

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Martin Downing](#) on Sun, 25 Feb 2001 18:44:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

From thread: <http://cow.physics.wisc.edu/~craigm/idl/archive/msg03957.htm> I  
a.. Date: Wed, 29 Nov 2000 16:30:54 -0500

Richard Tyc wrote:

>  
> I need to apply a smoothing type kernel across an image, and calculate the  
> standard deviation of the pixels masked by this kernel.  
>  
> ie. lets say I have a 128x128 image. I apply a 3x3 kernel (or simply a  
> mask) starting at [0:2,0:2] and use these pixels to find the standard  
> deviation for the center pixel [1,1] based on its surrounding pixels, then  
> advance the kernel etc deriving a std deviation image essentially.  
> I can see myself doing this 'C' like with for loops but does something  
exist  
> for IDL to do it better or more efficiently ?  
>  
> Rich

I was wandering through new Craig's IDL archive site (which is brilliant by  
the way) and came across this question asking for an efficient way of  
calculating the local standard deviation in an array. It seemed to me that  
the thread had not reached a full solution so perhaps some of you might be  
interested in this method which is very fast. It is based on the crafty  
formula for variance:

variance = (sum of the squares)/n + (square of the sums)/n\*n

[ apologies if this is going over old ground !]

```
function IMAGE_VARIANCE , image, halfWidth, MEAN=av_im, $  
    NEIGHBOURHOOD=NEIGHBOURHOOD,$  
    POPULATION_ESTIMATE=POPULATION_ESTIMATE  
;  
;+  
; NAME:  
; IMAGE_VARIANCE  
;  
; PURPOSE:  
; This function calculates the local-neighbourhood statistical variance.  
; I.e. for each array element a the variance  
; of the neighbourhood of +- halfwidth is calculated.  
; The routine avoids any loops and so is fast and "should" work for any  
dimension of array  
;  
;  
; CATEGORY:  
; Image Processing  
;  
;
```

```

; CALLING SEQUENCE:
;
; Result = IMAGE_VARIANCE(Image, HalfWidth)
;
; INPUTS:
; Image : the array of which we calculate the variance. Can be any dimension
; HalfWidth: the half width of the NEIGHBOURHOOD, indicates we are
looking at a
;     neighborhood +/- N from the pixel in each dimension
;
; OPTIONAL INPUTS:
; Parm2: Describe optional inputs here. If you don't have any, just
; delete this section.
;
; KEYWORD PARAMETERS:
;
; NEIGHBOURHOOD: calculate for the NEIGHBOURHOOD only not the central pixel.
;
; POPULATION_ESTIMATE: return the population estimate of variance, not the
sample variance
;
; OUTPUT:
; returns an array of same dimensions as input array in which each pixel
represents the local variance centred at that position
;
; OPTIONAL OUTPUTS:
; MEAN_IM: set to array of local area mean, same dimensionality as input.
;
; RESTRICTIONS:
; Edges are dealt with by replicating border pixels this is likely to
give an underestimate of variance in these regions
;
; PROCEDURE:
; Based on the formula for variance:
;  $var = (\text{sum of the squares})/n + (\text{square of the sums})/n*n$ 
;
; EXAMPLE:
; Example of simple statistical-based filter for removing spike-noise
;
;     var_im = image_variance(image, 5, mean=mean_im, /neigh)
;     zim = (image-mim)/sqrt(var_im)
;     ids = where(zim gt 3, count)
;     if count gt 0 then image[ids] = mean_im[ids]
;
; MODIFICATION HISTORY:
; Written by: Martin Downing, 30th September 2000
; m.downing@abdn.ac.uk
;-

```

```

; full mask size as accepted by SMOOTH()
n = halfWidth*2+1

; this keyword to SMOOTH() is always set
EDGE_TRUNCATE= 1
; sample size
m = n^2
; temporary double image copy to prevent overflow
im = double(image)
; calc average
av_im = smooth(im, n, EDGE_TRUNCATE=EDGE_TRUNCATE)
; calc squares image
sq_im = temporary(im)^2
; average squares
asq_im = smooth(sq_im, n, EDGE_TRUNCATE=EDGE_TRUNCATE)

if keyword_set(NEIGHBOURHOOD) then begin
; remove centre pixel from estimate
; calc neighbourhood average (removing centre pixel)
av_im = (av_im*m - image)/(m-1)
; calc neighbourhood average of squares (removing centre pixel)
asq_im = (asq_im*m - temporary(sq_im))/(m-1)
; adjust sample size
m = m-1
endif

var_im = temporary(asq_im) - (av_im^2)
if keyword_set(POPULATION_ESTIMATE) then begin
var_im = var_im *( double(m)/(m-1))
endif

return, var_im

end

```

-----  
Martin Downing,  
Clinical Research Physicist,  
Orthopaedic RSA Research Centre,  
Woodend Hospital,  
Aberdeen, AB15 6LS.  
m.downing@abdn.ac.uk

Richard Tyc wrote:

```

>
> WOW, I need to look at these equations over about a dozen times to see
what
> is going on ?
>
> I have been struggling with the variance of an nxn window of data,
INCLUDING
> central pixel
>
> ;mean of the neighboring pixels (including central)
> mean=smooth(arr,n)
> ;square deviation from that mean
> sqdev=(arr-mean)^2
> ;variance of an nxn window of data, INCLUDING central pixel
> var=(smooth(sqdev,n)*n^2-sqdev)/(n^2-1)
>

```

Almost right. Try:

```
var=smooth(sqdev,n)*n^2/(n^2-1)
```

but this still won't yield exactly what you're after, but maybe you're after the wrong thing ;)

What this computes is a smoothed box variance, not a true box variance, since the mean you are using changes over the box (instead of subtracting the mean value at the central pixel from each in the box, we subtract the box mean value at *that* pixel). Usually, this type of variance is a more robust estimator, e.g. for excluding outlier pixels, etc. (in which case you probably should exclude the central pixel after all to avoid the chicken and egg problem with small box sizes). If you really want the true variance, you're probably stuck with for loops, preferably done in C and linked to IDL.

This reminds me of a few things I've been thinking about IDL recently. Why shouldn't *all* of these smooth type operations be trivially feasible in IDL. Certainly, the underlying code required is simple. Why can't we just say:

```
a=smooth(b,n,/VARIANCE)
```

to get a true box variance, or

```
a=smooth(b,n,/MAX)
```

to get the box max. Possibilities:

\*MEAN (the current default)

\*TOTAL (a trivial scaling of mean),

\*VARIANCE

\*MEDIAN (currently performed by the median function, in a addition to its normal duties. To see why this is strange, consider that total() doesn't have an optional "width" to perform neighborhood filtering).

\*MIN

\*MAX

\*MODE

\*SKEW

etc.

To be consistent, these should all operate natively on the input data type (float, byte, long, etc. -- like smooth() and convol() do, but like median() does not!), and should apply consistent edge conditions activated by keywords. These seem like simple enough additions, and would require much reduced chicanery.

While I'm on the gripe train, why shouldn't we be able to consistently perform operations along any dimension of an array we like with relevant IDL routines. E.g., we can total along a single dimension. All due respect to Craig's CMAPPLY function, but some of these things should be much faster. Resorting to summed logarithms for multiplication is not entirely dubious, but why shouldn't we be able to say:

```
col_max=max(array,2,POS=mp)
```

and have mp be a list of max positions, indexed into the array, and rapidly computed? While we're at it, why not

```
col_med=median(array,2,POS=mp)
```

IDL is an array based language, but it conveniently forgets this fact on occassion. Certainly there are compatibility difficulties to overcome to better earn this title, but that shouldn't impede progress.

JD

--

J.D. Smith | WORK: (607) 255-6263  
Cornell Dept. of Astronomy | (607) 255-5842  
304 Space Sciences Bldg. | FAX: (607) 255-5875  
Ithaca, NY 14853

-----  
Martin Downing,  
Clinical Research Physicist,  
Orthopaedic RSA Research Centre,  
Woodend Hospital,  
Aberdeen, AB15 6LS.

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Craig Markwardt](#) on Sun, 25 Feb 2001 20:39:02 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

"Martin Downing" <martin.downing@ntlworld.com> writes:

>  
> I was wandering through new Craig's IDL archive site (which is brilliant by  
> the way) and came across this question asking for an efficient way of  
> calculating the local standard deviation in an array. It seemed to me that  
> the thread had not reached a full solution so perhaps some of you might be  
> interested in this method which is very fast. It is based on the crafty  
> formula for variance:  
>  $\text{variance} = (\text{sum of the squares})/n + (\text{square of the sums})/n*n$   
>  
> [ apologies if this is going over old ground !]

Hi Martin--

This looks like a great way to do things. Now if we could only  
translate those bloomin' British English spellings! NEIGHBOURHOOD  
indeed. :-)

As a side note, I believe that your message would have actually shown  
up as a reply to Richard's or JD's messages in the archive, if you had  
placed the Message-ID of their article in the References header of  
yours. I've tried it here. We'll see if it works.

Craig

--

-----  
Craig B. Markwardt, Ph.D.      EMAIL: [craigmnet@cow.physics.wisc.edu](mailto:craigmnet@cow.physics.wisc.edu)  
Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response  
-----

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [John-David T. Smith](#) on Mon, 26 Feb 2001 03:00:04 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Martin Downing wrote:

>  
> From thread: <http://cow.physics.wisc.edu/~craigm/idl/archive/msg03957.htm> I

> a.. Date: Wed, 29 Nov 2000 16:30:54 -0500  
 > Richard Tyc wrote:  
 >>  
 >> I need to apply a smoothing type kernel across an image, and calculate the  
 >> standard deviation of the pixels masked by this kernel.  
 >>  
 >> ie. lets say I have a 128x128 image. I apply a 3x3 kernel (or simply a  
 >> mask) starting at [0:2,0:2] and use these pixels to find the standard  
 >> deviation for the center pixel [1,1] based on its surrounding pixels, then  
 >> advance the kernel etc deriving a std deviation image essentially.  
 >> I can see myself doing this 'C' like with for loops but does something  
 > exist  
 >> for IDL to do it better or more efficiently ?  
 >>  
 >> Rich  
 >  
 > I was wandering through new Craig's IDL archive site (which is brilliant by  
 > the way) and came across this question asking for an efficient way of  
 > calculating the loacal standard deviation in an array. It seemed to me that  
 > the thread had not reached a full solution so perhaps some of you might be  
 > interested in this method which is very fast. It is based on the crafty  
 > formula for variance:  
 >  $\text{variance} = (\text{sum of the squares})/n + (\text{square of the sums})/n*n$

Righto. I knew I was fishing for something like this. Except I think you mean:

$(\text{population}) \text{ variance} = (\text{sum of the squares})/n - (\text{square of the sums})/n*n$

Luckily, that's how you've coded it too. Sample variance (=population  
 variance\*n/(n-1)) is of course the more common case in science (as opposed to  
 gambling).

JD

P.S. I think I originally got the idea from sigma\_filter.pro, a NASA library  
 routine, dating back to 1991. It's chock-full of other good tidbits too.  
 Thanks Frank and Wayne!

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
 Posted by [Martin Downing](#) on Mon, 26 Feb 2001 09:52:18 GMT  
[View Forum Message](#) <> [Reply to Message](#)

>> interested in this method which is very fast. It is based on the crafty  
 >> formula for variance:  
 >>  $\text{variance} = (\text{sum of the squares})/n + (\text{square of the sums})/n*n$   
 >  
 > Righto. I knew I was fishing for something like this. Except I think you

mean:

>  
> (population) variance = (sum of the squares)/n - (square of the sums)/n\*n  
>  
> Luckily, that's how you've coded it too. Sample variance (=population  
> variance\*n/(n-1)) is of course the more common case in science (as opposed  
to  
> gambling).

Sigh - I hear what you are saying, but this was a misunderstanding. I  
\*tried\* to make its use unambiguous by making the default option the  
absolute variance of the array (n as the denominator) , or when  
POPULATION\_ESTIMATE is set then calculate an \*estimate\* of the population  
from which this dataset is assumed to be a SAMPLE [giving (n-1) as the  
denominator]. Judging by your reply I failed dismally!

You are right - POPULATION\_ESTIMATE is normally termed "sample stdev" and is  
the equivalent of IDL's variance(x) - but what they mean is that it is an  
estimator of the popn stdev! Still waiting to try it in the casinos :)

Martin

---

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Martin Downing](#) on Mon, 26 Feb 2001 10:00:10 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> Hi Martin--  
>  
> This looks like a great way to do things. Now if we could only  
> translate those bloomin' British English spellings! NEIGHBOURHOOD  
> indeed. :-)

Just getting my own back for having to \*relearn\* how to type COLOR, I tend  
to use "COL" when setting a graphics keyword to preserve my heritage, but  
have to relent when defining a new functions. You should just enter /NEIGHBO  
:)

>  
> As a side note, I believe that your message would have actually shown  
> up as a reply to Richard's or JD's messages in the archive, if you had  
> placed the Message-ID of their article in the References header of  
> yours. I've tried it here. We'll see if it works.

OK, my news reader has the header field "Followup-To:" (outlook express -  
cringe) is this equivalent to Message-ID?

Martin

---

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Martin Downing](#) on Mon, 26 Feb 2001 10:21:41 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> Hi Martin--  
>  
> This looks like a great way to do things. Now if we could only  
> translate those bloomin' British English spellings! NEIGHBOURHOOD  
> indeed. :-)

Just getting my own back for having to \*relearn\* how to type COLOR, I tend to use "COL" when setting a graphics keyword to preserve my heritage, but have to relent when defining a new functions! Accordingly you should just enter /NEIGHBO  
:)

>  
> As a side note, I believe that your message would have actually shown  
> up as a reply to Richard's or JD's messages in the archive, if you had  
> placed the Message-ID of their article in the References header of  
> yours. I've tried it here. We'll see if it works.

OK, the closest sounding header with my news reader (outlook express - cringe) is "Followup-To:" - is this equivalent to Message-ID?

so should I have pasted the subject of this line  
References: <3986B6E9.48105613@met.ed.ac.uk> <8m6ims\$jr8\$1@news.lth.se>

Martin

---

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Martin Downing](#) on Mon, 26 Feb 2001 10:21:54 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

>> interested in this method which is very fast. It is based on the crafty  
>> formula for variance:  
>>  $\text{variance} = (\text{sum of the squares})/n + (\text{square of the sums})/n*n$   
>  
> Righto. I knew I was fishing for something like this. Except I think you mean:  
>  
>  $(\text{population}) \text{ variance} = (\text{sum of the squares})/n - (\text{square of the sums})/n*n$   
>  
> Luckily, that's how you've coded it too. Sample variance (=population  
>  $\text{variance}*n/(n-1)$ ) is of course the more common case in science (as opposed  
to  
> gambling).

>> ; POPULATION\_ESTIMATE: return the population estimate of variance, not the  
>> sample variance

Sigh - you are right of course, but this was just a misunderstanding.  
Setting the Keyword above returns an \*estimate\* of the population from which  
the input sample is assumed to have come from, is more commonly referred to  
as SAMPLE variance. The default is the variance of the array, which is  
usually coined population variance. I agree that until I take up gambling  
the former should be the default!! Although I know what I'm doing,  
personally I find these single name terms very ambiguous, as you can see!!!  
Shame I never read my comments.....

Martin

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Martin Downing](#) on Mon, 26 Feb 2001 15:32:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

"John-David Smith" <jdsmith@astro.cornell.edu> wrote in message  
news:3A99C6B4.10549265@astro.cornell.edu...

>  
> P.S. I think I originally got the idea from sigma\_filter.pro, a NASA  
library  
> routine, dating back to 1991. It's chock-full of other good tidbits too.  
> Thanks Frank and Wayne!

Hi John,  
Just checked the file SIGMA\_FILTER.pro at  
<http://idlastro.gsfc.nasa.gov/ftp/pro/image/?N=D>  
I really must spend more time browsing these great sites.  
The code is similar, however it does not calculate the true variance under  
the mask  
they calculate for a box width of n, (ignoring centre pixel removal):

```
-----  
mean_im=(smooth(image, n) )  
dev_im = (image - mean_im)^2  
var_im = smooth(dev_im, n)/(n-1)  
-----
```

This is not the true variance of the pixels under the box mask, as each  
pixel in the mask is having a different mean subtracted.

i.e (read this as a formula if you can!)

$$\text{Pseudo\_Variance} = \text{SUM}_{ij} ( ( I(x+i,y+j) - \text{MEAN}(x+i,y+j) )^2 ) / (n-1)$$

instead of true variance:

$$\text{Variance} = \text{SUM}_{ij} ( ( I(x+i,y+j) - \text{MEAN}_{xy} )^2 ) / (n-1)$$

which can be reduced to :  $\{(\text{SUM } ij ( ( I(x+i,y+j))^2 ) - (\text{SUM } ij I(x+i,y+j) ) ^2)/n \}/(n-1)$

hence the non loop method we use below:

```
-----
; calc box mean
mean_im = smooth(image, n)
; calc box mean of squares
msq_im = smooth(image^2, n)
; hence variance
var_im = ( msq_im - mean_im^2) * (n/(n-1.0))
-----
```

cheers

Martin

PS: Sorry about my before-and-after-coffee postings this morning, outlook decided to post my replies whilst I was still pondering - how kind - I've killed that \*feature\* now :)

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE  
Posted by [Craig Markwardt](#) on Mon, 26 Feb 2001 18:51:17 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

"Martin Downing" <[martin.downing@ntlworld.com](mailto:martin.downing@ntlworld.com)> writes:

>  
> OK, my news reader has the header field "Followup-To:" (outlook express -  
> cringe) is this equivalent to Message-ID?  
>

After investigating it a little more, I would say, don't worry about it. I manipulated the articles to go into the correct thread this time, but it may be too much work. If you can add a "References:" header entry then fine, go ahead.

Craig

--

```
-----
Craig B. Markwardt, Ph.D.      EMAIL:  craigmnet@cow.physics.wisc.edu
Astrophysics, IDL, Finance, Derivatives | Remove "net" for better response
-----
```

---

Subject: Re: efficient kernel or masking algorithm ? UPDATE

On Monday, February 26, 2001 10:38:06 AM UTC-5, Martin Downing wrote:

> "John-David Smith" <jdsmith@astro.cornell.edu> wrote in message  
> news:3A99C6B4.10549265@astro.cornell.edu...  
>>  
>> P.S. I think I originally got the idea from sigma\_filter.pro, a NASA  
> library  
>> routine, dating back to 1991. It's chock-full of other good tidbits too.  
>> Thanks Frank and Wayne!  
>  
> Hi John,  
> Just checked the file SIGMA\_FILTER.pro at  
> <http://idlastro.gsfc.nasa.gov/ftp/pro/image/?N=D>  
> I really must spend more time browsing these great sites.  
> The code is similar, however it does not calculate the true variance under  
> the mask  
> they calculate for a box width of n, (ignoring centre pixel removal):  
> -----  
> mean\_im=(smooth(image, n) )  
> dev\_im = (image - mean\_im)^2  
> var\_im = smooth(dev\_im, n)/(n-1)  
> -----  
> This is not the true variance of the pixels under the box mask, as each  
> pixel in the mask is having a different mean subtracted.  
> i.e (read this as a formula if you can!)  
> Pseudo\_Variance = SUM ij ( ( I(x+i,y+j) - MEAN(x+i,y+j))^2 )/(n-1)  
>  
> instead of true variance:  
> Variance = SUM ij ( ( I(x+i,y+j) - MEANxy)^2 )/(n-1)  
> which can be reduced to : {(SUM ij ( ( I(x+i,y+j)^2 ) - (SUM ij  
> I(x+i,y+j) ) ^2)/n }/(n-1)  
> hence the non loop method we use below:  
> -----  
> ; calc box mean  
> mean\_im = smooth(image, n)  
> ; calc box mean of squares  
> msq\_im = smooth(image^2, n)  
> ; hence variance  
> var\_im = ( msq\_im - mean\_im^2 ) \* (n/(n-1.0))  
> -----  
>  
> cheers  
>  
> Martin  
>  
> PS: Sorry about my before-and-after-coffee postings this morning, outlook  
> decided to post my replies whilst I was still pondering - how kind - I've

> killed that \*feature\* now :)

n seems to mean two things in your code: in the smooth function it is the window width and in your final variance calculation line it means number of samples. These should not be the same. If n is window size then the final line should read:

; hence variance

```
var_im = ( msq_im - mean_im^2 ) * (n*n/((n*n)-1.0))
```

Right? Or did I misunderstand something?

---