IQA: A One-Step Automatic Image Processing System

Julian Bullitt, Jay Thornton, Image Science Laboratory, Polaroid Corporation, Cambridge, MA

Abstract

IQA, Image Quality Assured, is a unified image processing platform designed to enable the automatic enhancement of images from a wide variety of image sources to be displayed on a wide variety of hard and soft copy formats. IQA is designed as an alternative to a paradigm of manual iteration using menu driven image processing selections. Based on the choice and degree of image processing steps, IQA automatically achieves visually optimal images. IQA consists of an algorithm set (many employing the Discrete Cosine Transform) and an intelligent control manager for quality optimization from device characterization data stored in profiles.

The Art of Image Quality Optimization

For most of the history of image reproduction, techniques for optimizing the quality of the reproduced image have been available only to skilled craftsmen. Until recently, those tools were based on careful manipulation of the reproduction process. Even in the hands of an artisan, many problems in image reproduction, for example burning and dodging or selective color correction, were difficult or very expensive to overcome. Thus the industry looked forward to the blessings of electronic and digital imaging whereby those difficult problems would be easily corrected. Digital imaging, however, has proven to be both a blessing and a curse. The same electronic tools which are powerful enough to fix stubborn reproduction problems, are also powerful enough to turn an image into garbage in just a few key strokes. In the early electronic age, an artisan was still required to optimize reproduced image quality.

In the last decade these electronic tools have become available to less skilled users who are not experts. For a typical consumer, and for many business users, image optimization consists of several, iterative, trial and error loops. First the user loops trying to use the tools to make the image look right on a monitor. Next the user tries printing the image. The print usually shows other deficiencies, perhaps in color or tone, perhaps in sharpness or noise. In any case, the user often returns to looping on the monitor in attempting to tweak the image into a good print. The "Undo" and "Revert" commands are used frequently to recover from harmful tweaks.

IQA represents an automatic, image quality optimizing software platform which effectively incorporates the collective knowledge accrued from over 150 years of photographic imaging. The IQA platform has an internal representation of the relevant dimensions of image quality and a control system for performing the trade-offs among those dimensions in order to optimize image quality. IQA effectively shifts the burden of managing the many conflicting subgoals of image quality from the human to the computer. This leaves the human free to manage the higher levels of the task such as design and composition, much as Edwin Land envisioned Polaroid one-step photography.[1]

Image Quality

Before breaking the IQA platform down and explaining some of its components, let's focus on what we mean by "optimizing image quality." The dimensions underlying the perception of quality in imaging systems are: sharpness, graininess (noise), color, and tone (exposure, contrast, highlight and shadow, etc.). Of course, any particular image may be pleasing or not due to its content and composition, but the imaging system itself determines image quality by virtue of its impact on the four fundamental perceptual dimensions. It is the imaging system quality that IQA aims to maximize. Implicit in this goal is our belief that image quality can be effectively quantified.[2] We measure quality by constructing metrics that link objective measurements taken from images to the subjective perceptions of groups of observers taken in psychophysical scaling experiments.[3] We employ these image quality metrics to drive film development, to drive product development, and to drive specific imaging system optimization strategies.

We acknowledge other approaches that seek to determine, within a single image, artifacts or areas of visible grain; but these approaches are extremely compute intensive. We regard our system optimization approach as offerring the best value because, in our experience, it recovers almost all of the available image quality without the computational burden of extensive image-by-image analysis. Note that this is not to say that IQA algorithms are not adaptive, but only that all of their control parameters can be set based on system considerations, not individual image analysis.

IQA Components

IQA can be divided into three primary components (Fig. 1): a set of image processing **algorithms**, **profiles** which characterize the color and spatial performance of the image system peripherals, and an intelligent **control system** for using the profile information to determine the parameters for the algorithms and to achieve image quality optimization.

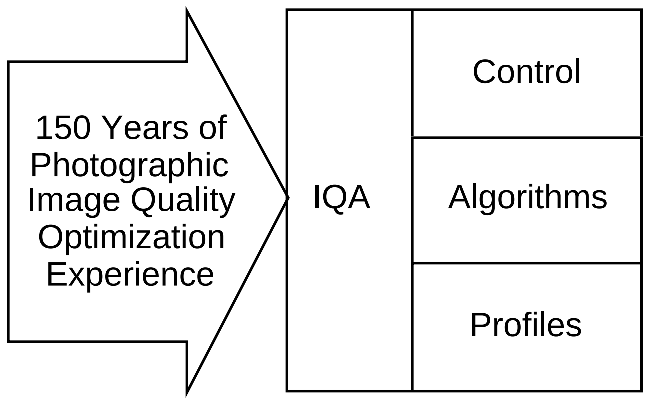

Figure 1.

The image processing algorithms consist of some modules addressing the various dimensions of image quality and other modules addressing common processing needs such as image resizing. Underneath the modules for image sharpening, noise reduction,[4] and resizing is a common engine based on the the Discrete Cosine Transform (DCT). The features and advantages of this engine will be the topic of a later paper.[5] The color management module delivers what is by now fairly conventional color matching between source images and the output images. Tone correction is accomplished by a pair of modules, one performing exposure correction and the other performing burning and dodging. These algorithms are related in that they both use a photographic model for scene exposure to correct the lightness and darkness of image areas. Exposure correction analyzes image features and adjusts the entire image.[6] Burning and dodging makes similar shifts in lightness or darkness, but makes local corrections rather than the global exposure correction. By making local corrections, burning and dodging can bring out detail that may have been tonally suppressed in the shadow or highlight portion of the image without affecting the overall image lightness.

Because there are no other presentations concerning the IQA profiles, they will be described here. The profiles used by IQA follow the format specified by the ICC (International Color Consortium) for containing the color characterization/transformation information of both input devices, such as electronic cameras and scanners, and output devices, such as printers and monitors. The ICC format is tagged, meaning that there are multiple data structures (tags), which are sequentially stored in the profile along with a directory of which tags are present and where they can be found in the file. The key tags (typically) contain three dimensional tables which relate device dependent color data (e.g. a specific device RGB) to device independent color data (e.g. CIE-LAB). In addition to specifying formats, and measurement units, the ICC also specifies a "reference viewing environment" (a D_{50} illuminated light booth) and a "reference reproduction medium" (ideal, 100% reflective print paper) in order to further clarify the intended appearance of images after color management based on these profiles. For more information visit the ICC Web site at: http://www.color.org.

All of the color management information in IQA profiles adheres to the ICC specification. Thus IQA profiles can be used as ICC profiles. To create a profile, this color information is collected by printing or scanning a large collection of different colors. The characterization data is then passed to programs which correlate the device dependent and independent color coordinates and record the resulting multi-dimensional tables in the profile.

The ICC format also accommodates private tags whose interpretation is not subject to standardization. IQA uses private tags to store characterization information supporting spatial processing and the control system which optimizes the algorithm parameters. This information is collected from measurement of test targets such as uniform patches (for noise properties) and edges at various orientations (for sharpness properties). The information in an IQA profile is the data necessary to optimize the performance of an imaging system. IQA benefits from that experience and essentially performs an "on-the-fly" quality optimization for the system composed of the particular input and output devices associated at that moment.

The third component of IQA, the intelligent control system, incorporates information obtained from the input and output profiles, from the image itself, and from the user's selections. This control system calculates optimal control values for all the processing modules acting in concert by simulating a virtual observer. For specific control values, the virtual observer represents a numerical prediction of expected image quality based on application of the quality metrics we have developed and validated in the context of both conventional and digital imaging. The control system is described more fully in the paper by Hultgren.[7]

IQA Architecture

In the preceding, we have deliberately used the term "IQA platform." "Platform" is meant to convey that the associated software provides a resource that could have many different forms of presentation in imaging products. The Polaroid product Before and After™ represents the platform with a user interface wrapped around it to become a stand-alone application. Polaroid PDC Studio™ demonstrates IQA support of our PDC 3000™ electronic camera in the form of an Adobe Photoshop™ plug-in (in addition to a stand-alone application). Closed systems such as Polaroid's Make-a-Print Express™ use IQA. IQA has appeared at the device driver level in our line of SprintScan™ 35mm scanners (where exposure adjustment has allowed outstanding performance in the scanning of negatives). IQA might also be made available as an operating system resource (like Apple ColorSync™). The above are all software implementations of IQA. The platform is also poised to benefit from hardware acceleration of its spatial processing engine, because JPEG image compression has driven down the cost of DCT ASICs.

The IQA extendible, multi-platform architecture is an object-oriented design implemented in C++ and available on

Windows 95™, Windows NT™, and MacOS™. IQA uses MMX™ technology when available. The engine can be thought of as a pipeline assembled from modules which perform either image processing operations or logical control of other modules. It is a pipeline in the sense that images are processed in independent strips and so output values for early strips are available while later strips are still being processed. The strip processing allows large images to be processed in a fairly small memory footprint. The pipeline uses a data flow architecture with distributed control so that new modules can be easily dropped into the pipeline since they include their own control logic. Programmers are presented with a simple API since all processing is done with a few calls, most of the control values are calculated internally, and all of the memory management is handled by IQA.

Conclusion

We have tried to provide an overview of problems addressed by IQA, the general approach Polaroid has taken to implementing an automatic image quality optimizing platform, and a foundation for the following papers which develop some of the component technologies in more detail.

Acknowledgements

The authors would like to thank the following members of the IQA team for helpful comments: Mike Antin, Bill Donovan, John Francis, Al Moyer, Rich Piccolo, and Mike Vigneau.

References

1. E.H. Land, H.G. Rogers & V.K. Walworth, "One-Step Photography," Neblette's Handbook of Photography and Reprography, J.M. Sturge, ed.. NY: Van Nostrand, 7th. ed., 1977.
2. C.J. Bartleson & F. Grum, Optical Radiation Measurements, Volume 5, Visual Measurements, NY, Academic Press, 1984.
3. J. Bowman, J. Bullitt, F.R. Cottrell, B.O. Hultgren, 42nd Annual Conference, pp 465-468, 1989.
4. I. Hajj-ahmad, M. Wober & Y. Yang, 50th Annual Conference, 432 1997.
5. I. Hajj-ahmad & M. Wober, 51st Annual Conference, paper 04B-P221, 1998.
6. J. Boyack & I. Hajj-ahmad, 51st Annual Conference, paper 29C-P044, 1998.
7. B.O. Hultgren, 51st Annual Conference, paper 05B-P207, 1998.