





The Art of
Image Processing
with Java



The Art of Image Processing with Java

Kenny A. Hunt



A K Peters, Ltd.
Natick, Massachusetts

Editorial, Sales, and Customer Service Office

A K Peters, Ltd.
5 Commonwealth Road, Suite 2C
Natick, MA 01760
www.akpeters.com

Copyright © 2010 by A K Peters, Ltd.

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the copyright owner.

Library of Congress Cataloging-in-Publication Data

Hunt, Kenny A.

The art of image processing with Java / Kenny A. Hunt.

p. cm.

Includes bibliographical references and index.

ISBN 978-1-56881-717-0 (alk. paper)

1. Image processing—Digital techniques. 2. Java (Computer program language)

I. Title.

TA1637.H87 2010

621.36'702855133—dc22

2010027302

Printed in India

12 11 10 09 08

10 9 8 7 6 5 4 3 2 1







Contents

1	Introduction	1
1.1	What Is Digital Image Processing?	1
1.2	Why Digital Image Processing?	5
2	Optics and Human Vision	13
2.1	Light	13
2.2	Camera Optics	15
2.3	Human Visual System	17
2.4	Exercises	23
3	Digital Images	25
3.1	Introduction	25
3.2	Color	25
3.3	Digital Images	37
3.4	Acquisition	43
3.5	Display	48
3.6	Exercises	48
4	Digital Images in Java	51
4.1	Overview	51
4.2	Image Structure	51
4.3	Java's Imaging Library	62
4.4	Exercises	77
5	Point Processing Techniques	81
5.1	Overview	81
5.2	Rescaling (Contrast and Brightness)	82
5.3	Lookup Tables	90
5.4	Gamma Correction	92
5.5	Pseudo Coloring	95
5.6	Histogram Equalization	99
5.7	Arithmetic Image Operations	107

5.8	Logical Image Operations	113
5.9	Alpha Blending	115
5.10	Other Blending Modes	116
5.11	Exercises	116
6	Regional Processing Techniques	123
6.1	Overview	123
6.2	Convolution	124
6.3	Smoothing	139
6.4	Edges	144
6.5	Edge Enhancement	155
6.6	Rank Filters	157
6.7	Template Matching and Correlation	166
6.8	Exercises	168
7	Geometric Operations	173
7.1	Affine Transformations	174
7.2	Custom Implementation	186
7.3	Nonlinear Transformations	191
7.4	Exercises	194
8	Image Printing and Display	197
8.1	Halftoning	197
8.2	Thresholding	199
8.3	Patterning	202
8.4	Random Dithering	204
8.5	Dithering Matrices	205
8.6	Error Diffusion Dithering	207
8.7	Color Dithering	212
8.8	Exercises	215
9	Frequency Domain	217
9.1	Overview	217
9.2	Image Frequency	217
9.3	Discrete Cosine Transform	222
9.4	Discrete Fourier Transform	228
9.5	Wavelets	249
9.6	Exercises	251
10	Image Compression	255
10.1	Overview	255
10.2	Run Length Coding	259
10.3	Hierarchical Coding	268

10.4 Predictive Coding	275
10.5 JPEG Case Study	280
10.6 GIF Case Study	284
10.7 Digital Data Embedding	286
10.8 Exercises	293
11 Morphological Image Processing	297
11.1 Components	298
11.2 Component Labeling	300
11.3 Dilation and Erosion	306
11.4 Opening and Closing	309
11.5 Component Representation	310
11.6 Component Features	313
11.7 Image Segmentation	317
11.8 Exercises	319
12 Advanced Programming	323
12.1 Overview	323
12.2 Concurrency	323
12.3 JAI	329
A Floating Point Rasters	333
B Scanners	335
References	339
References	339
Index	341



Introduction 1

1.1 What Is Digital Image Processing?

We must begin our journey by taking issue with the philosophical adage that “a picture is worth a thousand words.” It is my belief that a picture cannot begin to convey the depth of human experience and wisdom embedded in the words of Shakespeare, Dostoevsky, Dante, or Moses. A picture cannot convey with due precision the mathematical underpinnings of the discoveries of Galileo or Pascal nor can a picture give expression to the philosophy of Augustine, Plato, or Edwards. Nonetheless, while pictures do not carry the precision of written language, they do contain a wealth of information and have been used throughout the centuries as an important and useful means of communication. An *image* is a picture representing visual information. A *digital image* is an image that can be stored in digital form.

Prior to the advent of computation, images were rendered on papyrus, paper, film, or canvas using ink or paint or photosensitive chemicals. The non-digital images are prone to fading and hence suffer loss of image quality due to exposure to light or temperature extremes. Also, since non-digital images are fixed in some physical medium it is not possible to precisely copy a non-digital image. Throughout the annals of art history, forgers have attempted to copy paintings of well-known masters but usually fail due to their inability to precisely duplicate either a style or an original work. Han van Meegeren is one of the best known art forgers of the 20th century. His technique so closely mimicked the style and colors of the art masters that he was able to deceive even the most expert art critics of his time. His most famous forgery, *The Disciples at Emmaus*, was created in 1936 and was purportedly created by the well-known Dutch artist Johannes Vermeer. His work was finally exposed as fraudulent, however, at least in part by a chemical analysis of the paint, which showed traces of a plastic compound that was not manufactured until the 20th century!

Digital images, however, are pictures that are stored in digital form and that are viewable on some computing system. Since digital images are stored as binary data, the digital image never fades or degrades over time and the only way

to destroy a digital image is to delete or corrupt the file itself. In addition, a digital image can be transmitted across the globe in seconds and can be efficiently and precisely copied without any loss of quality.

Digital image processing is a field of study that seeks to analyze, process, or enhance a digital image to achieve some desired outcome. More formally, digital image processing can be defined as the study of techniques for transforming a digital image into another (improved) digital image or for analyzing a digital image to obtain specific information about the image.

From the cradle to the grave we are accustomed to viewing life through *digital* images. A parent's first portrait of their child is often taken before they are even born through the use of sophisticated ultrasound imaging technology. As the child grows, the parents capture developmental milestones using palm-sized digital video cameras. Portraits are sent over email to relatives and friends and short video clips are posted on the family's website. When the child breaks an arm playing soccer, the emergency-room physician orders an x-ray image and transmits it over the Internet to a specialist hundreds of miles away for immediate advice. During his lifetime the child will watch television images that have been digitally transmitted to the dish on top of his house, view weather satellite images on the Internet to determine whether or not to travel, and see images of war where smart bombs find their target by "seeing" the enemy.

Computer graphics is a closely related field but has a different goal than image processing. While the primary goal of computer graphics is the efficient generation of digital images, the input to a graphics system is generally a geometric model specifying the shape, texture, and color of all objects in the virtual scene. Image processing, by contrast, begins with a digital image as input and generates, typically, a digital image as output.

Computer vision, or machine vision, is another increasingly important relative of image processing where an input image is analyzed in order to determine its content. The primary goal of computer vision systems is the inverse of

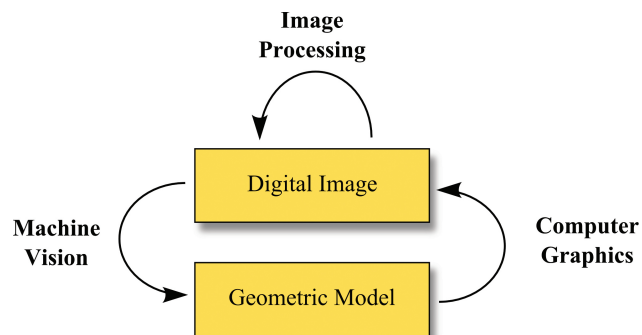


Figure 1.1. Disciplines related to image processing.




Figure 1.2. Image processing pipeline.

computer graphics: to analyze a digital image and infer meaningful information about the scene depicted by the image. Figure 1.1 illustrates the roles and relationships between each of these three disciplines where boxes represent a type of data while the connecting arrows show the typical input and output for a field of study.


A complete digital image processing system is able to service every aspect of digital image handling. Figure 1.2 shows the five typical stages in an image processing pipeline: image acquisition, image processing, image archival, image transmission, and image display. Image acquisition is the process by which digital images are obtained or generated. Image processing is the stage where a digital image is enhanced or analyzed. Image archival is concerned with how digital images are represented in memory. Image transmission is likewise concerned with data representation but places added emphasis on the robust reconstruction of potentially corrupted data due to transmission noise. Image display deals with the visual display of digital image data whether on a computer monitor, television screen, or printed page.

A visual example of the pipeline stages is given in Figure 1.3. During the image acquisition stage, an approximation of a continuous tone or analog scene is recorded. Since the captured image is an approximation, it includes some error which is introduced through sampling and quantization. During archival, a further degradation of quality may occur as the concern to conserve memory and hence conserve transmission bandwidth competes with the desire to maintain a high quality image. When the image is displayed, in this case through printing in black and white, image quality may be compromised if the output display is unable to reproduce the image with sufficient resolution or depth of color.


Construction of a complete image processing system requires specialized knowledge of how hardware architecture, the physical properties of light, the workings of the human visual system, and the structure of computational techniques affects each stage in the pipeline. Table 1.1 summarizes the most important topics of study as they correspond to each of the five primary stages in an image processing system. Of course a deep understanding of each of the listed areas of study is required to construct an efficient and effective processing module within any stage of the pipeline. Nevertheless, each stage of the processing pipeline raises unique concerns regarding memory requirements, computational efficiency, and image quality. A thorough understanding of the affects of each stage on image processing is required in order to achieve the best possible balance among memory, computation time, and image quality.




(a) Scene.



(b) Acquired.



(c) Archived.



(d) Displayed.

Figure 1.3. Effects of image processing stages on a processed image.

Processing Stage	Topic of Study
acquisition	physical properties of light human perception mathematical models of color
processing	software architecture data representation algorithm design
archival	compression techniques data representation
transmission	data representation transmission protocols
display	digital halftoning color models human perception

Table 1.1. Topics of study in image processing.

These five stages serve as a general outline for the remainder of this text. The image processing topics associated with each stage of the processing pipeline will be discussed with an emphasis on the processing stage which lies at the heart of image processing. By contrast, little coverage will be allocated to transmission issues in particular.

1.2 Why Digital Image Processing?

Digital images are used across an exceptionally wide spectrum of modern life. Ranging from digital cameras and cell phones to medical scans and web technology, digital image processing plays a central role in modern culture. This section provides examples of practical applications of image processing techniques. A general overview of these applications suffices to illustrate the importance, power, and pervasiveness of image processing techniques.

1.2.1 Medicine

Digital imaging is beginning to supplant film within the medical field. Computed tomography (CT) is a noninvasive imaging technique used to diagnose various ailments such as cancers, trauma, and musculoskeletal disorders. Magnetic resonance imaging (MRI) is a similarly noninvasive method for imaging the internal structure and function of the body. MRI scans are more amenable to diagnosing neurological and cardiovascular function than CT scans due to their greater contrast among soft tissue volumes. Figure 1.4 gives an example of both MRI and CT images where the MRI highlights contrast in the internal soft-tissue organs of a human pelvis while the CT image captures the internal skeletal structure of a human skull.

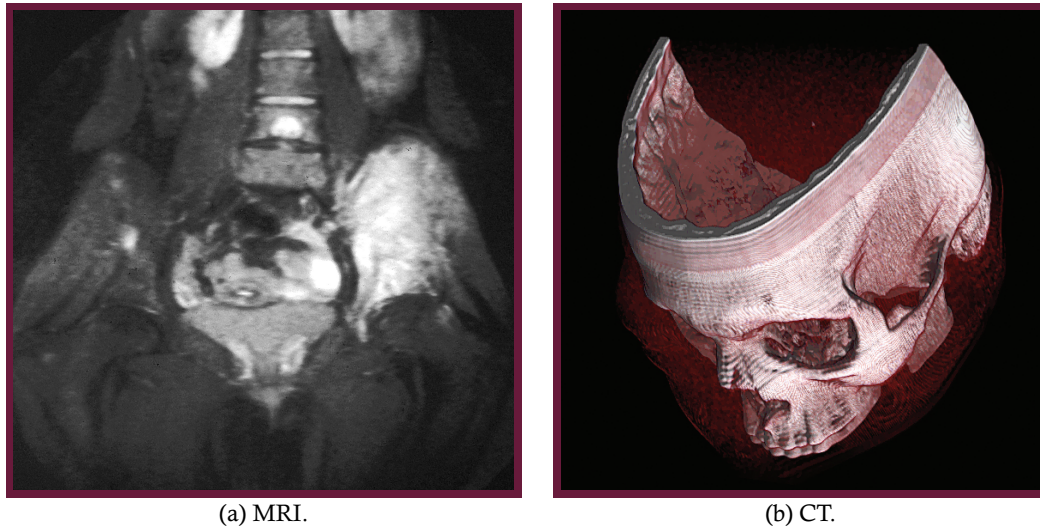


Figure 1.4. Medical images.

Since errors in the acquisition, processing, archival, or display of medical images could lead to serious health risks for patients, rigorous standards have been developed to ensure that digital images for medical use are properly archived and displayed. The Digital Imaging and Communications in Medicine (DICOM) is one such standard and has become the de facto standard for image processing in the health professions.

1.2.2 Biology

Biology is a natural science that studies living organisms and how they interact with the environment. Biological research covers a vast array of specialized subdisciplines such as botany, zoology, cell biology, microbiology, and biochemistry. Each of these disciplines relies to some degree on sophisticated computing systems to acquire and analyze large amounts of image-based data. These measurements ultimately provide information required for tasks such as deciphering complex cellular processes and identifying the structure and behavior of DNA.

Since image-based measurement is becoming increasingly vital to biological research, biologists must have basic knowledge in image processing to correctly interpret and process their results. Part (a) of Figure 1.5 shows a scanning electron microscope (SEM) image of a rust mite where the length of the mite is on the order of $60 \mu\text{m}$. Part (b) shows the structure of the eye of a fruit fly where each spherical sensor is on the order of $10 \mu\text{m}$ in diameter.

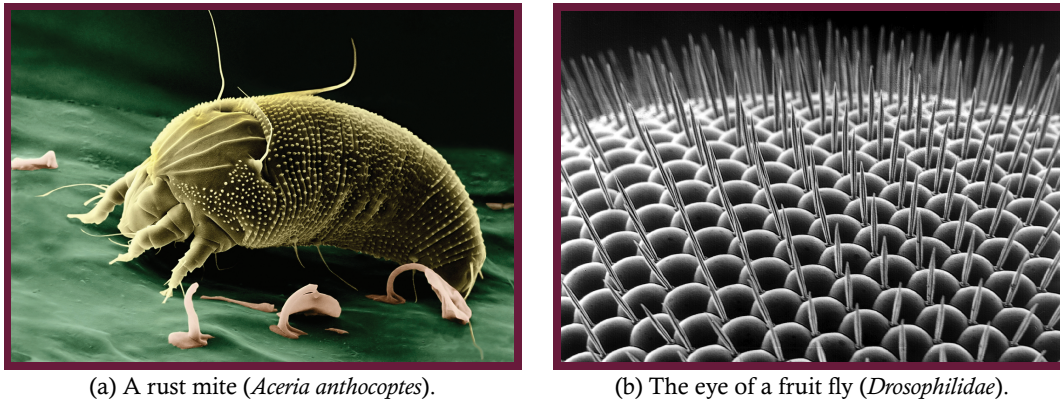
(a) A rust mite (*Aceria anthocoptes*).(b) The eye of a fruit fly (*Drosophilidae*).

Figure 1.5. Images in biology.

1.2.3 Biometrics

The security of national, corporate, and individual assets has become a topic of great importance in the modern global economy as terrorists, con men, and white-collar criminals pose an ongoing threat to society. When a person boards an airplane, enters credit card information over the Internet, or attempts to access medical records for a hospital patient, it is desirable to verify that the person actually is who they claim to be. The field of biometrics seeks to verify the identity of individuals by measuring and analyzing biological characteristics such as fingerprints, voice patterns, gait, facial appearance, or retinal scans. In most of these techniques, with the exception of voice recognition, the biological traits are obtained by the analysis of a digital image.

Biometrics has been used for decades in law enforcement to identify criminals from fingerprint images. Highly trained experts have traditionally performed fingerprint identification manually by comparing fingerprints of criminal suspects with fingerprints obtained from a crime scene. Systems are now commonly used to match fingerprints against large databases of suspects or known criminals. Specialized hardware is used to first acquire a digital image of an individual's fingerprint. Software is then used to analyze the image and compare it with a large database of known fingerprint images. Since the process is automated, it is possible to quickly search a very large database and quickly obtain accurate verification.


The use of palm scans is proving increasingly effective in the field of biometrics. A palm scanner is used to acquire an image of the blood flow through the veins of the hand in a completely non-invasive and contact-free fashion. Since the veins form a complex three-dimensional structure within a person's palm, individuals can be identified with extremely high accuracy, and forgery is extremely difficult.

1.2.4 Environmental Science


All life depends upon a healthy environment and the environmental sciences seek to understand the forces that affect our natural world. Environmental science is a broad and interdisciplinary field that includes the study of weather patterns (meteorology), oceans (oceanography), pollution as it affects life (ecology), and the study of the earth itself (the geosciences).

Data acquisition and analysis plays a key role in each of these fields since monitoring oceans, forests, farms, rivers, and even cities is critical to proper stewardship. Computer and imaging systems play an increasingly active and central role in these tasks. Satellite imaging is used to monitor and assess all types of environmental phenomena, including the effects of wildfires, hurricanes, drought, and volcanic eruptions. Motion-sensitive cameras have been installed in remote regions to monitor wildlife population densities. In recent years, these systems have discovered many new species and have even taken photographs of animals long believed extinct.

Figure 1.6 shows two enhanced satellite images of St. Louis, Missouri. The image in Figure 1.6(a) was taken during the great flood of 1993 while the image in Figure 1.6(b) was taken the following year. Environmental scientists tracked and measured the extent of the flood and the effect of the flood on terrain, vegetation, and city structures through sophisticated imaging systems and software.



(a) Satellite image in 1993.



(b) Satellite image in 1994.

Figure 1.6. Satellite images of the St. Louis flood. (Image courtesy of NASA/Goddard Space Flight Center Scientific Visualization Studio.)

1.2.5 Robotics

The field of robotics has made astounding progress in recent years. Robots now appear on the shelves of commercial toy stores, in industrial manufacturing lines, and in search and rescue missions. At the heart of most intelligent robots is a set of image processing routines that is able to process images gathered by the robot's "eyes" and determine how the robots should respond to their visually perceived environment. A team of robotics experts from the University of Southern Florida was brought in to assist in the search and rescue mission during the days after the World Trade Center collapse. These robots were specifically designed to navigate through dangerous situations looking for signs of life.

1.2.6 Professional Sports

Most professional sports leagues are developing computer systems to improve either the sports telecast or to assist umpires and referees throughout the game. The US Tennis Association, for example, uses specialized image processing systems to assist in making line calls. Officials were having increased difficulty with making correct calls as skilled tennis players can now generate 150 mile-per-hour serves and 100 mile-per-hour backhands.

Major League Baseball has also installed complex image processing systems to record the trajectory of each pitch made during a baseball game. Two cameras track the motion of the ball and are able to triangulate the position to within 1/2 inch accuracy over the entire trajectory of the pitch. A third camera is used to monitor the batter and determine the strike zone by computing the batter's knee-to-chest position. While the system is not used during game play it is used to augment television broadcasts. High-performance image processing algorithms superimpose the pitch trajectory and strike zone on instant replays. This gives sports fans an objective way to decide if the pitch was a ball or a strike. Major League Baseball does use the system to rate the performance of plate umpires in calling balls and strikes. At the conclusion of each game, the plate umpire is given a CD-ROM containing the trajectories of every pitch along with a comparison between the computer and umpire calls made.

Other sports have successfully used image-processing techniques for both decision-making and aesthetic purposes. Most major networks airing National Football League games superimpose yellow "first down" markers onto the playing field. These yellow stripes are obviously not actually on the field, but are applied using real-time image processing techniques. With the decreasing cost of computational power, it is to be expected that image processing will become more prevalent in all areas of professional sports.

1.2.7 Astronomy

Astronomers have long used digital images to study deep space over much of the electromagnetic spectrum: the Compton Gamma Ray Observatory captures digital images primarily in the gamma ray spectrum; the Chandra X-Ray Observatory and the Space Infrared Telescope Facility (also known as the Spitzer Space Telescope) provide coverage of the x-ray and infrared portions of the spectrum, respectively. The most well known telescope covering the visible portion of the spectrum is the Hubble Space Telescope, which was launched in 1990. The Hubble Telescope orbits the earth with a reflector-style optics system and a mirror of 2.4 meters in diameter. The focal length is 57.6 meters and it is able to take infrared images as well as images in the visible spectrum. Of course the images are digital since they must be transmitted to ground stations for viewing and analysis. The Hubble has produced some of the most remarkable images ever taken of created order.

Figure 1.7 is an image of the Antennae galaxies. These two galaxies are located in the constellation Corvus and are in the process of collapsing into a



Figure 1.7. Hubble Space Telescope image of the Antennae galaxies. (Image courtesy of NASA, ESA, and the Hubble Heritage Team.)





single galaxy. These galaxies are approximately 45 million light years away, and scientists predict that within 400 million years the two galaxies will have merged to form a single elliptical galaxy.

1.2.8 Conclusion

Ours is an increasingly visual culture and digital imaging is pervasive across nearly all professions, disciplines, and academic fields of study. The study of digital image processing will provide a foundation for understanding how best to acquire digital images, the nature of information contained within a digital image, and how to best archive and display images for specific purposes or applications.

Artwork

Figure 1.3. *“Migrant Mother”* by Dorothea Lange (1895–1965). Dorothea Lange was born in Hoboken, New Jersey in 1895 and devoted herself to portrait photography at a young age. After apprenticing with a photographer in New York City, she moved to San Francisco and worked predominantly with the upper class. After about 13 years she developed the desire to see things from a different point of view and Lange began shooting among San Francisco’s unemployed and documenting the increasing labor unrest. She was eventually hired by the Farm Security Administration (FSA) as a photographer and photojournalist. She is best known for her work with the FSA, which put a human face on the tragedy of the Great Depression and profoundly influenced the field of photojournalism in subsequent years. She died on October 11, 1965. Her most famous portrait is entitled “Migrant Mother,” which is shown in Figure 1.3. The image is available from the United States Library of Congress’s Prints and Photographs Division using the digital ID fsa.8b29516.



5.6 Histogram Equalization

Histogram equalization is a powerful point processing enhancement technique that seeks to optimize the contrast of an image. As the name of this technique suggests, histogram equalization seeks to improve image contrast by flattening, or equalizing, the histogram of an image.

A histogram is a table that simply counts the number of times a value appears in some data set. In image processing, a histogram is a histogram of sample values. For an 8-bit image there will be 256 possible samples in the image and the histogram will simply count the number of times that each sample actually occurs in the image.

Consider, for example, an 8-bit $W \times H$ grayscale image. There are 256 distinct sample values that *could* occur in the image. The histogram of the image is a table of 256 values where the i th entry in the histogram table contains the number of times a sample of value i occurs in the image. If the image were entirely black, for example, the 0th entry in the table would contain a value of $W \times H$ (since all the image pixels are black) and all other table entries would be zero. In general, for an N -bit $W \times H$ grayscale image where the i th sample is known to occur n_i times, the histogram h is formally defined by Equation (5.10):

$$h(i) = n_i \quad i \in 0, 1, \dots, 2^N \quad (5.10)$$

Histograms are typically normalized such that the histogram values sum to 1. In Equation (5.10) the histogram is not normalized since the sum of the histogram values is WH . The normalized histogram is given in Equation (5.11), where $\hat{h}(i)$ represents the probability that a randomly selected sample of the image that will have a value of i :

$$\hat{h}(i) = h(i)/(WH) = n_i/(WH), \quad i \in 0, 1, \dots, 2^N. \quad (5.11)$$

A histogram is typically plotted as a bar chart where the horizontal axis corresponds to the dynamic range of the image and the height of each bar corresponds to the sample count or the probability. Generally, the overall shape of a histogram doesn't convey much useful information but there are several key insights that can be gained. The spread of the histogram relates directly to image contrast where narrow histogram distributions are representative of low contrast images while wide distributions are representative of higher contrast images. Generally, the histogram of an underexposed image will have a relatively narrow distribution with a peak that is significantly shifted to the left while the histogram of an overexposed image will have a relatively narrow distribution with a peak that is significantly shifted to the right.

Figure 5.8 shows a relatively dark grayscale image and its corresponding histogram. The histogram is shifted to the left and has a relatively narrow distribution since most of the samples fall within the narrow range of approximately 25




Figure 5.8. An example histogram: (a) an 8-bit grayscale image and (b) its histogram.


to 80, while relatively few of the images samples are brighter than 128. The histogram for this example is indicative of an underexposed image—one that *may* be improved through histogram equalization.

Histogram equalization is a way of improving the local contrast of an image without altering the global contrast to a significant degree. This method is especially useful in images having large regions of similar tone such as an image with a very light background and dark foreground. Histogram equalization can expose hidden details in an image by stretching out the contrast of local regions and hence making the differences in the region more pronounced and visible.

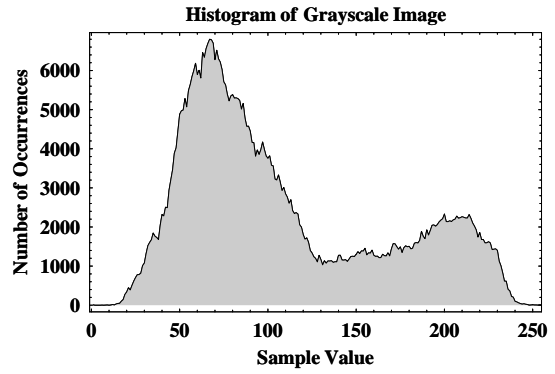
Equalization is a nonlinear point processing technique that attempts to map the input samples to output samples in such a way that there are equal amounts of each sample in the output. Since equalization is a point processing technique it is typically implemented through the use of a lookup table. For a source image, equalization computes the histogram of the source and then constructs a discrete cumulative distribution function (CDF) which is used as the lookup table. Given an N -bit image having histogram h , the normalized CDF \hat{C} is defined in Equation (5.12):

$$\hat{C}_j = \sum_{i=0}^j \hat{h}_i, \quad j \in \{0, 1, \dots, 255\}. \quad (5.12)$$


The cumulative distribution function essentially answers the question, “What percentage of the samples in an image are equal to or less than value J in the image?” Since the resulting CDF lies in the interval $[0 \dots 1]$ it can’t be directly used as a lookup table since the output samples must lie in the 8-bit range of $[0 \dots 255]$. The CDF is therefore scaled to the dynamic range of the output image. For an 8-bit image, each entry of the CDF is multiplied by 255 and rounded to obtain




(a) Underexposed image.



(b) Histogram of the underexposed image.



(c) Histogram-equalized image.



(d) Histogram of the equalized image.

Figure 5.10. Histogram equalization of an underexposed image.




(a) Low-contrast image.



(b) Histogram of the low-contrast image.



(c) Histogram equalized image.



(d) Histogram of the equalized image.

Figure 5.11. Histogram equalization of a low-contrast image.



Figure 5.12. Equalizing a color image.

Histogram equalization can also be done on color images by performing the grayscale technique on each separate band of the image. Care should be taken when doing this, however, since the colors will likely be dramatically altered as a result. If the tonal distributions are different among the red, green, and blue channels of an image, for example, the lookup tables for each channel will be vastly different and equalization will alter the color patterns present in the source. Histogram equalization of a color image is best performed on the intensity channel only, which implies that the equalization should be done on the brightness band of an image using the HSB or YIQ color spaces, for example.

Figure 5.12 demonstrates how equalization may alter the color balance of a source image. In this figure, the source image has a largely reddish tint. The histogram of the red channel is shifted to the high end while the histogram of the green channel is shifted to the low. When equalization is done on each of the red, green, and blue channels independently, the color balance is dramatically altered as seen in (b). When equalization is done on the intensity channel only, the chromaticity is retained but the brightness levels are altered, as shown in (c).

Since image histograms are so useful in image processing it is advisable to develop a `Histogram` class for the purpose of creating and processing histogram data. Listing 5.9 shows such a class, where the constructor accepts a `BufferedImage` source and the band from which to extract the histogram. The constructor then determines the maximum possible sample value and allocates a table of the proper size. The table is filled in by scanning the source and counting each sample that occurs in the image.

A histogram object can convert itself to an array via the `getCounts` methods, it can produce a cumulative distribution function as an array of doubles via the `getCDF` method, and it can produce a normalized histogram via the `getNormalizedHistogram` method. The histogram class can be used to implement a `HistogramEqualizationOp`, which is left as a programming exercise.

```
1 public class Histogram {
2     private int[] counts;
3     private int totalSamples, maxPossibleSampleValue;
4
5     public Histogram(BufferedImage src, int band) {
6         maxPossibleSampleValue = getMaxSampleValue(src, band);
7         counts = new int[maxPossibleSampleValue + 1];
8         totalSamples = src.getWidth() * src.getHeight();
9
10        for (Location pt : new RasterScanner(src, false)) {
11            int sample = src.getRaster().getSample(pt.col, pt.row, band);
12            counts[sample]++;
13        }
14    }
15
16    public int getNumberOfBins() {
17        return counts.length;
18    }
19
20    public int getValue(int index){
21        return counts[index];
22    }
23
24    public int[] getCounts() {
25        int[] result = new int[counts.length];
26        System.arraycopy(counts, 0, result, 0, counts.length);
27        return result;
28    }
29
30    public double[] getCDF() {
31        double[] cdf = getNormalizedHistogram();
32        for(int i=1; i<cdf.length; i++){
33            cdf[i] = cdf[i-1] + cdf[i];
34        }
35        return cdf;
36    }
37
38    public double[] getNormalizedHistogram() {
39        double[] result = new double[counts.length];
40        for(int i=0; i<counts.length; i++){
41            result[i] = counts[i] / (double)totalSamples;
42        }
43        return result;
44    }
45
46    private int getMaxSampleValue(BufferedImage src, int band) {
47        return (int)Math.pow(2, src.getSampleModel().getSampleSize(band)) - 1;
48    }
49 }
```

Listing 5.9. Histogram class.

Histograms have thus far been presented as one-dimensional constructs that process each channel of an image independently. Color histograms, by contrast, are three dimensional constructs that divide the images color space into volumetric bins. Each bin encloses a rectangular volume of colors in the space and the corresponding bin entry is a count of the number of pixels that are enclosed by the bin. A color histogram correctly links the individual samples of a pixel together in a dependent fashion.

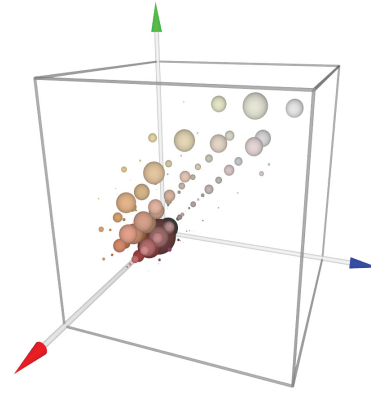
The size of the bins determines the resolution of the histogram, where smaller bins correspond to a greater resolution and larger bin sizes correspond to a lower resolution. While higher resolution histograms provide more accuracy in terms of characterizing the color distribution of an image, higher resolutions also lead to larger memory consumption and a corresponding lack of computational efficiency. Consider, for example, an 8-bit RGB color image such that each axis is divided into 256 bins. The color histogram has an impractical total of $256 \times 256 \times 256 = 16,777,216$ bins. The resolution of the histogram can be reduced by dividing each axis into only 5 bins such that the color histogram then has a computationally effective total of $5 \times 5 \times 5 = 125$ bins. Each axis of the color space may be given a resolution independently of the others such that the Y axis of the YCbCr color space may be given a higher resolution than either the Cb or Cr color spaces. Consider, for example, an 8-bit YCbCr image such that the Y axis is divided into 10 bins while each of the Cb and Cr axes are divided into 5 bins. The corresponding color histogram has a total of $10 \times 5 \times 5 = 250$ bins and has superior resolution in the intensity channel than the two chroma channels.

Figure 5.13 shows the RGB color histogram for two source images. The $12 \times 12 \times 12$ color histogram of part (a) is shown in (b) where the diameter of the spheres is directly proportional to the count of each volumetric bin and the color of each sphere corresponds to the color at the center of each bin. The $12 \times 12 \times 12$ color histogram of the source image shown in part (c) is given in (d). The large density of dark and red pixels of the source image in (a) is reflected in the distribution of data in the color histogram of (b) while the large density of light blue pixels is reflected in the distribution of the color histogram of part (d).

Since a color histogram provides a relatively coarse but concise characterization of an image, a color histogram can be used as a computationally effective means for comparing the similarity of two images. Color histograms are often used in content-based image retrieval (CBIR) systems to support efficient searches of large image databases. A CBIR system maintains a database of images that can be queried for similarity to a target image. A CBIR system allows users to find images in the database that are similar to the target image rather than searching the image database using metadata such as keywords. While two visually different images may have similar color histograms, the color histograms can be used as a coarse measure of similarity and serve as a preprocessing step in a CBIR query.




(a) Source image.



(b) RGB color histogram.



(c) Source image.



(d) RGB color histogram.

Figure 5.13. The color distribution of two source images as given by their $12 \times 12 \times 12$ RGB color histograms.

5.7 Arithmetic Image Operations

Two source images can be added, multiplied, one subtracted from the other or one divided by the other to produce a single destination. These operations are known as image arithmetic or image blending. Each of these methods is a point processing operation where corresponding samples from the two source images are combined into the destination. Image addition, for example, works by

Index

- 3CCD, 47
- 4-connected component, 299
- 8-connected component, 299

- AbstractDigitalImage (imaging), 52
- AC coefficient, 223
- acceleration, 329
- adaptive thresholding, 200
- AddBinaryOp (pixeljelly.ops), 112
- affine transform, 174
 - reflection, 175
 - rotation, 175
 - scaling, 175
 - shearing, 175
 - translation, 175
- AffineTransform (java.awt.geom), 183
- AffineTransformMapper (imaging.ops), 188
- AffineTransformOp (java.awt.image), 184
- aliasing, 220
- alpha blending, 115
- ArrayDigitalImage (imaging), 52
- aspect ratio, 42

- backward mapping, 179
- band, 38
- band pass filter, 243
- band stop filter, 243
- basis functions, 217
- Bayer filter, 47
- bicubic interpolation, 181
- bilinear interpolation, 180
- binary image, 37
- BinaryCACEncoder (pixeljelly.io), 271
- BinaryImageOp (pixeljelly.ops), 111
- bit shifting, 58

- blending mode
 - alpha blending, 115
 - darken only, 116
 - diffusion, 116
 - hard light, 116
 - lighten only, 116
 - screen, 116
- blurring, 139
- BufferedImage (java.awt.image), 62, 63
- BufferedImageOp (java.awt.image), 67
- BufferedImageOutputStream (pixeljelly.io), 290
- Butterworth filter, 240

- camera
 - aperture, 16
 - f-number, 16
 - f-stop, 16
 - optical zoom, 15
- CBIR, 106
- CCD, 44
- chain code, 311
 - absolute, 312
 - differential, 312
- channel, 38
- Chebyshev filter, 242
- circular indexing, 128
- clamping, 84
- classification, 297, 313
- closing, 309
- CMOS, 44
- coding redundancy, 256
- color
 - additive color model, 25
 - CMY color model, 28

- CMYK color model, 29
- color depth, 39
- color model, 25
- color separation, 29
- color space, 25
- gamut, 25
- HSB color model, 30
- L1 metric, 33
- L2 metric, 33
- RGB color model, 26
- similarity, 33
- subtractive color model, 26
- YCbCr color model, 32
- YIQ color model, 32
- YUV color model, 33
- Color (java.awt), 35
- color image, 37
- ColorModel, 62
- ColorModel (java.awt.image), 65
- ColorUtilities, 36
- Complex (pixeljelly.utilities), 244
- complex number, 229
 - magnitude, 230
 - phase, 230
- component, 298
 - area, 313
 - centroid, 315
 - circularity, 314
 - compactness, 314
 - eccentricity, 316
 - moments, 315
 - orientation, 316
 - perimeter, 313
- component labeling, 300
- compression ratio, 257
- computer graphics, 2
- computer vision, 2
- concurrency, 323
- ConcurrentOp (imaging.ops), 328
- cone, *see* photoreceptor
- connected path, 299
- connectivity, 298
- Constant area coding, 268
- contrast, 83
- convolution, 124
 - edge problem, 127
 - convolution theorem, 237
- ConvolutionOp (pixeljelly.ops), 137
- ConvolveOp (java.awt.image), 134
- cover image, 287
- cross correlation, 166
- CRT, 48
- cumulative distribution function, 100
- currying, 111
- DC coefficient, 223
- decoder, 256
- delta modulation, 278
- demosaicing, 47
- DICOM, 6
- digital image processing, 2
- DigitalImage (imaging), 51
- dilation, 306
- discrete cosine transform, 217, 222
- discrete Fourier transform, 217, 228
 - distributivity, 233
 - periodicity, 234
 - rotation, 233
 - translation, 233
- disjoint-set, 302
- dither matrix, 205
- dithering, 202
 - color, 212
 - error diffusion, 207
 - Floyd-Steinberg, 208
 - Jarvis-Judice-Ninke, 209
 - Sierra dithering, 209
 - Stucki, 209
- DitherMatrixOp (pixeljelly.ops), 206
- dynamic thresholding, 200
- edge, 144
 - edge detection, 144
 - edge enhancement, 155
 - edge map, 152
- encoder, 256
- energy compaction, 226
- erosion, 306
- error diffusion, 207
- ErrorDiffusionOp (pixeljelly.ops), 210
- exposure, 82
- extended padding, 128
- eye

- cornea, 18
- fovea, 18
- iris, 18
- optic nerve, 19
- perimetric angle, 18
- pupil, 18
- retina, 18
- rods and cones, 18
- visual axis, 18
- FalseColorOp (pixeljelly.ops), 98
- fast Fourier transform, 247
- feature, 313
- feature vector, 313
- FFT, 247
- flood filling, 300
- Floyd-Steinberg dithering, 208
- FloydSteinbergDitheringOp (pixeljelly.ops), 212
- forward mapping, 178
- frequency, 42
- gamma correction, 92
- GammaOp (pixeljelly.ops), 95
- Gaussian filter, 242
- geometric operations, 173
- GeometricTransformOp (imaging.ops), 188
- GIF, 284
- gradient, 145
- grayscale image, 37
- halftoning, 198
 - analog, 198
 - digital, 198
- high dynamic range imaging, 21
- high pass filter, 242
- histogram, 99
 - color, 106
- Histogram (pixeljelly.features), 104
- histogram equalization, 99
- homogeneous coordinate, 174
- homogeneous transformation matrix, 174
- ideal filter, 239
- image
 - acquisition, 43
 - continuous tone, 44
 - frequency, 42
 - quantization, 44
 - resolution, 42
 - sampling, 44
- image arithmetic, 107
- image blending, 107
- image compression, 255
- image frequency, 217
- ImageComponent (pixeljelly.gui), 76
- ImageDecoder (pixeljelly.io), 266
- ImageEncoder (pixeljelly.io), 264
- ImageIO (javax.imageio), 66
- ImagePadder (pixeljelly.utilities), 129
- ImageTiler (imaging.scanners), 324
- immediate mode, 62
- indexed image, 60
- interpixel redundancy, 255
- Interpolant (imaging.utilities), 186
- interpolation, 180
- InverseMapper (imaging.utilities), 187
- invertibility, 221
- InvertOp (pixeljelly.ops), 73
- Jarvis-Judice-Ninke dithering, 209
- Java Advanced Imaging (JAI), 329
- Java Media suite, 329
- JPEG, 280
- kernel, 124
- Kernel (java.awt.image), 134
- Kernel2D (pixeljelly.utilities), 135
- key element, 124
- least significant bit embedding, 287
- lens, 15
 - focal length, 15
 - magnification factor, 15
 - thin lens equation, 15
- light, 13
 - amplitude, 13
 - frequency, 13
 - wavelength, 13
- LinearArrayDigitalImage (imaging), 57
- Location (pixeljelly.scanners), 71
- logical operators, 113
- lookup tables, 90
- LookupOp (java.awt.image), 91

- lossless encoding, 259
- lossy encoding, 259
- low pass filtering, 238
- LSBSteganographer (pixeljelly.utilities), 292
- magic number, 266
- magnitude of gradient, 151
- MagnitudeOfGradientOp (pixeljelly.ops), 152
- Mask (pixeljelly.utilities), 159
- median filter, 157
- mezzotint, 204
- morphology, 297
- nearest neighbor interpolation, 180
- NearestNeighborInterpolant (imaging.utilities), 187
- nonlinear geometric transformations, 191
- normalized cross correlation, 167
- NullOp (pixeljelly.ops), 71
- opening, 309
- OpThread (imaging.ops), 326
- OrBinaryOp (pixeljelly.ops), 115
- packed pixel, 57
- PackedPixelImage (imaging), 60
- packing, 58
- patterning, 202
- periodic noise, 244
- photoreceptor
 - cone, 18
 - rod, 18
- pixel, 38
- pixel resolution, 42
- PluggableImageOp (pixeljelly.ops), 76
- predictive coding, 275
- Prewitt operator, 149
- process, 323
- pseudo coloring, 95
- psycho-visual redundancy, 255
- quadtree, 268
- quantization, 44
- random dithering, 204
- RandomDitherOp (pixeljelly.ops), 204
- range filter, 165
- rank filtering, 157
- RankOp (pixeljelly.ops), 162
- raster, 53
- raster scan, 56
- RasterScanner (pixeljelly.scanners), 71
- raw format, 257
- reconstruction, 182
- reconstruction filter, 48
- reflection, 175
- reflective indexing, 128
- ReflectivePadder (pixeljelly.utilities), 130
- regional processing, 81, 123
- registration, 173
- relative compression ratio, 257
- resampling, 182
- rescale, 82
 - bias, 82
 - gain, 82
- RescaleOp (java.awt.image), 90
- Roberts cross operators, 150
- rod, *see* photoreceptor
- root mean squared error, 257
- rotation, 175
- run length encoding, 259
- RunLengthEncoder (pixeljelly.io), 266
- salt and pepper noise, 157
- sample, 38
- sampling, 44
- savings ratio, 257
- scaling, 175
- scanning, 69
- segmentation, 297, 317
- separability, 133
- SeperableKernel (pixeljelly.utilities), 137
- sharpening, 155
- shearing, 175
- Sierra dithering, 209
- SimpleConvolveOp (pixeljelly.ops), 126
- single source, 67
- sinusoid, 218
- smoothing, 139
- Sobel operator, 150
- spectrum, 231
 - phase, 231
 - power, 231
- Steganographer (pixeljelly.utilities), 289

- steganography, 286
- stego image, 287
- Stucki dithering, 209
- System.arraycopy (java), 54

- template matching, 166
- thread, 323
- Thread (java), 324
- thresholding, 199
- TiledPadder (pixeljelly.utilities), 130
- translation, 175
- transparency, 115
- TwirlMapper (imaging.utilities), 193

- uniform filter, 139
- union-find, 302
- unpacking, 58

- vision
 - brightness adaptation, 20
 - glare limit, 20
 - instantaneous range, 20
 - photopic, 18

- scotopic, 18
- scotopic threshold, 20
- simultaneous contrast, 22

- wavelet, 249
 - Coiflet, 251
 - Daubechies, 251
 - Haar, 251
 - mexican hat, 251
- weighted blur filter
 - conical, 141
 - Gaussian, 141
 - pyramid, 140
- windowing, 235
 - Bartlett, 235
 - Blackman, 235
 - Hamming, 236
 - Hanning, 235
- WritableRaster (java.awt.image), 62, 65

- zero padding, 128
- ZeroPadder (pixeljelly.utilities), 129, 131