

Roadmap Guide to Image Analysis (One-day Tutorial)

John C. Russ

Materials Science and Engineering Dept., North Carolina State University, Raleigh, NC

0. An Overview of Image Processing and Analysis

0.A. Relationships

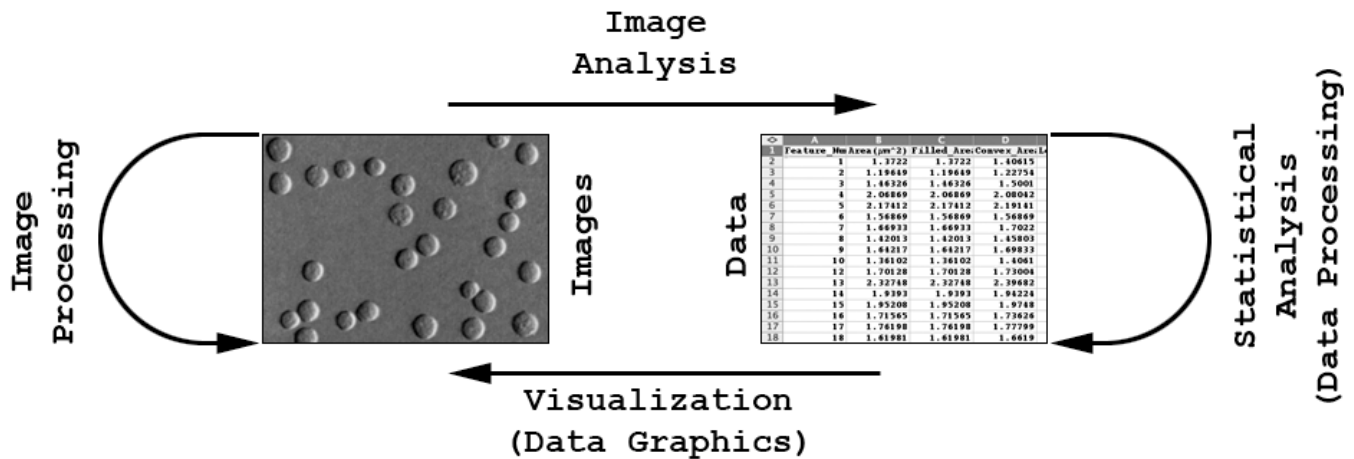


Image Processing operates on images and results in images, with changes intended to improve the visibility of features, or to make the images better for printing or transmission, or to facilitate subsequent analysis. **Image Analysis** is the process of obtaining numerical data from images. This is usually accomplished by a combination of measurement and processing operations. The data may subsequently be analyzed statistically, or used to generate graphs or other visualizations. Those topics are not covered in these notes.

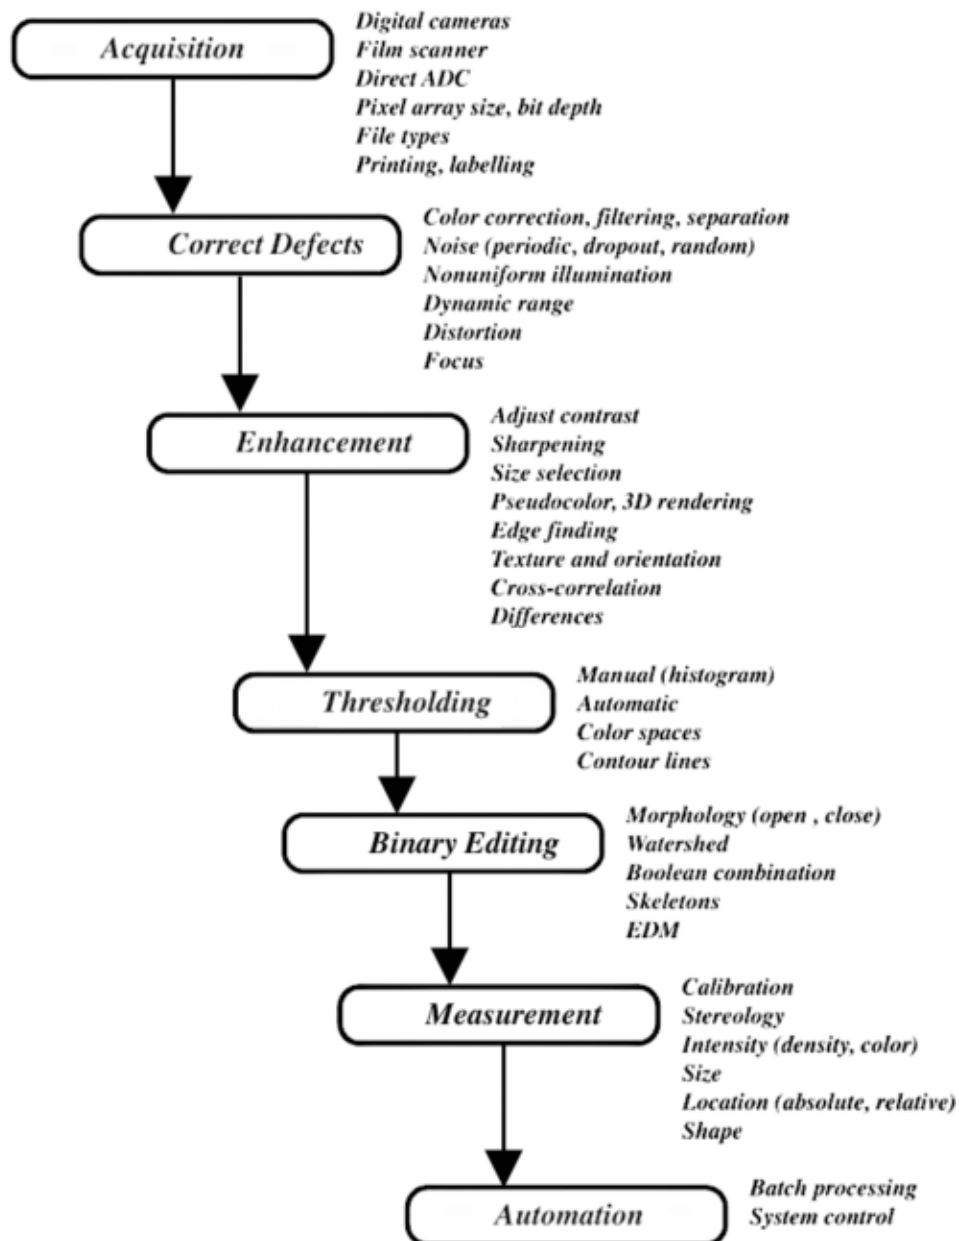
0.B. Work flow

The procedures in the following sections are generally applied in the order shown, as appropriate to a given image and final purpose (i.e., skip those steps that are not required in a particular application, but work from the top down). For instance, for a grey scale image you would skip the steps that pertain to color images; for images with good, uniform contrast you would skip steps that correct nonuniform illumination, adjust contrast, or increase local contrast; etc. For each topic, examples are shown to help explain the procedures and show results on a wide variety of images. Of course, original images should be archived and all processing and measurement steps documented.

The process of image analysis - obtaining meaningful numeric information from images - typically requires several steps performed in sequence. Many software programs contain at least the basic tools, ranging from NIH-Image, through Photoshop with Fovea Pro, to expensive solutions such as Matlab or Image Pro Plus (which also controls hardware), but it is important for the user to understand what each technique does and when to use them. The following roadmap offers a condensed guide to the most common procedures. In the great majority of typical situations, this roadmap should cover the required steps. It is important to understand what the various procedures do, because the names used in different programs vary widely and the same process may go by different names, or (worse yet) a name may be used for different operations.

The emphasis in this presentation is on when and why to apply various techniques to images. For a deeper explanation of how the various algorithms function, please refer to John Russ "The Image Processing Handbook," 5th edition, CRC Press, Boca Raton FL, 2006 (isbn 0-8493-7254-2). Algorithms discussed there and shown here are programmed into Fovea Pro, from Reindeer Graphics <<http://ReindeerGraphics.com>>, as plug-in modules for Adobe Photoshop (both PC and Mac) and compatible

programs (which include Photoshop Elements, Paint Shop Pro and Image Pro Plus). The CD with those plugins also contains a detailed hands-on step-by-step tutorial on the use of the routines, and many of the images used here.



Photoshop is used as a platform for these examples for a variety of reasons. It is a widely available, professionally supported software package that presents the same user interface on Mac and PC computers, can open a wide variety of image formats, and provides a simple yet quite complete interface that makes it possible to add functionality to the basic program (for more on programming the plugin modules, see “Introduction to Image Processing and Analysis” by John and Chris Russ, CRC Press, 2008, which includes the necessary software tools on CD). Demonstration copies of Photoshop and the Reindeer Graphics plugins can be downloaded from the web, and will run for several weeks to allow you to experiment with these procedures.

This tutorial does not attempt to teach the use of Photoshop itself, or try to help the reader locate each menu selection. If you are not familiar with Photoshop itself, you can download an illustrated guide to

the program and its proper use for scientific images from my website (www.drjohnruss.com/downloads/PS_intro.pdf).


Some of the examples shown toward the end of these notes are specific for Photoshop CS4 (Photoshop version 11), but the majority of the procedures can be applied using any version from 7 on, so if you have an older version it should work fine. Also note that Adobe offers significant price breaks for academic users, and licenses for multiple installations.

1. The program environment

1.A. Image acquisition, storage and printing

Devices (cameras, scanners, etc.)

It is important to understand the type of camera (color, monochrome, video, digital, tube-type, CCD, CMOS, etc.) or scanner (single pass, three pass, bit depth, reflective or transparency, etc.) and the characteristic types of artefacts that each produces (noise - both periodic and random, real resolution - as opposed to the number of stored pixels, color response, etc.). As a general rule of thumb, you want as many pixels as possible (consistent with the actual optical resolution) so that features are adequately defined for measurement. In the two examples below, two images of the same region of the specimen (a set of spheres that vary by a factor of nearly 40:1) are compared. The 2 Megapixel image allows the smallest objects to be confidently counted, and the larger sizes to be measured. The smaller image (approximately the resolution of a video camera) does not.



Enlarged portions of images captured with different numbers of pixels.

File types


TIFF files are a good choice on most computer platforms and for most software packages. Adobe PSD is also cross-platform and accommodates multiple layers, so that processing can be documented and annotation applied nondestructively. Avoid absolutely (!) any type of lossy compression (JPEG, wavelet, fractal) since you cannot be sure what details will be lost or changed that might be important in subsequent analysis.



Loss of detail in an images saved as jpeg compared to the original (40:1 compression).

Printing hardcopy

Dye-sub, ink jet, and laser printers may be useful depending on the intended use. Printer type, ink and paper affect image quality (balance against the uses for prints: draft, report, publication, poster, etc.). Correct color rendition requires ICC profiles for all output devices (screen and printer). The dynamic range of prints is less than that of the screen display, which is less than that of film.



Original



Dye-Sub



Ink Jet



Laser

Enlarged comparison of detail with different printer types.

Storage

Writing image files to CD-R, or storing on network RAID drives are popular tools, but you need a good database program and some discipline to make this work. Adobe Bridge has good facilities for using key words to locate images. QBIC (Query by image content) tools that can find an image by comparison or example are becoming practical but most image “cataloging” tools rely on a tiny thumbnail and human recognition.

1.B. Basic image display and manipulation

Managing multiple image windows

The “frontmost” or selected image window is usually the one being processed or edited. Some operations require two images (e.g., subtraction), and different programs offer various methods to designate how this is done. Most programs provide “Undo” capability (Photoshop saves up to 20 steps in its history), some create new copies with each operation.

Selection using ROI and wand tools

Various means of selecting regions of interest (rectangular or circular regions, arbitrary outlines, multiple disjoint areas, selection with a “wand” or region growing tool) may be available for interactive (manual) measurement or to restrict processing regions.

Drawing on and labeling images

Tools for adding text, lines and other labels are often very useful, with selection of colors, fonts, etc., as required. Keeping the text, etc., separate from the original image pixels is usually a good idea, preserving image details and also allowing for better quality graphics.

2. Correction of image defects

2.A. Color images

Is color correction required?

Locating points that should be neutral (equal amounts of red, green and blue) allows color adjustment in many cases. The example shows the Photoshop Adjustment Curves and the result (the dark, medium and bright points selected are indicated).



Neutral color adjustment.


The neutral color method makes independent adjustments to the red, green and blue channels, and does not compensate for overlap in the filters and response of the camera. If true colors in the image are important, and a reference standard with at least three known colors has been recorded under the same conditions, it is possible to solve for the constants in a set of equations of the form:

$$red = a_{11} * R + a_{12} * G + a_{13} * B$$

$$green = a_{21} * R + a_{22} * G + a_{23} * B$$

$$blue = a_{31} * R + a_{32} * G + a_{33} * B$$

This tristimulus correction is shown in the example. The measured RGB intensities in the red, green and blue reference areas were used to calculate the correction matrix, which was then applied to the entire image.

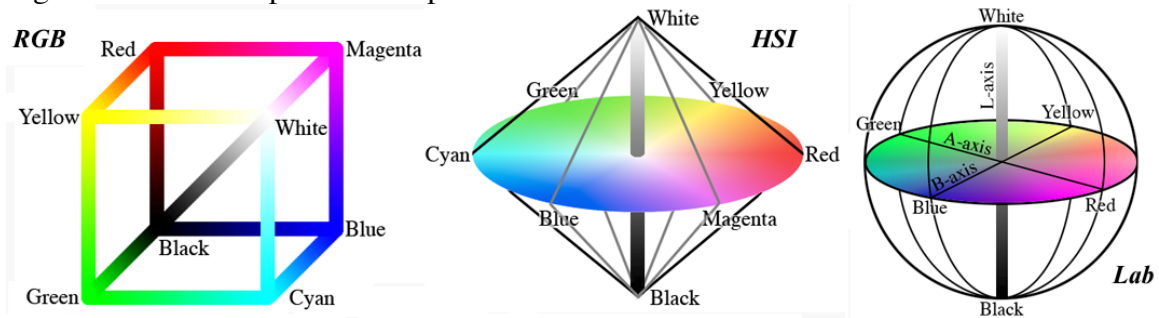


Color Space Blend (Tristimulus Values)			
Channel:	R	G	B
Result			
Red =	2.14387	-0.27217	-0.40766
Green =	-0.29021	1.61355	-0.38053
Blue =	-0.22500	-0.45712	1.80457

Tristimulus color correction.


Color filtering and separation to improve contrast

If color information is important to distinguish different structures, then applying color filters, or separating the image into color channels (RGB or HSI) may improve the contrast and simplify selection. Processing in HSI or LAB space avoids problems introduced when RGB channels are used.





Diagrams of color coordinate systems.

In some cases, selecting a single color channel provides good contrast and permits analysis. In the first example, the green channel is useful to reveal the fine capillaries in the eye. In the second, the hue channel shows the lipid spheres in the custard.




Extracting a single color channel from an image.

Color channels are also useful for merging individual greyscale images. This is commonly done to create stereo pair anaglyphs, by placing the left eye image into the red channel and the right eye image into the green and blue channels.




Combining stereo images to form a color anaglyph.

In general, any three greyscale images may be combined to form a color representation, without regard to whether the colors are accurately represented or even whether the images represent visible light information. The example shows three infrared telescope images taken at different wavelengths and combined.



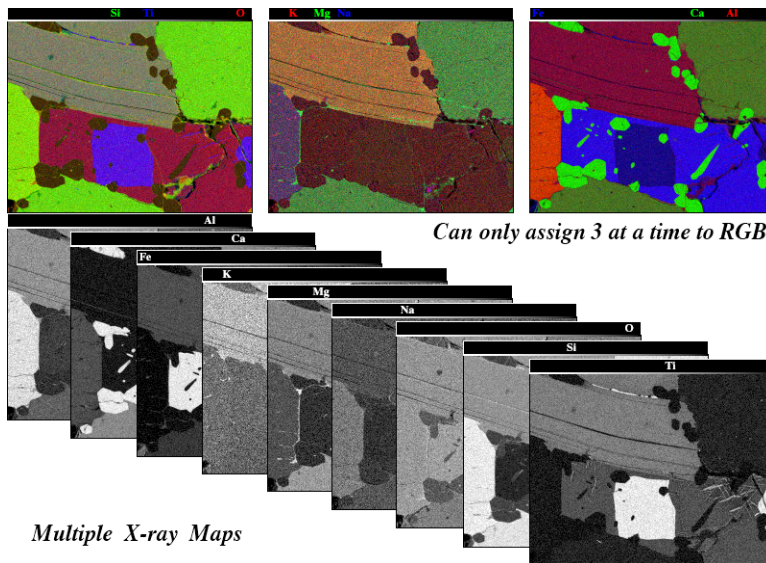
Combining three greyscale images to produce color.

It is possible and often useful to switch between different color coordinate spaces. In the example shown, three monochrome images (from a Landsat satellite) are combined as RGB channels to produce a color representation, and then a higher resolution image (from a Spot satellite) is substituted for the intensity channel leaving hue and saturation unchanged to produce a final result.



Replacing the intensity channel in a merged RGB image with a higher resolution image.

Merging multiple channels to produce a color image is only possible for three channels at a time, because human vision has just three kinds of color-sensitive cones. When more than three images are present, various combinations are possible.




Selecting different sets of three X-ray maps from a set of nine.


Reduction from color to greyscale

Color can be removed from an image in several ways: reducing saturation to zero, calculating the monochrome intensity or luminance, selecting a color channel, or calculating an optimum grey scale axis by principal components. The example shows a color image in which the details of the fingerprint are difficult to separate from the printed pattern on the paper. Removing the saturation leaves an

unsatisfactory greyscale image. Combining the channels using arbitrary weighting (Red=-65, Green=+185, Blue=+5, Contrast=-35) is an improvement but requires user manipulation. The principal components result does not.




Original image and the result of removing saturation.




Arbitrary weights applied to the color channels, and the principal components result.

Principal components analysis finds the set of axes in color space that correspond to the most significant combinations of pixel values in the various channels, and is useful for clarifying the structures in subtle color images and identifying the phases present in complex multichannel data sets. The example shows stained tissue before and after applying the principal components method.



Application of Principal Components analysis to a color image.

For the nine X-ray maps shown above, principal components analysis allows determining the linear combinations that have the greatest significance. In the example shown, the three principal axes are assigned to the red, green and blue display channels, revealing the location of the various chemical phases present in the specimen. The colocalization plot shows a projection onto the two major principal axes, with the clusters indicating the presence of six phases.





Assigning the three principal axes to RGB channels, and a colocalization plot.

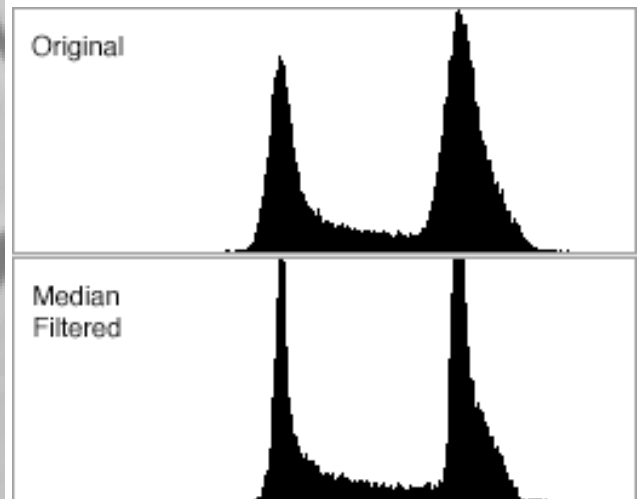

2.B. Noisy images

Random speckle noise

Random noise is commonly removed with either Gaussian smoothing (an optimum low pass filter) or median filtering. A Gaussian filter can be approximated by applying a kernel of weights to multiply by each pixel and its neighbors. Most image processing programs provide a built-in Gaussian smoothing (or blur) function with an adjustable radius. In the example, a 2-pixel standard deviation was used. Note that the result reduces the random noise but at the cost of blurring edges. Median filters are almost always preferred for noise reduction, as they retain edge sharpness and position. The neighborhood size (ideally a circular region for isotropic noise) controls the size of noise removed (and the processing time). In the example, a 2 pixel radius was used. Reduction of random noise also sharpens the peaks in the image histogram.





Original image with random noise and the Gaussian smoothed result.



Median filtered result and the change in the image histogram.


While it does not shift edges, the median filter does round corners and fills in gaps narrower than the radius used. The hybrid median avoids these problems, by performing multiple ranking operations in different neighborhoods around each pixel and combining the results.



Original image, conventional median, and hybrid median.

Shot noise (black and/or white spots) removal

Dropout pixels (common in some types of microscopy such as interference microscopes, or from dust on a scanned negative) are typically filled in with a median filter. In the example, a hybrid median replaces the random pixels set to black or white with the most likely value taken from a nearby pixel.



Removing shot noise with a hybrid median.

In color images, the meaning of a median value changes. Some programs use just the brightness or average of the red, green and blue values to perform the ranking of pixels, but a more general method employs a vector median. The plots the color coordinates of each pixel in the neighborhood and finds the one for which the sum of vector distances to all of the others is least. In the example, random noise and dropout pixels resulting from defects in the camera sensor are removed without erasing the fine lines and other detail in the image.



Applying a hybrid color vector median to remove random and shot noise.

Periodic noise removal

Filtering in Fourier transform space is very efficient as a way to remove electronic interference, halftone patterns from printing, vibration, some types of scan line noise, etc. The key to understanding the Fourier space methods is that in the power spectrum, the location of points corresponds to the frequency and orientation of the various sinusoids that together form the image, and show the amplitude of each sinusoid. Removing a “spike” removes the periodic noise without affecting other information in the

image. In the example, an image composed of just three sinusoids results in a Fourier transform with just three spikes (by convention each is plotted twice and the transform has rotational symmetry). Creating a filter or mask to isolate one frequency removes the corresponding lines from the image.





Image containing three sinusoids, and the Fourier transform power spectrum.



Mask to remove one spike, and the resulting inverse transformed image.

In the same way, spikes corresponding to periodic noise in a real image can be removed. In the example, fine details such as the girl's hair are preserved when the noise spikes are filtered out.






Image with periodic noise and the Fourier transform power spectrum.




Noise removal filter and resulting image.

A halftone printing pattern is another form of periodic noise that can be removed using the Fourier transform method. In the example, the power spectrum shows a pattern of regularly spaced spikes that correspond to the spacing and shape of the printing pattern. Removing them and applying a Butterworth filter to attenuate the high frequencies produces an image in which the details of the cat are clearly represented.




Halftone print and the FT power spectrum showing multiple spikes.



Power spectrum after applying the filter, and the inverse-transformed result.

Scan line noise removal

Even/odd scan line noise from video cameras, and some scan line artefacts from AFM, can be eliminated with a median filter using a tailored non-isotropic neighborhood. In the example, de-interlacing is applied to align the even and odd scan lines, which result in this case from the motion of the specimen as the video camera records the image. Then a median filter is applied in a neighborhood consisting of each pixel and its neighbors above and below.



Removing interlace and scan-line noise.

2.C. Nonuniform image illumination

Image measurement and analysis requires that the same features or objects look the same wherever they occur in the image. Nonuniform brightness resulting from variations in illumination, vignetting in the optics, curved specimen surfaces or thickness variations, etc., require correction before further steps can be taken.

Is a measured background image available?

A background image records variations in illumination, optical vignetting, etc. In the example, removing the stamp from the copystand allowed capturing an image of the flat, gray base, which was then subtracted pixel by pixel from the original to produce an image with uniform illumination.





Subtracting the background image produces a result with uniform lighting.


Is background visible throughout the image?

Automatic leveling based on fitting a polynomial to selected regions, or automatically to either a bright or dark background is fast and applies to gradual brightness variations typical of nonuniform illumination. In the first example, a manually marked region of interest (ROI) selects the area surrounding the stamp so that a polynomial background can be fit, which is then subtracted to level the image brightness.




Fitting and subtracting a background function to a selected region.

In the next examples, the brightest or darkest pixels in each region of the image are automatically selected to generate and subtract a polynomial background. This method assumes, of course, that there is background well distributed throughout the image area, and that it is locally either darker or lighter than the features.




Automatic leveling by fitting a polynomial to the brightest points.



Automatic leveling by fitting a polynomial to the darkest points.

Correcting varying contrast across an image


This is particularly applicable to sections of varying thickness, and is an extension of the polynomial fitting approach. In the example, varying specimen thickness causes a spatial variation in contrast. Fitting polynomials to both the brightest and darkest pixels in each region of the image allows stretching the local contrast between those limits to make the contrast uniform.




Automatic leveling of the contrast by polynomial fitting.

Are the features small in one dimension?


Rank-based leveling removes the features and allows creation of a background for leveling. This method can handle images in which the brightness varies in a more complex way with position than the polynomial models can accommodate, but requires that the features be small enough to be removed by erosion or dilation (called grey scale morphology). In the first example, replacing each pixel with its darkest neighbor removes the rice grains to produce a background that can be subtracted.







Generating and removing a background by grey scale morphology.

Rather than just eroding the features by dilating the background, it is generally preferable to follow the erosion with a dilation (the combination is called an opening). In this example removing the complex background from the diffraction pattern allows seeing smaller peaks in a plot of mean brightness vs. radius.






Original image, background produced by grey scale morphology, and net result.




Plots of mean brightness vs. radius for the original and leveled images.

When to subtract or divide by the background

This depends on whether the acquisition device is linear (e.g., CCD cameras) or logarithmic (e.g., film, vidicon tubes). Or, just try both and use whichever gives level results. The key to understanding this method is that if the features of interest can be selectively removed from the image, then the difference between the original image and the processed one is just those features of interest, and they can be isolated by removing the background. In this example, because the camera recorded intensities linearly rather than logarithmically, the background was divided into the original.




Original image and erosion to remove particles.



Dilation to restore size of underlying structures, and result of dividing by the background.

Using the color vector method introduced above for median filtering, color morphology can also be performed. In the example, selecting the red of the blood vessel as the target color allows generating a background image in which each pixel is replaced by the neighbor that is farthest from the target red color in color space coordinates. Subtracting this background isolates the blood vessels.



Using color vector ranking (color morphology) to generate a background for removal.

2.D. Expanding image contrast

Only after performing the preceding operations which alter the brightness and changes the image histogram), should adjustments to image contrast be performed. The histogram is an important tool for judging image contrast during acquisition. Ideally, the histogram should show pixels at all brightness levels from dark to bright, without any clipping at the ends (which means that possible information in bright or dark areas has been lost).




Image with a good histogram showing well-distributed brightness values.





Image whose cumulative histogram (red line) shows clipping of about 70% of the pixels to white.



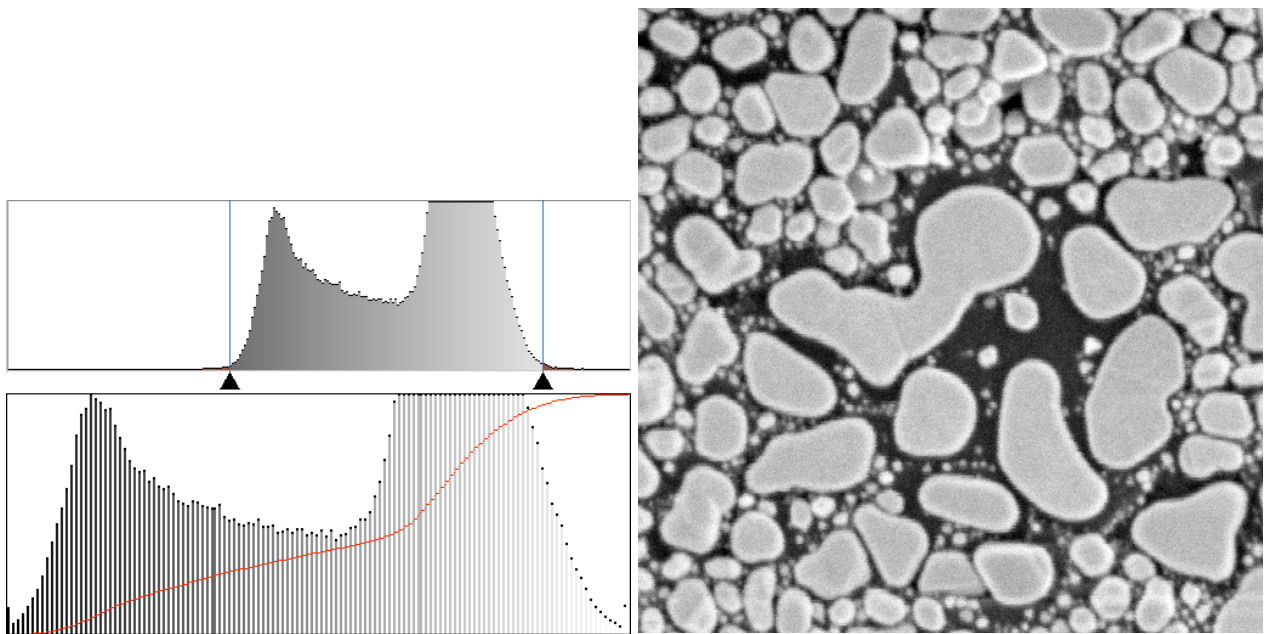


Image whose histogram shows low contrast, with values covering only half the range.

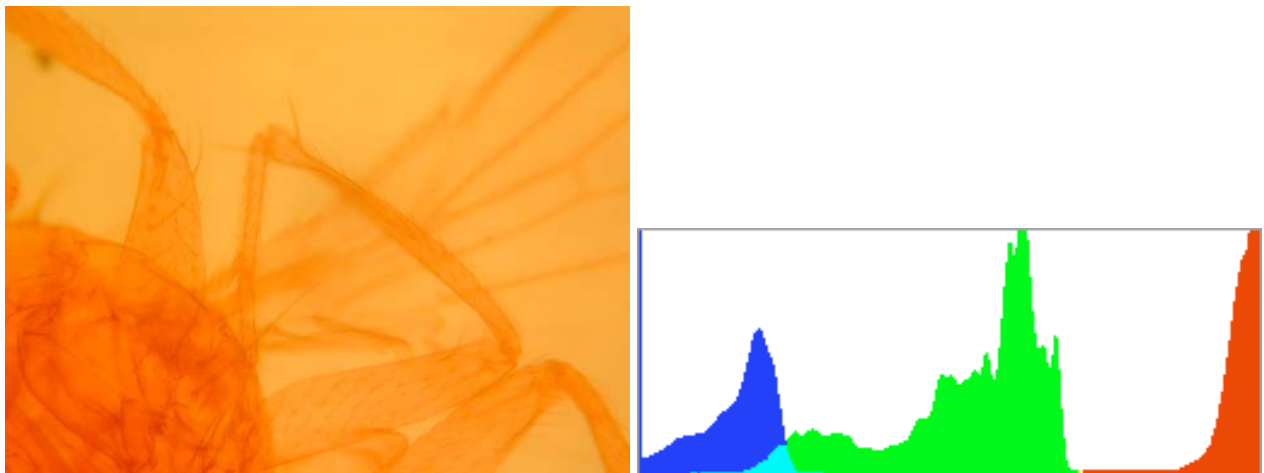
Linear expansion for grey scale and color images

Maximizing the contrast by setting dark and bright limits using the image histogram stretches the brightness over the full display range. There are just as many values with no pixels as in the original, but they are distributed throughout the brightness range. Automatic methods generally allow specifying the fraction of the pixels that can be set to exact white and black (typically 0.1 or 0.2%).




Setting dark and bright limits to stretch the histogram and increase image contrast.


When performing contrast adjustments on color images, it is important to work in HSI or Lab space and to adjust only the intensity or luminance channel, leaving the color information unchanged. In this example, automatically stretching the individual RGB channels produces color shifts, which should be avoided.



Original image, with histograms of the red, green and blue channels.




Color aliasing resulting from stretching the RGB channels individually.




Stretching just the L channel expands contrast and preserves color information.

Non-linear adjustments (gamma, equalization)



Nonlinear functions selectively expand contrast in one part of the grey scale range by contracting contrast elsewhere. This can sometimes be used to compensate for the characteristics of the acquisition device. Simple adjustment of a slider to set the medium gray value, or construction of a complete brightness transfer function, can be used. Manual adjustments to make the image “look good” are inconsistent and discouraged. Shaping the histogram to uniform (“equalization”) or Gaussian (as shown in the second example) are usually preferred. Again, for color images only the intensity or luminance channel is altered.





Original image and its histogram



Setting the mid-grey level to expand the bright range and show detail in dark areas.



Setting the mid-grey level to expand the bright range and show detail in light areas.



Creating a custom brightness transfer function and its result.



Automatic adjustment by shaping the histogram to a Gaussian reveals detail in both bright and dark areas.

Too great a dynamic range

Even in 8 bit images, and especially with a greater camera dynamic range, it is not possible to view or print images that cover the full range without clipping, while preserving the local detail. A range compression method known as homomorphic filtering produces superior results. This reduces the overall dynamic range of the image while preserving or increasing local contrast so that images can be better displayed or printed.



A wide dynamic range image and the result of homomorphic compression.

Negative images

Human vision is logarithmic, and can detect brightness variations better sometimes in the negative image.



The same image viewed as a positive and negative.

2.E. Distorted or foreshortened images

Making pixels square

Particularly with video cameras and analog to digital converters, adjustment may be needed to make dimensions the same vertically and horizontally.

Perspective distortion (non perpendicular viewpoint)

The light microscope has a shallow depth of field and usually the viewpoint is perpendicular to the sample, but in the electron microscope it is common to have tilted specimens which results in trapezoidal distortion. This must be corrected to permit meaningful measurements and even to facilitate

proper image processing. In the example, the perspective distortion of the building is corrected so that all of the windows are correctly seen to be the same size, but the street sign at the lower right has been distorted to an ellipse because it is not in the same plane as the building.



Image and the result after correcting for perspective distortion.

2.F. Focus problems


Shallow depth of field


Multiple images with different focus settings (but aligned and at the same magnification) can be combined by keeping the in-focus pixels from each (defined by having the greatest local variance).



Two images with different focus settings and their best-focus combination.

This method is easily extended to combine many images, each with a shallow depth of field, to form an extended focus composite, as shown in the example (the fixed curvature of microscope optics and the fact that focusing is accomplished by changing the distance from objective lens to specimen makes this particularly straightforward).







A sequence of three images and their extended focus composite.

Deconvolution of blurred focus

Dividing by the Fourier transform of the system point-spread function (PSF) can dramatically improve image resolution. The PSF can be measured, calculated, or estimated interactively. The effects of image noise are limited by apodization or the introduction of a Wiener constant. The example shows the use of an image of a single star representing the blur (due to incorrect mirror curvature in the Hubble telescope) to deconvolve a blurred image. It is also possible when no measured PSF is available to approximate one. The results shown are not as good as using a measured PSF but still much better than the original blurred image.




Blurred image and the measured PSF.




Results obtained by deconvolving with the measured PSF, and with an estimated one.

Deconvolution is capable of improving the sharpness and detail in many images, including ones taken using the automatic focusing algorithms in high-end digital cameras, as shown in the example.



Fragment of an image before and after focus deconvolution

It is also possible to remove blur due to motion, as shown in the example. The underlying assumption for deconvolution is that the blur - whether due to motion, defocus, etc. - is the same everywhere in the image.



Removal of motion blur by deconvolution.

2.G. Tiling large images

Shift and align multiple fields of view

Assuming perfect stage motion and no image distortion, stitching multiple fields together is easy. Several current microscopes produce images of entire slides by this method. In many real cases auto alignment and local scale adjustments are needed and require 10-20% overlap between fields, as shown in the example.



Four images taken by rotating the camera position on a tripod.



The merged composite result.


3. Enhancement of image detail

After corrections for defects and limitations in the as-acquired images have been made, enhancements to improve the visibility of some of the information in the image may be carried out. Generally this is done by suppressing some of the currently unneeded information so that what is left can be better seen.


3.A. Poor local contrast and faint boundaries or detail

Local equalization

Increasing the local contrast within a moving neighborhood is a powerful non-linear tool for improving the visibility of detail. The examples show the use of contrast and variance equalization.




Local contrast equalization removes the printed background to reveal fine details.





Local variance equalization brings out the fine scratches and other marks.

Sharpening (high pass filters)

Classical sharpening operations (e.g., the Laplacian) increase the local contrast at edges. They are equivalent to FFT-based high-pass filters. Directional derivatives and other neighborhood convolutions are applied in specific cases. In the first example, the important variation is only in the horizontal direction, so a first derivative can be applied to enhance the visibility of the detail. In the second example, applying the derivative parallel to the linear markings hides them.





Horizontal derivative applied to the central region of the image brings out faint lines.





Applying a first derivative parallel to the rolling marks hides them so the grain structure can be seen.

A directional first derivative only affects detail in a specific orientation. In the example, this is good for the cloud bands but note that the contrast varies around the limb of the planet. A second derivative (called the Laplacian) produces uniform contrast for detail in all orientations. Adding the Laplacian result back to the original (called Sharpening) is commonly used to improve the visual appearance of images.







First derivative oriented perpendicular to the cloud bands.



Laplacian and Sharpening filters.

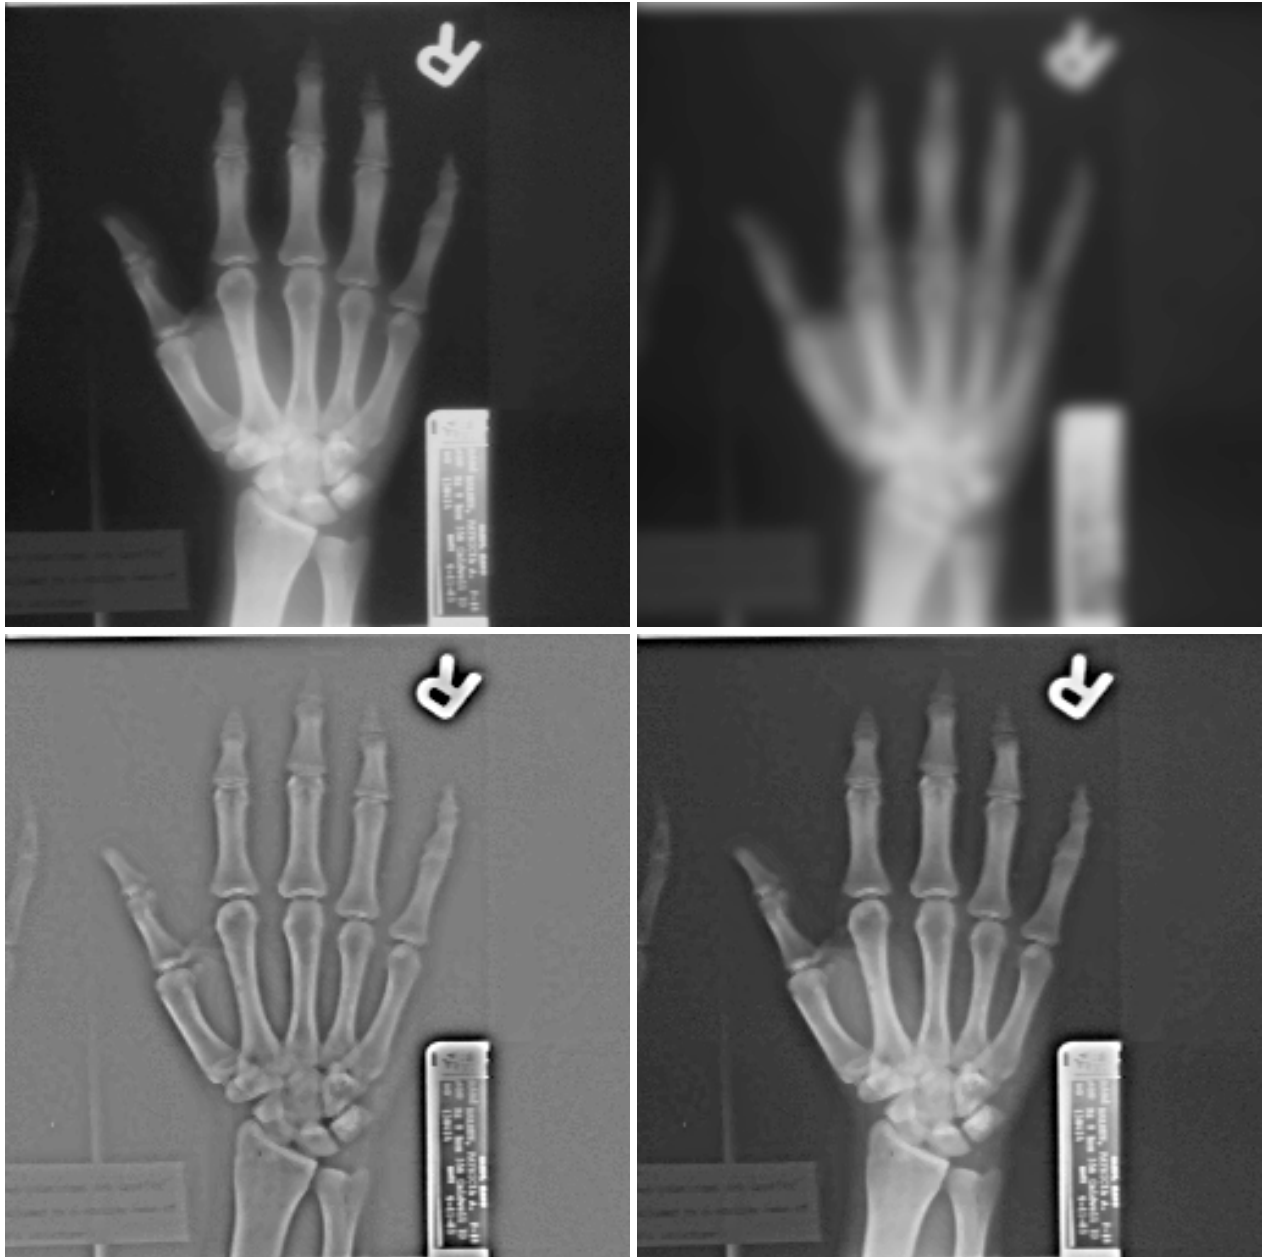
Derivative, Laplacian and Sharpening filters (and Gaussian smoothing) are examples of convolutions that are often performed by multiplying pixel values in a small neighborhood by weights and summing the result to obtain a new value for the central pixel, and repeating this for every pixel in the image. But convolutions can also be performed (with identical results but many fewer arithmetic operations) in the Fourier domain. This leads to the use of the terms high-pass (retaining high frequencies) and low-pass (retaining low frequencies). The example shows results for both cases.



High pass (ideal inverse) and Low pass (Gaussian) filters and their results.




Unsharp mask and difference of Gaussians

The basic unsharp mask operation is a powerful tool derived from long-established photographic darkroom technique. The example illustrates the procedure as a sequence of blurring to remove the detail of interest, subtracting that “unsharp” image from the original to isolate the details, and then adding that back to the original.



Step-by-step unsharp mask (original, blurred copy, difference, added to original).

Sharpening methods increase the visibility of noise, so the Difference of Gaussians technique, equivalent to a band-pass filter in Fourier space, should be used for noisy images, as shown in the example.




Original image, Laplacian and Difference of Gaussians.

Just as the median filter is a ranking operation for noise removal that is analogous to, but often better than, a Gaussian smooth, so too can rank operations be used for detail enhancement. Subtracting the result of a median filter or grey scale opening or closing from the original image leaves just the detail that was removed by the ranking operation. As for the unsharp mask, this detail may optionally be added back to the original. The advantage over the unsharp mask is the elimination of “halos” around edges, and more uniform enhancement in both light and dark regions.




Rank sharpening.




Comparison of unsharp mask (left) and rank sharpening (right)

Feature selection based on size


The Top Hat filter is closely related to the Difference of Gaussians, but based on ranking rather than smoothing. By comparing the brightest or darkest values in the central region (the “crown”) with that in the surrounding annulus (the “brim”), it selects features of a designated size and minimum separation distance. It can be used to select bright or dark objects for retention or removal.



Compare the maximum pixel value in the inner neighborhood (the "crown") to the maximum pixel value in the annular surrounding "brim."





Feature is detected




Feature is ignored

The operation of the top hat filter.



Locating the dust particles with a top hat filter


The top hat filter is also very useful for locating peaks ("spikes") in FFT power spectra. In the example, just the spikes are retained so that the inverse transform preserves the periodic structure while eliminating the random noise.



Original image, Fourier transform, mask generated by the top hat, retransformed result.

Pseudo-color (Cluts), pseudo-3D, and other display tools


The use of false- or pseudo-color look-up tables can make small brightness differences visually evident, depending on the choice of colors. Rendering images as though they were physical surfaces also assists in recognizing and interpreting detail. The examples show various combinations of these techniques.




Original image, pseudo-color, and geometric rendering.



Phong rendering, combined Phong + geometric rendering, combined pseudo-color + geometric rendering.



Original, pseudo-color, and geometric rendering.




Phong rendering, combined Phong + pseudo-color, combined pseudo-color+geometric rendering.


3.B. Are feature edges important?

Edge enhancement with derivative operators

Locating edges in images has been a major area of algorithm development because of the importance of edges in human vision, and for measurement. A variety of methods range from the simple Sobel to more advanced techniques like the Frei and Chen use multiple linear (or convolution) operators. Nonlinear (e.g., using the difference between the brightest and darkest pixels in a small neighborhood) and statistical operators (e.g., the local variance) are also useful. The examples illustrate several of these methods. It must also be noted that further steps after these operations may be required to isolate a structure for measurement, as shown in the last example.



Original and gradient vector (Sobel) edge delineation.





Original and Canny edge delineation.





Original and Canny edge delineation.



Sobel and Frei & Chen edge delineation.




Original image and variance operator



Thresholding and skeletonization to delineate grain boundaries for measurement.

Increasing edge sharpness and region uniformity


Maximum likelihood techniques can assign doubtful pixels to regions and create abrupt transitions that facilitate thresholding and segmentation.




Enlarged image region showing application of Kuwahara maximum likelihood operator.

3.C. Converting texture and directionality to grey scale or color differences


In many images, structures are discernible visually based on textural rather than brightness of color differences. Processing tools such as the Sobel orientation operator or a wide variety of local texture measurements based on entropy, fractal dimensions, statistical properties, etc., can convert these variations to brightness differences for thresholding.




Applying a variance operator to a test image.




Original image and fractal dimension texture isolating smooth structures.



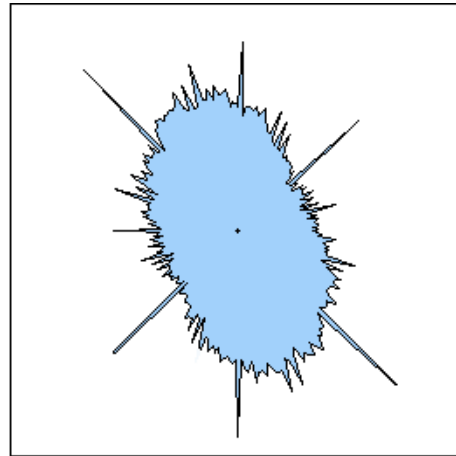

Range operator and outlines showing the selected structures.




Test image and the result of the Sobel orientation operator (scaled from 0 to 180°).




Fiber image and the result of the Sobel orientation operator (scaled from 0 to 360°).




The orientation values plotted as a histogram and a rose plot showing preferred fiber orientation.



Fiber image and result of color coding the Sobel orientation values to show fiber direction.



Original image and the Sobel orientation values.





Sobel edge delineation image, and assigning the orientation to the hue channel.

3.D. Fourier-space processing

Isolating and measuring periodic structures or signals

The FFT power spectrum can facilitate the selection of regular structures in the image and their measurement. Besides its use for removing periodic noise from images (e.g., electronic interference or halftone printing moiré patterns), this facilitates averaging and measuring repetitive structures (e.g., TEM images of atomic lattices). The spacing in the original image is calculated as the image width divided by the radius to the corresponding point in the Fourier power spectrum. In the examples, this is referenced to the (arbitrary) image dimensions as the subject of spatial calibration has not yet been covered.



Hi-resolution TEM image, Fourier transform power spectrum, and measurement of spacing in image.







Image of muscle tissue and measurement of the band spacing in the Fourier transform power spectrum.

Location of specific features


Cross-correlation with a target image is a powerful object finder. The examples show locating a target in a complex image.



Cross-correlation of the letter "A" (inset at left) locates similarly shaped letters (marks at right).




Step-by-step procedure: Original image with target (inset); cross-correlation result; top hat filter applied; threshold applied to select matches; counting the points determines the number of particles; convolution shows the image of the target particles without the complex background or dirt.



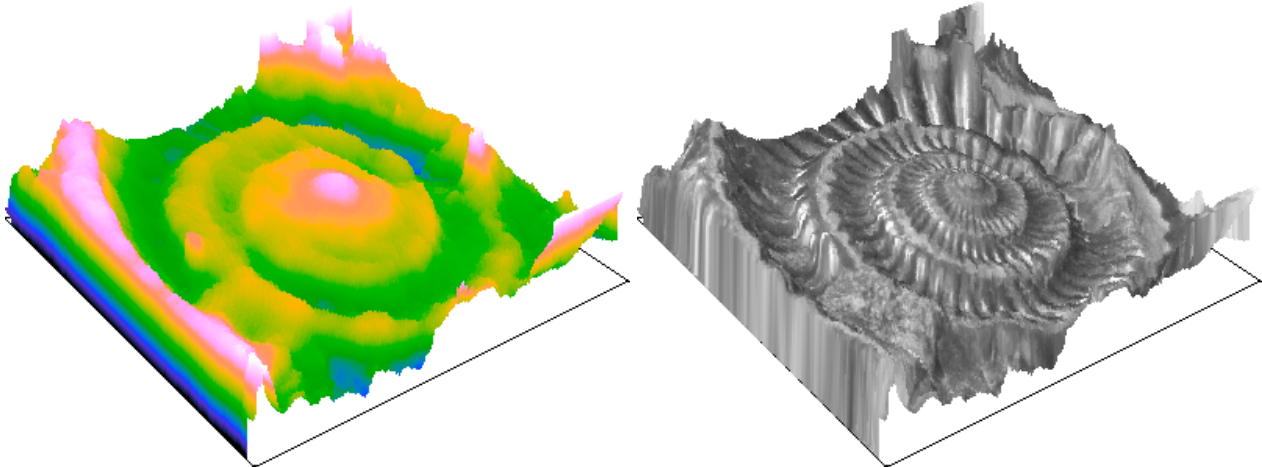

Cross-correlation result using one of the blisters as a target to locate the others.

Autocorrelation (correlating an image with itself) measures the representative size of structure present in a complex image. The example shows the autocorrelation result (inset) and the thresholded contour that corresponds to the mean feature size.



Autocorrelation results.

Cross-correlation is also used to match points, e.g., for image alignment or stereo pair parallax measurement.



Stereo pair parallax measurement: original image; parallax determined by cross-correlation; elevation data rendered and false-colored; geometric rendering with original image contrast superimposed.

3.E. Detecting image differences

Alignment

Alignment of images - both shift and rotation - consisting of either different color channels, serial sections, etc., may be performed manually or automatically, either based on the image contents or on discrete fiducial marks.

Subtraction and ratioing

Combining multiple images by subtraction or ratioing, and displaying the results in RGB or HSI color channels, is an effective way to eliminate lighting variations on curved surfaces, thickness variations in sections, variations in stain concentration, etc. It also makes it easy to display differences that might otherwise be overlooked.



The difference between these two images reveals even the motion of the hour hand.


4. Thresholding of image features

4.A. Thresholding using the histogram


Thresholding is used to define the pixels of interest (representing structure or objects in the image) as distinct from the background. Most typically it is based on finding a range of brightness or color that distinguishes the objects from the background.

Manual settings

Interactive setting of thresholds can be used to produce binary and contour images. This is where most image measurement errors arise, because of inconsistent human judgment.



Setting a manual threshold at the lowest point in the histogram (a common, but mistaken approach) produces a plausible but incorrect result.



An automatic method based on the histogram statistics produces the correct 60-40 result.

Automatic methods

A wide range of automatic methods for thresholding are available, many originally developed for reading printed text on paper. Selecting the appropriate technique is based on prior knowledge (that there are just two phases present, or that boundaries should be smooth).

Differentiating eq. (7) ...

$$\frac{\delta t}{\delta \gamma} = D(\beta) \cos \gamma, \quad \frac{\delta \theta}{\delta \gamma} = 1,$$

$$\frac{\delta t}{\delta \beta} = D'(\beta) \sin \gamma, \quad \frac{\delta \theta}{\delta \beta} = 1.$$

Using the Jacobian and the th ...

$$dr d\theta = [D(\beta) \cos \gamma - D'(\beta) \sin \gamma]$$

Performing a change of variables (7), and (8), we have

$$f(r, \phi) = \frac{1}{2} \int_0^{2\pi} \int_{-\gamma_m}^{\gamma_m} P_{\beta+\gamma} [D(\beta)]$$

$$\times [r \cos(\beta + \gamma - \phi) - D$$

$$\times [D(\beta) \cos \gamma - D'(\beta)$$

Fan Beam Projection $R_p(\gamma)$

fan-beam geometry. Fan-beam projections are rotation angle β and sampling angle α .

are generally satisfied. In particular, the holds when the source-to-origin distance

Differentiating eq. (7) ...

$$\frac{\delta t}{\delta \gamma} = D(\beta) \cos \gamma, \quad \frac{\delta \theta}{\delta \gamma} = 1,$$

$$\frac{\delta t}{\delta \beta} = D'(\beta) \sin \gamma, \quad \frac{\delta \theta}{\delta \beta} = 1.$$

Using the Jacobian and the th ...

$$dr d\theta = [D(\beta) \cos \gamma - D'(\beta) \sin \gamma]$$

Performing a change of variables (7), and (8), we have

$$f(r, \phi) = \frac{1}{2} \int_0^{2\pi} \int_{-\gamma_m}^{\gamma_m} P_{\beta+\gamma} [D(\beta)]$$


$$\times [r \cos(\beta + \gamma - \phi) - D$$

$$\times [D(\beta) \cos \gamma - D'(\beta)$$


Fan Beam Projection $R_p(\gamma)$

fan-beam geometry. Fan-beam projections are rotation angle β and sampling angle α .

are generally satisfied. In particular, the holds when the source-to-origin distance




Automatic thresholding of printed text. Note that the histogram does not have a minimum in this case.




Automatic thresholding of oil droplets based on the knowledge that they have smooth boundaries.

Images often require processing before thresholding using the techniques illustrated above, and may need additional processing afterwards as shown in these examples. Morphological operations (such as the skeletonization shown) and Boolean operations (such as the AND NOT combination shown) will be illustrated further in the next section.



Original image; brightness variation corrected; automatic thresholding applied; skeletonization used to delineate grain boundaries for measurement.




Test image; variance operator applied; high texture regions thresholded.




Gaussian blur applied; medium grey region thresholded; Boolean AND NOT used to combine thresholded binary images to select just the smooth, medium brightness region.

Select a color range by example, or with the histogram


In either RGB or HSI space, a range of colors can be specified to select structures for binarization. In most real cases, this involves more than a single color channel or a simple Boolean combination of individual color channels.



Stained potato: in this case the red channel can be processed individually to define the cells and the green channel can be processed individually to define the starch granules.




Candy: in this case the bright primary colors of the objects make color selection in RGB space possible (examples show the green and yellow candies, respectively).



Thresholding by selecting a range of hue and saturation values.


Sometimes it is easier to threshold the background than the features, if the objects vary in color or brightness but the background is uniform. In the example, the background is selected (based on its low saturation) and removed, after which the features can be separated (using a morphological watershed operation) for measurement.




Test image, background removed, features separated.

Contour lines (iso-brightness lines)


These lines are continuous, used to measure smoothness of boundaries, or to create contour maps.



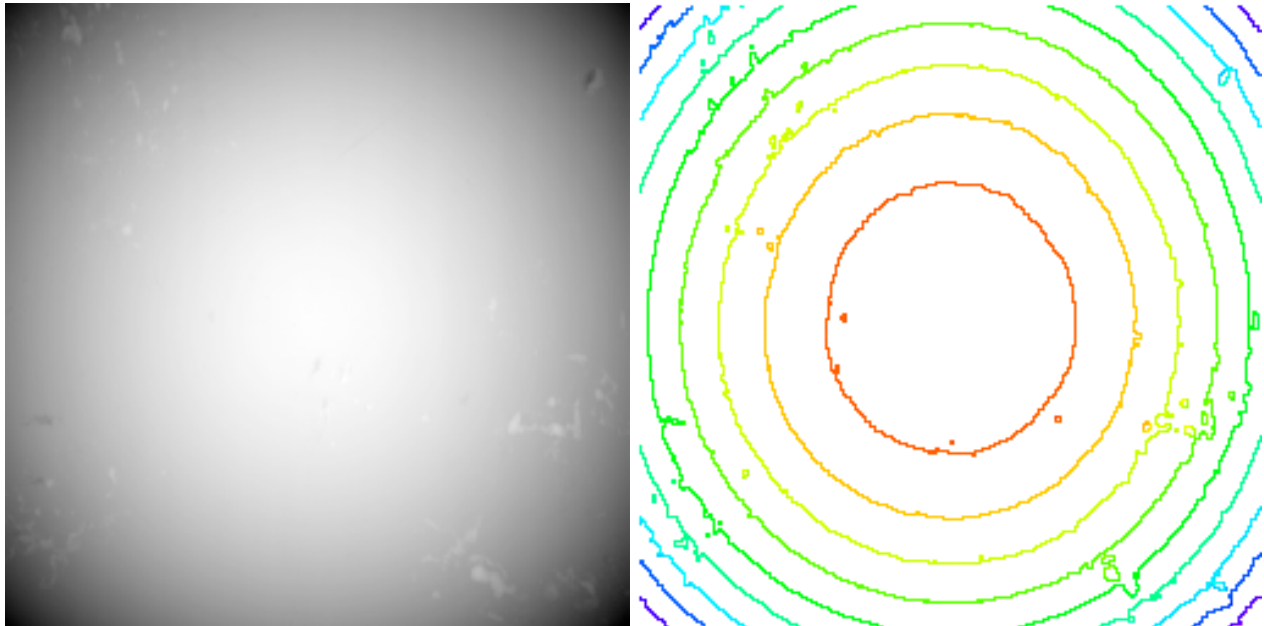
The total length of the contour lines drawn around the pores in the bread can be measured to calculate the surface area of the pores.



Drawing contour (iso-brightness) lines



Combining contour lines with geometric rendering and Phong shading.




Elevation data (scanned probe image) for a ball bearing, and a color-coded set of contour lines that reveal the roughness and an out-of-round condition.

4.B. Marking features manually

Region growing

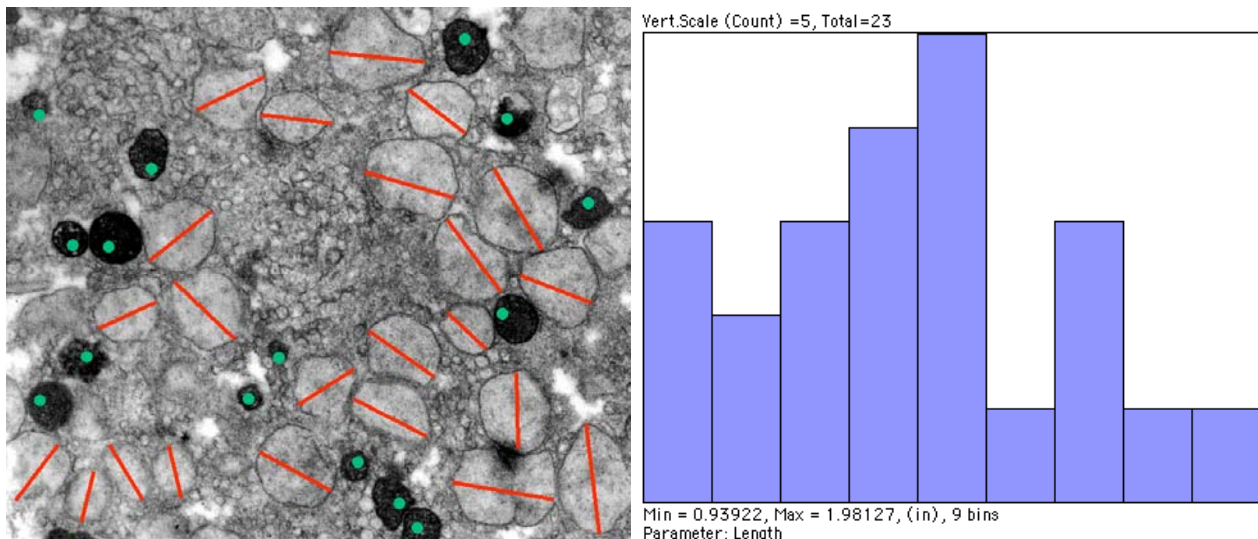
Rather than histogram-based methods, a “wand” approach allows human interaction to select features for measurement.



Selecting one region manually and then all similar regions automatically.

Manual lines, points, etc

Most systems provide some mechanism for the user to mark lines, points, circles, or other features onto the image which are counted and/or measured, when it is impractical to process the image to permit thresholding and automatic measurements. This combines the power of human recognition and independent knowledge with the computers ability to accurately count and measure.




Marking features in a unique color allows the computer to measure and count.

5. Binary image processing


5.A. Removing extraneous lines, points or other features

Erosion/dilation with appropriate coefficients to remove lines or points


The use of combinations of erosion, dilation, opening and closing permit selective correction of binary image detail when thresholding leaves artefacts behind. In the first example, thresholding the intermediate grey phase includes pixels that straddle the boundary between the bright and dark regions. An opening (erosion followed by dilation) removes the lines leaving the larger regions unchanged. In the second example, a test image with isolated black and white pixels, thin lines, and narrow gaps, has conditional erosion and dilation applied.



Original image, brightness leveled, thresholded for intermediate grey phase, classical opening applied.




Original image and conditional opening applied to pixels with more than 7 opposite colored neighbors.




Opening applied to pixels with more than 3 opposite neighbors, and a classic closing (applied to all pixels).

EDM based opening to remove small features or protrusions

Erosion + dilation based on the Euclidean distance map (EDM) is much faster and more isotropic than traditional pixel-based techniques. In the example, the iterative classic and conditional erosions produce various shapes, whereas thresholding the Euclidean distance map produces a circle, smaller than the original by an exact distance.






Circle eroded 30 iterations using conditional value of 0 (classic erosion), 1, 2, and 3 respectively, and the result of erosion using the EDM.

EDM based closing to fill in gaps


Dilation and erosion using the Euclidean distance map fills in gaps while retaining the shape of complex features. The example shows filling in the spaces in a lamellar structure and removing isolated dark features using classical closing and opening, compared to the same operations with the EDM.




Original image, thresholded binary, classic morphological result, EDM-based result.

5.B. Separating features that touch

The watershed, based on the EDM, is a powerful tool for separating touching convex shapes. The operation can be visualized by considering a rendered surface based on the EDM. Rain falling on this surface runs down each mountain to the sea, but where two mountains touch, water will arrive from two directions. Removing the pixels along these watershed lines separates the features. The peaks of the mountains (called UEP or ultimate eroded points) mark the centers of inscribed circles in each feature, and have a value that corresponds to the radius of that circle. They can be counted and used for measurement.



Operation of the watershed.



Test image, and the watershed result with UEPs superimposed.




Image of sand grains, automatically thresholded, watershed applied, and outlines overlaid on original.

Additional steps are sometimes required to make watershed segmentation possible. In the first example, it is necessary to fill in the outlines before applying the watershed. In the second, only those holes within the cells should be filled, not those between cells. The distinction is possible based on the shape (roundness) of the holes, treated as separate features. Selecting the round ones and ORing them with the cells fills the holes and enables a successful watershed.






Image of ice crystals in ice cream, automatic thresholding of boundaries, filling of holes, and watershed.



Original image of touching cells, automatically thresholded binary, and inverted image showing holes.

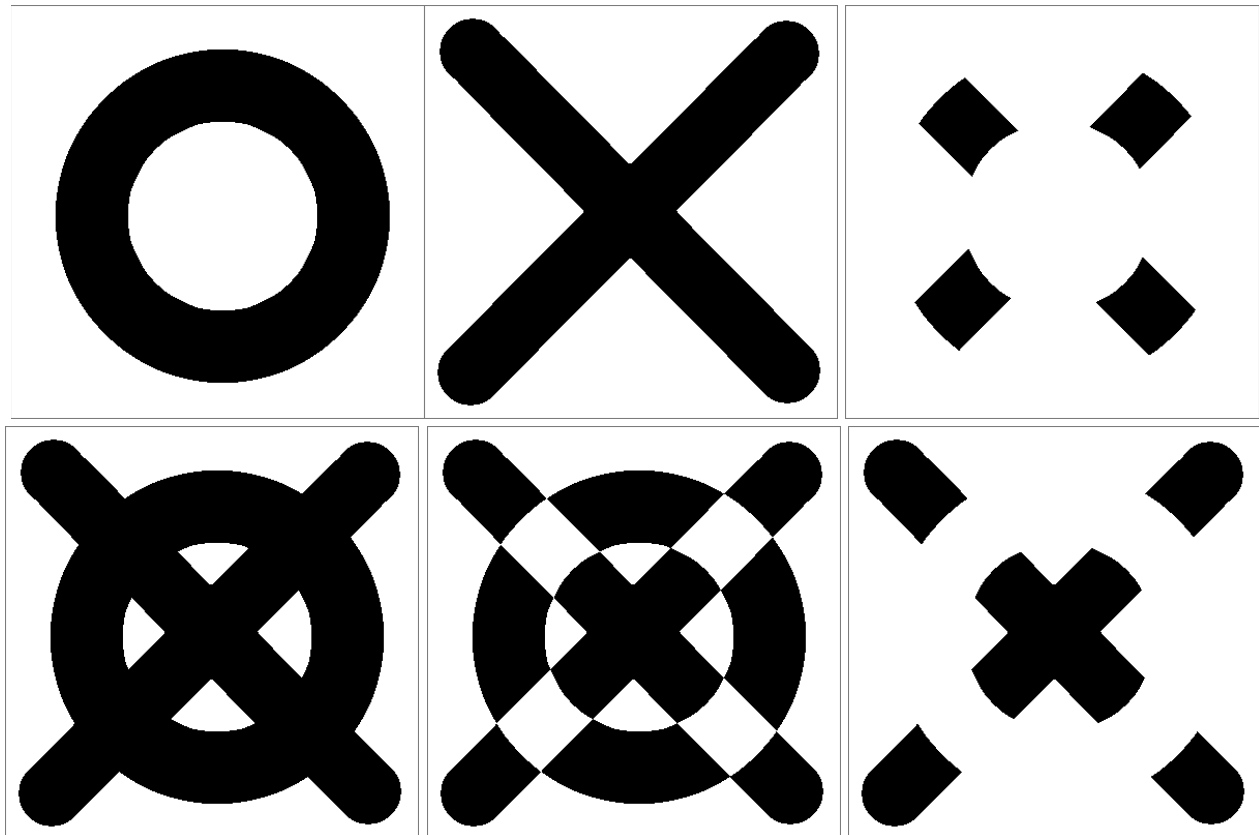


Holes within cells selected with Formfactor > 0.6, OR to combine with thresholded binary, and outlines of watershed results superimposed on the original.

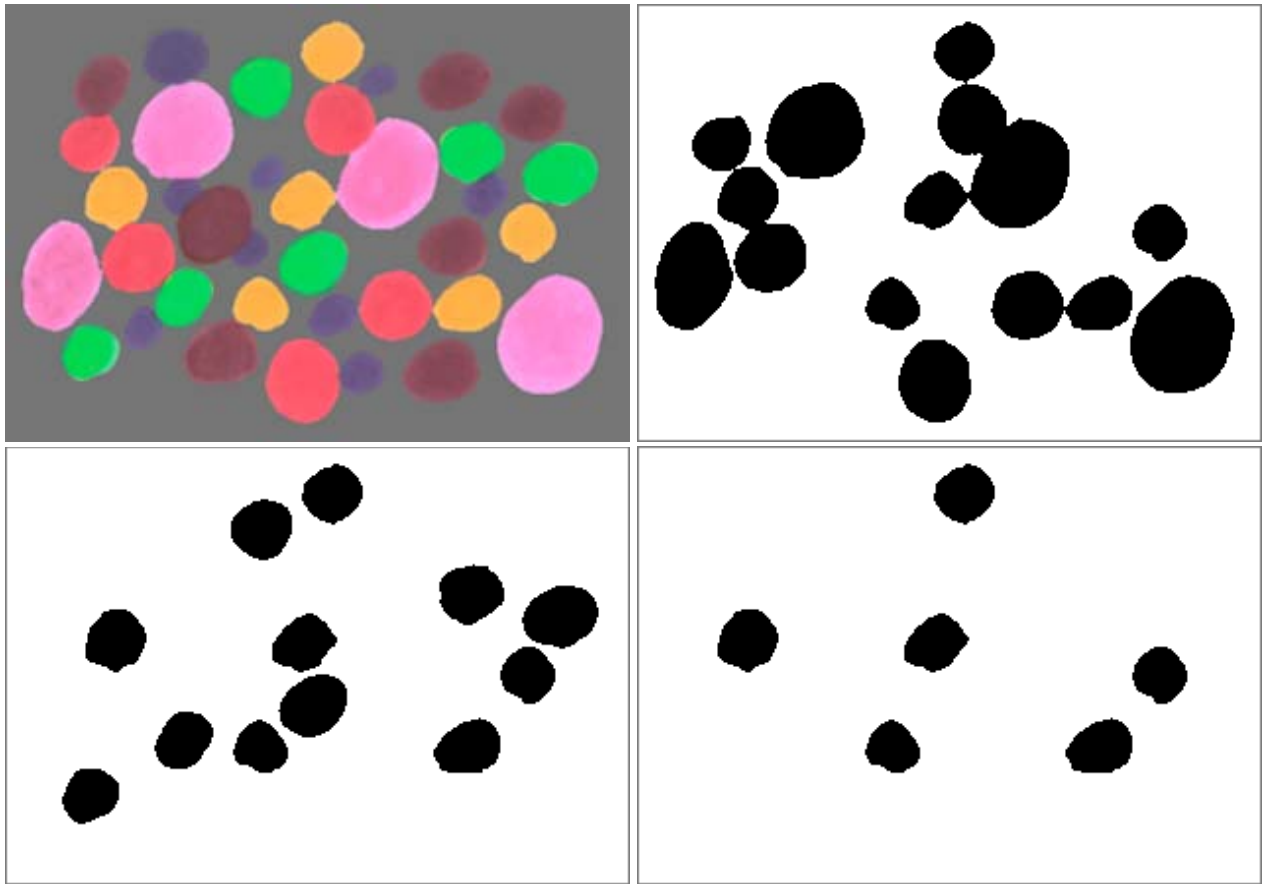
5.C. Combining multiple images of the same area to select features

Boolean logic to apply multiple criteria

For different colors, or images thresholded and processed differently, Boolean logic can combine the information and isolate the structures of interest. The Boolean functions are AND, OR, Exclusive-OR and NOT, which can be combined in many ways.

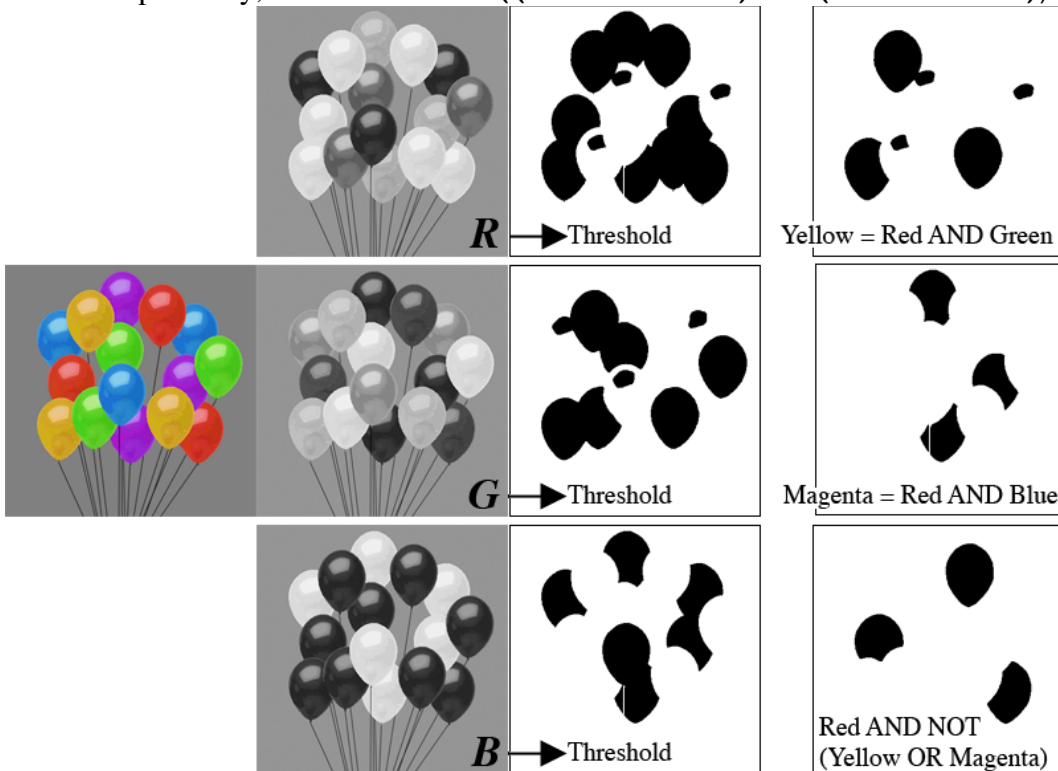


Two test shapes, and the results of Boolean AND, OR, Ex-OR and AND NOT operations.



Test image, red channel thresholded (red, yellow and magenta dots), green channel thresholded (yellow and green dots), AND result selecting yellow dots.

In more complex cases, multiple operations are required and parentheses are important to specify the order of operations. In this example, the red balloons can be selected using the bright areas in red, green and blue channels respectively, as **Red AND NOT ((Red AND Green) OR (Red AND Blue))**.



Boolean logic applied to select red balloons.





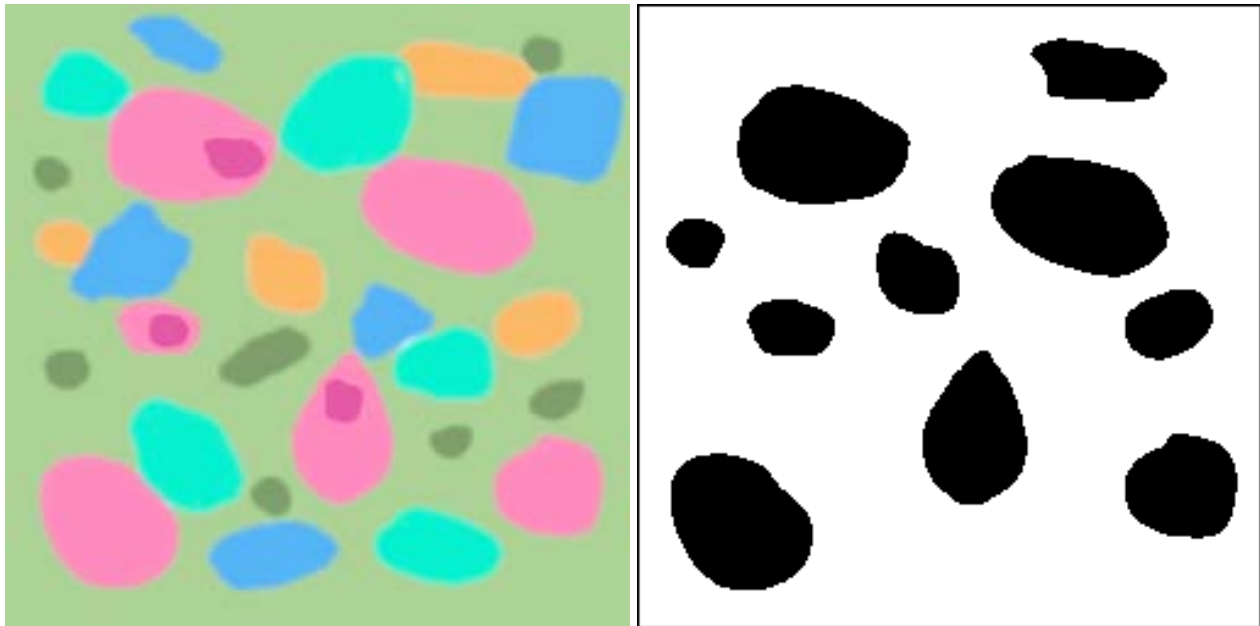
Image of gold particles, grey-scale opening to generate background, divided into original to isolate particles.




Thresholded organelles, thresholded particles, particles AND NOT organelles.

Using markers to select objects


Unlike the conventional pixel-based logic, this method selects entire features based on markers in a second image. In Boolean logic based on entire features (“Feature-AND”) the order of images is important, unlike the conventional pixel-based Boolean operators, which commute ($A \text{ AND } B = B \text{ AND } A$).




Test image and thresholded features with high red intensity.



Thresholded dark spots, and red features marked by dark spots.




Original image and thresholded cells.




Thresholded green nuclei and cells selected by nuclei as markers (outlines superimposed on original).

Feature selection by markers can also be used to implement the stereological disector, by swapping the foreground/background colors so that matched features are erased and unmatched ones are kept.




Two focal plane images, and the thresholded and watershed binaries.




Selection of the particles common to both layers, and those that are not (counted to implement the disector).

Region outlines as selection criteria

Outlines and boundaries are important markers to find adjacent features, as shown in the example.



Test image, thresholded orange objects, thresholded substrate, boundary along substrate.

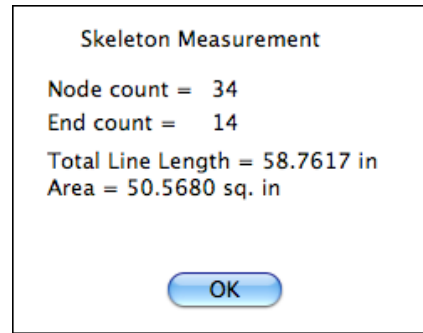



Boolean AND showing portions of boundary touching objects, and marker selection, showing objects that are touching the boundary.

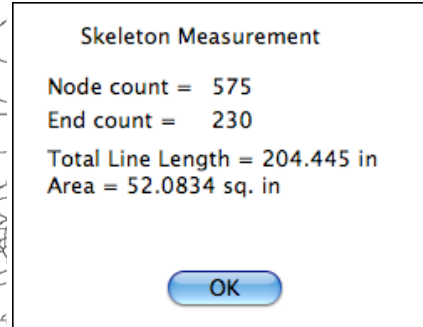


5.D. Feature skeletons provide important shape characterization

Grain boundary, cell wall, and fiber images

Broad thresholded lines can be thinned to single pixel width. Grain boundaries and cell walls form tessellations without end points, so pruning of branches with ends is an appropriate clean-up method. Removal or measurement of short branches based on length also provides a powerful tool. Several previous examples have illustrated this process. For overlapped fiber images, the number of fibers can be determined as half the number of end points, and the mean length as the total length divided by number. For images in which some fibers extend off the borders, the result is still correct (the other “half fiber” would be counted in another field of view).






Example of crossing fibers, skeletons, and end count.



The fiber density is half the number of ends divided by the image area.

Measuring total skeleton length, number of ends, number of branches

Removing node points (pixels with more than two neighbors) separates portions of the skeleton for measurement. Counting the number of ends gives an important piece of information about feature shape.



Original image, skeleton (with nodes removed), distribution of segment lengths.




Image of a gear, thresholded binary with skeleton superimposed showing 49 end points.

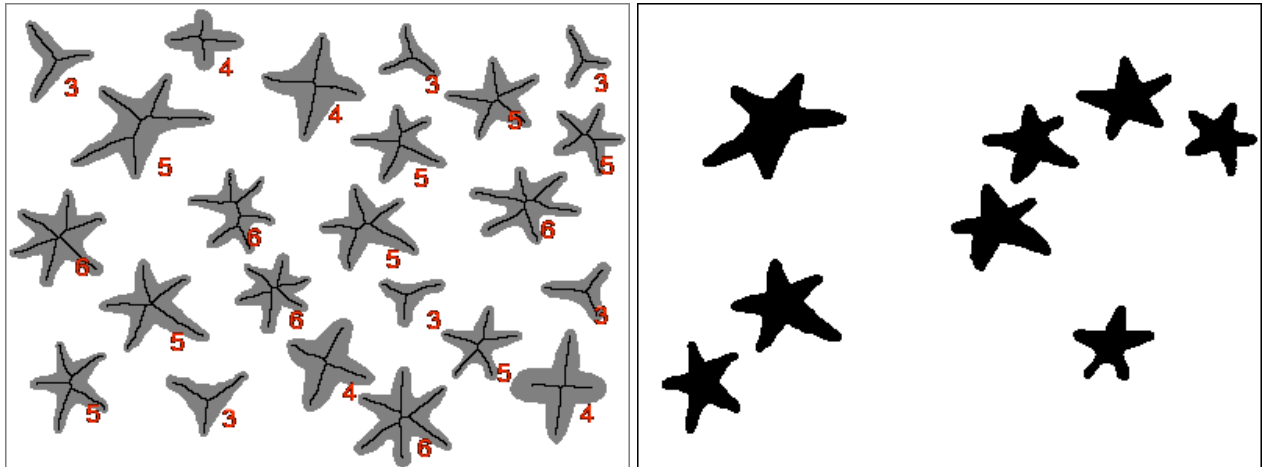



Image of multiple shapes, each with its skeleton superimposed and labeled with the number of end points, and selecting just those stars with five points.


5.E. Using the Euclidean Distance Map

Distances from boundaries or objects





Assigning values from the EDM to features enables the selection or measurement of feature position relative to irregular boundaries. Sampling the EDM with the skeleton provides width measurement for irregular structures (the EDM along the feature centerline records the radii of inscribed circles at every point).



Image, thresholded interior, EDM encoding distance from boundary.

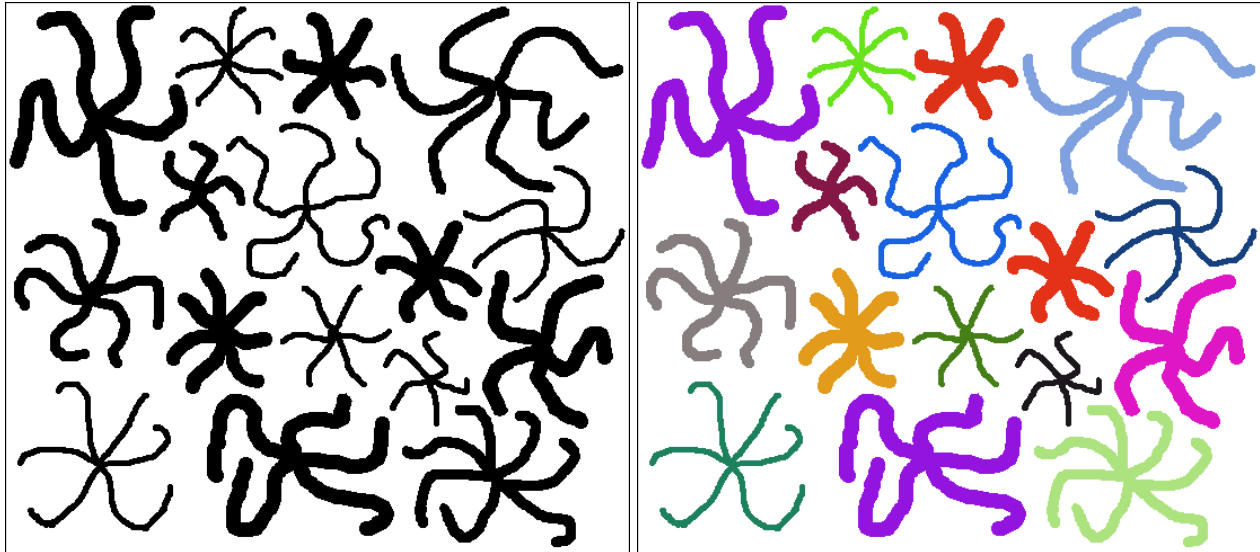


Objects within region, EDM values assigned to pixels, plot of number vs. distance from boundary.




Measuring the width of an irregular object: original, EDM (pseudo-colored for clarity), EDM values assigned to the skeleton; Excel display of the distribution of widths for the feature.

Complex features can be color coded according to measured properties, including ones derived from the skeleton and EDM. In this example, visual examination of the black features does not readily reveal which ones (if any) share the properties of number of arms, length of arms, and thickness of arms. Color coding the red, green and blue channels according to each of those properties immediately shows which pairs are the same.




Shading the red, green and blue channels according to the number of end points, mean branch length and mean branch width, respectively, allows the two pairs of objects that share those properties to be recognized immediately since they have the same color.

This example combines the skeleton, EDM and marker technology to perform a complex measurement. The original image represents a dendritic pattern such as roots, neurons, or a crack in a brittle material. The visual impression is that the terminal branches of the pattern are shorter as they move farther from the central body. Using the skeleton end points as markers to select the disconnected segments isolates the terminal branches. Assigning to these branches the EDM values of the space around the central body, then then plotting the length of each terminal branch against the minimum EDM value (the distance from the closest point to the central body) confirms the correlation.



Original image, the skeleton, and the terminal branches (disconnected segments selected by end points as markers).



The EDM of region around central cell (pseudo-colored for clarity), EDM values assigned to terminal branches, plot of minimum EDM value (closest point to central cell) vs. length of each branch, showing correlation.

6. Measurements

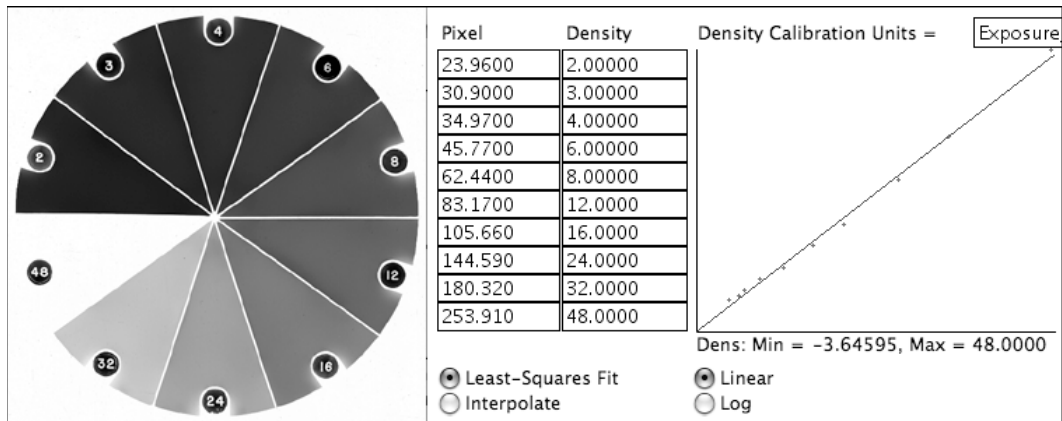
6.A. Calibration

Calibrating image dimensions

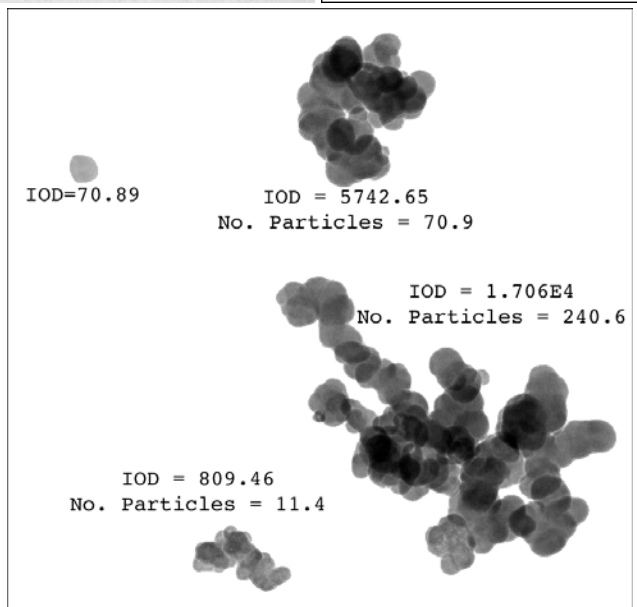
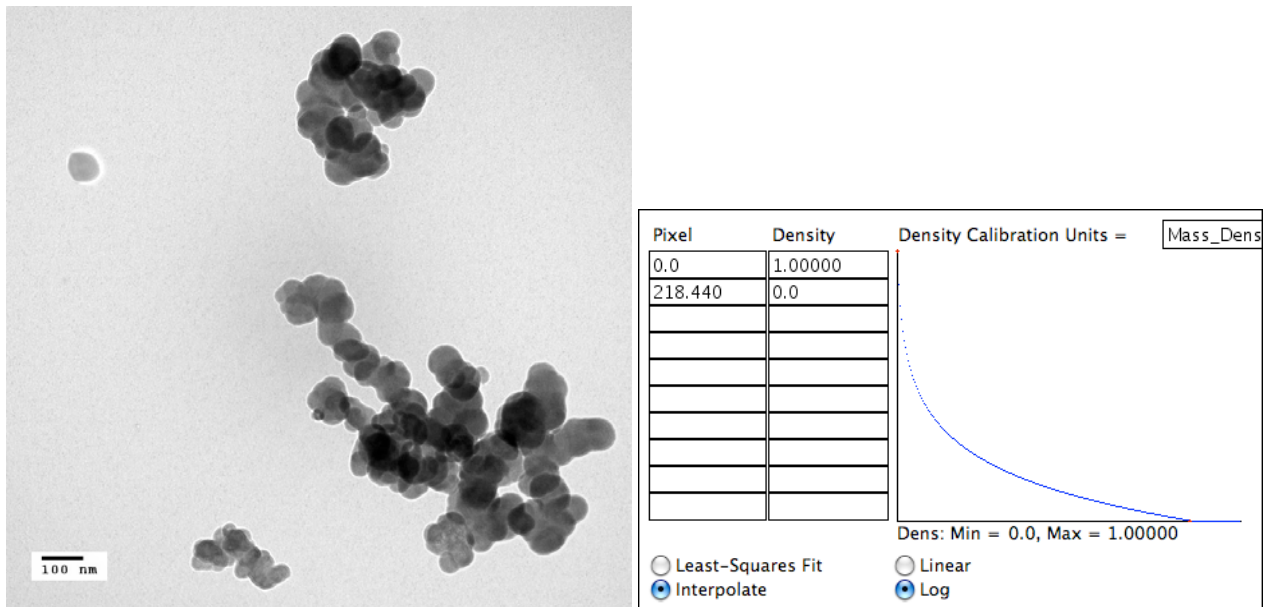
This is typically done by acquiring an image of a stage micrometer or some standard object.

Calibrating density/greyscale values

This requires density standards, which are readily available at macro scales but hard to make for microscopes. Recalibration is needed often because of variation in line voltages, lamp aging, optical variations, camera instabilities, etc. The example shows fitting a calibration curve to multiple points.




Some measurements can be performed using internal standards. The example shows determining the approximate number of particles in each cluster by measuring the integrated optical density of each cluster and ratioing that to the value for a single particle. The density calibration curve simply assigns zero density to the background intensity, and an arbitrary value of 1 to black, and the interpolates Beers' law (exponential absorption).



Original image, calibration curve and measurements of integrated optical density used to estimate number of particles in each cluster.

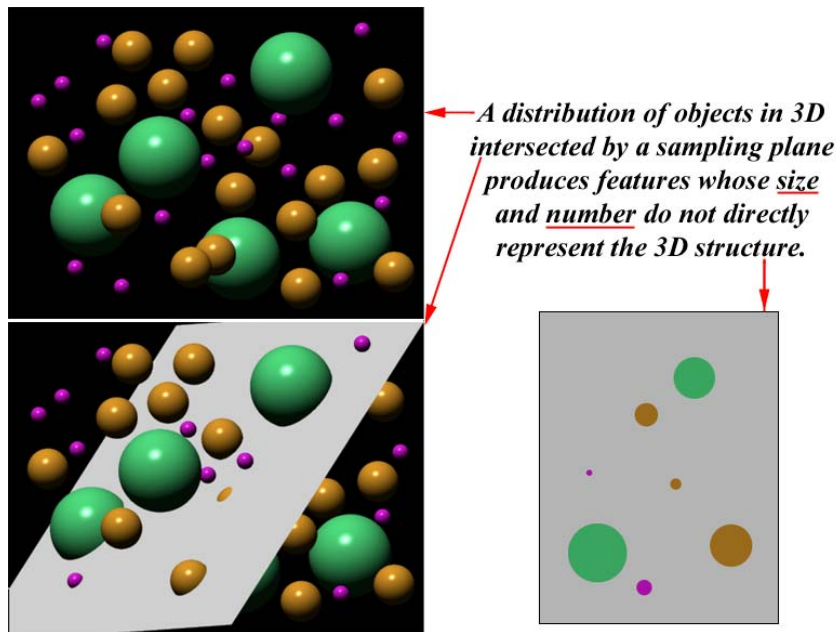
6.B. Intensity Profiles

Profiles of intensity, calibrated density, or color information often provide efficient ways to extract data for analysis. The example shows a plot of brightness vs. horizontal position, averaging the pixels vertically in the selected region. Plots of hue, saturation, etc. are also useful in some instances.



6.C. Stereological measurements

Apparently simple measurements such as the size of objects seen in a section, or even their number, do not directly report the important three-dimensional information.



Relationships between 3D structure and 2D images make it easy to correctly measure metric properties such as volume fraction, surface area, length, etc., as well as topological properties such as number and connectivity. Modern stereology emphasizes the best ways to perform sampling and sectioning to avoid bias. Most stereological techniques utilize appropriate grids superimposed on images and counting of intersections, which can often be automated. The basic 3D measurements made from 2D section images are:


Volume Fraction V_V
 = Area Fraction A_A
 = Line Fraction L_L
 = Point Fraction P_P

Surface Area per Unit Volume S_V
 = $(4/\pi) B_A$ (Boundary Length / Area)
 = $2 P_L$ (Intersection Points / Line Length)

Length per Unit Volume L_V
 = $2 P_A$ (Intersection Points / Area)
 = $2 P / (L \cdot T)$ where L is Line Length and T is Section Thickness


Number per Unit Volume) N_V
 Determined by Counting Ends (Disector)

For the example image, the grey cells can be measured to determine the volume fraction either by measuring the area fraction or by generating a point grid and counting the number of points that “hit” the phase. The latter method has the advantage that counting independent events (i.e., the grid must be sparse compared to the structure), the precision is known (the standard deviation of the result is just the square root of the number of hits). This makes it possible to estimate the number of fields of view that should be examined to obtain the required precision.



Original image, and binary image with superimposed grid. A Boolean AND is used to locate the hits. The area fraction measurement for the grey phase is 16.9%. The point count is $14/90 = 15.6\%$.


Surface area per unit volume can be measured by isolating the boundary of interest in the image (the example has three types of interface: between two white cells, between two grey cells, and between white and grey), and either measuring the length of the line or generating a line grid and counting intersections. In the example, the white-grey boundary has a measured total length of $551.3 \mu\text{m}$ in an image area of $5454.6 \text{ square } \mu\text{m}$, for a calculated surface area of $0.1287 \text{ square } \mu\text{m per cubic } \mu\text{m}$. The grid, with a total length of $233.3 \mu\text{m}$, produces 16 hits with the boundary, for a calculated surface area of $0.1371 \text{ square } \mu\text{m per cubic } \mu\text{m}$. The use of the cycloid grid on vertical section images produces measurements that are not biased by possible anisotropy in the structure.



Outline of the white-grey boundary with a superimposed grid and the intersections highlighted.

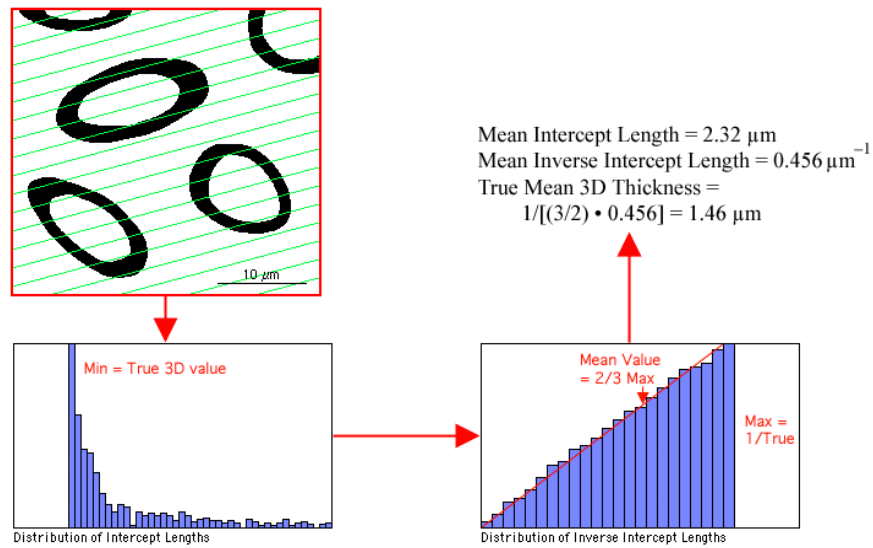
It is also possible to combine human recognition with computer methods by generating the grids in a contrasting color onto an image, selecting a unique pencil color, and marking points. The computer can count the points as well as generating grids and reporting the properties.

Measurement of anisotropy is often performed by generating a set of parallel lines and measuring the mean intercept length as a function of orientation, as shown in the example.



Elongated cell structure and intercept length measurement.

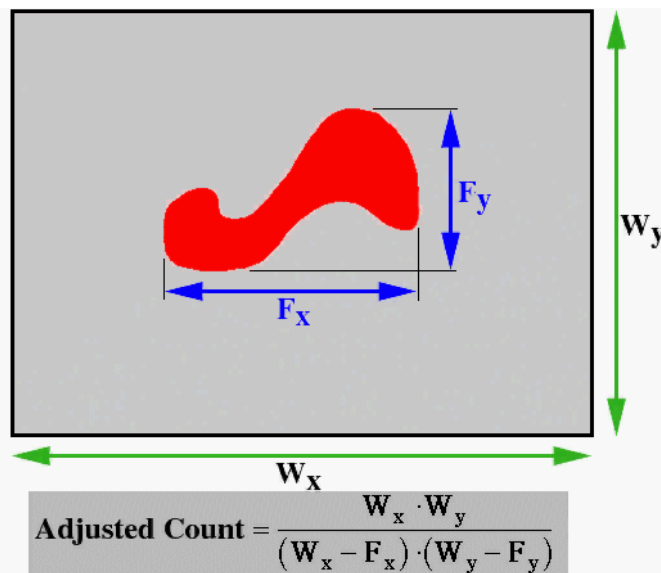
Stereological relationships make it possible to measure some dimensions that are not seen directly in the image. For example, sections through lamellae or tube walls may not be perpendicular, and so the actual thickness is not seen. But a known relationship between the mean inverse intercept length and the true (even unseen) three-dimensional thickness allows it to be determined.



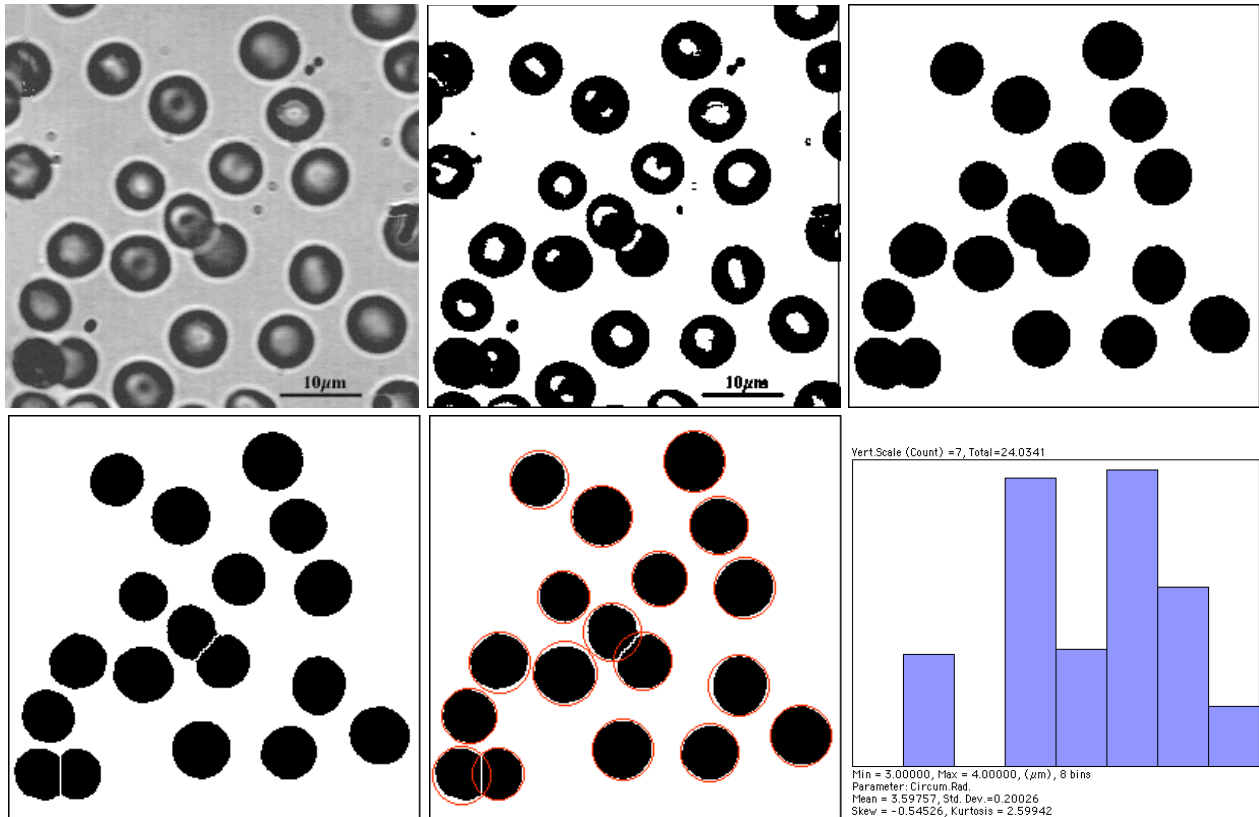
The disector measurement for determining the number of objects per unit volume, irrespective of their size or shape, was shown above.

6.D. Counting features

Counting features in a two dimensional image seems simple, but if the image is a sample of a larger total field, it is necessary to correct for edge effects by counting features that intersect two edges (e.g., right and bottom) and ignoring ones that intersect the other two (top and left). That works for counting but not for measurement, since features that intersect any edge cannot be seen in their entirety and consequently cannot be measured. The solution is to count those features that lie within the image field and can be measured according to an adjusted count. That compensates for the probability that similar features would touch an edge and be unmeasurable.




A typical sequence of operations that count and measure features is shown in the example. Eighteen objects are present within the image area but they represent 24, correcting for edge effects.

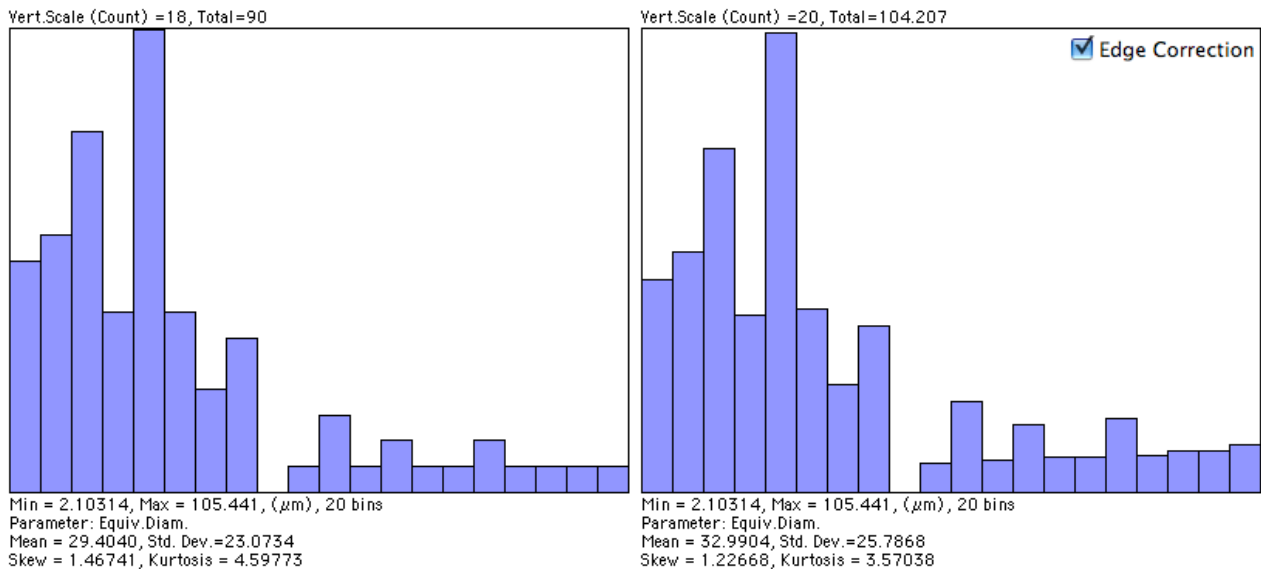


Original image, automatically thresholded binary, filling holes and rejecting small features and ones that intersect any edge, watersheded to separate touching features, measurement of circumscribed circle radius, plot with edge correction.

When feature sizes vary significantly, making the adjusted count correction for edge effects has a significant influence on the results. In the example, size distributions for the features are shown with and without the correction. The mean size changes by 10% and the shape of the distribution is altered.



Original image and binary image obtained by leveling brightness and automatic thresholding.



Measurement results without and with edge correction.

6.E. Measuring features

Feature measurements can be grouped into four categories:

- Size (area, length, breadth, perimeter, etc.)
- Shape (topology, dimensionless ratios, fractal dimension)
- Intensity (density, color)
- Location (absolute coordinates or distances from other objects)

There are many possible choices within each category. As an example of size measurements, the example shows the measurement of the length of grains of rice.


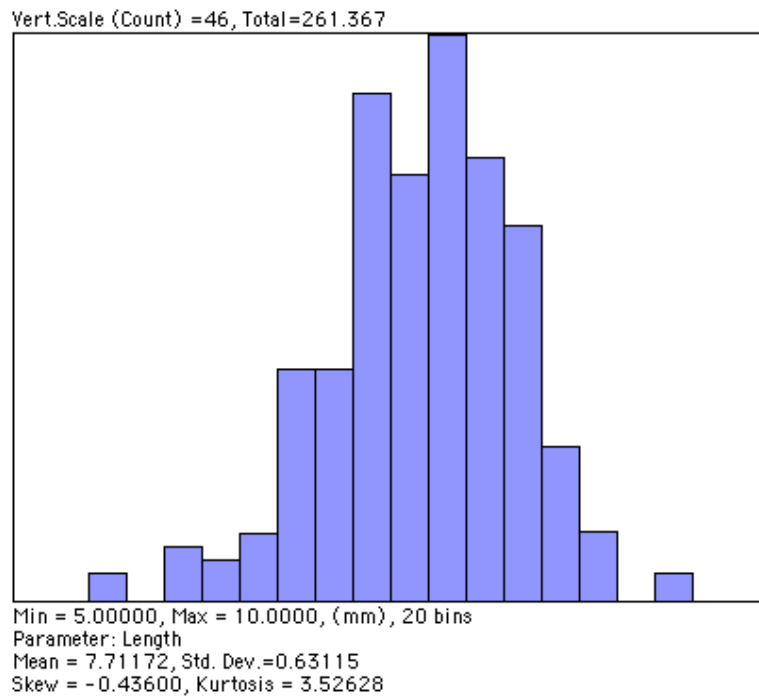


Image of rice grains and thresholded binary.

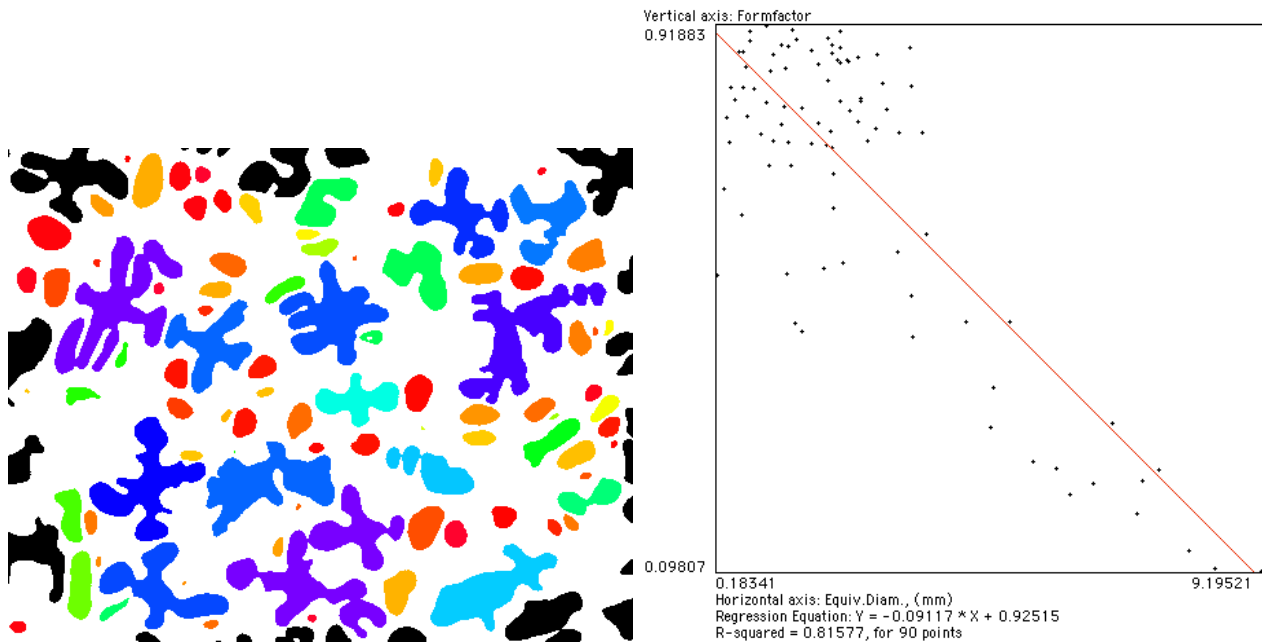


Size distribution of rice grains.

Shape is a more complicated topic than size. A wide variety of shape parameters are defined as formally dimensionless ratios and combinations of size measurements. There is no consistency in the definitions or naming of these, but a few examples are listed below.


- Formfactor = $4 \pi \text{ Area} / \text{Perimeter}^2$
- Roundness = $4 \text{ Area} / \pi \text{ Length}^2$
- Aspect Ratio = Length / Breadth
- Elongation = Fiber Length / Width
- Curl = Length / Fiber Length
- Convexity = Convex Perim / Perimeter
- Solidity = Area / Convex Area
- Compactness = Equiv. Diam./Max Diam.
- Modification Ratio = Inscribed Diam./Max Diam.
- Extent = Area / Bounding Rectangle Area

The formfactor is widely used as it is sensitive to the boundary irregularity of objects. In the example, a regression plot of size (equivalent circular diameter) vs. shape (formfactor) shows a strong correlation.




Features from a previous example color coded according to formfactor, and a plot of formfactor vs. size.

Unfortunately, dimensionless shape parameters such as formfactor are not very complete, nor human-friendly descriptors of shape. For example, each of the following figures has the same measured formfactor, but to a human they appear quite different.




Some different shapes with the same formfactor.

The computer measurements that seem to relate most closely to what people mean by shape are the topological shape, captured by the skeleton shown above with its number of end points and loops, and by the fractal dimension of the boundary. In the example below, several shapes are labeled with their measured fractal dimension, ranging from 1.0 for a regular or Euclidean object upwards toward 1.999 as the boundary becomes increasingly irregular.



Shapes labeled with their fractal dimension.

Measurements of position offer several different possible locations for an object, including the centroid and the center of the circumscribed circle. In many situations, the absolute position of an object (e.g., on a microscope slide or latitude and longitude) is not of as great interest as the distances from other objects. In particular, this can reveal clustering or self-avoidance. An ideally random distribution of objects has a mean nearest-neighbor distance equal to $0.5 \cdot (\text{Area} / \text{Number})^{1/2}$. The ratio of the measured mean nearest neighbor distance to that random value is a measure of the degree of departure from randomness. The example shows several cases, some of which are not easy to judge visually.



Four arrangements of 100 points. From left to right the ratio of measured nearest neighbor distance to that for a random distribution is 0.999, 1.778, 0.544, and 1.204.




Image of lipid droplets in custard, and a plot of nearest neighbor distances. The plot has a Poisson shape as expected for a random distribution, and the ratio of measured to ideal is 1.048.

Measurement data may be analyzed statistically and used to compare samples. It is also used for feature classification. This generally requires measuring enough known objects to build a set of rules for performing the classification. In the example shown, a very simple set of rules can identify the letters from A through E, regardless of the font, size and orientation.

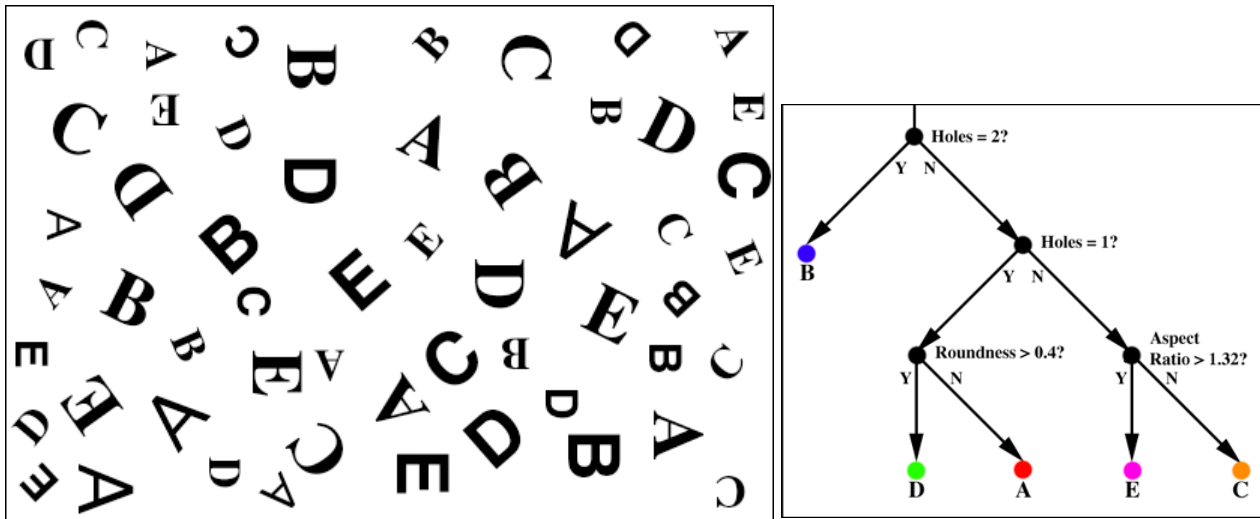


Image of letters in different sizes and fonts, and a set of rules for identifying them.

7. Automation and batch processing

Once a suitable procedure has been worked out for a particular application - the type of image and desired information - it can usually be automated for routine use. Extension Panels for Photoshop CS4 offer organization and automation. Actions and Scripts build automation sequences and support batch processing. Logging of results permits further analysis (e.g., in Excel or a statistics program). Graphics and diagrams can be saved in an html document.