$$

where t is the time of projection of the negative on the paper, $paperF$ is the contrast function of the paper used, and the $projNeg$ is the negative projected by the enlarger on the paper ($proj$), scaled and cropped ($scalecrop$).

The image of the negative projected by the enlarger ($projNeg$) is affected by the enlarger itself: the intensity of the light it produces; the optical characteristics of its lens (such as its optical distortion and the aperture used to make the print); and the geometric distortion produced when the plane of the film and the plane of the papers are not perfectly parallel. Also, if the negative bulges (something common in large negatives that are not held in place with glass carriers) it will also distort the projected image. This image is then cropped and scaled to the desired size. Cropping can be done at two different stages: in the negative carrier and at the paper level (using an easel)² The function $paperF$ will depend upon the type of developer and temperature used.

Let us assume that cropping is done at the carrier level (the image is first enlarged, then cropped), and that the projection of the negative does not affect the contrast function of the paper, nor vice-versa. Thus:

$$print \equiv proj(scale(paperF(crop(neg)), t)) \quad (1)$$

We are not concerned with recovering the time of exposure nor the paper’s contrast function, but rather, the effective contrast function given a particular time of exposure (i.e., we want the shifted curve as it was applied to the negative for this particular print). Then:

$$print \equiv proj(scale(contrast(crop(neg))))$$

²When the cropping is done in the negative carrier, the edges of the material used to crop the negative interact with the negative itself, resulting in a unique pattern along the edge of the printed image. We do not consider such effect in this paper.

We define the inverted print as the uncropped, unscaled, unprojected print:

$$invertedPrint \equiv scale^{-1}(proj^{-1}(print))$$

Thus:

$$contrast(crop(neg)) = invertedPrint$$

and $proj$, $crop$, and $scale$ (and their inverses) can be estimated using the registration and remapping algorithms typically used in panorama making [10]. The $invertedPrint$ will be an approximation to the negative (due to light diffusion of the enlarger, quantization errors in scaling, difficulties to perfectly match the distortion of the lens, etc.). As a consequence, it is very difficult to match the pixels on the negative to the $invertedPrint$ in order to estimate $contrast$. Instead, we take advantage of a property that the histogram of the $invertedPrint$ should be the same as the histogram of $contrast(crop(neg))$:

$$hist(contrast(crop(neg))) = hist(invertedPrint)$$

The $contrast$ function can be estimated using the histogram matching algorithm from [11]. We can then match the original print from the negative using:

$$matchPrint = contrast(neg) \quad (2)$$

If necessary, $matchPrint$ can be cropped to match the base print using the function $crop$: $crop(matchPrint)$. The $contrast$ function can also be estimated using a more advanced algorithm such as the two-scale tone matching in Bae et al. [9] (this method suffers from artifacts that histogram matching does not). As described above, the functions $crop$, $scale$, $proj$ and $contrast$ have been approximated. Fortunately their errors are relatively independent of each other: $scale$ and $crop$ operate on the geometrical boundaries of the projected image; $proj$ depends on the relative position of each pixel with respect to each other; and $contrast$ depends on the tonality of the original print with respect to the negative and can be approximated even if the pixels do not match perfectly.

2.1. Implementation Details

The Equation (2) is the basis for our print matching function, illustrated in Figure 1. We will assume that pixels values range within $[0, 1]$. In practice it is always desirable to operate on 16 bit images to avoid loss of tonal range.

We start with a negative image N , and a base print image B with a simple alpha channel A that covers the entire image (every pixels of B has an alpha channel value of 1). We find a set of pairs of control points between N and B ; each pair will correspond to the same feature in N and B and will assist in finding a function that maps B into an approximation of N . This set can be found automatically with the SIFT algorithm, which can find control points between two images even if they differ significantly in contrast and size [12], or it can be done manually by common pairs of points³. N , B and the set of control

³As few as a dozen control points might be necessary, but a larger number of points will be beneficial to detect erroneous pairs, and the effect of deformations present in the input images that are not considered by this model—such as the projection or scanning of a warped negative

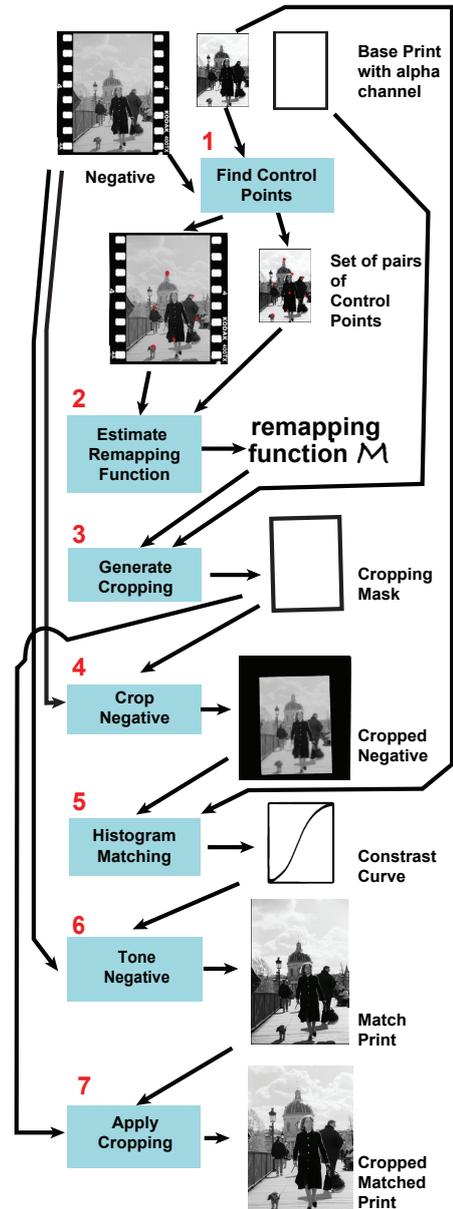


Figure 1: Diagram illustrating the major steps to compute the $matchPrint$ function. The numbers indicate the order of the operations.

points are used to perform feature matching and to compute \mathcal{M} , a 2-D perspective transform that takes into account the potential barrel or pin-cushion distortion of the lens (using a low order polynomial)—see [10] for details. We can now define the cropped negative N_c as:

$$N_c[x, y] = \begin{cases} N[x, y] & \text{if } \mathcal{M}(A)[x, y] = 1 \\ \text{undefined} & \text{otherwise} \end{cases}$$

Hence both N_c and $\mathcal{M}(B)$ will cover the same area of the photograph. We use the histogram matching algorithm from [11] to compute the function \mathcal{C} such that:

$$histogram(\mathcal{M}(B)) = histogram(\mathcal{C}(N_c))$$

The match print image `matchPrint` is equivalent to applying the contrast function \mathcal{C} to the negative image, i.e. for all $[x, y] \in N$:

$$\text{matchPrint}(N, B)[x, y] = \mathcal{C}(N)[x, y]$$

with an alpha channel $\mathcal{M}(A)$. Any point (x, y) of the match print with $\mathcal{M}(A)(x, y) = 1$ has a corresponding point in the base print B .

An example of the use of the `matchPrint` function is shown in Figure 2. The original negative looks relatively flat. The base print is scaled and remapped to match the negative, and the cropping of the negative is determined. The contrast function is calculated (shown in Figure 3) and applied to the negative, which is then cropped. The tonality of the new print closely matches the one of the base print. There is more detail in the clouds, and in the dark areas (such as her dress, under the coat), details that are lost in the scan of the base print. In this example the cropping is important: it has straightened the print, and cropped part of the sky.

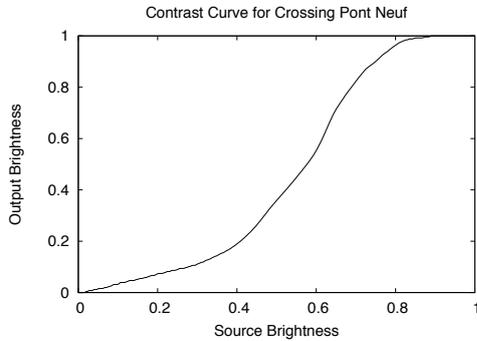


Figure 3: Computed contrast curve for *Crossing Pont Neuf* in Figure 2.

The `matchPrint` function is useful even when only smaller sections of the original image are used as a base print, as long as such region has a good tonality range (in this case no cropping is computed). Also, the function performs reasonably well when two prints are used as input (it is always preferable to use a negative scan, but that is not always possible). Figure 4 shows the use of `matchPrint` to a scan of a gelatin print by Lewis Hine and a base print that corresponds to a small section of a scan (from [13]).

2.2. Restoring Photographs by Combining Images

Sometimes it is necessary to restore a small section of a print, and there exist other versions of the same image that can be used to replace such damage. The `matchPrint` function can be used to match a section of one version to the section of the other, so they can be combined.

We assume two input images, the image to restore I_0 , and another image I_1 potentially of different size; I_0 has an alpha channel that describes usable areas of each image (a value of 1 in the alpha channel means the area is usable, and a value of 0 means the area should be replaced by the other image). This alpha channel is created by a human operator who determines the areas to be replaced.

Using a image blending algorithm (such as [14] that uses grassfire transforms to find seams for Laplacian pyramid blending, which we call *blend*) we combine I_0 and `matchPrint`(I_0, I_1) (respecting their corresponding alpha channels):

$$\text{restore}(I_0, I_1) = \text{blend}(I_0, \text{matchPrint}(I_0, I_1))$$

The quality of the restoration will primarily depend upon two factors: the first one is how well the tonality of the area to use matches the tonality of the area to be restored; the second is how well the detail in the area to use matches the detail in the area to be restored. They will be primarily affected by the relative difference in size between both images and the level of detail present in the desired region to be replaced.

Figure 5 shows an example of the use of `restore`. The Library of Congress scan of *Sadie Pfeifer...* (Figure 4) has what appears to be water damage in the left-bottom corner. The scan from [13] is used to restore the depicted area. The `restore` function performs well in this case because the area to restore is away from the center of attention of the original image, and it is not in sharp focus. Figures 6 and 7 show another example, where a damaged and badly scanned print, but of high resolution, is matched to a lower resolution scan from a book.

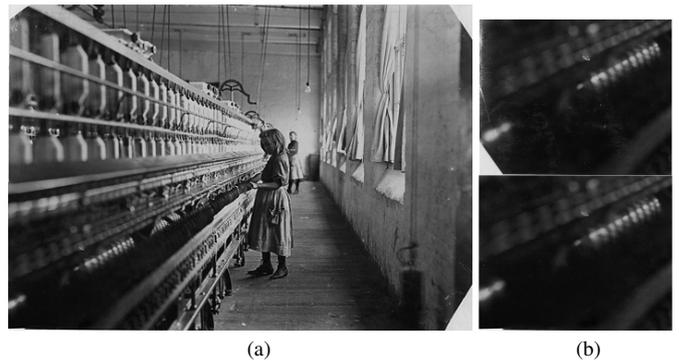


Figure 5: Restored *Sadie Pfeifer...* The scan of this photo from the Library of Congress (see Figure 4a) has water damage in the bottom left corner (shown in the top of (b)). Fortunately it is an area of low detail. A region from a scan from [13] (shown in Figure 4b) is used to restore it. The region on the top of (b) is the damaged area; the one on the bottom is the region used to replace it.

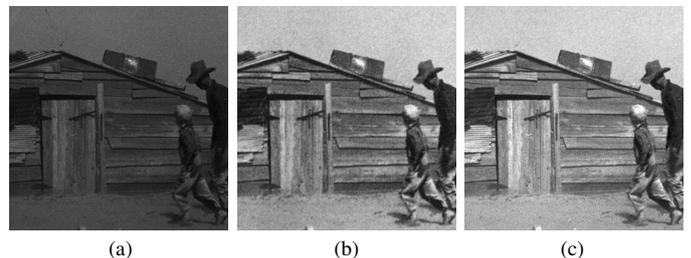


Figure 7: Detail of images in Figure 6. The level of detail of the original scan (a) is still present in the `matchPrint` (c), and the scratch in the top-left corner in (a) has been replaced in the `matchPrint` with the same area from the base print (b).

3. Dodging and Burning

Dodging and burning involves selectively projecting areas of the negative for a longer or shorter time, respectively. Let



Figure 2: *Crossing Pont Neuf* by D.M.German; used with permission. (a) Negative image looking flat. (b) Base print showing lack of detail in the shadows (e.g., dress under her coat) and highlights (e.g., areas of the sky). (c) The negative showing the area to be cropped. (d) matchPrint; we can note how the shadows and highlights have been recovered, while matching the overall contrast of the base print.

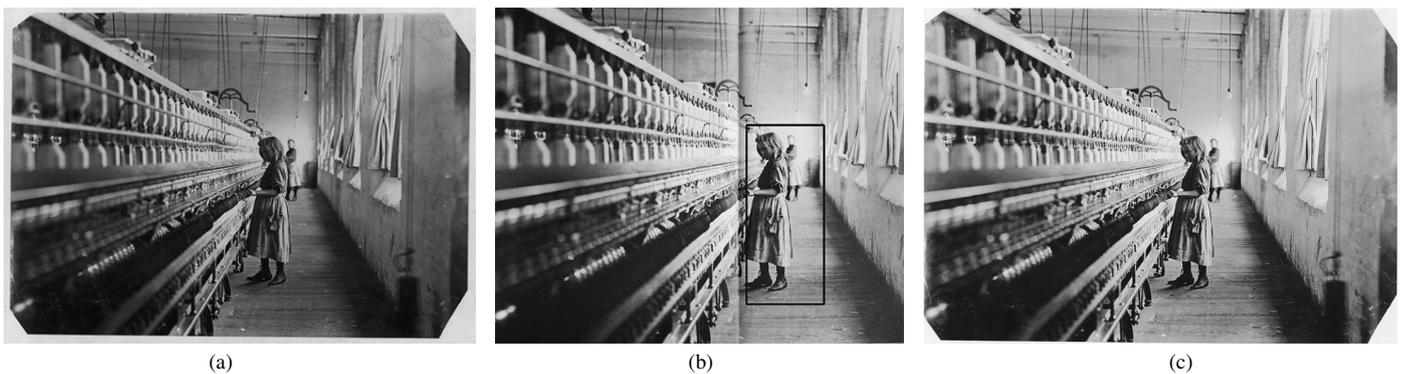


Figure 4: *Sadie Pfeifer, 48 inches high has worked half a year* (1908) by Lewis Hine; in the public domain. An example of a matchPrint created using only a section of the original image as the base print. (a) Scan of a print by the Library of Congress (nclc 01455). (b) Scan from a book [13] (a double page spread, hence the darker area in the middle; it has been slightly blurred to reduced the effect of halftoning printing). Only the boxed area is used as the base print. (c) The resulting matchPrint.

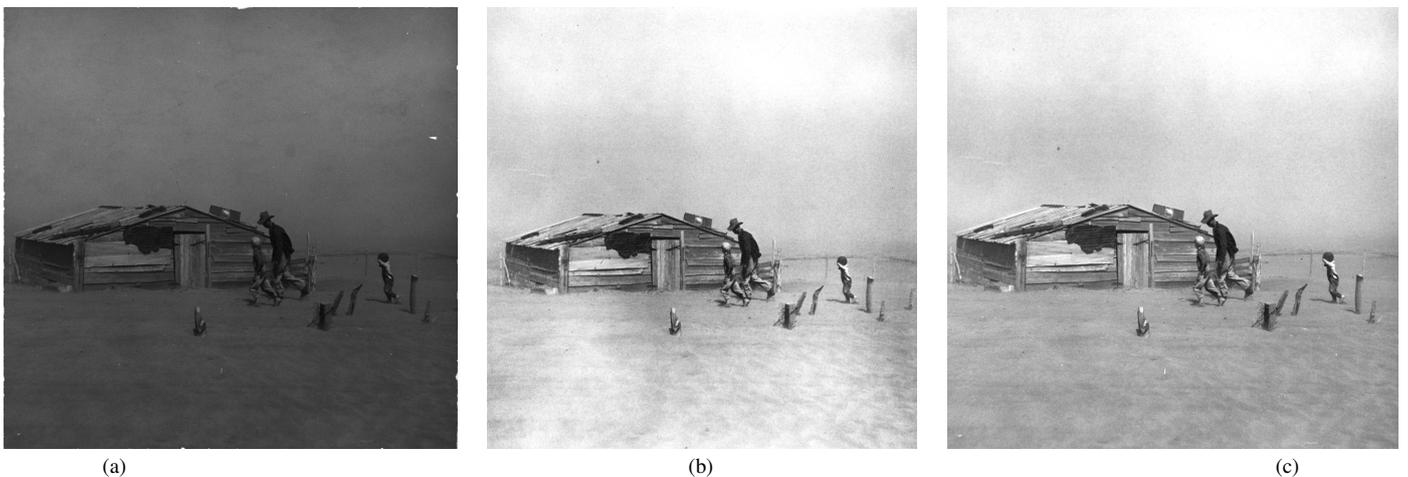


Figure 6: *Dust Storm, Cimarron Country* (1937) by Arthur Rothstein; in the public domain. (a) High resolution scan from dirty gelatin print, Library of Congress, USA (ppmsc.00241). (b) Scan from book [15] crediting a gelatin print also from the Library of Congress, USA. (c) Restored image where the sky of the matchPrint has been replaced with the sky from the base print and has been slightly cropped to remove the edge imperfections.

us assume we are interested in matching a negative (*neg*) to a dodged-and-burned base print (*dbBase*).

3.1. Model

Using Equation (1) we can model a dodged-and-burned print as

$$dbPrint \equiv proj(scale(paperF(crop(neg), T)))$$

where T is a matrix of the same size as neg indicating the projection time for each pixel. Without loss of generality, let us assume that there is no cropping or scaling:

$$dbPrint \equiv proj(paperF(neg), T)$$

As the $paperF$ function is increasingly monotonic, we can rewrite $dbPrint$ as the sum of a print at any reasonable time t , plus a dodge-and-burn contribution for each pixel (negative for dodging and positive for burning):

$$dbPrint \equiv proj(paperF(neg, t)) + dbContrib \quad (3)$$

The $matchPrint$ function can be used as an estimation of the left addend in equation (3) using the neg and $dbBase$ as inputs:

$$print_m = matchPrint(neg, dbBase) = proj(paperF(neg, t))$$

$$dbPrint \equiv print_m + dbContrib \quad (4)$$

Thus:

$$dbContrib \equiv dbPrint - print_m \quad (5)$$

We can replace $dbPrint$ with $dbBase$ (the print we are trying to match). Unfortunately $dbBase$ and neg are unlikely to be perfectly aligned due to quantization errors, optical projection, and scanning distortions (if they were perfectly aligned, $dbPrint$ would be identical to $dbBase$). We address this problem by assuming that, when blurred (e.g. using a Gaussian blur), both sides of Equation (5) should be the same:

$$dbMask \equiv blur(dbContrib) = blur(dbBase) - blur(print_m) \quad (6)$$

The size of the blurring kernel depends on the resolution of the image and its level of detail. The goal of this step is to diffuse any harsh edges present in the difference of the two images to make it possible to compare the luminance of each region.

The $blur(dbContrib)$ is similar to the dodge-and-burn mask used by print makers (and we will refer to it as $dbMask$). The $dbMask$ indicates, for every area of the image, the level of over and under exposure required to create the final print. Values below zero indicate dodging, and values over zero burning. We now need to apply this mask to $print_m$. To do so we multiply each point in $dbMask$ (equation 6) by its corresponding point in the negative (we use $*$ to represent this operation) and match the histogram of the resulting image to the expected contribution of the mask by computing the function f_c such as:

$$hist(dbContrib) = hist(f_c(dbMask * neg))$$

This step helps to compensate the loss of contribution that results from multiplying $dbMask$ by negative. Unfortunately we

do not know at this point $dbContrib$; we approximate f_c by assuming a) that the histogram of both sides remains the same after the images have been blurred; and b) that $f_c \circ blur = blur \circ f_c$. We therefore calculate f_c such that:

$$hist(blur(dbContrib)) = hist(f_c(blur(dbMask * neg)))$$

resulting in:

$$dbContrib = f_c(dbMask * neg) \quad (7)$$

which we can combine with Equations (4,6,7) and results in

$$dbPrint \equiv print_m + f_c((blur(dbBase) - blur(print_m)) * neg) \quad (8)$$

In summary, We can approximate the dodged-and-burned base print by computing a $matchPrint$ of the negative and the base print, and by estimating the function f_c .

3.2. Implementation Details

The Equation (8) becomes the foundation for our method to reproduce a dodged-and-burned print, illustrated in figure 8.

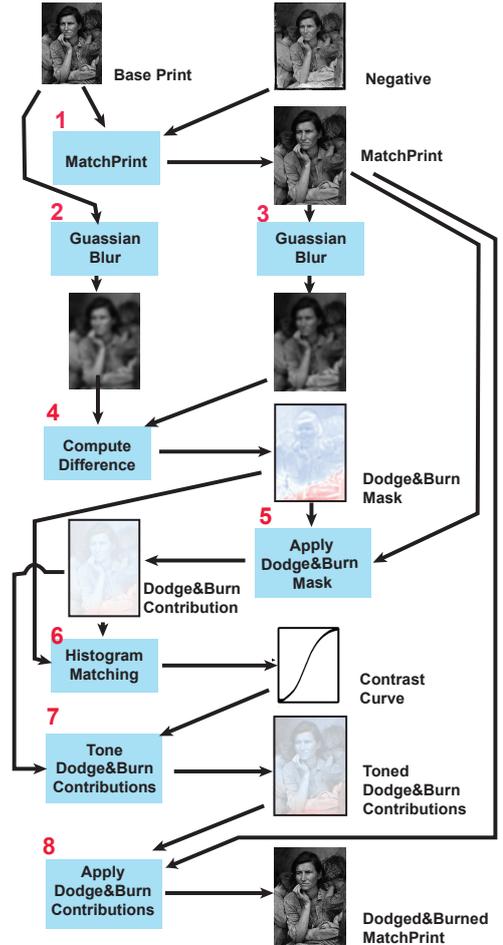


Figure 8: Diagram illustrating the major steps to compute the dodged-and-burned matchPrint function. The numbers indicate the order of the operations. For illustration purposes the mask and the contributions are shown in red for burning—positive values—and blue for dodging—negative values—and are presented darker than they are.

We start with two images, the base print B and the negative scan N . We first compute the $M = \text{matchPrint}(N, B)$, and in the process preserve the remapped base print $B_m = \mathcal{M}(B)$ where \mathcal{M} is the 2-D perspective transform used by `matchPrint`.

We now compute, using the histogram matching algorithm in [11] (in a manner similar to the way it was computed for `matchPrint`), the function \mathcal{C}_m such that, for all $[x, y] \in M$:

$$\begin{aligned} \text{histogram}(G(B_m[x, y]) - G(M[x, y])) = \\ \text{histogram}(\mathcal{C}_m(N[x, y] * (G(B_m[x, y]) - M[x, y]))) \end{aligned}$$

As described in section 3.1, G is a Gaussian blur where the size of the kernel depends on the size of the image. In our experiments a kernel size of 0.3% of the largest dimension of the image has provided good results. Finally, the desired dodged-and-burned image DB is computed, for all $[x, y] \in M$:

$$\begin{aligned} DB[x, y] = M[x, y] \\ + \mathcal{C}_m(N[x, y] * (G(B_m[x, y]) - G(M[x, y]))) \end{aligned}$$

We exemplify our method using the photograph *Roadside Stand Near Birmingham* by Walker Evans (1936). The scanned negative originates from the Library of Congress (LC-USF342-T01-008253, 5145×4173 pixels, including margin) and the base print is a scan of a print also from the Library of Congress (LC-DIG-ppmsc-00239, 2275×2136 pixels). Figures 9 and 10 show the negative scan, the base print, the `matchPrint`, and the burned-and-dodged `matchPrint`. This print has significant dodging and burning, and it is depicted in Figure 11. For this photograph the areas outside the store has been burned and the faces of the man and the girl inside the store are still visible.

Figure 12 shows the application of the `MatchDodgedAndBurnedPrint` function to *Migrant Mother* by Dorothea Lange. Notice how the base print shows burning in the lower area (below the baby's body). In this case there is no significant dodging in the image.

4. Comparing Different Versions of the Same Image Using `matchPrints`

The `matchPrint` function can also be used to compare two different prints of the same negative. Comparing two versions of an image can assist historians trying to identify the differences between them due to alterations in the negative or the print (either due to damage such as scratches or manipulation of the image)

For example, an early print of Dorothea Lange's *Migrant Mother* is presented in Figure 13. This print has important historical value, because it was made before Lange altered the negative (as described in the note of the bibliographic entry of this print; see also [17] as cited by [18]). The difference, although significant, is not obvious. But when `matchPrint` is applied to this image (using the negative as the base print), and we subtract one from the other, the difference becomes obvious: Lange was so bothered by the thumb of the mother that she decided

to remove it from the negative (these types of alterations were performed by lightly drawing on the emulsion side of the negative using a soft graphite pencil). The difference of the images shows also scratches that are in one, but not both two images. In this case we can observe a large scratch along the face in the older print.

Similarly, it is possible to compare different croppings of the same image. In Figure 14, the edge of four different versions of *Migrant Mother* (found in different photography books and at the Library of Congress) are shown.

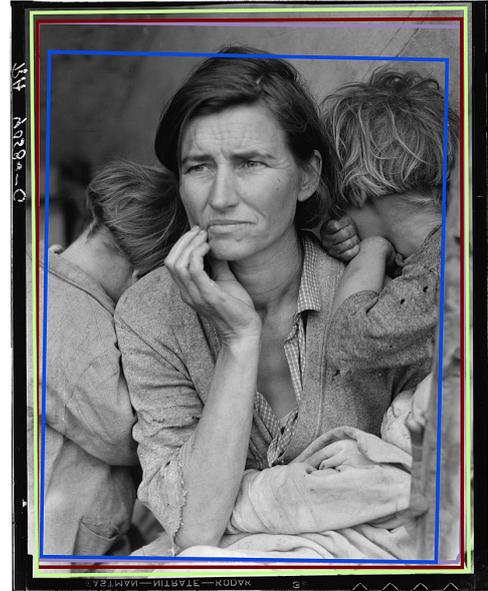


Figure 14: The colored rectangles around the edge of the image show each of four different croppings of *Migrant Mother* from various sources.

5. Quantitative Evaluation

We present in this section a set of metrics that attempt to quantify the effects of these algorithms. Their purpose is to evaluate the quality of the output images with respect to the input ones. In particular we are interested to measure how much closer to the base print, the `matchPrints`, and the dodged-and-burned `matchPrints` are than the negative, and how much information has been lost or gained in these processes.

Some of these metrics operate on the intensity values of the image, and others are based in information theory. Where applicable these metrics have been normalized so that they are invariant to the number of pixels in the domain.

Histogram metrics: $Mean, \sigma$ One simple method to compare two images is to compare their histograms. If two images have a similar mean, they have a similar brightness, and if they have a similar σ (its standard deviation), then they have a similar contrast.

Complexity of image: R_α The complexity of an image can be defined with respect to the number of elements (regions) R needed to extract a given ratio α of information in an



Figure 9: *Roadside Stand Near Birmingham* by Walker Evans (1936). (a) Negative from the Library of Congress; in the public domain. (b) Base print, a scan from a print, also from the Library of Congress; its cropping attempts to center and balance the print. (c) matchPrint, its foreground is still very bright. (d) Burned-and-dodged matchPrint, its foreground and the inside of the store are darker, but the faces of the man and girl remain still visible (they are black in the base print).

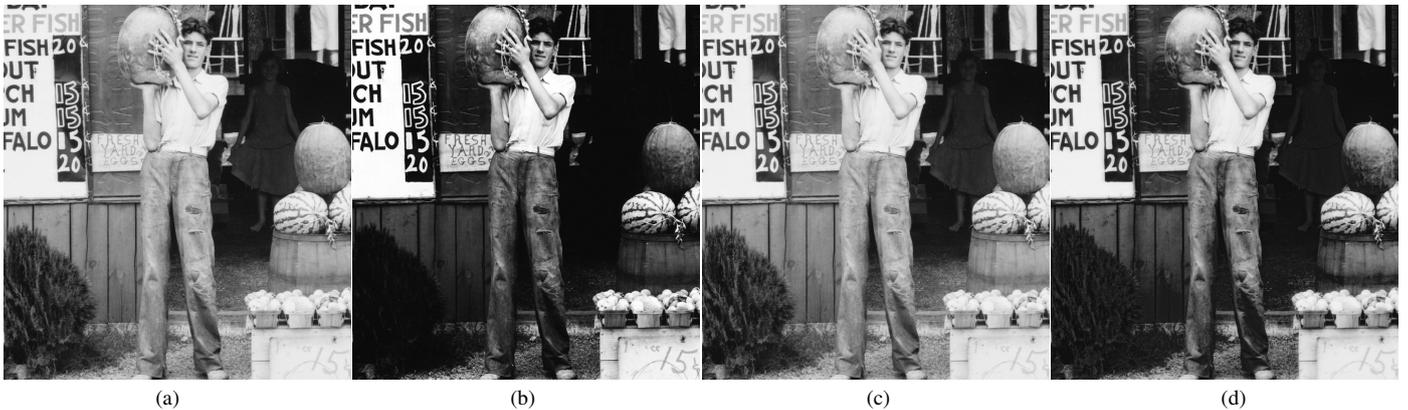


Figure 10: Close up of the images in Figure 9. (a) Negative, (b) base print, (c) matchPrint, and (d) dodged-and-burned matchPrint images. The girl is still visible in the dodged-and-burned image (d), but not in the base print (b).



Figure 11: (a) Dodging mask, (b) dodging contribution, (c) burning mask, and (d) burning contribution for *Roadside Stand Near Birmingham*. The dodging version is presented as the absolute value of pixels below zero, and the burning as the positive values only. Notice in particular how the burning applies mainly to the foreground and the two kids, while the dodging tries to brighten slightly the areas in the windows of the store and the top half of the print. For presentation purposes the pixels in these masks have been inverted: white means no contribution, and black, full contribution.

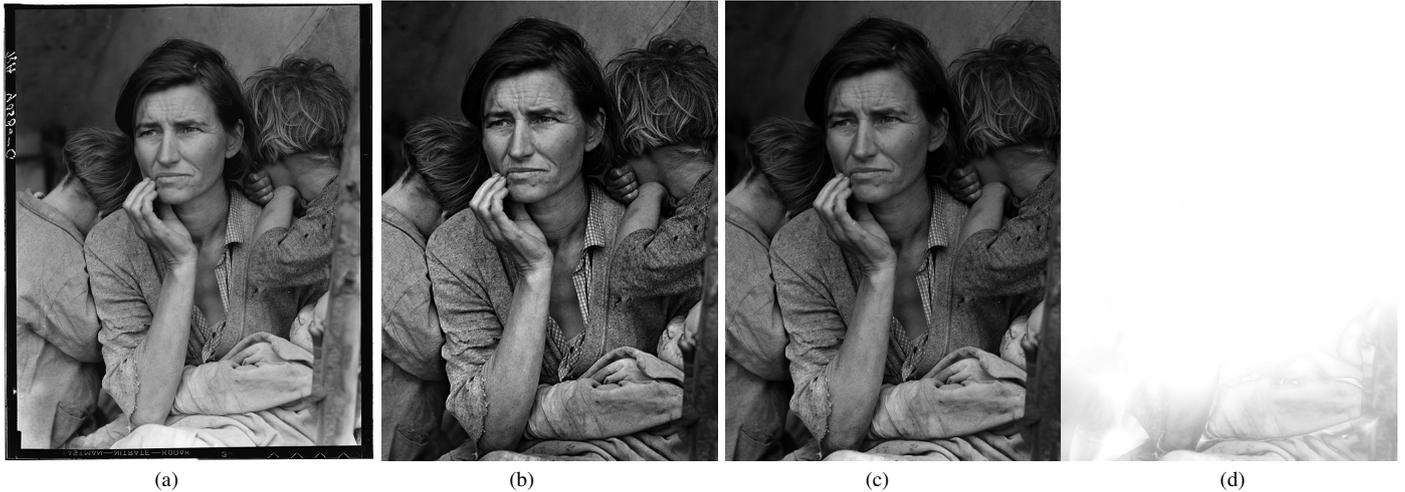


Figure 12: *Migrant Mother* (1936) by Dorothea Lange; in the public domain. (a) Scan of the negative done by the Library of Congress of the United States (LC-DIG-fsa-8b29516); it measures 6849×8539 pixels including margin. (b) Base print, created by Mike Johnston [16]; it measures only 364×504 pixels. (c) Dodged-and-burned matchPrint. (d) Burning contribution. The bottom of the image has been darkened. This image required no significant dodging.

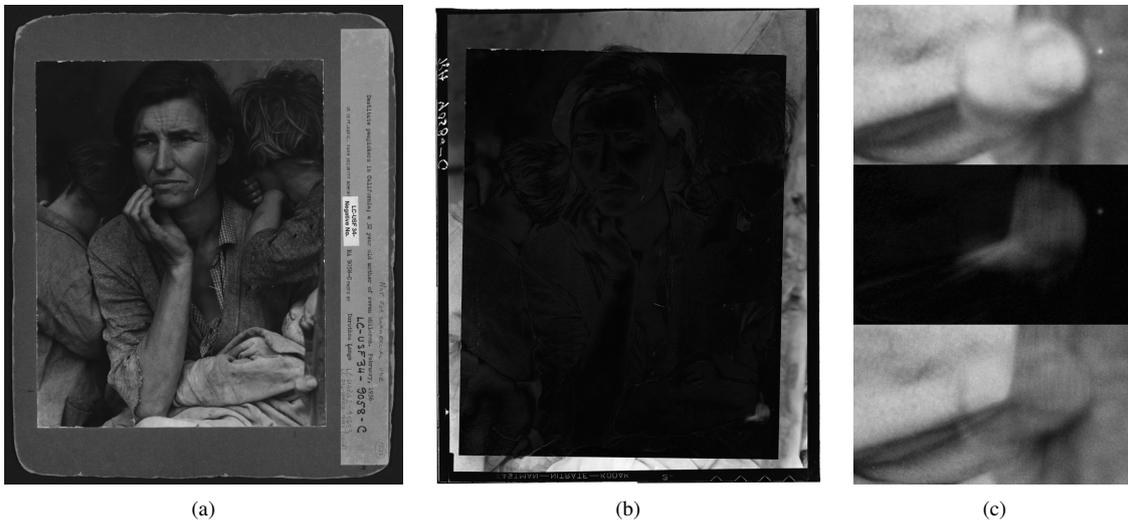


Figure 13: Comparing different versions of the same image. (a) Early version of *Migrant Mother*; this print was made before its negative was altered; the difference is not obvious when visually comparing this version with those in Figure 12. (b) Difference of the altered negative (from Figure 12) with the matchPrint of such negative and this print. (c) Close up of the bottom right of the image: at the top the version from (a) where the thumb is clearly visible, (b) shows the difference, and (c) shows the altered negative without the thumb.

image [19]. It is related to the complexity in describing the image and its structure from an informational point of view. The computation of R is guided by the maximization of information building regions with maximum intensity homogeneity. In this context, the gradient of a region is implicitly captured by the divergence between the histograms of its subregions since the algorithm chooses the spatial division that maximizes the Jensen-Shannon divergence [20] (i.e., information gain). Thus, higher values presuppose that the image has a lot of detail and a wide range of intensities, while lower values correspond to an image with few details and a short dynamic range. We use it to compare the loss or gain in detail between the images.

Table 1 shows these metrics when computed in the images *Crossing Pont Neuf* (CP), *Roadside Stand Near Birmingham*

(RS), and *Migrant Mother* (MM). The relative difference between the negative and the base print represents a measure of the amount of correction needed on the negative. As expected, for the three images, the mean of the dodged-and-burned matchPrint is closer to the base print than the negative. For CP and MM the standard deviation is preserved; but for RS it is reduced more, which is expected: the original negative has a vast dynamic range and the job of the print maker was to reduce some of the brightness in the foreground. With respect to the R_α , in all cases the dodged-and-burned image has a larger value; we can interpret this as a measure of quality: the dodged-and-burned images have more informative regions than either the base print or the negative. Notice how the dodged-and-burned versions of RS and MM have significantly improved the matchPrint, while the opposite is not true for CP (i.e., for this image dodge-and-burn does not appear to be necessary).

<i>Image</i>	<i>Mean</i>	<i>Mean_n</i>	σ	$R_{0.25}$
CP_n	34011	1.000	196	62
CP_b	38282	1.126	193	67
CP_m	36427	1.071	193	90
CP_{db}	36443	1.071	194	82
RS_n	32487	1.000	181	454
RS_b	24070	0.741	156	373
RS_m	35928	1.106	186	481
RS_{db}	26620	0.819	164	616
MM_n	24626	1.000	179	919
MM_b	19078	0.775	174	963
MM_m	19374	0.787	173	710
MM_{db}	17090	0.694	171	983

Table 1: Values of mean, mean_n (relative mean with respect to the negative), standard deviation σ , and complexity $R_{0.25}$ for the negative, base print, matchPrint, and dodged-and-burned matchPrint images (*n*, *b*, *m*, and *db* indexes, respectively) of *Crossing Pont Neuf* (CP), *Roadside Stand Near Birmingham* (RS), and *Migrant Mother* (MM).

Objectively measuring the improvement gained by dodge-and-burning an image is difficult. On one hand, dodge-and-burn tries to remove some distracting information from some regions of the negative; on the other hand, it tries to amplify the information in others. The following metrics are used to compare the effect of the improvement of this step. Each of the metrics below compares specific properties of the images.

CMI_α The Color Mutual Information gives us the degree of dependence between the color values (intensity domain) and the regions of the image R_α (spatial domain), and is interpreted as a measure of the information or saliency associated with an specific intensity value. It has been defined in [21] as a new information theoretic measure evaluating style and aesthetics aspects of artistic works. High values of CMI_α express a high dependence or correlation between an intensity value and a given region, and identify the most relevant intensities, that is, the intensity values conveying the most information. On the other hand, the lower values correspond to the intensities distributed uniformly in the image. CMI_α permits us to analyze the changes in the importance of each intensity value showing the images with a normalized thermal scale representation (from low—cold—to high—warm—values). By inspecting the CMI_α images we can visually identify the areas that are most likely to attract the attention of the viewer.

RMSE The Root Mean Square Error measures the average of the error understood as the difference between an estimator and the true values. In our case, it represents the displacement between two images with respect their intensity values and we would also expect that RMSE, with respect to the negative, is lower in the output images than in the base print image.

PSNR The Peak Signal Noise Ratio is a ratio to quantify the noise that can appear in a new representation of the data.

It is based on the RMSE and expressed in a logarithmic decibel scale where higher PSNR correspond to higher reproduction quality. This will measure a “fidelity” distance between our images.

MI The Mutual Information expresses the shared information or the degree of correlation between two random variables and is maximum when the corresponding structures overlap (measured in bits) [20]. Thus, MI can be seen as a measure of the structural or compositional similarity and can qualitatively be thought of as a measure of how well one image explains the other. We compute it from the image histograms. Higher values imply higher correlation.

Figures 15, 16 and 17 show the Color Mutual Information $CMI_{0.25}$ of *Crossing Pont Neuf*, *Roadside Stand Near Birmingham* and *Migrant Mother* using a thermal scale. Compare these images to those in Figures 2, 9 and 11. In *Crossing Pont Neuf* the matchPrint and the dodge-and-burned matchPrint enhance information in the areas that are more likely to attract the attention of the viewer: the dog, the sky and the woman—the effect is particularly prominent in her dress. In *Roadside Stand Near Birmingham* the amount of information in the matchPrint and the dodged-and-burned matchPrints is very similar to the negative one. We interpret this as a sign of quality: even though the images tone has significantly changed, the information in the print remains relatively constant. However, the base print shows more information in the foreground boxes than any of the two generated images are able to convey (these regions required more selective burning in this area than what our algorithm was able to achieve).

Migrant Mother shows a different story. The dodged-and-burned image shows lose of information in the area of the body of the child that corresponds to the burning mask (see burning mask in Figure 12). In other words, the print maker has removed information from this area, making it less relevant (and less distracting) to the eyes of the viewer. Instead, her face (delineated by her hair) is the most relevant part of the image.

Observing the effect of the matchPrint algorithms in these images, we can see the different changes that the print maker has achieved. In *Crossing Pont Neuf* the goal is to increase the attention to the main regions of the image; in the second it is to maintain as much detail as possible from the original negative while dramatically changing the luminosity of some of its areas, and in *Migrant Mother* the goal is to remove attention from an area of the image.

Table 2 shows RMSE, PSNR, and MI for each of the three sets of photographs. We show the comparison of the negative against each of the other 3 versions of the images: base print, matchPrint, and dodged-and-burned matchPrint. With respect to RMSE, we would expect the base print to have a higher value than any of the other two matchPrints. For all three images, the RMSE drops from the matchPrint to its dodged-and-burned version, as expected too (this step adds noise). Something similar happens with PSNR and MI: their values of the matchPrint are higher than with respect to the base print,

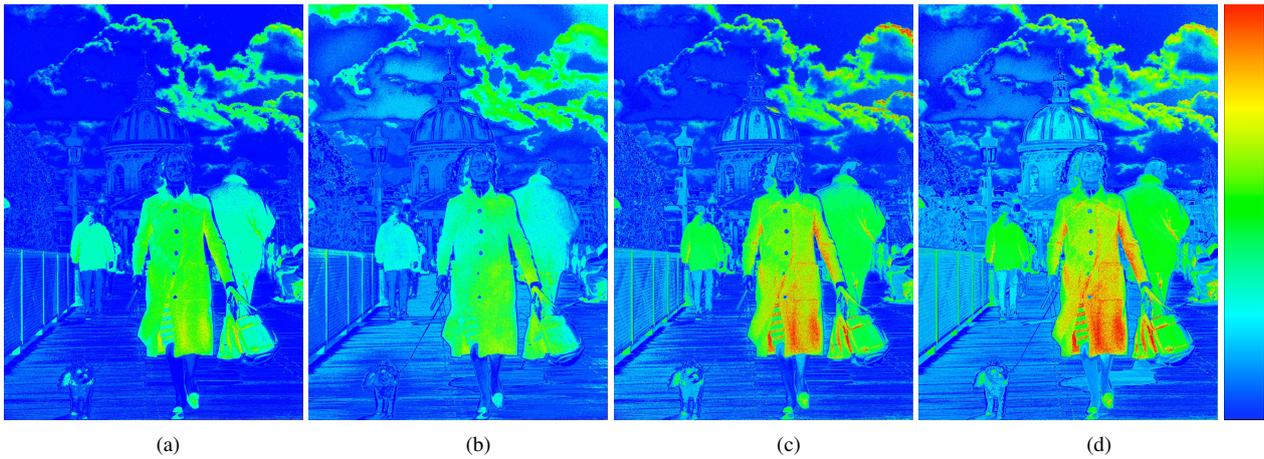


Figure 15: Color Mutual Information $CMI_{0.25}$ for *Crossing Pont Neuf* using a normalized thermal scale—the lowest value (0) corresponds to the blue and the highest (1) to the red. (a) Negative (b) Base print. (c) `matchPrint`, and (d) `dodged-and-burned matchPrint`. Both the `matchPrint` and the `dodged-and-burned matchPrint` are an improvement over the negative, particularly in the areas of the dog, the dress, and the sky—which are the most prominent regions in the photograph.



Figure 16: Color Mutual Information $CMI_{0.25}$ for *Roadside Stand Near Birmingham* using a normalized thermal scale—the lowest value (0) corresponds to the blue and the highest (1) to the red. (a) Negative showing a very good distribution of information. (b) Base print has no visible detail in the inside of the store where the girl is not recognizable. (c) `matchPrint` and (d) `dodged-and-burned matchPrint` have removed information from the shadows, but they are still well defined.

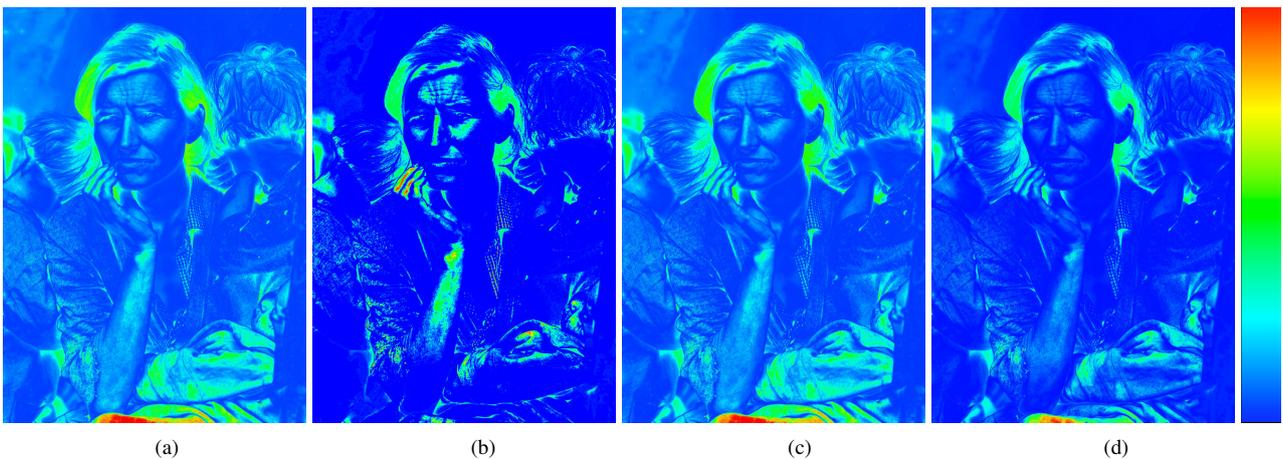


Figure 17: Color Mutual Information $CMI_{0.25}$ for *Migrant Mother* using a normalized thermal scale—the lowest value (0) corresponds to the blue and the highest (1) to the red: (a) negative, (b) base print, (c) `matchPrint`, and (d) `dodged-and-burned matchPrint`. Both `matchPrint`s show a more detailed information distribution than the base print (which is significantly smaller in size). It is also interesting to see that there is less information in the clothes of the child in her arms; this can be interpreted as an attempt of the print maker to remove the relevance of this area in the `dodged-and-burned print`.

<i>Images</i>	<i>RMSE</i>	<i>PSNR</i>	<i>MI</i>
(CP_n, CP_b)	10942	31.094	2.397
(CP_n, CP_m)	6161	41.072	6.474
(CP_n, CP_{db})	7911	36.730	4.289
(RS_n, RS_b)	13788	27.078	1.864
(RS_n, RS_m)	3871	49.142	6.638
(RS_n, RS_{db})	8760	34.957	3.243
(MM_n, MM_b)	8323	35.846	1.597
(MM_n, MM_m)	5394	43.381	6.638
(MM_n, MM_{db})	8304	35.888	4.128

Table 2: Measures of RMSE, PSNR, and MI for the pairs of images (negative against base print, `matchPrint` and dodged-and-burned `matchPrint` respectively) for *Crossing Pont Neuf* (*CP*), *Roadside Stand Near Birmingham* (*RS*) and *Migrant Mother* (*MM*).

but they drop after the images are dodged-and-burned. In *CP* and *RS*, RMSE and PSNR show less noise in the burned-and-dodged version than in the base print; we interpret this result as a sign of confidence that the data in the dodge-and-burned print is of higher quality (less noise) than the one in the base print. For *MM* we were surprised that the RMSE and PSNR values are so close to each other for the base print and the dodge-and-burned `matchPrint`. This could be the result of using a very low resolution image as the base print (18 times smaller in each dimension).

As expected, the MI column shows, for the three images, that the `matchPrint` has high level of correlation, and that this correlation drops in the dodged-and-burned version but it is still higher than the base print.

6. Discussion and Future Work

As described in the previous section, the methods presented in this paper show a significant improvement in the tonality of the high-resolution scan such that it approximates the tonality of the scanned print used as a reference. The main restriction of these methods is that they require at least two versions of the same image (one with higher detail—usually a negative—and the other with better tonality—usually a scan of a print). This method will not be useful when only one scan is available.

We believe that a skilled operator of image editing software (such as Adobe Photoshop[®]) will be capable of performing the same functionality than the methods presented in this paper. We believe that, as potential future work, it would be valuable to perform an experiment in which match prints are created by the methods herein, and by a human operator, and then evaluated by individuals to determine if one version is better than the other. These comparisons could also be done in terms of the time and skill required by the operator to perform the task.

It is possible that, given enough time, the skilled operator will prepare match prints that are superior. Unfortunately such experts are expensive, and some collections are composed of a large number of images that can be scanned in bulk. These methods can be used as a first-approximation for the majority of the collection. Afterwards, a hybrid method can be used for

the most important images of the collection: first compute a match print from the scans, and second, the human operator uses this match print as a starting point and fine-tune it as necessary, saving time.

It is also important to mention that some prints are intended to be “softer” (less sharp) than its negative. This is done for several purposes, such as an artistic decision (e.g. giving the impression of softer skin of the subject or to remove the effect of grain in the negative) or to compensate for some defects of the negative (such as the presence of dust and scratches). A match print computed from such print will not preserve such qualities. In this case the problem is the reverse, how to improve the tonal range of the print using the information in the negative, without affecting the level of detail of the print (scans of prints commonly show loss of tonal range in the shadows and highlights). This is another area that can be further explored in the future.

With respect to other future work, we believe it is worth exploring if the tonality of the `matchPrints` can be improved by partitioning the image into regions (instead of computing one contrast curve for the entire image, the image is divided into smaller regions and each has its own contrast curve). Another approach might be to compute the tonality of a pixel using only pixels surrounding it.

These methods can also be extended to be applicable to color images, both toned black and white (such as selenium and sepia), and color. The problems are similar: the scan and the print “look” different, but one wants to recover the printer’s intent, to be able to match the colour negative to the colour print.

It is necessary also to evaluate the results of the dodge-and-burned `matchPrint` with respect to other automatic toning methods (such as a direct application of the two-scale toning method [9]).

7. Conclusions

We have described a method to extract from the scan of a print (the base print) the decisions made by its print maker; namely the cropping of the negative, the selection of the brightness and contrast of the print, and any dodge-and-burning applied to it. We use this information to improve a scanned negative of the same image to create a version (which we call dodged-and-burned `matchPrint`) that is visually similar than the print, but with higher level of detail and tonality than the scan of the print.

We have also shown the usefulness of `matchPrint` in the comparison of multiple variants of the same image.

To evaluate the effectiveness of the methods presented, we have proposed various metrics that compare the tonality of the `matchPrints` and the original scans, and the distribution of information of each image. They show that the `matchPrints` have a tonality that is close to the base prints, while they retain the richness of information present in the negative. The Color Mutual Information metric is used to visually compare the effects of dodging-and-burning in the different versions of

the image, showing that in some cases dodging-and-burning is used highlight some areas and to reduce the relevance of others.

Acknowledgments

This research has been funded by the Natural Sciences and Engineering Research Council of Canada and by grant number TIN2007-68066-C04-01 of the Ministry of Education and Science of Spain. We thank the anonymous reviewers for their constructive comments and suggestions.

References

- [1] Double Take: A Comparative Look at Photographs, C.N. Potter, 1981.
- [2] T. Rudman, The Photographer's Master Printing Course, 2nd Edition, Reed International Books LTD, 1996.
- [3] A. Adams, The Print, Bulfinch, 1995.
- [4] R. E. Sobol, System and method for automatically cropping graphical images, United States Patent 7034848, assignee Xerox Corp (April 2006).
- [5] B. D. Kimbell, D. L. Dalton, M. L. Rudd, System and method for automatically cropping graphical images, United States Patent 6456732, assignee Hewlett-Packard Company (Sept 2002).
- [6] A. Santella, M. Agrawala, D. DeCarlo, D. Salesin, M. Cohen, Gaze-based interaction for semi-automatic photo cropping, in: CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems, ACM, New York, NY, USA, 2006, pp. 771–780.
- [7] J. Geigel, F. K. Musgrave, A model for simulating the photographic development process on digital images, in: SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques, 1997, pp. 135–142.
- [8] E. Reinhard, M. Stark, P. Shirley, J. Ferwerda, Photographic tone reproduction for digital images, in: SIGGRAPH '02: Proc. of the 29th annual conference on Computer graphics and interactive techniques, 2002, pp. 267–276.
- [9] S. Bae, S. Paris, F. Durand, Two-scale tone management for photographic look, in: SIGGRAPH '06: ACM SIGGRAPH 2006 Papers, 2006, pp. 637–645.
- [10] R. Szeliski, Image alignment and stitching: a tutorial, Found. Trends. Comput. Graph. Vis. 2 (1) (2006) 1–104.
- [11] R. C. Gonzalez, R. E. Woods, Digital Image Processing, 2nd Edition, Prentice Hall, 2002.
- [12] D. G. Lowe, Distinctive image features from scale-invariant keypoints, Int. J. Comput. Vision 60 (2) (2004) 91–110.
- [13] H.-M. Kotzle (Ed.), Photo Icons. The story behind the pictures, Taschen, 2005.
- [14] Y. Xiong, K. Turkowski, Registration calibration and blending in creating high quality panoramas, in: WACV '98: Proceedings of the 4th IEEE Workshop on Applications of Computer Vision (WACV'98), IEEE Computer Society, Washington, DC, USA, 1998, p. 69.
- [15] N. Rosenblum, A World History of Photography, Abbeville Press, New York, USA, 1984.
- [16] M. Johnston, Print Offer: Dorothea Lange's Migrant Mother, The Online Photographer. http://theonlinephotographer.blogspot.com/2006/12/print-offer-dorothea-langes-migrant_15.html, accessed March 2008 (Dec 2006).
- [17] F. J. Hurley, Portrait of a Decade, Baton Rouge: Louisiana State University, 1972.
- [18] L. Wells, Photography, a Critical Introduction, Routledge, 1972.
- [19] J. Rigau, M. Feixas, M. Sbert, An information theoretic framework for image segmentation, in: IEEE International Conference on Image Processing (ICIP '04), Vol. 2, IEEE Press, 2004, pp. 1193–1196.
- [20] T. M. Cover, J. A. Thomas, Elements of Information Theory, Wiley Series in Telecommunications, 1991.
- [21] J. Rigau, M. Feixas, M. Sbert, Informational dialogue with Van Gogh's paintings, in: P. Brown, D. W. Cunningham, V. Interrante, J. McCormack (Eds.), Computational Aesthetics 2008. Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging, Eurographics Association, 2008, pp. 115–122.