

# Image Processing: The Fundamentals

Maria Petrou

Costas Petrou



A John Wiley and Sons, Ltd., Publication

# Contents

Preface	xxiii
<b>1 Introduction</b>	<b>1</b>
Why do we process images?	1
What is an image?	1
What is a digital image?	1
What is a spectral band?	2
Why do most image processing algorithms refer to grey images, while most images we come across are colour images?	2
How is a digital image formed?	3
If a sensor corresponds to a patch in the physical world, how come we can have more than one sensor type corresponding to the same patch of the scene?	3
What is the physical meaning of the brightness of an image at a pixel position?	3
Why are images often quoted as being $512 \times 512$ , $256 \times 256$ , $128 \times 128$ etc?	6
How many bits do we need to store an image?	6
What determines the quality of an image?	7
What makes an image blurred?	7
What is meant by image resolution?	7
What does “good contrast” mean?	10
What is the purpose of image processing?	11
How do we do image processing?	11
Do we use nonlinear operators in image processing?	12
What is a linear operator?	12
How are linear operators defined?	12
What is the relationship between the point spread function of an imaging device and that of a linear operator?	12
How does a linear operator transform an image?	12
What is the meaning of the point spread function?	13
Box 1.1. The formal definition of a point source in the continuous domain	14
How can we express in practice the effect of a linear operator on an image?	18
Can we apply more than one linear operators to an image?	22
Does the order by which we apply the linear operators make any difference to the result?	22
Box 1.2. Since matrix multiplication is not commutative, how come we can change the order by which we apply shift invariant linear operators?	22

Box 1.3. What is the stacking operator? . . . . .	29
What is the implication of the separability assumption on the structure of matrix $H$ ? . . . . .	38
How can a separable transform be written in matrix form? . . . . .	39
What is the meaning of the separability assumption? . . . . .	40
Box 1.4. The formal derivation of the separable matrix equation . . . . .	41
What is the “take home” message of this chapter? . . . . .	43
What is the significance of equation (1.108) in linear image processing? . . . . .	43
What is this book about? . . . . .	44
<b>2 Image Transformations</b> . . . . .	<b>47</b>
What is this chapter about? . . . . .	47
How can we define an elementary image? . . . . .	47
What is the outer product of two vectors? . . . . .	47
How can we expand an image in terms of vector outer products? . . . . .	47
How do we choose matrices $h_c$ and $h_r$ ? . . . . .	49
What is a unitary matrix? . . . . .	50
What is the inverse of a unitary transform? . . . . .	50
How can we construct a unitary matrix? . . . . .	50
How should we choose matrices $U$ and $V$ so that $g$ can be represented by fewer bits than $f$ ? . . . . .	50
What is matrix diagonalisation? . . . . .	50
Can we diagonalise any matrix? . . . . .	50
<b>2.1 Singular value decomposition</b> . . . . .	<b>51</b>
How can we diagonalise an image? . . . . .	51
Box 2.1. Can we expand in vector outer products any image? . . . . .	54
How can we compute matrices $U$ , $V$ and $\Lambda^{\frac{1}{2}}$ needed for image diagonalisation? . . . . .	56
Box 2.2. What happens if the eigenvalues of matrix $gg^T$ are negative? . . . . .	56
What is the singular value decomposition of an image? . . . . .	60
Can we analyse an eigenimage into eigenimages? . . . . .	61
How can we approximate an image using SVD? . . . . .	62
Box 2.3. What is the intuitive explanation of SVD? . . . . .	62
What is the error of the approximation of an image by SVD? . . . . .	63
How can we minimise the error of the reconstruction? . . . . .	65
Are there any sets of elementary images in terms of which <i>any</i> image may be expanded? . . . . .	72
What is a complete and orthonormal set of functions? . . . . .	72
Are there any complete sets of orthonormal discrete valued functions? . . . . .	73
<b>2.2 Haar, Walsh and Hadamard transforms</b> . . . . .	<b>74</b>
How are the Haar functions defined? . . . . .	74
How are the Walsh functions defined? . . . . .	74
Box 2.4. Definition of Walsh functions in terms of the Rademacher functions . . . . .	74
How can we use the Haar or Walsh functions to create image bases? . . . . .	75
How can we create the image transformation matrices from the Haar and Walsh functions in practice? . . . . .	76
What do the elementary images of the Haar transform look like? . . . . .	80
Can we define an orthogonal matrix with entries only +1 or -1? . . . . .	85
Box 2.5. Ways of ordering the Walsh functions . . . . .	86
What do the basis images of the Hadamard/Walsh transform look like? . . . . .	88

What are the advantages and disadvantages of the Walsh and the Haar transforms?	92
What is the Haar wavelet?	93
<b>2.3 Discrete Fourier transform</b>	94
What is the discrete version of the Fourier transform (DFT)?	94
Box 2.6. What is the inverse discrete Fourier transform?	95
How can we write the discrete Fourier transform in a matrix form?	96
Is matrix $U$ used for DFT unitary?	99
Which are the elementary images in terms of which DFT expands an image?	101
Why is the discrete Fourier transform more commonly used than the other transforms?	105
What does the convolution theorem state?	105
Box 2.7. If a function is the convolution of two other functions, what is the relationship of its DFT with the DFTs of the two functions?	105
How can we display the discrete Fourier transform of an image?	112
What happens to the discrete Fourier transform of an image if the image is rotated?	113
What happens to the discrete Fourier transform of an image if the image is shifted?	114
What is the relationship between the average value of the image and its DFT?	118
What happens to the DFT of an image if the image is scaled?	119
Box 2.8. What is the Fast Fourier Transform?	124
What are the advantages and disadvantages of DFT?	126
Can we have a real valued DFT?	126
Can we have a purely imaginary DFT?	130
Can an image have a purely real or a purely imaginary valued DFT?	137
<b>2.4 The even symmetric discrete cosine transform (EDCT)</b>	138
What is the even symmetric discrete cosine transform?	138
Box 2.9. Derivation of the inverse 1D even discrete cosine transform	143
What is the inverse 2D even cosine transform?	145
What are the basis images in terms of which the even cosine transform expands an image?	146
<b>2.5 The odd symmetric discrete cosine transform (ODCT)</b>	149
What is the odd symmetric discrete cosine transform?	149
Box 2.10. Derivation of the inverse 1D odd discrete cosine transform	152
What is the inverse 2D odd discrete cosine transform?	154
What are the basis images in terms of which the odd discrete cosine transform expands an image?	154
<b>2.6 The even antisymmetric discrete sine transform (EDST)</b>	157
What is the even antisymmetric discrete sine transform?	157
Box 2.11. Derivation of the inverse 1D even discrete sine transform	160
What is the inverse 2D even sine transform?	162
What are the basis images in terms of which the even sine transform expands an image?	163
What happens if we do not remove the mean of the image before we compute its EDST?	166
<b>2.7 The odd antisymmetric discrete sine transform (ODST)</b>	167
What is the odd antisymmetric discrete sine transform?	167

Box 2.12. Derivation of the inverse 1D odd discrete sine transform . . . . .	171
What is the inverse 2D odd sine transform? . . . . .	172
What are the basis images in terms of which the odd sine transform expands an image? . . . . .	173
What is the “take home” message of this chapter? . . . . .	176
<b>3 Statistical Description of Images</b> . . . . .	<b>177</b>
What is this chapter about? . . . . .	177
Why do we need the statistical description of images? . . . . .	177
<b>3.1 Random fields</b> . . . . .	<b>178</b>
What is a random field? . . . . .	178
What is a random variable? . . . . .	178
What is a random experiment? . . . . .	178
How do we perform a random experiment with computers? . . . . .	178
How do we describe random variables? . . . . .	178
What is the probability of an event? . . . . .	179
What is the distribution function of a random variable? . . . . .	180
What is the probability of a random variable taking a specific value? . . . . .	181
What is the probability density function of a random variable? . . . . .	181
How do we describe many random variables? . . . . .	184
What relationships may $n$ random variables have with each other? . . . . .	184
How do we define a random field? . . . . .	189
How can we relate two random variables that appear in the same random field? . .	190
How can we relate two random variables that belong to two different random fields? . . . . .	193
If we have just one image from an ensemble of images, can we calculate expectation values? . . . . .	195
When is a random field homogeneous with respect to the mean? . . . . .	195
When is a random field homogeneous with respect to the autocorrelation function? . . . . .	195
How can we calculate the spatial statistics of a random field? . . . . .	196
How do we compute the spatial autocorrelation function of an image in practice? . . . . .	196
When is a random field ergodic with respect to the mean? . . . . .	197
When is a random field ergodic with respect to the autocorrelation function? . . . . .	197
What is the implication of ergodicity? . . . . .	199
Box 3.1. Ergodicity, fuzzy logic and probability theory . . . . .	200
How can we construct a basis of elementary images appropriate for expressing in an optimal way a whole set of images? . . . . .	200
<b>3.2 Karhunen-Loeve transform</b> . . . . .	<b>201</b>
What is the Karhunen-Loeve transform? . . . . .	201
Why does diagonalisation of the autocovariance matrix of a set of images define a desirable basis for expressing the images in the set? . . . . .	201
How can we transform an image so its autocovariance matrix becomes diagonal? . . . . .	204
What is the form of the ensemble autocorrelation matrix of a set of images, if the ensemble is stationary with respect to the autocorrelation? . . . . .	210
How do we go from the 1D autocorrelation function of the vector representation of an image to its 2D autocorrelation matrix? . . . . .	211
How can we transform the image so that its autocorrelation matrix is diagonal? . . . . .	213

How do we compute the K-L transform of an image in practice? . . . . .	214
How do we compute the Karhunen-Loeve (K-L) transform of an ensemble of images? . . . . .	215
Is the assumption of ergodicity realistic? . . . . .	215
Box 3.2. How can we calculate the spatial autocorrelation matrix of an image, when it is represented by a vector? . . . . .	215
Is the mean of the transformed image expected to be really 0? . . . . .	220
How can we approximate an image using its K-L transform? . . . . .	220
What is the error with which we approximate an image when we truncate its K-L expansion? . . . . .	220
What are the basis images in terms of which the Karhunen-Loeve transform expands an image? . . . . .	221
Box 3.3. What is the error of the approximation of an image using the Karhunen-Loeve transform? . . . . .	226
<b>3.3 Independent component analysis</b> . . . . .	234
What is Independent Component Analysis (ICA)? . . . . .	234
What is the cocktail party problem? . . . . .	234
How do we solve the cocktail party problem? . . . . .	235
What does the central limit theorem say? . . . . .	235
What do we mean by saying that “the samples of $x_1(t)$ are more Gaussianly distributed than either $s_1(t)$ or $s_2(t)$ ” in relation to the cocktail party problem? Are we talking about the <i>temporal</i> samples of $x_1(t)$ , or are we talking about all possible versions of $x_1(t)$ at a given time? . . . . .	235
How do we measure non-Gaussianity? . . . . .	239
How are the moments of a random variable computed? . . . . .	239
How is the kurtosis defined? . . . . .	240
How is negentropy defined? . . . . .	243
How is entropy defined? . . . . .	243
Box 3.4. From all probability density functions with the same variance, the Gaussian has the maximum entropy . . . . .	246
How is negentropy computed? . . . . .	246
Box 3.5. Derivation of the approximation of negentropy in terms of moments . . . . .	252
Box 3.6. Approximating the negentropy with nonquadratic functions . . . . .	254
Box 3.7. Selecting the nonquadratic functions with which to approximate the negentropy . . . . .	257
How do we apply the central limit theorem to solve the cocktail party problem? . . . . .	264
How may ICA be used in image processing? . . . . .	264
How do we search for the independent components? . . . . .	264
How can we whiten the data? . . . . .	266
How can we select the independent components from whitened data? . . . . .	267
Box 3.8. How does the method of Lagrange multipliers work? . . . . .	268
Box 3.9. How can we choose a direction that maximises the negentropy? . . . . .	269
How do we perform ICA in image processing in practice? . . . . .	274
How do we apply ICA to signal processing? . . . . .	283
What are the major characteristics of independent component analysis? . . . . .	289
What is the difference between ICA as applied in image and in signal processing? . . . . .	290
What is the “take home” message of this chapter? . . . . .	292

<b>4 Image Enhancement</b>	<b>293</b>
What is image enhancement?	293
How can we enhance an image?	293
What is linear filtering?	293
<b>4.1 Elements of linear filter theory</b>	<b>294</b>
How do we define a 2D filter?	294
How are the frequency response function and the unit sample response of the filter related?	294
Why are we interested in the filter function in the real domain?	294
Are there any conditions which $h(k, l)$ must fulfil so that it can be used as a convolution filter?	294
Box 4.1. What is the unit sample response of the 2D ideal low pass filter?	296
What is the relationship between the 1D and the 2D ideal lowpass filters?	300
How can we implement in the real domain a filter that is infinite in extent?	301
Box 4.2. $z$ -transforms	301
Can we define a filter directly in the real domain for convenience?	309
Can we define a filter in the real domain, without side lobes in the frequency domain?	309
<b>4.2 Reducing high frequency noise</b>	<b>311</b>
What are the types of noise present in an image?	311
What is impulse noise?	311
What is Gaussian noise?	311
What is additive noise?	311
What is multiplicative noise?	311
What is homogeneous noise?	311
What is zero-mean noise?	312
What is biased noise?	312
What is independent noise?	312
What is uncorrelated noise?	312
What is white noise?	313
What is the relationship between zero-mean uncorrelated and white noise?	313
What is iid noise?	313
Is it possible to have white noise that is not iid?	315
Box 4.3. The probability density function of a function of a random variable	320
Why is noise usually associated with high frequencies?	324
How do we deal with multiplicative noise?	325
Box 4.4. The Fourier transform of the delta function	325
Box 4.5. Wiener-Khinchine theorem	325
Is the assumption of Gaussian noise in an image justified?	326
How do we remove shot noise?	326
What is a rank order filter?	326
What is median filtering?	326
What is mode filtering?	328
How do we reduce Gaussian noise?	328
Can we have weighted median and mode filters like we have weighted mean filters?	333
Can we filter an image by using the linear methods we learnt in Chapter 2?	335
How do we deal with mixed noise in images?	337

Can we avoid blurring the image when we are smoothing it? . . . . .	337
What is the edge adaptive smoothing? . . . . .	337
Box 4.6. Efficient computation of the local variance . . . . .	339
How does the mean shift algorithm work? . . . . .	339
What is anisotropic diffusion? . . . . .	342
Box 4.7. Scale space and the heat equation . . . . .	342
Box 4.8. Gradient, Divergence and Laplacian . . . . .	345
Box 4.9. Differentiation of an integral with respect to a parameter . . . . .	348
Box 4.10. From the heat equation to the anisotropic diffusion algorithm . . . . .	348
How do we perform anisotropic diffusion in practice? . . . . .	349
<b>4.3 Reducing low frequency interference . . . . .</b>	<b>351</b>
When does low frequency interference arise? . . . . .	351
Can variable illumination manifest itself in high frequencies? . . . . .	351
In which other cases may we be interested in reducing low frequencies? . . . . .	351
What is the ideal high pass filter? . . . . .	351
How can we enhance small image details using nonlinear filters? . . . . .	357
What is unsharp masking? . . . . .	357
How can we apply the unsharp masking algorithm locally? . . . . .	357
How does the locally adaptive unsharp masking work? . . . . .	358
How does the retinex algorithm work? . . . . .	360
Box 4.11. Which are the grey values that are stretched most by the retinex algorithm? . . . . .	360
How can we improve an image which suffers from variable illumination? . . . . .	364
What is homomorphic filtering? . . . . .	364
What is photometric stereo? . . . . .	366
What does flatfielding mean? . . . . .	366
How is flatfielding performed? . . . . .	366
<b>4.4 Histogram manipulation . . . . .</b>	<b>367</b>
What is the histogram of an image? . . . . .	367
When is it necessary to modify the histogram of an image? . . . . .	367
How can we modify the histogram of an image? . . . . .	367
What is histogram manipulation? . . . . .	368
What affects the semantic information content of an image? . . . . .	368
How can we perform histogram manipulation and at the same time preserve the information content of the image? . . . . .	368
What is histogram equalisation? . . . . .	370
Why do histogram equalisation programs usually not produce images with flat his- tograms? . . . . .	370
How do we perform histogram equalisation in practice? . . . . .	370
Can we obtain an image with a perfectly flat histogram? . . . . .	372
What if we do not wish to have an image with a flat histogram? . . . . .	373
How do we do histogram hyperbolisation in practice? . . . . .	373
How do we do histogram hyperbolisation with random additions? . . . . .	374
Why should one wish to perform something other than histogram equalisation? . .	374
What if the image has inhomogeneous contrast? . . . . .	375
Can we avoid damaging flat surfaces while increasing the contrast of genuine tran- sitions in brightness? . . . . .	377



How can we enhance an image by stretching only the grey values that appear in genuine brightness transitions? . . . . .	377
How do we perform pairwise image enhancement in practice? . . . . .	378
<b>4.5 Generic deblurring algorithms</b> . . . . .	<b>383</b>
How does mode filtering help deblur an image? . . . . .	383
Can we use an edge adaptive window to apply the mode filter? . . . . .	385
How can mean shift be used as a generic deblurring algorithm? . . . . .	385
What is toboggan contrast enhancement? . . . . .	387
How do we do toboggan contrast enhancement in practice? . . . . .	387
What is the “take home” message of this chapter? . . . . .	393
<b>5 Image Restoration</b> . . . . .	<b>395</b>
What is image restoration? . . . . .	395
Why may an image require restoration? . . . . .	395
What is image registration? . . . . .	395
How is image restoration performed? . . . . .	395
What is the difference between image enhancement and image restoration? . . . . .	395
<b>5.1 Homogeneous linear image restoration: inverse filtering</b> . . . . .	<b>396</b>
How do we model homogeneous linear image degradation? . . . . .	396
How may the problem of image restoration be solved? . . . . .	396
How may we obtain information on the frequency response function $\hat{H}(u, v)$ of the degradation process? . . . . .	396
If we know the frequency response function of the degradation process, isn't the solution to the problem of image restoration trivial? . . . . .	407
What happens at frequencies where the frequency response function is zero? . . . . .	408
Will the zeros of the frequency response function and the image always coincide? . . . . .	408
How can we avoid the amplification of noise? . . . . .	408
How do we apply inverse filtering in practice? . . . . .	410
Can we define a filter that will automatically take into consideration the noise in the blurred image? . . . . .	417
<b>5.2 Homogeneous linear image restoration: Wiener filtering</b> . . . . .	<b>419</b>
How can we express the problem of image restoration as a least square error estimation problem? . . . . .	419
Can we find a linear least squares error solution to the problem of image restoration? . . . . .	419
What is the linear least mean square error solution of the image restoration problem? . . . . .	420
Box 5.1. The least squares error solution . . . . .	420
Box 5.2. From the Fourier transform of the correlation functions of images to their spectral densities . . . . .	427
Box 5.3. Derivation of the Wiener filter . . . . .	428
What is the relationship between Wiener filtering and inverse filtering? . . . . .	430
How can we determine the spectral density of the noise field? . . . . .	430
How can we possibly use Wiener filtering, if we know nothing about the statistical properties of the unknown image? . . . . .	430
How do we apply Wiener filtering in practice? . . . . .	431

**5.3 Homogeneous linear image restoration: Constrained matrix inversion** 436

If the degradation process is assumed linear, why don't we solve a system of linear equations to reverse its effect instead of invoking the convolution theorem? . . . 436

Equation (5.146) seems pretty straightforward, why bother with any other approach? . . . . . 436

Is there any way by which matrix  $H$  can be inverted? . . . . . 437

When is a matrix block circulant? . . . . . 437

When is a matrix circulant? . . . . . 438

Why can block circulant matrices be inverted easily? . . . . . 438

Which are the eigenvalues and eigenvectors of a circulant matrix? . . . . . 438

How does the knowledge of the eigenvalues and the eigenvectors of a matrix help in inverting the matrix? . . . . . 439

How do we know that matrix  $H$  that expresses the linear degradation process is block circulant? . . . . . 444

How can we diagonalise a block circulant matrix? . . . . . 445

Box 5.4. Proof of equation (5.189) . . . . . 446

Box 5.5. What is the transpose of matrix  $H$ ? . . . . . 448

How can we overcome the extreme sensitivity of matrix inversion to noise? . . . . . 455

How can we incorporate the constraint in the inversion of the matrix? . . . . . 456

Box 5.6. Derivation of the constrained matrix inversion filter . . . . . 459

What is the relationship between the Wiener filter and the constrained matrix inversion filter? . . . . . 462

How do we apply constrained matrix inversion in practice? . . . . . 464

**5.4 Inhomogeneous linear image restoration: the whirl transform** . . . . 468

How do we model the degradation of an image if it is linear but inhomogeneous? . . 468

How may we use constrained matrix inversion when the distortion matrix is not circulant? . . . . . 477

What happens if matrix  $H$  is really very big and we cannot take its inverse? . . . . 481

Box 5.7. Jacobi's method for inverting large systems of linear equations . . . . . 482

Box 5.8. Gauss-Seidel method for inverting large systems of linear equations . . . . 485

Does matrix  $H$  as constructed in examples 5.41, 5.43, 5.44 and 5.45 fulfil the conditions for using the Gauss-Seidel or the Jacobi method? . . . . . 485

What happens if matrix  $H$  does not satisfy the conditions for the Gauss-Seidel method? . . . . . 486

How do we apply the gradient descent algorithm in practice? . . . . . 487

What happens if we do not know matrix  $H$ ? . . . . . 489

**5.5 Nonlinear image restoration: MAP estimation** . . . . . 490

What does MAP estimation mean? . . . . . 490

How do we formulate the problem of image restoration as a MAP estimation? . . . 490

How do we select the most probable configuration of restored pixel values, given the degradation model and the degraded image? . . . . . 490

Box 5.9. Probabilities: prior, a priori, posterior, a posteriori, conditional . . . . . 491

Is the minimum of the cost function unique? . . . . . 491

How can we select then one solution from all possible solutions that minimise the cost function? . . . . . 493

Can we combine the posterior and the prior probabilities for a configuration  $x$ ? . . 493

Box 5.10. Parseval's theorem . . . . . 496

How do we model in general the cost function we have to minimise in order to restore an image? . . . . .	499
What is the reason we use a temperature parameter when we model the joint probability density function, since its does not change the configuration for which the probability takes its maximum? . . . . .	501
How does the temperature parameter allow us to focus or defocus in the solution space? . . . . .	501
How do we model the prior probabilities of configurations? . . . . .	501
What happens if the image has genuine discontinuities? . . . . .	502
How do we minimise the cost function? . . . . .	503
How do we create a possible new solution from the previous one? . . . . .	503
How do we know when to stop the iterations? . . . . .	505
How do we reduce the temperature in simulated annealing? . . . . .	506
How do we perform simulated annealing with the Metropolis sampler in practice? . . . . .	506
How do we perform simulated annealing with the Gibbs sampler in practice? . . . . .	507
Box 5.11. How can we draw random numbers according to a given probability density function? . . . . .	508
Why is simulated annealing slow? . . . . .	511
How can we accelerate simulated annealing? . . . . .	511
How can we coarsen the configuration space? . . . . .	512
<b>5.6 Geometric image restoration . . . . .</b>	<b>513</b>
How may geometric distortion arise? . . . . .	513
Why do lenses cause distortions? . . . . .	513
How can a geometrically distorted image be restored? . . . . .	513
How do we perform the spatial transformation? . . . . .	513
How may we model the lens distortions? . . . . .	514
How can we model the inhomogeneous distortion? . . . . .	515
How can we specify the parameters of the spatial transformation model? . . . . .	516
Why is grey level interpolation needed? . . . . .	516
Box 5.12. The Hough transform for line detection . . . . .	520
What is the “take home” message of this chapter? . . . . .	526
<b>6 Image Segmentation and Edge Detection . . . . .</b>	<b>527</b>
What is this chapter about? . . . . .	527
What exactly is the purpose of image segmentation and edge detection? . . . . .	527
<b>6.1 Image segmentation . . . . .</b>	<b>528</b>
How can we divide an image into uniform regions? . . . . .	528
What do we mean by “labelling” an image? . . . . .	528
What can we do if the valley in the histogram is not very sharply defined? . . . . .	528
How can we minimise the number of misclassified pixels? . . . . .	529
How can we choose the minimum error threshold? . . . . .	530
What is the minimum error threshold when object and background pixels are normally distributed? . . . . .	534
What is the meaning of the two solutions of the minimum error threshold equation? . . . . .	535
How can we estimate the parameters of the Gaussian probability density functions that represent the object and the background? . . . . .	537

What are the drawbacks of the minimum error threshold method? . . . . .	541
Is there any method that does not depend on the availability of models for the distributions of the object and the background pixels? . . . . .	541
Box 6.1. Derivation of Otsu's threshold . . . . .	542
Are there any drawbacks in Otsu's method? . . . . .	545
How can we threshold images obtained under variable illumination? . . . . .	545
If we threshold the image according to the histogram of $\ln f(x, y)$ , are we thresholding it according to the reflectance properties of the imaged surfaces? . . . . .	545
Box 6.2. The probability density function of the sum of two random variables . . . . .	546
Since straightforward thresholding methods break down under variable illumination, how can we cope with it? . . . . .	548
What do we do if the histogram has only one peak? . . . . .	549
Are there any shortcomings of the grey value thresholding methods? . . . . .	550
How can we cope with images that contain regions that are not uniform but they are <i>perceived</i> as uniform? . . . . .	551
Can we improve histogramming methods by taking into consideration the spatial proximity of pixels? . . . . .	553
Are there any segmentation methods that take into consideration the spatial proximity of pixels? . . . . .	553
How can one choose the seed pixels? . . . . .	554
How does the split and merge method work? . . . . .	554
What is morphological image reconstruction? . . . . .	554
How does morphological image reconstruction allow us to identify the seeds needed for the watershed algorithm? . . . . .	557
How do we compute the gradient magnitude image? . . . . .	557
What is the role of the number we subtract from $f$ to create mask $g$ in the morphological reconstruction of $f$ by $g$ ? . . . . .	558
What is the role of the shape and size of the structuring element in the morphological reconstruction of $f$ by $g$ ? . . . . .	560
How does the use of the gradient magnitude image help segment the image by the watershed algorithm? . . . . .	566
Are there any drawbacks in the watershed algorithm which works with the gradient magnitude image? . . . . .	568
Is it possible to segment an image by filtering? . . . . .	574
How can we use the mean shift algorithm to segment an image? . . . . .	574
What is a graph? . . . . .	576
How can we use a graph to represent an image? . . . . .	576
How can we use the graph representation of an image to segment it? . . . . .	576
What is the normalised cuts algorithm? . . . . .	576
Box 6.3. The normalised cuts algorithm as an eigenvalue problem . . . . .	576
Box 6.4. How do we minimise the Rayleigh quotient? . . . . .	585
How do we apply the normalised graph cuts algorithm in practice? . . . . .	589
Is it possible to segment an image by considering the <i>dissimilarities</i> between regions, as opposed to considering the similarities between pixels? . . . . .	589
<b>6.2 Edge detection</b> . . . . .	591
How do we measure the dissimilarity between neighbouring pixels? . . . . .	591

What is the smallest possible window we can choose? . . . . .	592
What happens when the image has noise? . . . . .	593
Box 6.5. How can we choose the weights of a $3 \times 3$ mask for edge detection? . . . . .	595
What is the best value of parameter $K$ ? . . . . .	596
Box 6.6. Derivation of the Sobel filters . . . . .	596
In the general case, how do we decide whether a pixel is an edge pixel or not? . . . . .	601
How do we perform linear edge detection in practice? . . . . .	602
Are Sobel masks appropriate for all images? . . . . .	605
How can we choose the weights of the mask if we need a larger mask owing to the presence of significant noise in the image? . . . . .	606
Can we use the optimal filters for edges to detect lines in an image in an optimal way? . . . . .	609
What is the fundamental difference between step edges and lines? . . . . .	609
Box 6.7. Convolution of a random noise signal with a filter . . . . .	615
Box 6.8. Calculation of the signal to noise ratio after convolution of a noisy edge signal with a filter . . . . .	616
Box 6.9. Derivation of the good locality measure . . . . .	617
Box 6.10. Derivation of the count of false maxima . . . . .	619
Can edge detection lead to image segmentation? . . . . .	620
What is hysteresis edge linking? . . . . .	621
Does hysteresis edge linking lead to closed edge contours? . . . . .	621
What is the Laplacian of Gaussian edge detection method? . . . . .	623
Is it possible to detect edges and lines simultaneously? . . . . .	623
<b>6.3 Phase congruency and the monogenic signal . . . . .</b>	<b>625</b>
What is phase congruency? . . . . .	625
What is phase congruency for a 1D digital signal? . . . . .	625
How does phase congruency allow us to detect lines and edges? . . . . .	626
Why does phase congruency coincide with the maximum of the local energy of the signal? . . . . .	626
How can we measure phase congruency? . . . . .	627
Couldn't we measure phase congruency by simply averaging the phases of the harmonic components? . . . . .	627
How do we measure phase congruency in practice? . . . . .	630
How do we measure the local energy of the signal? . . . . .	630
Why should we perform convolution with the two basis signals in order to get the projection of the local signal on the basis signals? . . . . .	632
Box 6.11. Some properties of the continuous Fourier transform . . . . .	637
If all we need to compute is the local energy of the signal, why don't we use Parseval's theorem to compute it in the real domain inside a local window? . . . . .	647
How do we decide which filters to use for the calculation of the local energy? . . . . .	648
How do we compute the local energy of a 1D signal in practice? . . . . .	651
How can we tell whether the maximum of the local energy corresponds to a symmetric or an antisymmetric feature? . . . . .	652
How can we compute phase congruency and local energy in 2D? . . . . .	659
What is the analytic signal? . . . . .	659
How can we generalise the Hilbert transform to 2D? . . . . .	660
How do we compute the Riesz transform of an image? . . . . .	660

How can the monogenic signal be used? . . . . .	660
How do we select the even filter we use? . . . . .	661
What is the “take home” message of this chapter? . . . . .	668
<b>7 Image Processing for Multispectral Images</b> . . . . .	<b>669</b>
What is a multispectral image? . . . . .	669
What are the problems that are special to multispectral images? . . . . .	669
What is this chapter about? . . . . .	670
<b>7.1 Image preprocessing for multispectral images</b> . . . . .	<b>671</b>
Why may one wish to replace the bands of a multispectral image with other bands? . . . . .	671
How do we usually construct a grey image from a multispectral image? . . . . .	671
How can we construct a single band from a multispectral image that contains the maximum amount of image information? . . . . .	671
What is principal component analysis? . . . . .	672
Box 7.1. How do we measure information? . . . . .	673
How do we perform principal component analysis in practice? . . . . .	674
What are the advantages of using the principal components of an image, instead of the original bands? . . . . .	675
What are the disadvantages of using the principal components of an image instead of the original bands? . . . . .	675
Is it possible to work out only the first principal component of a multispectral image if we are not interested in the other components? . . . . .	682
Box 7.2. The power method for estimating the largest eigenvalue of a matrix . . . . .	682
What is the problem of spectral constancy? . . . . .	684
What influences the spectral signature of a pixel? . . . . .	684
What is the reflectance function? . . . . .	684
Does the imaging geometry influence the spectral signature of a pixel? . . . . .	684
How does the imaging geometry influence the light energy a pixel receives? . . . . .	685
How do we model the process of image formation for Lambertian surfaces? . . . . .	685
How can we eliminate the dependence of the spectrum of a pixel on the imaging geometry? . . . . .	686
How can we eliminate the dependence of the spectrum of a pixel on the spectrum of the illuminating source? . . . . .	686
What happens if we have more than one illuminating sources? . . . . .	687
How can we remove the dependence of the spectral signature of a pixel on the imaging geometry and on the spectrum of the illuminant? . . . . .	687
What do we have to do if the imaged surface is not made up from the same material? . . . . .	688
What is the spectral unmixing problem? . . . . .	688
How do we solve the linear spectral unmixing problem? . . . . .	689
Can we use library spectra for the pure materials? . . . . .	689
How do we solve the linear spectral unmixing problem when we know the spectra of the pure components? . . . . .	690
Is it possible that the inverse of matrix $Q$ cannot be computed? . . . . .	693
What happens if the library spectra have been sampled at different wavelengths from the mixed spectrum? . . . . .	693

What happens if we do not know which pure substances might be present in the mixed substance? . . . . .	694
How do we solve the linear spectral unmixing problem if we do not know the spectra of the pure materials? . . . . .	695
<b>7.2 The physics and psychophysics of colour vision . . . . .</b>	<b>700</b>
What is colour? . . . . .	700
What is the interest in colour from the engineering point of view? . . . . .	700
What influences the colour we perceive for a dark object? . . . . .	700
What causes the variations of the daylight? . . . . .	701
How can we model the variations of the daylight? . . . . .	702
Box 7.3. Standard illuminants . . . . .	704
What is the observed variation in the natural materials? . . . . .	706
What happens to the light once it reaches the sensors? . . . . .	711
Is it possible for different materials to produce the same recording by a sensor? . .	713
How does the human visual system achieve colour constancy? . . . . .	714
What does the trichromatic theory of colour vision say? . . . . .	715
What defines a colour system? . . . . .	715
How are the tristimulus values specified? . . . . .	715
Can all monochromatic reference stimuli be matched by simply adjusting the intensities of the primary lights? . . . . .	715
Do all people require the same intensities of the primary lights to match the same monochromatic reference stimulus? . . . . .	717
Who are the people with normal colour vision? . . . . .	717
What are the most commonly used colour systems? . . . . .	717
What is the <i>CIE RGB</i> colour system? . . . . .	717
What is the <i>XYZ</i> colour system? . . . . .	718
How do we represent colours in 3D? . . . . .	718
How do we represent colours in 2D? . . . . .	718
What is the chromaticity diagram? . . . . .	719
Box 7.4. Some useful theorems from 3D geometry . . . . .	721
What is the chromaticity diagram for the <i>CIE RGB</i> colour system? . . . . .	724
How does the human brain perceive colour brightness? . . . . .	725
How is the alchne defined in the <i>CIE RGB</i> colour system? . . . . .	726
How is the <i>XYZ</i> colour system defined? . . . . .	726
What is the chromaticity diagram of the <i>XYZ</i> colour system? . . . . .	728
How is it possible to create a colour system with imaginary primaries, in practice? .	729
What if we wish to model the way a particular individual sees colours? . . . . .	729
If different viewers require different intensities of the primary lights to see white, how do we calibrate colours between different viewers? . . . . .	730
How do we make use of the reference white? . . . . .	730
How is the <i>sRGB</i> colour system defined? . . . . .	732
Does a colour change if we double all its tristimulus values? . . . . .	733
How does the description of a colour, in terms of a colour system, relate to the way we describe colours in everyday language? . . . . .	733
How do we compare colours? . . . . .	733
What is a metric? . . . . .	733
Can we use the Euclidean metric to measure the difference of two colours? . . . . .	734

Which are the perceptually uniform colour spaces? . . . . .	734
How is the <i>Luv</i> colour space defined? . . . . .	734
How is the <i>Lab</i> colour space defined? . . . . .	735
How do we choose values for $(X_n, Y_n, Z_n)$ ? . . . . .	735
How can we compute the <i>RGB</i> values from the <i>Luv</i> values? . . . . .	735
How can we compute the <i>RGB</i> values from the <i>Lab</i> values? . . . . .	736
How do we measure perceived saturation? . . . . .	737
How do we measure perceived differences in saturation? . . . . .	737
How do we measure perceived hue? . . . . .	737
How is the perceived hue angle defined? . . . . .	738
How do we measure perceived differences in hue? . . . . .	738
What affects the way we perceive colour? . . . . .	740
What is meant by temporal context of colour? . . . . .	740
What is meant by spatial context of colour? . . . . .	740
Why distance matters when we talk about spatial frequency? . . . . .	741
How do we explain the spatial dependence of colour perception? . . . . .	741
<b>7.3 Colour image processing in practice</b> . . . . .	<b>742</b>
How does the study of the human colour vision affect the way we do image processing? . . . . .	742
How perceptually uniform are the perceptually uniform colour spaces in practice? . . . . .	742
How should we convert the image <i>RGB</i> values to the <i>Luv</i> or the <i>Lab</i> colour spaces? . . . . .	742
How do we measure hue and saturation in image processing applications? . . . . .	747
How can we emulate the spatial dependence of colour perception in image processing? . . . . .	752
What is the relevance of the phenomenon of metamerism to image processing? . . . . .	756
How do we cope with the problem of metamerism in an industrial inspection application? . . . . .	756
What is a Monte-Carlo method? . . . . .	757
How do we remove noise from multispectral images? . . . . .	759
How do we rank vectors? . . . . .	760
How do we deal with mixed noise in multispectral images? . . . . .	760
How do we enhance a colour image? . . . . .	761
How do we restore multispectral images? . . . . .	767
How do we compress colour images? . . . . .	767
How do we segment multispectral images? . . . . .	767
How do we apply <i>k</i> -means clustering in practice? . . . . .	767
How do we extract the edges of multispectral images? . . . . .	769
What is the “take home” message of this chapter? . . . . .	770
<b>Bibliographical notes</b> . . . . .	<b>775</b>
<b>References</b> . . . . .	<b>777</b>
<b>Index</b> . . . . .	<b>781</b>