

ESO Phase 3 Data Release Description

Data Collection	XSHOOTER
Release Number	1
Data Provider	ESO, Science Data Quality Group
Document Date	2026-05-07
Document version	2.7
Document Author	Reinhard Hanuschik, Wolfgang Hummel, Burkhard Wolff

Abstract

This is the release of reduced 1D spectra from the XSHOOTER¹ spectrograph, ECHELLE (SLIT) mode (as opposed to the IFU mode which is not included). All spectra have been reduced under the assumption of point-like sources. This release is an open stream release, it includes so far archived XSHOOTER data and will be continued into the future

The selected data cover the vast majority of the entire XSHOOTER data archive, from the begin of operations in October 2009 until present. The data have been reduced with the XSHOOTER pipeline, version `xshoo/2.3.12` and higher. All data have their instrument signature removed: they have been de-biased, flat-fielded, wavelength-calibrated, order-merged, extracted, sky-subtracted and finally flux-calibrated. Telluric absorption has *not* been corrected for. The pipeline output products come in the ESO 1D standard binary table², along with some ancillary files.

The processing is performed by the Science Data Quality Group in an automated process. The pipeline processing uses the archived, closest-in-time, quality-controlled, and certified master calibrations. It is important to note that the reduction process itself is automatic, while the quality assessment and certification of the master calibrations is human-supervised. The overall data content grows with time as new data are being acquired and processed (with a delay of up to a few weeks).

The data format follows the ESO 1D spectroscopic standard² for Phase 3 data products. Each spectrum is a multi-column binary table. There is one product file for each set of input raw files. XSHOOTER products always come separately for each arm: UVB, VIS, or NIR.

This data release offers science-grade data products, with the instrumental signature and sky background removed, flux-calibrated, with error estimates and quality flags, and with a list of known shortcomings. They are considered to be ready for scientific analysis. They are expected to be useful for any kind of medium-to-high resolution spectroscopic research, including abundance and line profile studies, and radial velocity studies. For studies of energy distributions the researcher should keep in mind that the XSHOOTER instrument is a medium-to-high resolution spectrograph, designed for stability and throughput. It is usually not operated to achieve high-precision spectrophotometry, unless the wide slit is used. Slit losses have to be expected, and they may differ from arm to arm. There has been no attempt to correct for telluric absorption lines; this might affect certain science applications.

The telluric standard stars, measured frequently during most nights in order to provide a representative reference for telluric absorption lines, have been included in the data release as stand-alone science-grade spectra. They can serve two purposes: providing high-quality reference spectra for telluric correction, and providing science-grade spectra of early-type and solar-type stars, often repeated at various epochs.

Disclaimer. Data have been pipeline-processed with the best available calibration data. However, please note that the adopted reduction strategy may not be optimal for the original scientific

¹ <https://archive.eso.org/docurl?XSHOOTER>

² <http://www.eso.org/sci/observing/phase3/p3sdpstd.pdf>

purpose of the observations, nor for the scientific goal of the archive user.

Release Content

The XSHOOTER release is a stream release. The overall data content is not fixed but grows with time as new data are being acquired and processed. The data are tagged "XSHOOTER" on the ESO archive user interface³.

How to organize

The downloaded data come with their technical archive names. The first useful step to do is the organization of the data on the local disk. The README file coming with every download contains the information necessary to properly rename the files and establish their association. The fundamental unit of the renaming process is to capture column #2 in the README file (technical name starting with ADP) and column #3 (original name) and use something like 'mv \$2 \$3'. The file type is in column #4, find the possible values listed in Table 6. The spectral product is of type SCIENCE.SPECTRUM while the other files are of type ANCILLARY.<subtype>. Once you have renamed the files, they carry the names as also listed in the ASSON1, ASSON2 etc. header keys of the first extension of the file. See further down for a description of the file types and their content.

Then, the easiest way to identify all files belonging together is the timestamp in their name, e.g. 2009-10-05T07:45:26.840, which is common to all file names belonging together.

Data Selection

Data selection is rule-based. It is organized along the following criteria:

- instrument=XSHOOTER (or SHOOT as used until March 2010);
- observing technique (DPR.TECH) = ECHELLE,SLIT,<obsmode> with obsmode being one of STARE, NODDING , or OFFSET (**not**: ECHELLE,IFU and ECHELLE,SLIT,MAPPING⁴);
- category (DPR.CATG) = SCIENCE or CALIB (for the telluric standard stars);
- type (DPR.TYPE) = OBJECT (for the science data) or STD,TELLURIC (for the telluric standard stars).

No selection has been made on the basis of the observing mode (visitor or service), nor on settings. XSHOOTER ECHELLE settings are defined by the combination of arm (UVB, VIS, or NIR), binning (1x1, 1x2, or 2x2), and slit width. The MAPPING data have not been included because they are 2D scans of extended objects and require user decisions for the background subtraction and the extraction strategies.

We have processed only those data for which certified master calibrations exist in the archive. Depending on their type, these master calibrations exist at daily frequency or less. All master calibrations used here were certified, meaning checked for quality and proper registration of instrument effects. For XSHOOTER, the selection pattern for master calibrations was complete from the begin of operations on, covering both Service Mode (SM) and Visitor Mode (VM) science data and all settings.

Note that the fact that a certain set of XSHOOTER raw files does not have a product in this release does not necessarily imply a quality issue with the raw data. There are sometimes acquisition patterns for which the processing jobs can be predicted to fail (e.g. NODDING with one raw file; see "Rejected or failed processing" below). In some of those cases the raw data would probably process fine with a fine-tuned strategy.

³ http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form

⁴ The observing mode MAPPING was introduced in December 2015.

In general, the master calibrations were processed with different (earlier) pipeline versions than the science data in this stream. The most significant change in the XSHOOTER pipeline was related to version `xshoo-2.0` and higher, reflecting the results of a major science-grade review in 2012 and 2013. This change was implemented in 2013-01. All science data acquired later than that date have been processed with essentially the same pipeline version as their master calibrations. All science data acquired earlier have still been processed with the current pipeline version but their master calibrations were processed with earlier versions. See “Master calibrations” below for a discussion of the main differences.

Science data with the PROG.ID starting with 60. or 060. have been de-selected, considering them as test data. On the other hand, telluric standard stars have (almost) always the PROG.ID 60.A-9022(C) and have always been included in the processing scheme.⁵

Data taken at daytime (with obviously wrong ‘SCIENCE’ tag) have been ignored. Otherwise the header tag ‘SCIENCE’ has been blindly accepted from the raw data (originally defined by the PI), thus including sometimes standard stars intended by the PI for use as flux calibrator (or as telluric standard star) but labelled wrongly. This is often evident from the OBJECT header key. Also, there are rare cases when test observations were executed under the SCIENCE label. Some very short exposures with no signal fall into that category but have not been suppressed.

There is no data rejection based on quality. Likewise, we have not considered OB grades⁶ for selection: the observations might have any grade between A and D (if taken in SM), or X (in VM). The availability, or non-availability, of a particular file in this release does not infer any claim about the data quality.

For the observing modes NODDING and OFFSET, we have combined all raw files taken in one template and from one arm in a single product file. Note that we have combined only within a template; multiple observations of the same OB, or observations across OBs, *have not been co-added*. In the automatic quality checks of the products, we have flagged those rare cases when a template was aborted and an additional quality issue occurred (see section ‘Data Quality’).

For observing mode STARE, we have processed every single input raw file separately for data taken before December 2025. Later observations have been combined on the OB level (as NODDING and OFFSET data) since single X-Shooter OBs are only rarely used to study time variability.

We have included the telluric standard stars in the current release. They are acquired by the observatory staff during night-time, often in slots about every 2 hours. They match the science data of the previous 2 hours in setup (read mode, slit width, DIT for NIR) and airmass (within about +/- 0.2). If observatory-provided, they are always taken in NODDING mode (until 2013-04 in STARE mode). The strategic goal is to provide high-SNR spectra with about the same telluric absorption signature as the science data. The rules about time and airmass match are motivated by proximity of the PWV⁷ value causing a certain telluric absorption line depth. The setup match rule is followed in order to match the line width.

While the flux-calibrated reduced spectra of these standard stars might be useful as a first step towards a telluric correction of the science data, we also think they might be useful as stand-alone, high-SNR spectra of early-type, or solar-type, stars. Most of these targets are taken from the Hipparcos Catalogue.

Because of their special acquisition pattern, there are some quality remarks, see the section ‘Data quality’.

⁵ Users interested in searching only for XSHOOTER telluric standards should use this PROG.ID on the interface.

⁶ As given by the observatory staff to assess the match with user-defined constraints, for SM data.

⁷ PWV=precipitable water vapour.

Release Notes

Depending on the observing mode, the data reduction uses the standard XSHOOTER pipeline recipes `xsh_scired_slit_<mode>`, with mode being one of `stare`, `nod` or `offset`. Find a description of these science recipes in the Pipeline User Manual⁸.

Find the pipeline version used for processing in the header of the product file under “PROCSOFT”.

All recipe parameters, including the extraction methods, were set to the pipeline defaults. The only exceptions are purely technical⁹. Find more details in the User Manual, sections 10.15-10.17 (for version 12.1).

Pipeline Description

Information about the XSHOOTER pipeline (including downloads and manual) can be found under the URL <http://www.eso.org/sci/software/pipelines/>.

Data Reduction and Calibration

The main reduction steps of the XSHOOTER Echelle science and telluric standard star data are the following:

STARE observations. STARE observations taken before December 2025 have been reduced individually; multiple STARE observations from the same template are not stacked. Later observations have been combined if coming from the same template. Reduction steps are:

- The input spectrum gets bias-corrected (UVB and VIS arm), or dark-corrected (NIR arm).
- Cosmic rays are detected on the input spectrum using Laplacian edge detection.
- Only for observations from December 2025 onwards: If there is more than one input science frame, then they are combined into a single frame using the median (`stack-method = median`).
- The order table is used to locate the inter-order regions where the background is fit with a polynomial, which is then subtracted from the science data.
- The master flat frame is divided into the science frame.
- The sky background is subtracted (`sky-subtract=TRUE`, `sky-method=MEDIAN`).
- The science frame is rectified order by order and wavelength-calibrated, thus passing from pixel-pixel coordinates to wavelength-spatial coordinates.
- The spectrum is localized (`localize-method=MANUAL`). A simple sum extraction is done on the 2D rectified orders. No optimal extraction was performed (although the pipeline offers it) since it has artefacts.
- The extracted 1D orders are merged into a single spectrum.
- The spectrum is finally flux calibrated, either with a master response curve (until acquisition date 2015-03-31) or with the closest-in-time nightly response curve (after that date, see below).

The localization was done with the method `MANUAL` (`localize-slit-position=0.0`, `localize-slit-hheight=2.0`). This is the most robust method but not always the correct one (if the object is off-centre). The position of the peak signal can easily be checked on the preview plot #2 (for VIS arm data).

NODDING observations. They get combined within the same template, per arm; the pipeline assumes constant transparency and equal exposure times for all input frames. Reduction steps are:

- All input spectra are bias-corrected (UVB and VIS arm). N input frames are combined (alternating: object top, object bottom) where N must be even.
- Cosmic rays are detected on each input spectrum using Laplacian edge detection.
- The recipe combines the science frames taken at the same position (if e.g. sequence was

⁸ Under the XSHOOTER pages in <http://www.eso.org/sci/software/pipelines/>.

⁹ Namely: `generate-SDP-format=TRUE` (to generate the output format of the spectra).

AA BB, then data are combined as A B) doing median stacking (`stack-method=median`).

- Pairs are subtracted (e.g. A-B).
- The master flat frame is divided into each subtracted science frame pair.
- Each science frame pair is rectified order by order and wavelength-calibrated, thus passing from pixel-pixel coordinates to wavelength-spatial coordinates.
- Since `correct-sky-by-median=TRUE`, the recipe calculates and subtracts the median pixel value for each wavelength column in the rectified frame from the column pixel values.
- For each subtracted science frame pair, the recipe forms [(A-B) - shifted (B-A)] and then combines them using `combinenod-method=MEAN`.
- The standard extraction method is used to extract the spectrum (`extract-method=NOD`).
- The extracted 1D orders are merged into a single spectrum.
- The spectrum is finally flux calibrated, either with a master response curve (until acquisition date 2015-03-31) or with the closest-in-time nightly response curve (after that date, see below).

In case of odd number of input files, the pipeline rejects the last one and processes all others. Nodding with one input frame always fails.

The pipeline versions earlier than `xshoo/2.7.0` had a bug with incorrect handling of the nod throw¹⁰. The pipeline always assumed 5.0 arcsec, independent of the actual value. The effect was that for nod throw < 5 arcsec, the negative parts of the nodded signal were (partly) included in the extraction mask, and the extracted fluxes - and the corresponding SNR - were compromised. Data with nod throw smaller than 3.5 arcsec were strongly affected (amounting to a few percent of all data), the other data only marginally if at all (most data have the default 5.0 arcsec).

All nodding data with nod throw <= 3.5 arcsec have been reprocessed and are available as corrected versions. Data with a higher nod throw have not been reprocessed since the impact is marginal. Data with a nod throw larger than 5 arcsec are not affected. Data acquired after 2016-11-01 (any nod throw) have been properly processed right away. Data in other observing modes are not affected.

OFFSET observations. SKY and OBJECT frames get combined within the same template, per arm; the pipeline assumes constant transparency and sky brightness, and equal exposure times for all input frames. Reduction steps are:

- Cosmic rays are detected on each input frame using Laplacian edge detection.
- SKY frames are subtracted from OBJECT frames in pairs.
- The master flat frame is divided into each sky-subtracted science frame.
- Each sky-subtracted science frame is rectified order by order and wavelength-calibrated, thus passing from pixel-pixel coordinates to wavelength-spatial coordinates.
- The rectified science frames are stacked (`combinend-method=MEAN`).
- The standard extraction method is used to extract the spectrum.
- The extracted 1D orders are merged into a single spectrum.
- The spectrum is finally flux calibrated, either with a master response curve (until acquisition date 2015-03-31) or with the closest-in-time nightly response curve (after that date, see below).

Master Calibrations used for data reduction

Please check Sect. 7 of the Pipeline User Manual for the description of calibration data.

Table 1. Set of master calibrations used for data reduction

¹⁰ Nodding throw (or nod throw for short) = displacement between the two nodding positions; by default 5 arcsec but users can modify that value in the phase-2 OB design.

Type* (pro.catg)	name (first part)	Mandatory**/ optional	content
MASTER_BIAS_UVB or VIS	XS_MBIA	optional (but always provided)	master bias: created from 5 raw bias frames; removes bias level and bias structure.
MASTER_DARK_NIR	XS_MDRK	optional (mandatory for STARE)	master dark: created from 1 or 3 raw dark frames; removes dark level and dark structure.
ORDER_TAB_EDGES_SLIT_<arm> ORDER_TAB_AFC_SLIT_<arm>	XS_POES; XS_POTA	mandatory; optional, in addition	table with traces of order edges; same, corrected for flexure
MASTER_FLAT_SLIT_<arm>	XS_MFSL	mandatory	master flat: created from 5 (VIS) or 10 (UVB, NIR) raw flats; used for: defining order edges, removing gain noise, blaze function, slit noise; introducing lamp energy distribution. In UVB, two lamps are used, a D2 and a quartz lamp.
DISP_TAB_<arm> DISP_TAB_AFC_<arm>	XS_PDT2; XS_PDTA	optional; optional, in addition	dispersion table same, corrected for flexure
XSH_MOD_CFG_OPT_2D_<arm> XSH_MOD_CFG_OPT_AFC_<arm>	XS_PMC2; XS_PMCA	mandatory; optional, in addition	physical model configuration file same, corrected for flexure.
SPECTRAL_FORMAT_TAB_<arm>	XS_GSFT	mandatory	Static spectral format table
BP_MAP_RP_<arm>	XS_GBPM	optional (but always provided)	Static bad pixel map
Until 2015-03-31: MRESPONSE_MERGE1D_SLIT_<arm>	XS_MRSP	mandatory	Static response curve used for flux calibration; calculated from selected sets of standard star measurements; removes lamp response and remaining instrument signature
As of 2015-04-01 and later: RESPONSE_MERGE1D_SLIT_<arm>	XS_PRMF	mandatory	Response curve (closest-in-time) used for flux calibration; removes lamp response and remaining instrument signature
ATMOS_EXT_<arm>	XS_GEXT	mandatory	Static table with atmospheric extinction
SKY_LINE_LIST_<arm>	XS_GSLL	optional (for STARE mode)	static table with sky line positions

* arm = UVB|VIS|NIR

** if missing, pipeline would fail

Response curves (all three modes). Two different flux calibration schemes have been used:

- **2009-10-01 ... 2015-03-31:** master response curves (in all 3 arms);
- **2015-04-01 .. current:** nightly response curves from standard stars (in all 3 arms).

For almost every night used for XSHOOTER science, there exists a flux standard measurement. On the other hand XSHOOTER is neither designed for, nor operated under strictly photometric conditions, and for that reason the flux calibration cannot aim at photometric accuracy, but is geared towards removing the remaining instrumental signature (mainly from the spectral energy distribution of the flat lamp). For this purpose master response curves are sufficient, which are carefully compiled from individual (nightly) response curves, giving an accuracy of the flux scale in

the science products of about +/- 10% differentially (meaning the instrumental effects are removed with that accuracy). Still the photometry could be off by larger amounts, due to uncontrolled transparency variations and slit losses. Because of the increasing instability of the UVB flat-field lamp, the concept of master response curves was given up and replaced by nightly response curves for all data taken after 2015-03-31. See section 'Flux calibration' for more details.

Wavelength scale. The XSHOOTER Echelle products have a topocentric wavelength scale; no corrections for barycentric or heliocentric motion have been applied. The corresponding values have been calculated by the pipeline and are stored in the header (HIERARCH.ESO.QC.VRAD.BARYCOR and ...HELICOR, in km/s). The wavelength scale refers to air.

Telluric absorption lines. No correction for telluric absorption lines has been applied. For almost all spectra, one or more telluric standard stars have been measured within 2 hours of the observation, and within 0.2 in airmass. The telluric correction is not provided by the pipeline because it would require assumptions and judgements on the science spectrum.

For most of the telluric standard stars a fully reduced spectrum is also available in this release¹¹. The user may want to select the best suitable one and apply standard techniques to process it into a template for telluric absorption line correction with the same spectral resolution and about the same PWV value as the science data.

You could also download the raw data from the ESO archive through the 'calSelector' service, along with the master calibration files, process them and apply the product to the science spectra. Alternatively, users may want to visit the ESO *skytools* web page¹² for appropriate tools.

Be aware that a telluric standard star spectrum and a science spectrum, even though formally matching in slit width and complying with the operational match rules (within 2 hours in time, and 0.2 in airmass), not necessarily match in FWHM of the telluric lines. If the exposure-time integrated seeing for a science or a standard star spectrum was better than the slit width, the effective resolution will depend on the seeing rather than on the slit width.

Bad pixel map. We used *static* bad pixel maps for the processing, selected per arm and binning¹³.

Normal daytime calibrations and AFC frames (automatic flexure compensation). During the lifetime of XSHOOTER three different association schemes were used for ORDER_EDGE, DISP_TAB, and XSH_MOD_CFG_OPT master calibrations (dates are acquisition dates):

- **2009-10-01 ... 2011-03-08:** the products from the daytime calibrations have been taken, i.e. no AFC data (cf. Table 1).
- **2011-03-09 ... 2012-06-30:** only the products from the AFC frames (XS_POTA, XS_PDTA, XS_PMCA) have been used.
- **2012-06-30 ... now:** products from both daytime calibrations and AFC frames are associated.

A careful investigation of the shifts in X and Y direction to be compensated has shown that an effective difference in quality was not noticeable. Hence we have decided to keep the historical association scheme and process the data accordingly.

Master calibration names and recipe parameters used for reduction. The product header contains a list of all used master calibrations, look for keys "HIERARCH ESO PRO REC1 CAL<n> NAME" and "... CATG", with the index n. The pipeline parameters and their values are listed as "HIERARCH ESO PRO REC1 PARAM<n> NAME" and "... VALUE".

Products. The XSHOOTER pipeline creates a large number of intermediate and final product files.

¹¹ Best identified on the archive page by PROG_ID=60.A-9022(C).

¹² <http://www.eso.org/sci/software/pipelines/skytools/>

¹³ The pipeline also supports dynamic bad pixel maps derived from flats, their advantage being a closer time match, at the price of far less controlled quality.

The final product in the spectroscopic data format combines information from the following products:

- extracted spectrum (de-biassed, flat-fielded, sky-removed, extracted, wavelength-calibrated, rebinned but not fluxed), including errors and quality flags;
- flux-calibrated spectrum (same as above but also fluxed), including errors and quality flags.

These products are combined into the final binary table FITS file which is the delivered main data product, with the wavelength scale as first column and all other products as further columns (see 'File Structure' below).

There is also an ancillary FITS file:

- 2D (not extracted but rectified) spectrum, *not flux-calibrated*, with the 2D information of the signal available for quality checks.

There are further ancillary files for each product:

- a text file with OB-related information about the product file, including OB classification (if available) and comments and QC information (*Table 2*); discontinued for data taken from December 2025 onwards;
- one or two (in the case of VIS arm) overview plots, displaying the spectrum, a smoothed version, the unfluxed version, and the SNR; for the VIS arm, the second graphical file contains an overview plot of all three arms, discontinued for data taken from December 2025 onwards.

All XSHOOTER spectra come separately by arm, like the raw data. Typically there are three spectra (from the UVB, VIS and NIR arm).

The three products have essentially no spectral gap. The UVB product has the nominal signal truncated beyond 556nm, to suppress a spurious pseudo-absorption at 570nm arising from artifacts in the flat field (the suppression is controlled by the parameter `cut-uvb-spectrum=TRUE`).

Table 2. Content of ANCILLARY.README text file

parameter	values	meaning
OB related information:		
SM_VM	SM or VM	Data taken in Service Mode or Visitor Mode; VM data are less constrained in terms of OB properties, have no user constraints defined and therefore no grades (formally graded X meaning 'unknown')
OB_GRADE	A/B/C/D; X	Immediate grade given by night astronomer, considering ambient conditions checked against user constraints
OB_COMMENT S	Free text	Any optional comments added by the night astronomer, together with the <i>approximate</i> UT hh:mm (truncated after 200 characters). [Note that OBs might have been executed several times during the night, with or without comments. In those cases the user should always carefully check that the listed comment applies to the spectrum with the closest <i>previous</i> timestamp.]
QC related information:		
QCFLAG	e.g. 0010010	QC flag composed of 7 bits, see

parameter	values	meaning
		<i>Table 4.</i> Discontinued for data taken from December 2025 onwards.
QC_COMMENT	Free text	Automatically added comment if a quality issue is discovered by the processing system (there are no human-provided comments)

Rejected or failed processing. Under the following conditions the pipeline processing always fails, and no data products exist¹⁴:

- NODDING data sets with only one input file;
- OFFSET data sets with no OBJECT frames;
- OFFSET data sets with no SKY frames.

The possible reasons for the existence of such data could be:

- Aborted templates (then the data are likely to be useless).
- Extended sources: sometimes observations of extended sources are made with the OFFSET template in a quasi-mapping mode (but with the incorrect DPR.TYPE keyword OFFSET), without OBJECT-SKY pairs; unfortunately the acquisition template then does not encode properly the information of the extended nature of the source in the headers, nor of the intended strategy of the observations. The proper reduction of these data would be possible with knowledge of the science case, which however is beyond the scope of this project.

Furthermore there are occasionally processing failures. A typical case is a high number of nodding input frames (20 or more) for which the pipeline sometimes, after long iterations, did not converge. No major attempts were made to understand and fix the situation, unless it was obvious.

Note that only in very rare cases we have rejected a product that was successfully created by the pipeline. Even in case of heavy saturation we have decided to deliver the product to the archive, since the saturated pixels are marked in the QUAL column (*Table 5*), and the unsaturated regions might still be useful. Also, products with large parts of negative flux have not been rejected

Only in the following cases we have rejected a product:

- the template has been aborted AND another issue showed up (e.g. saturation, or negative fluxes);
- the product column FLUX is entirely 0.

Flux Calibration

One of the main improvements of the XSHOOTER pipeline over time was the flux calibration. Initially (before pipeline version `xshoo-2.x`) the response curves, as derived from almost nightly flux standard stars, had a poor algorithmic quality (e.g. uncontrolled splines). With the improvements of version `xshoo-2.x`, these issues have been solved, and the response curves are of good quality.

For the years 2009 until 2015, the response curves showed very good stability, and we decided to use *master response curves*. The main reason is that we are aiming at a *flux calibration*, not at a *photometric calibration*. XSHOOTER data have been taken under conditions which were not photometrically controlled, and depending on the slit width used, they suffer from slit losses. Without the goal of a photometric calibration, the use of daily, close-in-time response curves would only be reasonable if instrument components were observed to vary on short timescales. This was not the case until about 2015-03.

After that date, however, presumably due to ageing flat lamp, the stability of the UVB arm response curves degraded. In particular the region at 370 nm, close to the Balmer jump, is dominated by the transition between the two flat-field lamps and is particular sensitive to variability. It turned out

¹⁴ Such data could probably be processed interactively, e.g. with *Reflex*.

that the stability criterion of 10% was often violated after that date.

We have therefore decided to switch the flux calibration to the nightly response curves for all data taken on and after 2015-04-01. This approach provides optimal protection against short-term variability of the UVB lamps.

Although this instability affected only the UVB arm, we decided for consistency to switch the flux calibration to the nightly scheme *for all three arms*. The data from 2015-04-01 until 2016-11-01 that were already processed, archived and available for download, have been reprocessed. Data from 2016-11-02 on are processed with the same schema. There was no reason to reprocess the pre-2015-04-01 data, they are unaffected by this issue.

Master response curves. The procedure for the *master response curves* follows the scheme developed for UVES. These are based on carefully compiled individual response curves. The selection criteria were:

- good transparency¹⁵ (in order not to introduce additional bias from unstable atmospheric conditions),
- good algorithmic quality¹⁶.

In order to make use of the `xshoo-2.x` improvements, we started the compilation with data from 2013-01 and later. The master response curve for the NIR arm, averaged from all compiled response curves, and its +/- 10% envelope, are displayed in *Figure 1*.

We then took this master response curve and compared it to individual response curves from the earlier epochs. Again the good (in the above sense) response curves were selected interactively¹⁷ and averaged. That average, together with the +/-10% envelope, is plotted for the UVB arm in *Figure 2*. Then the existing master response curve was multiplied by a fudge factor (to account for differences in the normalization procedure for the master flats at that time), it is displayed as the blue line. It is obvious that adopting that master response curve (originally defined for the 2013 time range) also for the earlier epochs is possible within the +/-10% range. It has small systematic variations, possibly due to instrumental drifts (change of coating properties or flat-field lamp energy distribution) and possibly also to different pipeline versions, but it is well-defined in terms of spline robustness and also step size (in particular in the VIS and NIR telluric windows).

We have tested the validity of our approach with XSHOOTER observations of the flux standard star Feige 110. These data have been observed in NODDING mode, with the widest slit size 5", and have been processed with the corresponding science recipe, including the final step of the flux calibration with the master response curves for the three arms. A comparison of the flux calibrated XSHOOTER spectrum (which should in this case have photometric quality) to the tabulated energy

¹⁵ Defined by the +/-10% envelope of the best response curves, see *Figure 1*.

¹⁶ Bumps and wiggles in the response curve are unacceptable if they do not correspond to instrument features but are due to imperfect records of the standard star energy distribution, or to algorithmic flaws.

¹⁷ Amounting to several hundred individual curves, per arm.

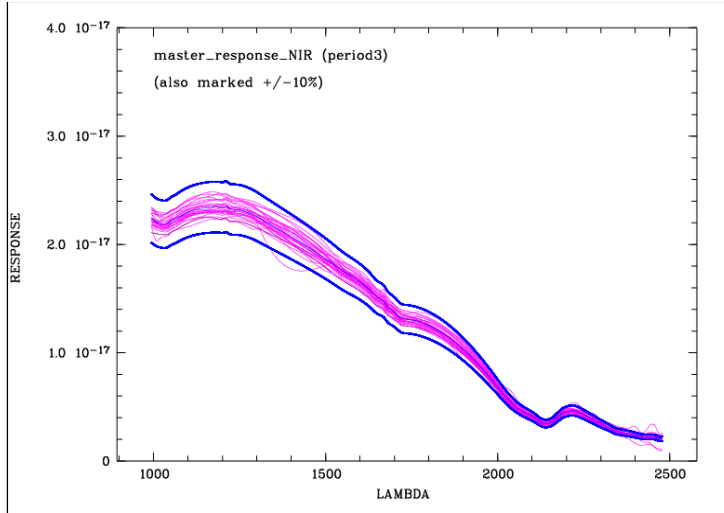


Figure 1. Selection of good response curves, and construction of master response curve. For the NIR arm, we show the selected individual response curves taken after 2013-01 (in purple), and the master response curve (thin blue line) derived from averaging them, with its +/-10% envelope in bold.




Figure 2. Propagation of master response curve back into 2009-2012. For the UVB arm, we show the average of all selected response curves from the pre-2013 period (plotted as thin red line), and its +/-10% envelope. The scaled master response curve from 2013 is overplotted in blue.

distribution is reassuring (Figure 3). We did also a comparison to a flux-calibrated UVES spectrum of the same star (Figure 4). Again, the agreement is very good.

Versions of master response curves. We have identified several periods of validity for the master response curves, separated by breakpoints. One breakpoint occurred on 2013-01-15. It is related to the onset of using the pipeline versions `xshoo/2.x` for the calibrations, in particular for the creation of the master flats. Those versions use a different normalization procedure than the versions `xshoo/1.x` used before. Because both the science data and the standard stars (used for computing the response curves) are flat-fielded, the master response curve and the flat-fields used need to be aligned in terms of pipeline algorithms.

The second breakpoint applies to the NIR arm only and reflects the exchange of the NIR flat lamp, along with its significantly changed chromatic energy distribution.

Table 3. Versions of master response curves

Period	static master response curves	end of period ("breakpoint") defined by:
2009-10 (begin of	UVB, VIS, NIR:	pipeline versions <code>xshoo/1.x</code> , old

Period	static master response curves	end of period (“breakpoint”) defined by:
operations) – 2013-01-15	XS_GRSF_090930A_<arm>.fits	normalization algorithm for master flats
2013-01-15 – 2015-03-31 (for UVB and VIS); 2013-01-15 – 2013-11-05 (NIR)	XS_GRSF_130125A_<arm>.fits	upgrade to xshoo/2.x, implementing a new normalization for the master flats, and a much better definition across telluric windows
2013-11-06 – 2015-03-31 (NIR)	XS_GRSF_131106A_NIR.fits	exchange of NIR flat lamp, modified chromatic energy distribution
2015-04-01 - now	No master response curves	Association schema changed to nightly response curves




Figure 3. The XSHOOTER Feige 110 spectrum from 2012-06-03 (all three arms), processed with the master response curves and plotted in blue, compared to the tabulated energy distribution (in green, together with the 10% confidence range in purple which appears in the log plot as +/-0.04).




Figure 4. The same spectrum as above, compared to a UVES spectrum of Feige 110 overplotted in red (the UVES spectrum comes in several chunks and extends to about 1000nm). The green line is again the tabulated energy distribution.

Data Quality

Master calibrations. All used master calibrations have been quality-reviewed and certified at the time of acquisition. The most important parameters for the quality of the data products are the SNR of the master flats, and the *rms* of the dispersion solution. The SNR of the master flats was always high enough to be dominated by the fixed-pattern (gain) noise, which is important to not compromise the SNR of the science data. The *rms* of the wavelength dispersion was higher in the first few years (with pipeline versions `xshoo/1.x`) than later, due to a lower number of lines found (Figure 5).




Figure 5. Evolution of the two QC parameters *n_lines* (number of lines used for the dispersion solution, plot #1), and *rms* (plot #4), for the VIS arm

SNR. There is a column “SNR” in the product table that is calculated from the signal and the corresponding error. It has no independent information but is provided for convenience. Its mean

value is written as header key SNR. There are also the header keys HIERARCH.ESO.QC.FLUXn.SN (n=1...2 or 3) that describe the SNR in various spectral windows, defined in the key comment.

SPEC_RES. The header key SPEC_RES contains the nominal resolving power, as derived from the arc lamp calibrations of the same slit width. The actual spectral resolving power of the science data has not been measured but could actually be higher, depending on seeing conditions and exposure time¹⁸. A better estimate would be available from the half-width of the telluric absorption lines.

Telluric standard stars. Although processed in the same way as the science spectra, spectra of the telluric standard stars have some characteristic quality features. The goal of the nighttime staff is to achieve a high-SNR spectrum while trying to avoid saturation. With the actual flux per pixel depending on the currently prevailing extinction and also on seeing, it is often unavoidable to have a first attempt ending up in saturated, or underexposed, spectra. Then, follow-up exposures are taken to optimize the flux per pixel. Optimization is usually done per arm, with the exception of the UVB arm which is not relevant for telluric line correction and is not optimized. The outcome of this procedure is often a sequence of spectra of the same target with exposure times getting shorter, or longer, or – if already optimized for this arm but not yet for the other arm – repeated identical exposure times. We offer all of these spectra and leave the final judgement and selection to the user who should keep in mind that even heavily saturated¹⁹ telluric star spectra exist in the archive. On the other hand many of the optimally exposed telluric standard star data have a very high SNR and offer the full scientific potential of exquisite spectra of early-type, or solar-type, stars.

Table 4. QC flags (for data taken before December 2025).

¹⁸ For a point-like source, the actual resolution is between the value corresponding to the seeing disk, exposure-time integrated, and the value corresponding to the slit width.

¹⁹ In certain regions only, not everywhere.

Bit	Content (0 if YES, otherwise 1)	Motivation	Header keys (HIERARCH ESO ...)
#1 - seeing	seeing < slit width?	flux losses if seeing disk truncated	QC.SEEING = mean of TEL.FWHM.START and TEL.FWHM.END
#2 - airmass	below threshold value 1.8?	higher probability of flux losses for higher airmass (if the ADC did not work properly)	QC.AIRM.MEAN = mean of TEL.AMBI.START and TEL.AMBI.END
#3 - flux calibration	flux fully registered?	'yes' only if 5" slit used, otherwise always 'no'	QC.SLITW
#4 - mean reduced counts (<i>not flux</i>)	≥ 0.1	If very small or negative \rightarrow extraction went wrong, or just indicates low signal	QC.COUNTS.MEAN = mean of column FLUX_REDUCED
#5 - mean SNR	≥ 0	If negative \rightarrow extraction went wrong	QC.SNR.MEAN = mean of column SNR
#6 - saturated pixels	number of saturated pixels below 1000 (UVB, VIS) or 20,000 (NIR)	quality issue if too many saturated pixels	QC.NPIXSAT
#7 - aborted template	TPL.EXPNO = TPL.NEXP?	if not, data reduction is likely suffering, e.g. from upcoming clouds during observation	TPL.EXPNO, TPL.NEXP

QC flag. For data taken before December 2025, discontinued afterwards. The header key "QCFLAG" in the XSHOOTER products describes automatically assigned quality flags. It is composed of seven binary bits (

Table 4). For each bit, the value 0 means "no concern".

The final QC flag is composed of all binary values. For instance, 0000000 is a pristine product, 1011100 is a product that might have some slit losses, and also has negative flux indicating a localization or extraction issue. The seeing check #1 is a rather coarse quality check and meant to be indicative only (the two header keys TEL.FWHM.START and ...END refer to the first raw file). The same is true for the airmass check (#2); the underlying values are also from the first raw file. Taken together, they indicate mainly how reliable the flux scale is. With #1 and/or #2 being 1, it becomes likely that the flux scale for spectra from the different arms may show jumps. Only for #3 being 0 the flux calibration is likely to be free from slit losses, but still not necessarily be of photometric quality since the transparency is not strictly controlled for XSHOOTER observations.

Flags #4 and #5 assess the extraction quality. They are not entirely independent, but not redundant either. Their underlying values are averaged across the whole spectral range. Flag #6 indicates possible saturation issues, while #7 is a technical indicator of template abortion, which might alert the user to unusual observing conditions that might explain abnormal product properties. Template abortion almost always indicates degradation of observational conditions, or instrumental problems, and results often in data sequences which are incomplete at least for good automatic processing, or violate quality assumptions of the pipeline. In case there were no quality issues discovered other than flag #7 being set, we have delivered the product to the archive. If flag

#7 was found to be set, and another quality issue out of #4, #5 or #6 was found, we have decided to not deliver the product since these products were always found to be useless.

Science products. There has been some internal quality control on the pipeline processing of the science data, monitoring:

- the quality of the associations of calibration data (checking that the master calibrations used are not more than a few days away from the science data);
- on-demand QC reports and quick-look overviews.

This information has largely been used to improve and fine-tune the reduction process. An individual one-by-one inspection of the products has not been done.

Quality flags. The XSHOOTER pipeline supports pixel-based quality flags (not to be confused with the QC flag that applies to the whole file), propagated through the entire calibration and reduction procedure. They come as co-added bit codes and are available per wavelength bin in the spectrum table as column “QUAL” (see Figure 6 below). The possible values for the bit codes are defined in the User Manual (sect. 11.3). Some important values are listed here:

Table 5. List of important quality flags

$2^{12} = 4096$	Pixel saturation (UVB or VIS)
$2^{19} = 524288$	Missing data
$2^{20} = 1048576$	Extrapolated flux in NIR
$2^{21} = 2097152$	Raw pixel value zero or negative (saturation in NIR)
$2^{22} = 4194304$	Interpolated flux during standard extraction




Figure 6. Quality flags in the 2D product of the NIR arm. The colour map has been chosen to display bits 2^{20} (extrapolated flux) and 2^{21} (raw pixel zero or negative, white patches).

The logarithm of the quality flags is plotted in blue as lowest panel in the preview plots, labeled “log Q”²⁰.

Previews. Quick look plots are delivered as ancillary files along with the main products. There are two plots:

1. the main preview plot (see Figure 7 and Figure 8), one per arm;
2. the overview plot (Figure 9), only for the VIS arm (not available for data taken from December 2025 onwards).

Process quality control. The quality of the data reduction is monitored with quality control (QC) parameters.

QC parameters are used to monitor the reduction quality. The most obvious check is the SNR versus signal control plot (the signal being expressed as the mean across the entire reduced spectrum, before flux calibration, and the SNR being the mean over the entire spectral range).

Figure 10 shows an example for NODDING, VIS arm. The blue crosses represent the low gain mode

²⁰ With bit 2^{22} being suppressed in the plots (but not in the data) in order to avoid swamping the plots.

which is often used to obtain very high SNR of bright targets by combining many input spectra. For that mode, the record SNR found so far is in excess of 1000, quite remarkable for a product that has been automatically pipeline-processed without any fine-tuning. The red filled circles represent the “workhorse” high-gain mode, which is able to deliver SNR of up to 300 or more. The spread in the achieved SNR for a given value of *mean_reduced* is likely to come from the spectral slope of the products. This is illustrated by the next figure which displays these parameters for three selected observing runs. Their SNR slopes are well confined, due to their targets having roughly the same spectral slopes.




Figure 7. Main preview plot for observations before December 2025, featuring: 1. from top to bottom the flux-calibrated spectrum (column FLUX, both at original resolution in black and as smoothed version in red); the unfluxed spectrum (column FLUX_REDUCED); the SNR; log Q where Q is the pixel quality flags; 2. small spectral windows at full resolution; 3. a set of related QC parameters, bottom; 4. labels for identification (top).

XSHOOTER science product preview
 Product: SCI_SLIT_FLUX_IDP_UVB
 Raw file: XSHOO.2023-12-17T04:20:36.644.fits
 MJD-OBS: 60295.180979682

Target: WDJ0529-3401
 OB ID: 3680546
 OB NAME: WDJ0529-3401
 TPL ID: XSHOOTER_slit_obs_AutoNodOnSlit
 RUN ID: 109.2383.012

PRO TECH: ECHELLE,SLIT,NODDING
 Arm: UVB
 Slit: 1.3x11
 Binning: 2 x 2




Figure 8. Preview plot for observations taken from December 2025 onwards. Top: the flux-calibrated spectrum, at original resolution in blue and averaged in orange. Middle: the signal-to-noise spectrum. Bottom: three spectral windows of the flux-calibrated spectrum.




Figure 9. Second overview plot, not available for data taken from December 2025 onwards, featuring the (smoothed) signal in all three arms (1), and the collapsed signal in cross-dispersion direction, as derived from the 2D product (2). This plot can be used to check the flux scale (effect of slit losses) and the cross-dispersion profile. It also has the slit widths labelled.




Figure 10. Mean SNR vs. mean reduced spectrum, for all NODDING products of the VIS arm (a total of 5214 data points between 2009-10-01 and 2013-12-31). We distinguish between data for the high-gain mode (the “standard” mode), the low-gain mode (in both cases all binnings and slits), and high-gain mode, 1x1 bin, 0.9x11 slit (the most frequently used setting). The data plotted in this figure are accessible under the URL given below.




Figure 11. Same as Figure 10, for three selected runs (again NODDING products from the VIS arm only). Runs #2 and #3 have delivered data for young stellar objects (panel C), while run #1 belongs to panel D (stellar evolution).

Known issues and features

Issues

Localization and extraction. Sometimes STARE data might have the object off-center, with the effect that sky and target get confused by the pipeline which uses pre-defined localization and extraction windows (see Figure 12). The spectrum then has negative values²¹. The method used for localization of STARE targets is MANUAL which is the most robust one, but sometimes the wrong choice. Please check the User Manual.




Figure 12. Cross-dispersion profile (all examples are cut out from the second overview plot), collapsed in X direction. Both profiles come from the STARE mode. Left: target is properly centred, no issue with localization and extraction. Right: Source is off-center (center being at 0), pipeline creates a negative spectrum by confusing sky and target. The difference in FWHM is caused by the seeing.

Incorrect STARE mode: before the MAPPING template was introduced (December 2015), the STARE mode was sometimes used for effectively the same arrangement of scans across extended sources. These data cannot easily be identified. Their products are very likely of poor quality because of poor background subtraction and wrong extraction strategy. The user should check suspicious STARE products for their cross-dispersion profile and for other products from the same OB: if they exist, and if they show different sources at varying positions, the products are mapping data and should be ignored.

For NODDING and OFFSET data the extraction method assumes a well-centered point source. If true, the cross-dispersion profile is well-defined as a single positive signal and two correlated negative signals (Figure 13a, b). Quite some spectra actually contain complex sources (multiple targets, complex background, extended sources, Figure 13c). This is one of the main sources for abnormal spectra (negative flux, negative SNR). *The pipeline does not recognize these situations, and the user should always carefully check the 2D spectra in case of doubts.*

On the other hand the pipeline can safely extract spectra without continuum (from an emission-line object), with a cross-dispersion profile in Figure 13d.

The extraction issue with NODDING data having a nod throw of less than 5 arcsec has been fixed (at least for the data being strongly affected, the ones with nod throw ≤ 3.5 arcsec; see 'Data reduction and calibration'). Users should replace their downloaded spectra if processed with pipeline versions earlier than `xshoo/2.7.0`. The value of the nod throw is read from the key "HIERARCH ESO SEQ NOD THROW".




Figure 13. Cross-dispersion profile in NODDING mode. From left to right: a) the target is properly centred and well-behaved. The two negative signals come from the nodding pattern. b) The target is properly centred and well-behaved. The negative peaks appear double reflecting a jittered NODDING pattern. c) Bad profile shape,

²¹ The extracted spectrum might still be useful after correcting for the wrong sign.

useless extraction. Possible reasons are: complex background, invalid NODDING pattern, extended/multiple sources. d) No continuum signal, but extraction looks ok and spectrum likely to be valid (from an emission-line object).




Figure 14. Impact of bad target positioning in the slit. The pipeline produces negative net signals (visible in all 3 arms, plot #1) because of bad centring (cross-dispersion profile in plot #2).




Figure 15. The same target, taken immediately after the acquisition error was noted by the nighttime staff, the target has been measured again, this time properly. The target is a telluric standard star, but it is representative for alignment issues of any target.

In Figure 14 and Figure 15 we illustrate the effect of bad centring of a point source. The example is a telluric standard star, but the argument applies to a science target in the same way.




Figure 16. Cross-dispersion profile in OFFSET mode. a: proper behaviour (object and sky well-defined). b: target off-centre, the extraction still works fine. c: odd behaviour (object and/or sky not well defined).

OFFSET data depend critically on the observer's definition of the Observing Block. The pipeline blindly trusts the labels OBJECT and SKY in the raw data. The user should always check carefully if

the OFFSET products have a cross-dispersion profile like in Figure 16 a, b (well behaved) or in c (ill-defined). Remember that OFFSET data without OBJECT, or without SKY, are suppressed.

It is not possible to tell from the header if the source is point-like or not. Furthermore it is impossible to check whether the OBJECT and SKY labels are correctly set by the PI. *If in doubt, check the 2D ancillary files and the QC plot for the VIS arm displaying the cross-dispersion profile.*

Saturation and gaps. Saturated (or more generally, bad) pixels in the raw data get ignored by the pipeline. If there are too many, so that no good pixels can be found by the pipeline for a given wavelength bin, the flux in this wavelength range is set to zero. In the QUAL extension of the 2D ancillary spectra, these regions can be easily spotted: dark pixels have bit code 512, pixels affected by cosmics have bit code 16 or 32; bit code 4096 indicates saturated pixels²². In the extracted main output file the QUAL column has the bit-code $2^{19}=524288$ (“missing data”) (see Figure 17 Figure 17). Outside of these gaps, the data might still be useful.




Figure 17. Spectral gaps due to saturation. They are marked by the bit code 2^{19} (blue in lowest panel of the QC plot).

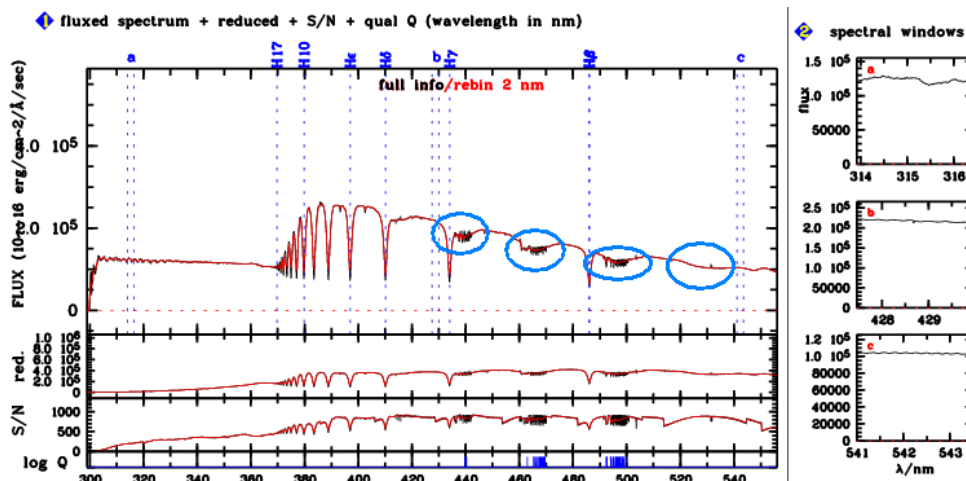


Figure 18. Ripples and bumps as numerical artefacts in very high-SNR spectra, typically spectra of telluric standard stars. The blue ellipses mark regions with some saturated or nearly saturated pixels.

²² Find the complete list of bit codes in the Pipeline User Manual.




Figure 19. Artefacts in the first orders of VIS spectra.

Partial saturation. If spectral regions have some saturated pixels but are not entirely saturated, patterns might occur as illustrated in Figure 18 (blue circles). These ripples tend to display at regular intervals of about 30 nm in the UVB regions longer than about 400 nm.

Spectral slopes in the VIS arm. Another artefact typical of high-SNR (telluric) spectra is displayed in Figure 19 (red circle). These bumps are caused by uncontrolled variations in the energy distribution of the flat-field lamp. If existing they display in the first two spectral orders, and the continuum slope in that region is incorrect. Often they do not show up in the :FLUX_REDUCED spectrum. More in the Reflex tutorial sect 7.5, “Instrument Response”.

Spikes in NIR spectra. Lots of artificial spikes exist in many NIR spectra. They are due to a pipeline issue, namely an insufficient handling of the bad pixels at the steps of re-sampling and frame-stacking. This pipeline issue will be fixed in the future.

APERTURE key filled wrongly. Due to a software bug, the APERTURE key in the products was filled with the slit length (always 11 arcsec) instead of the slit width, if processed with pipeline version xshoo/2.6.0 or earlier.

TUCD1 key filled wrongly. The XSHOOTER wavelength scale refers to air wavelength. The TUCD1 key was filled incorrectly with “em.wl” which stands for vacuum wavelengths, for all data downloaded before about 2017-08. The content of TUCD1 has now been corrected to “em.wl;obs.atmos”, and all data downloaded in 2017-08 and later are fine.

Negative fluxes. In the NIR, negative fluxes might appear in raw data and are the result of saturation (Figure 20). Affected spectral pixels are flagged by bit $2^{21} = 2097152$, and the impact on the spectrum are troughs embedded between peaks (unsaturated regions). Also, it might happen in rare cases that the reset anomaly causes negative pixels which then are perfectly valid, except for the wrong sign.




Figure 20. Spectral regions affected by negative (saturated) raw pixels in NIR spectra. The affected regions are visible as artificial depressions and flagged blue in the lowest panel of the QC plot.

Periods of failed ADCs. In the nights from 2022-07-03 to 2022-07-11 the ADCs failed and the atmospheric dispersion was not corrected for. The UVB-arm response function acquired in this period is subject to flux losses and UVB-arm observations cannot be flux-calibrated. As a consequence the observations from these nights are not covered by the current data release.

NIR-arm Wavelength Calibration Issue: It is found, that the NIR-arm suffers from a wavelength calibration issue. The issue is confined to the spectral order edges of the three most red orders at: 2458 nm, 2273 nm, 2075 nm and 1925 nm and can amount to up to three wavelengths bins or 0.18 nm at 2273 nm. One of the reasons is that there are no or not enough arc Th-Ar lines in the calibration spectra to constrain the wavelength solution in the order-overlapping regions. This effect causes a problem with the telluric correction. The UVB and VIS-arm spectra are not affected.

Improved bad pixel handling: The new default value of decode-bp is set to 2140143615=(2³¹-1)-1048576-2097152-4194304. In other words, all pixel flags are considered bad except for those pixels flagged as having an “extrapolated flux”, zero or negative values in the raw data, or are interpolated during the standard extraction of the one-dimensional spectrum.

IDPs **from 2023-04-01 on** are processed using improved recipe command line parameters: The science and telluric standard star recipes for STARE and OFFSET are called with `--decode-bp=2140143615`, while those for NODDING are called with `--decode-bp=1737490431`. The science and telluric standard star recipes for UVB and VIS are called with `--removecrhsingle-niter=4`, while those for the NIR arm are called with `--removecrhsingle-niter=0`. The new default for any recipe is `--removecrhsingle-sigalim=5`.

Features

Pick noise. At certain times (e.g. during most of 2010-03) the VIS arm detector showed a strong pick noise pattern that the pipeline cannot remove from the data. The impact is a high-frequency noise that is easily spotted in the spectral windows plot (2). See an example in Figure 21.

Steps at begin of order. Occasionally there are a few steps at the begin of an order, apparently induced by an imperfect flat-fielding (Figure 22).

Ozone absorption bands. In the extreme blue part of the UVB spectra (shorter than 320 nm) there are the ozone (O₃) absorption bands visible in very high-SNR spectra (Figure 23).




Figure 21. Pick noise in VIS data, which stands out prominently in the close-up plots.




Figure 22. Steps at begin of an order. In this example the feature at 360nm is rather pronounced. Such features appear at the begin of an order (the orders are best visible in the S/N panel).




Figure 23. Ozone absorption bands in the extreme UVB.

2D extracted spectrum not flux-calibrated. The ancillary 2D file (product category is SCI_SLIT_MERGE2D_<arm>) is not flux-calibrated although such a product would be useful and could be created by the pipeline. If users wish to have a flux-calibrated version of the 2D file, they should calculate the ratio of the columns :FLUX (the flux-calibrated 1D signal) and :FLUX_REDUCED (the uncalibrated 1D signal). That ratio function should then be expanded into a 2D file and get multiplied into the SCI_SLIT_MERGE2D file, resulting in a flux-calibrated 2D spectrum.

Intensity jump at 2.25 μ . In the reddest part of the NIR arm data, the sky background might have a gradient along the slit. Then the assumption by the pipeline of a flat sky background is no longer valid, which results in a sudden jump of the signal which is however an artefact and entirely due to this pipeline effect (Figure 24). Best is to ignore data beyond 2.25 μ if this step occurs. This feature does not occur at high S/N, nor for spectra acquired in NODDING mode. It might also occur as a downward jump.

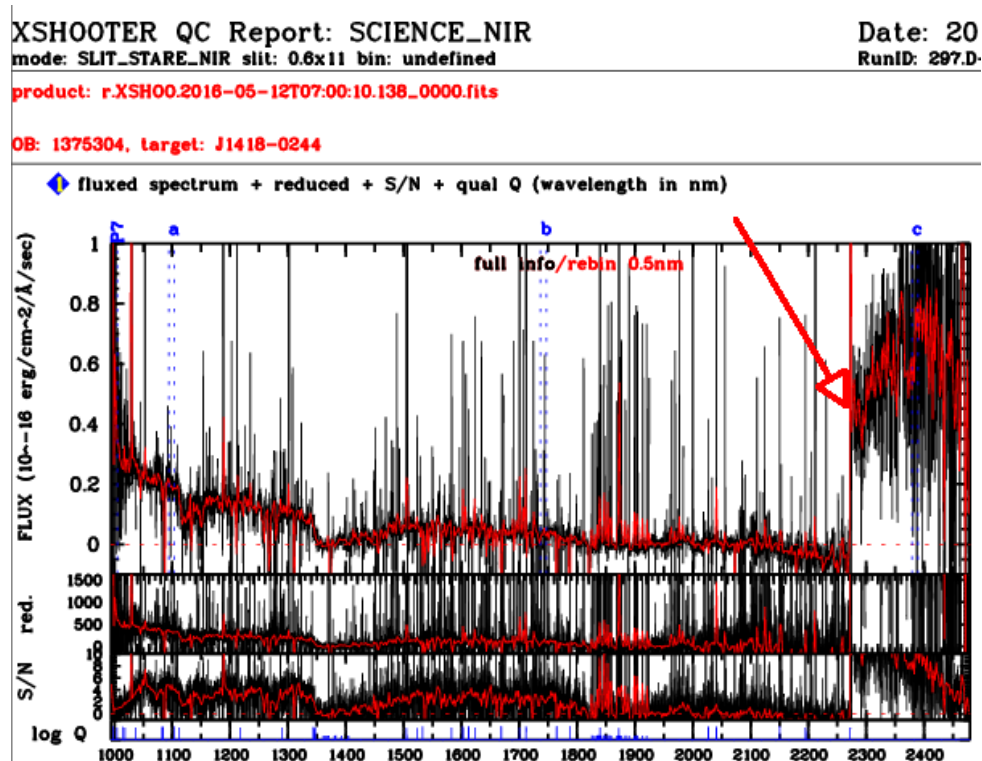


Figure 24. Intensity jump at 2.25 μ for STARE mode. For the last echelle order, the sky cannot be properly subtracted by the pipeline, and the signal is mostly due to the sky.

Slit losses. With the requirement that telluric standard stars should match their “parent” science data in airmass (within 0.2), there are occasionally telluric standard star spectra taken at high airmass but without careful alignment taking into account atmospheric dispersion effects. A particularly strong example of the resulting slit losses is displayed in Figure 25.

products: r.XSH00.2015-01-06T05:24:04.907_tpl_0019.fits + UVB + NIR (from telluric std star)

OB: 200193515, target: Hip020600

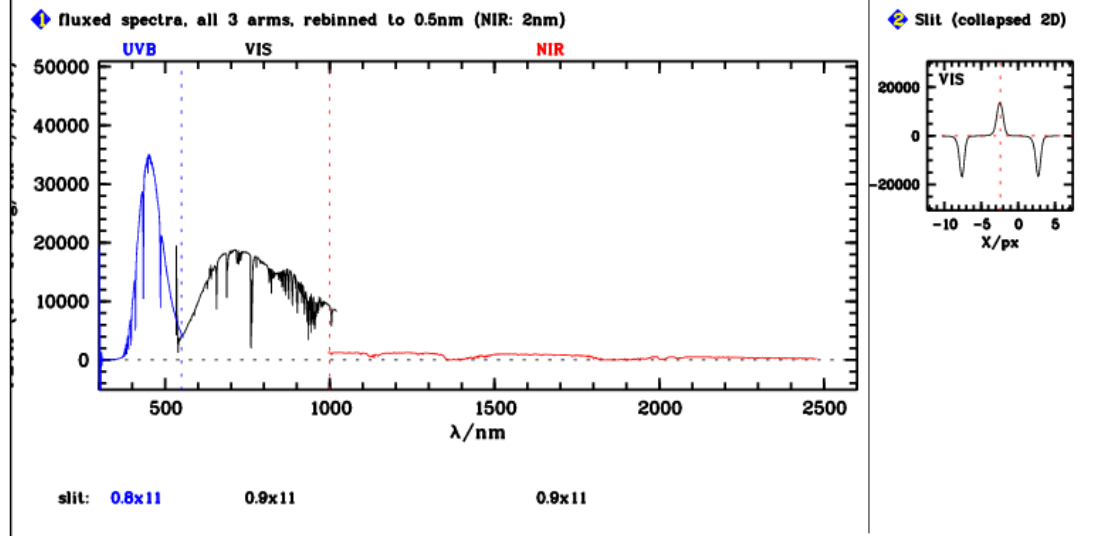


Figure 25. Slit losses due to high-airmass observations (2.37 in this case) of a telluric standard star without atmospheric dispersion correction. The blue parts of the UVB and the VIS spectra are strongly affected.

Data Format

Files Types

The primary XSHOOTER Echelle product is the flux-calibrated spectrum, in binary spectroscopic data format:

product category HIERARCH.ESO.PRO.CATG and PRODCATG	Description
SCI_SLIT_FLUX_IDP_<arm> SCIENCE.SPECTRUM	flux-calibrated product (arm=UVB, VIS or NIR) from science observations
TELL_SLIT_FLUX_IDP_<arm> SCIENCE.SPECTRUM	flux-calibrated product (arm=UVB, VIS or NIR) from a telluric standard star

The primary product has an ancillary FITS file delivered with it (useful for quality assessment; in 2D image format):

SCI_SLIT_MERGE2D_<arm> ANCILLARY.2DSPECTRUM	Order-merged 2D product (science observations) before extraction and flux calibration (arm=UVB, VIS or NIR)
TELL_SLIT_MERGE2D_<arm> ANCILLARY.2DSPECTRUM	Order-merged 2D product (telluric standard) before extraction and flux calibration (arm=UVB, VIS or NIR)

Naming convention before December 2025.

The following naming convention applies to the ORIGFILE product name for data taken before December 2025 (the ancillary FITS file has the same name, with XS_SRE2 replacing XS_SFLX): e.g. XS_SFLX_922346_2013-08-09T23:33:54.611_S1.5x11_1x1_VIS_NOD.fits has the components:

ORIGFILE component	XS	SFLX	922346	2013-08-09T23:33:54.611	S1.5x11_1x1_VIS_NOD.fits
refers to ...	XSHOOTER	product type (SFLX where S stands for science and FLX for fluxed, or RE2 for reduced, 2D)	OB ID	timestamp of first raw file archival	setup string: S1.5x11 for the slit; 1x1 for the binning (dropped in NIR); VIS for the arm; NOD for the technique.

The same scheme is also used for the telluric standard star spectra. Their OB ID is almost always 200193515, i.e. purely technical. Their RUN_ID is 60.A-9022(C), also purely technical. Both values mark these data as being provided by the Observatory.

The ancillary files (always delivered with the primary 1D spectrum) have the following ORIGFILE names:

Table 6. Naming conventions of ANCILLARY files

type	example	rule
ANCILLARY.2DSPECTRUM	XS_SRE2_998129_2013-09-26T08:49:39.394_S5.0x11_1x1_UVB_STA.fits	Same as primary spectrum, starting with XS_SRE2 (for science, reduced 2D)
ANCILLARY.PREVIEW	r.XSH00.2013-09-26T07:56:43.243_tpl_0000.fits_1.png	Technical filename of primary spectrum, with _1.png appended; the VIS products have also a second preview file,

		called _2.png
ANCILLARY.README	XS_SFLX_998167_2013-09-26T07:56:46.364_S0.6x11_NIR_NOD.txt	Same as primary spectrum, with fits replaced by txt (only for data taken before December 2025)

The user may want to read the ORIGFILE header key and rename the archive-delivered FITS files.

Naming convention from December 2025 onwards.

The ORIGFILE product name for observation from December 2025 onwards follows a naming convention which is

XSHOOTER_<PRO CATG>_<DATE-OBS>.fits

Using the header keywords HIERARCH ESO PRO CATG and DATE-OBS. The value of DATE-OBS can deviate by a millisecond from the time stamp given in the PROV1 header keyword. The preview files follow the same naming convention with the file type 'fits' replaced by 'png'.

File structure

The primary XSHOOTER product XS_SFLX comes as binary FITS table in multi-column format. The columns are labeled as follows:

Table 7. Internal structure of the XSHOOTER spectrum.

column	label	content
#1	WAVE	wavelength in nm
#2	FLUX	extracted, wavelength-calibrated, sky-subtracted, flux-calibrated SCIENCE signal in physical units [$\text{erg s}^{-1} \text{cm}^{-2} \text{\AA}^{-1}$]
#3	ERR	Error of FLUX (same units)
#4	QUAL	Quality flag
#5	SNR	Signal-to-noise ratio = FLUX/ERR
#6	FLUX_REDUCED	extracted, wavelength-calibrated, sky-subtracted, but not fluxed SCIENCE signal in counts
#7	ERR_REDUCED	corresponding error (not fluxed)

The difference between FLUX and FLUX_REDUCED is the flux calibration. The FLUX_REDUCED data might become useful if the user wishes to apply his own response curve instead of the master response curve. The SNR column is provided for convenience. The QUAL column contains the quality flag per wavelength bin, propagated throughout the reduction.

File Size

The (unbinned) XSHOOTER products have about 1.5 MB size. The ancillary 2D files have the same spectral coverage and come as 15-29 MB files (depending on arm). Files are always uncompressed.

Acknowledgment text

According to the ESO data access policy, all users of ESO data are required to acknowledge the source of the data with an appropriate citation in their publications. Since processed data downloaded from the ESO Archive are assigned a Digital Object Identifier (DOI), the following statement must be included in any publications making use of them:

Based on data obtained from the ESO Science Archive Facility with DOI(s):

<https://doi.eso.org/10.18727/archive/71>

Science data products from the ESO archive may be distributed by third parties, and disseminated via other services, according to the terms of the [Creative Commons Attribution 4.0 International license](#). Credit to the ESO provenance of the data must be acknowledged, and the file headers preserved.