Publication Year	2023
Acceptance in OA	2025-02-28T13:25:49Z
Title	Gaia Data Release 3. Processing and validation of BP/RP low-resolution spectral data
Authors	De Angeli, F., Weiler, M., MONTEGRIFFO, Paolo, Evans, D. W., Riello, M., Andrae, R., Carrasco, J. M., Busso, G., Burgess, P. W., Cacciari, C., Davidson, M., Harrison, D. L., Hodgkin, S. T., Jordi, C., Osborne, P. J., PANCINO, Elena, ALTAVILLA, Giuseppe, Barstow, M. A., Bailer-Jones, C. A. L., BELLAZZINI, Michele, Brown, A. G. A., CASTELLANI, Marco, Cowell, S., Delchambre, L., DE LUISE, Fiore, Diener, C., Fabricius, C., Fouesneau, M., Frémat, Y., Gilmore, G., Giuffrida, G., Hambly, N. C., Hidalgo, S., Holland, G., Kostrzewa-Rutkowska, Z., van Leeuwen, F., Lobel, A., MARINONI, Silvia, Miller, N., Pagani, C., Palaversa, L., PIERSIMONI, Anna Marina, Ragaini, S., RAINER, Monica, Richards, P. J., Rixon, G. T., Ruz-Mieres, D., SANNA, Nicoletta, Sarro, L. M., Rowell, N., SORDO, Rosanna, Walton, N. A., Yoldas, A.
Publisher's version (DOI)	10.1051/0004-6361/202243680
Handle	http://hdl.handle.net/20.500.12386/36331
Journal	ASTRONOMY & ASTROPHYSICS
Volume	674

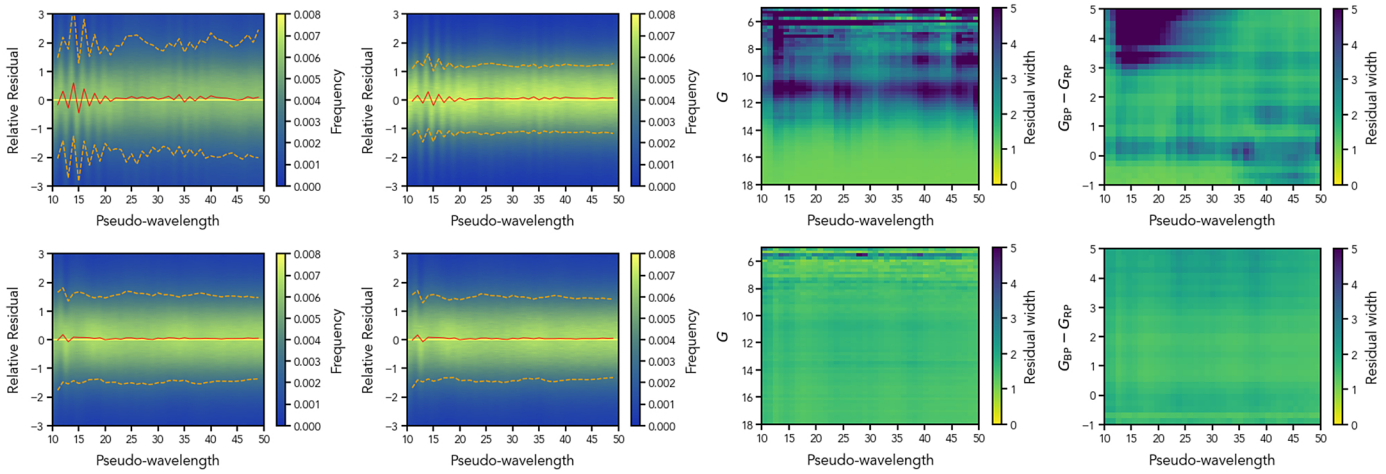

Fig. 8. Relative residual distribution for a subset of the calibrators covering the G-band magnitude range [5, 18]. The *first row* of plots shows the BP results, while the *bottom row* shows RP. In each row, the *first plot* shows the distribution of relative residuals vs. AL coordinate in the range [10, 50] where most of the flux is observed. In the *second plot*, the same distribution is shown including only data from sources in the magnitude range [13, 17]. In these first two plots, the 2D histogram is normalised to the number of measurements in each column and the relative number of sources is shown by the colour bar. The red line shows the median value, while the orange dashed lines show the 15.865 and 84.134 percentiles. The *following two plots* show the robust width of the distribution of relative residuals defined as the difference between the 84.134 and 15.865 percentiles divided by two vs. G-band magnitude and $G_{BP} - G_{RP}$ colour and AL coordinate for the entire magnitude range covered by this subset.

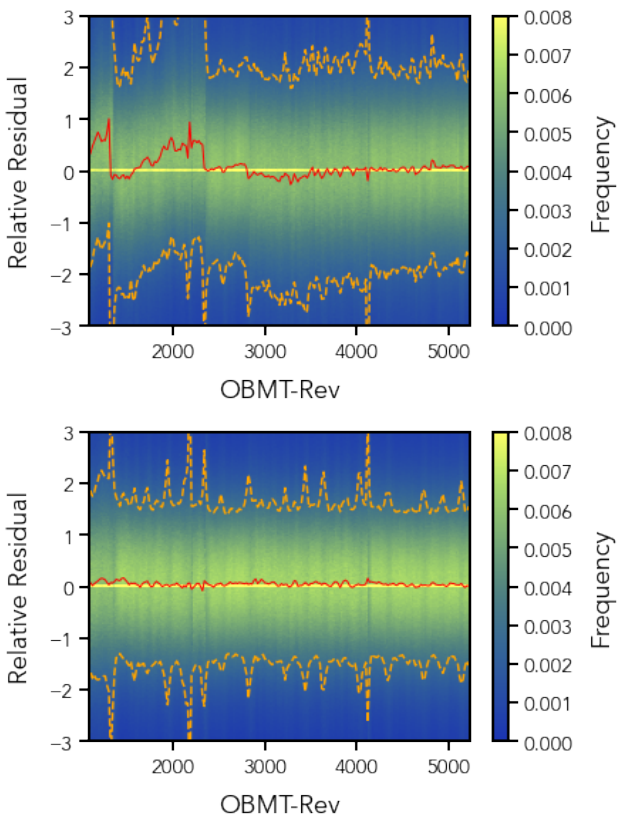

Fig. 9. Relative residual distribution for a subset of the calibrators covering the magnitude range [5, 18]. The *top panel* shows the BP residuals, while the *bottom* one shows the RP residuals. Only samples with AL coordinate in the range [10, 50] are included in this plot. The 2D histogram is normalised to the number of measurements in each column.

in the BP data due to the contamination in the early phases of the mission. Considering the long period of time with minimal contamination available, we decided to ignore all BP data collected before the decontamination event that took place shortly before OBMT-Rev 2340 when generating the final catalogue of mean spectra.

3.3.3. Convergence

Convergence of the iterative process was monitored by looking at different parameters: the median standard deviation of the solutions, the overall absolute change in parameters, and the average χ^2 of the residuals for a subset of the calibrators were all considered.

Each least squares solution for a calibration unit is assigned a standard deviation. The normalised median standard deviation of all least squares solutions over the OBMT-Rev range [3000, 4000] grouped by gate and window class combination versus iteration number is shown in Fig. 10. Each panel shows a combination of photometer (BP/RP), gate, and window class as indicated in the label. There are some configurations where the evolution of the median standard deviation is not monotonically decreasing, particularly in the first few iterations. If the calibration of each configuration were solved independently, one would expect the corresponding standard deviation to decrease in subsequent iterations. However, in the iterative process described in Sect. 3.3, all calibration units are linked together by the common catalogue of reference spectra that is updated at each source update. For this reason, the fact that the standard deviation does not decrease for all configurations is not a sign of a lack of convergence over all.

Overall convergence is assessed looking at the absolute relative change in the values of model parameters A_k between two subsequent iterations. Figure 11 shows how these evolve during the iterations for different nominal combinations of gate and window class in BP (top panels) and RP (bottom panels). Given the large number of parameters, only results for ROW4 are shown here, with other rows showing similar trends. The curves in each panel show the median value over the central part of the spectrum in different colours depending on the index j. The overall absolute relative change in calibration parameters is at or below 1% well before iteration 50 for the central part of the

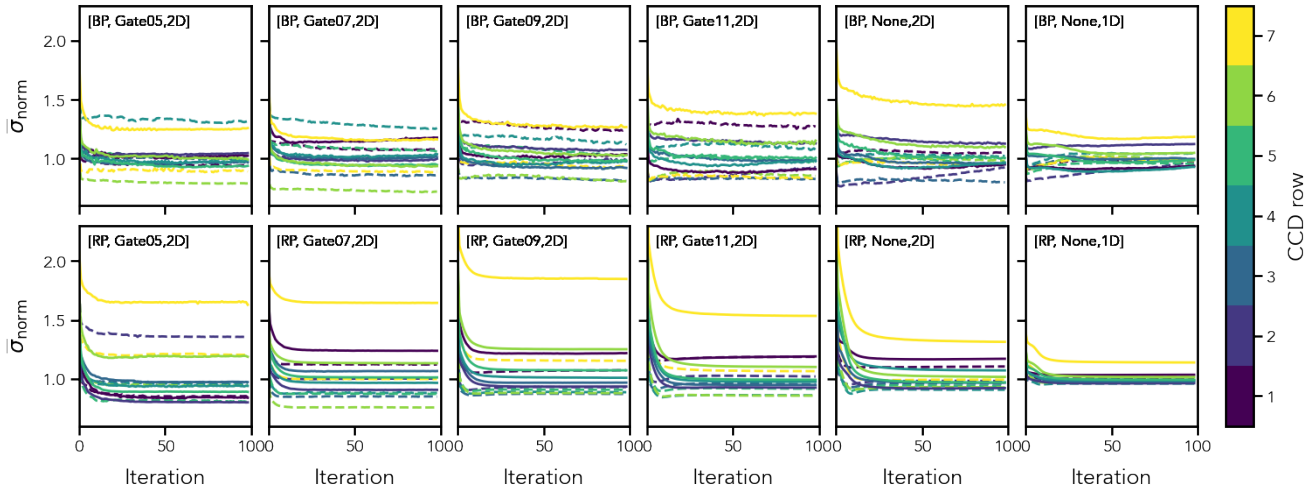

Fig. 10. Median standard deviation for all solutions covering the OBMT-Rev range [3000, 4000], normalised to the median standard deviation of all calibrations obtained for the same photometer (BP/RP), gate, and window class at iteration 50 (by that iteration the system seems to have become quite stable). *Top panels*: BP solutions, one panel per nominal combination of gate and window class. *Bottom panels*: RP solutions. Different colours indicate different CCD rows and solid and dashed lines are used for the preceding and following FoV, respectively.

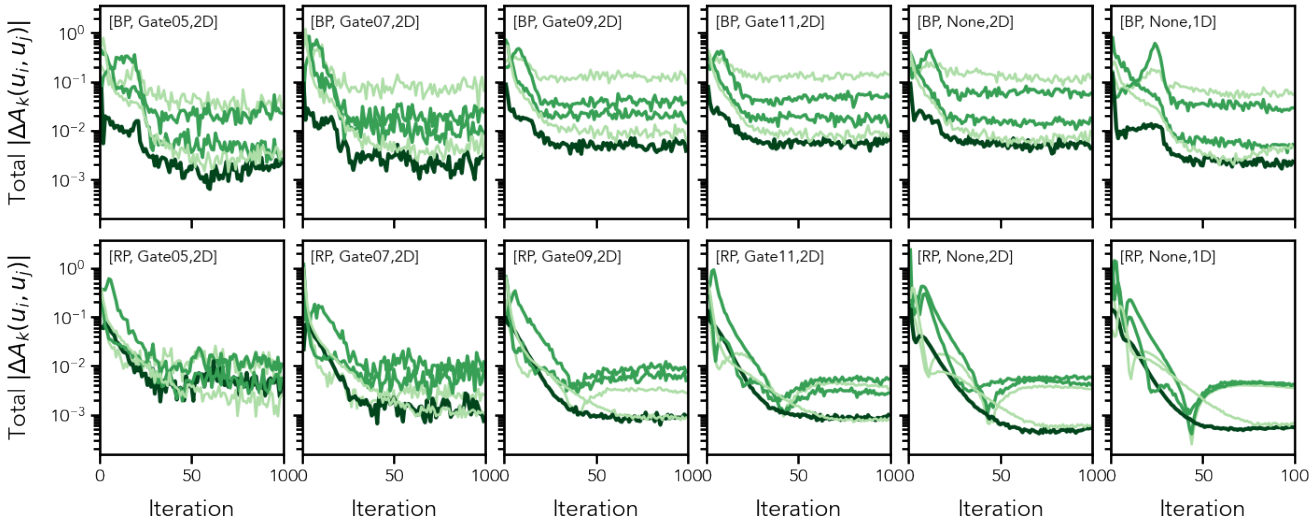

Fig. 11. Absolute relative change in the values of model parameters between two subsequent iterations for all solutions covering the OBMT-Rev range [3000, 4000] in a logarithmic scale. The relative change for each parameter is computed as the absolute difference between the values at two subsequent iterations, normalised by the value of the same parameter at the preceding iteration. *Top panels*: BP solutions, one panel per nominal combination of gate and window class. *Bottom panels*: RP solutions. Different colours indicate different values of the index j with the darkest line showing $j = 0$ and lighter colours being used for $j = \pm 1$ and $j = \pm 2$. The median value over the central part of the spectrum ($25.0 < u_i < 35.0$) is plotted.

spectra and for $j = 0$. For BP there seem to be larger relative changes (at about the 10% level) in the wings of the spectra and for $j \neq 0$. This is not completely unexpected and is probably due to correlations between the parameters.

Finally, Fig. 12 shows the evolution through the iterations of the normalised median χ^2 for the same random subset of the calibrators used for which residuals where shown in Sect. 3.3. In this plot, the normalised median χ^2 value at each iteration is obtained by dividing the corresponding median χ^2 by the value at the first iteration. The χ^2 value for each epoch spectrum is given as the sum of squared residuals between the observed spectrum and the predicted spectrum divided by the observed flux error. It is important to point out that the normalised median χ^2 shown here is not the quantity that is being minimised within the iterative process, which will be the sum of squared residuals for all observations of all calibrators within each calibration unit when solving the instrument model and the sum of squared residuals for all observations of each calibrator when solving the source update step. The increase in late iterations for BP shown in Fig. 12 could be due to changes in the distribution of χ^2 caused by the iterations trying to catch a few extreme outliers at the expense of slightly degrading the residuals for other sources.

There are indications from both the standard deviation and χ^2 analyses that in late iterations the solutions start diverging. We have mentioned a possible cause but this is not fully understood. The additional weighting introduced to give more leverage to blue sources seems to have an effect in this respect. Alternative strategies are being considered for future data releases. From the analysis of all criteria, iterations 55 and 40 were finally adopted for BP and RP, respectively, to proceed with the generation of a

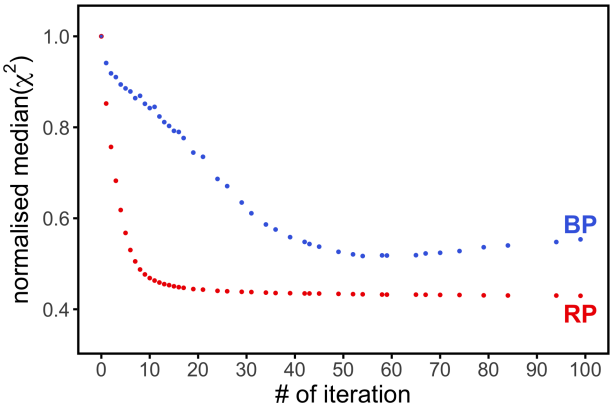

Fig. 12. Normalised median χ^2 for a subset of about 50 K calibrator sources with respect to the iteration number. Blue and red symbols show the BP and RP residuals, respectively.

reference catalogue of mean spectra to be used for the calibration of the CALONLY data.

3.4. Mean spectra representation

Once the internal reference system has been established by the flux and LSF calibration and calibration solutions are available covering all calibration units, a final source update is run including all observed spectra to generate the catalogue of mean spectra that are released as part of *Gaia* DR3. The algorithms described in this section have been applied only to this last run of the source update.

3.4.1. Internal reference system

The flux and LSF calibration procedure described in Sect. 3.3 leads to the definition of an internal reference system. This can be seen as an average instrument. The monitoring of intermediate results during the iterative process showed that in late iterations some of the spectral features in mean spectra assumed a smoother, shallower shape with respect to what is observed in the predicted and observed spectra. In order to ensure that the reference system and the corresponding mean spectra remain as close as possible to the actual instrument and to the actual data, we decided to instead use a specific epoch instrument and to represent the final mean spectra as observed in this system. The epoch instrument was chosen somewhat arbitrarily to be the one corresponding to CCD row 7 for BP and row 5 for RP at a time equal to 4500 in OBMT-Rev.

To avoid having to invert the instrument model to derive mean spectra directly in this new system, we computed a transformation matrix T where each row k contains the coefficients that need to be applied to the canonical Hermite function bases to reproduce the prediction of the kth basis in the chosen epoch instrument. These are the result of a fit of each predicted basis function, which is obtained by applying Eq. (1) to a mean spectrum where only one coefficient has a value equal to 1 while all others are 0, with the same set of 55 Hermite function bases. In the new system, the mean spectra are defined by the array of coefficients $\boldsymbol{b'}$ computed by multiplying the transformation matrix by the array of coefficients in the starting reference system $\boldsymbol{b}$, that is $\boldsymbol{b'} = T\,\boldsymbol{b}$. The covariance matrix of the source update least squares solution also needs to be converted by computing $C' = T\,C\,T^T$ where C is the covariance matrix in the starting reference system and C' is the covariance matrix in the new system.

3.4.2. Bases function optimisation

As described by Carrasco et al. (2021, see Sect. 5) and introduced in Sect. 3.3, the source mean BP/RP spectrum is described as a combination of basis functions. At the start of the calibration process, little is known about the instrument and therefore a generic set of basis functions is used throughout the initialisation phase. Hermite functions, that is, Hermite polynomials multiplied by a Gaussian, were used in this stage: they provide an orthonormal set of basis functions, are centred around zero, and allow to increase details and range by adding higher order bases. These Hermite functions also tend to zero for sufficiently high absolute values of the independent variable. This resembles the behaviour of BP/RP spectra where the combination of CCD efficiency and response ensures that the measured flux tends to zero for increasing distance from the source location.

We denote the $n-$th Hermite function $\varphi_n(x)$. In order to make the Hermite functions efficient in representing the BP/RP spectra, a linear transformation between the pseudo-wavelength and the argument of the Hermite functions is required. This transformation includes a shift $\Delta\theta$ such that the Hermite functions are centred approximately on the centre of the spectra, and a scaling factor Θ that adjusts the width of the Hermite functions to the width of the spectra to be represented. Furthermore, a suitable number of Hermite functions needs to be chosen. The BP/RP spectrum of a source s, $f_s(u)$, is then represented by the linear combination

$$f_s(u) = \sum_{n=0}^{N-1} b_{s,n}\,\varphi_n\left(\frac{u-\Delta\theta}{\Theta}\right). \tag{4}$$

In Eq. (1) the mean spectrum $f_s(u)$ appeared as $\sum_{n=0}^{N-1} b_{s,n}\,\varphi_n$. Here we have made explicit the transformation of the pseudo-wavelength u into the argument of the Hermite functions φ_n. The values of Θ, $\Delta\theta$, and N cannot be chosen independently from each other. Since the pseudo-wavelength range covered by most BP/RP spectra is $[0, 60]$, a value of $\Delta\theta$ of around 30 is required to centre the Hermite functions on the spectra. Furthermore, the linear combination of Hermite functions need to cover the range from -30 to 30. Increasing the number of Hermite functions used in the representation results in the coverage of a wider range of arguments, while increasing the scaling factor results in a reduction of the range of arguments (Carrasco et al. 2021). To find a suitable combination, we first determined the values of N for $\Delta\theta = 30$ for values of Θ from 2 to 3.5 such that the local minimum or maximum at the largest value of u of the $N-1$th basis function is close to 30. For all resulting combinations of Θ and N, a fixed number of five iterations of the instrument calibration was performed. A random subset of approximately 50 000 internal calibrators was used for this purpose. The total residuals in the epoch spectra were then computed and compared for different combinations of parameters. The distribution of the residuals versus various parameters were analysed to select the final combination of parameters. In both BP and RP, $N = 55$ is used, implying that 55 coefficients will be available for each BP/RP spectrum in *Gaia* DR3. The values for Θ and $\Delta\theta$ are slightly different for BP and RP, with $\Theta = 3.062231$ for BP and 3.020529 for RP, and $\Delta\theta = 30.00986$ for BP and 30.00292 for RP. The slight deviation from round numbers is simply the result of adjusting the parameters to the smallest and largest values in pseudo-wavelength in the set of internal calibrators used.

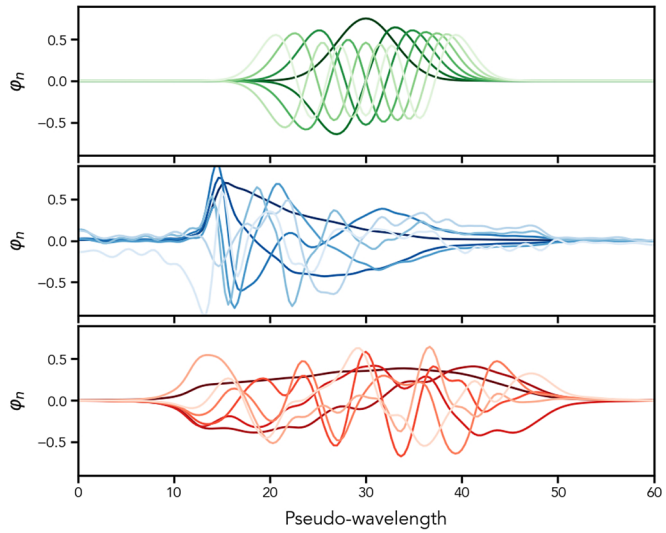

Fig. 13. Comparison between the first few canonical Hermite function (*top panel*), BP (*middle panel*), and RP (*bottom panel*) optimised bases.

Once the catalogue of mean spectra for the calibrators is established based on the set of standard Hermite functions, the set of bases can be optimised to improve the efficiency of the representation. This is achieved when most of the information is contained in the coefficients for the bases with the lowest indices and allows us to reduce the number of coefficients required to describe each spectrum by dropping coefficients that are within the noise.

The optimisation algorithm used normalised mean spectra for the subset of calibrators already used to define the best configuration for the standard Hermite functions. L2 normalisation was used to ensure equal weights for sources of different magnitude in the decomposition. The N coefficients representing each of these sources in the canonical set of bases are normalised with respect to their l_2-norm and are used to populate a matrix $M \times N$ where M is the number of sources. Singular value decomposition of this matrix gives the orthogonal matrix V that represents a rotation of the canonical Hermite bases into a new set of optimised bases.

Figure 13 shows the first few bases in the canonical Hermite function set (in the top panel) and in the optimised BP and RP sets of bases (in the following two panels). Darker shades are used for bases with lower indices. The first optimised bases, being tailored to the actual spectra, reproduce the average spectrum and exhibit the imprint of the transmission curve. Higher order bases become increasingly complex with narrower wavy structures required to fit the sharpest features in the spectra.

3.4.3. Truncation

As explained in Sect. 3.4.2, by expressing the mean spectra in terms of an optimised set of basis functions, a particular spectrum is essentially described by a small number of basis functions with low indices. The coefficients corresponding to higher order basis functions have small absolute values, and, taking their errors into account, are close to zero. Their effect in representing an BP/RP spectrum is therefore essentially adding noise, which manifests itself in wavy structures in the sampled spectrum. It is therefore of interest to suppress the insignificant high-order coefficients and with it, reduce the noise on the spectra.

A simple criterion to decide whether a number of high-order coefficients is insignificant or not has been suggested by Carrasco et al. (2021). The criterion is based on the standard deviation of the M coefficients with the highest indices, that is, the coefficients with indices ranging from $N - M$ to $N - 1$. All coefficients are normalised by their standard errors. We then compute the standard deviation of the M normalised coefficients with the highest order. If this standard deviation remains below a specified threshold, the M coefficients are considered insignificant. As threshold we use a multiple x of the standard error of the standard deviation. For the standard deviation of a set of M samples from a standard normal distribution we assume the simplified expression of $1/\sqrt{2(M-1)}$, and a mean of one. Thus, if the standard deviation of the M normalised coefficients with highest indices is smaller than $1 + x/\sqrt{2(M-1)}$, the coefficients are assumed to be consistent with being zero, and can be truncated. We adopted a value of $x = 2$, and for each BP/RP spectrum, progressively increasing values of $M > 2$ were tested for truncation until the standard deviation of the M coefficients exceeds the threshold for some M. If the truncation threshold is never reached, that is, all coefficients are considered to be consistent with being zero, the full number of $N = 55$ is kept. This happened for a small number of sources, in particular for BP spectra of faint and very red sources, where the flux in the BP spectrum is so low that it is indeed essentially consistent with being only noise.

This criterion makes two simplifications. First, the assumed mean and standard deviation is inaccurate for very small numbers of M. However, the resulting overestimation of the truncation threshold is on the level of a few per cent in the worst case, and has no significant impact on the truncation levels. Second, the truncation ignores correlations between the errors on the coefficients. For sources for which the optimised basis was constructed, the correlations are indeed very low, and the negligence is justified. This is the case for the vast majority of sources. On the other hand, for sources for which the optimised basis is less efficient, correlations might be larger, and the truncation unreliable. This is in particular the case for extremely red sources, or sources with spectral energy distributions that are very different from typical stellar spectral energy distributions, such as quasi-stellar objects (QSOs) or sources with strong emission lines. In the latter case, the truncation is to be used with caution, as it might affect the representation of narrow spectral features.

In the following, we illustrate the effect of truncation for four example cases. First, we consider the case of a typical, bright star ($G \approx 11.5$ mag and $G_{\rm BP} - G_{\rm RP} \approx 1.0$ mag) in Fig. 14. The top panels compare the sampled BP and RP spectra, represented by all 55 coefficients, and by the number of coefficients considered significant according to the procedure described above. These numbers of coefficients are 35 and 15 for BP and RP, respectively, for this example source. No difference in the sampled spectra is visible to the eye, although the number of basis functions used in the representation of the sampled spectrum is significantly smaller. The bottom panels of the figure illustrate the truncation process. The black symbols show the values of the coefficients normalised to their errors. The red curve shows the standard deviation of the M normalised coefficients, starting from $M = 3$ on the right-hand side. The blue shaded region is the cone given by $1 \pm 2/\sqrt{2(M-1)}$. When the red curve remains below the upper limit of the blue cone, the corresponding higher order coefficients are considered insignificant.

The truncation becomes more significant for noisier spectra. As a second example, we therefore consider a source with a similar colour as the first example, but fainter magnitude

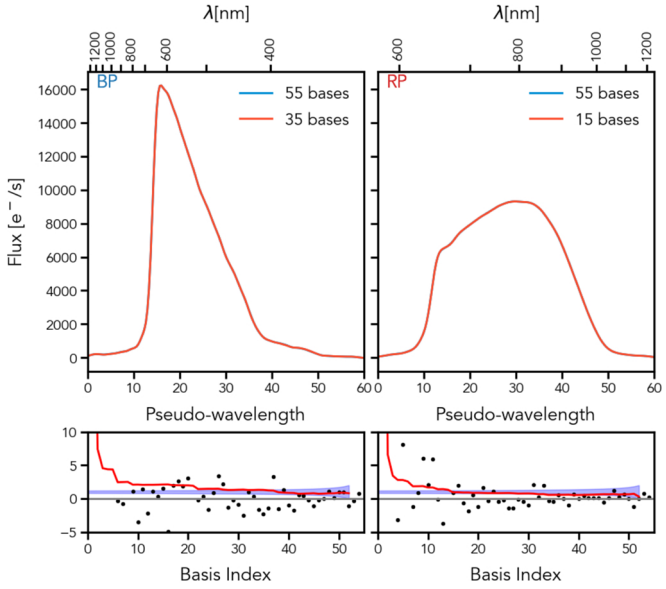

Fig. 14. Sampled BP (*left*) and RP (*right*) spectra are shown in the *top panels* for source Gaia DR3 6210089815971933056 ($G \approx 11.5$ mag and $G_{BP} - G_{RP} \approx 1.0$ mag). Each panel contains two curves: a blue curve showing the non-truncated spectrum using all 55 coefficients, and a red curve showing the truncated spectrum. The number of coefficients used for each spectrum is given in the label within the plot. The *bottom panels* show the truncation assessment. This is run independently for BP and RP. The black circles indicate the coefficients normalised by their formal errors, and the red line shows the standard deviation of the M normalised coefficients, starting from $M = 3$ on the right-hand side. The blue shaded region is the cone given by $1 \pm 2/\sqrt{2(M-1)}$.

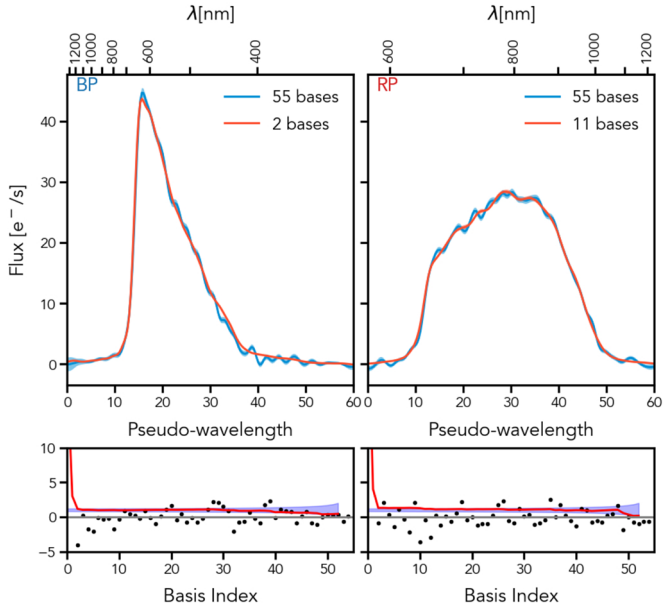

Fig. 15. Illustration of the effects of truncation on the mean spectra of source Gaia DR3 6776463197626299392 ($G \approx 18.1$ mag and $G_{BP} - G_{RP} \approx 1.0$ mag). We refer to the caption of Fig. 14 and the text for details.

($G \approx 18.1$ mag and $G_{BP} - G_{RP} \approx 1.0$ mag; see Fig. 15). In this case, more coefficients are in agreement with being zero, and the number of significant coefficients is only 2 and 11 for BP and RP, respectively. Truncating the representation of the spectra at these numbers of basis functions maintains the general shape of the spectra, but suppresses the wavy patterns introduced by the noisy higher index coefficients.

We also show examples for sources with emission lines. The first case is a bright source ($G \approx 11.5$ mag) with multiple emission lines in BP and RP, shown in Fig. 16. Here, the truncation criterion is not even reached for $M = 3$, as all coefficients are required to represent the complex spectra for this source. In similar cases, the number of significant coefficients should have been set to 55, but as the cases where $M <= 2$ were not tested for the truncation criterion, 53 is the maximum number returned by the algorithm. Therefore, the use of all 55 coefficients is recommended in cases where the number of significant coefficients is 53.

Finally, we consider a faint QSO with emission lines as an example. Figure 17 shows the BP and RP spectra of a QSO ($G = 18.7$ mag and $G_{BP} - G_{RP} = 0.5$ mag), with all 55 coefficients, and with the truncated representation, using 3 and 11 coefficients in BP and RP, respectively. The spectral energy distribution from SDSS is shown for comparison. In particular, the strong emission line visible in the SDSS spectrum coincides with a line in the BP spectrum. This line is removed by the truncation process. The truncation in the case of complex spectral shapes might therefore be too strong.

The truncation procedure was also tested by the subsystem dedicated to the estimation of astrophysical parameters within the DPAC analysis pipeline, referred to as Apsis (see

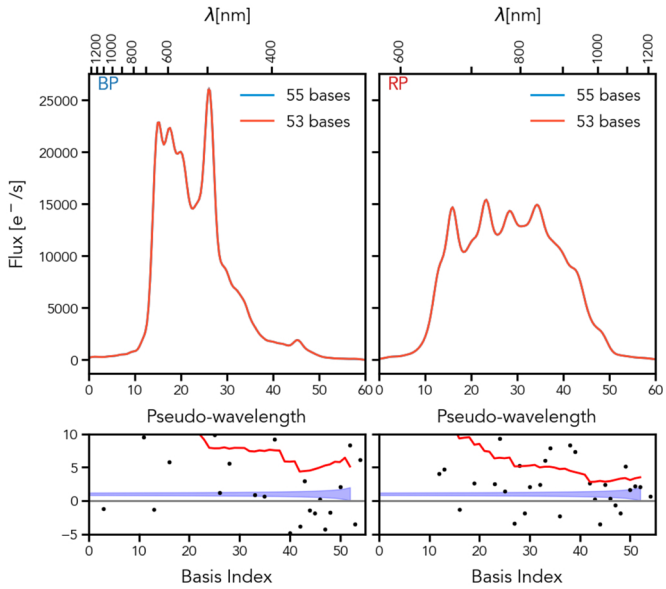

Fig. 16. Illustration of the effects of truncation on the mean spectra of source Gaia DR3 3032940844556081408 ($G \approx 11.5$ mag). We refer to the caption of Fig. 14 and the text for details.

Creevey et al. 2023). Most Apsis modules found that the truncation would have a negative impact on the quality of the scientific results based on the emission lines of quasars and certain types of stars. These tests were conducted at a very early stage when no external calibration was available, such that the conclusions were uncertain and most Apsis modules considered not truncating the coefficients to be the safer option. In the extreme case of ultracool dwarfs, which are very red and very faint stars, the truncation was found to have a positive impact and has been employed specifically for the Apsis module ESP-UCD which focuses on