

initiated FDF follow-up investigations in other wavelength ranges. The most extensive of these follow-up studies is a mapping of the FDF at radio wavelengths which is being carried out in co-operation with colleagues from the MPI for Radioastronomy, Bonn, at the VLA.

At present our FDF project is still work in progress. While our deep photometric study can probably be completed within the next few months, the spectroscopic programme and the follow-up investigations may keep us (and other interested groups) busy for years to come. We hope – and are optimistic – that these studies will eventually result in new and important insights into the evolution of our universe, which will perhaps be reported in future issues of *The Messenger*.




Figure 7: FORS spectrum of the Lyman forest of the FDF quasar Q0103-260.



# Astronomical Image Resampling

N. DEVILLARD, ESO

## 1. Introduction

Resampling an image is a standard operation used in astronomical imaging for various tasks: changing the scale of an image to superimpose it on another, shifting an image by a non-integer offset, rotating an image by an arbitrary angle, or deforming an image to counter detector or optical deformations, are the usual operations in need of image resampling. In a more general way, this operation is referred to as *image warping* in the digital image processing world.

## 2. Sampling a Signal

The stellar light landing on an astronomical detector is a continuous optical signal. It is sampled by the detector at precise positions, yielding a regular grid of intensity values also known as pixels (for *picture elements*, often abbreviated to *pel*). The initial signal carries by definition an infinite amount of information (because of its continuity), but it has been reduced to a finite number of values by the detector system. The sampling theory has proved that it is possible to reconstruct the initial signal from

the knowledge of its samples only, provided that a certain number of assumptions are fulfilled.

In our case, we will assume that this is always the case, i.e. that the pixel sampling frequency is always greater than twice the greatest spatial frequency of the image (Shannon or Nyquist sampling). This is true for most ground-based telescopes because the instruments have been designed so, but this does not apply to some HST instruments for example. In that case, it is still possible to retrieve the signal, modulo some assumptions. This is not discussed in this paper.




Figure 1: Top left is the original signal, top right its Fourier transform. Bottom left is the sampled signal and bottom right the Fourier transform of the sampled signal.

## 3. Sampling Theory

This section provides some background about the sampling theory in general. For convenience, a simple 1-dimensional signal  $S(t)$  will be used as a conventional signal, because in this case the same rules apply to images understood as 2-dimensional signals (an intensity as a function of the pixel position on the detector).

The Fourier transform is a very convenient tool to perform an academic study of this case. It also helps to see the signal in another space to understand exactly what is happening during the various operations performed on it.

The representation of a signal  $S(t)$  in Fourier space is a distribution that has no energy outside of a certain frequency range, i.e. it has non-zero values

only inside a given range. This representation is symmetrical for real-life signals, so the non-zero support is completely defined by the cut-off frequency  $f_c$ . Figure 1 shows what a time signal might look like, and what its representation in Fourier space would be. Now, what happens when this continuous signal is sampled? According to Fourier theory, sampling a signal with an ideal sampler is equivalent to convolving it with a delta function, which is equivalent in Fourier space to multiplying the signal by a comb function (a regularly-spaced collection of diracs).

The idea for reconstruction is that we would like at some point to reconstruct completely the continuous signal, and then re-sample it, i.e. take new samples at arbitrary positions. The sampling theory enables that the right assumptions are fulfilled. What it means in practice in Fourier space, is that side replica of the sampled signal spectrum should be removed (the central part of the spectrum isolated).

This is easily achieved in Fourier space by multiplying the sampled spectrum by a door function (a function that is 0 everywhere but 1 where the central spectrum is, see on figure 1 the box around the central spectrum replica). Unfortunately, this ideal door function corresponds to a *sine* function in the real space, which has infinite support and is thus not implementable in real-life. In other words, perfect reconstruction can only be achieved mathematically. To build a system that would achieve the same, we need to have an infinite quantity somewhere in the system, which does not belong to our real-life domain.

This ideal door function can be approximated by functions that do not have an infinite support, though. But as approximations are not truly the function, the reconstruction will not exactly stick to the real signal. The amount of errors in the reconstructed signal will depend on how the approximation is done. The approximated door function is usually called a *kernel* in the literature.

The search for “good” kernels has led to a veritable zoo of functions fit for one purpose or another. Many authors in very different fields have found kernels that have some interesting properties that pair ideally with some characteristics of the signals they studied. The filter domain is one of the key fields in digital electronics, whether it applies to sound machines (stereo systems), to images (television) or any kind of sampling machine. Various kernels are reviewed in the next section.

## 4. Image-Space Interpolation

Another way of reconstructing the missing parts of the signal between samples, is to stay in the real space and create the missing information based

on e.g. some assumptions on the signal smoothness. In electronics, the following reconstruction schemes are often used for reconstruction:

### 4.1 Nearest neighbour

Maybe the simplest idea is to say that the signal value is the one of the closest sample. This is known as a zero-order blocking filter, or nearest-neighbour interpolation. If we study that carefully, we see that this is strictly equivalent to applying a box kernel to the real space, thus a sine function in the Fourier space.

It means that the central replica of the sampled spectrum will be multiplied by a sine function, which is actually a very bad approximation of the ideal box kernel. This leads to obvious artefacts like jagged edges in images or granularity in sound. You can actually hear that kind of artefact when using noisy digital telephone lines like mobile phones: your voice appears deformed as if it were ringing.

This is due to the fact that the mobile phone receives voice samples one by one, and if the line is cut for a tiny amount of time (some samples are missing), it keeps emitting the same sound until it receives more information. This comes back to sub-sampling the voice, and reconstructing it with a very bad reconstruction approximation. The voice spectrum is folded and you get that kind of nasty artefacts.

Nearest-neighbour interpolation is the cheapest way of reconstructing a signal and also the dirtiest. It might be OK if you need speed (e.g. to display a zoomed image), but certainly not if you want to preserve the smoothness of the input signal.

### 4.2 Linear interpolation

One idea is to draw slope segments between samples. This corresponds strictly to a linear approximation of the values between known samples. In electronics, this is also referred to as a first-order blocking filter. Now what is the quality of a linear interpolation in real space, if we try to compare it to other kernels in Fourier space?

Drawing line segments between samples is strictly equivalent to convolving the input signal with a triangle function (demonstration left to the reader). Convolving the signal with a triangle function is then equivalent to multiplying its spectrum by the Fourier transform of a triangle, which is a *square sinc* function.

A *square sinc* is better than a *sinc*: it is never negative and goes to zero much faster, which means that other replica of the central spectrum are less likely to be included in the reconstructed spectrum. However, it is still a very bad approximation of the ideal box

function. It obviously attenuates high frequencies and tends to deform low-frequencies. The ringing goes towards zero but not that fast, which tends to include other replica in the reconstructed signal.

Linear interpolation is another cheap way of reconstructing a signal and might be a good candidate when artefacts are not so much an issue. For every-day digital pictures (television or photos), this is still valid, but for signals like astronomical images it is more reasonable to use a smoother reconstruction.

### 4.3 Spline interpolation

To preserve a certain smoothness of the signal between samples, one can use a set of smooth functions such as splines. Describing the maths behind these functions is outside the scope of this article, it is enough to know that they tend to reproduce quite well signal variations, and they are expensive to compute.

What do they correspond to in Fourier space? A spline function is usually obtained by convolving a box function with itself a certain number of times. How many times the convolution is applied determines the level of the spline functions, which are also referred to as N-order blocking filters in electronics.

In Fourier space, that means the N-th power of the *sinc* function. As N goes to higher degrees, the kernel tends to look more and more like a gaussian shape with steep slopes. This is a much better approximation of an ideal box function and usually yields the best visual results with normal images.

Because they are expensive to compute, splines are usually used in their third order (cubic splines). You can find that kind of interpolation for images in all respectable image-processing software packages (like PhotoShop or GIMP). For astronomical images, this is still a questionable method due to the noise in images, bringing very high frequencies which are not so well handled by cubic splines.

### 4.4 Fourier-based kernels

These interpolation methods are based on some study happening in Fourier space, but they are nonetheless implemented in the image space without ever having the need to convert the whole image by use of FFT or equivalent means.

As mentioned above, interpolation kernels can be found by hundreds in various fields. The most appropriate for image processing are quoted here:

- Lanczos2 or Lanczos3: these kernels are based on cosine functions, they are shown to respect quite well the low-frequency parts of the signal but




Figure 2: Three different interpolation schemes.

smooth a little bit the high-frequencies. This is an issue for noisy images such as astronomical images.

- Hann and Hamming kernels are also based on truncated cosine functions. They also tend to smooth out high frequencies and induce some rippling in the image space, but produce very acceptable quality for usual images.

- Hyperbolic tangent kernel: this is one of the best-known approximations of the ideal box function, obtained by multiplying symmetric hyperbolic tangent functions. It preserves quite well both high and low frequencies (as expected for a faithful approximation of a box function) and induces almost no rippling in the image space.

The latter kernel (hyperbolic tangent) is the one implemented in the jitter imaging recipe in the ISAAC pipeline. More information can be found about its definition, implementation, and evaluation in the corresponding article: <http://www.eso.org/projects/dfs/papers/jitter99/>

#### 4.5 Examples

Figure 2 shows an example of the nearest neighbour, linear, and spline interpolations of the same signal.

Figure 3 shows various examples of image interpolations applied to zoom an example image. Top left is nearest neighbour, top right is linear, bottom left is cubic spline, and bottom right is a hyperbolic tangent based interpolation kernel.

### 5. Spatial Transformations

Now that we know how to create new samples out of a sampled signal, the next question is: where do we want to place these new samples? This of course depends on the nature of the operation to be performed on the image.

Shifting an image by a non-integer offset for example, is a typical operation in need of pixel interpolation. Obviously, the worst case is when the shift must be done up to exactly half a pixel. This is the place that is the furthest from the real known samples, thus having the most invented properties.

Scaling an image (i.e. zooming it in or out) is also a common operation that involves interpolation. In your favourite image displayer (RTD or Saoimage), when you request to double the zoom on your image, the software behind actually performs a pure pixel replication, which is the same as nearest-neighbour interpolation. While this is very bad for signal properties, it is acceptable for a user display since you can very easily see where artefacts are. If you need to scale an image by a

non-integer factor e.g. to superimpose it with another image taken with a different pixel scale on the sky, you will need more accurate interpolation methods.

Correcting an image for detector deformations (pincushion or barrel, depending on the orientation) will most likely require a 2d polynomial of second degree at least. A true mapping of the deformation is more easily expressed in radial co-ordinates and yields complicated deformation formulas. Whatever the deformation you choose to correct for, you will end up with a formula that explains how to compute the one-to-one correspondence between a pixel in the deformed image and a pixel in the corrected image.

Other deformations can easily be implemented if they can be expressed as a one-to-one relationship between pixels in the original image, and pixels in the deformed (warped) image.

### 6. Possible Implementations

There are only two ways of implementing a resampling scheme for images: direct or reverse. The direct way is going from the original image to the deformed one, literally spraying the input pixels onto an output grid. The reverse way is going from the deformed image back to the original one, computing samples only where they are needed.




Figure 3: Various interpolation schemes. Top left is nearest neighbour, top right is linear, bottom left is cubic spline and bottom right is hyperbolic tangent.

## 6.1 Direct warping

Direct warping is taking pixels from the original image and trying to throw them onto an output grid that represents the output, warped image. The main drawbacks of this method are the following:

- The output grid needs to keep track of how many pixels have landed in each section, to be able to correctly normalise the output value. This can be a nightmare to implement efficiently.

- There might be some deformations that will leave holes in the output grid. If for some reason the output grid is denser than the input one, the risk is to have so many holes that the output image does not make sense. To avoid that, one can oversample the input image to make it of similar density to the output image, but then the memory consumption is growing by the same amount. For large images, this is not a solution.

- If there is no possible assumption about the way input pixels are thrown onto the output grid, it is impossible to guarantee that the output pixels will be written sequentially. This makes it hard to optimise I/O accesses, and slow down to the extreme the processing of large images, even for simple transformations.

Direct warping is usually considered too dangerous to implement because of potential holes in the resulting image, and because of speed concerns. Image processing is a domain that needs optimisation no matter what the underlying hardware is. A pixel operation that needs a millisecond to run will have you wait more than 17 minutes in front of your screen if you have to process an ISAAC image, or 71 hours if you are processing a VST image.

The best optimisations in the image processing domain usually reach about a hundred clock ticks to achieve. On a modern PC running at 500 MHz, considering that the image processing software is alone running on the CPU, you will still have to wait 2 seconds per ISAAC image and 54 seconds per VST image. And that is for CPU time only, a bad I/O optimisation forcing the soft-

ware to go to the disk regularly would multiply these figures by a factor thousand at least.

Direct warping is at the heart of the “drizzle” method implemented for HST image reconstruction. Drizzle makes use of a convolved linear interpolation scheme that brings more artefacts than simple linear interpolation. If big images are involved, or tricky (non-linear) deformations, or bad pixels, or large amounts of noise, this method is likely to create false information in the output image. The same is valid for almost all astronomical image processing packages: most of them use a linear reconstruction scheme for default interpolation method.

## 6.2 Reverse warping

Reverse warping is considering first the output grid as the image to obtain. The size of the grid can be determined by observing the transformation function between deformed and corrected image. For example: if the transformation is a scaling by a factor 2, the output grid will simply be 2 times bigger than the input image, containing 4 times more pixels.

The method is looping over all output pixels, computes by means of a reverse transformation which are the contributing pixels from the original image, applies an interpolation scheme as described above, and writes the result to the output image.

This has several advantages:

- The output grid cannot have holes, since all output pixels are reviewed.

- The output grid will be sequentially visited, allowing to cache the results to avoid going to disk too often.

- Because of the sequential nature of the operation (no matter what the transformation is), it is relatively easy to optimise the pixel accesses in a generic way, helping to achieve optimised but also portable code (i.e. a software that is fast no matter what the underlying hardware and OS are).

- The implementation is relatively easy to write and read, allowing a better maintenance.

If you implement a generic interpolation kernel scheme that allows the user to select the kernel to use, reverse warping becomes an easy task to implement and allows for large gains in speed in the resulting software.

This is the method implemented in the warping tool offered in the eclipse package (<http://www.eso.org/eclipse/>), and at the heart of the ISAAC imaging pipeline. The default kernel is the hyperbolic tangent mentioned above.

## 7. Conclusions

Astronomical image resampling is a complex operation that involves recent theorems in the field of signal and image processing, together with information theory knowledge to be carried out properly.

No information was given about the various methods that can be used to identify the transformation between two images. This is a complete research domain in itself, refer to the appropriate literature for more information. The emphasis in this article has been put on the various methods that can be used to interpolate pixels in an image, and efficient ways to implement them.

Usual interpolation schemes are shown to be insufficient in the usual noisy astronomical images, adding aliasing and other various artefacts into the images. There are more precise kernels such as the hyperbolic tangent, that are more suitable for astronomical image handling. It might also be a good idea to look into pre-processing filters to apply to the input images before trying to re-sample them.

If you care about the quality of your images whenever you have to apply resampling operations, you should query your favourite data reduction package to check out what kind of interpolation scheme is actually implemented behind the scenes. Choosing linear interpolation is rarely a good solution in the case of noisy images.

# OTHER ASTRONOMICAL NEWS

## Portugal to Accede to ESO (from ESO Press Release 15/00, 27 June 2000)

The Republic of Portugal will become the ninth member state of the European Southern Observatory.

On Tuesday, June 27, during a ceremony at the ESO Headquarters in Garching (Germany), a corresponding Agree-

ment was signed by the Portuguese Minister of Science and Technology, José Mariano Gago, and the ESO Director General, Catherine Cesarsky, in the presence of other high officials from Portugal and the ESO member states .

Following subsequent ratification by the Portuguese Parliament of the ESO Convention and the associated protocols, it is foreseen that Portugal will formally join this organisation on January 1, 2001.