July 25, 2006.

How do we choose a smoothing kernel for SHARP?

John E. Vaillancourt

johnv@oddjob.uchicago.edu

This is fairly standard stuff but I borrow heavily from Dave Cole's thesis() where he has put together an excellent summary of key aspects of map-making. For simplicity I will restrict the discussion to only one-dimension. It is very straightforward to add a second dimension.

This argument is really best produced with pictures. If we get much deeper into this then I can spend time putting together electronic drawings. But it is a lot of work that I do not think is yet necessary.

I have been mathematically explicit in areas where I have thoroughly worked-out the steps. In others I simply sketch (hand-wave) the general ideas where I have at least checked that I am probably not completely crazy.

1. Statement of Problem

From a grid of sampled points on the sky we wish to construct the best estimate of the real continuous sky image. That is, we want each point in our new map's dense grid to have the value we would have measured if we had actually pointed the polarimeter and telescope at that point on the sky.

Assuming that we do not attempt any fancy super-resolution algorithms then the best we can hope to extract from our data is the convolution of the sky with our telescope's beam. That is, if $s(x, y)$ is the distribution of sky brightness, and $b(x, y)$ is the point spread function of our instrument then we will measure

$$t(x) = b(x) * s(x) \tag{1}$$

In reality we do not directly measure this value. Our goal is to recover $t(x)$ after sampling it at discrete points.

2. Fourier analysis

If the discrete sampling is a regular pattern which continues to infinity in both directions then we can write our measurements as

$$t'(x) = \mathrm{III}\left(\frac{x}{l}\right) \cdot t(x), \tag{2}$$

the product of the telescope image $t(x)$ with the picket fence sampling function

$$\mathrm{III}(x) = \sum_{i=-\infty}^{\infty} \delta(x - i). \tag{3}$$

Using the convolution theorem to transform to Fourier space we have

$$T'(u) = l\ \mathrm{III}(lu) * T(u), \tag{4}$$

$$\text{where} \quad T(u) = B(u) \cdot S(u), \tag{5}$$

and $T'(u)$, $T(u)$, $B(u)$, and $S(u)$ are the Fourier transforms of $t'(x)$, $t(x)$, $b(x)$, and $s(x)$, respectively. The function $T'(u)$ is composed of an infinite number of copies $T(u)$, one centered at every value of $u = i/l$ for all integer i.

A discussion of the Nyquist criterion is not really critical here, but I will add it for completeness. Consider that $B(u)$ is finite with half-width r. (This must be true because it is the Fourier transform of $b(x)$, which is itself the Fourier transform of the finite telescope/instrument aperture. For a diffraction limited instrument this implies that $r \sim D/\lambda$, where D is the telescope diameter.) Since $B(u) = 0$ for $u > r$ we also have $T(u) = 0$ for $u > r$. If l is sufficiently small then the copies of $T(u)$ in $T'(u)$ do not overlap:

$$l < \frac{1}{2r} \approx \frac{\lambda}{2D} \tag{6}$$

In other words, the copies do not overlap if the sampling spacing satisfies the Nyquist criterion; l is less than half a beam width. We are very close to this criterion for SHARC/P ($l \approx 4.7''$, $1/r \approx 10''$) so hereafter I will assume it is true.

In this case we can recover $T(u)$ by masking away all but the central copy of $T'(u)$ with a boxcar of width l: $\Pi(u) = 1$ for $|u| < l/2$, $\Pi(u) = 0$ for $|u| > l/2$. Then

$$T(u) = \Pi(u) \cdot T'(u) \tag{7}$$

The inverse Fourier transform yields

$$t(x) = \mathrm{sinc}(\pi l x) * t'(x) \tag{8}$$

So recovering $t(x)$ only requires convolving the sampled data with $\mathrm{sinc}(\pi l x)$. For computational simplicity we can replace the sinc function with a gaussian of FWHM $\approx l$. This has little practical effect on the final answer. Note that if $t'(x)$ represented our irregularly gridded data then equation (8) would represent the method currently employed in `sharcsolve` and `sharp_combine`.

Aside 1: The spatial resolution resulting from this method is the larger of $2l$ or the beamsize $1/r$. That is, unless you are also Nyquist sampled, decreasing the beam size does not result in increased spatial resolution. Neither does sampling beyond the Nyquist limit.

Aside 2: In two-dimensions we have a choice of pillbox shapes to mask out the central copy. The Fourier transform of a rectangular pillbox is a two-dimensional sinc function: $\mathrm{sinc}(\pi l_x x)\,\mathrm{sinc}(\pi l_y y)$. A circular pillbox results in an Airy function. If we are truly Nyquist sampled then these two choices yield the same answer, since $T(u)$ cuts-off before the edge of any mask shape.

Aside 3: If $T(u)$ truly cuts-off at $u > r$ and the Nyquist criterion holds ($r < 1/2l$) then the mask function can be allowed to cut-off at any u: $r < u < 1/l - r$. Additionally, this cut-off does not need to be sharp as in a boxcar. One could use a mask which goes smoothly from unity at $u = r$ and $u = -r$ to zero at $u = 1/l - r$ and $-1/l + r$ (e.g. a cosine or Hanning function). The transform of such a mask is wider than a sinc (or Airy) function and has reduced sidelobes. A gaussian would be an even better approximation to this function than to a sinc or Airy function.

Finally we must consider that our assumption of a regular grid is not true in the case of SHARP. Here I will simply state (without proof) that the regularity of the grid is critical to the Fourier method if we wish to achieve perfect image reconstruction. The irregular grid causes significant aliasing which contaminates all spatial frequencies. It is not possible to simply mask out these frequencies.

Re-gridding

One way around the problem of an irregular sample grid is to interpolate onto a regular grid before carrying out the convolution operation. The success of this method is highly dependent upon the accuracy of the initial interpolation. However, it will be sufficient if the mean separation of the sampled data is smaller than the new grid spacing. In the case of SHARP this is dependent upon the number of co-added files and the relative amount of sky rotation captured.

For the sake of argument let's say we have sufficient data for an accurate regular grid interpolation. Further assume we use a weighted average to interpolate from our sample grid onto a regular grid:

$$z'(x_o) = \frac{\sum_{i=1}^{n} z(x_i)\, w_i}{\sum_{i=1}^{n} w_i} \tag{9}$$

where $z(x)$ is the data, $z'(x)$ is the smoothed data, and w_i is a decreasing function of distance from the data point x_i, to the new grid point x_o. If w_i were gaussian then equation (9) would represent the method currently employed in `sharcsolve` and `sharp_combine`.

We can re-write equation (9) as a convolution of the data $z(x)$ with a kernel $w(x)$

$$z'(x) = w(x) * z(x) \tag{10}$$

We set $t'(x) = z'(x)$ and follow with the Fourier method convolution:

$$t(x) = f(x) * z'(x) = [f(x) * w(x)] * z(x) \tag{11}$$

where $f(x) = \mathrm{sinc}(\pi l_g x)$ or an appropriate gaussian replacement. Note that the width of this function is given by the re-gridded pixel pitch l_g, not the instrument pixel pitch l. If $f(x)$ and $w(x)$ are both gaussian then their convolution can be replaced with a single equivalent gaussian. Again, this is the algorithm currently employed in `sharcsolve` and `sharp_combine`.

Now What?

The big question that drove me to compose this document was how to choose the widths of the smoothing kernels. The Fourier method makes it clear that the width should be a function of the pixel-to-pixel spacing (or pitch). As long as the Nyquist criterion is satisfied, the width of the smoothing kernel is dependent only on the pixel pitch and is completely independent of the beam size.

Similarly, if we have interpolated onto a regular grid then we should smooth with a kernel of width given by the pixel separation on the new grid regardless of the original instrument parameters.

But what do we do if we do not have a regular grid? If we smooth before using the Fourier method then the width of the interpolation kernel $w(x)$ does not appear to be dependent on any instrumental parameter such as pixel pitch or beam size. However, it seems reasonable

that one should consider these parameters when choosing both the width of $w(x)$ and the regular grid spacing l_g.

In the end, whichever width and pitch we choose, the final smoothing kernel $f(x) * w(x)$ will have a width given by the quadrature sum of the other two widths. I do not expect that this will differ much from a kernel given by the original instrument pixel pitch. To wave my hands here, perhaps it is best to choose $l^2 \approx l_g^2 + l_w^2$, where l_w is the width of the interpolating function $w(x)$. This would then be equivalent to simply using the Fourier method, ignoring the fact that the sampling grid is irregular.

Displaying the result

Lastly I should point out that the maps saved in the output FITS file resulting from our analysis is not $t(x)$ but

$$t_{\rm fits}(x) = \mathrm{III}\left(\frac{x}{l_g}\right) \cdot t(x). \tag{12}$$

The finite sampling has nothing to do with the analysis or smoothing algorithm chosen but is simply a result of the digital nature of an electronic file. It is not necessary for the FITS data pitch to be equivalent to l_g, it is chosen for computational convenience. The image we show in a map depends on the algorithm used to display $t_{\rm fits}(x)$. FITS viewers tend to use individual pixels for each value of $x = il_g$. I generally have IDL use a default interpolation algorithm (triangulation ?) to generate smooth contours from $t_{\rm fits}(x)$, but that can be modified with some work.