

Digital Image Processing Laboratory: Pointwise Operations and Gamma

February 7, 2021

Introduction

This laboratory will discuss various types of point operations on images. This generally refers to a remapping of pixels at one intensity to a different intensity. Point operations make up an important class of techniques for image enhancement.

In Section 1, we will introduce a graphical tool called a *histogram*, which shows the distribution of pixel intensities in an image. An enhancement technique called *histogram equalization* will be introduced in Section 2.

In Section 4, we will discuss an important parameter called *gamma* which describes the nonlinear behavior of certain display devices. You will determine the gamma of your computer monitor, and then apply a point operation known as *gamma correction* to compensate for this nonlinearity.

In some exercises you will be asked to display gray scale images in Python. If X is an image that takes on the values $\{0, 1, \dots, 255\}$, then it may be displayed by using the following commands.

```
import numpy as np
import matplotlib.pyplot as plt
from matplotlib import cm
from PIL import Image

gray = cm.get_cmap('gray', 256)
im = Image.open('race.tif')
x = np.array(im)
plt.imshow(x, cmap=gray)
```

1 Histogram of an Image

The *histogram* of a digital image shows the distribution of its pixel intensities. Let X be a gray scale (achromatic) image. A single pixel is denoted as X_s where $s \in S$ is a position in the lattice S . Generally, we will assume that X_s takes on the discrete values $0, \dots, L - 1$

Questions or comments concerning this laboratory should be directed to Prof. Charles A. Bouman, School of Electrical and Computer Engineering, Purdue University, West Lafayette IN 47907; (765) 494-0340; bouman@ecn.purdue.edu

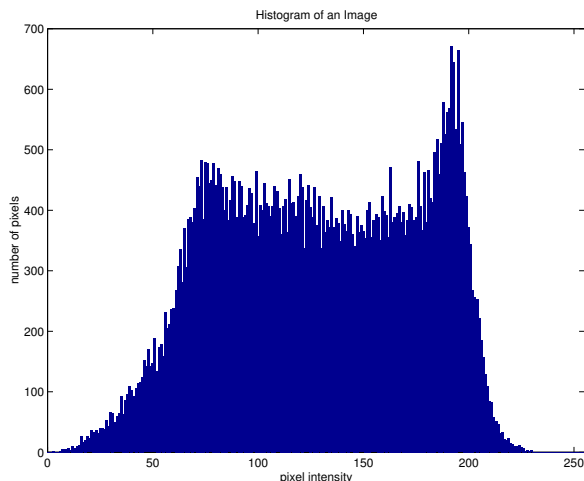


Figure 1: Histogram of an 8 bit image

where typically $L = 256$. The histogram of the image X is then given by

$$h(i) = \sum_{s \in S} \delta(X_s - i) .$$

So $h(i)$ computes the number of pixels that take on the value i . A typical histogram of an 8 bit image is shown in Figure 1.

Now we will display the histogram of two images.

1. Download the two gray scale images *race.tif* and *kids.tif* from the lab home page.
2. Use the Python function [\[matplotlib.pyplot.hist\]](#) to display the histograms of each image. **Note:** The *hist* command requires the image pixel values to be oriented in a single column. Using the syntax


```
plt.hist(x.flatten(), bins=np.linspace(0, 255, 256))
```

 the first argument orients the image x into a column vector, and the second argument specifies the bin centers.
3. Print or export each image and its labeled histogram.

Section 1 Report:

Hand in the two images and their labeled histograms.

2 Histogram Equalization

Histogram equalization is a common image enhancement technique. The objective is to transform the image so that the output histogram is uniform over the full range of gray values.

To approach this problem, think of an image as a collection of i.i.d. random variables, each with cumulative distribution and density of

$$\begin{aligned} F_x(x) &= \text{Prob}\{X \leq x\} \\ p_x(x) &= \frac{d}{dx} F_x(x) . \end{aligned}$$

Now consider forming the following random variable.

$$Y = F_x(X) \tag{1}$$

Here Y is the result of mapping the random variable X through its own cumulative distribution function. The cumulative distribution of Y can be easily computed.

$$\begin{aligned} F_y(y) &= \text{Prob}\{F_x(X) \leq y\} \\ &= \text{Prob}\{X \leq F_x^{-1}(y)\} \\ &= F_x(F_x^{-1}(y)) \\ &= 0 \quad \text{for } y < 0 \\ &= y \quad \text{for } 0 \leq y \leq 1 \\ &= 1 \quad \text{for } y > 1 \end{aligned}$$

This shows that Y has a uniform distribution on the interval $[0, 1]$. Therefore, the histogram of an image can be equalized by mapping the pixels through their cumulative distribution function $F_x(x)$.

In order to implement the histogram equalization method, we must estimate the cumulative distribution function for the image. We can do this using the image's histogram. Let $h(i)$ be the image's histogram formed by computing the number of pixels at gray level i . Typically, the pixels take on the values $i = 0, \dots, L - 1$ where $L = 256$ is the number of discrete levels that a pixel can take on. The cumulative distribution function can then be approximated by

$$\hat{F}_x(i) = \frac{\sum_{j=0}^i h(j)}{\sum_{j=0}^{L-1} h(j)}$$

Here the normalization term assures that $F_x(L - 1) = 1$. By applying the concept of (1), we equalize a pixel X_s at the position $s \in S$ where S is the set of positions in the image.

$$Y_s = \hat{F}_x(X_s)$$

However, Y_s has a range from 0 to 1 and may not extend over the maximum number of gray levels. To correct these problems, we first compute the minimum and maximum values of Y_s .

$$\begin{aligned} Y_{max} &= \max_{s \in S} Y_s \\ Y_{min} &= \min_{s \in S} Y_s \end{aligned}$$

And then we use these values to form Z_s , a normalized version of Y_s .

$$Z_s = (L - 1) \frac{\hat{F}_x(X_s) - Y_{min}}{Y_{max} - Y_{min}} \tag{2}$$

Histogram equalization is then the transformation from X_s to Z_s .

1. Write a Python function `def equalize(X)` that equalizes the histogram of the image X using equation (2).
2. Download the image *kids.tif* from the lab home page.
3. Compute and plot the function $\hat{F}_x(i)$ for the image *kids.tif*.
4. Apply your *equalize* function to the *kids.tif* image. Print or export the equalized image for your report.
5. Plot the histogram of the equalized image and print or export the figure. Make sure the axes and title are labeled.
6. Compare the image and its histogram to the originals, which were generated in Section 1.

Section 2 Report:

1. Hand in the function *equalize.m*.
2. Hand in a labeled plot of $\hat{F}_x(i)$ for the image *kids.tif*.
3. Hand in a labeled plot of the of the equalized image's histogram.
4. Hand in the equalized image.

3 Contrast Stretching

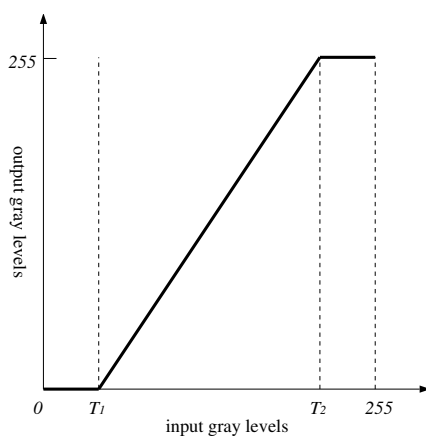


Figure 2: Contrast Stretching Transformation

Another useful image enhancement technique is *contrast stretching*. An example is depicted in Figure 2. The horizontal axis shows all possible intensities of the original image, and the vertical axis shows the intensities of the transformed image. This particular

transformation maps the “darker” pixels in the range $[0, T_1]$ to a level of zero (black), and similarly maps the “lighter” pixels in $[T_2, 255]$ to white. Then the pixels in the range $[T_1, T_2]$ are “stretched out” to use the full scale of $[0, 255]$. This can have the effect of increasing the contrast in an image.

If a pixel transformation can be described by a one-to-one function, $y = f(x)$, then it can be shown that the input and output histograms are approximately related by the following:

$$h_{out}(y) \approx \frac{h_{in}(x)}{|f'(x)|} \Big|_{x=f^{-1}(y)} . \quad (3)$$

Since x and y need to be integers in (3), the evaluation of $x = f^{-1}(y)$ needs to be rounded to the nearest integer.

The pixel transformation shown in Figure 2 is not a one-to-one function. However, equation (3) still may be used to give insight into the effect of the transformation. Since the regions $[0, T_1]$ and $[T_2, 255]$ map to the single points 0 and 255, we might expect “spike” behavior at the points 0 and 255 in the output histogram. The region $[1, 254]$ of the output histogram will be directly related to the input histogram through equation (3).

First, notice from $x = f^{-1}(y)$ that the region $[1, 254]$ of the output is being mapped from the region $[T_1, T_2]$ of the input. Then notice that $f'(x)$ will be a constant scaling factor throughout the entire region of interest. Therefore, the output histogram should approximately be a “stretched” and rescaled version of the input histogram, with possible spikes at the endpoints.

1. Write a Python function *stretch* that will perform the pixel transformation shown in Figure 2. The input and output images should be of type *uint8*. It should have the syntax

```
output = stretch(input, T1, T2) .
```

2. Download *kids.tif* from the lab home page and read it into Python. Display the image, and compute its histogram.
3. Now use your *stretch* function to increase the contrast of the image. Choose *T1* and *T2* so that the output image has a histogram which spans the full range of values from 0 to 255. Display the output image and its histogram. Label both of them and print or export them for your report.
4. Compare the image and its histogram to the originals, which were generated in Section 1.

Section 3 Report:

1. Hand in your code for *stretch*.
2. Hand in the transformed image and its histogram.

4 Gamma (γ)

There have recently been major advances in display technology, but for a long time the most widely used display devices have been cathode ray tubes (CRT). These devices exhibit a nonlinear relationship between the voltage applied at the cathode and the light intensity produced on the phosphor screen. Therefore if pixel values of an image represent (linearly-scaled) light intensity, these values need to be pre-scaled to obtain an accurate reproduction of light intensity on a CRT display. Liquid crystal display (LCD) technology has very different physical properties, but for important reasons we will discuss, most LCD displays use correction tables to mimic the response of a CRT.

For a CRT, the relationship between applied voltage (V) and light intensity (I) on the screen is described by a *power law*.

$$I = V^\gamma \quad (4)$$

The numerical value of the exponent is known as the *gamma* (γ) of the display, and a typical CRT gamma is around 2.5. If we relate equation (4) to the pixel values for an 8 bit image, we get the following relationship,

$$y = 255 \left(\frac{x}{255} \right)^\gamma \quad (5)$$

where x is the original pixel value, and y is the pixel light intensity produced by the display. This relationship is illustrated in Figure 3.

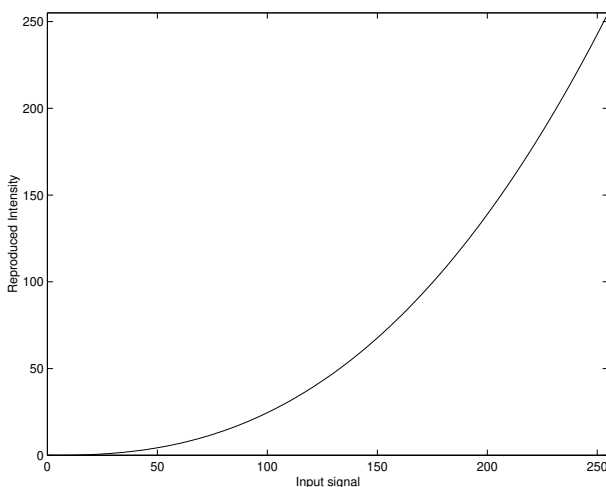


Figure 3: Nonlinear behavior of a display device having a γ of 2.2.

Because of this power law relationship, an image representing linearly-scaled light intensity will appear mostly too dark on the display, especially for small and mid range pixel values. To compensate for this effect, a process called *gamma correction* is performed which consists of applying the inverse of equation (5) to the pixel values (i.e. raise to the power $1/\gamma$) before the display. This is a straightforward pixel transformation, as we discussed in Section 3.

There is another very important reason for transforming pixel values from a linear scale to a “gamma corrected” scale. The human visual system is generally more sensitive to

changes in light intensity at low background illumination levels than at high background illumination. If we quantize an image (e.g. 0,1,...,255), it would be desirable from a coding efficiency standpoint for the change in perceived *brightness* between quantization levels to be uniform across the entire range. In fact quantizing a linearly scaled image to 8 bits results in very easily perceived gray level increments in the darker pixel range, producing *false contours* in dark regions.

By a complete coincidence, the gamma corrected scale we introduced in the context of CRT display properties very nearly matches a *perceptually uniform* scale associated with the human visual system. Since images can be most efficiently quantized in this scale, most digital image capture devices (e.g. digital cameras, scanners, etc.) perform “gamma correction” internally before storing the image data. Therefore most digital images in daily life (e.g. JPEG, MPEG video) are already in a gamma corrected scale.

Situations sometimes arise where we need to consider this gamma transformation, for example if the gamma associated with an image capture device significantly mismatches the gamma of the display device. Also, some image processing applications, especially in computer graphics, are most appropriately performed in a linear light intensity scale. In the following exercises, you will first measure the gamma value of your monitor, and then perform gamma adjustments on a couple digital photographs.

4.1 Setting the Black Level and Picture of Your Monitor

Before you can accurately determine the gamma for your computer monitor, it is essential that you properly set the monitor parameters called *Black Level* and *Picture*, also respectively known as *brightness* and *contrast*. Standard symbols for these controls are shown in Figure 4.

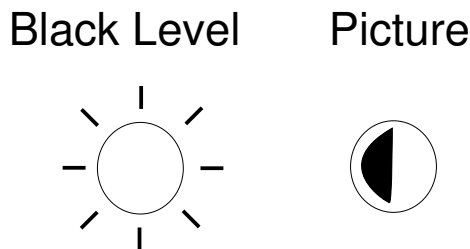


Figure 4: Video Monitor Controls

The following is the proper procedure for setting the Black Level and Picture. See [1] for a detailed explanation of these parameters.

1. Set the Picture control to approximately half-way.
2. Adjust the Black Level just to the point where there is no background illumination of the phosphors in regions of the screen that should be black. In other words, the black regions should be as dark as possible, but any increase in Black Level will cause those phosphors to illuminate. This maximizes the contrast ratio of the monitor. Note: Once the Black Level is properly set, it should not need to be changed again.
3. Adjust the Picture to obtain a comfortable brightness.

4.2 Determining the Gamma of Your Computer Monitor

In this section, you will determine the gamma of your monitor by using a simple experiment that exploits properties of the human visual system. The method works by visually comparing a checkerboard pattern to a fix gray level. The checkerboard pattern is produced by displaying a repeated pattern of black and white dots with the values shown in Fig. 5. When viewed from a distance, this checkerboard pattern appears as a uniform gray level due to the low pass blur from the optics of the eye. The perceived intensity of the checkerboard is given by

$$I_c = (I_{255} + 0)/2$$

where I_{255} is the intensity produced by a full intensity dot of level 255. The intensity of this checkerboard is then compared to a fixed gray level produce by an input $0 \leq g \leq 255$. Assuming a standard gamma model for the behavior of the monitor, the gray level g produces a perceived intensity of

$$I_g = I_{255}(g/255)^\gamma .$$

We can then calculate the value of γ for the monitor by determining the gray level g which makes $I_g = I_c$.

In order to visually match the two gray values, you will produce a patch containing stripes of constant gray interlaced with stripes of checkerboard pattern. Such a patch is shown in Fig. 6. You will then adjust the gray level until the stripe pattern disappears when viewed from a distance.

255	255	0	0
255	255	0	0
0	0	255	255
0	0	255	255

Figure 5: Fundamental Unit of Checkerboard Pattern

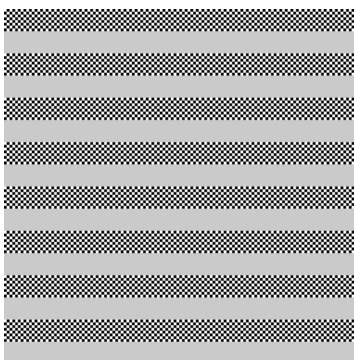


Figure 6: Array Pattern for Determining Gamma

Do the following to determine the value of gamma for your video monitor.

1. Set the Black Point of your monitor as described in Section 4.1.
2. Write a Python function that will display a 256×256 array containing horizontal stripes which alternate between a checkerboard pattern and a constant gray level, similar to Figure 6.
 - Since the gray level strips must be adjusted to match the checkerboard, the gray level should be an input to your function.
 - Each stripe should be 16 pixels in height.
 - Each block of the checkerboard pattern should consist of four pixels, as shown in Figure 5.
 - Place the appropriate commands within the function to display the array.
3. Initially display the array using a gray level of 127. Then adjust the gray level until the stripes appear to have the same intensity.
4. Record the value of the matching gray level that produces the best intensity match between the stripes.
5. Print or export the image corresponding to the matching gray level.
6. Derive an analytical expression for γ in terms of the matching gray level and calculate the value of γ for your monitor.

Do the regions appear to have the same intensity in the printed version?

Section 4.2 Report:

1. Hand in your image corresponding to the matching gray level.
2. Hand in a derivation of the expression which relates the matching gray level to the value of γ .
3. Hand in the values of the measured gray level and the measured γ .

4.3 Gamma Correction

As we covered, *gamma correction* is used to compensate for the nonlinear behavior of a display device. Most often images are already encoded in gamma corrected form, and will appear fine when displayed on most video monitors. However, if an image is stored with a linear scaling it becomes necessary to correct the image. If the value of gamma for your monitor is known, then the correction process consists of applying the inverse of equation (5).

1. Determine the value of gamma for the monitor you are currently using by following the procedure of Section 4.2.
2. Download the image *linear.tif* from the lab home page. This image is encoded on a linear scale. Use Python to display the image, and include the image in your report.
3. Gamma correct *linear.tif* for your monitor.
4. Use Python to display the gamma corrected form of *linear.tif*, and print or export the image for your report.

Section 4.3 Report:

1. Hand in the original and corrected images. Label them and indicate the value of gamma that was used to correct the image.
2. Hand in the formula you used to transform the original image.

Sometimes an image is encoded in gamma corrected format, but not for a gamma corresponding to your monitor. To accurately reproduce the image, it is necessary to change the gamma correction for the image.

1. Download the image *gamma15.tif* from the lab home page. This image is already gamma corrected for a $\gamma = 1.5$.
2. Derive the procedure for transforming *gamma15.tif* so that it will be displayed properly on your monitor.
3. Use Python to display the corrected version, and print or export the image for your report.

Section 4.3 Report:

1. Hand in the corrected image. Be sure it is labeled in the report.
2. Document the procedure you used to change the gamma correction of the original image.

References

- [1] C. A. Poynton, "Black Level" and "Picture", 1997.