Subject: Re: Multi-band sampling strategy
Posted by Maxwell Peck on Wed, 18 Aug 2010 06:35:28 GMT
View Forum Message <> Reply to Message

On Aug 18, 1:10 pm, JessW <jess.wal...@gmail.com> wrote:
> Hi all,
>
> I'm using an IDL script in ENVI to retrieve the values of a multiband
> image at a user-specified number of pixels. The script works as
> intended, but is quite slow; sampling an 8-band image at ~ 38,000
> pixel locations takes
> 52.35 minutes. The issue is the following loop:
>
> FOR i= 0L, nSamplePix - 1 DO BEGIN
> pixValues[i,*] = ENVI_GET_SLICE(fid=fid,line=(index[1,i]-1),$
> pos=bandsSelected,xs=(index[0,i]-1), xe=(index[0,i]-1))
> ENDFOR
>
> If anyone can suggest an alternate, faster sampling strategy--
> preferably one that doesn't include the purchase of a new, faster
> computer--I would be extremely grateful.
>
> Thanks,
>
> Jess

I'm not clear on what you're doing but have you tried envi_get_data
with the dims keyword, or if that is not appropriate maybe creating a
point or polygon ROI (ENVI_CREATE_ROI/ENVI_GET_ROI_DATA)?

Max

Subject: Re: Multi-band sampling strategy
Posted by JessW on Wed, 18 Aug 2010 17:26:07 GMT
View Forum Message <> Reply to Message

Hi Max,

Thanks for responding, and sorry about the fuzzy description of what
I'm doing. Let me try to clarify.

I am tracking the reflectance change of particular pixels across a
time-series of Landsat data. So for my current stack of 8 Landsat
scenes spanning April - October, the program returns 8 reflectance
values for a given pixel address x,y. ENVI_GET_SLICE seemed ideal for
this task because it can sample a single pixel location through a
number of bands; however, as I mentioned, it is very slow when it

loops over a large number of pixels to sample.

ENVI_GET_DATA appears to work on only one band, so it doesn't seem
quite as suitable for what I'm trying to do. I'll have to investigate
the use of ENVI_CREATE_ROI; thanks for that suggestion.

Jess

Subject: Re: Multi-band sampling strategy
Posted by penteado on Wed, 18 Aug 2010 18:09:16 GMT
View Forum Message <> Reply to Message

On Aug 18, 2:26 pm, JessW <jess.wal...@gmail.com> wrote:
> Hi Max,
>
> Thanks for responding, and sorry about the fuzzy description of what
> I'm doing. Let me try to clarify.
>
> I am tracking the reflectance change of particular pixels across a
> time-series of Landsat data. So for my current stack of 8 Landsat
> scenes spanning April - October, the program returns 8 reflectance
> values for a given pixel address x,y. ENVI_GET_SLICE seemed ideal for
> this task because it can sample a single pixel location through a
> number of bands; however, as I mentioned, it is very slow when it
> loops over a large number of pixels to sample.
>
> ENVI_GET_DATA appears to work on only one band, so it doesn't seem
> quite as suitable for what I'm trying to do. I'll have to investigate
> the use of ENVI_CREATE_ROI; thanks for that suggestion.
>
> Jess

If your image is not too large to fit in memory, it seems that it
might be better to just get the whole image as a 3D array, once, then
take out the subset you want directly, just with indexing. If the
image is too large, but the section you want fits into a small
rectangle, you could retrieve just that region, defined by the max and
min of your indexes, as a 3D array, then pick the points inside it.

Subject: Re: Multi-band sampling strategy
Posted by wita on Wed, 18 Aug 2010 19:46:17 GMT
View Forum Message <> Reply to Message

Dear Jess,

The procedure that you are now using is fetching a complete slice of your image
for each sample location. I do not know how large your image is, but if I take
a typical TM scene (6000x6000 pixels), then many of your 38,000 pixels will
actually be positioned in the same slice. So you are probably fetching the
same slice a couple of times from your image just to take out another sample
location. The more sample locations you have, the worse the speed penalty gets.

Instead, what you could do to speed this up is first sort out which sample
locations fall on which slice and group them together. There is actually a
very efficient way to accomplish this in IDL by using the Histogram() command
in combination with the REVERSE_INDICES. If you have never heard of this then
please have a look at the histogram tutorial on David Fannings site to
understand the trick.

So in pseudo code it could look like this, where I assume that SampleLines,
SampleCols store the line/column numbers of your sample positions:

```
h = Histogram(SampleLines, MIN=0, REVERSE_INDICES=ri)
FOR j=0, N_Elements(h)-1 DO BEGIN
  IF ri[j+1] GT ri[j] THEN BEGIN
     SampleColsinLine = samplecols[ri[ri[j]:ri[j+1]-1]]
     SampledSpectra = Get_Sample_Spectra(fid, j, SampleColsinLine)
  ENDIF
ENDFOR
```

with the function Get_Sample_Spectra() looking like this:

```
FUNCTION get_sample_spectra, fid, line, samplecols
   Slice = ENVI_Get_Slice(FID=fid,line=line, pos=bandsSelected)
   RETURN, Slice[samplecols, *]
END
```

This has the additional advantage that you are progressing through the file
in the order in which it is stored on disk. This automatically takes advantage
of the read-ahead capabilities of operating systems and disk drives.

Creating point ROIs as Maxwell suggest is, in my experience, also becoming
very slow for a large number of samples. The suggestion by Paulo is inherently
limited by the size of your memory. Moreover, it is very inefficient for small
sample sizes.

Hope this helps.

Allard

Subject: Re: Multi-band sampling strategy
Posted by JessW on Wed, 18 Aug 2010 20:58:02 GMT
View Forum Message <> Reply to Message

Hi Allard,

Interesting. You are correct; I'm using random pixel locations across an entire Landsat scene, so many of the pixels may fall in the same slice. I had assumed that since I specify a unique line/column address in the ENVI_GET_SLICE command, IDL retrieves only the values for that one pixel. But obviously if the command retrieves the entire slice prior to defining the column selection, it makes for an inefficient process, as you point out.

I've tangled with histograms and reverse indices before; the experience left me with a splitting headache, so it may take me some time to successfully apply your suggestion to my program. I'll post a follow-up message when I have results.

Thanks very much for taking the time to reply!

Jess

Subject: Re: Multi-band sampling strategy
Posted by Maxwell Peck on Wed, 18 Aug 2010 21:11:49 GMT
View Forum Message <> Reply to Message

On Aug 19, 6:58 am, JessW <jess.wal...@gmail.com> wrote:
> Hi Allard,
>
> Interesting. You are correct; I'm using random pixel locations across
> an entire Landsat scene, so many of the pixels may fall in the same

> slice. I had assumed that since I specify a unique line/column
> address in the ENVI_GET_SLICE command, IDL retrieves only the values
> for that one pixel. But obviously if the command retrieves the entire
> slice prior to defining the column selection, it makes for an
> inefficient process, as you point out.
>
> I've tangled with histograms and reverse indices before; the
> experience left me with a splitting headache, so it may take me some
> time to successfully apply your suggestion to my program. I'll post a
> follow-up message when I have results.
>
> Thanks very much for taking the time to reply!
>
> Jess

Jess,

Allard's approach is certainly better. I was unclear initially if you were subsetting 'rectangles' or random point locations. As Allard suggests point ROI's can get slow if you have a lot.. It all really depends on how many points we're talking.

Max

Subject: Re: Multi-band sampling strategy
Posted by JessW on Wed, 25 Aug 2010 18:13:07 GMT
View Forum Message <> Reply to Message

The belated follow-up...

Allard's 'pseudo-code' code worked perfectly.

Extracting ~38,000 points with ENVI_GET_SLICE using
- previous extraction method: 52.35 minutes
- histograms: 3.5 minutes

Wow! I had no idea histograms could deliver that kind of efficient performance. It looks like I'll just have to keep the aspirin handy from now on.

Thanks again

Jess

Subject: Re: Multi-band sampling strategy

Posted by wita on Tue, 31 Aug 2010 09:08:41 GMT
View Forum Message <> Reply to Message

Hi Jess,

good to hear that it works so well.

regards,

Allard