

NOVEMBER 2005
VOL. 16, NO. 2

*Newsletter of the SPIE Electronic
Imaging Technical Group*

*Gabriel Marcu, Apple Computer, Inc.
Editor/Technical Group Chair*

Contents

Multiscale contrast enhancement for medical images, A. Monica Trifas, John M. Tyler, and Oleg S. Pinykh*, Department of Computer Science, Louisiana State University (LSU) at Baton Rouge, LA; *Department of Radiology, LSU Health Sciences Center at New Orleans, LA 2

Demosaicking techniques for multispectral cameras using mosaic focal-plane-array technology, Gaurav Baone and Hairong Qi, Advanced Imaging and Collaborative Information Processing (AICIP), Department of Electrical and Computer Engineering 3

Automatic conversion technique from data dependent triangulation to SVG B-splines, G. Messina, E. Ingrà, * S. Battiato* and G. Di Blasi*, STMicroelectronics, AST- Imaging Group, Catania, Italy; *Department of Computer and Mathematic Sciences, IPLAB Group, University of Catania, Italy 4

Resolution and light sensitivity tradeoff with pixel size, Joyce Farrell and Feng Xiao*, Stanford University, Stanford, CA; *Agilent Technologies, Santa Clara, CA 5

Half-toning via Perona-Malik diffusion and stochastic flipping, Jianhong (Jackie) Shen, School of Mathematics, University of Minnesota 6

ISO dynamic range interpretation: risky business, Don Williams and Peter D. Burns, Kodak Research Laboratories, Eastman Kodak Company 7

Image noise reduction and sharpening filter, Miliwoje Aleksic, Maxim Smirnov, and Sergio Goma, ATI Technologies 8

Gray tracking for LC displays, Gabriel Marcu, Apple Computer, Inc. 9

Target tracking using soliton patterns in a large array optoelectronic system, Adrienne Raglin, Army Research Laboratory 12

ELECTRONIC IMAGING

Volumetric depth peeling for virtual arthroscopy

David Borland, John P. Clarke*, and Russell M. Taylor II, Department of Computer Science, *Department of Radiology, University of North Carolina at Chapel Hill

Three dimensional (3D) display of medical data sets is an increasingly useful tool. Viewing 3D reconstructions of objects from magnetic resonance imaging (MRI) and computerized tomography (CT) data is more natural and intuitive than mentally reconstructing these objects from orthogonal slices of the data, especially with the increase in size of datasets due to improving scanner resolutions. When displaying such data sets using volume rendering, appropriate selection of the transfer function is critical for determining which features of the data will be displayed. For virtual arthroscopy, however, no transfer function can enable views like that of a socket from within the ball of a joint because the ball itself blocks the view.

Volumetric depth peeling (VDP) was developed to enable the rapid survey of joints for pathology by automatically culling occluding voxels, thus enabling textbook-like illustrations of joint surfaces that are not possible even with clinical arthroscopy (see Figure 1). VDP borrows conceptually from previous work,¹ but is extended to the more general case of ray-based volume rendering, retaining full transfer-function control.

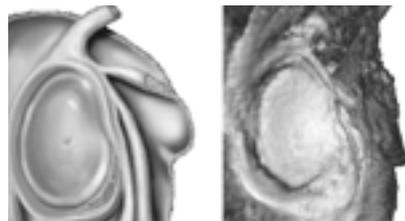


Figure 1. Textbook cut-away of a shoulder socket compared with volumetric depth peeling.

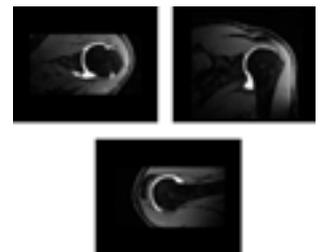
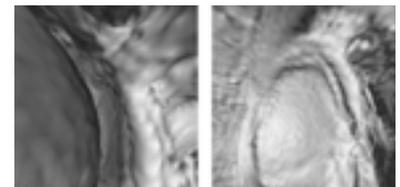


Figure 2. Three orthogonal MRI slices.

Figure 3. Standard virtual arthroscopy versus volumetric depth peeling.



Arthroscopy

Arthroscopy is the process of inserting an arthroscope (a small camera) into a joint through an incision in the skin and so providing views of the interior of the joint. The technique is useful for diagnosing derangements within the joint and for performing surgery to treat them. However, arthroscopy is an expensive, invasive procedure that requires skilful manipulation of arthroscopic tools confined to the cramped spaces within joints. For pre- and post-operative diagnosis of joint derange-

Continues on page 11.

Multiscale contrast enhancement for medical images

A. Monica Trifas, John M. Tyler, and Oleg S. Pinykh*, Department of Computer Science, Louisiana State University (LSU) at Baton Rouge, LA; *Department of Radiology, LSU Health Sciences Center at New Orleans, LA

Here we present a method of image enhancement that uses multiscale methods for contrast manipulation. The method is based on the Laplacian pyramid and 2D wavelets. The basic idea in multiscale enhancement is to decompose the image into components that represent individual details, and to improve the contrast by operating on these components rather than on the original image. In our research, the image is decomposed according to the Laplacian pyramid transform.

A pyramid is a multiscale representation built through a recursive method that leads naturally to self-similarity. To build a Gaussian pyramid, the original image is convolved with a lowpass filter and subsampled with a factor of two; the filter-subsample operation is repeated recursively to produce a sequence of images. To obtain a Laplacian pyramid of the same image, a bandpass filter is used rather than a lowpass filter. A Laplacian pyramid is a complete representation of an image in the sense that one can perfectly reconstruct the original given the coefficients in the pyramid. The reconstruction process is straightforward: we simply expand each image up to the full size of the original image using an interpolation filter, and then sum all of the interpolated images.

A Gaussian-like weighting function is used to compute the predicted value for each pixel. This function is centered on the pixel itself. The weighting function is convolved with the image and the predicted values of all pixels are obtained as a result of the convolution operation. Then, the lowpass filtered image is subtracted from the original.

Contrast improvement is achieved by modifying the coefficients of the Laplacian pyramid. Small coefficients represent subtle details. These are amplified to improve the visibility of the corresponding details. The strong density variations have a major contribution to the overall dynamic range, and these are repre-

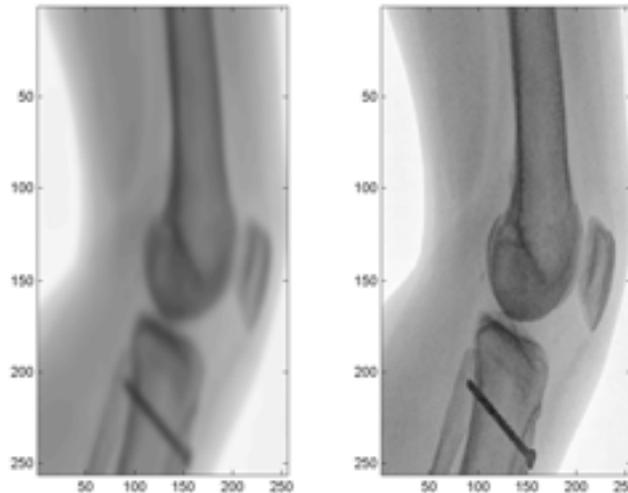


Figure 1. Example of the contrast enhancement of a medical image using a technique incorporating the Laplace pyramid.

sented by large coefficient values. They can be reduced without risk of information loss, and by compressing the dynamic range, overall contrast resolution will improve.

The human visual system is sensitive to the different spatial frequencies in an image. In particular, the plots for human visual frequency indicate that some frequencies are more visible than the others, and some are not important at all. Removing certain frequencies can help emphasize the others (keeping the total image 'energy' the same), and improving the quality of the image.

To optimize the method, we have used a set of medical images with visible details. Based on these, we have constructed and quantified an ideal image-frequency profile, i.e. a one that corresponds to the most 'balanced, natural' image.

We have studied the effects of the most common artifacts (such as blurring and noise) on the frequency content of each image. First we compared (using statistical parameters) the multiscale decomposition corresponding to a

blurred or sharp medical image with that corresponding to an 'ideal' image. Based on this comparison, we computed the values of the coefficients to be applied to the components of the blurred or extremely sharp image in different frequency bands. For each image, we built a Laplacian pyramid decomposition, then multiplied each level of the pyramid with pre-selected coefficients. Finally, we reconstructed the image using the new levels stored in the Laplacian pyramid. An example of the typical improvement is shown in Figure 1. Thus, we demonstrated that a Laplacian pyramid can be used to produce distinguishable image frequency profiles, and have applied these to the contrast enhancement of medical images.

A. Monica Trifas, John M. Tyler, and Oleg S. Pinykh*

Department of Computer Science, Louisiana State University (LSU) at Baton Rouge, LA

*Department of Radiology, LSU Health Sciences Center at New Orleans, LA
E-mail: mtrifas@hotmail.com

References

1. P. J. Burt and E. H. Adelson, *The Laplacian Pyramid as a Compact Image Code*, **IEEE Trans. on Communications**, COM-31 (4), pp. 532-540, 1983.
2. S. Dippel, M. Stahl, R. Wiemker, and T. Blaffert, *Multiscale Contrast Enhancement for Radiographies: Laplacian Pyramid Versus Fast Wavelet Transform*, **IEEE Trans. on Medical Imaging** 21 (4), pp. 343-353, 2002.
3. P. Vuytsteke and E. Schoeters, *Image Processing in Computer Radiography*, **Int'l Symp. on Computerized Tomography for Industrial Applications and Image Processing in Radiology**, Berlin, Germany, 1999.

Demosaicking techniques for multispectral cameras using mosaic focal-plane-array technology

Gaurav Baone and Hairong Qi, Advanced Imaging and Collaborative Information Processing (AICIP), Department of Electrical and Computer Engineering

The advent of the digital camera has put an end to the laborious processes required for traditional film-based photography. Instantaneous results and simple post-processing options have made digital cameras popular and irreplaceable. However, the final image that we perceive as the photograph is rather different than what the device actually captures. The camera consists of an array of photo-sensors arranged systematically in a specific pattern. This pattern, known as the color filter array (CFA) pattern, controls the placement of various wavelength-specific sensors throughout the image plane.

To form a color image, three types of sensors, each sensitive to one of the primary colors of the visible spectrum (red, green and blue), are used. Instead of piling up all the three sensors at each pixel location in the image plane, the color filter array allows controlled placement of only one sensor of each type at each pixel. That is, at each pixel location we have only one band value is sensed and two are missing. This problem is solved by using a technique called *demosaicking*, which uses neighborhood information to estimate the missing pixel values and thus forms the required three-band color image (see Figure 1). This color camera technology is often referred to as mosaicked technology.

The success of this technology in color cameras motivated us to apply it to multispectral image-acquisition systems. Existing multispectral cameras use expensive and mechanically delicate equipment like image spectrometers to capture images. The extension of this technology to multispectral cameras would require a new filter array that can accommodate more than three spectral bands. This problem has been solved by Miao et al.,¹ who have developed a generic method to generate mosaic filter array patterns for any given number of spectral bands in an image (see Figure 2). However, having given the solution to this problem, the primary problem of recreating the multispectral image from the registered band values still remains.

Color cameras use interpolation-based strategies to recreate the final color image from the sparsely-distributed three-band values, and we have successfully extended these to the multispectral case. However, the problem with interpolation-based methods in multispectral images is that, due to the increased number of spectral bands, the distribution of the missing band values in the mosaicked image becomes

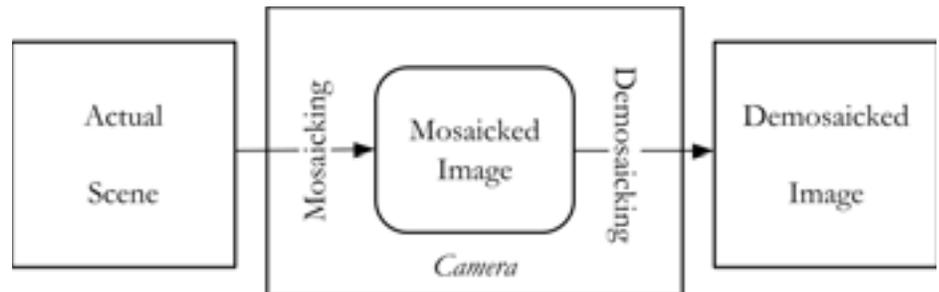


Figure 1. Block diagram of the mosaicked technology.

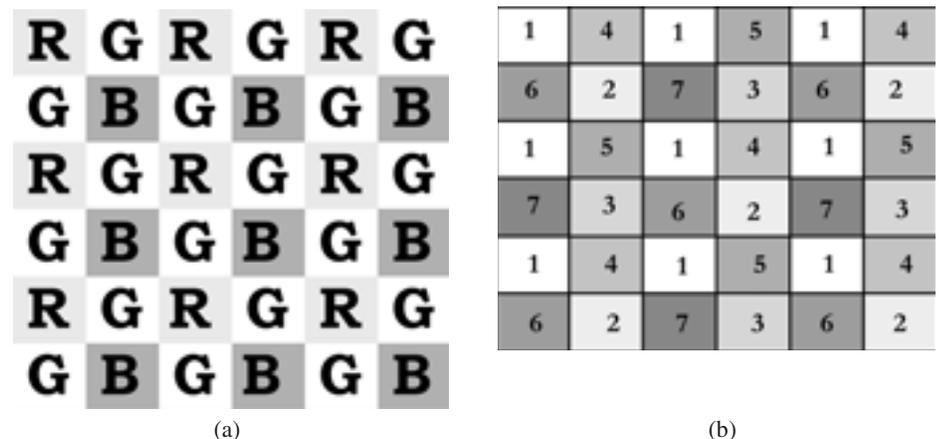


Figure 2. Mosaic filter array patterns for different number of spectral bands: (a) three band CFA (Bayer), (b) seven band filter array.

even more sparse and complicated. This affects the correlation between neighboring pixels, which in turn affects the final output.

Another problem with the interpolation-based methods is that they have no provision for removing external noise and degradation. To tackle these problems we have developed a novel demosaicking approach that treats the demosaicking problem as a classic case of the image restoration problem. In our case, the original image is the actual scene that has to be captured by the camera. This means that we do not know anything about the original image, which goes through external noise and degradations and is then captured in the form of the mosaicked image by the filter array. This is then treated as degraded, and our goal is to restore it to form a degradation-free multispectral image. Missing pixels are considered to be part of the image degradation.

As there is no prior information about the original image, we adopt the *maximum a-posteriori probability* (MP) method to solve this problem. The technique assumes that the original image is from an ensemble whose distribution is assumed to be a multivariate Gaussian.² The problem then focuses on maximizing the *a-posteriori* probability: the conditional probability of the original image given the observed image. Maximizing this will give us the maximum likelihood of the image chosen being the original for the given degraded image. The optimization problem is solved using the well-known gradient-descent method.

We have tested our algorithm on seven-band synthetic multispectral images, corrupting the images with noise and other types of degradation. These corrupted images were used as in-

Continues on page 10.

Automatic conversion technique from data dependent triangulation to SVG B-splines

G. Messina, E. Ingrà,* S. Battiato* and G. Di Blasi*, STMicroelectronics, AST- Imaging Group, Catania, Italy; *Department of Computer and Mathematic Sciences, IPLAB Group, University of Catania, Italy

The conversion of raster images into vectorial shapes is a challenging branch of computer graphics. Here we present a technique to convert surfaces, obtained through a data-dependent triangulation (DDT), into Bezier curves by using a Scalable Vector Graphics (SVG) file format. The method takes as input a data-dependent triangulation obtained from a raster image. Characteristics of the triangles are taken into account to trace a map of the boundaries, and the estimated triangle barycenters are connected together. Thus the conversion of the resulting polylines into curves is performed. After the curves have been simplified and closed, the final representation is obtained by sorting the surfaces in a decreasing order.

SVG language and vectorization

The SVG standard allows the representation of complex graphical scenes via a collection of graphic vectorial-based primitives. These have several advantages over classical raster images, including scalability, resolution independence, and so forth. SVG format could find useful application in the world of mobile imaging devices, where cameras must be matched to displays of limited color, size, and resolution.

The DDT has been used to approximate local pixel neighborhoods using triangles.¹ The triangulation replaces the input raster image with a set of triangles according to a specific cost function that is able to implicitly detect the edge details. The overall perceptual error is then minimized by choosing a suitable triangulation. On the other hand, the DDT is strictly connected to the original pixel positions: as a result, the number of triangles is larger than

the number of pixels. This triangulation could be directly managed by SVG primitives. However, although the quality achieved in this way is rather good, the size of the resulting files may be very large. For example, if we consider a 1600×1200 RGB raster image, the uncompressed size of the file is about 5Mb and the resulting DDT (in SVG format) is about 255Mb.

Proposed technique

Once the triangulation has been performed, the algorithm extracts from the DDT only those triangles that are not equilateral and those equilateral triangles with nearest neighbours in a different colour. In this way, only the boundaries have been considered: the triangles have been synthesized as single points using their estimated barycenters. The construction of the adjacent lists of barycenters is fundamental to permit the correct interconnection of these points. The lists have been created taking into consideration the (x,y) position of each point and by sorting first along the rows and then down the columns. The barycenters have then been connected together by following the boundaries along the wind rose directions: the outer boundaries are connected together, and then the inner boundaries are processed, until the whole triangle set has been elaborated.

After these areas have been created, the conversion to Bezier curves is performed. The points found in the previous step are used as control points and new ones are created as contour points. A simplification of the curves is then applied by removing useless intermediate points. Finally the resulting surfaces are sorted

using an approximation of the area and saved in SVG format. This operation is necessary because the SVG viewer shows the last layer in foreground, thus the bigger areas must be put in background. In Figure 1 some illustrations of the pipeline steps are shown.

Results

The technique has been compared with other raster-to-vector conversion methods^{2,3} and the software performs well in terms of both perceptual and measured quality. Details of the experiments can be found on the SVG UniCT group page.⁴

G. Messina, E. Ingrà,* S. Battiato*, and G. Di Blasi*

STMicroelectronics

AST- Imaging Group, Catania, Italy

E-mail: giuseppe.messina@st.com

*Department of Computer and Mathematic Sciences

IPLAB Group, University of Catania, Italy

E-mail: elisaingra@gmail.com, {battiato, gdiblasi} @dmi.unict.it

References

1. S. Battiato, G. Barbera, G. Di Blasi, G. Gallo, and G. Messina, *Advanced SVG Triangulation-Polygonalization of Digital Images*, **Proc. SPIE 5670** (1), 2005.
2. S. Battiato, A. Costanzo, G. Di Blasi, G. Gallo, and S. Nicotra, *SVG Rendering by Watershed Decomposition*, **Proc. SPIE 5670** (3), 2005.
3. S. Battiato, G. Di Blasi, G. Gallo, G. Messina, and S. Nicotra, *SVG Rendering for Internet Imaging*, **Proc. IEEE/CAMP2005 (Int'l Workshop on Computer Architectures for Machine Perception)**, Palermo, 2005.
4. <http://svg.dmi.unict.it/>

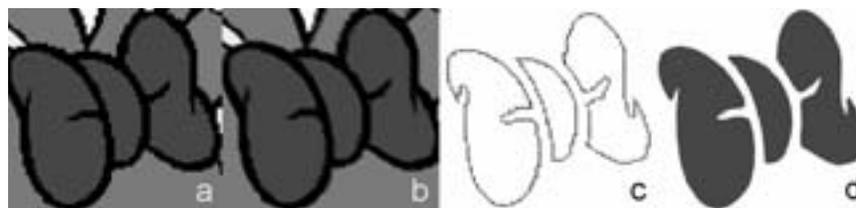


Figure 1. (a) Input raster image. (b) Data-dependent triangulation obtained from the input image. (c) Extraction of a region of interest from triangle boundaries (to show the effectiveness of the approach). (d) Final Bezier curves.

Resolution and light sensitivity tradeoff with pixel size

Joyce Farrell and Feng Xiao*, Stanford University, Stanford, CA; *Agilent Technologies, Santa Clara, CA

Digital cameras are now a standard feature in cellular phones, driving the market for complementary metal oxide semiconductor (CMOS) imagers that can fit within a small form factor. Given that the size of a CMOS imaging sensor array is fixed, the only way to increase sampling density and spatial resolution is to reduce pixel size. But reducing pixel size reduces the light sensitivity. Hence, under these constraints, there is a tradeoff between spatial resolution and light sensitivity.

We define two metrics to characterize the tradeoff between spatial resolution and light sensitivity as a function of pixel size.^{1,2} First, we characterize spatial resolution by the Nyquist sampling frequency of an imaging sensor. Second, we characterize photometric sensitivity by the minimum scene illuminance (lux) required to reach a signal-to-noise ratio (SNR) of 30dB at an exposure of a thirtieth of a second. Our human experimental measurements³ show that sensor SNR must be 30dB in order to render photon noise invisible.

We use the Image Systems Evaluation Toolbox (ISET) to calculate these metrics for imaging sensors with the same dye size but with

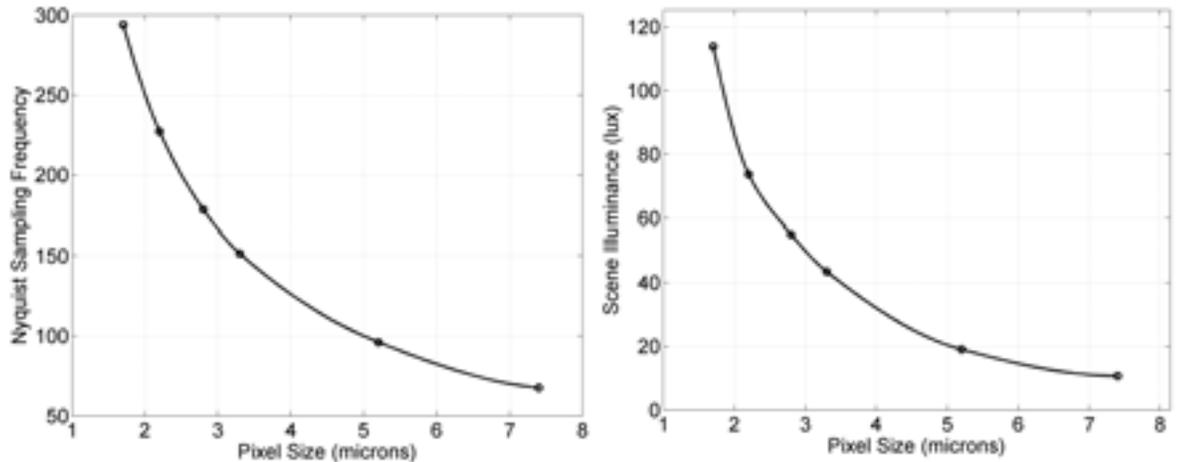


Figure 1. (a) Nyquist sampling frequency increases with shrinking pixel size. (b) Minimum scene illuminance needed to reach 30dB SNR at exposure duration of 1/30s increases with shrinking pixel size.

different pixel sizes.⁵ Table 1 lists the optics and sensor parameters used to simulate the effects of pixel size on an imaging array with fixed dye size.

Tradeoffs

Figure 1a shows how the sensor resolution, estimated by the Nyquist sampling frequency (ISO 12233 standard)⁶ depends on pixel size. Figure 1(b) shows that the minimum scene illuminance (lux) required to reach a pixel SNR of 30dB at an exposure duration of 1/30s depends on pixel size.

The consequences of these tradeoffs are illustrated visually using the ISET digital cam-

era simulation software (Figure 2). We simulated images captured with monochrome sensors containing various pixel sizes. The mean scene luminance was set to 20cd/m², simulating indoor tungsten illumination. The simulated exposure duration was 100ms. This figure illustrates the visibility of photon noise in the 2 μ m-pixel image and the visibility of spatial aliasing in the 6 μ m-pixel image.

We plan to use these types of simulated images in psychophysical experiments. We will use pairwise preference judgments to generate iso-preference curves for camera images with

Continues on page 11.



(a) 2 μ m pixel



(b) 4 μ m pixel



(c) 6 μ m pixel

Figure 2: Simulated camera images that illustrate the tradeoff between resolution and light sensitivity for a 512 \times 512 imaging array with different pixel sizes and fixed exposure duration. The images are scaled to have an equal maximum display value.

Halftoning via Perona-Malik diffusion and stochastic flipping

Jianhong (Jackie) Shen, School of Mathematics, University of Minnesota

Halftoning is the key process governing most binary or ternary printing devices. The major task involved is easy to describe: to convey faithfully a multitude of shades or colors using only a few elementary ones, e.g., black ink dots. In some sense, halftoning simulates quantum physics. The latter reveals that the continuum of the material world we observe actually emerges from the discrete, or quantum, building blocks of basic particles and their discrete states. Halftoning attempts to reverse Mother Nature's engineering by designing models and algorithms to express smoothly varying tones via only a few discrete colors.

For maximal clarity, we focus here only on the halftone process of expressing continuous-tone (or contone) grey shades by simply turning black dots on or off, as in most black and white inkjet printers. We describe here a novel method for halftoning an image using a new error-diffusion algorithm. For a given contone image u in $[0,1]$, the method shows how to design its halftone version, b , which is binary from $\{0,1\}$ at each pixel. In case of printing, one may for convenience assume that $b=0$ deposits an ink dot, while $b=1$ leaves the spot blank.

Error diffusion

Suppose at a pixel α , the contone value $u(\alpha)=0.75$, while the halftoned value $b(\alpha)=0$ or 1. The error $e(\alpha)=u(\alpha)-b(\alpha)=0.75$ or -0.25 , which in either case is non-negligible. This typical scenario differentiates halftoning from other design tasks in computer graphics or computer aided design (CAD), for which pointwise approximation errors often diminish when one employs high-order Fourier modes, polynomials, splines, or wavelets. In contrast, pointwise error evaluation appears pointless for halftoning. As for the human vision system (HVS), the errors must be blurred to prevent them being detectable by the naked eye.

This understanding inspired one of the greatest halftoning methods: invented by Floyd and Steinberg,¹ it is called error diffusion (ED). At a current pixel α , the halftone error $e(\alpha)=u(\alpha)-$

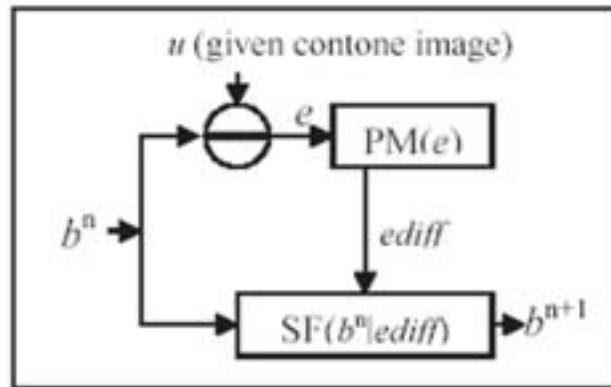


Figure 1. Flowchart of the progressive two-step halftone algorithm.

$b(\alpha)$ is distributed to its neighboring pixels' β s in such a manner that the errors cancel out each other locally. Thus, a typical ED algorithm often relies upon four entities: a path visiting all the pixels, a pixelwise decision rule for converting $u(\alpha)$ to $b(\alpha)$, local windows into which halftone errors are diffused, and their distribution weights.

Novel two-step progressive algorithm

The new halftoning algorithm² is iterative, with two steps at each iteration. Suppose at step n , the current halftone image is b^n with error field e^n . First one diffuses the error field $ediff=PM(e^n)$, where PM stands for the Perona-Malik diffusion process to be explained later. Next, the diffused error field $ediff$ induces a stochastic flipping (SF) strategy that can further polish the halftone image:

$$b^{n+1}=SF(b^n | ediff)$$

So the iteration continues until it converges.

The characteristics of the new algorithm are: independence from particular choices of visiting paths, local diffusion windows, or diffusion weights; being progressive instead of aiming at single-pass completion; and allowing straightforward parallel implementation.

Perona-Malik diffusion (PM)

Diffusion is ubiquitous—e.g. heat diffusion and Brownian motion³—and is often homogeneous

and isotropic: i.e., the same everywhere and in every direction, as quantified by the celebrated heat equation: $u_t=D \Delta u$. For intelligent image enhancement, Perona and Malik⁴ devised an image-adapted anisotropic diffusion mechanism in the form of:

$$u_t = \nabla \cdot (D(\|\nabla u\|) \nabla u)$$

where the diffusivity $D(\cdot)$ depends upon the input image u so that diffusion across edges is discouraged, crucial for not messing up different objects in images.

The current algorithm employs a revised version of the PM diffusion for error diffusion: $ediff=PM(e^n)$ at each step n , so that error flows are confined within each object, as the HVS does.

Moreover, the PM diffusion is a parallel process in contrast to most, which are sequential.

Stochastic flipping (SF)

The diffused error field $ediff$ contains valuable information about the performance of the current halftone version b^n . At any pixel α , a smaller $ediff(\alpha)$ signifies good performance of the current halftone, while a larger one flags deficiency. This key qualitative observation leads to the quantitative strategy of stochastic flipping $b^{n+1}=SF(b^n | ediff)$, which constitutes the other important half of the algorithm.

The flowchart of this progressive two-step halftone algorithm is depicted in Figure 1.

Jianhong (Jackie) Shen

School of Mathematics
University of Minnesota, MN
E-mail: jhshen@math.umn.edu

References

1. R. Floyd and L. Steinberg, *Proc. Soc. Info. Disp.* **17**, pp. 75-77, 1976.
2. J. Shen, *Proc. SPIE (on Electronic Imaging)* 2006, to appear.
3. T. F. Chan and J. Shen, *Image Processing and Analysis - Variational, PDE, wavelets, and stochastic methods*, SIAM Publisher, Philadelphia, 2005.
4. P. Perona and J. Malik, *IEEE Trans. Pat. Anal. Mach. Intell.* **12**, pp. 629-639, 1990.

ISO dynamic range interpretation: risky business

Don Williams and Peter D. Burns, Kodak Research Laboratories, Eastman Kodak Company

The term dynamic range is as old as signal analysis itself. Ask anyone actively involved in the optical or imaging sciences to define it for image capture and you are sure to get an opinion. While many will include words like *minimum*, *maximum*, *tonal range*, *accurately detect*, or *reliably detect*, others will offer the ubiquitous *20 times log₁₀ of the maximum signal to dark noise ratio* (...or is it 10 times?). Others, citing bit depth, may confuse signal encoding with detector performance. It is truly a cluttered landscape, and evokes an observation that this is science in action, at its messy best.

Two ISO electronic imaging standards for capture device dynamic range metrology have been issued. Both ISO 15739 (digital still camera noise) and ISO 21550 (film scanner dynamic range) adopt signal-to-noise ratio (SNR) threshold criteria for determining dynamic range endpoints. There is little doubt that SNR is an appropriate specification metric. Beyond this, though, these standards provide an opportunity for understanding alternative meanings for 'signal' and 'noise' measurement in terms of their spatial, and possibly temporal, components.¹

Intended for digital image capture, the concepts are easily extendable to signal types other than those in the optical sciences. These standards are not perfect but, by drawing upon them, the reader is challenged to think of a statistical interpretation of dynamic range, vis-à-vis SNR. One can adopt a signal-detection perspective and interpret the dynamic range evaluation in terms of various levels of risk. This can be generalized for various signal-specific spatial frequency bands, and noise correlation statistics.

To start, consider the following definition of dynamic range:

Dynamic Range: *The extent of energy over which a digital capture device can reliably detect signals, reported as either a normalized ratio (xxx:1) or in equivalent log optical-density units.*

The operative words in this definition are *reliably detect*. The reliability (think probabil-

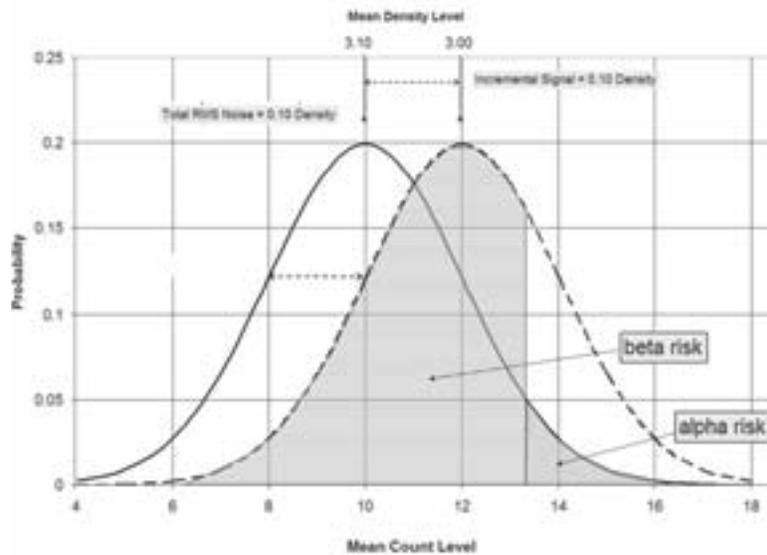


Figure 1. Two signal distributions.

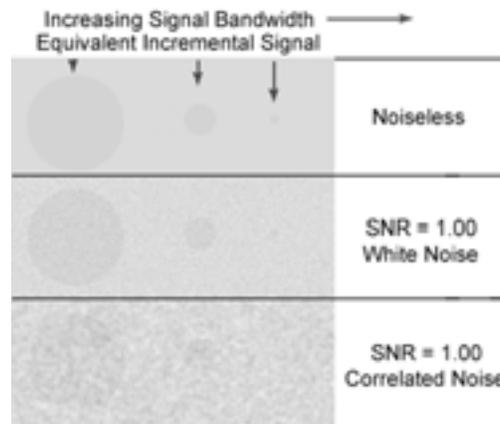


Figure 2. Signal and noise of varying bandwidth.

ity) of detecting any given signal is a function of not only the signal strength, but also the ambiguity that noise introduces. Adopting the ISO 21550 criterion, the density levels where the incremental SNR is equal to 1.0 define the dynamic range endpoints. Because it adopts an incremental signal criterion, its utility lies in quantifying how well a given object intensity, I_o , can be distinguished from another intensity of an arbitrarily small difference, ΔI . In the context of a noise source, it will answer questions like, "How well can this capture device distinguish between an optical density of say 3.00 and 3.10?" The Gaussian distributions of

Figure 1 help to convey this concept. For simplicity, the abscissa numerical values represent digital count value.

The distributions of Figure 1 are consistent with the limiting SNR of 1.0 for dynamic range: a mean incremental signal of 0.10 along with an RMS noise level of 0.10. If a single or average pixel value falls within the alpha risk region, one may accept that the value is not part of the 3.10 density population. What is more pronounced with this example is the large beta risk region associated with the SNR of 1.0. That is, the risk in deciding that a single pixel value (think high frequency) belongs to the 3.10 population when in fact it belongs to the 3.00 density population. Indeed, it is this risk that some may find unacceptable with respect to incremental signal detection reliability. Two ap-

proaches to this problem are to either set a higher SNR criterion for the dynamic range threshold or to reconsider what is meant by signal for a digital image.

The previous discussion, and the above standards, adopt a pixel-centric view of both signal and noise characteristics. The signal to be detected is a small difference in a single image sample, pixel value. The rms noise is calculated for distributions of individual pixel values.

Consider instead the image intensity distribution and physical extent of important 'signals' in images. Important signals (objects) will span more than a signal pixel. This can be interpreted in terms of an input signal spatial frequency spectrum. Correspondingly, image noise is not completely described by its pixel-by-pixel variance of rms error, as it is in the current standards. The spatial correlation for noise sources in image capture systems also implies spatial frequency content. It is this spatial frequency aspect to dynamic range evaluation that may well influence future standards. Figure 2 demonstrates how images with differing signal frequency content, but identical low-frequency signal-to-total noise values can appear dramatically different.

It is likely that future standards for imaging

Continues on page 10.

Image noise reduction and sharpening filter

Milivoje Aleksic, Maxim Smirnov, and Sergio Goma, ATI Technologies

One of the major trends in wireless communications is the incorporation of digital cameras into cellular phones. Typical resolution has increased from VGA (720×480) to 2Mpixel last year, and is likely to increase to 6Mpixel this year. Market analysts expect that sales of CMOS sensors in cellular phones will surpass the sales of CMOS sensors in digital still cameras (DSC) this year. The cell phone consumer expects to enjoy similar image quality to a DSC, despite cell phone camera limitations (like low lens quality and high noise). The key problem is how to display and capture this type of image with reasonable quality, and also how to compress them to minimize the storage space and product cost.

In general image processing there are a lot of different types of filters for improving visual image quality, but these do not typically consider compression size. Our solution is based on the bilateral filter, the original version of which combines two kernels: the *spatial*, which processes the pixels based on their geometrical distance, and *photometric*, which takes into account the perceptual similarity between the currently processed pixel and the pixels in its vicinity.^{1,2} This filter was introduced to reduce noise while preserving the edge structure.

Our approach to bilateral filtering enhances this basic concept in order to allow the filter to sharpen the image while at the same time reducing its noise level. We use two spatial kernels: one low-pass, and one high-pass. As a rule, though this is not a strict requirement, the two kernels are selected to be complementary.

$$W_s^{LP}[n, m] + W_s^{HP}[n, m] = 1$$

Here $W_p[n, m]$ is a photometric kernel calculated adaptively as a function of Euclidian distance D^2 between the current pixel \vec{x}_{ij} and its neighbors:

$$W_p[n, m] = f(D^2(\vec{x}[i, j], \vec{x}[i-n, j-m]))$$

Bilateral filtering combines the special and photometric kernels in one:

$$\vec{y}[i, j] = \frac{\sum_{n=-\frac{M-1}{2}}^{\frac{M-1}{2}} \sum_{m=-\frac{M-1}{2}}^{\frac{M-1}{2}} W_p[n, m] \times W_s[n, m] \times \vec{x}[i-n, j-m]}{\sum_{n=-\frac{M-1}{2}}^{\frac{M-1}{2}} \sum_{m=-\frac{M-1}{2}}^{\frac{M-1}{2}} W_p[n, m] \times W_s[n, m]}$$

N and M represent vertical and horizontal sizes of the special kernel correspondingly, and are assumed to be odd numbers. Note that \vec{x}_{ij} and \vec{y}_{ij} are scalars for grey-scale images and multi-component (usually three-component)

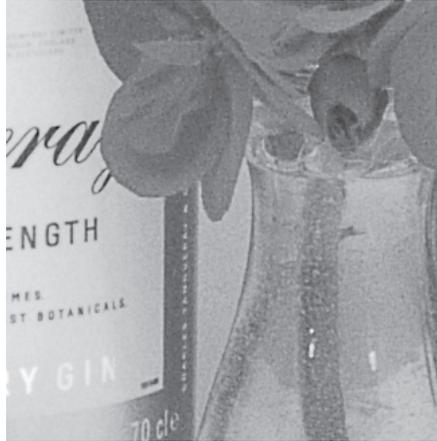


Figure 1. ISO400 detail, before use of the new filter.



Figure 2. ISO400 detail, after use of the new filter.

vectors for color images. The scaling denominator in the equation reflects the requirement for the filter to preserve the DC component of the input signal.

Please note that the scaling denominator in this equation is now a function of the input signal and, therefore, must be calculated for each pixel of the input image.

Traditional bilateral filter implementation assumes that both spatial and photometric kernels have Gaussian shape, and, the Euclidean distances between pixel values are calculated in CIE-lab perceptual color space although different kernel shapes and different methods of calculating the Euclidean distances are not precluded.

The adaptive noise removal/sharpening kernel is calculated as sum of low-pass (noise re-

duction) and high-pass (sharpening) components as:

$$W[n, m] = \{K_{sharp} \times W_p^{HP}[n, m] \times W_s^{HP}[n, m] - b_{HP} \times U\} + \{K_{LP} \times W_p^{LP}[n, m] \times W_s^{LP}[n, m]\}$$

Here, $W_p^{LP}[n, m]$ and $W_p^{HP}[n, m]$ are low- and high-pass photometric kernels respectively, b_{HP} and K_{LP} are filter normalizing coefficients, and U is kernel of a pass-through 2D filter.

$$U = \begin{bmatrix} 0 & 0 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 0 & 0 \end{bmatrix}$$

The low-pass photometric kernel is also calculated in a simplified way: it uses a flat window with a threshold parameter instead of regularly used Gaussian function. Normalizing the low-pass coefficient ensures that the low-pass part of the resulting kernel passes the DC component unchanged and it is consistent with the classical bilateral filtering equation:

$$K_{LP} = \frac{1}{\sum_n \sum_m W_p^{LP}[n, m] \times W_s^{LP}[n, m]}$$

The high-pass photometric kernel exploits an equation similar to its low-pass counterpart. The coefficient normalization procedure, however, is different: i.e. the middle element of the high-pass kernel is shifted so that the kernel's frequency response would be zero for the DC signal components:

$$b_{HP} = \sum_n \sum_m K_{sharp} \times W_p^{HP}[n, m] \times W_s^{HP}[n, m]$$

The adaptation process uses three parameters: the *noise removal threshold*, which is used to classify variations in the input image as noise or as edges; the *sharpening threshold*, which determines variation levels qualifying for edge sharpening; and the *sharpening coefficient*, which sets the desired amount of sharpening.

The results of applying this modified bilateral filter are simultaneous noise reduction and sharpening, with JPEG compression improved by a factor of 1.5 to 2.

Milivoje Aleksic, Maxim Smirnov, and Sergio Goma
ATI Technologies
Toronto, Ontario, Canada
E-mail: maleksic@ati.com

References:

1. C. Tomasi and R. Manduchi, *Bilateral Filtering for Gray and Color Images*, Proc. 1998 IEEE Int. Conf. Computer Vision, pp. 839-846.
2. Michael Elad, *On the Origin of the Bilateral Filter and Ways to Improve It*, IEEE Trans. Image Processing 11, pp. 1141-1151, October, 2002.

Gray tracking for LC displays

Gabriel Marcu, Apple Computer, Inc.

TFTLCD (thin-film transistors liquid crystal display) devices have proliferated in recent years: this is mostly due to their having high brightness, contrast, and sharpness; being almost entirely free of geometric image distortion; compact, and light; exhibiting no flicker; and consuming little power. In these transmissive displays, color is produced by an additive mixture of light passing through the triplets of red, green, and blue filters that form each pixel. The color produced by the display depends on the backlight source, the LC light modulator element, and the color filters.

One of the colors that is most difficult to reproduce using LC displays is gray. This is due to both the asymmetry of the transfer function of the R, G, and B channels and the higher sensitivity of the human visual system to color differences in grays rather than in saturated colors. Correct gray reproduction requires a very precise control of the balance between the color channels. In particular, on the twisted nematic (TN) displays with a significant asymmetric RGB native response (the most common type of LC found in the portable devices today), the grays tend to be bluish, an effect that is more noticeable for middle and darker grays.^{1,2} A typical color shift measured in differences of correlated color temperature can reach more than 3000K across the entire dynamic range of these displays. For displays where color reproduction is critical, correcting this color shift is important. This article presents a method of doing this.

The idea is to perform the gamma and gray-tracking correction in a single process using both the luminance and chrominance information of the individual RGB input channels. This method represents a further refinement of the method introduced in Reference 3. For a target gamma and white point, this technique involves computing the RGB values input to the display such that the resulting gray has the minimum color difference to the target color. The method ensures that the solution is optimal, in the sense that no other set of RGB values will be closer to the desired color.

The algorithm runs as follows. First a target gamma correction is specified in luminance steps, $Y_j, j = 1, \dots, N$. A target white point (or gray point along the gray levels) is specified as a set of chromaticity coordinates $(x, y)_j, j = 1, \dots, N$. With this, the gray tracking target is completely specified as a set of target grays, $(x_p, y_p, Y_p), j = 1, \dots, N$. Then the native response of the panel is measured in luminance and chromaticity for each channel, $(x_r, y_r, Y_r)_k, (x_g, y_g, Y_g)_k,$

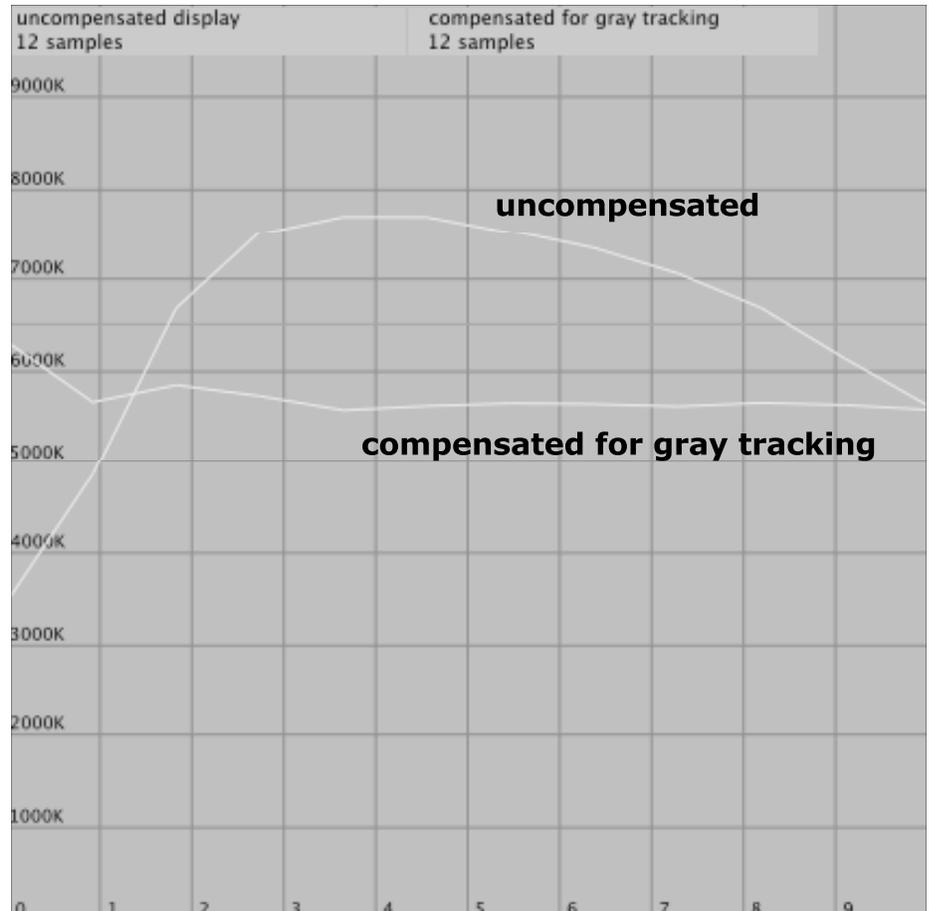


Figure 1. The correlated color temperatures of panels with and without gray tracking compensation. The target white point (set to 5700K) of the panel with compensation is much more stable across its dynamic range than that without.

$(x_b, y_b, Y_b)_k, k = 1, \dots, M$. Then, the algorithm runs a searching procedure to find—from all combinations of red, green, and blue values—the one that minimizes the chromatic color difference to the target gray.

The algorithm can be described in pseudo code in the following steps:

1. Select the target white, $W_{[ij]} = (Y_p, x_p, y_p)$, to which the panel is corrected. The gamma correction is done here, the number of values on which gamma correction is specified depends on the panel.
2. For each input $W_{[ij]}, j = 1$ to N {

Set current minimum color difference, D , to a large value for $r = 1$ to $M, g = 1$ to $M, b = 1$ to M

Compute color $C = \text{AdditiveMixture}(r+g+b);$

```

 $\Delta E = C - W$ 
if(  $\Delta E \leq D$  ) {
     $R_{[ij]} = r, G_{[ij]} = g, B_{[ij]} = b;$ 
     $D = \Delta E;$ 
}
}

```

The method works fine for most LCD panels, even for most with moderate crosstalk correlation between the R, G, B channels. For panels with severe crosstalk, the additive mixing formula may lead to inaccurate results and a replacement of this formula with an accurate prediction function may be required.⁴

Continues on page 10.

Gray tracking for LC displays

Continued from page 9.

The gamma and gray tracking corrections for several TFT-LCD panels were computed and evaluated. A typical result, presented in the form of dependency of the correlated color temperature on the input level, is shown in Figure 1. It can be observed that the bluish color shift is corrected and the white point is stable across almost the entire dynamic range. The fluctuation of the correlated color temperature at low gray levels is caused by the instrument used to measure the color.

It was found that the panels with corrected gray tracking render the images more naturally than those using independent R, G, B gamma correction for the same target gamma. In particular, the grays and flesh tones benefit from this correction. The images' bluish cast was removed and the naturalness of the images was restored.

It was found that the gray-tracking compensation was effective in removing even the color banding for adjacent gray levels (caused by the time-frame modulation). This was possible because the gray-tracking correction was done for each individual gray level and was independent of the color correction of adjacent levels. The only constraint was on the monotony of each of the resulted R, G, and B correction curves. However, due to the limited number of bits for coding the input R, G, and B values, not all color banding can be removed: in some high-brightness displays color banding may still be visible in the dark grays. To completely correct for these artifacts, the R, G, and B input values should be coded with more than 8bits/channel (the number of input entries may remain unchanged).

It is important to note that the gray-tracking individual RGB gamma corrections may sometimes conflict, with each correction imposing a set of curves different from the set imposed by the other. In our experience, precise gray tracking is far more important than exact gamma correction, so we sacrifice the correctness of RGB individual-channel gamma correction for perfect gray tracking and perfect gamma correction for grays.

This method can accommodate the color shift from various sources such as time modulation, flare, wavelength dependency, and chromaticity variation with the input level. This means the method is applicable to other display technologies, and is not limited to LC displays.

Gabriel Marcu

Apple Computer, Inc.
Cupertino, CA
E-mail: marcu@apple.com

References

1. Gabriel Marcu, Wei Chen, Kok Chen, Peter Graffagnino, and Olav Andrade, *Color characterization issues for TFT-LCD displays*, **Proc. SPIE 4664**, January 2002.
2. Louis D. Silverstein, *Initial Investigation of Optiva TCF Material for LCD Color Correction*, **Optiva Inc. Internal Technical Report**, February 28, 2002.
3. G. Marcu and K. Chen, *Gray tracking correction for TFTLCDs*, **10th IS&T/SID Color Imaging Conf.**, Scottsdale, AZ, 2002.
4. M. H. Brill and L. D. Silverstein, *Isoluminous Color Difference Metric for Application to Displays*, **Proc. SID 2002**, p. 809, Boston, May 2002.

Target tracking using soliton patterns in a large-array optoelectronic system

Continued from page 12.

ated at the head and foot of the figure. The final resulting image, captured after 175τ , shows the trajectory of the solitons as the figure moved across the image.

Numerical simulations for this optoelectronic model show that a soliton pattern can be generated for moving objects within images with stationary backgrounds, thus establishing a trajectory indicating the movement of the objects within the scene. Target-tracking information can therefore be obtained from the resulting image.

Adrienne Raglin

Army Research Laboratory
Adelphi, MD
E-mail: raglin@arl.army.mil

References

1. M. A. Vorontsov and B. A. Samson, *Parallel nonlinear optoelectronic image processing for real-time motion detection*, **Optical Engineering** **38** (3), pp. 558-563, 1999.
2. W. Lu, S. L. Lachinova, and R. G. Harrison, *Detection and tracking of small moving objects in image sequences by use of nonlinear spatiotemporal optical systems*, **Optics Lett.** **29** (8), 2004.
3. A. J. Raglin and M. Chouikha, *Competitive Dynamics and Pattern Formation in a Large Array of Opto-electronic Feedback Circuit System*, **Proc. SPIE 5298**, pp. 349-361, 2004.

Demosaicking techniques for multispectral cameras

Continued from page 3.

puts to the MAP-based demosaicking block, and the results demonstrated that the method successfully performs demosaicking, while at the same time reducing noise and degradation. Comparisons have shown that the MAP method performs better with traditional interpolation-based methods. We have also tested the demosaicked multispectral images for real-world multispectral applications like target recognition. We have shown that the mosaicked technology does not negatively affect classification accuracy. In fact, in some cases, accuracy can be improved because of degradation removal.

Gaurav Baone and Hairong Qi

Advanced Imaging and Collaborative
Information Processing (AICIP)
Department of Electrical and Computer
Engineering
University of Tennessee, Knoxville TN
E-mail: gaurav_baone@yahoo.com

References

1. L. Miao, H. Qi, and W. Snyder, *A generic method for generating multi-spectral filter arrays*, **IEEE Int'l Conf. on Image Processing (ICIP)**, Singapore, October 24-27, 2004.
2. H. Andrews and B. Hunt, **Digital Image Restoration**, 1st edition, Prentice Hall Inc., 1977.

ISO dynamic range interpretation: risky business

Continued from page 7.

performance evaluation will be based on a more complete interpretation of SNR. For dynamic range, this would require an approach that accounts for not only different risk levels but also spatial, and perhaps temporal, characteristics of signal and noise. This would be consistent with a spatial frequency view of information content.² This approach could include the modulation transfer function for describing signal capture, and the noise power spectrum for image noise evaluation.

Don Williams and Peter D. Burns

Kodak Research Laboratories
Eastman Kodak Company
Rochester, NY
E-mail: {don.williams,
peter.burns}@kodak.com

References

1. P. D. Burns and D. Williams, *Low-frequency MTF estimation for digital imaging devices using slanted-edge analysis*, **Proc. SPIE 5294**, pp. 93-101, 2004.
2. J. C. Dainty and R. Shaw, **Image Science**, Academic Press, ch. 10, 1974.

Volumetric depth peeling for virtual arthroscopy

Continued from cover.

ments, it is less invasive to view MRI images.

By viewing slices of MRI data along the three orthogonal axes (see Figure 2), clinicians can diagnose joint pathologies without the use of invasive arthroscopic procedures. Although joint pathologies can be determined from these images, diagnosis requires tedious manual selection of slices and a difficult mental re-construction of 3D structures from slices along one axis, correlated with features along the other two axes.

Virtual arthroscopy enables real-time evaluation of joint surfaces in 3D, removing this tedious slice manipulation and difficult 3D mental reconstruction. However, due to the close proximity of surfaces in a joint, radiologists find existing methods lacking for obtaining desired views of features of interest.

Volumetric depth peeling

VDP was developed to enable radiologists to obtain views external to the space between the bone and cartilage surfaces ('joint space') for viewing entire areas of interest. Such views are impossible using standard volume rendering because the surfaces occlude each other from viewpoints outside the joint space.

VDP extends standard volume rendering by

automatically culling occluding voxels between the viewpoint and the features of interest, enabling unobstructed views of entire regions within the structure being examined. VDP decouples occlusion calculation from the volume rendering transfer function, enabling independent optimization of settings for rendering and occlusion. No pre-segmentation of the dataset is required, and VDP handles multiple layers of occluding material, as well as self-occlusion.

A comparison of virtual arthroscopy using both standard volume rendering and virtual arthroscopy using VDP is shown in Figure 3. On the left is an image from within the joint space, with the humeral head on the left and the shoulder socket and related structures on the right. Images such as this demonstrate the restricted choice of views when using standard volume rendering. The physician would like to be able to see the entire surface of the socket, a view only available from a position within the humeral head. However, the material within the humeral head will occlude such a view.

On the right is an image created using VDP. The viewpoint is moved to a position within the humeral head, and the occluding voxels of the humeral head are automatically culled, en-

abling a view of the entire upper portion of the socket. Note the clearly visible torn ligament circled in white. VDP also enables flying through the socket, turning around, and viewing the surface of the humeral head.

Full studies of VDP for virtual arthroscopy and virtual urography are currently underway. VDP has also been used for displaying CTs of bone fractures, brain aneurysms, and non-medical datasets such as CTs of engine blocks. A more in-depth discussion of the technique will be presented in a forthcoming publication.

David Borland, John P. Clarke*, and Russell M. Taylor II

Department of Computer Science

*Department of Radiology

University of North Carolina at Chapel Hill
Chapel Hill, NC

E-mail: borland@cs.unc.edu

References

1. Z. Nagy and R. Klein, *Depth-Peeling for Texture-Based Volume Rendering*, 11th Pacific Conf. on Computer Graphics and Applications, p. 429, 2003.
2. D. Stoller and P. Tirman, *Diagnostic Imaging: Orthopaedics*, Amirsys, 2004.

Exciting changes coming for SPIE publications

Beginning in January 2006, two new publications will take the place of oemagazine as well as our Technical Group newsletters. SPIE Newsroom, a dynamic news website integrated with spie.org, will cover the latest technical developments in optics and photonics. SPIE Professional, a quarterly print magazine published exclusively for Society members, will emphasize career trends and industry insights associated with the optics and photonics profession, as well as important Society news and information. This is the final issue of the Electronic Imaging Technical Group newsletter.

The website will provide a steady stream of technical and industry news articles organized around technical communities that match SPIE members' areas of expertise and interest, including electronic imaging, biomedical optics, industrial sensing, defense and security, nanotechnology, and others. Readers can subscribe to monthly e-alerts to be directed to the latest technical articles in their field. SPIE Professional will offer members a new perspective on their industry and their diverse roles in it.

Both of these new efforts will launch in January 2006 - keep an eye on spie.org for the latest updates.

Resolution and light sensitivity tradeoff with pixel size

Continued from page 5.

different pixel size, scene luminance, and exposure duration.

Joyce Farrell and Feng Xiao*

Stanford University, Stanford, CA

*Agilent Technologies, Santa Clara, CA

E-mail: feng_xiao@agilent.com

Catrysse, and B. Wandell, *A simulation tool for evaluating digital camera image quality*, **Proc. SPIE 5294**, pp. 124-131, 2004.

5. <http://www.imageval.com>.
6. D. Williams and P. D. Burns, *Low-Frequency MTF Estimation for Digital Imaging Devices using Slanted Edge Analysis*, **Proc. SPIE 5294**, pp. 93-101, 2003.

References

1. T. Chen, P. Catrysse, A. El Gamal, and B. Wandell, *How Small Should Pixel Size Be?*, **Proc. SPIE 3965**, pp. 451-461, 2000.
2. P. Catrysse and B. Wandell, *Roadmap for CMOS image sensors: Moore meets Planck and Sommerfeld*, **Proc. SPIE 5678**, pp. 1-13, 2005.
3. F. Xiao, J. E. Farrell, and B. Wandell, *Psychophysical thresholds and digital camera sensitivity: The thousand photon limit*, **Proc. SPIE 5678**, pp. 75-84, 2005.
4. J. E. Farrell, F. Xiao, P.

Table 1: Optics and Sensor Parameters

Optics and Sensor	Parameter Settings					
Pixel Sizes (μm)	1.7	2.2	2.8	3.3	5.2	7.4
Read noise (e-)	10	15	25	30	35	43
Dark noise (e-/sec)	80	100	120	150	200	240
Conversion gain ($\mu\text{V}/\text{e-}$)	60	40	30	25	18	13
Voltage swing (V)	0.7	0.8	0.9	1.0	1.1	1.2
Fill factor	50%					
Peak QE (550nm)	0.65					
Dye size (mm)	0.512					
Imaging lens F#	2.8					
Integration time (sec)	0.033					
Scene illumination	D65 light					

Target tracking using soliton patterns in a large-array optoelectronic system

Adrienne Raglin, Army Research Laboratory

Target tracking is a major requirement for surveillance, security, and obstacle avoidance. Previous systems have used differencing between frames or registration of an object's optic flow to perform target tracking. Recent research in spatiotemporal nonlinear dynamics has provided an additional approach that uses optoelectronic systems that combine the parallel nature of optics and the computational strength of electronics. With the use of application-specific hardware these systems can potentially exploit their parallel-processing capability for high-resolution image processing including edge enhancement and motion tracking.^{1,2}

We presented in Reference 3 a large array of optoelectronic feedback circuit systems that produce a variety of patterns including the soliton. A soliton can be considered as a localized state visualized as a bright circular spot. The soliton behaves in a similar way to a neural-network node: the stronger the stimulus is above a threshold value, the stronger the response of the soliton. The soliton pattern of an object can provide a different representation of the target of interest within an image, indicating the location of the target and marking its trajectory as it moves within the scene.

In the model for this system (see Figure 1) an initial image is applied to the system at $t=0\tau$, and an injected image is applied to the system for a set time interval. This mimics the real-

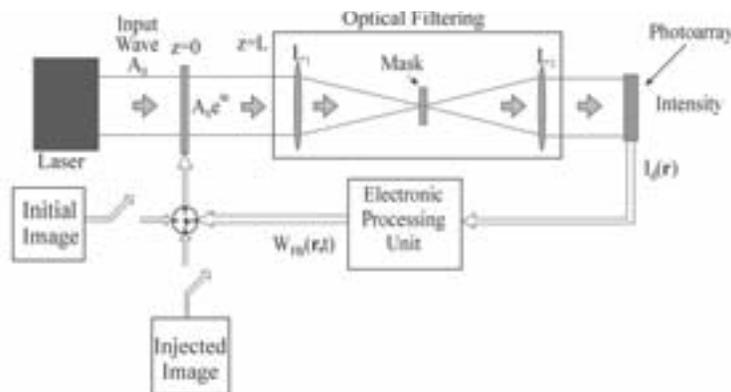


Figure 1. Schematic model for a large-array optoelectronic feedback circuit that generates a series of spatial patterns.

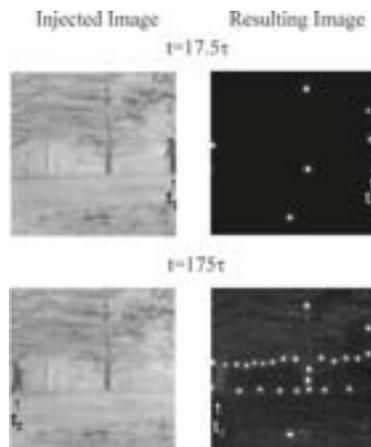


Figure 2. Simulation results: two resulting images, captured at $t_1=17.5\tau$ and $t_2=175\tau$, show solitons at the top and bottom of the figure as it moves across the image.

time capture of an image from a camera that has a set update rate. The laser acts as a light source and the image is transformed from an electrical to an optical signal through interaction with a phase-only spatial light modulator (SLM) that modulates the wavefront phase. After transmission through the SLM the wave is diffracted. It is then transmitted through a Fourier filter that consists of two confocal lenses with a Gaussian-shaped mask at the joint focal plane. The optical field from this filter is registered onto a photoarray, the intensity distribution of which is used as the input

to the electronic processing unit that applies a nonlinear map to obtain a feedback signal. This feedback signal is applied to the SLM closing the loop.

The change in the phase (which represents the changes to the image) due to the feedback signal is described by the following equation:

$$t \frac{\partial \phi(r,t)}{\partial t} + \phi(r,t) = D \nabla_{\perp}^2 \phi(r,t) + K W_{10}(r,t)$$

where t is the system response time, D is the diffusion coefficient,

∇_{\perp}^2 is the transverse Laplacian in the x - y plane, and K is the feedback gain coefficient.

In the simulation, an infrared image of a single figure walking through a wooded area was used (see Figure 2). Solitons are gener-

Continues on page 10.

SPIE Society of Photo-Optical
Instrumentation Engineers

P.O. Box 10 • Bellingham, WA 98227-0010 USA

Change Service Requested

DATED MATERIAL

Non-Profit Org.
U.S. Postage Paid
Society of
Photo-Optical
Instrumentation
Engineers