Subject: Re: Medical Imaging Question
Posted by Axel vom Endt on Thu, 05 Aug 1999 07:00:00 GMT
View Forum Message <> Reply to Message

David Fanning wrote:
>
> I mean by this that I want to "see" 4094 shades of gray.

I'm not sure about the reference for this, but I read somewhere that the human eye cannot even distinguish between 256 levels of gray. Isn't that the reason to use the fancy color tables in IDL?

From what I understand about medical imaging the 4096 levels are meant to have some headroom, so that you can zoom into some parts of the intensity range to enhance certain features of the image without losing resolution.

I think that these fudged gray scales have only been used in the old days of the PC, when you had to use 256 colors out of 262144 on your VGA card, i.e. 2^6 different shades each of r,g,b and therefore only 64 levels of pure gray. I've done that once, but the results haven't been impressive.

Hope that helps

Axel

Subject: Re: Medical Imaging Question
Posted by davidf on Thu, 05 Aug 1999 07:00:00 GMT
View Forum Message <> Reply to Message

Mirko (mirko_vukovic@notes.mrc.sony.com) writes:

> davidf@dfanning.com (David Fanning) wrote:
>>
>> I've got a 12-bit Dicom image. I want to display this in
>> such a way that there is a direct correlation between
>> the pixel value (0 to 4094) and the representation of
>> that pixel value on the display. How do I do that?
>>
>
> Can you please expand on that? What do you mean by direct correlation?
> For a 4-bit image would color be a good correlation
> (16 values <=> 16 colors/intensities)?

I mean by this that I want to "see" 4094 shades of gray.
Now I know there are only 256 "pure" shades of gray, so

I presume that I have to fudge a gray somehow. For example,
the pure gray (128, 128, 128) could be fudged by something
like (128, 128, 135). The latter is not really a gray, but
it's almost gray and presumably it "looks" different than
the pure gray. My question really boils down to this: Is there
some standard way to "fudge" gray values between the pure
gray values, so that it appears as though I have 4094
shades of gray? And if so, what is the algorithm to
go from the pixel representation on the display to the
real pixel value?

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting
Phone: 970-221-0438 E-Mail: davidf@dfanning.com
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Toll-Free IDL Book Orders: 1-888-461-0155

Subject: Re: Medical Imaging Question
Posted by mirko_vukovic on Thu, 05 Aug 1999 07:00:00 GMT
View Forum Message <> Reply to Message

In article <MPG.12134dfd2c8c96a7989871@news.frii.com>,
davidf@dfanning.com (David Fanning) wrote:

>
> I've got a 12-bit Dicom image. I want to display this in
> such a way that there is a direct correlation between
> the pixel value (0 to 4094) and the representation of
> that pixel value on the display. How do I do that?
>

Can you please expand on that? What do you mean by direct correlation?
For a 4-bit image would color be a good correlation
(16 values <=> 16 colors/intensities)?

Mirko

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

Larry Busse (ljb@ljbdev.com) writes:

> In MR and CT, images are usually displayed with a gray map that is
> appropriate for the particular area being imaged or diagnostic
> procedure. The look-up tables are sometimes referred to as
> "lung-window", or "bone-window", or "soft-tissue-window". These
> correspond to different window/level settings where
> window = (WhiteValue - BlackValue) and
> level = (WhiteValue + BlackValue)/2.

Let me see if I understand this correctly. Are you saying
that I might have sliders that would select a "window" of
data. Say between the values of 1000 and 3500, and that
what I would see on my display would be something like
this:

```
TV, BytScl(image, Max=3500, Min=100, Top=!D.Table_Size-1)
```

In other words, the gray scale values could be a portion
or window onto the entire data universe. If this is so,
how do you usually implement such a sliding window into
your data?

Many thanks to all (including bashful e-mail senders)
for their comments. :-)

Cheers,

David
--
David Fanning, Ph.D.
Fanning Software Consulting
Phone: 970-221-0438 E-Mail: davidf@dfanning.com
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Toll-Free IDL Book Orders: 1-888-461-0155

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

David,

I'm no expert on DICOM but I've worked in enough Radiology departments

to shed a little light on your question.

As you probably know, the DICOM standard was developed to allow the many different kinds of medical imaging equipment (CT, MRI, Ultrasound, Nuclear Medicine, etc.) to provide data in a format that could be handled by workstations and computer networks throughout hospitals. I think the notion of digitizing images to 12-bits (or 4096 gray levels) was done to preserve the full dynamic range of older recording media such as x-ray film. The doctors basically didn't want to loose any information (contrast resolution or spatial resolution) just because the data was going to be digitized and sent over a computer network.

In MR and CT, images are usually displayed with a gray map that is appropriate for the particular area being imaged or diagnostic procedure. The look-up tables are sometimes referred to as "lung-window", or "bone-window", or "soft-tissue-window". These correspond to different window/level settings where
window = (WhiteValue - BlackValue) and
level = (WhiteValue + BlackValue)/2.

and the curve from black to white can be linear of some other curve. If the window is greater than 256 and the display only capable of 256 levels, then you won't see every level.

All this meant to say that I don't think it's necessary (or possible) to see 4096 shades of gray on the screen but it is important to preserve all the info in the data because looking for a tumor in the lung will require much different display settings that looking for a tumor in a bone.

I hope this helps. For more general info on DICOM try
http://www.rsna.org/REG/practiceres/dicom/

Good Luck,

David Fanning wrote:
>
> Mirko (mirko_vukovic@notes.mrc.sony.com) writes:
>
>> davidf@dfanning.com (David Fanning) wrote:
>>>
>>> I've got a 12-bit Dicom image. I want to display this in
>>> such a way that there is a direct correlation between
>>> the pixel value (0 to 4094) and the representation of
>>> that pixel value on the display. How do I do that?
>>>
>>
>> Can you please expand on that? What do you mean by direct correlation?

>> For a 4-bit image would color be a good correlation
>> (16 values <=> 16 colors/intensities)?
>
> I mean by this that I want to "see" 4094 shades of gray.
> Now I know there are only 256 "pure" shades of gray, so
> I presume that I have to fudge a gray somehow. For example,
> the pure gray (128, 128, 128) could be fudged by something
> like (128, 128, 135). The latter is not really a gray, but
> it's almost gray and presumably it "looks" different than
> the pure gray. My question really boils down to this: Is there
> some standard way to "fudge" gray values between the pure
> gray values, so that it appears as though I have 4094
> shades of gray? And if so, what is the algorithm to
> go from the pixel representation on the display to the
> real pixel value?
>
> Cheers,
>
> David
>
> --
> David Fanning, Ph.D.
> Fanning Software Consulting
> Phone: 970-221-0438 E-Mail: davidf@dfanning.com
> Coyote's Guide to IDL Programming: http://www.dfanning.com/
> Toll-Free IDL Book Orders: 1-888-461-0155

--
Larry Busse   Phone: (606)344-9433
L.J.B. Development   Fax: (606)341-8486
3384 Madison Pike   mailto:larry@ljbdev.com
Fort Wright, KY 41017   http://www.ljbdev.com
~~~~~~ ~~~~~

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

Struan Gray (struan.gray@sljus.lu.se) writes a lengthy
verson of "it ain't gonna happen":

Thanks for this information, Struan. This is basically
what I thought, but didn't have the background to
support my beliefs. I'm going to write it up as an
article, if you don't mind, so that I don't have to keep
answering this question over and over. :-)
~~~~~~

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting
Phone: 970-221-0438 E-Mail: davidf@dfanning.com
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Toll-Free IDL Book Orders: 1-888-461-0155

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))
