THE FUTURE IMAGE REPORT

THE FUTURE IMAGE REPORT BUYER'S GUIDE TO DIGITAL POINT-AND-SHOOT CAMERAS

Point-And-Shoot Goes Digital: Filmless Images For Everyone.

By R. David Etchells

There's a revolution happening in photography – the filmless revolution! Digital imaging technology is moving farther and farther into territory once the exclusive province of film. The advantages are numerous: Images are immediately available to incorporate into reports, presentations or other documents. Operating costs are virtually zero. (No film or processing, remember?) Images can be stored and transmitted electronically. Finally, digital images can be manipulated using powerful software tools to enhance their appearance or clarify the information they contain.

Until recently, most of the activity in this filmless revolution has been happening in high-end cameras for studio photography. With price tags ranging from $10,000 to $50,000, you clearly need a pretty serious application in order to take the plunge for one of those! Now though, a number of developments at the "low end" have led to very capable devices in the "around $1,000" category. While clearly lacking some of the features and image quality of their big brothers, these more affordable devices are opening up broad new application areas for digital imaging.

On this disk, we'll take a look at the capabilities and applications of the new generation of portable, "Point-and-shoot" digital cameras. And we'll take a first-hand look at the latest models from all the key manufacturers.

What is a Digital Point-and-shoot Camera?

We call the cameras we'll be reviewing here "Point-and-shoot" because that's about all it takes to use them: This is a designation long used to describe a range of conventional film-based cameras designed for low cost and ease of use by nonprofessional photographers. The basic idea is a camera that essentially runs itself: The photographer need only choose the subject, point the camera, and push the button.

In the film-based world, some Point-and-shoots have sophisticated features such as zoom lenses and automatic focusing. The majority of models sold though, are very simple devices with "fixed focus" and a non-zoom lens. So it is with most digital devices in this category. The devices tested on this disk generally have fixed-focus lenses, holding good focus over a range from roughly 4 feet to infinity, although some come equipped with a close-up accessory. Most have built-in electronic flashes, and variable-speed electronic shutters. A few models have zoom-lenses and/or autofocus, and we can doubtless expect to see more such devices in the future.

With the digital devices, the relative paucity of features reflects the price constraints manufacturers face. The sensors and other necessary electronic components make even the most basic devices fairly expensive. If more advanced features were added, the prices would head beyond $1,000 and affordability would suffer. As electronic components continue their inexorable downward march on the price curve though, it is likely we'll see both less expensive entry-level units as well as more feature-packed models.

For now, the main differences between competing devices have to do with the flexibility of the lens systems, resolution and image quality, amount of on-board image storage, battery life, flash flexibility, and ease of interfacing. We'll explore each of these parameters in more detail later.

What can I do with a Digital Point-and-shoot?

While applications for Digital Point-and-shoot (DPS) cameras are numerous and varied, it's important to consider their limitations as well as their capabilities. There's no question their resolution and image quality are limited, as well as their focusing and focal length flexibility. Thus, you clearly would *not* use a DPS to create a high-quality marketing brochure, a full-size poster image, or a digital portrait meant to be hung on a wall. On the other hand, their convenience is unmatched for applications where image sizes will generally be small, or where image clarity isn't the foremost concern.

To understand different application areas, we find it useful to think in terms of three different uses of images, which we'll outline below:

Images as Data

Are you looking for a camera mainly as a way to capture essential information about people, objects, or situations? (After all, a picture *is* still worth a thousand words!) If you are, then your needs fit into the "Images as Data" category. Digital Point-and-shoot cameras are finding widespread application in real estate marketing, identification (using image databases to record and store images of people or objects), insurance documentation (again, linked to a digital image database system), and medical/dental environments. We've heard of at least one oral surgeon using a DPS camera with close-up and special flash attachments to document before/after pictures of patients' mouths. The low cost and immediate feedback available with DPS devices will undoubtedly lead to other, similar uses in the near future. We address several of these uses in our test images, contained elsewhere on this CD.

Images as Communication

Do you want to use the "instant images" of digital photography to send information to others you are working with, or increase the impact of your business presentations? If so, you fit into the category of "Images as Communication." While one might argue this is really just "images as data" in another guise, the usage here is different. An example would be a clothing manufacturer and his overseas client, emailing back and forth annotated images of a sample to make sure they both understand exactly how the final product should turn out. The image is a key part of the communication flow in a collaborative decision making or other work process, where it does indeed replace a thousand words.

Images have never been common in business communications, no doubt at least in part because they were difficult to include. With DPS cameras, modern photo-capable laser printers, and image-handling capability built into most word processors, things could begin to change dramatically. As digital cameras become more common, and better integrated with desktop business tools, images will also become a larger and larger part of business communications.

In many companies, digital images are already beginning to show up in trip reports (just how big *was* the crowd around the trade show booth?), field investigation reports, internal product announcements and so on. Now that DPS camera prices are dropping below capital-authorization thresholds, more and more companies are taking advantage of the existing image-handling capabilities of their laser printers and word processors.

Images as Memories
This is really the consumer market (YOU, at home), where conventional Point-and-shoot cameras are common. DPS cameras are definitely an appealing gadget for the well-heeled enthusiast. Many home users already have color inkjet printers, and the pleasure of cranking out pictures of family and friends from your own printer is pretty attractive. (Not to mention pictures of class field trips, school plays, family trips, etc.) While a little pricey for many budgets, prices are steadily declining, and we expect a lot of DPS cameras will end up as adjuncts to the family computer system.

Images as Business
If you run a service bureau, a DPS camera could be a supremely logical addition to your service offerings (look at the recent Kinko's/Apple deal with QuickTake cameras and color LaserWriters.) Having a DPS camera available for customer rental not only provides revenue on its own, but can drive demand for other imaging services. Again, the cost of the devices is low enough that virtually *any* shop can afford one.

Why not just use a Video Camera?

Many people are confused by the distinction between a *digital* camera, and a conventional video camera connected to a *digitizer*. Both systems leave you with digital files on your computer, so what's the difference? The answer is really twofold: 1) convenience and 2) image quality.

The convenience factor may not be obvious at first glance. Certainly, you can carry a small camcorder around about as easily as most (but not all!) DPS cameras, so portability isn't a huge issue. The big difference with a DPS camera is that once you've captured the images into the camera, you're pretty well done: Getting them out of the camera is a fairly routine, mechanical process. With a video camera however, not quite. First your computer needs to be equipped with video digitizing hardware (we'll explain this in greater detail below); but furthermore it's as though you have to capture the image twice. First, you capture a brief motion sequence onto tape. Then, when you get back to your computer, you need to spool through the tape to acquire just the specific frames you want to use as still images. This can be a time-consuming process, as you hunt back and forth on the tape to find the sections containing the images you want. So-called "still video" devices, the earliest electronic still image cameras, were developed to address this issue. A few are still available on the market, but digital cameras have gained the upper hand, and virtually obsoleted still video. That's because both still and motion video devices require computers equipped with digitizing hardware, and both are subject to image quality limitations for the following reason.

There are significant differences in image quality between DPS cameras and their video cousins, particularly camcorders. The basic problem is that the image data in a video camera is output as an analog signal, covering a very limited *dynamic range* (the difference between the highest and lowest possible signal values.) The constrained dynamic range of video signals usually means video images have a fairly limited range of tonal gradation. Video signals are also limited in the amount of detail they can carry, a problem made worse when the signal is recorded on tape. The reasons for this are beyond this brief discussion, but the basic fact is that an image encoded into a video signal will almost always contain less information than one taken directly from the sensor.

All these factors add up to a dramatic difference in both convenience and image quality between most digital cameras and digitized video.

Making the Choice – What to Look For

As simple as they are, choosing the right DPS camera involves a number of decisions. Key issues are computing platform, image quality, on-board image capacity, exposure versatility, feature set, and included software. We'll lead you through each of these areas in this section.

Computing Platform
This is an obvious first choice, but is becoming less of an issue than it once was. Most of the cameras we tested could be interfaced to both Mac and Windows platforms.

Image Quality
Image quality is a complex subject involving exposure accuracy, color purity, optics, and image compression techniques. We'll cover this whole area in greater detail later. For now, the main thing to keep in mind is that there's really no substitute for direct experience: Look at the sample images on this disk that correspond to your expected usage – you may be surprised by the differences between cameras!

On-Board Image Capacity
Depending on your application, the amount of on-board image storage could be pretty important. There is generally a tradeoff between image capacity and image quality however, so don't get too excited by a large image capacity until you check to see how much actual memory the camera contains compared to the resolution of the images. One camera may promise storage for twice as many images as another, yet have no more actual storage space. The consequence would likely be a loss of image quality. On-board camera memory capacities range from 1 to 4 megabytes, while image storage can vary from 8 to 96 images. Again, look at the sample images on this disk to decide what the *usable* storage capacity of each device would be for your application.

Several cameras allow you to make your own decision about the memory/cost tradeoff, by allowing you to store images on plug-in cards. They generally use industry-standard "PC" (formerly PCMCIA) memory cards to extend the on-board memory. This approach allows each user to add as much or as little memory to the camera as they feel they need.

Exposure Versatility
One notable *dis*advantage of digital cameras is their inability to handle radically different lighting conditions just by using a different type of film: With a digital camera, the camera *is* the film, so whatever you buy is what you'll have to live with. This means you'll want to be sure your camera's exposure system (shutter, variable lens opening, supplementary flash) can handle a wide enough range of conditions.

In this category, look at a camera's "equivalent ISO number," which is a measure of how sensitive the camera's sensor is to light. A higher number means the camera can successfully capture images with less light. Most DPS cameras have equivalent ISO ratings in the 80-200 range.

Beyond the ISO rating, look for a wide range of shutter speeds and lens apertures. Typical numbers here are shutter speeds from 1/30 to 1/200 of a second, and lens openings from f2.8 to f16. The wider the range spanned by both these sets of numbers the better. For the most part, we found the cameras tested did poorly in low-light situations. See the Portrait/Outdoor images for examples.

Also look at the capability of the built-in flash if one exists. Key factors here are the distance range over which the flash will be useful, and the number of "modes" it has. You'll generally find fairly limited flash power, evidenced by maximum working distances of as little as 10 feet. Flash "modes" refer to different ways the camera's light sensors can control the flash circuitry. In the most basic mode, the camera "stops down" (closes) its lens aperture and runs the flash at full power so most of the light hitting the subject is provided by the flash. Some cameras offer a "fill" mode for their flash. In this mode, the camera sets the shutter speed and lens opening based on the amount of light coming from the scene overall. In "fill" mode, only enough flash illumination is used to "fill in" the shadows and bring the illumination up to the minimum required. The result is often a more natural, evenly-lit subject.

Feature Set
Besides the basic parameters listed above, you'll find a variety of features on the cameras that may have varying degrees of usefulness in your application:

Close-Up and Wide-Angle
The ability to take close-up pictures is critical to many applications. Many DPS cameras have a "fixed-focus" lens, in which the focal point is set so everything from about 4 feet on out is in focus. This is handy if you need to photograph subjects in that range, but pretty useless if you have to take a picture of something only a few inches wide. Some DPS units have optional attachments to improve their close-up performance: The Apple QuickTake includes a clever snap-on adapter that simultaneously shortens the focus of the lens, adjusts the field of view of the viewfinder, and angles the light from the flash to point more directly at the close-in subject. The Kodak DCS 40, Logitech Pixtura, and Dycam C-10 each have built-in adapter rings allowing them to accept standard photographic close-up lenses sold by third parties. All of the Chinon-based cameras (Chinon,

Dycam C-10, Ritz Dakota) have a "macro" mode available in their auto-focusing system for real close-up work.

Wide-angle photography is another important ability if you need it. (This could be particularly important for real estate or insurance photography, where shooting conditions may be cramped.) Again, some cameras have aftermarket accessories available to expand their field of view, so you will want to take this into account when purchasing. (We're aware of products for both the Apple and Kodak units, products for other cameras may follow.)

Several cameras feature zoom lenses, with a range of focal lengths running from a mild wide-angle (roughly equivalent to a 35mm lens on a 35mm camera) out to a mild telephoto (slightly more than the equivalent of a 100mm lens on a 35mm camera). After spending quite a bit of time with all the cameras, we found the zoom capability especially useful: Expect to find this feature showing up on more and more cameras in the future.

Battery Life

We're not sure if this is a "feature," but it's important, and this seemed a logical place to include it. You'll find a wide variation between cameras in how long their batteries last. Take note of whether the battery life is specified with standard alkaline cells, or lithium cells: The lithium cells last *much* longer, but cost *much* more, and may be harder to find. Also, can the camera run off rechargeable batteries? If you plan to take a lot of pictures, the ability to use rechargeable cells could save significant money in the long run. (Note too, that the mere presence of an external power adapter may not mean the camera can use rechargeables, and probably does *not* mean the unit will recharge batteries while it is plugged in.)

External "booster" battery packs are also available for some cameras, increasing their working time in the field manyfold. Keep this in mind if you plan to spend long periods on the road, far from power outlets. If a "booster" pack isn't offered, can you change or replace batteries in the field?

External Power Adapter

While we're talking about power, how about an external power adapter? You'll find DPS cameras generally consume much more power while communicating with the computer. If you have an option to run from a power adapter while downloading images, your batteries will last a *lot* longer. It's important to note here though, to be *certain* to use the correct power adapter with you camera – the wrong one could easily fry a thousand-dollar investment!

Self-Timer

A self-timer is simply a shutter delay which allows you to trigger the camera, then run around to be included in the picture before the shutter snaps. This is an easy addition from the standpoint of the technology required, but often omitted. If you plan to be part of any group photos, see that your camera has this feature!

Included Software

Image Download

The software included with the camera can be pretty important. In fact, this is an area where you'll find some of the greatest variation between cameras. In our reviews, we'll try to give you some idea of how the different units fare in this respect. Things to consider include how convenient it is to unload images from the camera. Can an entire camera-full of images be pulled onto your hard drive quickly for later processing? Or, is the camera held captive while each image is individually processed? Alternatively, can "thumbnail" images be viewed prior to downloading the full-size files? - This can save a lot of time in culling the few best shots from a large group.

Image Manipulation

Getting the images into the computer is just the beginning: Does the included software allow you to manipulate the images after you've captured them? How easily can you compensate for poor exposures, color casts, or misaligned images? How about minor retouching to remove blemishes, errant reflections, etc.?

File Formats

Particularly important is the software's file-export capability: Can the furnished software easily provide images in the file format you need? Your requirements may be different if you plan to capture images for multimedia presentations rather than printed output. Most camera software is fairly competent in this respect, with some manufacturers increasing their present level of capability. Virtually all packages provide for export of standard RGB TIFF files, which most document layout and word-processing applications will support as well. If you need formats beyond the standard TIFF, look into what the camera software supports more carefully.

Cataloging

One of the first things you'll learn once you start using a digital camera is how fast images pile up! (If you haven't already, you'll doubtless find yourself buying a *much* larger hard drive.) Having adequate storage space is only half the battle though: *Finding* images is another matter entirely! In recognition of this fact, some cameras include at least rudimentary image cataloging capability in their software package. (On the Mac, just the ability to create a "thumbnail" preview to display as the file icon is a big plus. A word of warning though: Hundreds of file icons with thumbnail images attached can really slow down access to folders on your hard drive. Turning off the "preview" option and using a cataloging program instead can really help system performance.)

Some camera software includes rudimentary cataloging capability, but in our experience, none of these packages offer the level of capability you'd need to manage more than a few dozen files. Accordingly, if you need to track large numbers of images, you should really look at one of the many image-database programs on the market.

Links to External Applications

One of the nicest characteristics of "desktop" applications is how easily you can move data or design elements from one application to another. The Macintosh has excelled in this respect for years, and the Windows platform is rapidly coming along. There are varying degrees of integration available though. Some software requires you to save a file to disk in one application before it can be imported into another. Other packages let you place an element from one application directly into a document of another. Sometimes, you can simply "drag and drop" an image from the camera software into a page layout program. Look at the camera software with your specific application in mind, and see what's required to move the images from the camera to the programs where you'll finally use them.

Image Quality

Image quality is one of the most important characteristics to consider in choosing a camera, and is impossible to evaluate from manufacturers' spec sheets. This underscores the importance of the test images on this disk: Ultimately, the only way to tell if a given camera will produce acceptable results in *your* application is to look at and try out images of similar subjects.

The reason image quality is so hard to get a handle on is twofold. First, there is currently no standard, objective scale by which to measure color accuracy in digital cameras. Second, all DPS devices use image compression technology to cram images into their limited memory. Depending on the algorithm used and the amount of compression applied, image quality can vary widely, even between devices using the same CCD sensor!

Let's look at some of the elements that contribute to image quality:

Color Quality

Color quality is a complex and generally misunderstood topic. While it is probably less an issue for typical Point-and-shoot applications than for high-end studio cameras, the large differences we found between the cameras tested suggest a detailed discussion of color quality would be useful.

Color quality is actually made up of two related but different parameters – *color purity* and *tonal balance*. People tend to assume color errors can be corrected easily in an image-editing program such as Photoshop, and so tend to discount their importance. This is generally true of *tonal* errors, but errors due to color *purity* are virtually impossible to fix.

Gray Balance and Color Accuracy (Tonal Errors)

As mentioned above, *tonal* errors in a digital image are fairly easily compensated for, often with very gratifying results. Of these, *gray balance* is both the most dramatic and the most easily corrected. While a full treatment of gray balance is beyond this particular discussion, we can nonetheless cover a few key points, and see the effect relatively simple adjustments can have. For a brief lesson on how to perform gray balance corrections in Photoshop, refer to the accompanying tutorial (entitled "Black, White and Gray: The Keys to Better Color") on this CD.

As the name suggests, gray balance refers to balancing the red, green and blue channels across the entire tonal range, bringing gray tones in the image into neutral color balance. Obviously, if the blue channel of a camera responded more than the red or green channels, pictures would have an overall blue cast. While gray-balance problems are sometimes as simple as this, generally they are more complex.

THE DISCERNING EYE.

A long but useful diversion. . . .

To understand color accuracy (and why I prefer to refer to it as "color purity"), let's look at how a digital camera sees color. At their heart, digital image input devices try to match the operation of the human eye. It's virtually impossible for any one colored object to mimic *exactly* the spectral characteristics of any other. "White" light consists of a continuous spectrum of component wavelengths or colors, and an object's color in the most absolute sense is the product of how strongly it reflects or absorbs each of these wavelengths. Devices that measure the full spectral characteristics of objects are called "spectrophotometers," and actually measure an object's response to light at many different individual wavelengths. The output of such a device is a curve like that of Figure 1, which shows a sample spectral response curve for a color we would call "red." Note that the curve is highest in the red portions of the spectrum, tapering off in the green and blue. This means the object is reflecting the most light in the red wavelengths, and proportionately less in the green and blue portions of the spectrum. You can readily understand how difficult it would be to match this exact response curve just by mixing various pigments together, each of which had its own unique color "signature!"

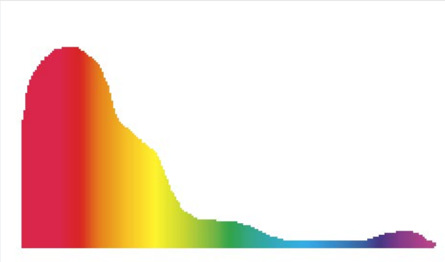

Figure 1

Fortunately, the human visual system allows us a convenient shortcut: Our eyes have only three types of color sensors (except for a very few people who actually have four, but that's a story for another time). Each of these sensor types is sensitive to a different portion of the color spectrum. Signals from all three types of sensors are combined by our brain to produce our final perception of color. The exact response curve for each human sensor type is complex, and will vary from person to person, but in a very general sense, they respond to red, green or blue light. There's a fair amount of crossover among the different sensors, but the red-green-blue breakdown is a pretty good approximation.

For our purposes, the important fact about the eye's color sensors is that they really don't know exactly what color of light is shining on them; all our brain can tell is that the green sensor, for instance, is sending it a signal of a particular strength. That signal could be produced by any of an almost infinite combination of colors of light at different intensities.

So what does this mean? In essence, it means we can trick the brain into seeing almost any color we want it to, just by mixing varying amounts of red, green or blue light together. This is the basis of the entire world of RGB scanners, computer displays, color television and, of course, digital cameras.

If we'll conveniently forget for the moment about the whole world of CMYK color (as is used on the printed page), we can see that a digital camera's job is to figure out what percentages of red, green and blue light to mix together for our eye/brain to think it's seeing the original color. All we need to do is take three separate exposures of a scene through red, green and blue filters in turn. This will tell us how much of each color we need to mix together to fool the eye into "seeing" the original color. Sounds pretty simple, doesn't it?

The problem in practice is that the spectral curves of color filters don't exactly match the curves of the sensors in the human eye. This means that they will respond differently from the way our eyes would to any given mixture of color, yielding incorrect combinations of red, green and blue. What's worse, most commercial filters "leak" varying amounts of the colors they are supposed to be filtering out. For example a red filter may, in fact, let moderate amounts of green light through, and the green filter might do the same with red light. The technical term for this is "cross-channel contamination," and it means that the RGB numbers our camera gives us may not be accurate, depending on which color the camera is looking at.

The color plate in Figure 2 shows what happens when this occurs. On the left is a fairly pure yellow. On the right is what this color would look like if the blue filter of the camera "leaked" significant amounts of red and green light (the "primary" colors that make up yellow). The camera has told the computer, "Hey, there's 10% blue here!" because that's how much red and green light leaked onto the blue sensor. As you can see, the result is a dull, muddy color that looks nothing like what we were hoping for.

Figure 2

If we want to get back to our original color, you might be tempted to say, "Well, let's just cut back the blue channel in Photoshop." This would be fine if the only object in the picture happened to be this particular shade of yellow. The problem is that a reduction in the blue channel will affect *everything* in the image, and will throw our other colors way out of balance.

What we've just learned is that an error in color purity *cannot* be adjusted with simple manipulations in Photoshop: Since the color errors are themselves a function of what color we're looking at, each color in the image needs to have its own, separate correction applied to it! This underscores the importance of having good color purity in any image-acquisition device. This is truly a case of GIGO (garbage in, garbage out), and you can waste huge amounts of time in image-editing applications trying to fix the problems.

Figure 3

In a digital camera, we care about color across the entire range of tonality, which means we need to concern ourselves with the gray balance in the shadows, the highlights and all points in between.

In practice, it is neither practical nor necessary to make gray-balance adjustments point by point across the whole tonal range. Usually, adjusting the balance in the highlights, shadows and midtones brings dramatic results with relatively little effort.

In Photoshop, these adjustments can be made with the Image/Adjust/Levels function, and correction curves can be saved for fairly automated application to groups of images captured with similar lighting. The color plates of Figure 3 show sample results. The image on the left is "as captured" by one of the cameras in the test. The image on the right is derived from the same file, but has had a rudimentary gray-balance adjustment performed on it. Note how much "cleaner" and brighter the colors are, almost as if a layer of grime had been wiped from the image.

While gray-balance adjustments can produce a dramatic improvement in color quality, they are by no means a panacea. While proper gray balance can substantially *brighten* colors, it will *not* correct color impurities or the sort discussed in the sidebar, resulting from poor-quality color filters. You also need to keep in mind that time spent adjusting gray balance is time lost to more productive work. A camera that looks like a bargain may prove otherwise if you have to spend 20 minutes color-correcting every image it produces.

Evaluating color purity and color accuracy.

Important as it is, color accuracy is a slippery parameter to quantify. Rather than trying for some sort of absolute standard for color accuracy (which would likely be difficult to interpret anyway), we opted here simply to provide consistently exposed test images that would help you make your own subjective evaluations. The "Dave Box" target includes a Macbeth™ color chart as a well-established and reasonably consistent color reference, readily available at most well-stocked camera stores.

We've also provided another comparison aid in the form of a large grid of color chips taken from each camera's test shot. We have found such grids very useful in our own evaluations, as we can use Photoshop's cut-&-paste capability to perform side-by-side comparisons very easily. To compare the color chips from one camera to those of the others, use Photoshop's marquee tool to select the row of chips of interest. Once you've selected them, holding down the option key while dragging the selection will make a "floating" copy of it. You can then click and drag the selection to position it next to the row of chips from any of the other cameras to get a very good idea of how well the different devices handle different colors. You'll want to crop the color chips very closely, to eliminate any black border around them. This lets you see the colors from one camera directly adjacent to those from another, making even very slight color shifts obvious.

Dynamic range

Dynamic range measures how wide a range of subject luminance (*e.g.,* brightness) a sensor can accurately reproduce. Most digital cameras can capture a wider range of brightness values than the printed page can reproduce. The key issue is how good a job the hardware and software does of compressing the full range of scene brightness into a range that output devices can reproduce. This capability, whether in a camera or a high-end scanner, is called tone compression, and it is an important characteristic of high-quality input devices.

Tone compression is as much an art as a science, and different devices use different input-to-output curves to accomplish it. Generally, you want to avoid losing either the highlights or the shadows, yet still maintain reasonable contrast in the midtones. Visually, photographers are accustomed to looking for tonal problems in very light and very dark subjects. Traditionally, people shoot white porcelain objects to study subtle highlight detail and dark camera bodies or electronic equipment to look at shadow detail. This is a valid approach, but subtle differences are often difficult to detect on the final printed output. On screen, without an assist from Photoshop, the problem is even greater. Especially in the shadow areas, the tonal response of CRTs is very poor.

The way to really see what's going on is to use the Photoshop Levels control to stretch the tonal range of the image in the areas you're interested in. For highlights, this is accomplished by moving the black-point slider well up into the midtones. This forces everything darker than the midtones on down to black, stretching the remaining tones over a wider range and emphasizing tonal differences in the brighter portions of the image. For shadow areas, reverse the procedure, moving the white-point slider down into the midtones. This forces everything from the midtones on up to white to go all the way to white, stretching the shadow tones across the full range and emphasizing subtle tonal differences there as well. Figures 4 – 6 show this process applied to an image with deep shadow. In Figure 4 we see the display from a Photoshop Levels adjustment. Note that all the image data are clumped on the far left, in the deep shadow. The illustration immediately above it shows the corresponding image; there's no apparent detail to be seen. In Figure 5, we've moved the highlight slider far to the left until it's just touching the right-hand edge of the histogram "lump." See how this brightens the screen display, pulling up detail where none was to be seen. Finally, Figure 6 shows how the brightness levels of the image have been "stretched" to cover the full tonal range. The image is the same as in Figure 5; the change has just been made permanent in the file. Note that if the gray balance of the camera is inaccurate, these procedures may also show color casts in either the highlights or the shadows that wouldn't be obvious to the unaided eye. (In other words, don't be surprised if your shadows end up looking yellow or red or whatever after you've tinkered with them to this extent.)

What do you look for? Whether in the highlights or shadows, you want a camera to reproduce tonal variations smoothly, without obvious "tonal breaks," "posterizing" or "quantization" (three different terms all meaning the same thing). The most common defects occur when highlights wash out to white or shadows plug up, going black in regions where there is still some tonal variation in the subject that you'd like to retain. These faults are sometimes hard to see, but Photoshop will smoke them out every time.

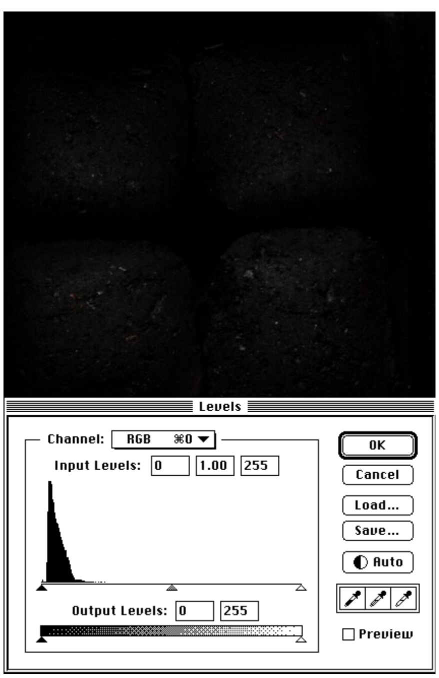

Figure 4

Figure 5

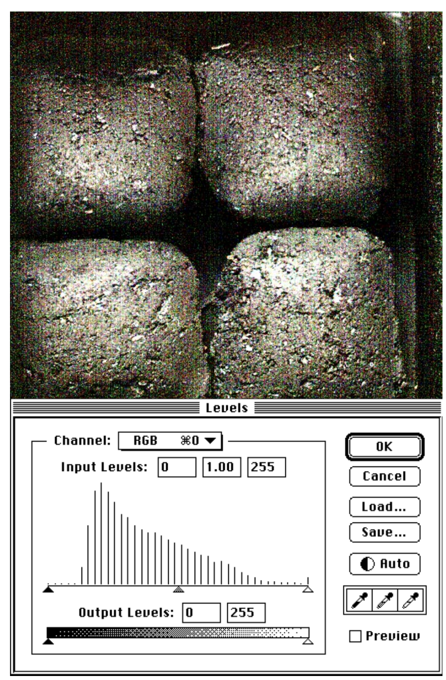

Figure 6

For the DPS cameras, the "deep shadow" test of charcoal briquettes in a black box was *almost* too severe a challenge. Most of the devices tested could only just barely discern the presence of the briquettes, let alone show any significant detail. Nonetheless, differences in performance are clearly evident in the shadow detail composite image.

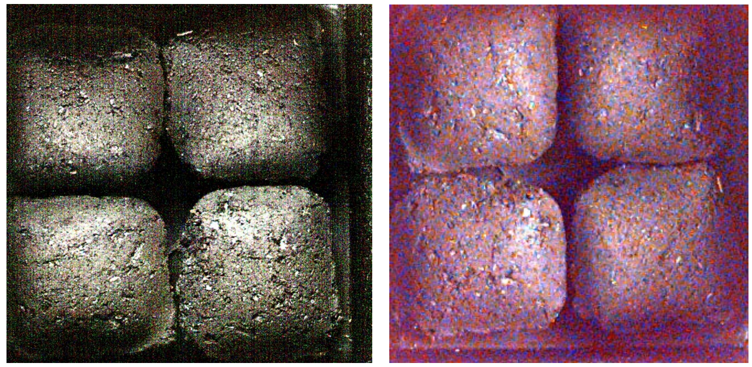

Figure 7

Camera limitations

In discussing techniques for finding camera limitations that aren't visible to the naked eye, a natural question is, "Who cares?" After all, if you can't see it, why worry? The answer is that sooner or later, you'll need to "push" an image in some way, perhaps to open up the shadows, boost midtone contrast or hold back the highlights a little. When you do this, any tonal imperfections will be magnified, along with whatever detail you're trying to bring out. If you've never shot a photo that wasn't perfectly exposed, then congratulations, and just skip this entire section. If you're subject to the same slings and arrows as the rest of us mortals, though, you owe it to yourself to look carefully at what a digital camera will do before you invest your hard-earned money. With a film camera, you can always use a different type of film to deal with poor exposure, a contrasty subject, etc. With a digital camera, what you have is what you've got.

Noise

The characteristic that most limits the overall tonal range of a digital camera is noise in the sensor array. Without getting overly technical, a noisy sensor means you'll see "grain" or "snow" in solid grays or colors. This effect is particularly pronounced in shadow areas. On the CD images, you can use the procedure described earlier to boost the shadows into an easily visible range. The darker steps of the Kodak gray scale will give you the most consistent reference to work with, but for subjective evaluation, the charcoal briquettes in the black box are perhaps the most useful.

Sensor noise usually appears as "snow" in the image, much as you would see on a television set experiencing poor reception. Depending on the construction of the camera system, the noise may appear as randomly distributed monochrome or colored flecks. (Certain sensor irregularities may result in some patterning to the noise flecks, but usually the distribution will be entirely random.)

You're likely to have gotten a good look at sensor noise in the previous experiment, while looking at shadow detail. In the process of pulling the shadow tones up into the visible range, you also would have greatly increased the visible effect of any sensor noise that might have been present. Figure 7 shows shadow detail samples from a studio camera with low noise on the left, and from one with high noise on the right.

In DPS cameras, noise effects are largely masked by the JPEG image-compression process. Noise is nonetheless a primary limitation, even if we don't see the effects directly.

Sensor and optical artifacts

As we discussed earlier, each image captured by a digital camera is actually three *separate* images (red, green and blue). In order for your picture to look right, all three of these images must be kept in perfect registration. If the registration between the tricolor images is off at any point, the result will be a color "artifact." (Artifact here is just another word for something in the image that came from the camera, rather than the scene.)

Array striping artifacts

Ideally, a digital camera would arrange to have separate red-, green- and blue-sensitive sensor elements staring at each pixel-size area of the scene. In most cases, this is prohibitively expensive, even for high-end studio cameras.

A popular, low-cost way of achieving single-shot capability in a digital camera is to "stripe" the sensor array with microscopic color filters. The result, as shown in Figure 8, is generally an array of RGB color triplets, much as you can see if you look closely at a color TV picture tube when it is operating. This approach trades off resolution for single-shot capture capability, but carries with it other penalties as well.

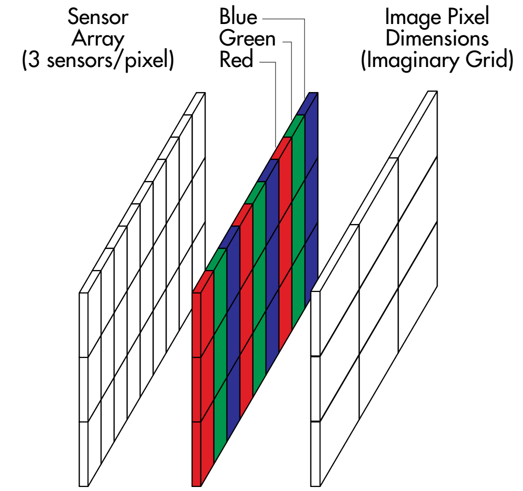

Figure 8: "Striped" CCD Sensor Array ("Single-Shot")

The most significant limitation of striped sensor arrays is that they are prone to generating color moiré and other artifacts when viewing small, high-contrast objects. (This situation frequently arises in product photography, where fine black type on labels can cause these kinds of problems.) In the Davebox image, you may be able to see such artifacts around type, or as a colored moiré on the resolution target with the repeating patterns of vertical lines. In print applications, as long as the file isn't magnified too much, these artifacts are somewhat hidden by the halftone printing process. If the image is magnified at all, though, these artifacts can be quite evident. In multimedia applications, these effects are almost always visible.

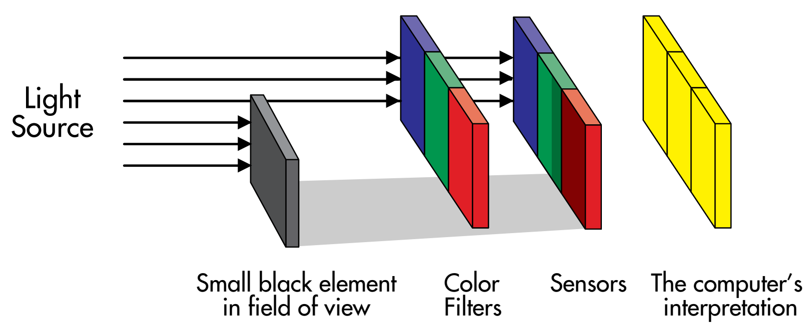

Figure 9: Striped Sensor/Part-III 1.1

On the right-hand side of the Davebox, you'll see an area with patterns of parallel lines, mentioned earlier. This is a test pattern for measuring resolution on high-end cameras. While this high-resolution test pattern isn't too useful for evaluating most of the DPS cameras, you'll undoubtedly notice some interesting color artifacts around the different blocks of lines. These false colors are caused by an interference pattern between the spacing of the lines in the test target and the spacing of the red-green-blue filters on the surface of each camera's imaging array. Differences in compression techniques between cameras make them more or less prone to producing such artifacts. While this pattern represents a particularly severe test, a camera that has trouble with it is also likely to have problems with things like venetian blinds, etc.

Overexposure "blooming"

A final limitation of CCD sensors is their reaction to severe illumination overloads. This is the reason for the inclusion of the shiny pot lid in the Dave Box target, which reflects light sources back into the camera lens. In the face of extremely high light overloads, some CCDs will "leak" charge from the overexposed elements into adjacent cells. This phenomenon is called blooming, and various methods are employed to prevent it. It most frequently shows itself as a colored fringe around specular highlights. Usually, the sensor will bloom differently in each of the red, green or blue channels, producing the colored fringes where one channel has bloomed more than the others. The impact of this for your particular work will depend on its nature: If you intend to do a lot of photography of chromed auto parts, blooming could be a big problem. On the other hand, it would be a complete nonissue in photos of bath towels. In any event, little if any blooming was evident in the DPS cameras tested.

Resolution(!)

Resolution is one of the most misunderstood, misrepresented and confusing parameters in the entire field of digital photography. It is also one of the hardest characteristics to specify in a precise, objective manner, particularly for the digital Point-and-shoot cameras.

One of the most important things to understand about resolution in digital cameras is that pixels are ***not*** resolution! ALL of the entry level Point-and-shoot cameras use large amounts of image compression to squeeze a reasonable number of images into their limited RAM. The complexity and variety of these compression schemes makes it nearly impossible to arrive at any meaningful correlation of pixel count with the actual ability of a camera to resolve detail. Once again, it's really important to look at the test images to determine how each camera *actually* performs.

For the DPS cameras, the starburst target in the upper right-hand corner of the Davebox is probably the best objective measure of resolution. As you progress toward the center of this target, the "spokes" become smaller and more closely spaced. At some point, they blur together into a formless blob. Higher-resolution cameras are able to resolve the individual spokes further into the center of the test pattern than lower-resolution ones.

Compression: 10 pounds into a 5-pound bag

At first glance image compression looks like magic. In fact, at second glance, it's pretty easy to convince yourself it *is* magic! How else would you describe a technology that crams *sixteen* 1-megabyte images into a single megabyte of actual storage?

Of course, nothing comes for free, not even with the spiffiest of high technology: While you can get away with pretty substantial amounts of image compression without people noticing, the level used in most DPS cameras is well beyond that point. Nonetheless, different devices succeed to varying degrees, depending on the details of their particular compression techniques.

The key to usable image compression is to throw away "unimportant" information in the images, and to take advantage of local areas of similarity within each image. For instance, you don't need 24 bits of information to tell you what color a particular piece of an image is, if it happens to be about the same as the pieces on either side of it: Just record the differences! If the differences are small, the information you'll need to store will be small as well. Also, the eye tends to be much more sensitive to certain kinds of detail in an image than to others. If we can find a way to throw out only the sort of detail our eyes aren't very sensitive to, we can reduce the size of the file without our eyes noticing what we've done.

This is the essence of all image compression schemes, but as you may suspect, the concept is easier to describe than to implement. While standard techniques exist, there is a lot of latitude as to how manufacturers can tweak them to produce the best results for their particular camera. Recent advances in compression technology go beyond standard techniques, and some manufacturers (notably Kodak) have taken advantage of them to capture finer detail with little increase in memory usage.

When looking at the images on this CD, the effects of image compression generally can be seen as a "blockiness" in areas of fine detail and high contrast. Look around the edges of objects contrasting strongly with their backgrounds: You'll see errors introduced by the compression process in the form of square blotches at the corners, and "stairstepping" down diagonal edges. You will see differences between cameras though, which may be important in your application. Note too, that it is important to look at areas with much "flatter" contrast as well: Sometimes, manufacturers tweak the compression methods in ways which throw out too much information in areas with subtle contrast differences. We found this particularly evident in the "insurance" images of the car engine. In looking at these images, you may need to use the midtone or "gamma" slider on the Photoshop Levels control to boost the brightness of the shadows. Looking closely though, we found the dark-colored engine parts and hoses to be a tough test for some cameras' compression schemes: In some cases, the three-dimensional hoses literally look "flat," with little apparent shape to them.

Digital Point & Shoots: Today into Tomorrow

In just the last 12 months, digital Point-and-shoot cameras have gone from virtually nowhere to one of the fastest-growing segments of digital imaging. Expect to see new models every few months or so into the foreseeable future, with increasing features and decreasing costs. In the meantime, the current crop of devices for the first time provide genuine utility at a cost within reach of most business budgets. (And many personal budgets as well.) Beyond practical uses, I have to say that the cameras are just plain a lot of *fun* as well: It was a sad day around the Etchells house when it was time to pack up all the cameras and send them back to the manufacturers!

LIGHTING NOTES FOR DPS CAMERAS

One of the things we struggled with in performing the camera tests was how to get reasonable lighting for some of the shots (such as the Davebox), without resorting to equipment or techniques everyday users of the devices wouldn't have available to them. The DPS cameras are generally set up to produce best results with "daylight" lighting. (That is, light with about the same color characteristics as average indirect daylight.) Most flash units produce this type of light, but typical home or office lighting does not. Most DPS cameras have a pretty limited flash capability, meaning supplementary light will be needed for many photographs. What to do?

While this doesn't relate directly to the evaluation of one camera against another, we felt lighting would be pretty important to your eventual use of the devices. Also, photographic lighting terminology can sound a bit arcane to the uninitiated. Hence, this diversion.

Our eyes are amazingly adaptable to different viewing conditions. We tend to perceive colors very similarly even in radically different lighting conditions, because our eyes can compensate for a lot of the differences caused by the color of the illumination. Cameras however, whether film or digital, work in absolutes: If the light is "warmer" or more yellow, the captured image will look yellow as well. Likewise if the lighting is "cooler" or more bluish, the resulting image will have a definite blue cast. You'd expect this to be what you'd want from a camera – a faithful copy of whatever it is looking at. Unfortunately, because our eyes tend to compensate so much for different lighting, the accurate images the camera shows us often look unnatural, with exaggerated color casts.

So what's the point of all this? Just that if you want your digital camera images to look anything like what you're photographing, you'll need to be certain your lighting has a color balance pretty close to daylight.

Color "Temperature"

How do you tell how closely your lighting matches daylight? By its color "temperature," of course. Say what?

To avoid getting into higher physics and talk of "blackbody radiators," lets just describe color "temperature" as a measure of the color cast of a light source: Lower numbers correspond to warmer or yellower light, higher numbers to cooler or more bluish light. The units are degrees Kelvin, or just Kelvin or "K" for short. "Average" daylight is considered to have a color temperature of somewhere around 5500K. Typical household incandescent bulbs produce light with a color temperature around 2800K. Common "photoflood" bulbs generally have color temperatures of 3200 or 3400K. If you're setting up for inside photography, you'll want the special photoflood bulbs with a blue coating on them. These bulbs have a color temperature of 5000K, which is close enough to daylight for most purposes. (In the "Davebox" tests, we tested the cameras with both on-camera flash and 5000K.)

A word of warning about photoflood bulbs though: They run *HOT!* BE SURE YOU ONLY USE THEM IN FIXTURES DESIGNED TO TAKE THE EXCESSIVE HEAT!! Also note that photoflood bulbs typically have lifetimes on the order of a few *hours*, so don't leave them on any longer than you absolutely need to.

This whole discussion on lighting may be superfluous for most people: The built-in flash units on most cameras are adequate for many situations. If you do need more control of the lighting though, and are working on a budget, try some 5000K photofloods. (And buy a bigger air conditioner!) In practice, we unfortunately found the photofloods of limited value, as even 2000 watts (!) of floodlights still produced a rather dark image. Nonetheless, if you are photographing a relatively small (and heat-tolerant!) object, photofloods will give you much more control over the lighting than will the on-camera flash.

Future Image, Inc. is a research and information firm specialized in digital photo-imaging, with offices in Burlingame (CA) Penfield (NY), and Nagoya (Japan). Publications include the monthly Future Image Report, the Digital Camera Reference CD Series, and a range of market research reports. Future Image can be reached by phone at: 1 800 PIXELS2, Fax: 415-579 05 66. Email: AGFuture@AOL.com Information on the full range of Future Image products is available through the company's web site: http://www.imagesite.com/futureimage