Ben Marriage, ben@met.ed.ac.uk writes:

> I believe that this gives different results to
> the inbuilt IDL routines.

   It's possible I'm missing something, but the result of my function
should be the same as yours, but multiplied by sqrt(8/9). You and IDL
are dividing by (N-1) in the definition of variance, I am dividing by
N. Statisticians love to argue about which is appropriate and where,
but for an image of the s.d. it doesn't matter much. The box smooth
always divides by N, but you can always add a renormalisation step and
it'll still be faster than nested loops.

   I use this as part of a routine to do so-called statistical
differencing of images, which brings out fine detail on a varying
background (it's like an unsharp mask weighted by the local image
statistics). In this case the the s.d. image is multiplied by an
arbitrary, user-selectable factor, so the difference between N and N-1
is irrelevant.

Struan