
Subject: Medical Imaging Question

Posted by [davidf](#) on Thu, 05 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi Folks,

Here is a medical imaging question that I am asked often enough that my standard "ain't gonna happen" answer is wearing thin, even for me. Perhaps it *would* happen if I understood a little more about it. In exchange for help, I'll write an article and make it available for everyone. :-)

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?

A quick search of the Internet suggests that I'm not the first person to think about this issue, but I can't seem to put my hands on the documents that will explain it to me. Any help is greatly appreciated.

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 information 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 statistics: it applies a more aggressive 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 Photography, 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

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 information 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 statistics: it applies a more aggressive 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 Photography, 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

Subject: Re: Medical Imaging Question
Posted by [levick](#) on Sun, 08 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

In article <37AA0BAA.99AD6EC5@bu.edu>,
Axel vom Endt <endt@bu.edu> wrote:
> 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?
>

Incorrect. Most studies find that the number of brightness jnd's is about 450.

Ex

Sent via Deja.com <http://www.deja.com/>
Share what you know. Learn what you don't.

Subject: Re: Medical Imaging Question
Posted by [Struan Gray](#) on Mon, 09 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Fanning, davidf@dfanning.com writes:

- > This is basically what I thought, but didn't
- > have the background to support my beliefs.

Well, the eye stuff I've culled from other sources. Fortunately this is not the only newsgroup with knowledgeable posters :-)

- > 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. :-)

Sure. You might want to emphasise that modelling the eye as an n-bit linear detector is a very dodgy approximation, and even using a fixed point representation (which is closer to logarithmic seeing) only gets you so far. The brain is very good at outwitting simplistic models, both by concentrating attention and by learning over time.

I think another poster said that medics don't need the whole 4096 levels by the time they've identified what they're interested in. Something like a chest X-ray is a classic survey problem, where you need lots of detail, both spatial and spectral, but don't know exactly where until after you've taken the data. This is one of the last redoubts of big pieces of photographic film, along with wide-area sky surveys for transient things like comets, where the problem is much the same.

Struan

Subject: Re: Medical Imaging Question
Posted by [pford](#) on Wed, 11 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Fanning (davidf@dfanning.com) wrote:
: 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, ByteScl(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

TV, ByteScl(image, Max=3500, Min=100, Top=!D.Table_Size-1)
is basically what is done. There is usually a LUT associated with it that
references the pixel value to the intensity or color that may not be
'linear.' The window, as explained earlier, depend on the target.

One of the problems with the above scheme with nuclear medicine images is
that there may be a few pixels that are several magnitude larger than all the
other pixel, therefore using a range 0-100, the max value is set at 100
and everything else falls into the range 0 to 10 for example. This can be
corrected by truncating the max pixel value. Unfortunately, the vendors
seem to be clueless how to do this other than manual trial and error
method.

Regards

Patrick Ford, MD
Baylor College of Medicine
pford@bcm.tmc.edu

Subject: Re: Medical Imaging Question
Posted by [David Fenyes](#) on Fri, 13 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello David,

- 1) Of course you can physically do this. One way is to find or build a special CRT (or any other technology) that can display that kind of greyscale. Of course you can dither the image in magnification, then optically (or neuronally) blur. Alternatively, change each pixel into an exposure time (meaning the pixel is on for e.g. 0-4.095 seconds), then expose a film in a dark room (or find a souped up laser filmer). Develop and place against a lightbox in a dark room to appreciate. Or write software to time average on the retina (like the HP48 calculator greyscale programs do).
- 2) Of course, you probably don't really want to do this. Radiologists, especially mammographers and chest radiologists, are famous for wanting all the dynamic range. However, they like to sit in total darkness, and often use blinders. Even so, they often use the hotlight, or have the tech expose to get the object of interest in their visual sweet spot. In otherwords, having all the dynamic range there on the film doesn't free you from physically windowing and leveling to see your object of interest.
For digital images such as CT, images are routinely printed with several window/level settings to look at bone, solid organs, excreting kidneys, brain, etc. On the scanner, they are constantly playing with the settings. This is no handicap. Like the lightboxes and the exposure settings for plain film, it's a natural way of making sure the important pathology patterns register on the neural networks in the retina and occipital lobe.
- 3) Naturally, you still want to keep the numerical image as accurate as possible while window/level mapping it onto the more practical 256 levels. That way, if you perform any computations on the image, you have all the precision. Computers don't suffer from the perception limitations that humans do :-)

Regards,

David

--

David Fenyes -- dfenyes@home.com

Subject: Re: Medical Imaging Question
Posted by [m218003](#) on Mon, 16 Aug 1999 07:00:00 GMT

In article <7os14l\$kr@gazette.bcm.tmc.edu>,
pford@bcm.tmc.edu (Patrick V. Ford) writes:
> David Fanning (davidf@dfanning.com) wrote:
[...]

> One of the problems with the above scheme with nuclear medicine images is
> that there may be a few pixels that are several magnitude larger than all the
> other pixel, therefore using a range 0-100, the max value is set at 100
> and everything else falls into the range 0 to 10 for example. This can be
> corrected by truncating the max pixel value. Unfortunately, the vendors
> seem to be clueless how to do this other than manual trial and error
> method.
>
>

maybe I am too loud here, but shouldn't this kind of problem be easily
recognized by standard statistical outlier tests? That almost screams for
Struan's beloved histogram function, doesn't it? If you need something more
sophisticated, it appears that this problem is related to the problem of
determining biomass burning fires on satellite images (there they are looking
for the hot spots you are trying to exclude). Basically, one would look for
outlier values and reject them only if no neighbouring pixel shows similarly
high values. But, of course, this takes some processing time...

Cheers,
Martin.

Martin Schultz -- MPI fuer Meteorologie, Bundesstr. 55, 20146 Hamburg
martin.schultz@dkrz.de

Subject: Re: Medical Imaging Question
Posted by [Patrick V. Ford](#) on Mon, 16 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

In article <MPG.1221c12492ca7015989895@news.frii.com>, David Fanning
<davidf@dfanning.com> wrote:

> Martin Schultz (m218003@modell3-d.dkrz.de) writes:
>
>>> One of the problems with the above scheme with nuclear medicine images is
>>> that there may be a few pixels that are several magnitude larger than all
>>> the
>>> other pixel, therefore using a range 0-100, the max value is set at 100
>>> and everything else falls into the range 0 to 10 for example. This can be
>>> corrected by truncating the max pixel value. Unfortunately, the vendors

>>> seem to be clueless how to do this other than manual trial and error
>>> method.
>>>
>>>
>> maybe I am too loud here, but shouldn't this kind of problem be easily
>> recognized by standard statistical outlier tests? That almost screams for
>> Struan's beloved histogram function, doesn't it? If you need something more
>> sophisticated, it appears that this problem is related to the problem of
>> determining biomass burning fires on satellite images (there they are
>> looking
>> for the hot spots you are trying to exclude). Basically, one would look for
>> outlier values and reject them only if no neighbouring pixel shows
>> similarly
>> high values. But, of course, this takes some processing time...
>
> This sounds like a Median filter to me. Fast, easy, and
> a hell of a lot easier than trying to work through the
> intricacies of the Reverse_Index keyword. :-(
>
> Cheers,
>
> David
>
> P.S. I have to confess that I *have* written a modified
> median filter function for a client recently, which only
> applies the filter to user-selected high pixels (cosmic
> rays screwing up the CCD camera in this case). To implement
> it properly I needed to use a Histogram. :-)

I know that this can be easily solved automatically since DesAcc inc did it with ImportAccess. Vendors don't want to spend effort on the non sexy things like this and spend their time of iterative reconstruction methods. The advantage of the histogram method is that it still permits quantification. The too hot area are usually not the area of interest.

Regards

Patrick Ford

Subject: Re: Medical Imaging Question
Posted by [davidf](#) on Mon, 16 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Martin Schultz (m218003@modell3-d.dkrz.de) writes:

>> One of the problems with the above scheme with nuclear medicine images is
>> that there may be a few pixels that are several magnitude larger than all the
>> other pixel, therefore using a range 0-100, the max value is set at 100
>> and everything else falls into the range 0 to 10 for example. This can be
>> corrected by truncating the max pixel value. Unfortunately, the vendors
>> seem to be clueless how to do this other than manual trial and error
>> method.

>>

>>

> maybe I am too loud here, but shouldn't this kind of problem be easily
> recognized by standard statistical outlier tests? That almost screams for
> Struan's beloved histogram function, doesn't it? If you need something more
> sophisticated, it appears that this problem is related to the problem of
> determining biomass burning fires on satellite images (there they are looking
> for the hot spots you are trying to exclude). Basically, one would look for
> outlier values and reject them only if no neighbouring pixel shows similarly
> high values. But, of course, this takes some processing time...

This sounds like a Median filter to me. Fast, easy, and
a hell of a lot easier than trying to work through the
intricacies of the Reverse_Index keyword. :-)

Cheers,

David

P.S. I have to confess that I *have* written a modified
median filter function for a client recently, which only
applies the filter to user-selected high pixels (cosmic
rays screwing up the CCD camera in this case). To implement
it properly I needed to use a Histogram. :-)

--

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 Tue, 17 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Pavel Romashkin (promashkin@cmdl.noaa.gov) writes:

> If the 4096 grayscale levels are needed for *visual*
> analyses then is it not possible to map them to color

- > scale - I'd imagine a human eye should be able to
- > recognize variety in color easier than in grayscale.

There's a technical problem, in that you need to find 4096 different points in colour space that can be represented on a monitor (easy), which allow adjacent colours in the 12-bit table to be differentiated by the human eye (fairly easy) and which form a 'logical' ordered series (oops).

However you solve that last hiccup, some of the transitions from one level to the next end up being more obvious to the viewer than others. Playing with a normal image and the built-in colour maps with well-defined bands shows this nicely: some maps make certain features really jump out at you, while others will emphasise something completely different. When done right it can be an aid to visualisation, but unless you allow the user to fiddle in realtime, you need to know which density areas are most important before you start.

Struan

Subject: Re: Medical Imaging Question
Posted by [davidf](#) on Tue, 17 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Pavel Romashkin (promashkin@cmdl.noaa.gov) writes:

- > I have been out of the loop here for a week - had to move :-(but got closer to
- > Ft. Collins now :-)

Out of the sophisticated, effete city and back with the simple folks. I like it. :-)

- > Pardon me if it is a silly question but I wanted to ask this a week ago just as this
- > thread started. If the 4096 grayscale levels are needed for *visual* analyses then
- > is it not possible to map them to color scale - I'd imagine a human eye should be
- > able to recognize variety in color easier than in grayscale.
- > I think it may not work only since the method is so well established that the client
- > will want to use conventional, well-known shades of gray that mean a lot to
- > pathologists.

Not a bad idea, probably, but in my experience color is anathema to medical researchers. I take a great deal of abuse in my IDL classes when I display the CTSCAN image in color and upside down. :-(

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 [Pavel Romashkin](#) on Tue, 17 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi David,

I have been out of the loop here for a week - had to move :-((((but got closer to Ft. Collins now :-)

Pardon me if it is a silly question but I wanted to ask this a week ago just as this thread started. If the 4096 grayscale levels are needed for *visual* analyses then is it not possible to map them to color scale - I'd imagine a human eye should be able to recognize variety in color easier than in grayscale.

I think it may not work only since the method is so well established that the client will want to use conventional, well-known shades of gray that mean a lot to pathologists.

However, if digital analyses are in mind, then does not masking work well (just the same thing that ENVI has)?

Cheers,

Pavel

Subject: Re: Medical Imaging Question

Posted by [davidf](#) on Tue, 17 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Stein Vidar Hagfors Haugan (steinhh@ulrik.uio.no) writes:

> Who needs a histogram? To get the n'th percentile of a data set:

>

> `npercentile=data((sort(data))(n*n_elements(data)/100))`

>

> So, how much money did you charge for this, David?

Uh, lots. But he's already been on the phone asking for a refund after reading this article. :-)

Cheers,

David

P.S. Let's just say that those who can, do. Those who can't, write books. The trick when you write this up, of course, is to act like it's so obvious that any idiot could have come up with it. That's what makes you sound like an expert. :-)

--

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 [steinhh](#) on Tue, 17 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

Who needs a histogram? To get the n'th percentile of a data set:

```
npercentile=data((sort(data))(n*n_elements(data)/100))
```

Oh.. sorry about packing the thing into a one-liner....

So, how much money did you charge for this, David?

Regards,

Stein Vidar

Subject: Re: Medical Imaging Question

Posted by [edward.s.meinel](#) on Wed, 18 Aug 1999 07:00:00 GMT

[View Forum Message](#) <> [Reply to Message](#)

In article <MPG.1214e6875d54157c98987f@news.frii.com>, davidf@dfanning.com (David Fanning) wrote:

>

> 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?
>

There's lots of ways to do that.

1)Sliders to set the max and min

2)Sliders to set the contrast (slope) and brightness (y-intercept)

3)Sliders or text-entry widgets to set the upper and lower penetration points

4)Select min and max on a plot of the histogram

Personally, I like #2 the least, but others like it because it works sort of like your monitor and TV controls. What I do is a combination of 1, 3, and 4. I have a plot of the histogram and below it are three text widgets -- one displays the cursor location in the histogram window (I suppose I could also include one that displays the number of bin members), the second is for entering either the minimum value or lower penetration point (if it is between 0.00 and 1.00 assume penetration point, otherwise it's the minimum value), and the third is for entering either the maximum value or upper penetration point. I can also select the minimum and maximum levels using the histogram window -- left mouse button selects the minimum, right mouse button selects the maximum.

Ed Meinel

Sent via Deja.com <http://www.deja.com/>
Share what you know. Learn what you don't.

Subject: Re: Medical Imaging Question
Posted by [Peter Clinch](#) on Wed, 18 Aug 1999 07:00:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Fanning wrote:

> Not a bad idea, probably, but in my experience color
> is anathema to medical researchers. I take a great
> deal of abuse in my IDL classes when I display the
> CTSCAN image in color and upside down. :-(

Colour isn't anathema: look at PET nuclear medicine scans...

But with a 4096 level x ray or mri representation the problem with colour is along the lines of "is blue bigger than red, or vice versa" You can solve this one way using "temperature" type scales (think this is what is done with PET, will ask next time I'm over at Nuclear Medicine), but these tend to have fewer values than 12 bit: is pale red "more" than deeper red? If so, how does it relate to deep yellow? And so on... And will one person's scheme match another? probably not! When selecting a background for our Uni web pages a pale cyan was an option. I described it as "warm", many others rejected it as "too cold"...

With a staright grey level, you know exactly where you are: more intensity, more value.

Pete.

--

Peter Clinch University of Dundee
Tel 44 1382 660111 ext. 33637 Medical Physics, Ninewells Hospital
Fax 44 1382 640177 Dundee DD1 9SY Scotland UK
net p.j.clinch@dundee.ac.uk <http://www.dundee.ac.uk/~pjclinch/>
