Spectral PSF Subtraction I: The SPSF Look-Up-Table Method

P. MØLLER, ESO

Introduction: Why Spectral PSF Subtraction?

The first FORS follow-up run of the STIS/NICMOS survey for Damped Lyα (DLA) galaxies was completed in August 1999. The purpose of the FORS spectroscopy is to search for Lyα emission lines from candidate DLA galaxies with projected distances from the underlying quasars as small as 0.30 arcsec. The candidate DLA galaxies are typically some 8–10 magnitudes fainter than the underlying quasars.

To be able to look for the faint Lyα emission lines hidden under bright quasar spectra, we shall need a spectral reduction code optimised to subtract the two-dimensional (henceforth 2D) signal of the quasars with extremely high accuracy. By a strike of good fortune, a 1D spectral extraction code was designed several years ago with the specific aim to provide the best possible determination of the Spectral Point Spread Function along the slit. As outlined below there are minor, but important, differences between a 2D imaging PSF and the 1D Spectral PSF. To avoid confusion, we shall in what follows write SPSF for the 1D Spectral PSF.

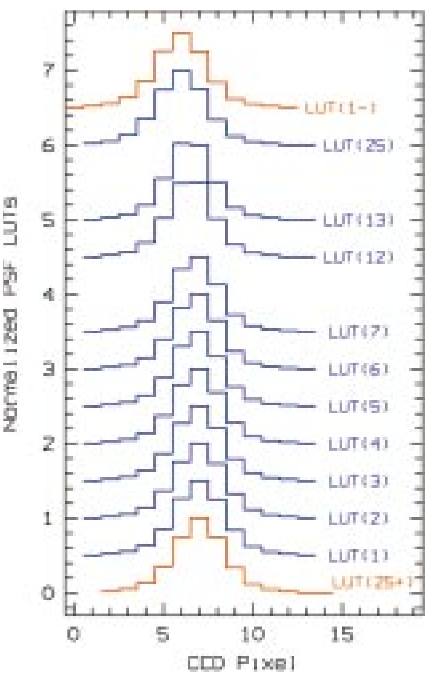

Figure 1: Attempt to represent the Look-Up-Table (LUT) used for the prediction of the Spectral PSF (SPSF) at any wavelength bin of the 2D spectrum. In this example, the increase in sampling (n) is 25 and the length of the SPSF (l) is 13. Each LUT entry (a row in Fig. 1) is a string of 13 numbers making up a single SPSF realisation (see Section 2).

The example presented above, the search for high redshift DLA galaxies, is merely one field in which SPSF subtraction is essential. The technique is potentially extremely useful in any field of astronomy dealing with separation of objects with small projected distances on the sky, or even objects covering each other. In this and in a following issue of *The Messenger* we are going to present three reports aimed at detailing the exact properties of the algorithm to be used for the reduction of our FORS spectra. This first report presents the details of the algorithm itself. The second report (Møller et al., 2000) presents, as a simple example of a research field in which exact SPSF subtraction can provide interesting new results, the serendipitous discovery of extended Lyα emission from the host galaxy of a radio-quiet high-redshift QSO. The third and last report will present the first confirmed identification of an intervening DLA galaxy from the STIS/NICMOS survey.

2. Maximum S/N Extraction of 1D Spectra

When 2D photon-counting detectors were first attached to spectrographs, it was immediately realised that by applying a proper weighting scheme to all pixels of the 2D spectrum, one may optimise the signal-to-noise (the S/N) of the extracted 1D spectrum. Determination of the optimal weights requires knowledge of the SPSF. The more precise guess of the actual SPSF one is able to obtain, the higher the resulting S/N of the extracted 1D spectrum. It is also true, however, that the improvement in S/N when going from a "fairly good SPSF" to the "exact SPSF" will be minor in most cases. Therefore, for the purpose of 1D extraction there is no strong incentive to optimise the model SPSF beyond a certain point. Extraction algorithms not optimised in this sense present no problem as far as the extraction of 1D spectra is concerned, but they will not be adequate for our purposes. This point is of some importance and will be detailed below.

In what follows we shall consider a 2D spectrum which falls roughly along the columns of a CCD. The rows will then (if we for now ignore the 2D distortion in wavelength space) represent bins in wavelength space.

A number of effects complicate the determination of the optimal weights:

• **The SPSF is wavelength dependent.** The seeing is typically better in the red than in the blue. For broadband imaging this means that a blue object will have a slightly larger PSF than a red object. In spectroscopy, it means that the SPSF will be broad in the blue end, and narrow in the red end of the spectrum.

• **Focus changes along the spectrum.** Most spectrographs are not simultaneously in focus over the entire wavelength range. One therefore has to choose a compromise for the best over-all focus. This may introduce a quite unpredictable change of the SPSF width along the spectrum.

• **The spectrum does not lie exactly along the detector rows (or columns).** Most spectrographs have a distorted spectrum which is "C" or "S" shaped. Hence, the SPSF is typically shifting from one column into several other columns along the way.

The SPSF itself is a continuous and smooth function. However, what we actually observe is a discontinuous step-function induced by the process of observing the continuous function with an array of finite sized pixels. To emphasise the difference (which is extremely important in what follows) we shall refer to the observed step-function as a "realisation" of the SPSF. For a given SPSF, there will be an infinite number of different realisations, all differing by shifts on the detector smaller than a pixel. An example of several different realisations of the same SPSF from a FORS1 spectrum are shown in Figure 1.

Horne (1986) reported on an extremely powerful and simple algorithm that, as long as the spectrum is aligned well with the CCD rows or columns, determines the SPSF realisations along the spectrum with very good accuracy. A generalisation of the Horne algorithm to the case of spectra *not* aligned well with the CCD was found by Marsh (1989). As pointed out above, optimal weights for maximum S/N 1D extraction routines are derived from the SPSF realisations, and that was the goal of the Horne and Marsh algorithms. In the limit of very bright objects, the optimal weights are all identical and equal to 1, for faint objects the SPSF realisations are needed to determine the weights. For *very* faint sources and for sources with strong absorption features, several additional problems come into play. First of all the SPSF is essentially impossible to determine where there are bright skylines and where there is strong absorption (such as in the Lyman forest of quasar spectra). Hence, some sort of smoothing or fitting algorithm must be used to interpolate between the poorly determined bins. The smoothing (or in the case of the

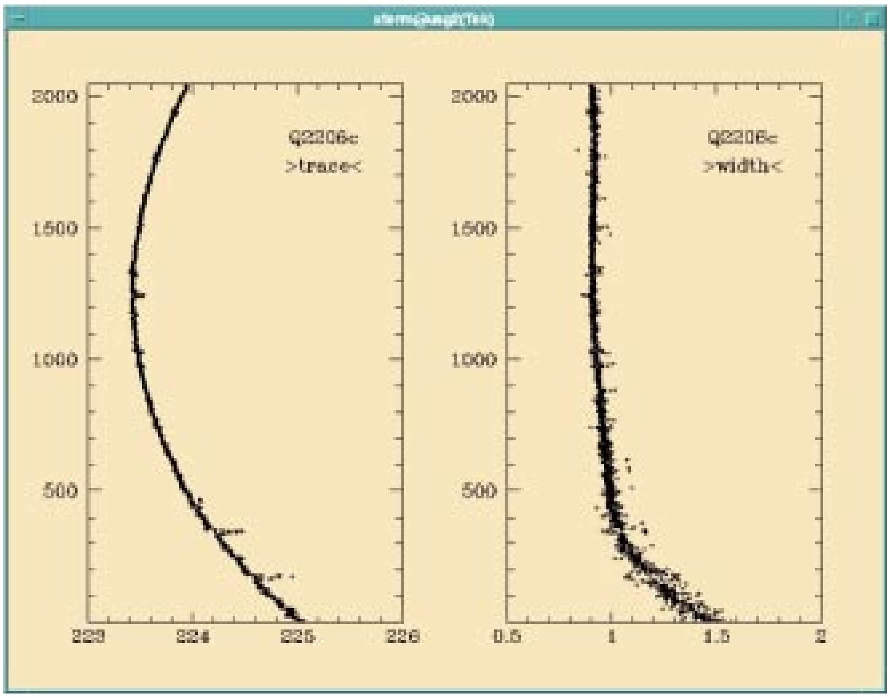

Figure 2: Output of the "trace" and "width" commands. The left frame plots the centroid of the SPSF and covers three CCD columns. The right frame plots the relative width of the SPSF as a function of wavelength (blue is down, red is up). The data are from FORS1/GRISM-600B taken under excellent seeing conditions, but note that we have rotated the FORS data 90 degrees.

Horne algorithm the polynomial fitting) introduces some imperfections in the weights. For the purpose of extracting the 1D spectrum, those imperfections are mostly of no consequence. For the purpose of searching for faint objects hidden under the bright spectrum, there are three specific problems. First, the SPSF determination is *in situ.* This means that the hidden object will be treated by the code as part of the SPSF, and it will be included in the data for the polynomial fitting. A narrow emission line then becomes somewhat harder to detect, a broad emission line will most likely be missed entirely, having been interpreted as a wavelength-dependent change in the SPSF. Second, most codes ignore the extended wings of the SPSF because they contain essentially no information (for the purpose of extracting a 1D spectrum). Fitting to low S/N wings *in situ* would in any case be a fit to mostly random noise. The wings are nevertheless important to remove if one wishes to search for faint extended emission. Third, since the main goal of our project is to search for faint emission in the Lyman forest of high redshift QSOs, we would expect major problems if we were to use the SPSF determined *in situ* on the basis of the few bits of QSO left between the many strong absorption lines in the forest.

All three problems listed above are related to, or made worse by, the fact that the SPSF determination is performed locally in each separate wavelength bin (*in situ*). Hence, the solution is to use a global SPSF determination algorithm. A code based on this principle was developed and installed at Copenhagen University Observatory in 1989. It has been used only by myself and collaborators, mostly for extraction of quasar spectra but lately also for extraction of faint galaxy spectra of early-type galaxies at intermediate redshifts (Treu et al., 1999). A short summary of the method was given in Møller & Kjærgaard (1992) but details of the algorithm have never previously been published.

The code was developed for the purpose of 1D spectral extraction, but an extra option was included: The possibility to subtract the 2D version of the extracted spectrum. The 2D spectrum-removal option was intended mostly as a check that the algorithm had worked properly. In addition, since the code included "cosmic" rejection via sigma clipping, it provided a possibility to check if cosmics had been overlooked by the code. The 2D-subtraction option has, however, already been successfully applied in the search for DLA galaxies (Møller 1999).

3. The SPSF Look-Up-Table Method

3.1 General Overview: The Four Steps

The algorithm is iterative and based on four simple steps: (i) Trace the spectrum (determine the centroid of the SPSF in each wavelength bin), (ii) determine the width of the SPSF in each wavelength bin, (iii) determine the amplitude of the SPSF in each wavelength bin, (iv) predict the realisation of the SPSF in each wavelength bin.

Figure 2 shows an example of the first two steps, the "centroid trace" and "width" determinations on a FORS1/Grism-600B 2D spectrum. The second and third steps in the above list are in reality performed simultaneously via a two-parameter minimum χ^2 fit as follows. For a given guess of the width, the predicted SPSF realisation is scaled linearly in width, and the best-fitting amplitude is determined. The best global two-parameter fit is found. It is easy to show that minimum χ^2 fitting of an assumed SPSF realisation to the actual data in each wavelength bin is mathematically equivalent to a maximum S/N 1D extraction. After completion of the first three steps, a new and better determination of the SPSF is possible, and a new iteration of the four steps can be performed. When the iterations converge towards a global minimum of the total χ^2, the process may be stopped and the fitted 2D spectrum subtracted from the input data frame.

The determination of the SPSF realisations is done on basis of the latest "trace". The trace command determines the centroid of the SPSF with high accuracy in each wavelength bin. The code then subdivides all pixels into n subpixels in the spatial direction (along the slit). Effectively it then builds up an empirical Look-Up-Table (LUT) of n different realisations of the SPSF. Each of those realisations correspond to the subset of rows in the input data frame where the trace command determined that the centroid fell inside the same given sub-pixel. In practical terms, the LUT is an array of ($l \times n$) numbers where l is the selected length of the SPSF (in CCD pixels) and n *is* the resolution into sub-pixels which is deemed necessary/possible for the given data. Typically we use $n = 25$, and will do so in the example given below.

3.2 The Core Algorithm: Creating the LUT

The final quality of the 2D spectral subtraction is depending critically upon how carefully the LUT was constructed. The first step of the LUT construction is the summing of all data-rows with the same sub-pixel trace position, which pass a set of additional quality conditions (S/N, no cosmics found in that bin). After the distribution of raw LUT data is completed, all n realisations are normalised. In Figure 1 we show (in blue) the normalised LUT determined from the example FORS spectrum used in Report 2. Clearly, the first entry (named LUT(1) in Figure 1) and the last entry (LUT(25)) are neighbour realisations in the same sense that e.g. LUT(3) and LUT(4) are, except that they are shifted by a full CCD pixel along the slit. This is exemplified by the two realisations reproduced in red in Figure 1. E.g. LUT(25+) is identical to LUT(25) but shifted by one pixel. As the last step in the LUT construction algorithm, this process of expanding the 2D LUT by "adding on" shifted 1D LUT realisations is repeated $(l \times n)$ times. By selection of the column corresponding to the leftmost pixel of LUT(1), the code obtains a single string of the $(l \times n)$ numbers making up the complete LUT. This string of numbers hence contain the entire LUT, only now re-ordered in a way so that it is a single smooth and continuous function. This function corresponds to the input SPSF at n times improved sampling, but still containing the observationally induced CCD pixel smoothing.

It can be seen that if a single, or a few, of the individual entries in the 2D LUT grid are missing (by unfortunate locations of cosmics, sky-lines and absorption lines), then a whole row will be missing in the 2D LUT grid. In the continuous high-resolution representation of the LUT, however, a missing entry will only result in a set of single missing values at each 25th sub-pixel. Those can easily be determined via interpolation between good values. After interpolation to determine unknown values, possibly some smoothing if the spectrum in general has low S/N, the process is reversed, and the 2D LUT is reconstructed from the interpolated/ smoothed high-resolution representation. This is the last step of a given iteration. If another iteration is started, the new LUT will be used.

4. Degeneracy of Results and Concluding Remarks

We have presented an algorithm that is specifically aimed at performing the best possible removal of a 2D spectrum to reveal faint emission-line objects "hidden" under a bright spectrum. A few notes should be added. **First:** As detailed in Section 2, an algorithm with this goal has to use a global approach to the determination of the SPSF. A global approach is clearly incompatible with the fact that the SPSF varies along the spectrum; some additional assumption is required. In the code described here, we have assumed that the SPSF at any wavelength can be obtained via a linear scaling in width of the same base SPSF. **Second:** Clearly, if two point sources are super-imposed at exactly the same position, there is no way the code can tell that there are in fact two different objects. It will regard the summed spectrum as that of a single point source. **Third:** If the emission-line object is not a point source, and if it extends across the bright point source, the code will try (as a default) to assign as much flux as possible to the point source. This means that part of the emission-line object will be assigned to the spectrum of the bright point source, and the code will "dig a hole" in the extended object at the position of the bright point source. In both of those two cases, there is hence a certain degeneracy in the final solution as to how much flux should be assigned to the two objects. This degeneracy can be broken in several ways. Typically, one may assume that either the spectrum of the point source is smooth and continuous in the given spectral region, or one may assume that the extended object has a smooth and continuous surface brightness distribution at the position of the bright point source. In the case of two compact and well separated objects no additional assumptions are needed, the solution will be unique.

Acknowledgements

I wish to thank Michael I. Andersen who supplied the routines that perform the optimal fitting of Chebyshev polynomials, and Sandra Savaglio who provided extensive comments on several other functionalities of this data-reduction tool. I am grateful to Keith Horne for many useful comments on an earlier version of this report.

References

Horne, K., 1986, *PASP*, **118**, 609.
Marsh, T.R., 1989, *PASP*, **101**, 1032.
Møller, P., Kjærgaard, P., 1992, *A&A*, **258**, 234.
Møller, P., 1999, in *Astrophysics with the NOT,* Karttunen, H. and Piirola, V. (eds.), University of Turku, p. 80-88.
Møller, P., et al., 2000, *The Messenger*, **99**, 33.
Treu, T., Stiavelli, M., Casertano, S., Møller, P., Bertin, G., 1999, *MNRAS*, **308**, 1037.

SPSF Subtraction II: The Extended Lyα Emission of a Radio Quiet QSO

P. MØLLER[1], S.J. WARREN[2], S.M. FALL[3], P. JAKOBSEN[4], J.U. FYNBO[5]

[1]*ESO, Karl-Schwarzschild-Str. 2, Garching, Germany*
[2]*Blackett Laboratory, Imperial College of Science, Technology and Medicine, London, UK*
[3]*Space Telescope Science Institute, Baltimore, USA*
[4]*Astrophysics Division, Space Science Department of ESA, ESTEC, Noordwijk, The Netherlands*
[5]*Institute of Physics and Astronomy, University of Århus, Århus, Denmark*

1. Introduction

A common trait of high-redshift radio galaxies is their extended Lyα emission. This emission-line nebulosity is generally aligned with the radio axis, and similar emission has been reported around radio-loud QSOs (Schneider et al. 1987, Heckman et al. 1991a,b, Hu et al. 1991). For radio-quiet QSOs, it appears that extended Lyα emission is much less common. Hu et al. (1991) found extended emission-line nebulosity around three of ten radio-loud QSOs, but from none of the seven radio-quiet QSOs in their sample. Steidel et al. (1991) and Bergeron et al. (1999) reported the detection of extended Lyα emission from two radio-weak QSOs, and Bremer at al. (1992) reported the so far only detection of extended Lyα emission from a radio-quiet QSO.

Two methods have been used in all of the work quoted above, narrow-band imaging and direct inspection of two-dimensional (2D) CCD spectra to look for regions of emission extending away from the spectrum of the QSO. A tool that would allow modelling and subtraction of the 2D spectrum of the QSO would not only improve the chances of detecting faint emission lines in the vicinity of the QSO, it could also allow more detailed analysis. Our first report on spectral PSF (SPSF) subtraction (Møller, 2000) described a technique for SPSF subtraction. In this second report, as a simple example of an application of the tool, we report the serendipitous discovery of extended Lyα emission from the host galaxy of the radio quiet $z = 2.559$ QSO Q2206-199.

2. Observations

On August 12, 1999, we obtained four FORS1/Grism 600B long slit spectra, each of 2000 sec integration time, of the radio quiet QSO Q2206-199. Our observing strategy was to obtain several spectra with different offsets along the slit, to minimise the effects of flat-field errors and other systematic effects. This is a crucial point to keep in mind when one is in search of extremely faint features close to, or on top of, a much brighter object.