
Subject: Re: Find minimums in an array...

Posted by [David Fanning](#) on Wed, 16 Oct 2002 18:50:37 GMT

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

Filipe (engpcris@livjm.ac.uk) writes:

- > I'm looking for a fonction or an algorithm who could find the X
- > minumun values in a array without having to use Sort (for a time
- > reason) or Min (updating the array everytime)
- >
- > For Example in a Array of 1000 elements I want to find the 10 smallest
- > values.
- >
- > Do you know any quick how to perfome it? (as I need to do this
- > operation many times in a loop)

Oh, my gosh! I believe HISTOGRAM can be used to solve this problem!!!

http://www.dfanning.com/tips/histogram_tutorial.html

Cheers,

David

P.S. Let's just say that *this* is what enlightenment feels like! :-)

--

David W. Fanning, Ph.D.

Fanning Software Consulting, Inc.

Phone: 970-221-0438, E-mail: david@dfanning.com

Coyote's Guide to IDL Programming: <http://www.dfanning.com/>

Toll-Free IDL Book Orders: 1-888-461-0155

Subject: Re: Find minimums in an array...

Posted by [James Kuyper](#) on Wed, 16 Oct 2002 19:30:55 GMT

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

David Fanning wrote:

- >
- > Filipe (engpcris@livjm.ac.uk) writes:
- >
- >> I'm looking for a fonction or an algorithm who could find the X
- >> minumun values in a array without having to use Sort (for a time
- >> reason) or Min (updating the array everytime)
- >>

>> For Example in a Array of 1000 elements I want to find the 10 smallest
>> values.
>>
>> Do you know any quick how to performe it? (as I need to do this
>> operation many times in a loop)
>
> Oh, my gosh! I believe HISTOGRAM can be used to solve
> this problem!!!
>
> http://www.dfanning.com/tips/histogram_tutorial.html

That should be:
http://

> P.S. Let's just say that *this* is what enlightenment
> feels like! :-)

I feel somewhat unenlightened. Could you be a little more specific?

Subject: Re: Find minimums in an array...
Posted by [David Fanning](#) on Wed, 16 Oct 2002 20:01:17 GMT
[View Forum Message](#) <> [Reply to Message](#)

James Kuyper (kuyper@saicmodis.com) writes:

>> P.S. Let's just say that *this* is what enlightenment
>> feels like! :-)
>
> I feel somewhat unenlightened. Could you be a little more specific?

Well, I just thought "ten lowest values", what a perfect
job for the histogram function. Divide your array into
values, then Histogram it, and starting looking in the
lowest bins. Keep going until you find 10 values! Find
the locations by using REVERSE_INDICES.... Oh, oh. Hang
on a second, I've got to go check with the guru...

http://www.dfanning.com/tips/histogram_tutorial.html

Cheers,

David

P.S. Let's just say when you have finally learned how
to use a hammer, everything looks like a nail. :-)

--

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

Subject: Re: Find minimums in an array...
Posted by [James Kuyper](#) on Wed, 16 Oct 2002 21:00:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Fanning wrote:

```
>  
> James Kuyper (kuyper@saicmodis.com) writes:  
>  
>>> P.S. Let's just say that *this* is what enlightenment  
>>> feels like! :-)  
>>  
>> I feel somewhat unenlightened. Could you be a little more specific?  
>  
> Well, I just thought "ten lowest values", what a perfect  
> job for the histogram function. Divide your array into  
> values, then Histogram it, and starting looking in the  
> lowest bins. Keep going until you find 10 values! Find  
> the locations by using REVERSE_INDICES.... Oh, oh. Hang  
> on a second, I've got to go check with the guru...  
>  
> http://www.dfanning.com/tips/histogram\_tutorial.html
```

I don't believe that the indices, within each bin, refer to the array entries in sort order. So you'll have to at least sort the elements of the last bin, to get precisely the 10 smallest items. Unless including that bin brings the count to precisely 10.

The efficiency of such an algorithm will depend upon the bin size. The smaller the bin size, the fewer items you'll have to sort in the final bin, but there will be a corresponding increase in the number of bins you'll have to look at. The proper choice will depend a lot upon the distribution of the data. For very clumpy data, this approach won't be significantly faster than the sort-based approach.

Subject: Re: Find minimums in an array...
Posted by [David Fanning](#) on Wed, 16 Oct 2002 21:55:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

James Kuyper (kuyper@saicmodis.com) writes:

> For very clumpy data, this approach won't be
> significantly faster than the sort-based approach.

If my reading of the tome is correct, for very clumpy data a *second* histogram is in order. :-)

But I was really fooling around more than I was offering a serious solution. I doubt anything that finds the 10 least minimums will be very fast in a loop. :-(

Cheers,

David

--

David W. Fanning, Ph.D.

Fanning Software Consulting, Inc.

Phone: 970-221-0438, E-mail: david@dfanning.com

Coyote's Guide to IDL Programming: <http://www.dfanning.com/>

Toll-Free IDL Book Orders: 1-888-461-0155

Subject: Re: Find minimums in an array...

Posted by [JD Smith](#) on Fri, 18 Oct 2002 16:04:04 GMT

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

On Wed, 16 Oct 2002 14:00:35 -0700, James Kuyper wrote:

> David Fanning wrote:

>>

>> James Kuyper (kuyper@saicmodis.com) writes:

>>

>>>> P.S. Let's just say that *this* is what enlightenment feels like!

>>>> :-)

>>>

>>> I feel somewhat unenlightened. Could you be a little more specific?

>>

>> Well, I just thought "ten lowest values", what a perfect job for the

>> histogram function. Divide your array into values, then Histogram it,

>> and starting looking in the lowest bins. Keep going until you find 10

>> values! Find the locations by using REVERSE_INDICES.... Oh, oh. Hang on

>> a second, I've got to go check with the guru...

>>

>> http://www.dfanning.com/tips/histogram_tutorial.html

>

> I don't believe that the indices, within each bin, refer to the array

> entries in sort order. So you'll have to at least sort the elements of

- > the last bin, to get precisely the 10 smallest items. Unless including
- > that bin brings the count to precisely 10.
- >
- > The efficiency of such an algorithm will depend upon the bin size. The
- > smaller the bin size, the fewer items you'll have to sort in the final
- > bin, but there will be a corresponding increase in the number of bins
- > you'll have to look at. The proper choice will depend a lot upon the
- > distribution of the data. For very clumpy data, this approach won't be
- > significantly faster than the sort-based approach.

Au contraire. SORT comes to mind, but is it the fastest method? It's certainly straightforward looking:

```
a=randomu(sd,1000)
```

```
s=sort(a)
vals=a[s[0:9]]
```

It's done some unnecessary work putting all those elements in order (true *selection* on an array is of order N , sorting is $N \log(N)$, but IDL doesn't have a true selection routine). Can we improve the speed with HISTOGRAM? The answer to that question is usually "yes", but it comes at the expense of readability:

```
h=histogram(a,NBINS=1000/10,REVERSE_INDICES=ri)
s=0L
for j=0L,n_elements(h)-1L do begin
    s=s+h[j]
    if s ge 10 then break
endfor
vals2=a[ri[ri[0]:ri[j+1]-1]]
vals2=(vals2[sort(vals2))][0:9]
```

I've just taken the first few bins of the histogram, such that we get at least 10 elements, and done a much smaller SORT on the data in those bins to find the smallest 10. Notice how I arranged for there to be 100 bins for this 1000-point vector, so that, for randomly distributed data anyway, you'd have 10 elements per bin on average. The run-time is not terribly sensitive to this number, as long as you divide the data up into some decent number of bins (more than 1 ;). E.g. I could just pick NBINS=100 always. Bottom line -- SORT is slow, HISTOGRAM is fast.

Also notice that I used a FOR loop to find the first bin where you've hit at least 10 elements cumulative. I expect this to be near the front, so this is much faster than a loop-free solution like:
`j=(where(total(h,/CUMULATIVE) ge 10))[0].`

Here are the timings (in sec):

1000-point Uniform Random Data:

```
a=randomu(sd,1000)
```

SORT METHOD: 0.000612

HISTOGRAM+SORT METHOD: 0.000151

A nice 4x speedup.

1,000,000-point, Poisson Distributed Data:

```
a=randomn(sd,1000000,POISSON=100)
```

SORT METHOD: 1.7388

HISTOGRAM+SORT METHOD: 0.1453

A 12x speedup.

100,000-point, Bizarrely Clumpy Data:

```
a=[4+randomu(sd,2500),10+randomu(sd,2500), $  
  150+randomu(sd,2500),200+randomu(sd,2500)]
```

SORT METHOD: 0.0966

HISTOGRAM+SORT METHOD: 0.0377

A 2.6x speedup.

As always, you're mileage may vary. Check the performance on your equipment, using your data.

Also, lest everyone think I'm out to turn their code into piles of HISTOGRAM spaghetti: just because a fast HISTOGRAM/REBIN/REFORM solution exists, doesn't mean you should always use it (though sometimes, HISTOGRAM is the simplest approach from the beginning). The SORT method above has much clearer intent. If you're dealing with a small data set, don't loop over it many times, or aren't in a time-critical section of the code, readability and maintainability may be more compelling than speed. I typically code for readability (with some care taken to avoid trivial FOR loops), and then make another pass optimizing critical sections of the code with these techniques (carefully commenting of course ;).

Good luck,

JD
