Deblending RVS Spectra

RVS Extraction

prepared by: C. Allende Prieto
approved by:
reference: GAIA-C6-SP-MSSL-CAP-003
issue: 1
revision: 0
date: December 17, 2008
status:

Abstract

This document describes an algorithm to *deblend* spectra with overlapping windows in the RVS focal plane. Presuming the onboard software will assign as many windows as overlapping objects, and that the shape of the point spread function in the direction across-scan is perfectly known, the flux from each individual source can be recovered by solving a linear system of equations. We perform a number of basic tests for the case of two overlapping sources, finding that the best results are obtained by solving the problem analytically. We also find that the current algorithm for assigning windows to overlapping objects is far from optimal: when the overlapping objects get close enough the signal in the window for the fainter object is dominated by noise, making it impossible to recover how much signal belongs to which source. A scheme with windows of similar sizes is demonstrated to perform much better.

Document History

Issue	Revision	Date	Author	Comment
1	0	2008-11-26	CAP	Draft
1	1	2008-12-12	CAP	Revision after Mark Cropper's comments
1	2	2008-12-17	CAP	Revision after discussions with Mark Cropper

Contents

1 Introduction

Gaia scans the sky continuously, and the dispersed RVS spectra drift through the CCDs, generating charge which is shifted in sync and read-off at the edges. RVS is a slitless spectrograph, and therefore the dispersed spectra may overlap on the focal plane.

Because the angle at which Gaia scans a given field will vary from visit to visit, two spectra which overlap in one visit will not overlap in others. Heavily blended spectra will be harder to disentangle, but can be eliminated from the calculation of average spectra in the Multiple Transit Analysis (MTA; GAIA-C6-SP-MSSL-HEH-012-2). Nevertheless, because each observation is unique (consider, e.g., a radial velocity variable star), there may be few or no transits in which a source's spectrum is fully isolated in crowded fields, and each transit's contribution to the final signal-to-noise ratio is valuable, we must make our best efforts to separate blended spectra in each transit.

Due to the impossibility to telemeter down to Earth the vast data rate generated by the $3+(2+9+2)\times 7+3\times 4=106$ CCDs in the Gaia Focal Plane Assembly (FPA), only small windows around the stellar sources are extracted and sent down to the ground, limiting the analysis of the background contribution as well as our ability to resolve overlapping spectra. Nonetheless, additional information, with higher angular resolution, is provided from the astrometric (AF) and photometric (BP/RP) instruments, which can be exploited to disentangle the contributions from individual overlapping sources.

Extracted spectra that are close enough in the AC (x) direction to overlap will suffer from blending. Spectra from sources for which there is no RVS spectrum extracted are treated under background subtraction in WP620. For the window size selected for RVS (10 pixels), an extracted object cannot contaminate significantly the spectrum of another object, unless their windows physically overlap: the expected FWHM of the AC PSF will vary between 2 and 4 pixels approximately (see §6) and therefore the RVS windows will cover at least 3σ from the center, leaving at most 10^{-3} of the flux outside the window. Thus, deblending will only be considered for RVS spectra with overlapping windows.

2 Impact of blending for RVS

The fraction of overlapping spectra can be estimated from the mean time interval between stars transiting a given pixel in a Gaia CCD. This quantity is available from a recent technical note by B. Holl (GAIA-C3-TN-LU-BH-001-01). In Table 2 of this document, we see the average time between transits (I_{col}) for six different stellar density bins and three magnitude limits. We select the figures for a limiting magnitude of 17.0, which are appropriate for RVS. We also adopt a spread factor of one, and will scale the results according to the RVS AC window size. Fig. 1 shows in black the fraction of the sky (A) versus the stellar density (n) in each bin for this

limiting magnitude.

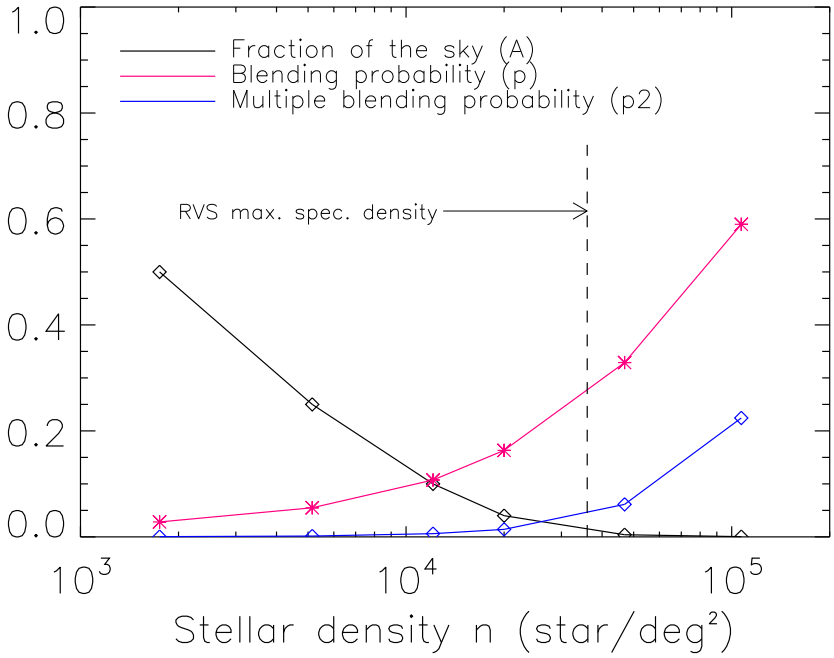

Figure 1: The black line shows the fraction of the sky as a function of stellar density (stars/deg^2) down to a limiting magnitude of 17.0 (from Table 2 of GAIA-C3-TN-LU-BH-001-01). The red line shows the expected fraction of blended RVS spectra. The blue line shows the expected fraction of multiple (> 2) blends. The maximum density that the RVS onboard software has to handle, as given in the specifications, is 36,000 stars/deg^2, marked in the graph with a dashed line.

This problem has been recently discussed (in two-dimensions) by Mignard (GAIA-CU4-TN-OCA-FM-035-1), and we refer the reader to this document for more details. The rate at which stars cross an RVS AC window (10 pixels) can be written

$$\lambda = \frac{10}{I_{col} C_{1\rightarrow 2}} \quad \mathrm{s}^{-1}, \tag{1}$$

where the factor $C_{1\rightarrow 2}$ accounts for the increase in the stellar density from one to two fields of view. The probability of finding an RVS spectrum overlapped with others, i.e. centered within the same 1-second long time interval, is given by

$$p = 1 - e^{-\lambda}, \tag{2}$$

which for small values of λ can be approximated $p \simeq \lambda$. The result corresponds to the red line in Fig. 1. Similarly, we can estimate the fraction of multiple blends (> 2 objects with overlapping windows)

$$p2 = 1 - e^{-\lambda} - \lambda e^{-\lambda}, \tag{3}$$

Table 1: Input parameters and estimated blending probability (p) for representative Gaia fields. I_{col}, $C_{1\to2}$, λ, and p are defined in the text (§2).

Sky fraction (A)	stellar density (n)	I_{col}	$C_{1\to2}$	λ	p
0.50	1.75×10^3	699	0.500	0.029	0.03
0.25	5.15×10^3	237	0.747	0.057	0.06
0.10	1.21×10^4	101	0.873	0.113	0.11
0.04	2.00×10^4	61	0.920	0.178	0.16
0.004	4.69×10^4	26	0.964	0.399	0.33
0.0004	1.07×10^5	11.4	0.984	0.892	0.59

which is also shown with a blue curve in Fig. 1. We also present the results in tabular form (Table 1). Weighting each bin with its stellar density and sky fraction, we find an average star crossing rate of about 0.11 s^{-1}, which implies an average blending rate for RVS of 11 %. The average incidence of multiple blendings (more than two stars) is about 0.5 %, of which less than one tenth will correspond to more than three stars. The blending rate will grow to 27 % in fields with 36,000 stars/deg^2, which is the highest density that the RVS is designed to handle before starting to miss sources. Higher blending rates will still happen at higher densities, but deblending will not always be possible due to the limited availability of windows imposed by the telemetry rate.

It is also possible to scale the estimates made by Marrese & Busso for BP/RP (GAIA-C5-TN-LEI-PM-003-1). Our definition of blending includes any overlap between the RVS spectral windows, which corresponds to the *worst case* considered in GAIA-C5-TN-LEI-PM-003-1. There are mainly three factors that need to be adjusted when going from the photometry to the RVS spectra: the spread in AC, the spread in AL, and the limiting magnitude with the corresponding change in stellar density. The spread in AC for RP/BP is $\sim$ 10–12 pixels, roughly the same as for RVS. The spread in AL is approximately 30 times smaller for RP/BP than for RVS ($30-40$ vs. 1100 pixels), and the stellar density increases by at least a factor 4 at $G<20$ than at $G<17$ (see Table 1 in GAIA-C3-TN-LU-BH-001-01; We estimated $\log_{10}N\propto0.18G_{\rm rvs}$ in GAIA-C6-SP-MSSL-CAP-002, while Brown reports $\log_{10}N\propto(0.16-0.23)G$ from simulations described in GAIA-C5-TN-LEI-AB-013 and GAIA-C5-TN-LAB-AKM-001-1). Taking these factors into account we can easily see that overlapping will be more severe for RVS than for RP/BP, and therefore our estimates are somewhat more optimistic than those by Marrese & Buso, who concluded that *strong crowding* will take place for about 9 % of the BP/RP sources.

3 Basics

A monochromatic point source would not produce an infinitely narrow image on the RVS detectors, but it will instead be spread over a number of pixels with a shape given by the Point Spread Function $PSF(x,y)$, where x corresponds to the across-scan direction (AC) and y to the

along-scan direction (AL). An observed spectrum will be the convolution of the stellar spectrum entering Gaia's primary mirror and the PSF. The PSF is shaped by a number of effects, including diffraction, smearing due to motions and time-delayed integration, pixelization, and aberrations. If this function can be written as the product of two functions, one of each dimension, $PSF(x,y) = P(x)L(y)$, the two-dimensional convolution can be separated into two one-dimensional convolutions. We assume this is case for RVS (see the discussion in GAIA-LL-046): the PSF component in the dispersion direction ($L(y)$ or Line Spread Function) is assumed to be a slow function of x – constant over a source's AC span– which simplifies the problem of deblending overlapping sources from two dimensions to just one.

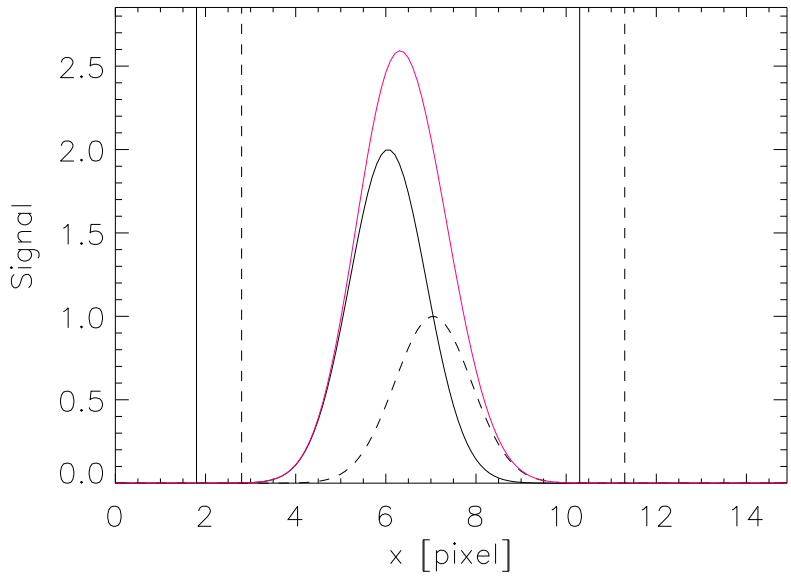

Figure 2: The individual flux from each of two overlapping sources in AC($\equiv x$) is shown in black (solid for the brighter and dashed for the fainter). The added signal is in red. If the signal is integrated over two windows of the same width, each centered on one of the peaks, the PSF is known perfectly, and the windows do not overlap exactly, it is possible to recover the contribution from each source.

At any given pixel RVS CCDs two or more sources may overlap in the across scan direction (AC or x). As explained above, we can disentangle their contributions without considering neighboring pixels in the AL (y) direction – which will correspond to neighboring frequencies in their corresponding spectra, but not necessarily the same frequencies for all sources. Fig. 2 illustrates the situation for two sources: one twice as bright as the other. The coadded signal is shown in red. A window of width $2w$ is assigned to each object (vertical bars), and we only know the total signal inside each window. Assuming the PSF in the AC direction, $P(x)$, is perfectly known and normalized

$$\int_{-\infty}^{\infty} P(x)dx = 1, \tag{4}$$

the signal profile across scan for one star only depends on its central intensity. If for a single object the signal profile in AC is written $I \times P(x)$, the signal in a window of width $2w$ centered on C_1, the central location of the first object, is then

$$S_1 = I_1 \int_{-w}^{w} P(x)dx + I_2 \int_{-w-(C_2-C_1)}^{w-(C_2-C_1)} P(x)dx. \quad (5)$$

In general, the PSF will vary across scan, there can be more than two overlapping objects, and the window size need not be the same for all spectra. For each of N overlapping windows we have

$$S_i = \sum_{j=1}^{N} I_j \int_{C_i-C_j-w_i}^{C_i-C_j+w_i} P(C_j,x)dx. \quad (6)$$

Eq. 6 constitutes a square linear system of equations from which we can obtain the solutions, the integrated signal for each source I_i, numerically.

Note that a similar system holds for the squared uncertainties

$$\sigma^2(S_i) = \sum_{j=1}^{N} \sigma^2(I_j) \left[\int_{C_i-C_j-w_i}^{C_i-C_j+w_i} P(C_j,x)dx \right]^2. \quad (7)$$

4 Window limits

The RVS windows are 10 pixels wide in AC, and therefore for an isolated spectrum $w = 5$ pixels. The electrons that have already been included in one spectrum will not be repeated in other, which implies that $2w$ will be smaller than 10 for blended spectra other than the brightest of the group. In such cases, the windows will still span an integer number of pixels, but they may cover different lengths on each side of the source's center, and of course the total window size need not be an even number of pixels.

To consider the most general case, we introduce two different semi-spans for each object's window: w_i^- and w_i^+, and their equivalent quantities in pixel units $n_i^- = w_i^-/p$ and $n_i^+ = w_i^+/p$, where p is the pixel size in the AC direction in the same units as w_i, assuming no gaps between pixels. **We will always assume that the windows assigned to each object are known exactly, i.e., that we know which exact pixels in the AC direction are coadded in a particular sample of a particular spectrum**.

Given that the window for an isolated object spans an even number of pixels, we presume that the Video Processing Unit (VPU) algorithm will center the window (even if asymmetric) between two pixels. Using a reference system in pixel units and with integer values centered in the pixels, this implies that an object which peaks at C_i will have a window centered at

$\lfloor C_i \rfloor + 1/2$ spanning between $\lfloor C_i \rfloor + 1/2 - n_i^-$ and $\lfloor C_i \rfloor + 1/2 + n_i^+$, where the $floor$ function $\lfloor C_i \rfloor$ gives the largest integer smaller than C_i[1].

Thus, for the RVS windows, and with x in pixel units, Eq. 6 can be rewritten

$$S_i = \sum_{j=1}^{N} I_j \int_{\lfloor C_i \rfloor + \frac{1}{2} - C_j - n_i^-}^{\lfloor C_i \rfloor + \frac{1}{2} - C_j + n_i^+} P(C_j, x) dx. \tag{8}$$

If the first source is the brightest, then $n_1^+ = n_1^- = 5$, and for the rest of the sources n_i^+, n_i^-, or both, will become < 5 so that pixels included in the spectrum of a brighter source are not included again. The top panel of Fig. 3 illustrates the window assigned to an object centered at $x = 7.8$, which covers $2.5 \leq x \leq 12.5$. The bottom panel of Fig. 3 shows a similar object which overlaps with a fainter source at $x = 12.1$, displaced in AL by about 10 pixels, and to which the VPU would assign a window $12.5 \leq x \leq 17.5$.

5 Gaussian PSF

In the case when the PSF is Gaussian and with approximately constant width across the blend

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(x-C)^2}{2\sigma^2}\right], \tag{9}$$

we can rewrite Eq. 8 explicitly as

$$S_i = \frac{1}{2} \sum_j I_j \left[erf\left(\frac{\lfloor C_i \rfloor + 1/2 - C_j + n_i^+)}{\sqrt{2}\sigma}\right) - erf\left(\frac{\lfloor C_i \rfloor + 1/2 - C_j - n_i^-)}{\sqrt{2}\sigma}\right)\right]. \tag{10}$$

This case has more than a purely academic interest. A complex shape PSF can be modeled as the sum of several Gaussians, slightly displaced from each other and with different heights. If m components are used, the PSF is characterized by not just one parameter as in a single Gaussian (σ) but by $3m$ parameters, the widths σ_k, positions δ_k, and heights h_k of all components

$$P(x) = \frac{1}{\sqrt{2\pi}} \sum_{k=1}^{M} \frac{h_k}{\sigma_k} \exp\left[-\frac{(x - C - \delta_k)^2}{2\sigma_k^2}\right], \tag{11}$$

where

$$\sum_{k=1}^{m} h_k = 1.$$

[1]Note that if instead of having integer values of the coordinates at the pixel centers we choose to have them between pixels, setting, for example, the coordinates origin at the lower left corner in Fig. 3 will lead to the same window limits without the summand 1/2.

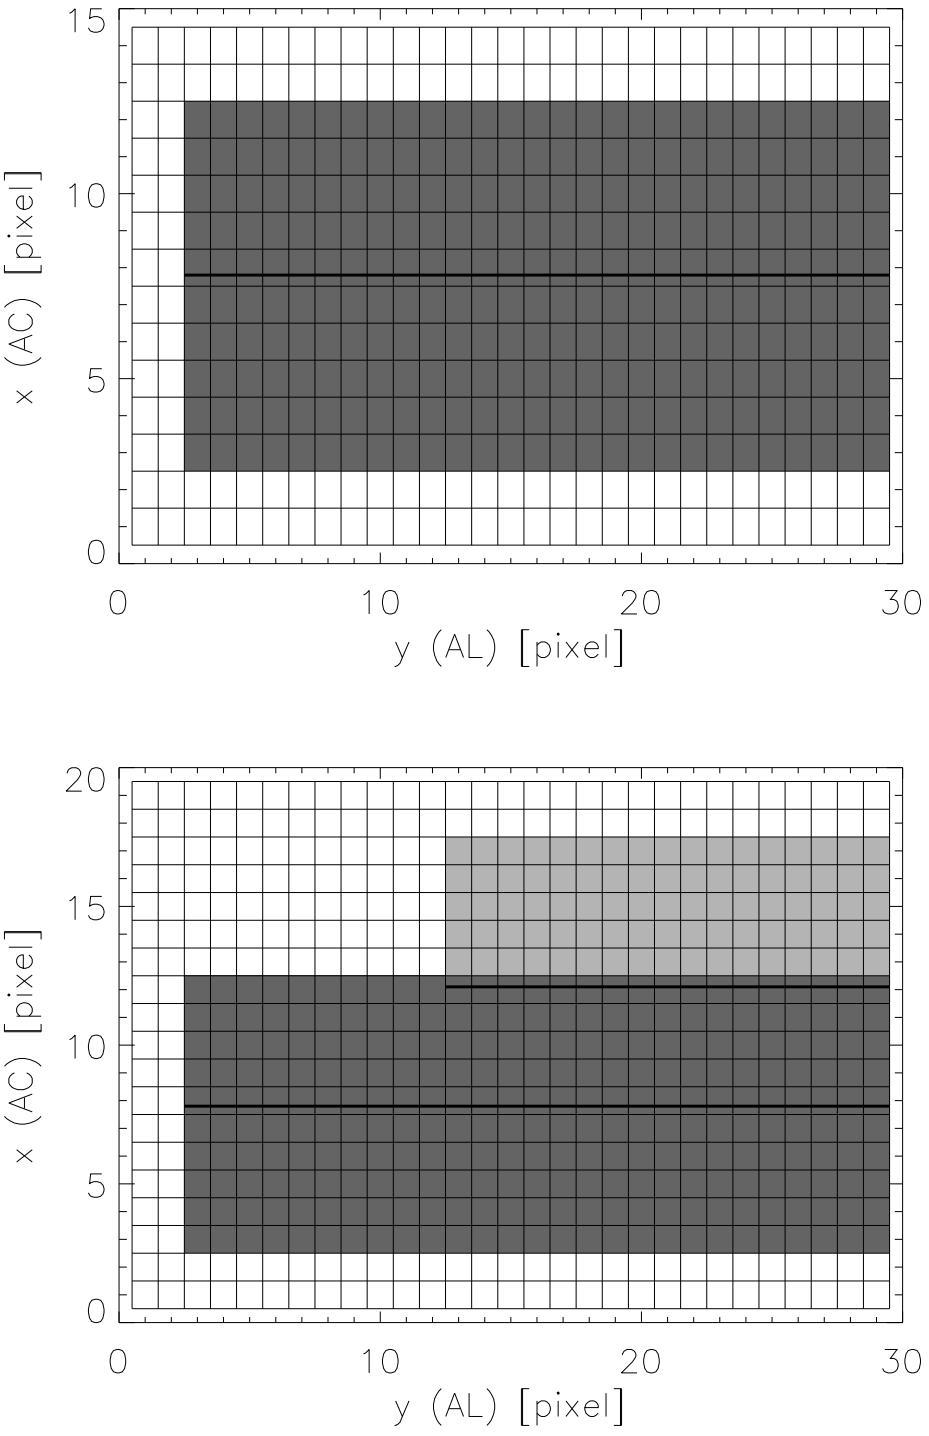

Figure 3: Schematic representation of RVS windows for an isolated spectrum (top panel) and for two overlapping objects (bottom panel).

In this case we have the system

$$S_i = \frac{1}{2}\sum_j I_j \sum_k h_k \quad \times \tag{12}$$
$$\left[erf\left(\frac{\lfloor C_i \rfloor + 1/2 - C_j - \delta_k + n_i^+)}{\sqrt{2}\sigma_k}\right) \quad - \quad erf\left(\frac{\lfloor C_i \rfloor + 1/2 - C_j - \delta_k - n_i^-)}{\sqrt{2}\sigma_k}\right)\right]$$

6 AC motion

The scanning law and other effects will cause the motion of objects in the focal plane to deviate somewhat from the CCD TDI (AL) direction. For a given CCD, a maximum side-to-side drift of about 4 pixels is expected. This can be effectively modeled as a convolution with a box-car function. There are many other effects at play (see, e.g. GAIA-C2-TN-OPM-CB-002-02), but the intrinsic width of the optical PSF (assumed with a FWHM of 2 pixels), and the misalignment between an object's motion and the TDI direction will likely be dominant.

As illustrated in Fig. 4, the convolution of a Gaussian profile with a FWHM of 2 pixels and a box-car with a width of 4 pixels is approximately similar to a Gaussian profile with a FWHM$\simeq$ 3.70 pixels. For simplification, will will consider only two extreme cases in the experiments described below: Gaussian profiles with a FWHM of 2 and 3.7 pixels.

7 Analytical solution

The system of linear equations in Eq. 6 can be written in matrix form as

$$[M] \ \mathbf{I} = \mathbf{S}, \tag{13}$$

where the element

$$M_{ij} = \int_{\lfloor C_i \rfloor + \frac{1}{2} - C_j - n_i^-}^{\lfloor C_i \rfloor + \frac{1}{2} - C_j + n_i^+} P(C_j, x)dx \tag{14}$$

of the square array $[M]$ corresponds to the fraction of the photons from object j that contribute to the signal in the window i. When one or more of the windows contains very little signal, the corresponding matrix elements will be very small, making the numerical solution of the system prone to significant errors. Fortunately, for the case of a small number of overlapping sources N, we can write explicitly the solution.

For the case of $N = 2$, Eq. 13 reads

$$\begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \end{pmatrix} = \begin{pmatrix} S_1 \\ S_2 \end{pmatrix}, \tag{15}$$

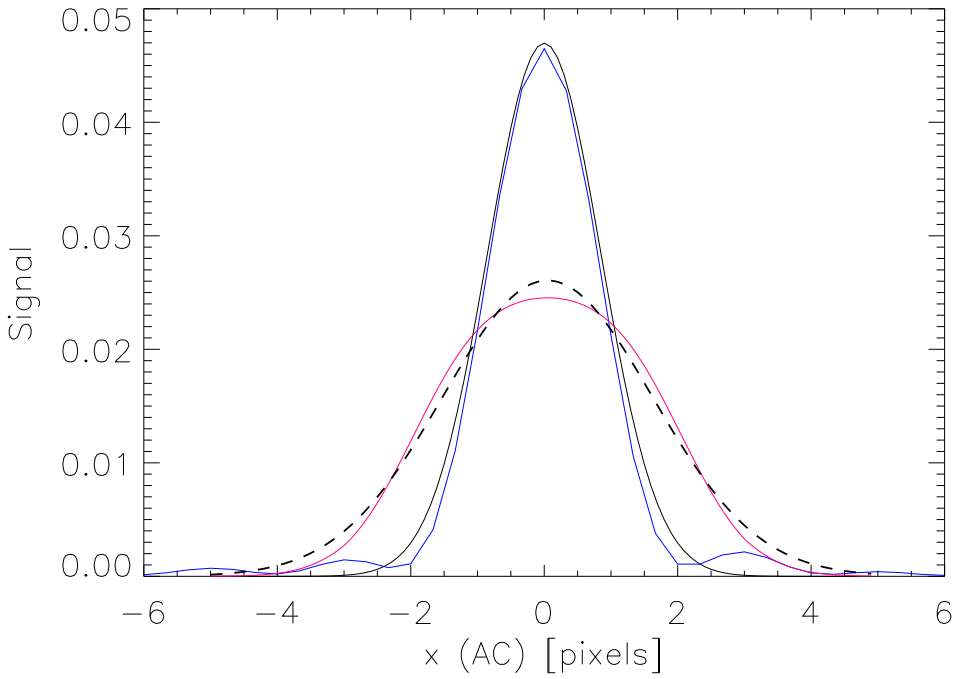

Figure 4: Extreme cases for the AC PSF. The minimum FWHM is 2 pixels, which corresponds to the black solid line. The real profile will differ from this one mainly in the wings: the blue line is a calculated PSF from CU2 for the CCD in the second strip (column) from the left and the second row from the top in the RVS focal plane. After convolving the 2-pixel FWHM Gaussian profile with a box-car accounting for the maximum AC displacement expected (4 pixels), we find the broken black line, which can be roughly approximated by a Gaussian with a FWHM of 3.7 pixels (red line).

and when $\Delta = M_{11}M_{22} - M_{12}M_{21} \neq 0$, the system has a unique solution

$$\begin{aligned} I_1 = &\quad (S_1 M_{22} - S_2 M_{12})/\Delta \\ I_2 = &\quad (S_2 M_{11} - S_1 M_{21})/\Delta. \end{aligned} \tag{16}$$

Similarly, one can painstakingly write the solutions for the case $N = 3$.

8 Testing the algorithm for $N = 2$

In order to test the proposed deblending algorithm we simulate the observation of two objects with overlapping windows in AC. First, we examine an ideal situation without noise, and then we move on to more realistic scenarios. The AC PSF is assumed to be Gaussian and the same for the two objects, and the *free* parameters are the intensity ratio and the distance between their AC central positions. The windows have a full width of 10 pixels for the primary objects, and they are reduced for fainter overlapping objects, as specified for the algorithm onboard (see Section 4.3.5.1 of the VPU specification, GAIA.ASU.SP.PLM.00024).

8.1 No noise in the spectra

For the first tests, we assume a total intensity for the bright object of 200 (in arbitrary units), and an intensity ratio of 2 between the two overlapping objects. We initially adopt an AC FWHM of 2 pixels, place the objects far enough in AC that their windows barely overlap, and then progressively reduce their separation. For each case, we solve the linear system in Eq . 6 using several methods: i) analytical (see Eq. 16), ii) singular value decomposition (SVD), and iii) Gaussian elimination.

Fig. 5 illustrates the signal associated with each source (blue for the bright source, and red for the fainter one) as well as the total added signal (black; only visible when distinct from the red/blue curves). The windows assigned to each object are also marked with vertical lines, and broken lines are used for the bright object, in order to make visible overlapping window limits. It is readily apparent that the signal in the second (faint object) window is quite small for separations smaller than the FWHM (lower-left panel).

Fig. 6 shows the signal levels that we recover for each object after running the deblending algorithm. The numerical solutions were obtained with the algorithms in IDL. SVD gives a poor performance, and while Gaussian elimination is robust down to nearly one FWHM, the best results are obtained with the analytical solution, which performs well down to about $0.4\times$FWHM. The detailed curves change somewhat for different choices of the FWHM and the intensity ratio, but the results are similar: the analytical solution performs best, Gaussian elimination next, and SVD last, with Gaussian elimination solutions degrading at a faster rate than SVD as the sources get closer. Tests using Cramer's rule and LU decomposition performed similar to SVD.

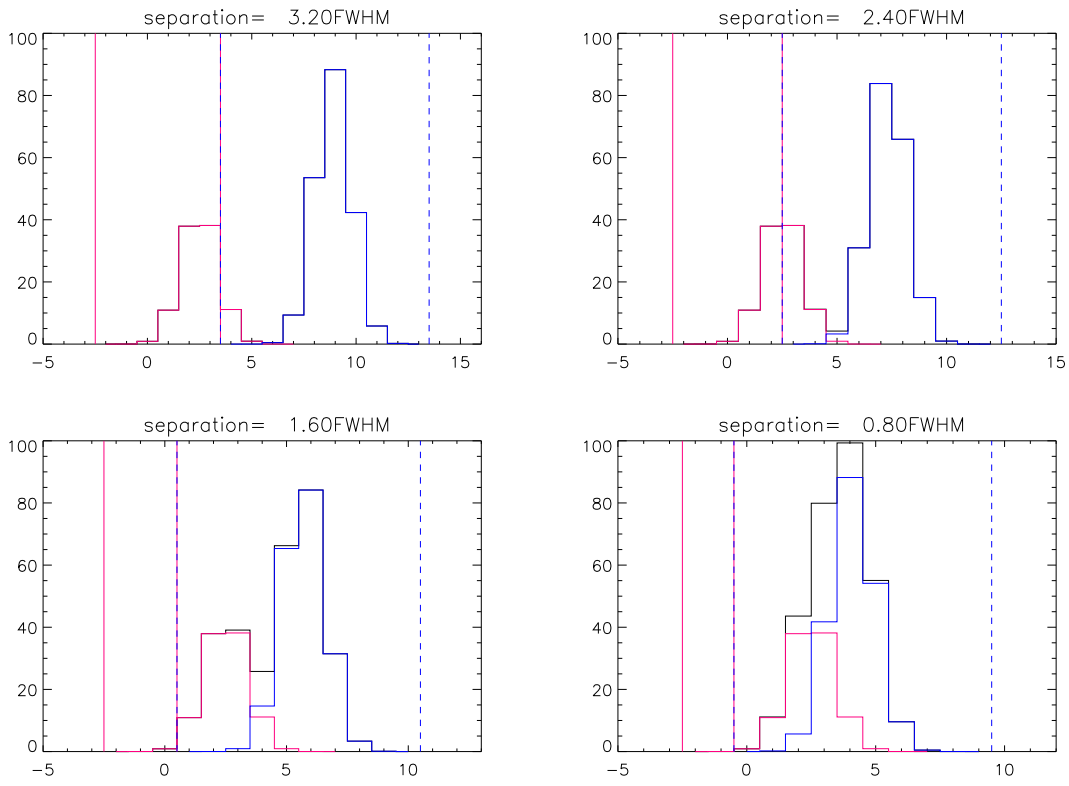

Figure 5: Signal from two objects with overlapping windows in AC. The total intensity, in arbitrary units, is 200 for the bright object (blue) and 100 for the fainter one. The objects share a common PSF, given by a 2 pixel-FWHM Gaussian profile. We test the deblending algorithm for different distances between the central positions of the objects and four cases are illustrated here. The windows assigned to each object are also shown with vertical bars; the brightest object gets a full 10-pixel wide window.

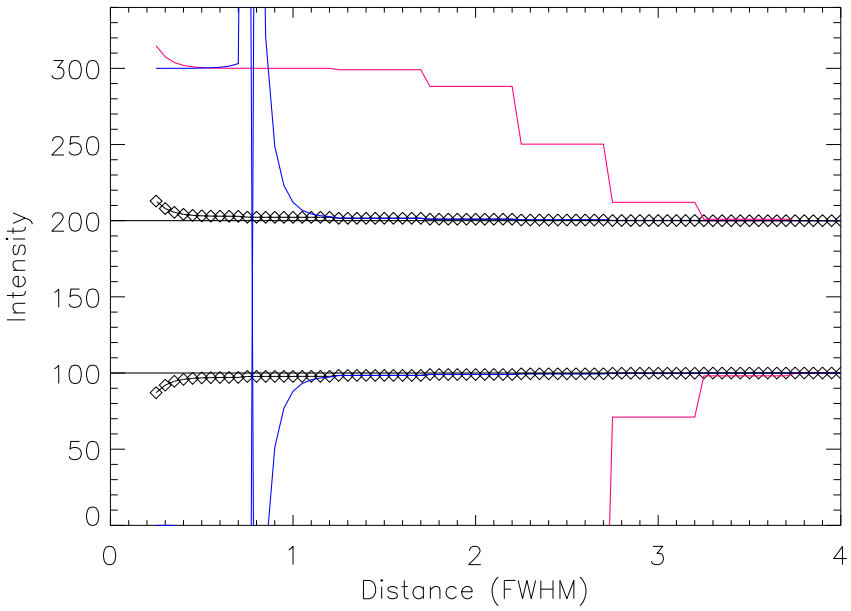

Figure 6: This figure shows the original intensities for two blended objects (black) and the values returned by the recovery algorithm as a function of source separation (across scan). The results for the analytical solution are shown with open symbols. The blue and red lines correspond to Gaussian elimination and SVD, respectively, which exhibit numerical problems for small source separations.

8.2 Accounting for noise

A more realistic assessment requires adding noise to the data. As described in the Appendix A, the most important contributors are shot (Poisson) noise and, for stars fainter than $G_{\rm rvs} \sim 12$ mag, readout noise. For simplification, we will only use one value of the readout noise, the one that applies in the regime where readout noise is important: 4.31 e^- per pixel.

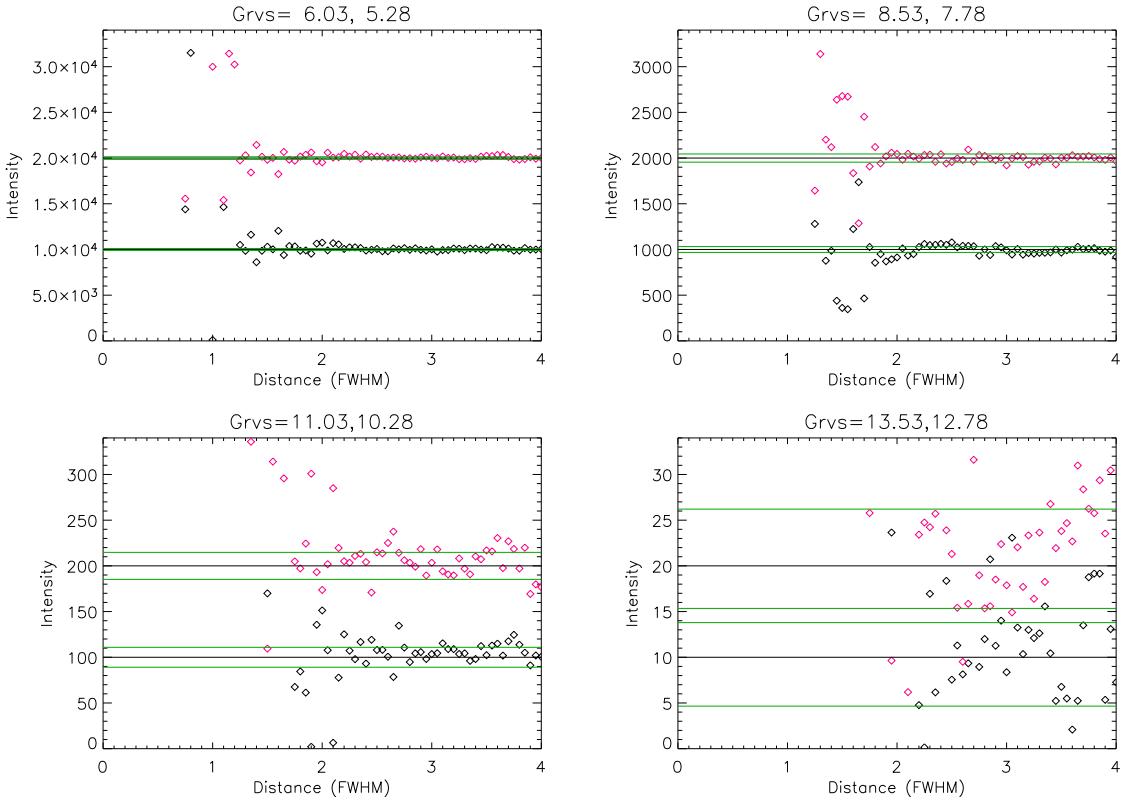

Figure 7: Original intensities for two blended objects (black) and the values returned by the recovery algorithm as a function of source separation (across scan). The results for the analytical solution are shown with open symbols: black for the bright source and red for the faint one. A readout-noise of 4.31 e^- has been included. The green lines mark the $\pm 1\sigma$ error expectation for each of the two sources.

Fig. 7 includes four particular simulations for the same signal ratio used in the noiseless test, and signal levels of the brightest star of 2×10^4, 2×10^3, 2×10^2, and 20 e^-. Note that in this case the signal units are not arbitrary, but they must be photo-electrons in order to match the units used for the readout noise. The results of the deblending procedure seem robust down to a separation between the sources which decreases as the signal-to-noise of the spectrum increases: the procedure breaks down for separations of about 1.2 FWHM, 1.7 FWHM, and 2.1 FWHM for a bright source with $G_{\rm rvs}$ of 6.0, 8.5, and 11.0, respectively. It is hard to draw conclusions for fainter sources, and more statistically robust tests are necessary.

9 Changing windows

Upon inspection of Fig. 5, it becomes clear that the window selection algorithm chosen for RVS presents a potential obstacle for our deblending algorithm to work optimally. With a full window assigned to one object, the secondary window ends up with essentially no signal, leading to nearly null matrix elements in Eq. 13. The problem becomes even more acute for blends involving more than two spectra, and there is a risk for the window of a third of fourth overlapping object to contain no data at all; in such case the system in Eq. 6 would be undetermined and our algorithm useless.

In the presence of noise, windows with little signal will still be affected by readout noise, leading to a signal-to-noise ratio that can be $<< 1$. An alternative window scheme that would divide the overlapping area among the sources in a more balanced fashion could improve the numerical stability of the solution and thus the algorithm's performance.

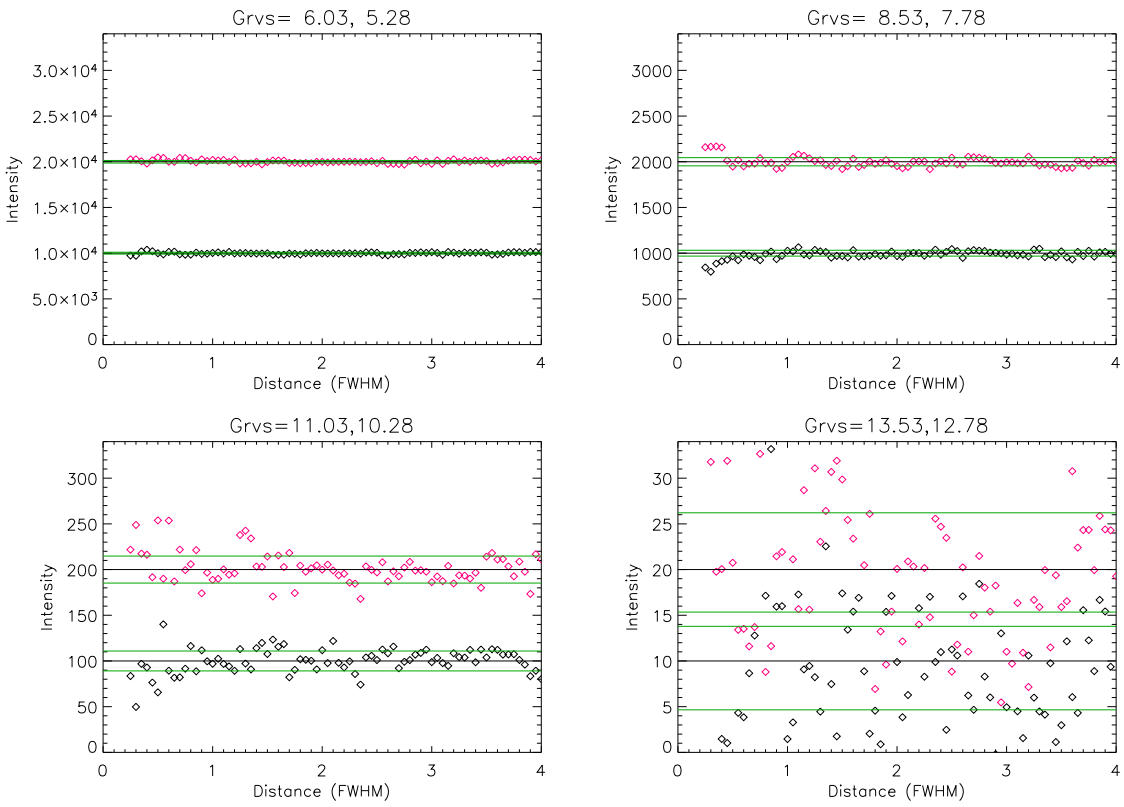

Figure 8: Same as Fig. 7, but with an updated window scheme that divides the overlapping area in AC more evenly.

We tested again for the case $N = 2$ making the window sizes as similar as possible for the two objects, i.e. identical to within 1 pixel. Fig. 8 is the equivalent of Fig. 7 with the updated windows, and reveals a very significant improvement in performance. Similar results are obtained

for different image ratios, as illustrated in Fig. 9, which shows the ratio of the errors in the recovered fluxes for the current and the proposed windows for four pairs of overlapping stars, two with identical magnitudes of 8.5 or 13.5, and two others where the fainter object of the blend is 11 times (2.6 mag) brighter. The errors were derived as the average error for 10 monte-carlo runs. The conclusion is clear: a very significant improvement is obtained by changing the window assignment scheme for blending sources located at less than about 4 pixels (2× FWHM).

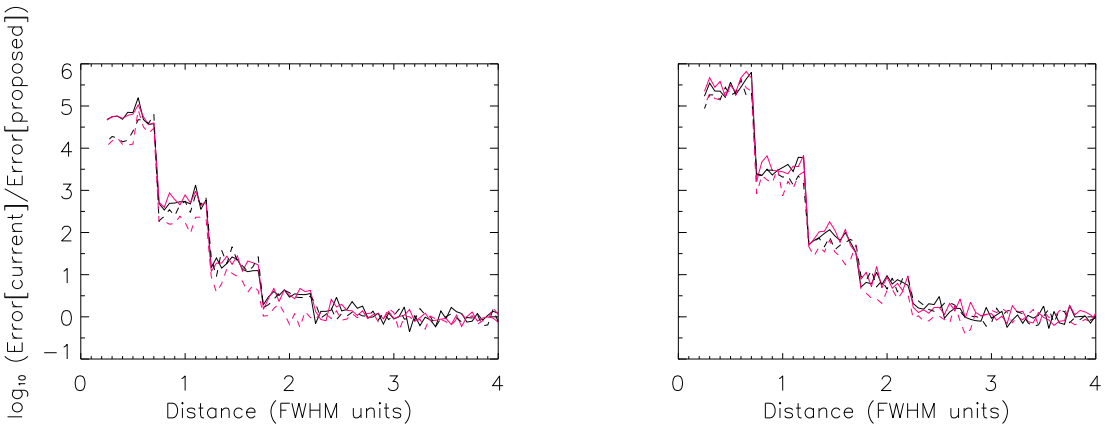

Figure 9: Ratio of the errors in the recovered fluxes for the current and the proposed window assignment schemes. The left-hand panel considers the case of a blend of two identical stars with $G_{\rm rvs} = 8.53$ (solid black and red lines), and another where the fainter source is $G_{\rm rvs} = 8.53$ (black dashed line) and the bright one is 2.6 magnitudes brighter (red dashed line). The right-hand panel considers the case of a blend of two identical stars with $G_{\rm rvs} = 13.5$ (solid black and red lines), and another where the fainter source is $G_{\rm rvs} = 13.5$ (black dashed line) and the bright one is 2.6 magnitudes brighter (red dashed line). Although the separation between is changed in steps of a tenth of pixel, the errors for the current scheme increase significantly every time that the window of the fainter source is reduced by one more pixel.

A key question to answer is how many sources would be lost if the current window assignment scheme is not changed. Repeating the arguments used in §1 to estimate the fraction of blended spectra for a narrower window of 4 pixels, we find that roughly 4 % of the RVS transits, about 160 million spectra, will be too blended to be disentangled with the current window scheme. It the light of these results, we highly recommend that the Gaia RVS windowing algorithm is modified.

Our tests include noise in the signal from the blended sources, but assume that their positions are known exactly. We have performed simulations considering the effect of random errors in the location of the sources at the level of 1/10 of a pixel, or about 20 mas. These experiments are described in more detail in the Appendix B, but the main result is that such errors only have an impact on bright sources, where the low intrinsic noise allows them to surface. For a difference of 2 magnitudes between the blended stars, the fainter source in the pair needs to be brighter

than $G_{rvs} \sim 11$ for the uncertainties in the sources' locations to have a measurable effect in the recovered magnitudes.

10 Summary

We present an algorithm to separate the spectra of two or more sources with overlapping RVS windows. This is a purely algebraic algorithm, which does not involve any assumption about the nature of the overlapping objects beyond that they are point sources. It works at a pixel level, separating the contributions to each sample in AL in sequence. Given that the fluxes in neighboring AL samples are correlated due to an AL PSF wider than 1 pixel, it could be helpful to smooth the resulting deblended spectra accordingly.

We show that with the current VPU window assignment scheme for overlapping sources, which gives priority to the brightest sources, the success of the algorithm is limited. However, the performance is greatly enhanced if the windows span is more evenly distributed among the sources.

The uncertainties in the signal are carefully considered, including readout noise, shot noise, and the limited accuracy in the exact location of the sources on the focal plane and in particular on the RVS CCDs. We examine how the uncertainties in the decomposed signals vary as a function of brightness and brightness ratio for the case of two overlapping spectra.

We estimate that about 10 % of the RVS spectra – 400 million – will be blended to some degree. If the window assignment scheme proposed here is adopted, we expect to successfully disentangle most of these sources, but some 160 million RVS spectra will be closely overlapped and cannot be deblended with the scheme currently in place.

We examine, in an appendix, the potential impact of changing the width of the extraction windows for isolated objects, concluding that the current choice of 10 pixels is appropriate. We also assess the impact of uncertainties in the estimated positions of the sources on the performance of the deblending algorithm, finding that they can cause a slight, but measurable, degradation for bright sources.

Acknowledgments: I thank Paola Marrese for useful discussions and Paola Sartoretti for providing simulated PSFs and assistance interpreting them.

APPENDIX A: Signal-to-noise ratio for the RVS spectra

In this appendix we will estimate the signal-to-noise ratio for RVS spectra of isolated objects as a function of window size. Larger windows will render more signal, but also a higher background and more frequent overlapping among spectra.

The nominal window width for isolated objects is 10 pixels ($n^- = n^+ = 5$), and the same width applies to the brightest objects in a blend. Fainter objects in a blend will end up with a reduced width.

For an in a window spanning an even number of pixels $2n$ ($n = n_- = n^+$) with a Gaussian PSF we see from Eq. 10 that the signal in the window is

$$S = I \times erf\left(\frac{\lfloor C \rfloor + 1/2 - C + n}{\sqrt{2}\sigma}\right). \tag{17}$$

We will ignore for now the errors induced by the uncertainties in the estimated central location of the source; these will become negligible for large windows. Then, we can arbitrarily select any value for C, and a convenient choice is to place the source exactly at the center of the window, $C = n + 1/2$, which leads to

$$S = I \times erf\left(\frac{n}{\sqrt{2}\sigma}\right). \tag{18}$$

The average background sky signal per pixel is about 4 times smaller than the GAIA-C6-SP-MSSL-CAP-003 readout noise, and the dark current several orders of magnitude smaller. The uncertainty in the signal is

$$\sigma(S) \simeq \sqrt{S + \sigma_0^2 + 2n \times (sky + dark)}, \tag{19}$$

where σ_0 is the readout noise (4.317 e^- for LR $G_{rvs} > 10$, 4.600 e^- for HR ($7 < G_{rvs} < 10$), and $\sqrt{10} \times 4.600 = 14.5465$ e^- for objects with $G_{rvs} < 7$), $sky = 60.7667/q$ e^- (with q being the number of samples, 1260 for HR and 420 for LR at the time of this writing), and $dark = 0.001325$ e^-. The relative error in the determined brightness of a star is then identical to the relative error of the signal in the integration window $\sigma(S)/S$. We will use a Gaussian PSFs of FWHM 2 and 3.7 pixels (see §6). Noting that for a Gaussian $\sigma = \text{FWHM}\sqrt{-1/8/\ln(1/2)}$ and that the G_{RVS} magnitude is related to the signal per sample in a extracted spectrum by $G_{RVS} = -2.5\log(I \times q) + 22.5866$ we can calculate the relative error as a function of window width and magnitude. Table 1 provides a handy reference to estimate the signal and the noise for RVS observations as a function of G_{rvs}. The resulting average signal-to-noise ratio per pixel is displayed in Fig. 10. We refer here to the final spectrum: once it has been collapsed in the AC direction either by on-the-chip binning, or by adding all the signal in the window using software.

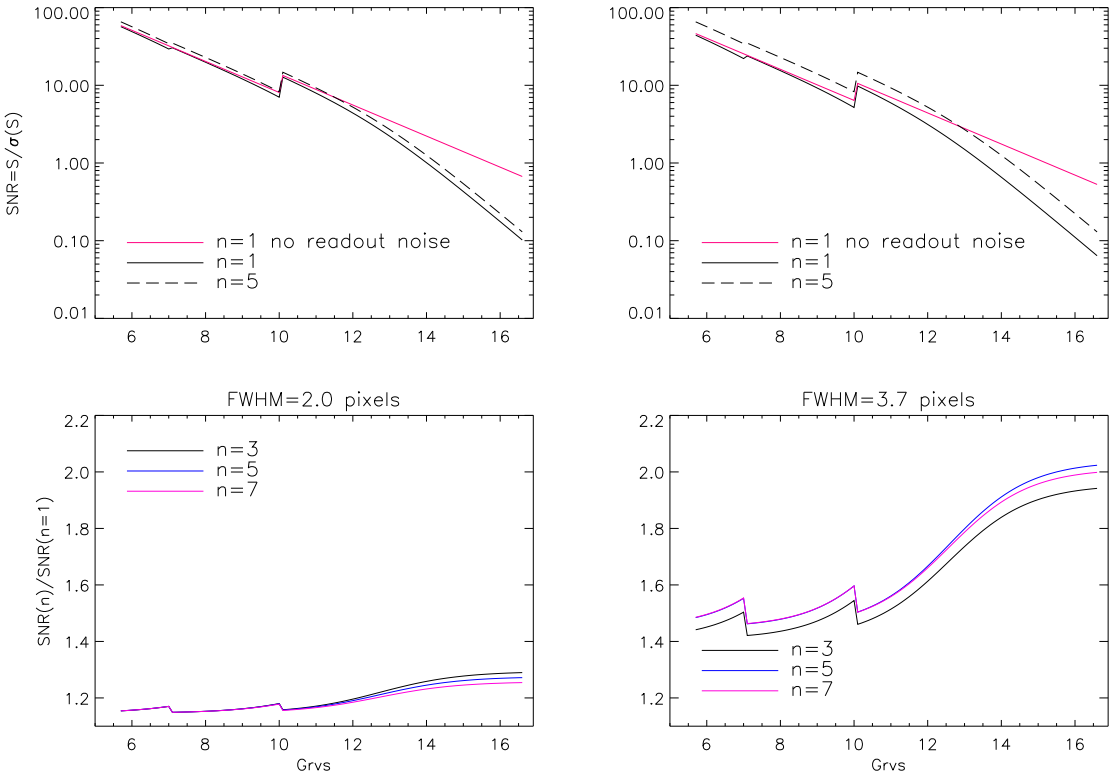

Figure 10: Schematic representation of RVS windows for an isolated spectrum (top panel) and for two overlapping objects (bottom panel).

Table 2: S/N per pixel as a function of brightness for 10-pixel wide windows ($n = 5$)

G_{rvs}	S/N per pixel	G_{rvs}	S/N per pixel
5.7	65.6	10.7	10.9
6.7	40.1	11.7	6.3
7.7	26.3	12.7	3.3
8.7	16.3	13.7	1.6
9.7	9.8	14.7	0.7
		15.7	0.3

The knee at $G_{\rm rvs} = 10$ mag corresponds to the change between the HR and LR modes, and a less-obvious knee is also present at $G_{\rm rvs} = 7$, when the readout noise increases tenfold. Note that for Grvs > 10, the pixel size in the spectral direction is increased by a factor 3 by binning on-the-chip. The readout noise (per pixel) is independent of the pixel size and of the window size in AC ($2n$). The red line corresponds to the case $n = 1$ with only shot (Poisson) noise, and shows that readout noise becomes a dominant contribution for stars fainter than ~ 12.

Changing the window size from 1 pixel to 5 leads to a modest improvement in the signal-to-noise ratio of the extracted spectrum for a PSF with a FWHM of 2 pixels (top left-hand panel in Fig. 10), and this effect is more significant for a PSF with a FWHM of 3.7 pixels, when relatively more flux is lost for a small $n = 1$ window. In more detail, the lower panels reveal that increasing the window size from $n = 1$ brings in slightly more signal, especially for a FWHM of 3.7 pixels, and no significant additional background for bright sources. It also shows how the sky contribution degrades slightly the spectra of the faint objects when $n > 5$, but this is a second-order effect.

For our Gaussian PSF with FWHM of 2 pixels, 0.04% of the flux escapes the $n = 3$ window, while only $\sim 10^{-9}$ escapes for $n = 5$, and that further reduces to $\sim 10^{-16}$ for $n = 7$. For this narrow PSF, our calculations indicate a marginally better performance on the faint end for small windows $n = 2 - 3$. However, we do not consider errors in the location of the source, which will degrade the signal-to-noise ratio for the smallest windows. For $n = 2$, an error in C of 1/2 pixel will result in underestimated fluxes by ~ 3 %, but this figure will be reduced to 0.1 % for $n = 3$. When the PSF FWHM is 3.7 pixels, about 6 % of the flux is lost with $n = 3$, about 10^{-3} when $n = 5$ and about 10^{-5} when $n = 7$.

All things considered, $n = 5$ (i.e. a total window full width of 10 pixels) seems a good choice for isolated objects.

APPENDIX B: Uncertain window locations

Unavoidably, the location of an object on the RVS CCD will be only known to some degree of approximation. For the brightest objects ($G_{\rm rvs} < 7$), this location can be measured using the RVS data themselves, but for the vast majority of the sources, this will be estimated based on a calibrated mapping between the astrometric fields and the RVS field. The RVS pixels are 30 μm long in AC, or about 0.2 arcsec. We can conservatively expect the location of a source during a transit to be known to about 10–20 mas.

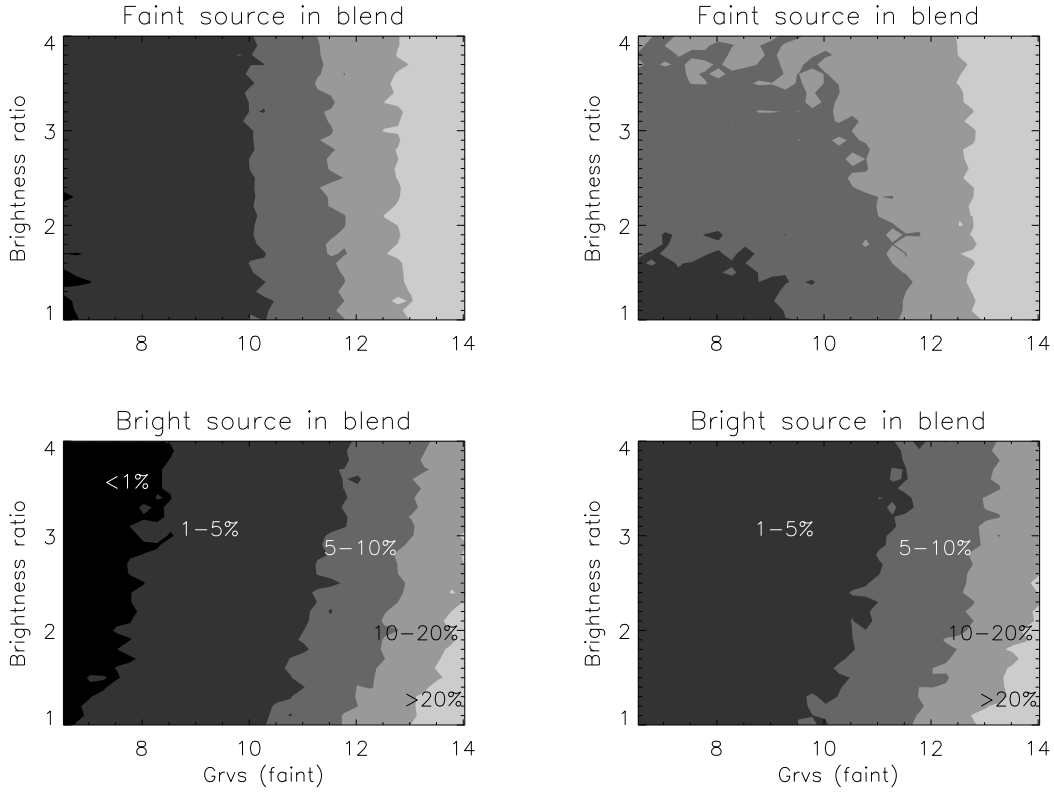

Figure 11: Map of relative uncertainties ($|I_{\rm recovered} - I_{\rm true}|/I_{\rm true}$) for the deblended fluxes using the proposed (even) windows. The two sources are placed at a separation of 1-FWHM, and the FWHM is 2 pixels. The left-hand panels correspond to the case when the uncertainties in the sources location are neglected, and the right-hand panels include uncertainties in the positions of the sources of 1/10 of a pixel.

We experiment introducing Gaussian errors with $\sigma = 0.1$ pixels in the expected central location of the sources before applying deblending. We performed a comprehensive set of tests for a separation of 1–FWHM between the two sources, and different brightnesses and brightness ratios. For each configuration, we simulated 100 transits, and computed mean errors. We did this exercise both for a PSF with a FWHM of 2 and 3.7 pixels (see the discussion in §6). The relative errors (ratio between mean difference between the true and the recovered signal and the

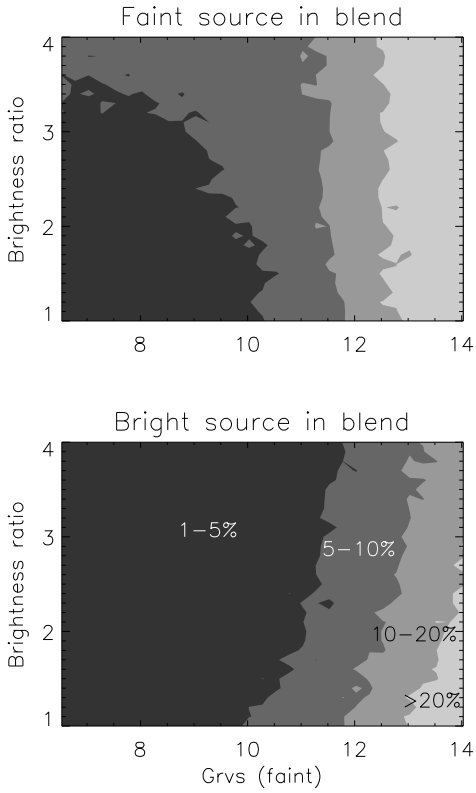

Figure 12: Similar to the right-hand panels in Fig. 11, but with a FWHM of 3.7 pixel. The separation between the two blended sources is again 1–FWHM, but note that the actual number of pixels that this translates to has changed. These results include uncertainties of 1/10 of a pixel in the positions of the sources.

true signal) for a FWHM of 2 pixels are displayed as contour maps in Fig. 11. The left-hand panels illustrate the case when errors in the positions are neglected and the right-hand panels include uncertainties in the sources' locations.

Within the range of parameters under consideration, the errors with which the flux of each object are recovered are mainly dependent on its own brightness. This is shown by the nearly vertical stripes in the top-left panel of Fig. 11. As we varied the magnitude of the faint source and the brightness (intensity) ratio, the magnitude of the bright source is $G_2 = G_1 - 2.5\log(I_2/I_1)$, which causes the tilted contours in the lower-left panel. As one might intuitively expect, absolute errors are similar for the two sources, and therefore the brighter one ends up with smaller relative errors than the fainter.

The 1/10 pixel (1σ) uncertainty in the source's position degrades the relative performance of the deblending algorithm for bright sources, but it has a modest impact for fainter ones, for which shot noise becomes dominant. The contours in Fig. 12 also consider uncertainties in the sources' location, but they are for the case of a FWHM of 3.7 pixels and, again, a source separation of 1 FWHM. As expected, the impact of the position errors is largest when the PSF is sharper, as readily apparent comparing the contours for the fainter source in the blend (lower panels).

Perhaps of more interest than the relative errors, we can examine the ratio between the measured errors and those expected for an isolated object, namely the contributions from shot and readout noise. This ratio is mapped, again for a separation between sources of 1–FWHM, in Fig. 13, where it is apparent that the deblending procedure introduces very little additional errors for stars fainter than $G_{\rm rvs} \sim 11$, unless they are blended with stars more than two magnitudes brighter. The errors increase twofold, threefold, and fivefold for blended stars with $G_{\rm rvs} \sim 10$, 8 and 5 magnitudes.

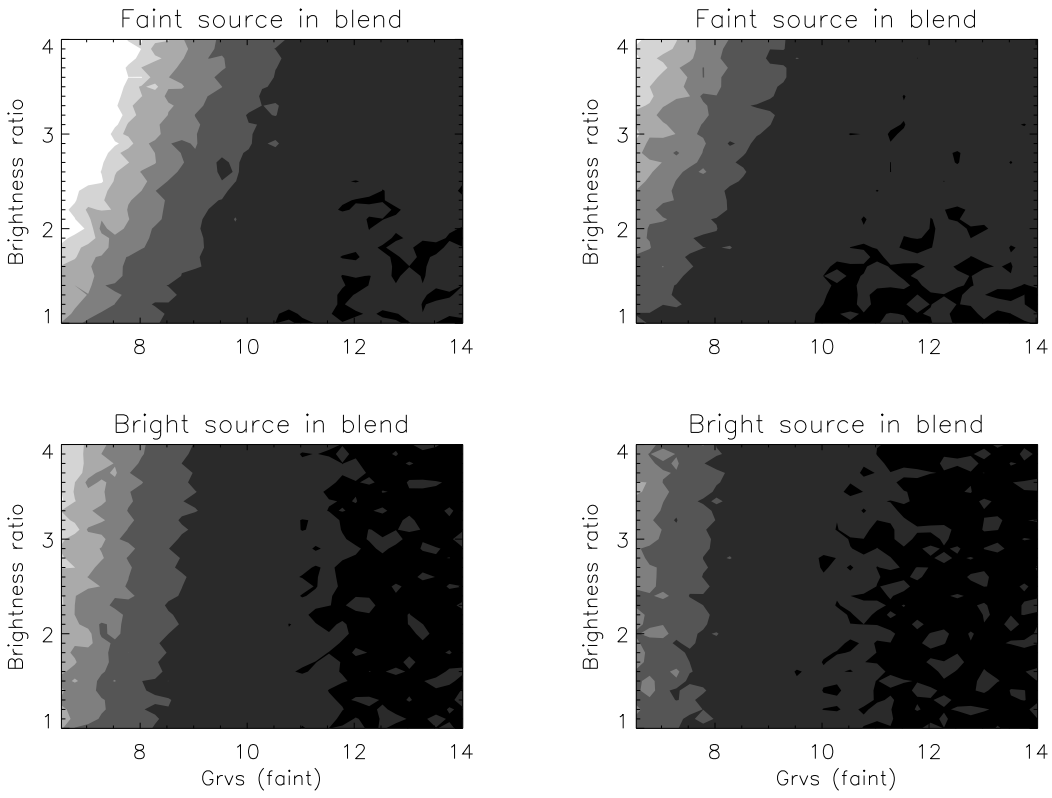

Figure 13: Ratio of the measured errors ($|I_{\text{recovered}} - I_{\text{true}}|$) to those expected for an isolated object (here approximated as $\sqrt{I + \sigma_0^2}$, where σ_0 represents the readout noise, taken to be 4.3 e^-). The two sources are placed at a separation of 1-FWHM, and the FWHM is 2 pixels (left-hand panels) or 3.7 pixels (right-hand panels). Uncertainties in the positions of the sources of 1/10 of a pixel are included. The grey scale indicates, from black to white, ratios of 0–1, 1–2, 2–3, 3–4, 4–5, 5–6, and > 6.