
Point Operations

Point operations perform a modification of the pixel values without changing the size, geometry, or local structure of the image. Each new pixel value $b = I'(u, v)$ depends exclusively on the previous value $a = I(u, v)$ at the *same* position and is thus independent from any other pixel value, in particular from any of its neighboring pixels.¹ The original pixel values a are mapped to the new values b by some given function f , i.e.,

$$b = f(I(u, v)) \quad \text{or} \quad b = f(a). \quad (4.1)$$

If, as in this case, the function $f()$ is independent of the image coordinates (i.e., the same throughout the image), the operation is called “global” or “homogeneous”. Typical examples of homogeneous point operations include, among others:

- modifying image brightness or contrast,
- applying arbitrary intensity transformations (“curves”),
- inverting images,
- quantizing (or “posterizing”) images,
- global thresholding,
- gamma correction,
- color transformations
- etc.

We will look at some of these techniques in more detail in the following.

In contrast to Eqn. (4.1), the mapping function $g()$ for a *nonhomogeneous* point operation would also take into account the current image coordinate (u, v) , that is,

$$b = g(I(u, v), u, v) \quad \text{or} \quad b = f(a, u, v). \quad (4.2)$$

A typical nonhomogeneous operation is the local adjustment of contrast or brightness used, for example, to compensate for uneven lighting during image acquisition.

¹ If the result depends on more than one pixel value, the operation is called a “filter”, as described in Chapter 5.

4.1 Modifying Image Intensity

4.1.1 Contrast and Brightness

Let us start with a simple example. Increasing the image’s contrast by 50% (i.e., by the factor 1.5) or raising the brightness by 10 units can be expressed by the mapping functions

$$f_{\text{contr}}(a) = a \cdot 1.5 \quad \text{or} \quad f_{\text{bright}}(a) = a + 10, \quad (4.3)$$

respectively. The first operation is implemented as an ImageJ plugin by the code shown in Prog. 4.1, which can easily be adapted to perform any other type of point operation. Rounding to the nearest integer values is accomplished by simply adding 0.5 before the truncation effected by the `(int)` typecast in line 8 (this only works for positive values). Also note the use of the more efficient image processor methods `get()` and `set()` (instead of `getPixel()` and `putPixel()`) in this example.

Prog. 4.1

Point operation to increase the contrast by 50% (ImageJ plugin). Note that in line 8 the result of the multiplication of the integer pixel value by the constant 1.5 (implicitly of type `double`) is of type `double`.

Thus an explicit type cast (`int`) is required to assign the value to the `int` variable `a`. 0.5 is added in line 8 to round to the nearest integer values.

```

1  public void run(ImageProcessor ip) {
2      int w = ip.getWidth();
3      int h = ip.getHeight();
4
5      for (int v = 0; v < h; v++) {
6          for (int u = 0; u < w; u++) {
7              int a = ip.get(u, v);
8              int b = (int) (a * 1.5 + 0.5);
9              if (b > 255)
10                 b = 255; // clamp to the maximum value ( $a_{\max}$ )
11                 ip.set(u, v, b);
12             }
13         }
14     }

```

4.1.2 Limiting Values by Clamping

When implementing arithmetic operations on pixel values, we must keep in mind that the calculated results must not exceed the admissible range of pixel values for the given image type (e.g., $[0, 255]$ in the case of 8-bit grayscale images). This is commonly called “clamping” and can be expressed in the form

$$b = \min(\max(f(a), a_{\min}), a_{\max}) = \begin{cases} a_{\min} & \text{for } f(a) < a_{\min}, \\ a_{\max} & \text{for } f(a) > a_{\max}, \\ f(a) & \text{otherwise.} \end{cases} \quad (4.4)$$

For this purpose, line 10 of Prog. 4.1 contains the statement

```
if (b > 255) b = 255;
```

which limits the result to the maximum value 255. Similarly, one may also want to limit the results to the minimum value (0) to avoid negative pixel values (which cannot be represented by this type of 8-bit image), for example, by the statement

```
if (b < 0) b = 0;
```

The above statement is not needed in Prog. 4.1 because the intermediate results can never be negative in this particular operation.

4.1.3 Inverting Images

Inverting an intensity image is a simple point operation that reverses the ordering of pixel values (by multiplying by -1) and adds a constant value to map the result to the admissible range again. Thus for a pixel value $a = I(u, v)$ in the range $[0, a_{\max}]$, the corresponding point operation is

$$f_{\text{inv}}(a) = -a + a_{\max} = a_{\max} - a. \quad (4.5)$$

The inversion of an 8-bit grayscale image with $a_{\max} = 255$ was the task of our first plugin example in Sec. 2.2.4 (Prog. 2.1). Note that in this case no clamping is required at all because the function always maps to the original range of values. In ImageJ, this operation is performed by the method `invert()` (for objects of type `ImageProcessor`) and is also available through the **Edit** \triangleright **Invert** menu. Obviously, inverting an image mirrors its histogram, as shown in Fig. 4.5(c).

4.1.4 Threshold Operation

Thresholding an image is a special type of quantization that separates the pixel values in two classes, depending upon a given threshold value q that is usually constant. The threshold operation maps all pixels to one of two fixed intensity values a_0 or a_1 , that is,

$$f_{\text{threshold}}(a) = \begin{cases} a_0 & \text{for } a < q, \\ a_1 & \text{for } a \geq q, \end{cases} \quad (4.6)$$

with $0 < q \leq a_{\max}$. A common application is *binarizing* an intensity image with the values $a_0 = 0$ and $a_1 = 1$.

ImageJ does provide a special image type (`BinaryProcessor`) for binary images, but these are actually implemented as 8-bit intensity images (just like ordinary intensity images) using the values 0 and 255. ImageJ also provides the `ImageProcessor` method `threshold(int level)`, with $level \equiv q$, to perform this operation, which can also be invoked through the **Image** \triangleright **Adjust** \triangleright **Threshold** menu (see Fig. 4.1 for an example). Thresholding affects the histogram by separating the distribution into two entries at positions a_0 and a_1 , as illustrated in Fig. 4.2.

4.2 Point Operations and Histograms

We have already seen that the effects of a point operation on the image's histogram are quite easy to predict in some cases. For example, increasing the brightness of an image by a constant value shifts the entire histogram to the right, raising the contrast widens

4 POINT OPERATIONS

Fig. 4.1

Threshold operation: original image (a) and corresponding histogram (c); result after thresholding with $a_{th} = 128$, $a_0 = 0$, $a_1 = 255$ (b) and corresponding histogram (d); ImageJ's interactive Threshold menu (e).

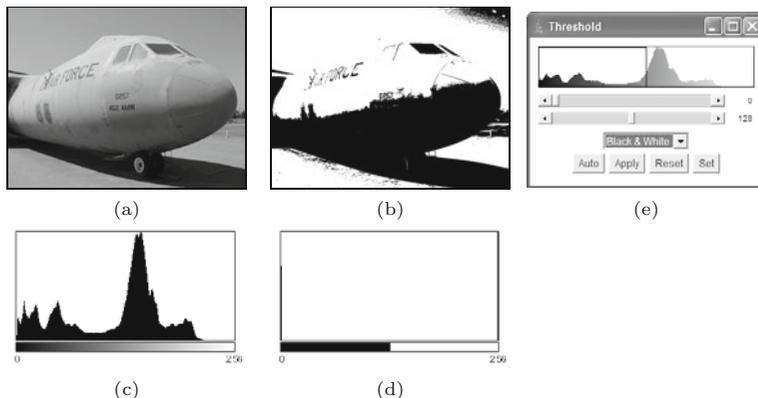
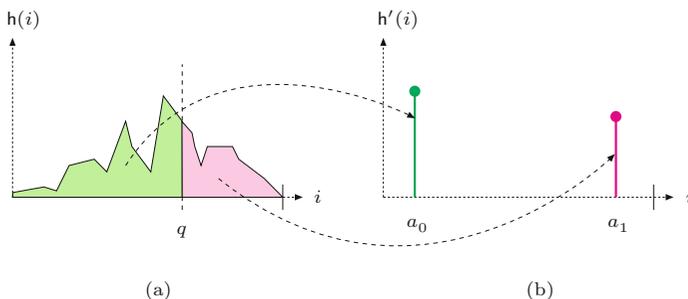


Fig. 4.2

Effects of thresholding upon the histogram. The threshold value is a_{th} . The original distribution (a) is split and merged into two isolated entries at a_0 and a_1 in the resulting histogram (b).



it, and inverting the image flips the histogram. Although this appears rather simple, it may be useful to look a bit more closely at the relationship between point operations and the resulting changes in the histogram.

As the illustration in Fig. 4.3 shows, every entry (bar) at some position i in the histogram maps to a *set* (of size $h(i)$) containing all image pixels whose values are exactly i .²

If a particular histogram line is *shifted* as a result of some point operation, then of course all pixels in the corresponding set are equally modified and vice versa. So what happens when a point operation (e.g., reducing image contrast) causes two previously separated histogram lines to fall together at the same position i ? The answer is that the corresponding pixel sets are *merged* and the new common histogram entry is the sum of the two (or more) contributing entries (i.e., the size of the combined set). At this point, the elements in the merged set are no longer distinguishable (or separable), so this operation may have (perhaps unintentionally) caused an irreversible reduction of dynamic range and thus a permanent loss of information in that image.

² Of course this is only true for ordinary histograms with an entry for every single intensity value. If *binning* is used (see Sec. 3.4.1), each histogram entry maps to pixels within a certain *range* of values.

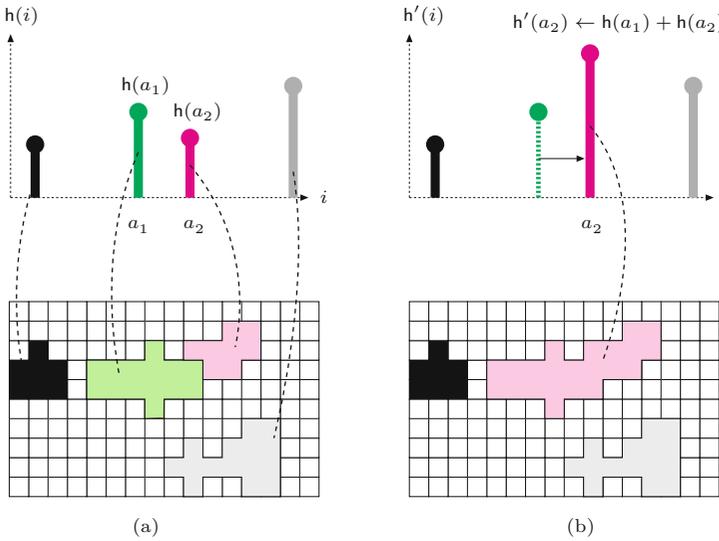


Fig. 4.3
Histogram entries represent sets of pixels of the same value. If a histogram line is moved as a result of some point operation, then all pixels in the corresponding set are equally modified (a). If, due to this operation, two histogram lines $h(a_1)$, $h(a_2)$ coincide on the same index, the two corresponding pixel sets merge and the contained pixels become indiscernable (b).

4.3 Automatic Contrast Adjustment

Automatic contrast adjustment (auto-contrast) is a point operation whose task is to modify the pixels such that the available range of values is fully covered. This is done by mapping the current darkest and brightest pixels to the minimum and maximum intensity values, respectively, and linearly distributing the intermediate values.

Let us assume that a_{lo} and a_{hi} are the lowest and highest pixel values found in the current image, whose full intensity range is $[a_{min}, a_{max}]$. To stretch the image to the full intensity range (see Fig. 4.4), we first map the smallest pixel value a_{lo} to zero, subsequently increase the contrast by the factor $(a_{max} - a_{min}) / (a_{hi} - a_{lo})$, and finally shift to the target range by adding a_{min} . The mapping function for the auto-contrast operation is thus defined as

$$f_{ac}(a) = a_{min} + (a - a_{lo}) \cdot \frac{a_{max} - a_{min}}{a_{hi} - a_{lo}}, \quad (4.7)$$

provided that $a_{hi} \neq a_{lo}$; that is, the image contains at least *two* different pixel values. For an 8-bit image with $a_{min} = 0$ and $a_{max} = 255$, the function in Eqn. (4.7) simplifies to

$$f_{ac}(a) = (a - a_{lo}) \cdot \frac{255}{a_{hi} - a_{lo}}. \quad (4.8)$$

The target range $[a_{min}, a_{max}]$ need not be the maximum available range of values but can be any interval to which the image should be mapped. Of course the method can also be used to reduce the image contrast to a smaller range. Figure 4.5(b) shows the effects of an auto-contrast operation on the corresponding histogram, where the linear stretching of the intensity range results in regularly spaced gaps in the new distribution.

4 POINT OPERATIONS

Fig. 4.4

Auto-contrast operation according to Eqn. (4.7). Original pixel values a in the range $[a_{lo}, a_{hi}]$ are mapped linearly to the target range $[a_{min}, a_{max}]$.

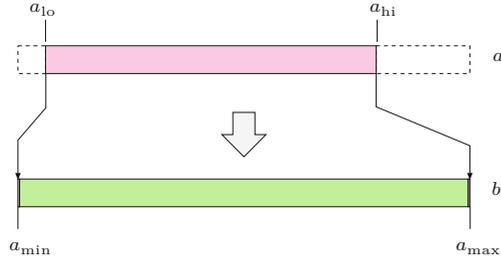
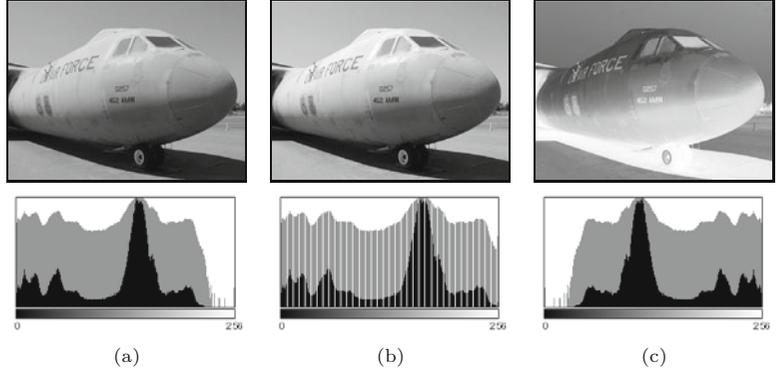


Fig. 4.5

Effects of auto-contrast and inversion operations on the resulting histograms. Original image (a), result of auto-contrast operation (b), and inversion (c). The histogram entries are shown both linearly (black bars) and logarithmically (gray bars).

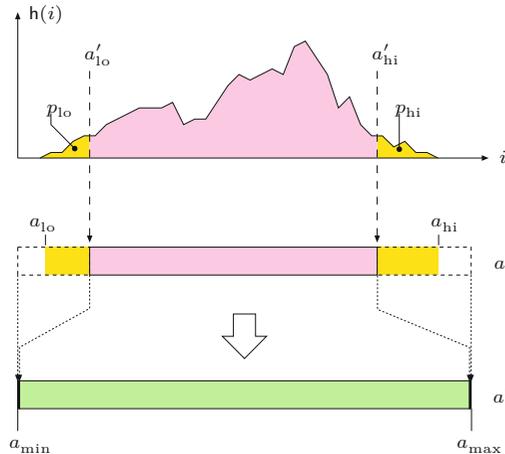


4.4 Modified Auto-Contrast Operation

In practice, the mapping function in Eqn. (4.7) could be strongly influenced by only a few extreme (low or high) pixel values, which may not be representative of the main image content. This can be avoided to a large extent by “saturating” a fixed percentage (p_{lo}, p_{hi}) of pixels at the upper and lower ends of the target intensity range. To accomplish this, we determine two limiting values a'_{lo}, a'_{hi} such that a predefined quantile q_{lo} of all pixel values in the image I are smaller than a'_{lo} and another quantile q_{hi} of the values are greater than a'_{hi} (Fig. 4.6).

Fig. 4.6

Modified auto-contrast operation (Eqn. (4.11)). Predefined quantiles (q_{lo}, q_{hi}) of image pixels—shown as dark areas at the left and right ends of the histogram $h(i)$ —are “saturated” (i.e., mapped to the extreme values of the target range). The intermediate values ($a = a'_{lo}, \dots, a'_{hi}$) are mapped linearly to the interval a_{min}, \dots, a_{max} .



The values a'_{lo} , a'_{hi} depend on the image content and can be easily obtained from the image's cumulative histogram³ H :

$$a'_{lo} = \min\{i \mid H(i) \geq M \cdot N \cdot p_{lo}\}, \quad (4.9)$$

$$a'_{hi} = \max\{i \mid H(i) \leq M \cdot N \cdot (1 - p_{hi})\}, \quad (4.10)$$

where $0 \leq p_{lo}, p_{hi} \leq 1$, $p_{lo} + p_{hi} \leq 1$, and $M \cdot N$ is the number of pixels in the image. All pixel values *outside* (and including) a'_{lo} and a'_{hi} are mapped to the extreme values a_{min} and a_{max} , respectively, and intermediate values are mapped linearly to the interval $[a_{min}, a_{max}]$. Using this formulation, the mapping to minimum and maximum intensities does not depend on singular extreme pixels only but can be based on a representative set of pixels. The mapping function for the modified auto-contrast operation can thus be defined as

$$f_{mac}(a) = \begin{cases} a_{min} & \text{for } a \leq a'_{lo}, \\ a_{min} + (a - a'_{lo}) \cdot \frac{a_{max} - a_{min}}{a'_{hi} - a'_{lo}} & \text{for } a'_{lo} < a < a'_{hi}, \\ a_{max} & \text{for } a \geq a'_{hi}. \end{cases} \quad (4.11)$$

Usually the same value is taken for both upper and lower quantiles (i.e., $p_{lo} = p_{hi} = p$), with $p = 0.005, \dots, 0.015$ (0.5%, ..., 1.5%) being common values. For example, the auto-contrast operation in Adobe Photoshop saturates 0.5% ($p = 0.005$) of all pixels at both ends of the intensity range. Auto-contrast is a frequently used point operation and thus available in practically any image-processing software. ImageJ implements the modified auto-contrast operation as part of the **Brightness/Contrast** and **Image > Adjust** menus (**Auto** button), shown in Fig. 4.7.

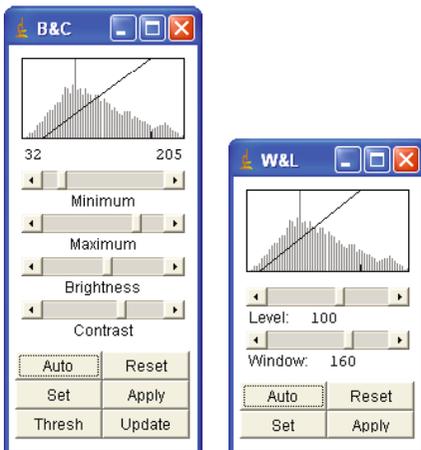


Fig. 4.7 ImageJ's **Brightness/Contrast** tool (left) and **Window/Level** tool (right) can be invoked through the **Image > Adjust** menu. The **Auto** button displays the result of a modified auto-contrast operation. **Apply** must be hit to actually modify the image.

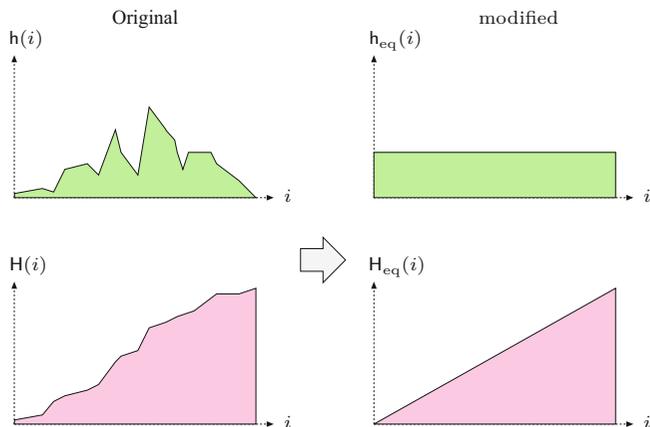
4.5 Histogram Equalization

A frequent task is to adjust two different images in such a way that their resulting intensity distributions are similar, for example, to use

³ See Sec. 3.6.

Fig. 4.8

Histogram equalization. The idea is to find and apply a point operation to the image (with original histogram h) such that the histogram h_{eq} of the modified image approximates a *uniform* distribution (top). The cumulative target histogram H_{eq} must thus be approximately wedge-shaped (bottom).

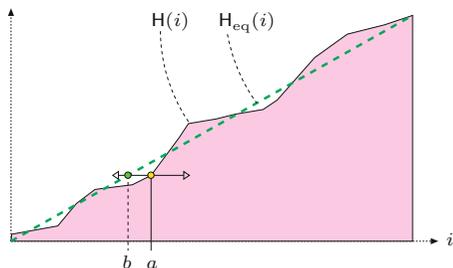


them in a print publication or to make them easier to compare. The goal of histogram equalization is to find and apply a point operation such that the histogram of the modified image approximates a *uniform* distribution (see Fig. 4.8). Since the histogram is a discrete distribution and homogeneous point operations can only shift and merge (but never split) histogram entries, we can only obtain an approximate solution in general. In particular, there is no way to eliminate or decrease individual peaks in a histogram, and a truly uniform distribution is thus impossible to reach. Based on point operations, we can thus modify the image only to the extent that the resulting histogram is *approximately* uniform. The question is how good this approximation can be and exactly which point operation (which clearly depends on the image content) we must apply to achieve this goal.

We may get a first idea by observing that the *cumulative* histogram (Sec. 3.6) of a uniformly distributed image is a linear ramp (wedge), as shown in Fig. 4.8. So we can reformulate the goal as finding a point operation that shifts the histogram lines such that the resulting cumulative histogram is approximately linear, as illustrated in Fig. 4.9.

Fig. 4.9

Histogram equalization on the cumulative histogram. A suitable point operation $b \leftarrow f_{\text{eq}}(a)$ shifts each histogram line from its original position a to b (left or right) such that the resulting cumulative histogram H_{eq} is approximately linear.



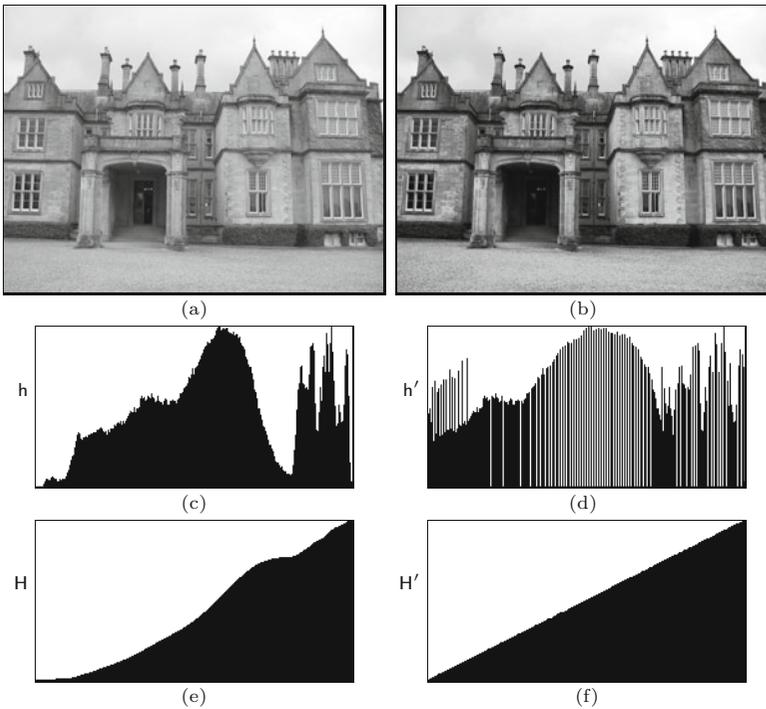
The desired point operation $f_{\text{eq}}()$ is simply obtained from the cumulative histogram H of the original image as⁴

$$f_{\text{eq}}(a) = \left\lfloor H(a) \cdot \frac{K-1}{M \cdot N} \right\rfloor, \quad (4.12)$$

⁴ For a derivation, see, for example, [88, p. 173].

4.5 HISTOGRAM EQUALIZATION

Fig. 4.10
Linear histogram equalization example. Original image I (a) and modified image I' (b), corresponding histograms h , h' (c, d), and cumulative histograms H , H' (e, f). The resulting cumulative histogram H' (f) approximates a uniformly distributed image. Notice that new peaks are created in the resulting histogram h' (d) by merging original histogram cells, particularly in the lower and upper intensity ranges.



for an image of size $M \times N$ with pixel values a in the range $[0, K-1]$. The resulting function $f_{\text{eq}}(a)$ in Eqn. (4.12) is monotonically increasing, because $H(a)$ is monotonic and K , M , N are all positive constants. In the (unusual) case where an image is already uniformly distributed, linear histogram equalization should not modify that image any further. Also, repeated applications of linear histogram equalization should not make any changes to the image after the first time. Both requirements are fulfilled by the formulation in Eqn. (4.12). Program 4.2 lists the Java code for a sample implementation of linear histogram equalization. An example demonstrating the effects on the image and the histograms is shown in Fig. 4.10.

Notice that for “inactive” pixel values i (i.e., pixel values that do not appear in the image, with $h(i) = 0$), the corresponding entries in the cumulative histogram $H(i)$ are either zero or identical to the neighboring entry $H(i-1)$. Consequently a contiguous range of zero values in the histogram $h(i)$ corresponds to a constant (i.e., flat) range in the cumulative histogram $H(i)$, and the function $f_{\text{eq}}(a)$ maps all “inactive” intensity values within such a range to the next lower “active” value. This effect is not relevant, however, since the image contains no such pixels anyway. Nevertheless, a linear histogram equalization may (and typically will) cause histogram lines to merge and consequently lead to a loss of dynamic range (see also Sec. 4.2).

This or a similar form of linear histogram equalization is implemented in almost any image-processing software. In ImageJ it can be invoked interactively through the **Process** \triangleright **Enhance Contrast** menu (option **Equalize**). To avoid extreme contrast effects, the histogram

4 POINT OPERATIONS

Prog. 4.2

Linear histogram equalization (ImageJ plugin). First the histogram of the image `ip` is obtained using the standard ImageJ method `ip.getHistogram()` in line 7. In line 9, the cumulative histogram is computed “in place” based on the recursive definition in Eqn. (3.6). The `int` division in line 16 implicitly performs the required floor (`[]`) operation by truncation.

```
1 public void run(ImageProcessor ip) {
2     int M = ip.getWidth();
3     int N = ip.getHeight();
4     int K = 256; // number of intensity values
5
6     // compute the cumulative histogram:
7     int[] H = ip.getHistogram();
8     for (int j = 1; j < H.length; j++) {
9         H[j] = H[j - 1] + H[j];
10    }
11
12    // equalize the image:
13    for (int v = 0; v < N; v++) {
14        for (int u = 0; u < M; u++) {
15            int a = ip.get(u, v);
16            int b = H[a] * (K - 1) / (M * N); // s. Equation (4.12)
17            ip.set(u, v, b);
18        }
19    }
20 }
```

equalization in ImageJ by default⁵ cumulates the *square root* of the histogram entries using a modified cumulative histogram of the form

$$\tilde{H}(i) = \sum_{j=0}^i \sqrt{h(j)}. \quad (4.13)$$

4.6 Histogram Specification

Although widely implemented, the goal of linear histogram equalization—a uniform distribution of intensity values (as described in the previous section)—appears rather ad hoc, since good images virtually never show such a distribution. In most real images, the distribution of the pixel values is not even remotely uniform but is usually more similar, if at all, to perhaps a Gaussian distribution. The images produced by linear equalization thus usually appear quite unnatural, which renders the technique practically useless.

Histogram *specification* is a more general technique that modifies the image to match an arbitrary intensity distribution, including the histogram of a given image. This is particularly useful, for example, for adjusting a set of images taken by different cameras or under varying exposure or lighting conditions to give a similar impression in print production or when displayed. Similar to histogram equalization, this process relies on the alignment of the cumulative histograms by applying a homogeneous point operation. To be independent of the image size (i.e., the number of pixels), we first define *normalized* distributions, which we use in place of the original histograms.

⁵ The “classic” linear approach (see Eqn. (3.5)) is used when simultaneously keeping the `Alt` key pressed.

4.6.1 Frequencies and Probabilities

The value in each histogram cell describes the observed frequency of the corresponding intensity value, i.e., the histogram is a discrete *frequency distribution*. For a given image I of size $M \times N$, the sum of all histogram entries $h(i)$ equals the number of image pixels,

$$\sum_i h(i) = M \cdot N. \quad (4.14)$$

The associated *normalized* histogram,

$$p(i) = \frac{h(i)}{M \cdot N}, \quad \text{for } 0 \leq i < K, \quad (4.15)$$

is usually interpreted as the *probability distribution* or *probability density function* (pdf) of a random process, where $p(i)$ is the probability for the occurrence of the pixel value i . The cumulative probability of i being any possible value is 1, and the distribution \mathbf{p} must thus satisfy

$$\sum_{i=0}^{K-1} p(i) = 1. \quad (4.16)$$

The statistical counterpart to the cumulative histogram H (Eqn. (3.5)) is the discrete *distribution function* $P()$ (also called the *cumulative distribution function* or cdf),

$$P(i) = \frac{H(i)}{H(K-1)} = \frac{H(i)}{M \cdot N} = \sum_{j=0}^i \frac{h(j)}{M \cdot N} = \sum_{j=0}^i p(j), \quad (4.17)$$

for $i = 0, \dots, K-1$. The computation of the cdf from a given histogram \mathbf{h} is outlined in Alg. 4.1. The resulting function $P(i)$ is (as the cumulative histogram) monotonically increasing and, in particular,

$$P(0) = p(0) \quad \text{and} \quad P(K-1) = \sum_{i=0}^{K-1} p(i) = 1. \quad (4.18)$$

This statistical formulation implicitly treats the generation of images as a random process whose exact properties are mostly unknown.⁶ However, the process is usually assumed to be homogeneous (independent of the image position); that is, each pixel value is the result of a “random experiment” on a single random variable i . The observed frequency distribution given by the histogram $h(i)$ serves as a (coarse) estimate of the probability distribution $p(i)$ of this random variable.

4.6.2 Principle of Histogram Specification

The goal of histogram specification is to modify a given image I_A by some point operation such that its distribution function P_A matches

⁶ Statistical modeling of the image generation process has a long tradition (see, e.g., [128, Ch. 2]).

Alg. 4.1

Calculation of the cumulative distribution function (cdf) $P(i)$ from a given histogram h of length K . See Prog. 4.3 (p. 73) for the corresponding Java implementation.

```

1: Cdf(h)
   Returns the cumulative distribution function  $P(i) \in [0, 1]$  for a
   given histogram  $h(i)$ , with  $i = 0, \dots, K-1$ .
2: Let  $K \leftarrow \text{Size}(h)$ 
3: Let  $n \leftarrow \sum_{i=0}^{K-1} h(i)$ 
4: Create map  $P: [0, K-1] \mapsto \mathbb{R}$ 
5: Let  $c \leftarrow 0$ 
6: for  $i \leftarrow 0, \dots, K-1$  do
7:      $c \leftarrow c + h(i)$                                 ▷ cumulate histogram values
8:      $P(i) \leftarrow c/n$ 
9: return  $P$ .
    
```

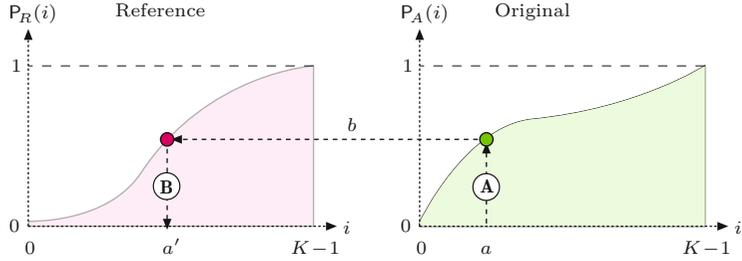
Fig. 4.11

Principle of histogram specification. Given is the reference distribution P_R (left) and the distribution function for the original image P_A (right). The result is the mapping function

$f_{\text{hs}}: a \rightarrow a'$ for a point operation, which replaces each pixel a in the original image I_A by a modified value a' . The process has two main steps:

- Ⓐ For each pixel value a , determine $b = P_A(a)$ from the right distribution function.
- Ⓑ a' is then found by inverting the left distribution function as $a' = P_R^{-1}(b)$.

In summary, the result is $f_{\text{hs}}(a) = a' = P_R^{-1}(P_A(a))$.



a reference distribution P_R as closely as possible. We thus look for a mapping function

$$a' = f_{\text{hs}}(a) \tag{4.19}$$

to convert the original image I_A by a point operation to a new image $I_{A'}$ with pixel values a' , such that its distribution function $P_{A'}$ matches P_R , that is,

$$P_{A'}(i) \approx P_R(i), \quad \text{for } 0 \leq i < K. \tag{4.20}$$

As illustrated in Fig. 4.11, the desired mapping f_{hs} is found by combining the two distribution functions P_R and P_A (see [88, p. 180] for details). For a given pixel value a in the original image, we obtain the new pixel value a' as

$$a' = P_R^{-1}(P_A(a)) = P_R^{-1}(b) \tag{4.21}$$

and thus the mapping f_{hs} (Eqn. (4.19)) is defined as

$$f_{\text{hs}}(a) = P_R^{-1}(P_A(a)), \quad \text{for } 0 \leq a < K. \tag{4.22}$$

This of course assumes that $P_R(i)$ is invertible, that is, that the function $P_R^{-1}(b)$ exists for $b \in [0, 1]$.

4.6.3 Adjusting to a Piecewise Linear Distribution

If the reference distribution P_R is given as a continuous, invertible function, then the mapping function f_{hs} can be obtained from Eqn. (4.22) without any difficulty. In practice, it is convenient to specify the (synthetic) reference distribution as a *piecewise linear* function $P_L(i)$; that is, as a sequence of $N+1$ coordinate pairs

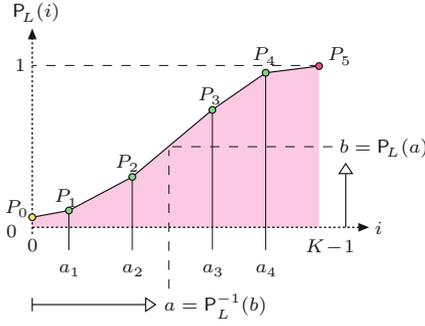


Fig. 4.12
Piecewise linear reference distribution. The function $P_L(i)$ is specified by $N = 5$ control points $(0, P_0)$, $(a_1, P_1), \dots, (a_4, P_4)$, with $a_k < a_{k+1}$ and $P_k < P_{k+1}$. The final point P_5 is fixed at $(K-1, 1)$.

$$L = ((a_0, P_0), (a_1, P_1), \dots, (a_k, P_k), \dots, (a_N, P_N)),$$

each consisting of an intensity value a_k and the corresponding cumulative probability P_k . We assert that $0 \leq a_k < K$, $a_k < a_{k+1}$, and $0 \leq P_k < 1$. Also, the two endpoints (a_0, P_0) and (a_N, P_N) are fixed at

$$(0, P_0) \quad \text{and} \quad (K-1, 1),$$

respectively. To be invertible, the function must also be strictly monotonic, that is, $P_k < P_{k+1}$ for $0 \leq k < N$. Figure 4.12 shows an example of such a function, which is specified by $N = 5$ variable points (P_0, \dots, P_4) and a fixed end point P_5 and thus consists of $N = 5$ linear segments. The reference distribution can of course be specified at an arbitrary accuracy by inserting additional control points.

The intermediate values of $P_L(i)$ are obtained by linear interpolation between the control points as

$$P_L(i) = \begin{cases} P_m + (i - a_m) \cdot \frac{(P_{m+1} - P_m)}{(a_{m+1} - a_m)} & \text{for } 0 \leq i < K-1, \\ 1 & \text{for } i = K-1. \end{cases} \quad (4.23)$$

where $m = \max\{j \in [0, N-1] \mid a_j \leq i\}$ is the index of the line segment $(a_m, P_m) \rightarrow (a_{m+1}, P_{m+1})$, which overlaps the position i . For instance, in the example in Fig. 4.12, the point a lies within the segment that starts at point (a_2, P_2) ; i.e., $m = 2$.

For the histogram specification according to Eqn. (4.22), we also need the *inverse* distribution function $P_L^{-1}(b)$ for $b \in [0, 1]$. As we see from the example in Fig. 4.12, the function $P_L(i)$ is in general not invertible for values $b < P_L(0)$. We can fix this problem by mapping all values $b < P_L(0)$ to zero and thus obtain a “semi-inverse” of the reference distribution in Eqn. (4.23) as

$$P_L^{-1}(b) = \begin{cases} 0 & \text{for } 0 \leq b < P_L(0), \\ a_n + (b - P_n) \cdot \frac{(a_{n+1} - a_n)}{(P_{n+1} - P_n)} & \text{for } P_L(0) \leq b < 1, \\ K-1 & \text{for } b \geq 1. \end{cases} \quad (4.24)$$

Here $n = \max\{j \in \{0, \dots, N-1\} \mid P_j \leq b\}$ is the index of the line segment $(a_n, P_n) \rightarrow (a_{n+1}, P_{n+1})$, which overlaps the argument value b . The required mapping function f_{hs} for adapting a given image with intensity distribution P_A is finally specified, analogous to Eqn. (4.22), as

Alg. 4.2

Histogram specification using a piecewise linear reference distribution. Given is the histogram h of the original image and a piecewise linear reference distribution function, specified as a sequence of N control points L . The discrete mapping f_{hs} for the corresponding point operation is returned.

```

1: MatchPiecewiseLinearHistogram( $h, L$ )
   Input:  $h$ , histogram of the original image  $I$ ;  $L$ , reference distribution function, given as a sequence of  $N + 1$  control points  $L = [(a_0, P_0), (a_1, P_1), \dots, (a_N, P_N)]$ , with  $0 \leq a_k < K$ ,  $0 \leq P_k \leq 1$ , and  $P_k < P_{k+1}$ . Returns a discrete mapping  $f_{\text{hs}}(a)$  to be applied to the original image  $I$ .
2:    $N \leftarrow \text{Size}(L) + 1$ 
3:   Let  $K \leftarrow \text{Size}(h)$ 
4:   Let  $P \leftarrow \text{CDF}(h)$  ▷ cdf for  $h$  (see Alg. 4.1)
5:   Create map  $f_{\text{hs}}: [0, K - 1] \mapsto \mathbb{R}$  ▷ function  $f_{\text{hs}}$ 
6:   for  $a \leftarrow 0, \dots, K - 1$  do
7:      $b \leftarrow P(a)$ 
8:     if ( $b \leq P_0$ ) then
9:        $a' \leftarrow 0$ 
10:    else if ( $b \geq 1$ ) then
11:       $a' \leftarrow K - 1$ 
12:    else
13:       $n \leftarrow N - 1$ 
14:      while ( $n \geq 0$ )  $\wedge$  ( $P_n > b$ ) do ▷ find line segment in  $L$ 
15:         $n \leftarrow n - 1$ 
16:       $a' \leftarrow a_n + (b - P_n) \cdot \frac{(a_{n+1} - a_n)}{(P_{n+1} - P_n)}$  ▷ see Eqn. 4.24
17:       $f_{\text{hs}}[a] \leftarrow a'$ 
18:   return  $f_{\text{hs}}$ .
```

$$f_{\text{hs}}(a) = P_L^{-1}(P_A(a)), \quad \text{for } 0 \leq a < K. \quad (4.25)$$

The whole process of computing the pixel mapping function for a given image (histogram) and a piecewise linear target distribution is summarized in Alg. 4.2. A real example is shown in Fig. 4.14 (Sec. 4.6.5).

4.6.4 Adjusting to a Given Histogram (Histogram Matching)

If we want to adjust one image to the histogram of another image, the reference distribution function $P_R(i)$ is not continuous and thus, in general, cannot be inverted (as required by Eqn. (4.22)). For example, if the reference distribution contains zero entries (i.e., pixel values k with probability $p(k) = 0$), the corresponding cumulative distribution function P (just like the cumulative histogram) has intervals of constant value on which no inverse function value can be determined.

In the following, we describe a simple method for histogram matching that works with discrete reference distributions. The principal idea is graphically illustrated in Fig. 4.13. The mapping function f_{hs} is not obtained by inverting but by “filling in” the reference distribution function $P_R(i)$. For each possible pixel value a , starting with $a = 0$, the corresponding probability $p_A(a)$ is stacked layer by layer “under” the reference distribution P_R . The thickness of each horizontal bar for a equals the corresponding probability $p_A(a)$. The bar for a particular intensity value a with thickness $p_A(a)$ runs from

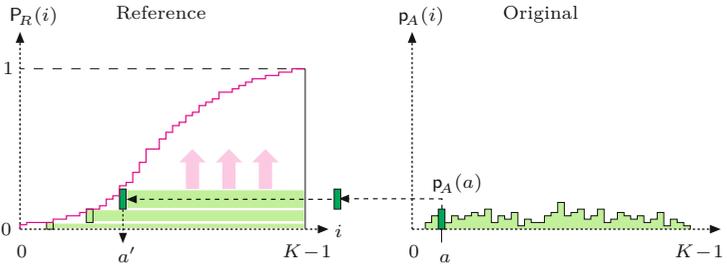


Fig. 4.13 Discrete histogram specification. The reference distribution P_R (left) is “filled” layer by layer from bottom to top and from right to left. For every possible intensity value a (starting from $a = 0$), the associated probability $p_A(a)$ is added as a horizontal bar to a stack accumulated “under” the reference distribution P_R . The bar with thickness $p_A(a)$ is drawn from right to left down to the position a' , where the reference distribution P_R is reached. The function $f_{\text{hs}}()$ must map a to a' .

right to left, down to position a' , where it hits the reference distribution P_R . This position a' corresponds to the new pixel value to which a should be mapped.

Since the sum of all probabilities p_A and the maximum of the distribution function P_R are both 1 (i.e., $\sum_i p_A(i) = \max_i P_R(i) = 1$), all horizontal bars will exactly fit underneath the function P_R . One may also notice in Fig. 4.13 that the distribution value resulting at a' is identical to the cumulated probability $P_A(a)$. Given some intensity value a , it is therefore sufficient to find the minimum value a' , where the reference distribution $P_R(a')$ is greater than or equal to the cumulative probability $P_A(a)$, that is,

$$f_{\text{hs}}(a) = \min\{j \mid (0 \leq j < K) \wedge (P_A(a) \leq P_R(j))\}. \quad (4.26)$$

This results in a very simple method, which is summarized in Alg. 4.3. The corresponding Java implementation in Prog. 4.3, consists of the method `matchHistograms()`, which accepts the original histogram (`hA`) and the reference histogram (`hR`) and returns the resulting mapping function (`fhs`) specifying the required point operation.

Due to the use of normalized distribution functions, the *size* of the associated images is not relevant. The following code fragment demonstrates the use of the `matchHistograms()` method from Prog. 4.3 in ImageJ:

```
ImageProcessor ipA = ... // target image I_A (to be modified)
ImageProcessor ipR = ... // reference image I_R

int[] hA = ipA.getHistogram(); // get histogram for I_A
int[] hR = ipR.getHistogram(); // get histogram for I_R

int[] fhs = matchHistograms(hA, hR); // mapping function f_hs(a)

ipA.applyTable(fhs); // modify the target image I_A
```

The original image `ipA` is modified in the last line by applying the mapping function f_{hs} (`fhs`) with the method `applyTable()` (see also p. 83).

4.6.5 Examples

Adjusting to a piecewise linear reference distribution

The first example in Fig. 4.14 shows the results of histogram specification for a continuous, piecewise linear reference distribution, as

Alg. 4.3

Histogram matching.

Given are two histograms: the histogram h_A of the target image I_A and a reference histogram h_R , both of size K . The result is a discrete mapping function $f_{hs}()$ that, when applied to the target image, produces a new image with a distribution function similar to the reference histogram.

```

1: MatchHistograms( $h_A, h_R$ )
   Input:  $h_A$ , histogram of the target image  $I_A$ ;  $h_R$ , reference histogram (the same size as  $h_A$ ). Returns a discrete mapping  $f_{hs}(a)$  to be applied to the target image  $I_A$ .
2:  $K \leftarrow \text{Size}(h_A)$ 
3:  $P_A \leftarrow \text{CDF}(h_A)$  ▷ c.d.f. for  $h_A$  (Alg. 4.1)
4:  $P_R \leftarrow \text{CDF}(h_R)$  ▷ c.d.f. for  $h_R$  (Alg. 4.1)
5: Create map  $f_{hs}: [0, K-1] \mapsto \mathbb{R}$  ▷ pixel mapping function  $f_{hs}$ 
6: for  $a \leftarrow 0, \dots, K-1$  do
7:    $j \leftarrow K-1$ 
8:   repeat
9:      $f_{hs}[a] \leftarrow j$ 
10:     $j \leftarrow j-1$ 
11:   while  $(j \geq 0) \wedge (P_A(a) \leq P_R(j))$ 
12: return  $f_{hs}$ .

```

described in Sec. 4.6.3. Analogous to Fig. 4.12, the actual distribution function P_R (Fig. 4.14(f)) is specified as a polygonal line consisting of five control points $\langle a_k, q_k \rangle$ with coordinates

$k =$	0	1	2	3	4	5	
$a_k =$	0	28	75	150	210	255	.
$q_k =$	0.002	0.050	0.250	0.750	0.950	1.000	

The resulting reference histogram (Fig. 4.14(c)) is a step function with ranges of constant values corresponding to the linear segments of the probability density function. As expected, the *cumulative* probability function for the modified image (Fig. 4.14(h)) is quite close to the reference function in Fig. 4.14(f), while the resulting *histogram* (Fig. 4.14(e)) shows little similarity with the reference histogram (Fig. 4.14(c)). However, as discussed earlier, this is all we can expect from a homogeneous point operation.

Adjusting to an arbitrary reference histogram

The example in Fig. 4.15 demonstrates this technique using synthetic reference histograms whose shape is approximately Gaussian. In this case, the reference distribution is not given as a continuous function but specified by a discrete histogram. We thus use the method described in Sec. 4.6.4 to compute the required mapping functions.

The target image used here was chosen intentionally for its poor quality, manifested by an extremely unbalanced histogram. The histograms of the modified images thus naturally show little resemblance to a Gaussian. However, the resulting *cumulative* histograms match nicely with the integral of the corresponding Gaussians, apart from the unavoidable irregularity at the center caused by the dominant peak in the original histogram.

Adjusting to another image

The third example in Fig. 4.16 demonstrates the adjustment of two images by matching their intensity histograms. One of the images is selected as the reference image I_R (Fig. 4.16(b)) and supplies the

4.6 HISTOGRAM SPECIFICATION

Prog. 4.3

Histogram matching (Java implementation of Alg. 4.3). The method `matchHistograms()` computes the mapping function `fhs` from the target histogram `hA` and the reference histogram `hR` (see Eqn. (4.26)). The method `Cdf()` computes the cumulative distribution function (cdf) for a given histogram (Eqn. (4.17)).

```
1  int[] matchHistograms (int[] hA, int[] hR) {
2      // hA ... histogram hA of the target image IA (to be modified)
3      // hR ... reference histogram hR
4      // returns the mapping fhs() to be applied to image IA
5
6      int K = hA.length;
7      double[] PA = Cdf(hA);           // get CDF of histogram hA
8      double[] PR = Cdf(hR);           // get CDF of histogram hR
9      int[] fhs = new int[K];          // mapping fhs()
10
11     // compute mapping function fhs():
12     for (int a = 0; a < K; a++) {
13         int j = K - 1;
14         do {
15             fhs[a] = j;
16             j--;
17         } while (j >= 0 && PA[a] <= PR[j]);
18     }
19     return fhs;
20 }
```

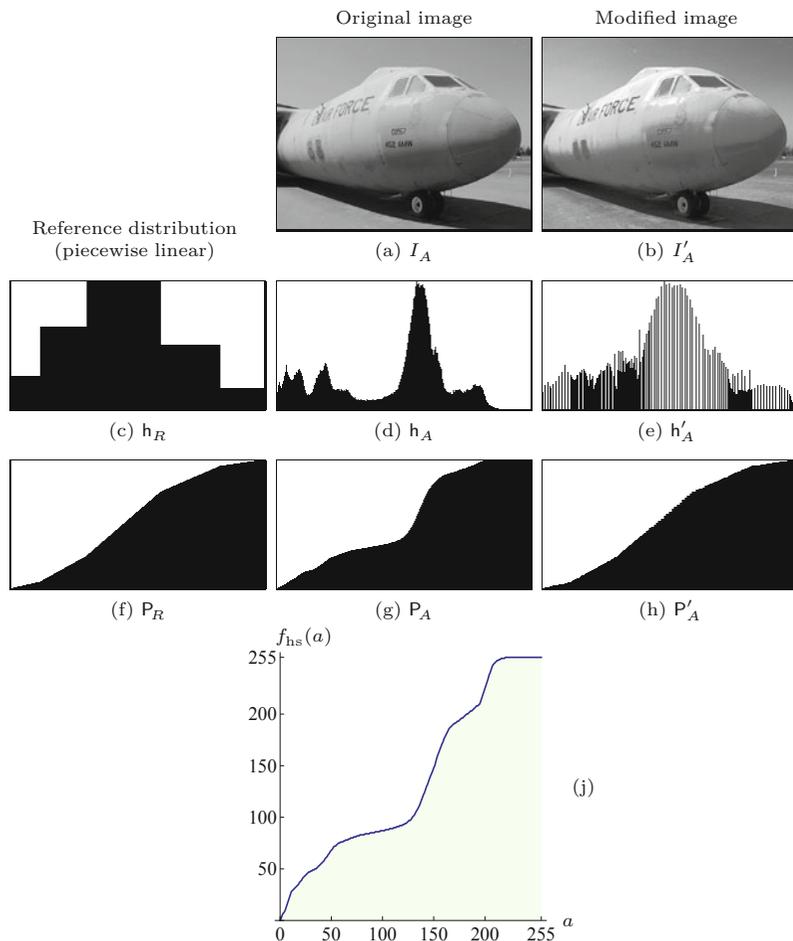
```
22 double[] Cdf (int[] h) {
23     // returns the cumul. distribution function for histogram h
24     int K = h.length;
25
26     int n = 0;                          // sum all histogram values
27     for (int i = 0; i < K; i++) {
28         n += h[i];
29     }
30
31     double[] P = new double[K];          // create CDF table P
32     int c = h[0];                        // cumulate histogram values
33     P[0] = (double) c / n;
34     for (int i = 1; i < K; i++) {
35         c += h[i];
36         P[i] = (double) c / n;
37     }
38     return P;
39 }
```

reference histogram h_R (Fig. 4.16(e)). The second (target) image I_A (Fig. 4.16(a)) is modified such that the resulting cumulative histogram matches the cumulative histogram of the reference image I_R . It can be expected that the final image $I_{A'}$ (Fig. 4.16(c)) and the reference image give a similar visual impression with regard to tonal range and distribution (assuming that both images show similar content).

Of course this method may be used to adjust multiple images to the same reference image (e.g., to prepare a series of similar photographs for a print project). For this purpose, one could either select a single representative image as a common reference or, alternatively, compute an “average” reference histogram from a set of typical images (see also Exercise 4.7).

Fig. 4.14

Histogram specification with a piecewise linear reference distribution. The target image I_A (a), its histogram P_A (g), and distribution function P_A (g); the reference histogram h_R (c) and the corresponding distribution P_R (f); the modified image $I_{A'}$ (b), its histogram $h_{A'}$ (e), and the resulting distribution $P_{A'}$ (h). Associated mapping function f_{hs} (j).



4.7 Gamma Correction

We have been using the terms “intensity” and “brightness” many times without really bothering with how the numeric pixel values in our images relate to these physical concepts, if at all. A pixel value may represent the amount of light falling onto a sensor element in a camera, the photographic density of film, the amount of light to be emitted by a monitor, the number of toner particles to be deposited by a printer, or any other relevant physical magnitude. In practice, the relationship between a pixel value and the corresponding physical quantity is usually complex and almost always nonlinear. In many imaging applications, it is important to know this relationship, at least approximately, to achieve consistent and reproducible results.

When applied to digital intensity images, the ideal is to have some kind of “calibrated intensity space” that optimally matches the human perception of intensity and requires a minimum number of bits to represent the required intensity range. Gamma correction denotes a simple point operation to compensate for the transfer characteristics of different input and output devices and to map them to a unified intensity space.

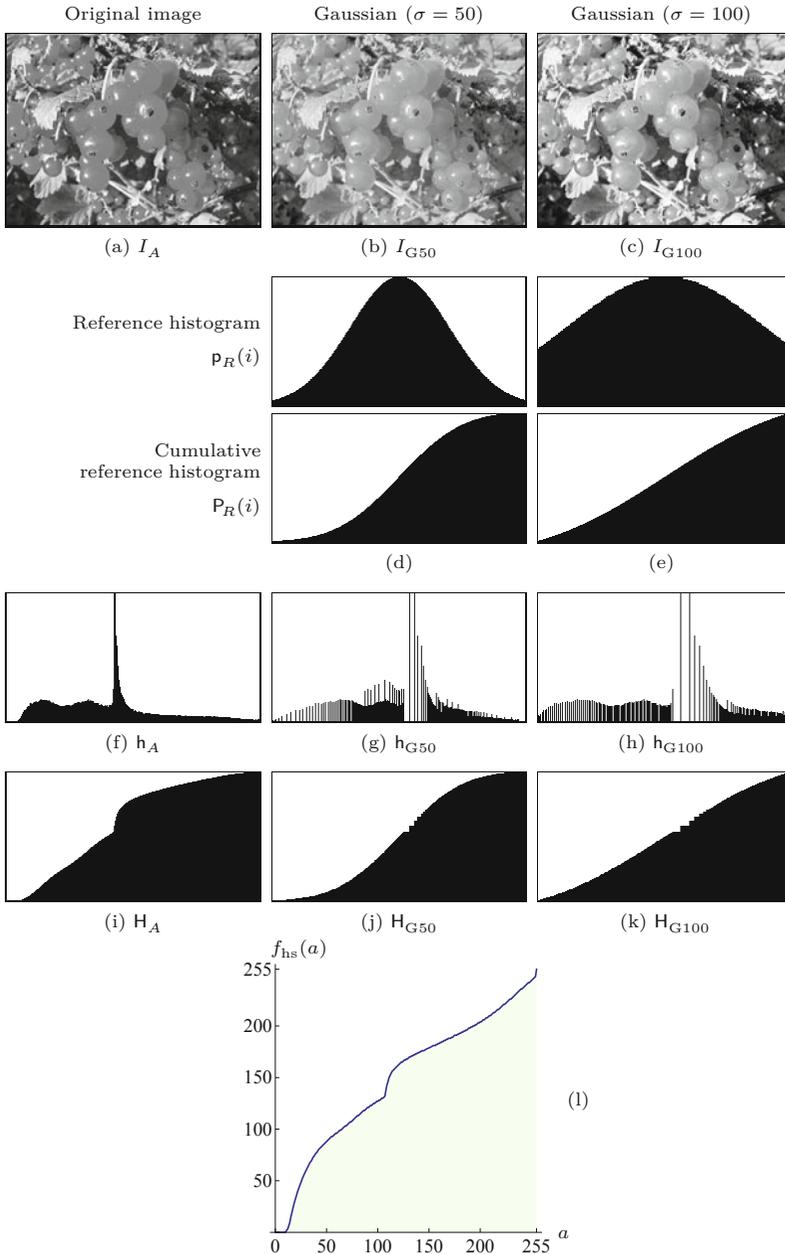


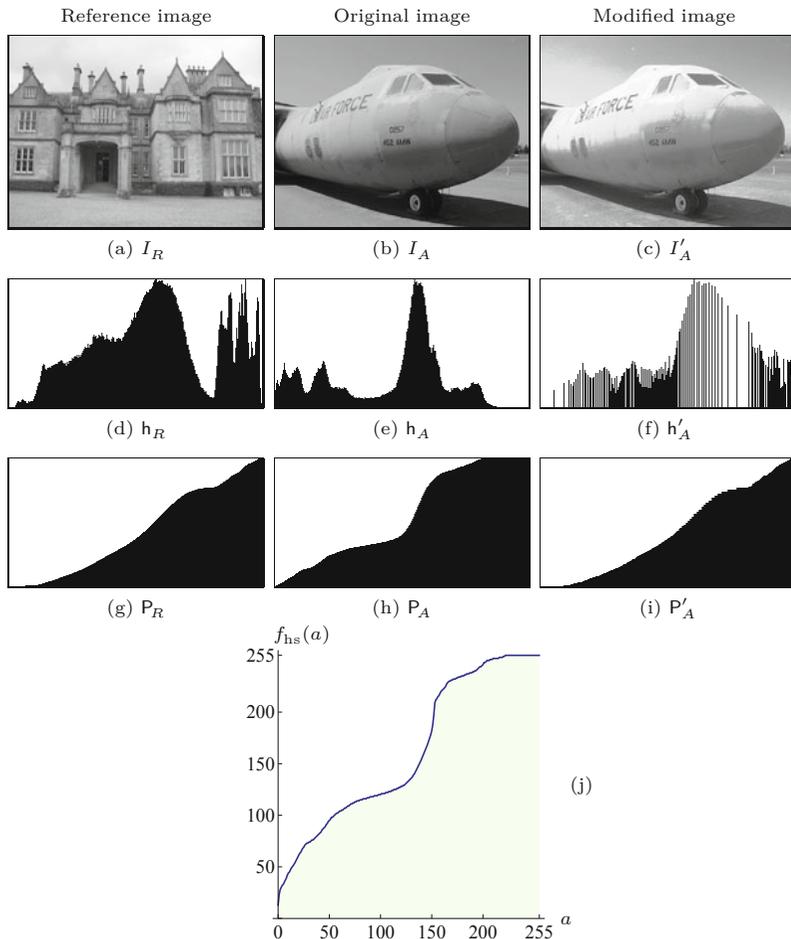
Fig. 4.15 Histogram matching: adjusting to a synthetic histogram. Original image I_A (a), corresponding histogram (f), and cumulative histogram (i). Gaussian-shaped reference histograms with center $\mu = 128$ and $\sigma = 50$ (d) and $\sigma = 100$ (e), respectively. Resulting matching, I_{G50} (b) and I_{G100} (c) with the corresponding histograms (g, h) and cumulative histograms (j, k). Associated mapping function f_{hs} (l).

4.7.1 Why Gamma?

The term “gamma” originates from analog photography, where the relationship between the light energy and the resulting film density is approximately logarithmic. The “exposure function” (Fig. 4.17), specifying the relationship between the *logarithmic* light intensity and the resulting film density, is therefore approximately *linear* over a wide range of light intensities. The slope of this function within this linear range is traditionally referred to as the “gamma” of the photographic material. The same term was adopted later in televi-

Fig. 4.16

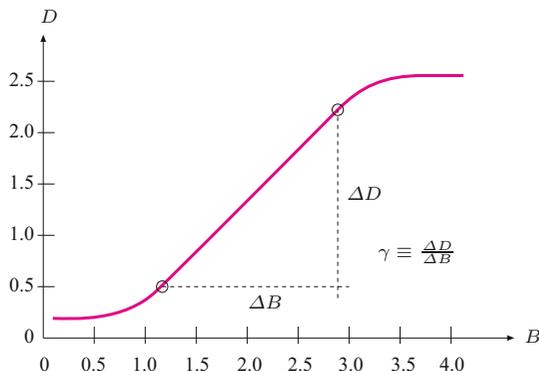
Histogram matching: adjusting to a reference image. The target image I_A (a) is modified by matching its histogram to the reference image I_R (b), resulting in the new image $I_{A'}$ (c). The corresponding histograms $h_A, h_R, h_{A'}$ (d-f) and cumulative histograms $H_A, H_R, P_{A'}$ (g-i) are shown. Notice the good agreement between the cumulative histograms of the reference and adjusted images (h, i). Associated mapping function f_{hs} (j).



sion broadcasting to describe the nonlinearities of the cathode ray tubes used in TV receivers, that is, to model the relationship between the amplitude (voltage) of the video signal and the emitted light intensity. To compensate for the nonlinearities of the receivers, a “gamma correction” was (and is) applied to the TV signal once before broadcasting in order to avoid the need for costly correction measures on the receiver side.

Fig. 4.17

Exposure function of photographic film. With respect to the logarithmic light intensity B , the resulting film density D is approximately linear over a wide intensity range. The slope $(\Delta D/\Delta B)$ of this linear section of the function specifies the “gamma” (γ) value for a particular type of photographic material.



Gamma correction is based on the exponential function

$$f_\gamma(a) = a^\gamma, \quad (4.27)$$

where the parameter $\gamma \in \mathbb{R}$ is called the *gamma* value. If a is constrained to the interval $[0, 1]$, then—independent of γ —the value of $f_\gamma(a)$ also stays within $[0, 1]$, and the function always runs through the points $(0, 0)$ and $(1, 1)$. In particular, $f_\gamma(a)$ is the identity function for $\gamma = 1$, as shown in Fig. 4.18. The function runs *above* the diagonal for gamma values $\gamma < 1$, and *below* it for $\gamma > 1$. Controlled by a single continuous parameter (γ), the power function can thus “imitate” both logarithmic and exponential types of functions. Within the interval $[0, 1]$, the function is continuous and strictly monotonic, and also very simple to invert as

$$a = f_\gamma^{-1}(b) = b^{1/\gamma}, \quad (4.28)$$

since $b^{1/\gamma} = (a^\gamma)^{1/\gamma} = a^1 = a$. The inverse of the exponential function $f_\gamma^{-1}(b)$ is thus again an exponential function,

$$f_\gamma^{-1}(b) = f_{\bar{\gamma}}(b) = f_{1/\gamma}(b), \quad (4.29)$$

with the parameter $\bar{\gamma} = 1/\gamma$.

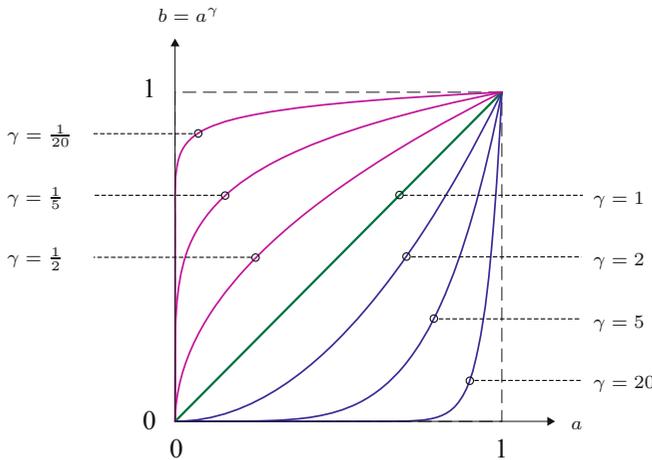


Fig. 4.18
Gamma correction function
 $f_\gamma(a) = a^\gamma$ for $a \in [0, 1]$ and
different gamma values.

4.7.3 Real Gamma Values

The actual gamma values of individual devices are usually specified by the manufacturers based on real measurements. For example, common gamma values for CRT monitors are in the range 1.8 to 2.8, with 2.4 as a typical value. Most LCD monitors are internally adjusted to similar values. Digital video and still cameras also emulate the transfer characteristics of analog film and photographic cameras by making internal corrections to give the resulting images an accustomed “look”.

In TV receivers, gamma values are standardized with 2.2 for analog NTSC and 2.8 for the PAL system (these values are theoretical; results of actual measurements are around 2.35). A gamma value of $1/2.2 \approx 0.45$ is the norm for cameras in NTSC as well as the EBU⁷ standards. The current international standard ITU-R BT.709⁸ calls for uniform gamma values of 2.5 in receivers and $1/1.956 \approx 0.51$ for cameras [76, 122]. The ITU 709 standard is based on a slightly modified version of the gamma correction (see Sec. 4.7.6).

Computers usually allow adjustment of the gamma value applied to the video output signals to adapt to a wide range of different monitors. Note, however, that the power function $f_\gamma()$ is only a coarse approximation to the actual transfer characteristics of any device, which may also not be the same for different color channels. Thus significant deviations may occur in practice, despite the careful choice of gamma settings. Critical applications, such as prepress or high-end photography, usually require additional calibration efforts based on exactly measured device profiles (see Sec. 14.7.4).

4.7.4 Applications of Gamma Correction

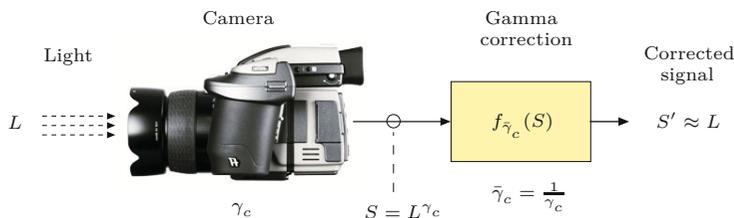
Let us first look at the simple example illustrated in Fig. 4.19. Assume that we use a digital camera with a nominal gamma value γ_c , meaning that its output signal s relates to the incident light intensity L as

$$S = L^{\gamma_c} . \quad (4.30)$$

Fig. 4.19

Principle of gamma correction.

To compensate the output signal S produced by a camera with nominal gamma value γ_c , a gamma correction is applied with $\bar{\gamma}_c = 1/\gamma_c$. The corrected signal S' is proportional to the received light intensity L .



To compensate the transfer characteristic of this camera (i.e., to obtain a measurement S' that is proportional to the original light intensity L), the camera signal S is subject to a gamma correction with the inverse of the camera's gamma value $\bar{\gamma}_c = 1/\gamma_c$ and thus

$$S' = f_{\bar{\gamma}_c}(S) = S^{1/\gamma_c} . \quad (4.31)$$

The resulting signal

$$S' = S^{1/\gamma_c} = (L^{\gamma_c})^{1/\gamma_c} = L^{(\gamma_c \frac{1}{\gamma_c})} = L^1$$

is obviously proportional (in theory even identical) to the original light intensity L . Although this example is quite simplistic, it still demonstrates the general rule, which holds for output devices as well:

⁷ European Broadcast Union (EBU).

⁸ International Telecommunications Union (ITU).

The transfer characteristic of an input or output device with specified gamma value γ is compensated for by a gamma correction with $\bar{\gamma} = 1/\gamma$.

In the aforementioned, we have implicitly assumed that all values are strictly in the range $[0, 1]$, which usually is not the case in practice. When working with digital images, we have to deal with discrete pixel values, for example, in the range $[0, 255]$ for 8-bit images. In general, performing a gamma correction

$$b \leftarrow f_{gc}(a, \gamma),$$

on a pixel value $a \in [0, a_{\max}]$ and a gamma value $\gamma > 0$ requires the following three steps:

1. Scale a linearly to $\hat{a} \in [0, 1]$.
2. Apply the gamma correction function to \hat{a} : $\hat{b} \leftarrow \hat{a}^\gamma$.
3. Scale $\hat{b} \in [0, 1]$ linearly back to $b \in [0, a_{\max}]$.

Formulated in a more compact way, the corrected pixel value b is obtained from the original value a as

$$b \leftarrow \left(\frac{a}{a_{\max}} \right)^\gamma \cdot a_{\max}. \quad (4.32)$$

Figure 4.20 illustrates the typical role of gamma correction in the digital work flow with two input (camera, scanner) and two output devices (monitor, printer), each with its individual gamma value. The central idea is to correct all images to be processed and stored in a device-independent, standardized intensity space.

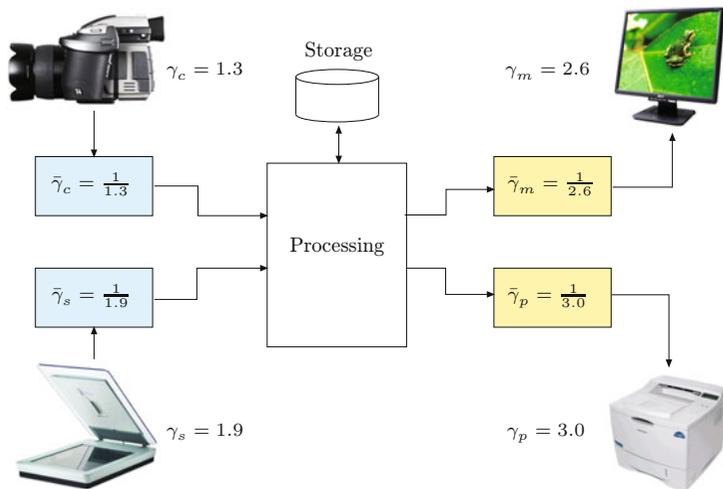


Fig. 4.20 Gamma correction in the digital imaging work flow. Images are processed and stored in a “linear” intensity space, where gamma correction is used to compensate for the transfer characteristic of each input and output device. (The gamma values shown are examples only.)

4.7.5 Implementation

Program 4.4 shows the implementation of gamma correction as an ImageJ plugin for 8-bit grayscale images. The mapping function $f_{gc}(a, \gamma)$ is computed as a lookup table (`Fgc`), which is then applied to the image using the method `applyTable()` to perform the actual point operation (see also Sec. 4.8.1).

Prog. 4.4

Implementation of gamma correction in the `run()` method of an ImageJ plugin. The corrected intensity values `b` are only computed once and stored in the lookup table `Fgc` (line 15). The gamma value `GAMMA` is constant. The actual point operation is performed by calling the ImageJ method `applyTable(Fgc)` on the image object `ip` (line 18).

```

1  public void run(ImageProcessor ip) {
2      // works for 8-bit images only
3      int K = 256;
4      int aMax = K - 1;
5      double GAMMA = 2.8;
6
7      // create and fill the lookup table:
8      int[] Fgc = new int[K];
9
10     for (int a = 0; a < K; a++) {
11         double aa = (double) a / aMax;      // scale to [0, 1]
12         double bb = Math.pow(aa, GAMMA);    // power function
13         // scale back to [0, 255]:
14         int b = (int) Math.round(bb * aMax);
15         Fgc[a] = b;
16     }
17
18     ip.applyTable(Fgc); // modify the image
19 }

```

4.7.6 Modified Gamma Correction

A subtle problem with the simple power function $f_\gamma(a) = a^\gamma$ (Eqn. (4.27)) appears if we take a closer look at the *slope* of this function, expressed by its first derivative,

$$f'_\gamma(a) = \gamma \cdot a^{(\gamma-1)},$$

which for $a = 0$ has the values

$$f'_\gamma(0) = \begin{cases} 0 & \text{for } \gamma > 1, \\ 1 & \text{for } \gamma = 1, \\ \infty & \text{for } \gamma < 1. \end{cases} \quad (4.33)$$

The tangent to the function at the origin is thus horizontal ($\gamma > 1$), diagonal ($\gamma = 1$), or vertical ($\gamma < 1$), with no intermediate values. For $\gamma < 1$, this causes extremely high amplification of small intensity values and thus increased noise in dark image regions. Theoretically, this also means that the power function is generally not invertible at the origin.

A common solution to this problem is to replace the lower part ($0 \leq a \leq a_0$) of the power function by a *linear* segment with constant slope and to continue with the ordinary power function for $a > a_0$. The resulting modified gamma correction function,

$$\bar{f}_{\gamma, a_0}(a) = \begin{cases} s \cdot a & \text{for } 0 \leq a \leq a_0, \\ (1+d) \cdot a^\gamma - d & \text{for } a_0 < a \leq 1, \end{cases} \quad (4.34)$$

$$\text{with } s = \frac{\gamma}{a_0(\gamma-1) + a_0^{(1-\gamma)}} \quad \text{and} \quad d = \frac{1}{a_0^\gamma(\gamma-1) + 1} - 1 \quad (4.35)$$

thus consists of a *linear* section (for $0 \leq a \leq a_0$) and a *nonlinear* section (for $a_0 < a \leq 1$) that connect smoothly at the transition point

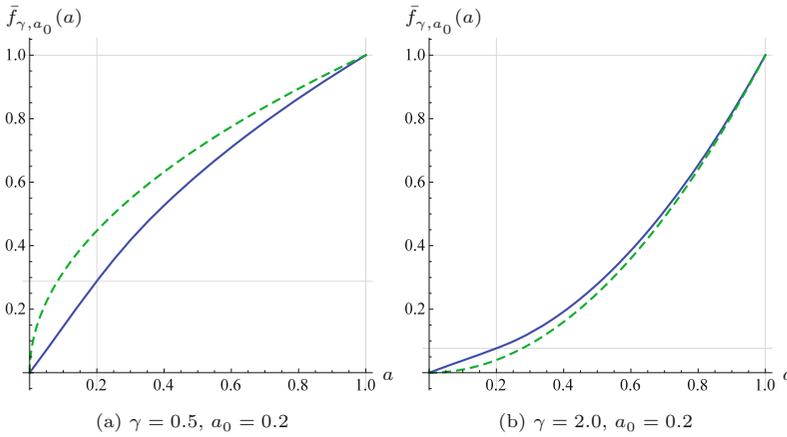


Fig. 4.21
Modified gamma correction. The mapping $\bar{f}_{\gamma, a_0}(a)$ consists of a linear segment with fixed slope s between $a = 0$ and $a = a_0$, followed by a power function with parameter γ (Eqn. (4.34)). The dashed lines show the ordinary power functions for the same gamma values.

$a = a_0$. The linear slope s and the parameter d are determined by the requirement that the two function segments must have identical values as well as identical slopes (first derivatives) at $a = a_0$ to produce a continuous function. The function in Eqn. (4.34) is thus fully specified by the two parameters a_0 and γ .

Figure 4.21 shows two examples of the modified gamma correction $\bar{f}_{\gamma, a_0}(a)$ with values $\gamma = 0.5$ and $\gamma = 2.0$, respectively. In both cases, the transition point is at $a_0 = 0.2$. For comparison, the figure also shows the ordinary gamma correction $f_\gamma(a)$ for the same gamma values (dashed lines), whose slope at the origin is ∞ (Fig. 4.21(a)) and zero (Fig. 4.21(b)), respectively.

Gamma correction in common standards

The modified gamma correction is part of several modern imaging standards. In practice, however, the values of a_0 are considerably smaller than the ones used for the illustrative examples in Fig. 4.21, and γ is chosen to obtain a good overall match to the desired correction function. For example, the ITU-BT.709 specification [122] mentioned in Sec. 4.7.3 specifies the parameters

$$\gamma = \frac{1}{2.222} \approx 0.45 \quad \text{and} \quad a_0 = 0.018, \quad (4.36)$$

with the corresponding slope and offset values $s = 4.50681$ and $d = 0.0991499$, respectively (Eqn. (4.35)). The resulting correction function $\bar{f}_{\text{ITU}}(a)$ has a *nominal* gamma value of 0.45, which corresponds to the *effective* gamma value $\gamma_{\text{eff}} = 1/1.956 \approx 0.511$. The gamma correction in the sRGB standard [224] is specified on the same basis (with different parameters; see Sec. 14.4).

Figure 4.22 shows the actual correction functions for the ITU and sRGB standards, respectively, each in comparison with the equivalent ordinary gamma correction. The ITU function (Fig. 4.22(a)) with $\gamma = 0.45$ and $a_0 = 0.018$ corresponds to an ordinary gamma correction with effective gamma value $\gamma_{\text{eff}} = 0.511$ (dashed line). The curves for sRGB (Fig. 4.22(b)) differ only by the parameters γ and a_0 , as summarized in Table 4.1.

4 POINT OPERATIONS

Fig. 4.22

Gamma correction functions specified by the ITU-R BT.709 (a) and sRGB (b) standards. The continuous plot shows the modified gamma correction with the nominal γ values and transition points a_0 .

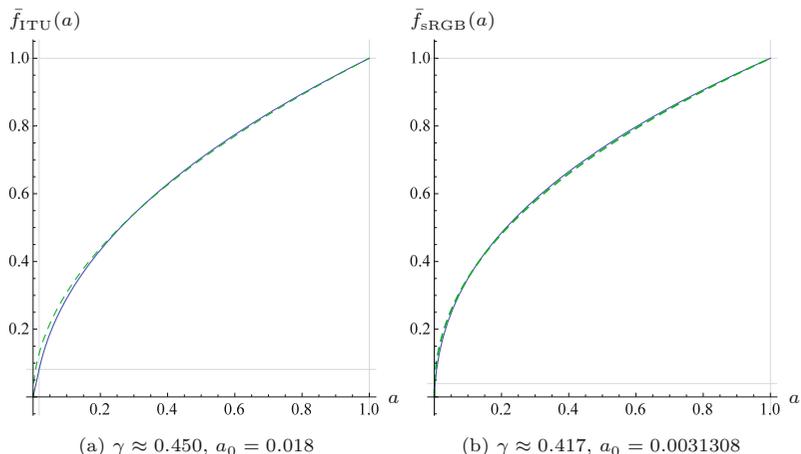


Table 4.1

Gamma correction parameters for the ITU and sRGB standards based on the modified mapping in Eqns. (4.34) and (4.35).

Standard	Nominal gamma value	a_0	s	d	Effective gamma value
	γ				γ_{eff}
ITU-R BT.709	$1/2.222 \approx 0.450$	0.018	4.50	0.099	$1/1.956 \approx 0.511$
sRGB	$1/2.400 \approx 0.417$	0.0031308	12.92	0.055	$1/2.200 \approx 0.455$

Inverting the modified gamma correction

To invert the modified gamma correction of the form $b = \bar{f}_{\gamma, a_0}(a)$ (Eqn. (4.34)), we need the inverse of the function $\bar{f}_{\gamma, a_0}(\cdot)$, which is again defined in two parts,

$$\bar{f}_{\gamma, a_0}^{-1}(b) = \begin{cases} b/s & \text{for } 0 \leq b \leq s \cdot a_0, \\ \left(\frac{b+d}{1+d}\right)^{1/\gamma} & \text{for } s \cdot a_0 < b \leq 1. \end{cases} \quad (4.37)$$

s and d are the quantities defined in Eqn. (4.35) and thus

$$a = \bar{f}_{\gamma, a_0}^{-1}(\bar{f}_{\gamma, a_0}(a)) \quad \text{for } a \in [0, 1], \quad (4.38)$$

with the *same* value γ being used in both functions. The inverse gamma correction function is required in particular for transforming between different color spaces if nonlinear (i.e., gamma-corrected) component values are involved (see also Sec. 14.2).

4.8 Point Operations in ImageJ

Several important types of point operations are already implemented in ImageJ, so there is no need to program every operation manually (as shown in Prog. 4.4). In particular, it is possible in ImageJ to apply point operations efficiently by using tabulated functions, to use built-in standard functions for point operations on single images, and to apply arithmetic operations on pairs of images. These issues are described briefly in the remaining parts of this section.

4.8.1 Point Operations with Lookup Tables

Some point operations require complex computations for each pixel, and the processing of large images may be quite time-consuming. If

the point operation is *homogeneous* (i.e., independent of the pixel coordinates), the value of the mapping function can be precomputed for every possible pixel value and stored in a lookup table, which may then be applied very efficiently to the image. A lookup table \mathbf{L} represents a discrete mapping (function f) from the original to the new pixel values,

$$\mathbf{F} : [0, K-1] \xrightarrow{f} [0, K-1]. \quad (4.39)$$

For a point operation specified by a particular pixel mapping function $a' = f(a)$, the table \mathbf{L} is initialized with the values

$$\mathbf{F}[a] \leftarrow f(a), \quad \text{for } 0 \leq a < K. \quad (4.40)$$

Thus the K table elements of \mathbf{F} need only be computed once, where typically $K = 256$. Performing the actual point operation only requires a simple (and quick) table lookup in \mathbf{F} at each pixel, that is,

$$I'(u, v) \leftarrow \mathbf{F}[I(u, v)], \quad (4.41)$$

which is much more efficient than any individual function call. ImageJ provides the method

```
void applyTable(int[] F)
```

for objects of type `ImageProcessor`, which requires a lookup table F as a 1D `int` array of size K (see Prog. 4.4 on page 80 for an example). The advantage of this approach is obvious: for an 8-bit image, for example, the mapping function is evaluated only 256 times (independent of the image size) and not a million times or more as in the case of a large image. The use of lookup tables for implementing point operations thus always makes sense if the number of image pixels ($M \times N$) is greater than the number of possible pixel values K (which is usually the case).

4.8.2 Arithmetic Operations

ImageJ implements a set of common arithmetic operations as methods for the class `ImageProcessor`, which are summarized in Table 4.2. In the following example, the image is multiplied by a scalar constant (1.5) to increase its contrast:

```
ImageProcessor ip = ... //some image
ip.multiply(1.5);
```

The image `ip` is destructively modified by all of these methods, with the results being limited (clamped) to the minimum and maximum pixel values, respectively.

4.8.3 Point Operations Involving Multiple Images

Point operations may involve more than one image at once, with arithmetic operations on the pixels of *pairs* of images being a special but important case. For example, we can express the pointwise *addition* of two images I_1 and I_2 (of identical size) to create a new image I' as

Table 4.2
ImageJ methods for arithmetic operations applicable to objects of type `ImageProcessor`.

<code>void abs()</code>	$I(u, v) \leftarrow I(u, v) $
<code>void add(int p)</code>	$I(u, v) \leftarrow I(u, v) + p$
<code>void gamma(double g)</code>	$I(u, v) \leftarrow (I(u, v)/255)^g \cdot 255$
<code>void invert(int p)</code>	$I(u, v) \leftarrow 255 - I(u, v)$
<code>void log()</code>	$I(u, v) \leftarrow \log_{10}(I(u, v))$
<code>void max(double s)</code>	$I(u, v) \leftarrow \max(I(u, v), s)$
<code>void min(double s)</code>	$I(u, v) \leftarrow \min(I(u, v), s)$
<code>void multiply(double s)</code>	$I(u, v) \leftarrow \text{round}(I(u, v) \cdot s)$
<code>void sqr()</code>	$I(u, v) \leftarrow I(u, v)^2$
<code>void sqrt()</code>	$I(u, v) \leftarrow \sqrt{I(u, v)}$

$$I'(u, v) \leftarrow I_1(u, v) + I_2(u, v) \quad (4.42)$$

for all positions (u, v) . In general, any function $f(a_1, a_2, \dots, a_n)$ over n pixel values a_i may be defined to perform pointwise combinations of n images, that is,

$$I'(u, v) \leftarrow f(I_1(u, v), I_2(u, v), \dots, I_n(u, v)). \quad (4.43)$$

Of course, most arithmetic operations on multiple images can also be implemented as successive binary operations on pairs of images.

4.8.4 Methods for Point Operations on Two Images

ImageJ supplies a single method for implementing arithmetic operations on pairs of images,

```
copyBits(ImageProcessor ip2, int u, int v, int mode),
```

which applies the binary operation specified by the transfer mode parameter `mode` to all pixel pairs taken from the *source image* `ip2` and the *target image* (the image on which this method is invoked) and stores the result in the target image. u, v are the coordinates where the source image is inserted into the target image (usually $u = v = 0$). The following code segment demonstrates the addition of two images:

```
ImageProcessor ip1 = ... // target image (I1)
ImageProcessor ip2 = ... // source image (I2)
...
ip1.copyBits(ip2, 0, 0, Blitter.ADD); // I1 ← I1 + I2
// ip1 holds the result, ip2 is unchanged
...
```

In this operation, the target image `ip1` is destructively modified, while the source image `ip2` remains unchanged. The constant `ADD` is one of several arithmetic transfer modes defined by the `Blitter` interface (see [Table 4.3](#)). In addition, `Blitter` defines (bitwise) logical operations, such as `OR` and `AND`. For arithmetic operations, the `copyBits()` method limits the results to the admissible range of pixel values (of the target image). Also note that (except for target images of type `FloatProcessor`) the results are *not* rounded but truncated to integer values.

ADD	$I_1(u, v) \leftarrow I_1(u, v) + I_2(u, v)$
AVERAGE	$I_1(u, v) \leftarrow (I_1(u, v) + I_2(u, v)) / 2$
COPY	$I_1(u, v) \leftarrow I_2(u, v)$
DIFFERENCE	$I_1(u, v) \leftarrow I_1(u, v) - I_2(u, v) $
DIVIDE	$I_1(u, v) \leftarrow I_1(u, v) / I_2(u, v)$
MAX	$I_1(u, v) \leftarrow \max(I_1(u, v), I_2(u, v))$
MIN	$I_1(u, v) \leftarrow \min(I_1(u, v), I_2(u, v))$
MULTIPLY	$I_1(u, v) \leftarrow I_1(u, v) \cdot I_2(u, v)$
SUBTRACT	$I_1(u, v) \leftarrow I_1(u, v) - I_2(u, v)$

Table 4.3

Arithmetic operations and corresponding transfer mode constants for `ImageProcessor`'s `copyBits()` method. Example: `ip1.copyBits(ip2, 0, 0, Blitter.ADD)`.

4.8.5 ImageJ Plugins Involving Multiple Images

ImageJ provides two types of plugin: a generic plugin (`PlugIn`), which can be run without any open image, and plugins of type `PlugInFilter`, which apply to a single image. In the latter case, the currently active image is passed as an object of type `ImageProcessor` (or any of its subclasses) to the plugin's `run()` method (see also Sec. 2.2.3).

If two or more images I_1, I_2, \dots, I_k are to be combined by a plugin program, only a single image I_1 can be passed directly to the plugin's `run()` method, but not the additional images I_2, \dots, I_k . The usual solution is to make the plugin open a dialog window to let the user select the remaining images interactively. This is demonstrated in the following example plugin for transparently blending two images.

Example: Linear blending

Linear blending is a simple method for continuously mixing two images, I_{BG} and I_{FG} . The background image I_{BG} is covered by the foreground image I_{FG} , whose transparency is controlled by the value α in the form

$$I'(u, v) = \alpha \cdot I_{BG}(u, v) + (1 - \alpha) \cdot I_{FG}(u, v), \quad (4.44)$$

with $0 \leq \alpha \leq 1$. For $\alpha = 0$, the foreground image I_{FG} is nontransparent (opaque) and thus entirely hides the background image I_{BG} . Conversely, the image I_{FG} is fully transparent for $\alpha = 1$ and only I_{BG} is visible. All α values between 0 and 1 result in a weighted sum of the corresponding pixel values taken from I_{BG} and I_{FG} (Eqn. (4.44)).

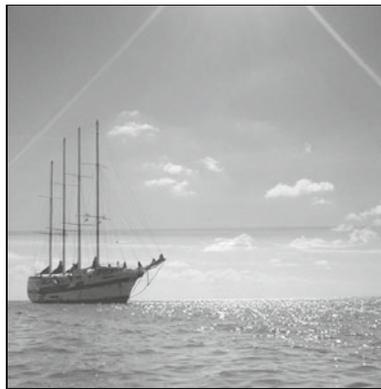
Figure 4.23 shows the results of linear blending for different α values. The Java code for the corresponding implementation (as an ImageJ plugin) is listed in Prog. 4.5. The background image (`bgIp`) is passed directly to the plugin's `run()` method. The second (foreground) image and the α value are specified interactively by creating an instance of the ImageJ class `GenericDialog`, which allows the simple implementation of dialog windows with various types of input fields.

Fig. 4.23

Linear blending example. Foreground image I_{FG} (a) and background image (I_{BG}) (e); blended images for transparency values $\alpha = 0.25, 0.50,$ and 0.75 (b–d) and dialog window (f) produced by `GenericDialog` (see Prog. 4.5).



(a) $I_{FG}, \alpha = 0.0$



(b) $\alpha = 0.25$



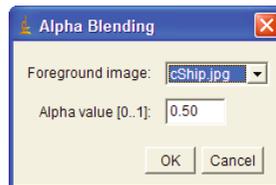
(c) $\alpha = 0.50$



(d) $\alpha = 0.75$



(e) $I_{BG}, \alpha = 1.0$



(f)

4.9 Exercises

Exercise 4.1. Implement the auto-contrast operation as defined in Eqns. (4.9)–(4.11) as an ImageJ plugin for an 8-bit grayscale image. Set the quantile p of pixels to be saturated at both ends of the intensity range (0 and 255) to $p = p_{lo} = p_{hi} = 1\%$.

Exercise 4.2. Modify the histogram equalization plugin in Prog. 4.2 to use a lookup table (Sec. 4.8.1) for computing the point operation.

Exercise 4.3. Implement the histogram equalization as defined in Eqn. (4.12), but use the *modified* cumulative histogram defined in Eqn. (4.13), cumulating the square root of the histogram entries. Compare the results to the standard (linear) approach by plotting the resulting histograms and cumulative histograms as shown in Fig. 4.10.

Exercise 4.4. Show formally that (a) a linear histogram equalization (Eqn. (4.12)) does not change an image that already has a uniform intensity distribution and (b) that any repeated application of histogram equalization to the same image causes no more changes.

Exercise 4.5. Show that the linear histogram equalization (Sec. 4.5) is only a special case of histogram specification (Sec. 4.6).

Exercise 4.6. Implement the histogram specification using a piecewise linear reference distribution function, as described in Sec. 4.6.3. Define a new object class with all necessary instance variables to represent the distribution function and implement the required functions $P_L(i)$ (Eqn. (4.23)) and $P_L^{-1}(b)$ (Eqn. (4.24)) as methods of this class.

Exercise 4.7. Using a histogram specification for adjusting *multiple* images (Sec. 4.6.4), one could either use one typical image as the reference or compute an “average” reference histogram from a set of images. Implement the second approach and discuss its possible advantages (or disadvantages).

Exercise 4.8. Implement the modified gamma correction (see Eqn. (4.34)) as an ImageJ plugin with variable values for γ and a_0 using a lookup table as shown in Prog. 4.4.

Exercise 4.9. Show that the modified gamma correction function $\bar{f}_{\gamma, a_0}(a)$, with the parameters defined in Eqns. (4.34)–(4.35), is C1-continuous (i.e., both the function itself and its first derivative are continuous).

4 POINT OPERATIONS

Prog. 4.5

ImageJ-Plugin (Linear_Blending). A background image is transparently blended with a selected foreground image. The plugin is applied to the (currently active) background image, and the foreground image must also be open when the plugin is started. The background image (bgIp), which is passed to the plugin's run() method, is multiplied with α (line 22). The foreground image (fgIP, selected in part 2) is first duplicated (line 20) and then multiplied with $(1 - \alpha)$ (line 21). Thus the original foreground image is not modified.

The final result is obtained by adding the two weighted images (line 23). To select the foreground image, a list of currently open images and image titles is obtained (lines 30–32). Then a dialog object (of type `GenericDialog`) is created and opened for specifying the foreground image (fgIm) and the α value (lines 36–46).

```
1 import ij.ImagePlus;
2 import ij.gui.GenericDialog;
3 import ij.plugin.filter.PlugInFilter;
4 import ij.process.Blitter;
5 import ij.process.ImageProcessor;
6 import imagingbook.lib.ij.IjUtils;
7
8 public class Linear_Blending implements PlugInFilter {
9     static double alpha = 0.5; // transparency of foreground image
10    ImagePlus fgIm;           // foreground image (to be selected)
11
12    public int setup(String arg, ImagePlus im) {
13        return DOES_8G;
14    }
15
16    public void run(ImageProcessor ipBG) { // ipBG = I_BG
17        if(runDialog()) {
18            ImageProcessor ipFG = // ipFG = I_FG
19                fgIm.getProcessor().convertToByte(false);
20            ipFG = ipFG.duplicate();
21            ipFG.multiply(1 - alpha); // I_FG ← I_FG · (1 - α)
22            ipBG.multiply(alpha); // I_BG ← I_BG · α
23            ipBG.copyBits(ipFG,0,0,Blitter.ADD); // I_BG ← I_BG + I_FG
24        }
25    }
26
27    boolean runDialog() {
28        // get list of open images and their titles:
29        ImagePlus[] openImages = IjUtils.getOpenImages(true);
30        String[] imageTitles = new String[openImages.length];
31        for (int i = 0; i < openImages.length; i++) {
32            imageTitles[i] = openImages[i].getShortTitle();
33        }
34        // create the dialog and show:
35        GenericDialog gd =
36            new GenericDialog("Linear Blending");
37        gd.addChoice("Foreground image:",
38            imageTitles, imageTitles[0]);
39        gd.addNumericField("Alpha value [0..1]:", alpha, 2);
40        gd.showDialog();
41
42        if (gd.wasCanceled())
43            return false;
44        else {
45            fgIm = openImages[gd.getNextChoiceIndex()];
46            alpha = gd.getNextNumber();
47            return true;
48        }
49    }
50 }
```