

A Digital Path to Better Pictures

James R. Milch
Eastman Kodak Company, Rochester, NY

Abstract

Digital image processing offers the photographer new opportunities to create better pictures. After a scene has been captured in digital form, digital processing can selectively modify portions of the picture. The photographer can divide parts to be retained and those to be changed in many different ways. This is particularly important when combining several photographs.

Amateur and commercial photographers will use digital methods differently, but in both cases the most value is derived when the entire work flow, from capture to delivery, is optimized for the new technology. Digital imaging introduces several new issues concerning image quality. However, with proper care, the quality of pictures generated digitally can exceed traditional results.

Introduction

For many years, digital image processing was a wonderful theory, practiced by patient professors with big computers. The advances in silicon technology have put a tool for making better pictures into the hands of every photographer. The power of digital processing stems from a few basic features of digital systems. These features are largely independent of the source of the image or its ultimate destination.

Consumers are just becoming aware of digital photography, but in commercial photography, film scanning and digital processing are already part of the standard work flow. Direct scene capture with a digital camera is growing in popularity in all sectors. No matter how the scene is captured, digital image processing offers new features and better pictures to photographers.

The "Digital" Nature of Digital Imaging

The distribution of illuminance in a scene is an analog phenomenon—it is a continuous function that varies smoothly in space. Most display and visualization media are also analog in nature. Digital imaging represents the image information in a discrete, or digital form. It is useful to think about the digital nature of a digital image in two parts. The first is spatial sampling and the second is intensity quantization.

There is no intrinsic spatial structure in an image, but, on a fine enough scale, neighboring regions have similar intensities. A digital image is created by *sampling* the image on this scale. Each sample, or pixel, is a measurement of intensity over a small region. The pixels are not "little boxes" on the image, but rather local measurements in regions of slowly changing intensity. The pixels are usually placed on a square grid, spaced by the pixel pitch (e.g., 300 pixels per inch).

The sampled values of intensity may take on any value, but the human eye can discern only 100-200 intensity levels. There is no loss incurred by "rounding", or *quantizing* the intensity to the nearest level. In many systems it is necessary to quantize before the right set of levels are known. In these cases, the only solution is to quantize to more levels (sometimes as many as 10000), and reduce the number of levels later. A careful analysis of noise and signal processing steps though the system is required to determine the optimum degree of quantization.

These two steps, sampling and quantizing, convert an analog image into a digital one. For color images, three intensity values (rather than one) are required, but there is no essential change. Stereo imaging, volume imaging, and motion imaging each introduce the need to handle more numbers, but the principles remain the same. Digitization converts a rather ephemeral object, an image, into a very concrete form--a finite collection of numbers, each of which is limited to a small set of integers. (To understand "ephemeral", think about sending an analog image from Boston to Chicago in less than one second, with reasonable fidelity, using neither sampling nor quantization, but purely analog means.)

Essential Problems with Digital Imaging

There are two important disadvantages to working with digital images. These are both "engineering problems", easily cured by careful system design and expensive components. There is no excuse for poor system design, but products are often cost constrained, so these problems are very real today. Equipment purchasers and development managers should not assume that "going digital" will provide a perfect solution—"trust, but verify." However, time is on the side of digital processing: increases in performance of computers, memories, and data transmission networks are lowering costs rapidly.

The first problem with digital imaging is that the spacing of the pixels and quantization levels must be close enough, or image quality will suffer. The solution to this problem is completely obvious—but sometimes impossible because of other system requirements. The real negative impact comes when the user is surprised and disappointed by poor image quality long after the image is captured.

Even if the pixel spacing and quantization levels are perfect, each pixel of the image must be measured and written out individually, leaving many opportunities for visible errors, or *artifacts.* This cause of poor image quality will be discussed further below.

The second problem is that a digital imaging system must store, process, and transmit every element of the image in a serial fashion, rather than simultaneously. Most analog processes (e.g., optical printing) deal with the whole image at once—including all of its color elements. This makes it much easier to build analog systems, which process images rapidly.

The Value of Digital Imaging

The good news about digital imaging can be condensed into two key points. They follow directly from the fact that the digital image is a finite collection of simple numbers. I can only touch on the many implications of these two points in this paper.

You Can Make a *Perfect* Copy of a Digital Image, Anytime, *Anywhere.*

Analog processes are always subject to noise. A copy of an analog image may be almost the same, but it will be subtly changed by the noise inherent in the copying process. Most copying processes also introduce spatial distortions in the image and variations in response as a function of intensity. Digital signals can always be broken down to a small number of physical levels, so that system noise causes no changes. Some digital techniques (e.g., lossy compression) introduce noise or confuse the signal, but this is done in a controlled fashion, to meet other system requirements.

Not only can you copy an image, but you can copy part of an image. That part may be a rectangular region, an oval, a certain color range, or the outline of a physical object in the image. It may also be one color plane, a portion of a Fourier space calculated from the image, or a portion of an image sequence. The selection of the portion to be copied can be automatic or derived from a user interaction with the image.

An integral part of copying images is composing a new image from pieces of other images. Copies of digital images don't "decay". They will sit quietly forever, stored in digital form, ready for use. They are always fresh. Pieces can be combined in many ways—simple overlays, addition, partially opaque addition, or addition under a mask. Images can be composed even if they don't come from the same kind of source. One can add an infrared image of a scene to a visible light image of the same scene to see details better through haze.

Finally, one of the most practical features of digital images is the great variety of methods available for storing and transmitting them. Having copied an image, you can store the numbers, which represent it as bits on a RAM chip, domains in magnetic media, or pits on optical media. You can transmit the bits over switched networks. Best of all, the technology needed to do this has been developed already at great cost for other purposes. Much of the equipment needed has already been installed and amortized. The day soon may come when digital imaging drives advances in digital technology, but so far it has been the other way around.

You Can Apply *Controllable* Changes to *Selected* Parts of a Digital Image.

Controlled modification of analog images is generally difficult. Some changes are easy, but most require specialized equipment and great skill. Also, different kinds of changes require different kinds of specialized equipment—and often different operators. The situation is quite different for digital images. Modifications can be made by manipulating the numbers representing the image according to formulas. As above, you can select just a part of the image and localize the modifications to that part. Some changes have clear precedents in the analog world. Examples of these are color shifts, unsharp masking, and contrast changes. Others are quite novel, such as distorting human heads (kids love it!) and extracting texture.

Compared to the chemical image processing which takes place inside photographic film or paper, digital image processing is flexible, controllable, and smart. It is flexible and controllable because it is independent of the image capture step and can be changed at will. It is smart because the parameters of the calculation can depend on image content. This might be local information (make the grass greener), neighborhood information (do not sharpen noisy edges), or global information (lower the contrast of images shot in full sunlight).

A very pleasing feature of digital image processing is that all of these modifications can be done with the same equipment—a programmable computer or digital signal processor. The use of a computer can make the controlled modification of images more accessible to the casual user. (However, this does not follow automatically!) From the user's point of view, modifications range from "utility" to "art". A utility operation is one that does not change the meaning of an image, but is forced on the user by the imaging system. Consider, for example, "fix the colors up, sharpen the image, and crop it to fit that printer over there." This is drudge work. In the best of worlds, utility operations are done by the system, invisibly to the user. Art is at the other end of the spectrum. An example of this might be, “lift the man out of the first image and drop him into the second image between the trees. Blend him in; add oranges hanging from the tree on the left; add some realistic clouds in the sky.” This work must always be directed interactively by a human. In the best of worlds, art operations are completed instantly and the image on the screen is just what will be printed.

In between utility and art are many intermediate operations. New tools and effects are invented every day. They often effect a complex change in the image under the guidance of the operator. Two popular approaches to this are “click and fix” and “ring-around”. “Click and fix” tools let the operator identify roughly a place in the image to act, and then use some intelligence to do the right thing in exactly the right place. A “ring-around” automatically generates half a dozen different changes, and lets the operator choose the best one. Often, the same correction can be done in a variety of ways. An artist can remove a scratch from an image in a few minutes using standard cloning tools. A semi-automatic scratch removal tool is faster and more reliable; the operator just clicks on the scratch, then the algorithm finds its boundaries and expertly covers it up. Perhaps the user would prefer that the software examine every image for scratches automatically, removing them before anyone even knows they were there.

There are many kinds of changes, which may be applied to varying degrees, to various parts of a digital image. This gives digital image processing tremendous power and flexibility. If the computer is fast and has enough memory, an experienced operator finds her greatest challenge to be keeping track of all the changes and managing their interactions. The coming generation of image processing software will address this issue by hiding more of the utility operations and making the art easier for the novice operator.

The Quality of Digital Images

Most analog pictures are captured for a specific purpose. They are processed and displayed using a well-known, standard procedure. Flexibility is one of the key features of digital imaging. A picture is carried through a chain of system elements. It starts as a scene, captured with a digital camera, or on photographic film, which is then digitized by a film scanner. The resulting digital image can be stored, transmitted, or modified, not just once but many times. Finally, the image is converted to a visible output, or sometimes analyzed for its information content. The effect of each step on image quality is not difficult to control. However, early steps often must be done without knowledge of the later steps. The processing that the image will get later is uncertain. This variability in the image chain makes it difficult to assure that the final picture will have the quality desired.

Digital image quality is a large and diverse subject.[1,2] The following discussion will touch on three topics: information content, color management, and artifacts.

Information Content

Pixels, signal-to-noise ratio, data storage, and transmission bandwidth are getting cheaper, but they are far from free. There is pressure at each step of the imaging chain to discard information. The system designer or digital image user must often answer questions like:

“How much information should I capture from the scene?”

“Will this element of the imaging chain carry enough information to give me a good enough image when I finally want to view it?”

“If I print this image on that printer, will I like the result?”

The theories required to answer these questions are well understood, but there are no general tools available to system designers for obtaining rapid answers. Experienced users have worked out the answers by trial and error, one case at a time. Imaging systems and software have not advanced to the point of relieving inexperienced users of the need to worry about such issues.

Most scenes contain far more information than any photographic or digital system can handle. An imaging system captures “enough” information from the scene, transmits and stores “enough” information, and creates an output image with “enough” information. If the three “enough’s” are roughly equal, and the result meets the customers’ expectations, the system is a success. Thus, the management of information content is a critical issue in the design of a digital imaging system.

The information content in a digital image is not a simple function of the number of bits used to store it. The usual digital image descriptors (512 by 768 pixel image, 24 bits per pixel) set an upper bound for the information content. Few actual images reach the upper bound. Much of the loss is introduced by the device used to digitize the scene (e.g., a digital camera), which introduces noise and smears fine structural details. This truth is captured in the aphorism, “all pixels are not created equal”.

Once the image is digitized, the information content is usually preserved. Sometimes, for artistic effect, a user intentionally discards information by blurring the image. Image compression may, or may not, preserve information content.[3] “Lossless” compression takes advantage of statistical redundancies in the image. It creates a “shorthand” that can represent that image with fewer (one-half to one-third as many) bytes of data. “Lossy” compression goes farther. This method analyzes the image and extracts the information, which is most visibly relevant. Then, it applies lossless compression to that information and discards the rest.

In the abstract, there is no way to calculate how much lossy compression is acceptable. For a specific system, with the input, output, and range of image processing specified, the effect of compression can be studied by trial and error. In this context, the creative expression "visually lossless compression" has been coined to describe compression that discards information, but so little that the user sees no effect.

Color Management

Copying and printing images is commonplace in the world of photography and printing. Because analog images are inherently local objects, matching the colors in the copy to those in the original can be done by inspection. Early applications of digital imaging fit into this same work flow, so the original image is usually at hand. More recently, the ability to transmit digital images has changed the work flow. Now it is common to deliver only a collection of numbers with the request for a print and to expect that the colors in the print will be right! First and foremost, this requires a way to interpret the meaning of the numbers in the digital image. Just as critical, there must be a simple means to convert those numbers to the desired colors on the print. These two points are the essence of color management.

People often assume that there is a standard way to represent the gradations of tone and color in a digital image. There are standards—many of them—each of which makes excellent sense in solving the specific problem for which it was designed. The simple coda "use three numbers to describe the colors a human eye can see" dissolves into hopeless confusion when practical issues, such as scene dynamic range, material gamuts, named colors, quantization error, and image compression, are taken into account.

The International Color Consortium has provided part of the solution. This voluntary group of companies who provide and use digital color has introduced a standard method for specifying the meaning of digital color and for translating the numbers in an image file to those needed for a specific printer. The rest of the solution must be built into an imaging system, making device calibration and characterization simple and routine. This is starting to occur.

Artifact-Free Reproduction

A digital imaging system that is designed well and implemented properly produces excellent images. Some design errors reduce image sharpness, dynamic range, or colorfulness. These effects are usually predictable. Other problems have a greater impact: they introduce detail into the image that really should not be there. These false details are called *artifacts*. They may appear as spots, shadows, edges, streaks, bands, mottle, worms, blocks, or rainbows.

No one likes artifacts. Product development teams do not like them because they always appear late in the program, after the design is locked down. They are rarely found in the images used for routine system testing. Customers do not like artifacts because they come and go. Usually they come just before the beginning of a major job and go when the service representative appears.

Artifacts are usually caused by small variations in component performance, which sometimes add up to a visible flaw in the image. Artifacts will, in the end, succumb to good system analysis and careful engineering, but they may cause long delays in an engineering project and an escalation in product cost.

Prospect

Digital photography is a young technology. Today, most digital pictures start as scenes captured on AgX film. By scanning the film or perhaps a print, the photographer can tap into new capabilities for improving or customizing the picture. Digital cameras are advancing rapidly in performance and declining in price. They provide unique advantages, particularly for the professional photographer. Digital cameras will penetrate consumer photography more slowly, as the broad infrastructure needed to support them takes root. In both areas, the real payoff of direct digital capture comes only when the user changes his picture-taking habits to take advantage of its unique features.

The image processing software available today is powerful, but rather primitive. The user needs a rather complete understanding of where the image came from, which operations to select, and how to prepare the image to get a good print. This is changing rapidly, as image-capable computers and imaging peripherals reach reasonable price levels. Lower prices create a broad market, justifying the investment in better software. Concepts such as resolution-independent imaging and just-in-time[4] are migrating from high-end, specialized systems to desktop software.

The silicon revolution has put the patient professors on the sideline and created new excitement in the world of photography. Silver and silicon will both find a place in this world. Together, they will put new tools into the hands of everyone who wants to record a moment in time and recreate the events, feelings, or meaning of that moment.

References

1. J. A. J. Roufs, "Perceptual Image Quality: Concept and Measurement", *Phillips J. of Res.*, **47 (1)**, pp. 35-62 (1992).
2. R. E. Jacobson, "An Evaluation of Image Quality Metrics", *Royal Photographic Society Symposium Proceedings Photographic Image Quality*, pp. 7-16 (Sept 1994).
3. M. Rabbani and P. W. Jones, *Digital Image Compression Techniques*, SPIE Tutorial Text Series, **TT7**, Bellingham, WA (1991).
4. J. R. Milch, "A Storage Format for Resolution Independent Imaging Systems", *Proceedings of ISEP*, (1996).