6.1 Lookup Table

The internal representation of a digital image is simply a matrix of numbers representing grayscale or color values. But when an image is displayed on a computer monitor we typically do not see a direct mapping of the image. An image is first passed through a *lookup table* (LUT) that maps the image intensity values to brightness values, Figure 6.1. If the lookup table is linear with unit slope and zero intercept then the image is directly mapped to the display, otherwise, the displayed image will not be an exact representation of the underlying image. For example, most computer monitors intentionally impose a non-linear LUT of the general form $D = I^\alpha$ (i.e., gamma correction), where α is a real number, and D and I are the displayed and image values. A variety of interesting visual effects can be achieved by simple manipulations of the functional form of the LUT. Keep in mind though that in manipulating the lookup table, the underlying image is left untouched, it is only the *mapping* from pixel value to display brightness that is being effected.

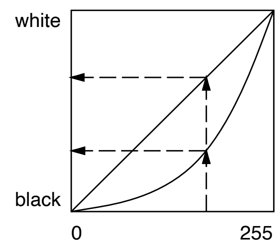

Figure 6.1 Lookup table

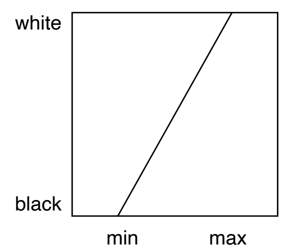

Figure 6.2 Autoscale

6.2 Brightness/Contrast

Perhaps the most common and familiar example of a LUT manipulation is to control the brightness or darkness of an image as shown in Figure 6.3. The bright and dark images of Einstein were created by passing the middle image through the LUTs shown in the same figure. The functional form of the LUT is a unit-slope line with varying intercepts: $g(u) = u + b$, with the image intensity values $u \in [0, 1]$. A value of $b > 0$ results in a brightening of the image and $b < 0$ a darkening of the image. Another common manipulation is that of controlling the contrast of an image as shown in Figure 6.4. The top image is said to be high contrast and the bottom image low contrast, with the corresponding LUTs shown in the same figure. The functional form of these LUTs is linear: $g(u) = mu + b$, where the relationship between the slope and intercept is $b = 1/2(1 - m)$. The image contrast is increased with a large slope and negative intercept (in the limit, $m \to \infty$ and $b \to -\infty$), and the contrast is reduced with a small slope and positive intercept ($m \to 0$ and $b \to 1/2$). The image is *inverted* with $m = -1$ and $b = 1$, that is white is mapped to black, and black is mapped to white. Of course, an image can be both contrast enhanced and brightened or darkened by simply passing the image through two (or more) LUTs.

Autoscaling is a special case of contrast enhancement where the minimum image intensity value is mapped to black and the maximum value is mapped to white, Figure 6.2. Autoscaling maximizes the contrast without saturating at black or white. The problem with this sort of autoscaling is that a few stray pixels can dictate the contrast resulting in a low-contrast image. A less sensitive approach is to sort the intensity values in increasing order and map the 1% and 99% intensity values to black and white, respectively. Although this will lead to a small amount of saturation, it rarely fails to produce a high-contrast image.

Figure 6.3 Brightness

6.3 Gamma Correction

Typically, high contrast images are visually more appealing. However a drawback of linear contrast enhancement described above is that it leads to saturation at both the low and high end of the intensity range. This may be avoided by employing a non-linear contrast adjustment scheme, also realizable as a lookup table (LUT) manipulation. The most standard approach is *gamma correction*, where the LUT takes the form:

$$g(u) \quad = \quad u^{\alpha}, \tag{6.1}$$

where $\alpha > 1$ increases contrast, and $\alpha < 1$ reduces contrast. Shown in Figure 6.5 are contrast enhanced (top: $\alpha = 2$) and contrast reduced (bottom: $\alpha = 1/2$) images. Note that with the intensity values u scaled into the range $[0,1]$, black (0) and white (1) are mapped to themselves. That is, there is no saturation at the low or high end. Gamma correction is widely used in a number of devices because it yields reasonable results and is easily parameterized. One drawback to this scheme is that the gray values are mapped in an asymmetric fashion with respect to mid-level gray (0.5). This may be alleviated by employing a *sigmoidal* non-linearity of the form

$$g(u) \quad = \quad \frac{1}{1+e^{-\alpha u+\beta}}. \tag{6.2}$$

In order that $g(u)$ be bound by the interval $[0,1]$, it must be scaled as follows: $(g(u) - c_1)/c_2$, where $c_1 = 1/(1+e^{\beta})$ and $c_2 = 1/(1+e^{-\alpha+\beta}) - c1$. This non-linearity, with its two degrees of freedom, is more versatile and can produce a more balanced contrast enhancement. Shown in Figure 6.6 is a contrast enhanced image with $\alpha = 12$ and $\beta = 6$.

Figure 6.4 Contrast

6.4 Quantize/Threshold

A digital image, by its very nature, is *quantized* to a discrete number of intensity values. For example an image quantized to 8-bits contains $2^8 = 256$ possible intensity values, typically in the range $[0, 255]$. An image can be further quantized to a lower number of bits (b) or intensity values (2^b). Quantization can be accomplished by passing the image through a LUT containing a step function, Figure 6.7, where the number of steps governs the number of intensity values. Shown in Figure 6.7 is an image of Einstein quantized to five intensity values, notice that all the subtle variations in the curtain and in his face and jacket have been eliminated. In the limit, when an image is quantized to one bit or two intensity values, the image is said to be *thresholded*. Shown in Figure 6.8 is a thresholded image of Einstein and the corresponding LUT, a two-step function. The point at which the step function transitions from zero to one is called the threshold and can of course be made to be any value (i.e., slid left or right).

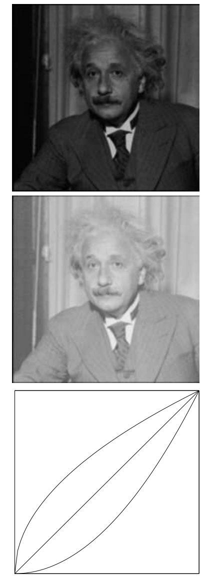

Figure 6.5 Contrast: Gamma

6.5 Histogram Equalize

The intensity values of a typical image are often distributed unevenly across the full range of 0 to 255 (for an 8-bit image), with most the mass near mid-gray (128) and falling off on either side, Figure 6.9. An image can be transformed so that the distribution of intensity values is flat, that is, each intensity value is equally represented in the image. This process is known has *histogram equalization* [7]. Although it may not be immediately obvious an image is histogram equalized by passing it through a LUT with the functional form of the cumulative distribution function. More specifically, define $N(u)$ as the number of pixels with intensity value u, this is the image histogram and a discrete approximation to the probability distribution function. Then, the cumulative distribution function is defined as:

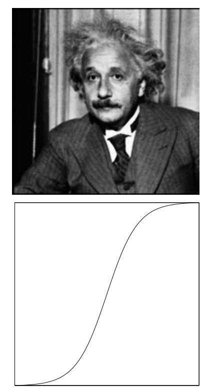

Figure 6.6 Contrast: Sigmoid

$$C(y) = \sum_{i=0}^{u} N(i), \tag{6.3}$$

that is, $C(u)$ is the number of pixels with intensity value less than or equal to u. Histogram equalization then amounts to simply inserting this function into the LUT. Shown in Figure 6.9 is Einstein before and after histogram equalization. Notice that the effect is similar to contrast enhancement, which intuitively should make sense since we increased the number of black and white pixels.

[7] Why anyone would want to histogram equalize an image is a mystery to me, but here it is in case you do.

In all of these examples the appearance of an image was altered by simply manipulating the LUT, the mapping from image intensity value to display brightness value. Such a manipulation leaves the image content intact, it is a non-destructive operation and thus completely invertible. These operations can be made destructive by applying the LUT operation directly to the image. For example an image can be brightened by adding a constant to each pixel, and then displaying with a linear LUT. Since such an operation is destructive it may not be inveritble, for example when brightening an 8-bit image, all pixels that exceed the value 255 will be truncated to 255.

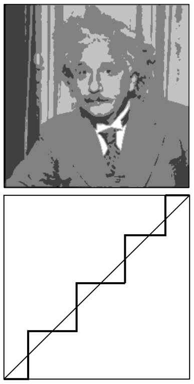

Figure 6.7 Quantize

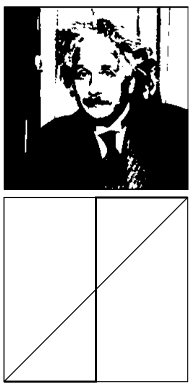

Figure 6.8 Threshold

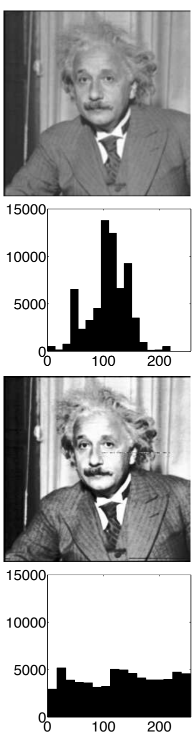

Figure 6.9 Histogram equalize