
Subject: bug in IDL's hanning() window-generating function

Posted by [bennetsc](#) on Wed, 01 Aug 2001 08:16:38 GMT

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

Perhaps this has been discussed or reported before, but if so, RSI has not yet fixed it.

I just noticed that IDL's hanning() has a bug described by Harris (1978) as common in software packages. The bug is easy to look for. To check, generate a window of even length. Then look at its coefficients. The array should *not* be exactly symmetrical. It should have only one highest coefficient next to, and to the right of, the "center" of the array, noting that the center of the time (or space) period covered by the discrete window actually falls at the largest window coefficient, which is immediately to the right of the "center" of the array, because of the even number of coefficients. IDL's hanning() produces a symmetric window with two identical coefficients at the center and zeros at the ends, which is incorrect. A Hann window of even length should only be zero at the left (i.e. starting) end, whereas the other end should be equal to the first coefficient in from the zero. (A Hann window of odd length should also be zero at the start, but not at the end. See further below.)

For example, "print, hanning(6)" gives

```
0.00000  0.345492  0.904509  0.904508  0.345491  0.00000
```

whereas the correct sequence is

```
0.00000  0.250000  0.750000  1.00000  0.750000  0.250000
```

Next, generate a window of odd length. An odd-length window should have its largest coefficient at its central element, but the next coefficient should be identical to it, followed by decreasing values out to the end of the window that are the mirror image of values going the other direction from the central element. This results in the last coefficient of the window still matching the second coefficient of the window. The duplication of the highest coefficient in the middle of the window, with the duplicated value being less than 1, seems to me to account for the time (or space) shift in sampling a given periodic function that one finds when using $2n + 1$ samples instead of $2n$ samples. In other words, the center of the time (or space) period covered by the discrete window falls halfway between the two largest data window array elements. Note that this means that the single highest coefficient of a typical window of even length will be 1, but the two highest coefficients of a window of odd length will be equal and somewhat less than 1; i.e. the value 1 would fall halfway between the two coefficients actually present at the discrete sampling interval.

For example, "print, hanning(5)" gives

0.00000 0.500000 1.00000 0.500000 0.00000

but the correct sequence should be

0.00000 0.345492 0.904509 0.904508 0.345491

Note that the third and fourth coefficients above are "identical" within the limits of precision and the default format used in the print statement, as are the second and fifth. Also, note that the same kind of error is produced regardless of the value of ALPHA coded (default is .5, which yields a Hann window) in the call to hanning(), excluding ALPHA=1., of course.

In any case, I think the point Harris made is that a discrete sampling of a window function should not taper to the same value at the end that it has at the beginning because to do so would include the first sample of the *next* period (windowed segment.) So IDL's hanning() gets it wrong for both even- and odd-length windows. :-)

Ref.:

Fredric J. Harris, On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, Proc. IEEE, vol. 66 no. 1, January 1978

Sciences

Scott Bennett, Comm. ASME LG, CFIAG
College of Oceanic and Atmospheric

Oregon State University
Corvallis, Oregon 97331

* Internet: sbennett at oce.orst.edu *

* "Lay then the axe to the root, and teach governments humanity. *

* It is their sanguinary punishments which corrupt mankind." *

* -- _The_Rights_of_Man_ by Tom Paine (1791.) *

Subject: Re: bug in IDL's hanning() window-generating function

Posted by [Harold Cline](#) on Fri, 03 Aug 2001 21:54:33 GMT

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

Just an FYI for everyone on this topic. ;-)

Thanks,
Harold Cline
IDL Product Manager

RELEASE_NOTE for IDL5.5:

HANNING:

The formula to compute the Hanning window has been changed from $0.5 * (1 - \cos(2 * \pi * k / (N - 1)))$ to $0.5 * (1 - \cos(2 * \pi * k / N))$.

Because of the change from $1/(N - 1)$ to $1/N$, the Hanning window will no longer appear symmetric in the time domain, but will give improved results for spectral estimation.

Additional notes:

When filters such as the Hanning are to be used as finite-impulse response (FIR) filtering in the time domain, it is useful to have the filter be symmetric, and go to zero at both ends. However, if you are filtering your time series in preparation for doing spectral analysis, the filter should not be symmetric, and should not go to zero at one of the endpoints. Thus, the new behavior of the HANNING function assumes it will be used for spectral analysis.

Reference: Oppenheim & Schaffer (Discrete-Time Signal Processing, 1999), p. 468-469.

Subject: Re: bug in IDL's hanning() window-generating function

Posted by [bennetsc](#) on Sat, 04 Aug 2001 06:03:42 GMT

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

In article <MPG.15d3c90ca203e28a989e4e@news.frii.com>,

David Fanning <david@dfanning.com> wrote:

> Scott Bennett writes:

>

>> They simply cheated my advisor and

>> refused to make good when confronted with the facts.

>

> I don't know. Pretty strong words for what appears

> from the facts you present to be a misunderstanding

> to me. But I know from personal experience that

Based upon the communications I had with RSI employees, I do not think it reasonable that there was any misunderstanding, except for the fact that no one here had informed me of a rate change for educational customers nor of a linkage of that rate change to elimination of full, new manual sets for each release. That was eventually explained to me. And when the RSI employee explained that that change had occurred, I replied that sending us the release upgrades (i.e.

CDROMS and _What's_New_in..._ books) that they had not sent us would be sufficient (though late.) The RSI employee said they didn't have the books anymore, just the CDROMs. It is not our fault if they are unable to give us exactly what we paid for. Nevertheless, they have the legal obligation to give us what we paid for. I offered what seemed to me a reasonable accommodation to *their* problem in the form of being willing to accept an adequate substitution for the contracted materials. They refused the accommodation that would have made it possible, in view of their disability, to meet their contractual obligation.

I see no evidence whatsoever of any misunderstanding. They obviously knew what the points at issue were. I think that the two employees that I communicated with on this matter also understood that we were not at all happy or satisfied with having paid for something we didn't get and couldn't use. It seems possible, in retrospect, that they also had an estimation of the risk involved in their contract violation.

> educational institutions are extremely difficult to
> work with. The left hand never seems to know what
> the right hand is doing. The people in charge of
> purchasing the software never seem to have a clue
> who is using it. Hence, information like where the
> manuals are and what the terms of the deal are
> are forever being lost.
>
> And the people at educational institutions can be
> difficult.

I can be difficult when someone attempts to shortchange us. I am also not unique in that respect.

>
> I'm still livid over a deal I did at a large university
> in Arizona. They wanted cheap software. I went to
> bat for them, stuck my neck WAY out to make it happen,
> and then found myself accused in this newsgroup
> of being a "software junkie", since I had clearly
> lowered the price to get them "hooked" on the
> software. Scheesh! After that experience I spent
> a couple of years saying to hell with educational
> institutions.

>
> It goes both ways.

>

Yes, it does. That's why each case must be examined

separately. In this case, no outside consultants were even involved. It was strictly between the customer and the vendor. The customer paid and the vendor defaulted.

So I'll finish my work using IDL 5.0. IMO, the amount of money for approximately three years' worth of service contracts, while large relative to a grant budget, is most likely too small to be worth the legal costs of contract enforcement, so RSI will probably get away with it this time.

FWIW, the RSI employee who sent me email yesterday has not responded to my reply to her yet, but perhaps she took the day off.

As for books, David may recall I was an early buyer of the first edition (second printing) of his excellent first IDL book. If I end up working with IDL again someday, I'll probably try to find a copy of whatever edition is his most recent at that time. For now, though, I have what I need to finish my thesis work using 5.0.

Sciences
Scott Bennett, Comm. ASMELG, CFIAG
College of Oceanic and Atmospheric
Oregon State University
Corvallis, Oregon 97331

* Internet: sbennett at oce.orst.edu *

* "Lay then the axe to the root, and teach governments humanity. *

* It is their sanguinary punishments which corrupt mankind." *

* -- _The_Rights_of_Man_ by Tom Paine (1791.) *

Subject: Re: bug in IDL's hanning() window-generating function
Posted by [bennetsc](#) on Sat, 04 Aug 2001 07:18:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

In article <3B6B1D98.1665BFFE@rsinc.com>,
Harold Cline <harold@rsinc.com> wrote:
> Just an FYI for everyone on this topic. ;-)
>
> Thanks,
> Harold Cline
> IDL Product Manager
>
> -----
> RELEASE_NOTE for IDL5.5:

> HANNING:
> The formula to compute the Hanning window has been changed
> from $0.5 * (1 - \cos(2\pi k / (N-1)))$ to $0.5 * (1 - \cos(2\pi k / N))$.
> Because of the change from $1/(N-1)$ to $1/N$, the Hanning window
> will no longer
> appear symmetric in the time domain, but will give improved results for
> spectral estimation.

That's very good to know.

>
> Additional notes:
>
> When filters such as the Hanning are to be used as
> finite-impulse response
> (FIR) filtering in the time domain, it is useful to have the filter be
> symmetric, and go to zero at both ends. However, if you are

First off, data filtering and data windowing (a.k.a. data tapering) are two different things. One can design an FIR data filter by modifying a spectral window to pass or stop certain frequency/wavenumber ranges as desired. So a Hann spectral window (i.e. `FFT(HANNING(length_of_desired_filter))`) could be modified and then used as a time-domain filter.

> filtering your
> time series in preparation for doing spectral analysis, the
> filter should
> not be symmetric, and should not go to zero at one of the
> endpoints. Thus,

That is correct. Note, however, that one often wishes to calculate a spectrum without filtering. In my work, for example, I have not filtered the data before applying the window and taking the transform because I want to see all of the resolvable spectrum. A data window is still necessary to prevent generation of spurious spectral power that would result from the break in "continuity" at the data segment boundaries.

IMHO, IDL's `DIGITAL_FILTER()` provides a much better filter than it is possible to contrive from a Hann spectral window, though it can take a bit of experimentation to find good values to specify for "A" (maximum level Gibbs ringing.) With the right value for "A", `DIGITAL_FILTER()` can even outperform a filter based upon the sinc function (FT of a boxcar window) at the same filter order. For those who want to try this, code a fairly wide filter size and code 43.93125 for the A argument to `DIGITAL_FILTER()`. Then generate a sinc-based filter with the same number of coefficients and compare the frequency responses of the two filters. IIRC, the sinc-based filter has a maximum

Gibbs ringing amplitude in the range of only 2%-3% and the maximum ringing amplitude from the filter you get from DIGITAL_FILTER() is about half that, while the transition ramps of the two filters are practically indistinguishable. (A look at the code for DIGITAL_FILTER() looks like it might be a sinc filter that's been tightened up a bit further by multiplying it by a Kaiser-Bessel data window. Pretty neat.)

A caveat here: certain configurations of arguments to DIGITAL_FILTER() yield filters that should preserve the mean of the data, but don't. This can easily be fixed by normalizing all the coefficients before convolving them with the data. For example,

```
filter = digital_filter(0., .00777778, 43.93125, 1125)
filter = filter / total(filter)
```

> the new behavior of the HANNING function assumes it will be used for
> spectral analysis.

>

> Reference: Oppenheim & Schaffer (Discrete-Time Signal Processing,
> 1999), p.
> 468-469.

>

Thanks for fixing it.

Sciences
Scott Bennett, Comm. ASMELG, CFIAG
College of Oceanic and Atmospheric

Oregon State University
Corvallis, Oregon 97331

* Internet: sbennett at oce.orst.edu *

* "Lay then the axe to the root, and teach governments humanity. *

* It is their sanguinary punishments which corrupt mankind." *

* -- _The_Rights_of_Man_ by Tom Paine (1791.) *

Subject: Re: bug in IDL's hanning() window-generating function

Posted by [bennetsc](#) on Tue, 07 Aug 2001 06:33:57 GMT

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

In article <9kg7jb\$bbt\$1@news.orst.edu>, I, <bennetsc@NOSPAMucs.orst.edu>, wrote:

> I have not filtered the data before applying the window and
> taking the transform because I want to see all of the resolvable
> spectrum. A data window is still necessary to prevent generation

> of spurious spectral power that would result from the break in
> "continuity" at the data segment boundaries.

I see I forgot to mention the other big reason for using a data window, which is to reduce leakage. Not using a data window is the same thing as using a rectangle (i.e. boxcar) window, which results in no leakage at all *iff* the only signals in the data are exactly at one or more of the Fourier frequencies for the time period being transformed and there is no noise in the data. In the real world, however, that is virtually never the case. Energy in any frequency band in a spectrum may lose and/or gain energy to/from any other band in that spectrum by leakage, which kind of runs counter to the whole point of analyzing a spectrum. The amplitude of leakage is worst for the rectangle window (i.e. no window) and is typically very much reduced for other windows.

Sciences

Scott Bennett, Comm. ASME LG, CFIAG
College of Oceanic and Atmospheric
Oregon State University
Corvallis, Oregon 97331

* Internet: sbennett at oce.orst.edu *

* ----- *

* "Lay then the axe to the root, and teach governments humanity. *

* It is their sanguinary punishments which corrupt mankind." *

* -- _The_Rights_of_Man_ by Tom Paine (1791.) *
