Subject: Re: Medical Imaging Question
Posted by Struan Gray on Fri, 06 Aug 1999 07:00:00 GMT
View Forum Message <> Reply to Message

David Fanning, davidf@dfanning.com writes:

> I want to "see" 4094 shades of gray.

This is hard. A really good monitor designed for photogrammetry might get you to 8-9 bits of true greyscale, especially if you use duo- or tri-toning as you suggested, but 12 bits is hopeless. The only way to display 12 linear bits is by writing onto sheet film with a well-calibrated laser-based film recorder, and hanging it on a lightbox - and even there you're pushing the envelope.

Then you have to deal with the human eye. Visual perception (as opposed to the physics of the retina) is logarithmic: a rough figure is that the eye-brain system can distinguish intensity variations of the order of 1%. With linear encoding of intensity you are wasting bits in the bright areas of your image (where the bit transitions are packed more densely than 1%) and losing detail in the shadows (where, at worst case the bits represent 50% variations in intensity). This, along with technical issues to do with the response of TV tubes, is why a gamma function helps, ie I = I^gamma looks more detailed and even than I alone, with gamma = 0.45 being so-called optimal for the human eye.

People say that 12-bit linear encoding is about the maximum the human eye can resolve, unless the scene subtends a very large angle and the viewer can dark-adapt their eyes to individual sub-scenes. 8-bit gamma encoding of this looks pretty good, but some infomation is lost and the best strategy for monitor-based viewing is to keep the 12-bit information and allow the user to scan and zoom around the image, creating locally-valid 8-bit gamma versions of the information as appropriate. In IDL this would be simple.

Your final option is to do some image processing. There are several well-established ways of enhancing detail in images with a large dynamic range. Photoshop users (and darkroom enthusiasts) are familiar with the unsharp mask. I use a technique called Statistical Differencing, which is essentially an unsharp mask weighted by the local statisitics: it applies a more agressive mask in areas which lots of small detail. Plotting in light-shaded form in 3D can be surprisingly effective, which is mathematically the same as the common trick of adding the local derivative to the original image. All these tricks help the viewer see detail at the expense of the local average intensity.

I culled these references from a discussion of the human eye and perception in rec.photo.digital, they might be worth a look if you're really interested:

The Reproduction of Colour (in Photograpy, Printing, and Television) by R.W.G. Hunt.

Illumination and Color in Computer Generated Imagery, by Roy Hall

Digital Color Management, by Edward Giorgianni and Thomas Madden

Color Appearance Models, by Mark Fairchild

Struan
(Apologies if you see more than one copy of this - my newsserver promised me it hadn't accepted the first one(s))
