Review of Digital Image Processing

We started the second part of the course with an introduction to digital image processing (DIP) and analysis. These have become exceptionally important topics in recent years as conventional analog-based imaging (e.g., using film) has been replaced by digital approaches. It now is possible to correct many imperfections in research and medical images, to enhance images for aesthetic purposes, and to extract quantitative data, using the computer-based processing and analysis techniques that we discussed.

Image processing and image analysis yield different results. Processing yields an image that is obtained from modification of a precursor image, which often is of inferior quality. Analysis yields numerical data that are extracted from images.

Digital images

We started by introducing the basic attributes of digital images. Digital images have discretely varying properties and are generated from an analog precursor with smoothly varying properties.

The analog-to-digital conversion process involves sampling and quantization, as shown in the schematic below. Sampling transforms the image into a finite, two-dimensional array of "samples." Quantization converts the values in the array into integers with values (e.g., 0 to $2^{12} -1 = 4095$) bounded by the bit depth (e.g., 12) of the detector. Adequate sampling is critical to ensuring that the digital image is an accurate representation of its analog counterpart.

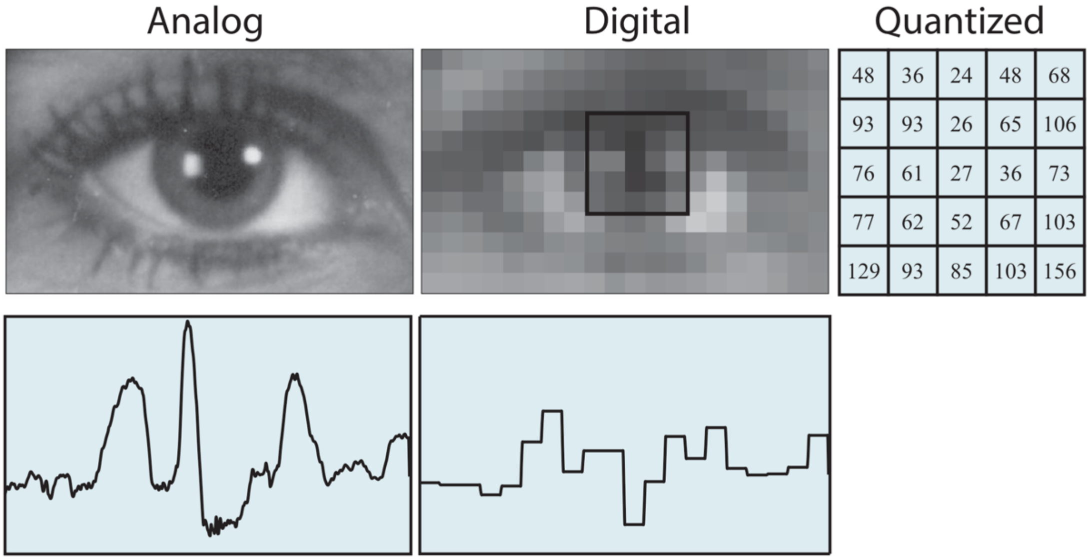

48	36	24	48	68
93	93	26	65	106
76	61	27	36	73
77	62	52	67	103
129	93	85	103	156

Detectors

Detectors used to generate digital images include point detectors, like the photomultiplier tube, which lack spatial sensitivity, and array detectors, like cameras, which are spatially sensitive. In general, detectors convert light into charge and then charge into voltage. Ultimately, the analog voltage is converted into a digital number.

In class, we discussed the charge-coupled device (CCD) and scientific complementary metal-oxide semiconductor (sCMOS) cameras. CCDs have long dominated scientific imaging but have largely been replaced by sCMOS sensors.

The essence of a CCD camera is shown in the schematic below. When light from an object is focused onto the camera, a photon-induced charge is generated in each “pixel” that is proportional to the intensity from the corresponding point in the object. The charge distribution in the detector array thus generates an image of the object.

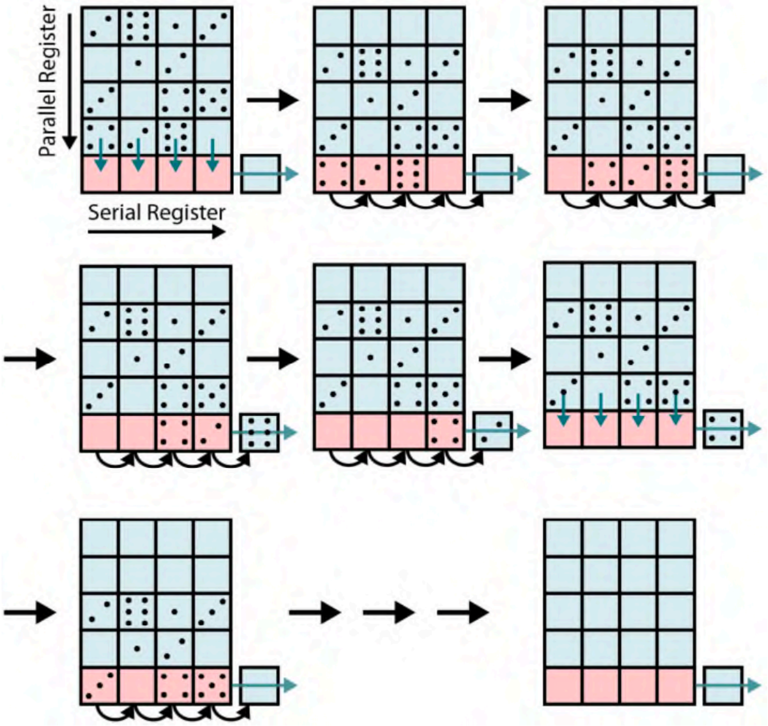

In the CCD camera, photons create “photoelectrons,” which first accumulate in part of the array known as the “parallel register.” After an exposure is complete, voltage sequences are used to move the photoelectrons in the parallel register towards a unique row – the “serial register.” In the serial register, charge is converted into an analog voltage. Finally, an analog-to-digital converter assigns a digital value to each voltage. In the sCMOS camera, charge to voltage conversion (and sometimes digitization) occur at each pixel. This makes sCMOS sensors potentially much faster than CCDs, and this is a major reason for their current dominance in microscopy applications based on cameras.

Image processing – restoration

Most of our discussion of DIP was focused on techniques that are directed at image restoration, which is the process of returning a degraded image to a more idealized, uncorrupted state. For example, restoration corrects defects resulting from imperfections in the illumination, optics, and detector.

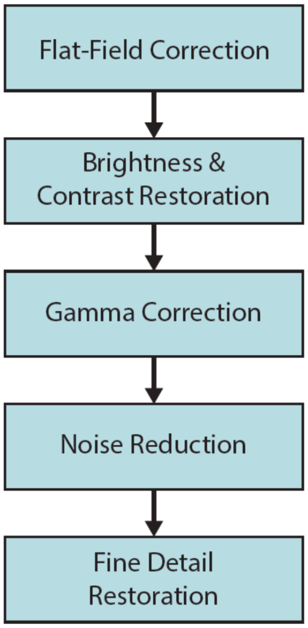

In class, we described an established restoration strategy (see flowchart at right), and you had an opportunity to implement many of these steps on sample images.

The first step in restoration is flat-field correction. We did not practice this process, which is directed at removing degradation from sources like detector imperfections, and uneven illumination.

We did apply the second step, brightness and contrast restoration, to real images. This process is directed at improving brightness (a measure of "lightness") and contrast (a measure of grayscale/color variation).

We discussed two tools – the image histogram and the input/output lookup table (LUT) – that play a key role in brightness and contrast restoration. Both are illustrated in the schematic below. The histogram is a graph of the allowed image intensities versus the number of pixels in the image corresponding to each intensity value. The histogram can be used to assess image quality and to identify steps that will aid in image rehabilitation. For example, an image with a tightly clustered histogram has relatively poor contrast because clustering is indicative of lack of variation. Appropriate corrective steps should lead to a broader histogram.

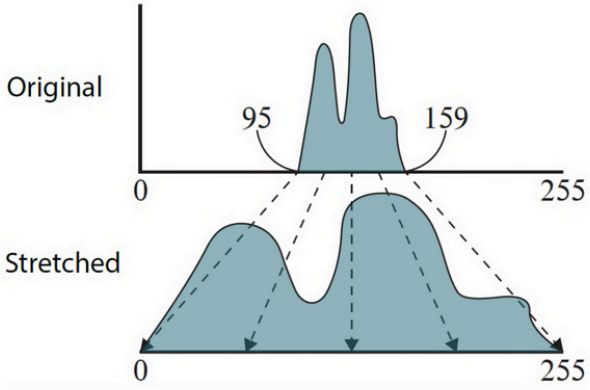

The LUT is a function/rule that transforms an input pixel value into a new output value. A very common LUT (shown in the schematic) is one that maps the image minimum into zero and the image maximum into 255 and ramps linearly between the minimum and maximum. This LUT is used in the contrast-enhancement technique known as histogram stretching, which we implemented on an image that was too bright and that lacked contrast.

The third step in restoration is gamma (γ) correction, which implements a non-linear LUT. Gamma correction is an especially useful method of enhancing dim features in an image without saturating (losing contrast) in brighter features. To achieve this goal, gamma correction is implemented with $\gamma < 1$.

The fourth and fifth steps in image restoration are noise reduction and final detail restoration, respectively. We discussed how both steps can be implemented using filters and applied popular filters to corrupted images. Filters that are used to reduce noise include spatial convolution filters, such as the mean and Gaussian convolution filters, and the median filter. All three filters reduce noise by creating an output pixel value that depends on values from multiple neighboring pixels. In the case of the mean filter, the output is a simple average. The effect is to reduce noise but increase blur. In the case of the median filter, the output is a median of neighboring pixel values. Again, the effect is noise reduction, but fine detail is better preserved by the median filter (at the cost of more extensive computation associated with sorting).

Filters used for fine detail restoration include the popular unsharp filter. This filter is based on generating an unsharp (smoothed) version of the image and then subtracting the smoothed image from the original image. This process removes low-frequency features and generates an image that retains high-frequency, finer-scale features, like edges.

(I skipped color correction this year, so you are not responsible for it. I have included the following two paragraphs in case you want to read them.)

We also discussed color correction, which tends to be necessary when an image is collected under conditions where there is a temperature-induced mismatch between the color output of the light source and the color balance calibration of the detector.

The subject of color correction also led to a discussion of color coordinates, including red, green, blue (RGB) and hue, saturation, and intensity (HSI) coordinates. The latter tend to be less familiar – HSI coordinates include a brightness (intensity) component that is largely decoupled from the color components (hue and saturation). Notably, HSI coordinates often are better suited to color image processing than the more familiar RGB coordinates. This is because HSI can be used to process only the intensity coordinate, and this avoids the introduction of color shifts.

Image processing – enhancement

Images can be enhanced, as well as restored. The distinction between enhancement and restoration is slightly murky but, in essence, enhancement differs from restoration in being a more subjective process that is directed at altering an image to influence impact on the observer. We did not focus as much on enhancement, which can be implemented using many of the same tools that are used to achieve restoration.

Image analysis – single particle tracking

We closed our discussion of DIP and analysis with a brief discussion of analysis, which is an immense topic that is directed at extracting numerical data from images. Specifically, we used single-particle tracking to determine organelle trajectories, speeds, and diffusion coefficients from movies generated using time-lapse fluorescence microscopy.