## ESO Phase 3 Data Release Description


| | |
|---|---|
| **Data Collection** | MUSE-DEEP |
| **Data Provider** | ESO, Quality Control Group |
| **Document Date** | 2023-01-03 |
| **Document version** | **1.5.2** |
| **Document Author** | Reinhard Hanuschik/Danuta Dobrzycka |

*Changes with respect to the earlier version:*
*- narrow-field mode NFM-AO added*
*- exposure times shorter than 10sec processed normally*
*- EXPOSURE_MAP added as of 2019-11-20 (ARCFILE)*


# Abstract


This is the release of reduced deep IFU datacubes from the MUSE[1] spectrograph, taken in the Wide Field Mode. MUSE, the **M**ulti-**U**nit **S**pectroscopic **E**xplorer, is an Integral Field Spectrograph located at the VLT UT4 telescope. The instrument samples the sky with 0.2 arcseconds spatial pixels in the currently offered Wide Field Mode with natural seeing (WFM-NOAO), and, since 2017, also assisted by the UT4 AO system GALACSI (WFM-AO). The Narrow-Field mode (NFM) has a FOV of 7.4"x7.4" and samples with 0.025 arcseconds spatial pixels. It is offered since P103 and is supported by laser tomography, as NFM-AO.

*Each deep datacube is combined from observations across OBs*[2]. Where multiple visits of the same target exist, with multiple OBs, the deep datacube combines the input files from these OBs with the goal to reach the maximum possible depth of the observations. There are also many OBs that visit a given target only once, and then no deep datacube exists. Therefore, the MUSE and the MUSE-DEEP releases are generally complementary. We have successfully combined deep datacubes with more than 120 input files. The deepest datacubes represent a total integration time of more than 30 hrs.

This release is an open stream release. The release covers the two MUSE Science Verification periods in June and August 2014, and data from the regular MUSE operations which started in September 2014. Data from the AO Science Verification period in August and September 2017 are also included. Depending on the availability of an end-of-run signal, new data are processed within a month or two after that signal, or with a larger delay in some cases.

The data have been reduced with the MUSE pipeline, version `muse-1.6.1` and higher[3]. The data reduction has two steps: removal of instrument signature, and combination of all products from that step into the deep datacube. Resampling has been done once, at the latest step. Error propagation is the same as for the OB datacubes. Sky correction is also the same, except for the case of crowded fields (globular clusters) where no sky correction is applied.

The Quality Control Group at ESO processes the data in an automated process. In an initial step there is an interactive selection of programmes and candidate targets. Then, each observation is pipeline-processed with time-matching, quality-controlled, certified and archived master calibrations. The reduction process is largely automatic. There is an automatic scoring process for the


[1] http://www.eso.org/sci/facilities/paranal/instruments/muse.html
[2] OB = Observing Block, a single pointing on the sky and the fundamental unit of the VLT observations. It can hold one or more object observations. Its maximum length is limited to one hour for operational reasons. If an observer needs to go deeper, he has to prepare and execute more than one OB. Combining these multiple OBs (visits) in a single data product (the deep datacube) is the goal of these observations and of this data release. The products of the individual OBs are the combined MUSE datacubes, also called OB datacubes here. They are available in a separate release, called MUSE.
[3] Weilbacher P.M., Streicher O., Palsa R.: "MUSE-DRP: MUSE Data Reduction Pipeline", http://ascl.net/1610.004 and Weilbacher P.M. et al. 2012SPIE.8451E..0BW

quality control, and a semi-automatic review and certification process for the data products, focusing on non-zero scores.

The data format follows the ESO science data products standard for datacubes[4] and is the same as for the OB datacubes.

This data release offers data products which are considered to be ready for scientific analysis, i.e. with instrument and atmospheric signatures removed, calibrated in physical units and including error estimates.

**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 purpose of the observations, nor for the scientific goal of the archive user. There might be cases where the selection of input data was not optimal to reach e.g. the highest possible spatial resolution.

This release description describes the specific aspects of the MUSE-DEEP processing, while the aspects common with the MUSE release are mentioned only briefly for conciseness. Their details can be found in the MUSE release description.

# Release Content

This release is a stream release. The data are tagged "MUSE-DEEP" in the ESO archive user interface[5].

The release starts with the two MUSE Science Verification periods in June and August 2014, and includes data from the regular MUSE operations which started in September 2014. When a signal is available that a run has been finished, new data are processed and added a month or two after that signal. If no such signal is available[6], the delay can be half a year, and even longer for Large Programmes when it is not obvious that the collection of data for a given target is finished. Datacubes for the NFM-AO mode have been added as of April 2019 (P103).

Although we try to be as careful as possible with the selection of completed datasets, rare cases might occur where data collection continues after our deep datacube has been processed and archived. In that case we replace the previous version by a newer deeper version, with the older version still being available on demand.

The names of all input raw files are recorded in the header of the corresponding data product[7].

The purpose of the deep combination is the **maximized signal contrast (SNR)**, with 2 related aspects:

- every source spectrum has a better SNR in a deep cube than in any individual cube,
- it is possible to detect fainter sources in a deep cube than in any individual cube.

In most cases, multiple visits of the same target have been designed by multiple OBs within the same programme, with the PI-intended goal to reach the maximum depth of the observations. In a few cases, we have found multiple visits designed by *different programmes*. While many of them are still designed by the same PI (in different periods) and represent the same *logical programme*, some of them are coming from different programmes and different PIs. We have decided to combine these "multi-PI" OBs in a single deep datacube. In these cases the data product might go even deeper than intended by the respective PIs. (We cannot guarantee to have discovered all of these cases.)

[4] Find it under http://www.eso.org/sci/observing/phase3.html under 'Quick links', 'ESO SDP standard'. The section 'Integral Field Spectroscopy: 3D Data Cubes' defines the structure of the MUSE data products.
[5] http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form
[6] This is particularly true for GTO data, carry-over runs and large programmes covering several periods.
[7] Header keywords PROV*i*.

## Data Selection

All input data qualifying for the MUSE processing were reviewed for the MUSE-DEEP project.

Mode and setting selection is the same as for MUSE:

- *instrument mode (INS.MODE) = WFM-NOAO-{E or N} and WFM-AO-{E or N}* and NFM-AO-N (there is no E for NFM-AO);
- *N* = 'nominal wavelength range' (480-930 nm), or *E* = 'extended' (465-930 nm).

The WFM-AO mode has a gap without signal between about 580 and 596 nm (N range), and 576 and 601 nm (E range), respectively, due to laser-induced sodium lines. If existing, we have also co-added data taken in AO and NOAO modes. This is justified because in general these programmes are designed to have matching seeing constraints. Note that in those cases there is no spectral gap (it is filled with NOAO data) but the SNR is lower across this range, and there might be steps in the spectral fluxes because of the different number of combined spectra within and outside the sodium range.

We used the following information sources for the candidate selection:

- programme titles and abstracts (scanning in particular for 'deep'),
- QC reports,
- target names.

The programme scan helped to identify the qualifying runs. We found that for the first year of MUSE operations about 50% of all programmes were advertised as going deep. By selecting all QC reports for those programmes (or runs) and sorting them by target name, we were able to safely identify all multi-OBs. In case of non-unique target names, or in complex situations where the targets were larger than the 1'x1' field of view of MUSE, the previews from the QC reports were used for a final decision.

**Applied guidelines for the selection:**

**1. Seeing.** Many combination candidates were taken in Visitor Mode (VM), in GTO time. Then, no OB grades are available, and the final selection of input files was based on an assessment of the measured seeing conditions. The rejection criteria we applied were relaxed, only strong deviations (i.e. by a factor 2 or so) from the requested conditions were used for rejecting input candidates[8]. In Service Mode, we effectively applied the same criteria. Often we accepted OBs graded C, if that grade was only due to a mild violation of the seeing constraint. If there were other problems with the data, as documented in the OB comments, these were taken into account (if found applicable).

**2. Photometry.** For the deep combination, input data with varying photometric conditions (PH, CL, TN, TK[9]) were accepted. Quite often the observations are not requiring photometric conditions. *Deep cubes are generally not suitable for precise photometry.* If required, more precise photometric information can be derived from single-OB datacubes taken under photometric conditions (PH).

**3. Cosmetics, in particular satellite trails.** We have rejected in a few extreme cases input files with strong satellite trails, but fainter ones were deemed acceptable since normally satellite trails (or generally transient sources) affect only a small portion of the FOV. Also, we did not notice any OB degraded because of satellite trails.

**4. Background.** We have trusted the scheduling decision at the telescope and have not rejected input candidates because of background criteria, with one exception: if the OB comment says "aborted due to increasing background", these data have been rejected for deep combination.

**5. Other issues.** On an individual basis we have rejected exposures with nightlog comments like "aborted because of derotator issue", unless it turned out that the data are ok.

**Previews from the MUSE processing.** In the process of target and OB selection, the information

[8] This strategy is consistent with what we found in some PI publications.

[9] PH-Photometric, CL-Clear, TN-Thin cirrus, TK-Thick cirrus

gathered with the OB-based combined datacubes from the MUSE project benefitted us a lot, so that we could apply our selection based on the full information of the FOV image and of the processing results.

**Combination by criteria other than by target.** In a few cases, the OB combination by target was inadequate, in particular for exposure time sequences, or if different pointings were collected in a single OB. These cases could be identified safely, and the final deep datacubes were then constructed using common pointings, and/or common exposure times.

**Products.** Any given input dataset (defined by target) consists of *N* OBJECT frames and *M* SKY frames, coming from at least 2 OBs. *N* must be at least 2, and its maximum value is about 130, due to the 2 TB memory available for processing[10]. *M* is often zero (many deep observations have no dedicated SKY pointings). The product is always 1 DEEP COMBINED datacube per target.

**Relation between MUSE and MUSE-DEEP releases.** For the runs which do not attempt to go deep, the COMBINED datacube in the MUSE release is the final product. Likewise, there are runs which have some targets with deep observations and others with a single visit. For those single OBs, the COMBINED datacube in the MUSE release is the final product. Of course, if there is a SINGLE datacube only (one exposure in one OB), this is the final product. Therefore, the MUSE and the MUSE-DEEP release together should both be queried for datacubes of a given target or a given run. Only if there is a MUSE-DEEP datacube, the corresponding OB-based MUSE datacubes are in principle obsolete for analysis, but may still be valuable for photometry, best-seeing analysis, multi-epoch variability studies and for cross-checks. See Table 1 for an overview of the types of MUSE datacubes.

If you need access to the single datacubes that participated in a combined datacube, there is a special download channel for them, as described in the MUSE release description.

Table 1. Types of MUSE and MUSE-DEEP datacubes

| **Product type** | **from input file** | **PRO.CATG** | **occurrence** |
|---|---|---|---|
| MUSE-DEEP: | | | |
| DEEP COMBINED | OBJECT | DATACUBE_DEEP | always |
| MUSE: | | | |
| COMBINED | OBJECT | DATACUBE_COMBINED | often |
| SINGLE | OBJECT | DATACUBE_SINGLE | rare |
| COMBINED | SKY | DATACUBE_SKY_COMB | rare |
| SINGLE | SKY | DATACUBE_SKY | often |

**Multiple run IDs.** Many MUSE programmes that go deep are split into different run IDs that need to be combined across periods. These data are unfortunately not marked by any metadata key to belong together. We have used several "fuzzy" criteria to identify them, e.g. common target names, OB naming schemes, programme titles, etc. The final confirmation was often only possible by the QC report of the FOV image.

# Release Notes

## Pipeline Description

Find the detailed description of the recipes in the Pipeline User Manual[11], section 9 (recipe reference). Find the pipeline version used for this processing in the header of the product datacube, under "HIERARCH ESO PRO REC1 PIPE ID". The version for the initial dataset was `muse_1.6.1`. Information about the MUSE pipeline (including downloads, manuals, cookbook) can also be found

[10] We have created a few datacubes with larger input datasets by wavelength stitching, see here.
[11] Under the MUSE link in http://www.eso.org/sci/software/pipelines/.

under that URL. The MUSE pipeline has been written mainly by Peter Weilbacher (see Weilbacher et al. 2012 http://adsabs.harvard.edu/abs/2012SPIE.8451E..0BW for a description, and Weilbacher et al. 2016 http://ascl.net/1610.004 for the code reference).

The QC pages[12] contain further information about the MUSE data, their reduction and the pipeline recipes for *calibration data*. Monitoring of MUSE performance and quality parameters is provided under the Health Check monitor[13].

## Data Reduction and Calibration

**Reduction steps, overview.** The data reduction uses a cascaded recipe scheme, with two main parts. It is the same for NOAO and AO data. AO data are reduced with the proper AO calibration data associated.

The first part works on individual input raw files. No combination is done at that stage. First, every input raw file (OBJECT or SKY) is pre-processed with the recipe `muse_scibasic`. Then, the SKY product files (if any) are further processed with the recipe `muse_create_sky` to create the SKY_LINES and SKY_CONTINUUM files for the later sky subtraction. The sky contribution is evaluated by considering the information on the instrument line spread function, which is contained in the LSF_PROFILES master calibration file. Next, the OBJECT product files are processed with the recipe `muse_scipost`, using the SKY products (if existing) for the sky subtraction[14].

After the `muse_scipost` step, all input OBJECT files have a PIXEL_TABLE product with the pixel coordinates stored in a table, and an IMAGE_FOV product (a 2D collapse) used for the alignment correction. *These products can be considered as being free (within known limitations) from instrumental artefacts. Therefore the next step is possible, the combination of data from potentially many OBs and different nights. This step aims at collecting as many signal photons as possible, while reducing the noise due to sky background and shot noise.* The pixel-table format guarantees that the signal from every single pixel is preserved and not compromised by numerical binning at an early step.

In the second part of the science cascade, all PIXEL_TABLEs which belong together (as defined by the initial target selection step) are combined. Two steps are necessary: first, the input IMAGE_FOVs are processed with `muse_exp_align` in order to measure the relative alignment of the input data, in order to detect and correct for possible alignment errors due to instrument wobble, see below. Then, finally, the input PIXEL_TABLEs are processed with `muse_exp_combine` which applies the alignment correction, and finally resamples the overlapping pixels in order to go deep. It is only at that last step that the input data are resampled. The output of that last step is the COMBINED DATACUBE called DATACUBE_DEEP, and the combined IMAGE_FOV_DEEP. Find the overview of the recipes in Table 2.

Table 2. Overview of MUSE-DEEP science reduction cascade.

| Recipe | Number in the figures | Applied to |
|---|---|---|
| `muse_scibasic` | 1 | single OBJECT or SKY |
| `muse_scipost` | 2 | single OBJECT |
| `muse_create_sky` | 2a | single SKY |
| `muse_exp_align` | 3 | whole input dataset |
| `muse_exp_combine` | 4 | whole input dataset |

**Reduction steps, details.** For the details about the reduction cascade we refer to the MUSE release description. We follow the same numbering scheme for easy reference, with annotations as required.

[12] http://www.eso.org/qc/MUSE/pipeline/pipe_gen.html
[13] http://www.eso.org/HC, select MUSE.
[14] Contrary to the MUSE project, the MUSE-DEEP release has no shallow datacubes based on SKY observations.

***Part 1, single pixel-table.***
1.1 `muse_scibasic`: same as for MUSE.

1.2 `muse_create_sky`: same as for MUSE.

1.3 `muse_scipost`: same as for MUSE, except for the last step which is:

- apply the astrometric solution.

There is no resampling into a single datacube (since this can never be a final product for MUSE-DEEP).

The pipeline parameters for this recipe are set to their default values, except for the following parameters:

- if no SKY observation is available, and if the processing method is not CROWDED (the standard case):
  `--skymethod=model` and `--skymodel_fraction=0.2`
- if no SKY observation is available, and if the processing method is CROWDED (an exceptional case):
  `--skymethod=none`
- if SKY is available:
  `--skymethod=subtract-model`

1.4 `muse_scipost` for SKY: not applied.

The processing method CROWDED has been implemented for the cases of crowded field observations (globular clusters) without SKY which are known in advance for MUSE-DEEP. The corresponding MUSE datacubes suffer from an over-subtraction of the SKY background which is determined on the OBJECT data, with the level of over-subtraction depending on the prevailing seeing. In this situation it seems a better strategy to not subtract sky at all. The data analysis of the final datacubes needs to be done with aperture photometry anyway.

***Part 2, combined datacube.***
2.1 In the second part the pipeline recipes work on the products (pixel-tables and FOV images) from all input files together. The recipe `muse_exp_align` is used to create a coordinate offset table for automatic exposure alignment. This step is particularly important for the deep processing since it corrects instrumental alignment errors which potentially are larger across OBs and across different nights than within a single OB.

In order to always have an alignment solution, the following parameters are used for WFM data:

- as default (if all FOV images align well, and if the processing method is not CROWDED):
  `--rsearch=5,3,2,0.8`
  `--threshold=10.`
  `--iterations=200000.`
  `--srcmax=120.`
- special case (if FOV images do not align well with the defaults):
  `--threshold` reduced to values lower than 10, until successful execution.
- special case (processing method CROWDED: needs more relaxed parameters because of the very high number of sources in the field):
  `--rsearch set to default`
  `--threshold=100.`
  `--iterations=20000.`
  `--srcmax=200`
  `--srcmin=2`
  `--step=5`

For the NFM-AO data we use `--srcmin=1` and `--srcmax=2 or 1.`

2.2 Finally the output OFFSET_LIST table from `muse_exp_align` and the pixel-tables are combined into the final combined datacube.

**Products.**
In Figure 1 and Figure 2, we show the entire processing scheme for the cases 'no SKY' and 'SKY'. In these figures, we use the following numbering scheme for the recipes:

| | |
|---|---|
| `muse_scibasic` | 1 |
| `muse_scipost` | 2 |
| `muse_create_sky` | 2a |
| `muse_exp_align` | 3 |
| `muse_exp_combine` | 4 |

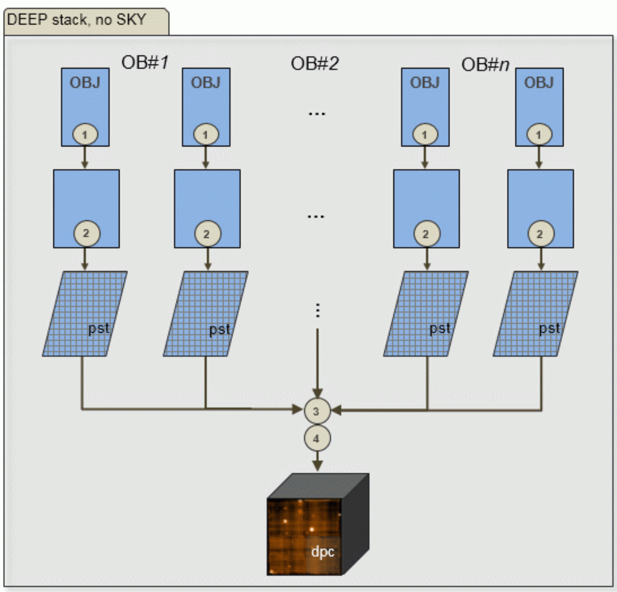


Figure 1. Reduction cascade for *N* input files from n OBs, no SKY. 'pst' marks the pixel-table products of `muse_scipost`. The final product is the deep combined datacube (dpc).

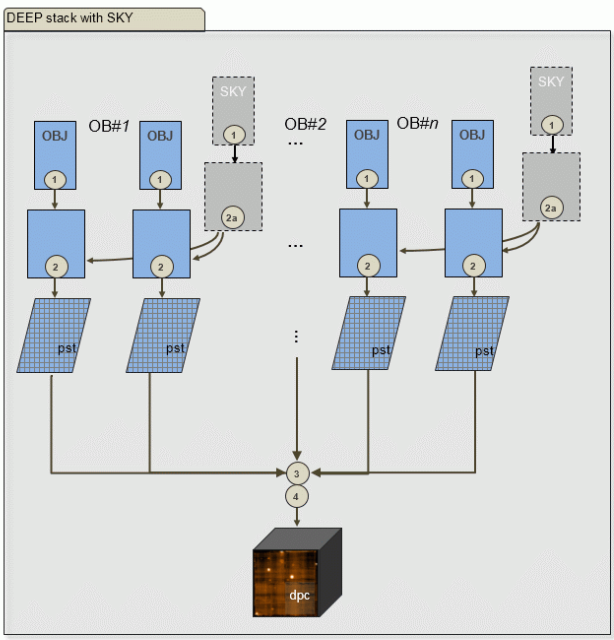


Figure 2. Same as above, for the case of *N* input OBJECT files and *M* SKY files.

The pipeline log files for all steps are stored in the text file that is delivered with each datacube.

While that information is technical, it might help with the understanding of the individual steps and might also serve as reference in case a user wants to redo certain reduction steps.

**Master Calibrations used for data reduction.** This is identical for MUSE and MUSE-DEEP. Check the MUSE release description.

**Wavelength scale.** The MUSE IFU products are wavelength calibrated. The wavelength scale is *barycentric*.

**Telluric absorption.** Telluric absorption lines have been corrected file by file and night by night with the STD_TELLURIC file that was derived from a standard star observation (the same as for the photometric calibration). The other comments in the MUSE release description apply here as well.

For the deep combination, it is not unusual to include observations from a considerable time span (90 days or more). The residuals of the corresponding telluric systems then do not overlap exactly in the barycentric rest frame, which might result in an additional broadening corresponding to +/-30 km/s at most.

**Flux calibration.** All comments in the MUSE release description apply for MUSE-DEEP as well. For the DEEP combination scheme, it is clear that the goal is to optimize the SNR, while a photometric accuracy cannot be guaranteed. The quality of the photometry in a COMBINED datacube (if observed under photometric conditions) is likely better than in a DEEP datacube, and should therefore be retrieved from there. *We have not suppressed any input file simply because of poor photometry.*

**Master calibration names and recipe parameters used for reduction.** Check the MUSE release description.

## Data format and metadata information

The final MUSE-DEEP science data product has two 3D image extensions:

- 3D datacube with 2 spatial dimensions and 1 wavelength axis, with flux-calibrated spatial pixels;
- 3D datacube with the errors.

The following additional FITS file are delivered together with the MUSE-DEEP datacube:

- 2D white-light image from the collapsed datacube, called IMAGE_FOV_DEEP;
- 2D exposure map, called EXPOSURE_MAP (as of processing date 2019-11-20 and later, check the ARCFILE timestamp).

The IMAGE_FOV is useful for previewing in image viewers like `rtd`. The exposure map is useful for complex co-addition patterns and can also be previewed in image viewers.

In addition, there is an associated text file delivered that contains the combined pipeline logs with all executions steps for all participating input files, and also the OB grades and comments for them.

There is a set of png files that serve both as QC plot and as preview of the FOV. There is always one for the final deep datacube, and *N* corresponding ones if *N* single files participated[15].

The spectra contain some header keywords added that are related to the QC process. They are listed in Table 3.

## Data Quality

**Master calibrations.** All comments from the MUSE release description apply.

[15] Remember that these individual datacubes are NOT delivered.

**QC, review and certification process.** The MUSE-DEEP datacubes have been reviewed and certified by a process involving both automatic scoring and human-supervised certification. Both the single products (output of `muse_scipost`) and the deep combined datacubes (output of `muse_exp_combine`) are exposed to the QC process.

Table 3. FITS keywords added

| 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; they have no user constraints defined and therefore no OB grades. |
| *QC related information:* | | |
| QCFLAG | e.g. 0000001000 | QC flag composed of 10 bits, see Table 4. |
| QC_COMM<n> | Free text | Comment about the number of combined files, and quality comments |

For the intermediate single products, the QC system scores parameters like

- NAXIS1/2/3 (the size of the product axes; anomalies indicate processing failures);
- NUM_SAT (number of saturated pixels in the raw file);
- maximum correction of wavelength scale by the `muse_scibasic` recipe;
- association quality (proximity of arclamp calibration).

For the deep combined datacubes, the QC parameters are:

- differential offset applied by the alignment procedure;
- number of sources found by the pipeline;
- time difference between first and last OB.

The measured values are compared to reference values and scored. A non-zero score flags a potential issue. All deep combined datacubes are inspected. QC comments are propagated to the datacube headers.

**QC flag.** Similar to the MUSE datacubes, the MUSE-DEEP datacubes have the header key "QCFLAG". It is composed of 10 bits (Table 4). The value 0 always means "OK, no concern". This schema is largely identical to the one for MUSE datacubes, except for their last bit #11 (dataset completeness) which has no meaning here. All comments about the score flags in the MUSE release description apply, except for:

Flag #10 refers to the alignment of the input data. Since the combination was always checked by eye, values 0 or 1 have no particular meaning and have been added for completeness only.

**QC plots and previews.** The QC and preview plots have been originally developed as quick-look plots for the process quality control. It was felt that they might also be useful to the archive user. They are delivered as associated files along with the products. There are two types of plots:

1. the QC plot for the deep combined datacube (Figure 3);
2. the QC plot for a single datacube (Figure 4).

**Process quality control.** The quality of the data reduction is monitored with quality control (QC) parameters, which are stored in a database. The database is publicly accessible and has a browser and a plotter interface[16].

[16] Browser: http://archive.eso.org/qc1/qc1_cgi?action=qc1_browse_table&table=muse_sci_deep
Plotter: http://archive.eso.org/qc1/qc1_cgi?action=qc1_plot_table&table=muse_sci_deep

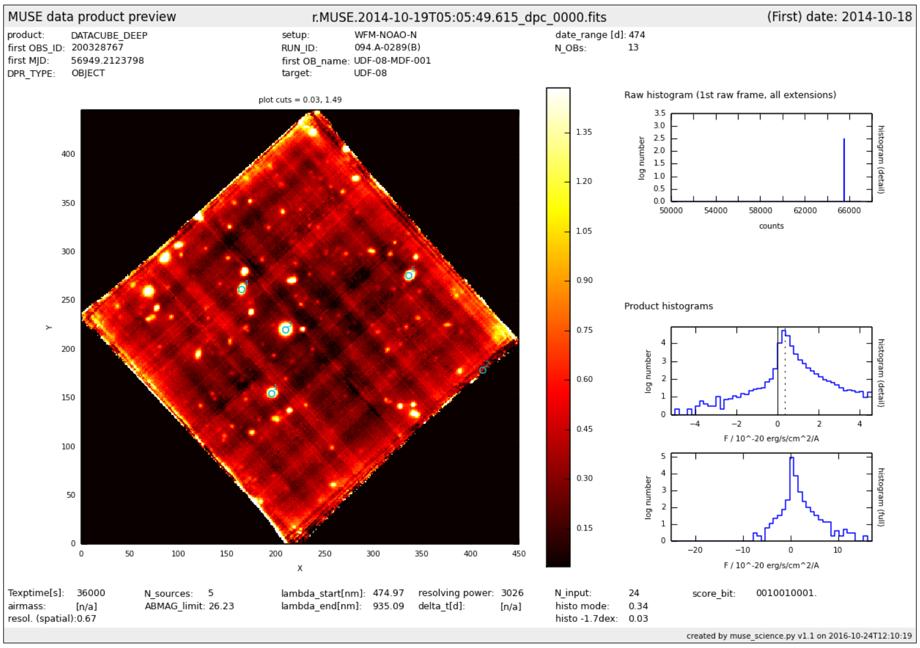

Figure 3. Main QC plot of the deep combined datacube, featuring: the preview (display of the IMAGE_FOV_DEEP file); the histogram of the 1st input raw frame (close-up of the range 50,000-65,000 ADU, as a saturation check); two product histograms (one as a close-up of fluxes around zero, to check the background subtraction; the other one is a histogram for the entire dynamic range of the datacube). **At bottom:** a set of QC parameters applicable to the product (Texptime = total exposure time of the datacube, N_sources = number of pipeline-detected sources, as marked on the display; ABMAG_limit = limiting magnitude (depth) of the datacube; mean_FWHM = median FWHM of point sources, for AO mode; N_input = number of input OBJECT files; histo mode = flux value for the maximum in the product histogram, also marked by the broken line; histo-1.7 = flux value where histogram value has fallen off by -1.7dex as compared to the mode; score_bit = QC flag as stored in the header, see Table 4. **On top:** some keywords read from the product file header, like first OB name and target name.

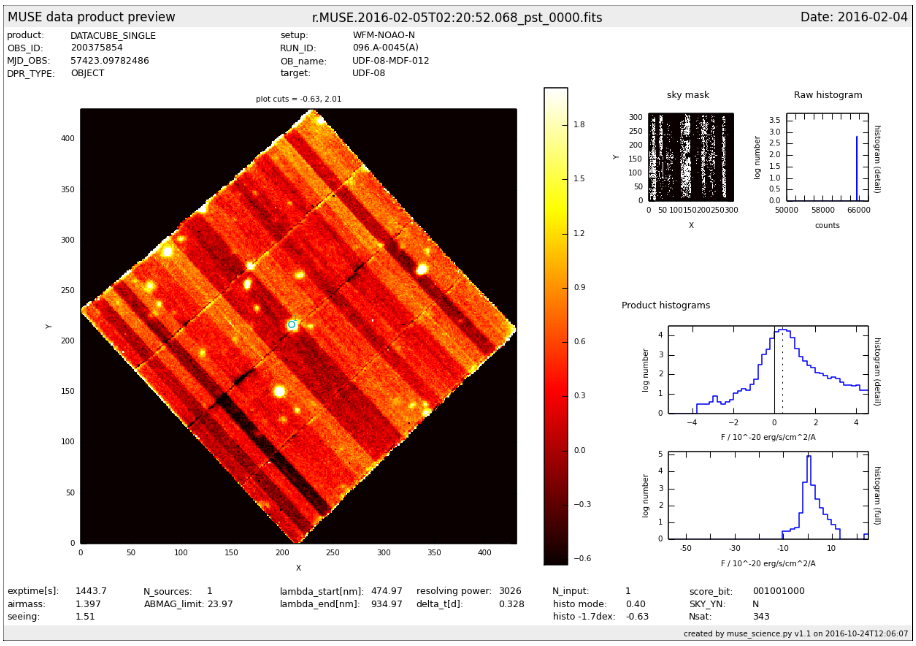

Figure 4. QC plot of one of the individual datacubes, for comparison to the plot displayed above. It shows the same properties and parameters as the previous figure, except for: exptime (exposure time of the raw file); SKY_YN: Y if this datacube has used a dedicated SKY observation for SKY subtraction; Nsat = number of saturated pixels. If SKY_YN=N, there is also a display of the sky mask used for the sky background fit.

| Bit | **Content** (if YES, value is 0, otherwise 1) | **Motivation** |
|---|---|---|
| #1 – master sky line fit | No pipeline error upon master sky fit? | Catch a pipeline error upon master sky fit ("master sky fit failed with error code 21: the iterative process did not converge."); if at least 1 `muse_scipost` product has that error, the value 1 is propagated to the deep cube. |
| #2 – OBJECT vs. SKY | This datacube comes from OBJECT frames? | Not applicable (always 0 for MUSE-DEEP) |
| #3 – SKY observation | A dedicated (user-defined) sky observation exists? | Not applicable (due to the nature of many deep targets, there is usually no quality difference between cases with SKY and without SKY). |
| #4 – arc calibration | Time difference within 1.5d (previous/this/next night)? | Usually, daytime calibrations come within 0.5 days after the science observation; if more than a day difference, probability for a mismatch is higher, affecting the wavelength scale error (very rarely violated); if at least 1 `muse_scipost` product has a score 1, the value 1 is propagated to the deep cube. |
| #5 – SKY_FLAT | Existing? | Not applicable (always 0 for MUSE-DEEP) |
| #6 – saturated pixels | Number of saturated pixels in all input raw frames lower than 300? | Flags cases with partial saturation (which cannot be directly discovered in the product datacube); if at least 1 `muse_scipost` product has a score 1, the value 1 is propagated to the deep cube. |
| #7 – number of sources | Number of sources found by the pipeline >0? | Not applicable (flag almost always 0 for MUSE-DEEP datacubes) |
| #8 – sky subtraction quality | HISTO_17 parameter >-20? (i.e. not negative) | Quality of sky subtraction: the issue of sky over-subtraction for crowded fields is solved for MUSE-DEEP cubes, hence this bit is almost always 0. |
| #9 – wavelength scale quality | LSHIFT_MAX | Quality of wavelength scale: maximum of residual correction done on sky lines, in Angstrom; should be <0.5 A (was 0.2 Å before 2019-10-01; if at least 1 `muse_scipost` product has a score 1, the value 1 is propagated to the deep cube. |
| #10 – alignment | Differential offset between individual observations <6e-5 deg (0.2arcs) *AND* all input frames matched? | Not applicable here (alignment is always done as careful as possible and checked by eye) |

Table 4. Definition of QC flags. Flags marked "not applicable" are included to align with the MUSE scheme. Find the up-to-date list under the URL http://www.eso.org/qc/PHOENIX/MUSE/score_bits_deep.txt.

The criterion for bit #9 has been relaxed with processing date 2019-10-01 from the previous threshold value 0.2 Å to 0.5 Å, due to confidence built by experience and a stable pipeline.

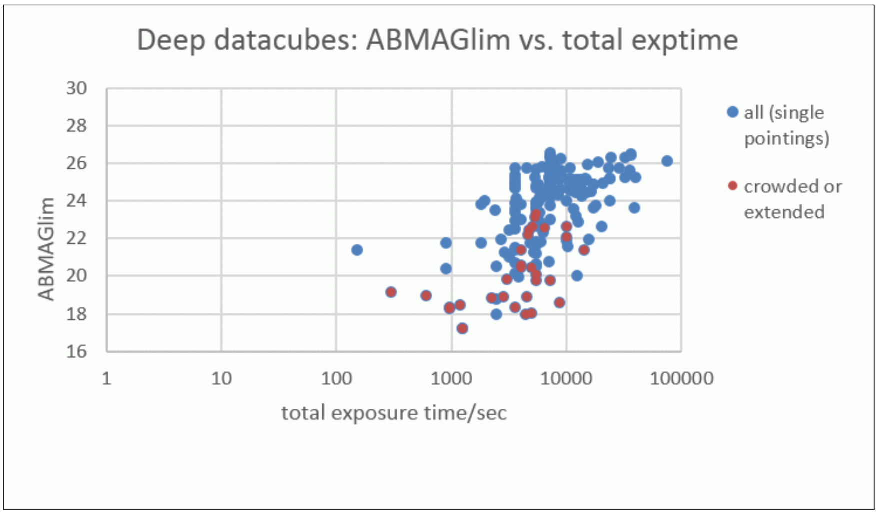


Figure 5. Limiting magnitude ABMAGlim vs. total exposure time, for all deep datacubes until 2015-09 that result from a single pointing (with small jittering). Data points from a crowded field or from extended sources are marked in red.

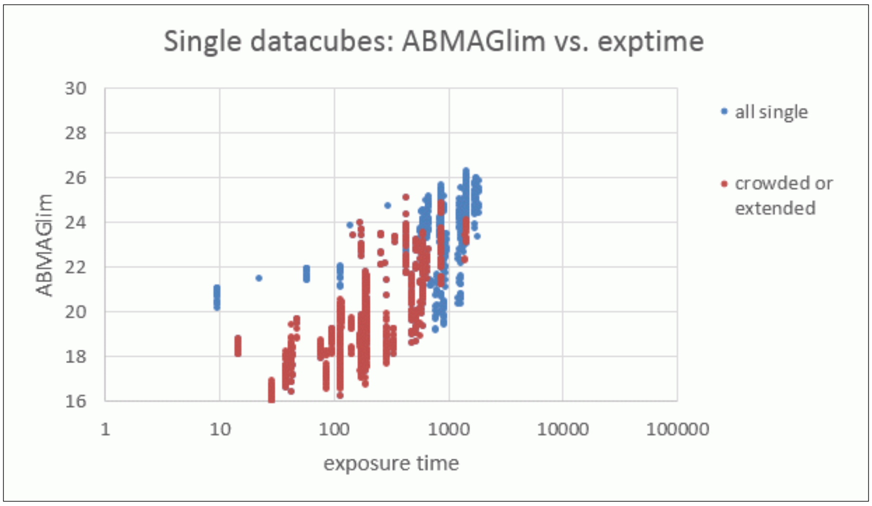


Figure 6. Limiting magnitude ABMAGlim vs. exposure time, for all single exposures until 2016-09 (as taken from the MUSE release). Data points from a crowded field or from extended sources are marked in red.

**Error propagation.** This is the same as for MUSE datacubes and is described in their release description.

**Limiting magnitude ABMAGlim.** Each datacube has a QC parameter ABMAGlim. Its exact definition is described in the MUSE release description. The deep datacubes are expected to have a correspondingly higher value of ABMAGlim than the single or the OB-based datacubes, except for pathological situations like crowded fields.

In Figure 5 we display this QC parameter for all deep datacubes, versus their total exposure time. We have selected only values for single pointings (excluding values for datacubes with several, partly overlapping pointings), because the limiting magnitude is a concept assuming applicability across the entire field of view. We have also identified those datacubes with a background that is presumably not dominated by background noise:

- targets are globular clusters ("crowded field", see example in Figure 7),
- targets have an extended, diffuse emission ("extended object", see Figure 8).

They are plotted in red. These datacubes cannot be expected to have their ABMAGlim improved with increasing exposure times.

A general trend towards ABMAGlim increasing with total exposure time is clearly visible. There is some saturation in the ABMAGlim values, they do not go beyond about 26.5 mag. The definition of ABMAGlim refers to the narrow noise peak, as seen in Figure 10. This figure illustrates the presence of residual pattern in the background due to incomplete correction of instrumental signature (slice-to-slice response) and sky emission line residuals, and the resulting effect on the background noise properties. The reduction of background noise falls short of the expected ~1/sqrt(exptime) scaling, presumably due to the superposition of the residual background pattern which is already intrinsic to the OB-based datacubes.

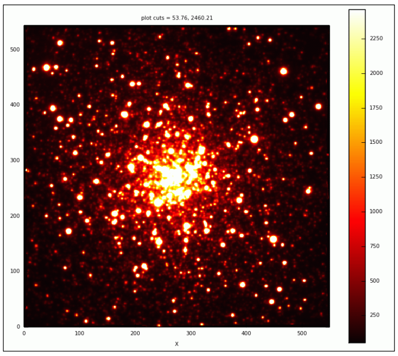

Figure 7. Crowded field: limiting magnitude is dominated by point sources rather than by background noise.

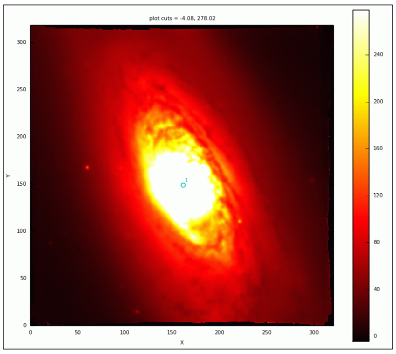

Figure 8. Extended source: limiting magnitude is dominated by diffuse object emission rather than by background noise.

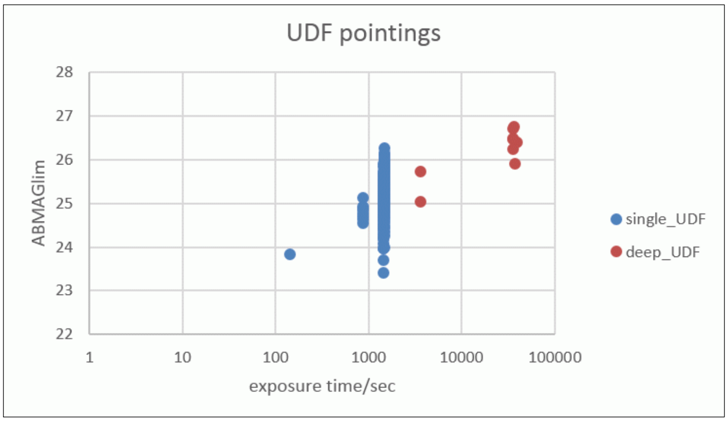


Figure 9. ABMAGlim plot for the single (blue) and deep (red) datacubes for pointings of the Hubble Ultra-Deep Field South[17]. See also Figure 13 and Figure 14.

As illustrated in Figure 6, the same parameters displayed for the single datacubes from the MUSE release show the systematic and expected trend. We have again marked the crowded or extended fields which are subject to the systematic effects. In particular the crowded fields get their background over-subtracted in the OB-based MUSE reduction scheme.

In Figure 9 we display the ABMAGlim values for a set of programmes designed to go deep, targeting at the Hubble UDF. One programme is collecting a total of 1 hour per pointing, the other one collects about 10 hours in each of 9 pointings. There is clearly the trend towards higher ABMAGlim values for longer exposure times.

[17] Participating runs: 094.A-0205B, 094.A-0289B, 095.A-0010A, 096.A-0045(A/B); PIs L. Wisotzki, R. Bacon

In Figure 10 we demonstrate how the background noise peak narrows upon going deep. The FOV image of the deep datacube (right) shows how many faint sources peak out of the narrow noise pool, sources which are not seen in the single (left) nor the OB-based datacubes (middle). This figure also demonstrates that the sky residuals at least partly cancel out upon deep combination.

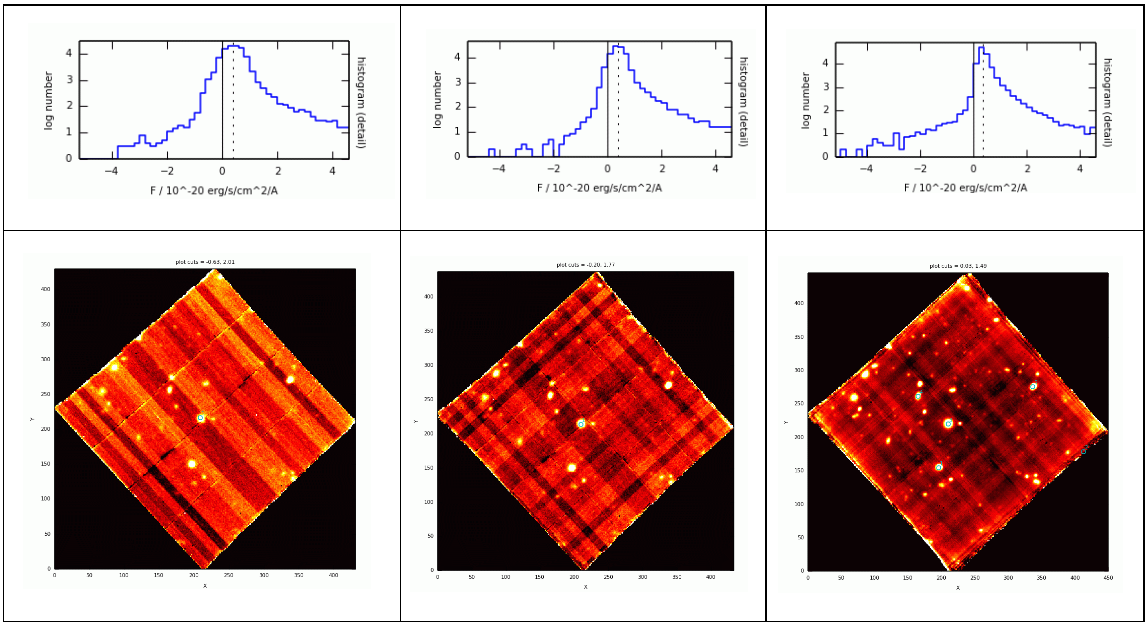


Figure 10. Datacube histograms (top) and FOV images (below) of the same UDF field, for a single 1500 sec exposure (left), an OB-based 3000 sec combined datacube (middle), and the final deep datacube, worth 10 hrs of exposure time, 13 OBs and collected over a time span of 474 days.

**Mapping deep datacubes and total exposure time.** In a few cases, deep exposures have been obtained for fields that are larger than the 1'x1' MUSE field of view. Often PIs have then designed OBs with e.g. four pointings that have some overlap. See a typical example in Figure 11. Whenever technically possible (in terms of total number of input files currently limited to about 125) and reasonable, we have combined those pointings in one single datacube.

Originally, the MUSE pipeline did not provide exposure maps. For situations like the one sketched in Figure 11, it is straightforward to derive the exposure map. For more complex situations (like in Figure 12, Figure 13 and Figure 14) it is best to obtain an overview of the pointings from the preview plots of each input exposure. In such complex situations, the effective exposure time per pixel is a weighted average (EXPTIME). This is then also true for ABMAGlim.

Since 2019-11-20 (processing date, check ARCFILE timestamp), we deliver an EXPOSURE_MAP as ancillary file for each deep datacube.

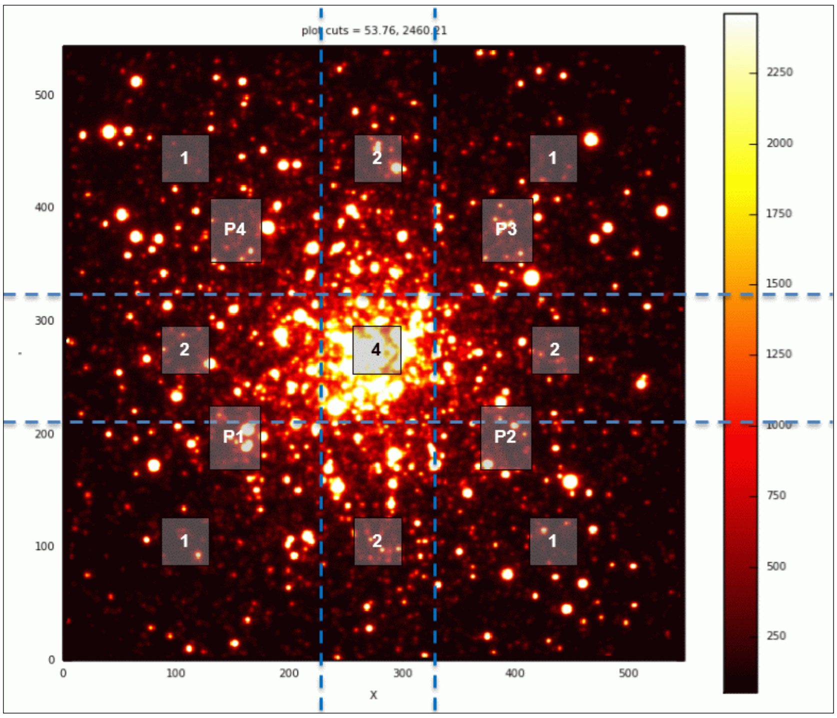


Figure 11. Sketch of a typical case of deep mapping. The single FOV of MUSE covers about 320x320 pixels. This globular cluster has been mapped in 4 partly overlapping pointings, their centres are marked as P1…P4. The exposure times per pixel, and thereby also the noise characteristics of ABMAGlim and SNR of extracted sources, depend on the source position in the FOV. If the effective exposure time for all exposures of pointing P1 is normalized to 1, then there are large fields (labelled 1) with effective exposure time 1, stripes with effective exposure time 2, and the central region with effective exposure time 4.

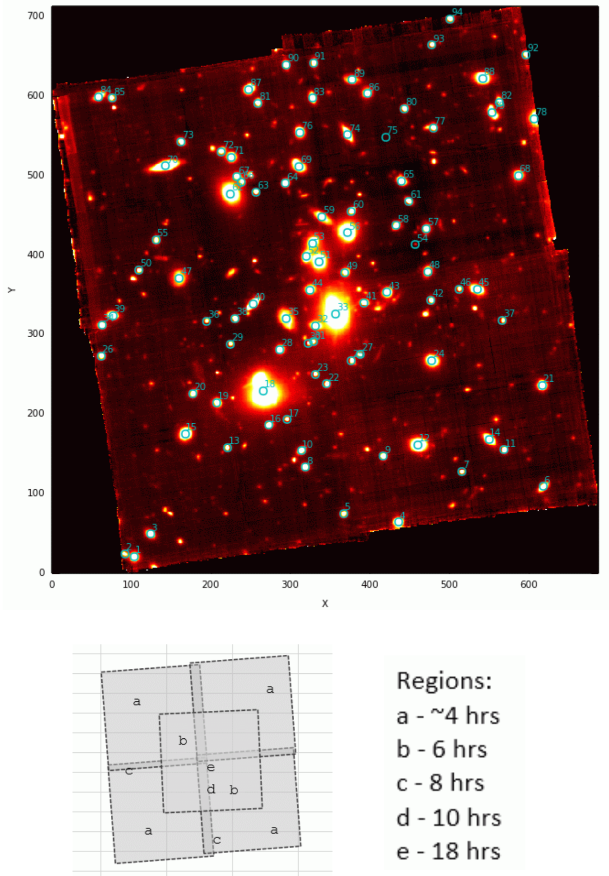


Figure 12. Complex 4+1 mapping[18] of Abell-2714.

At left we display the preview of the field, and at bottom left the exposure map. There is the 4-position mapping pattern as in Figure 11, marked ‘a’, and a central pointing ‘b’. The pointings ‘a’ received roughly 4 hrs, the pointing ‘b’ an additional 2 hrs. The exposure map reveals overlapping stripes ‘c’ and ‘d’, and a small central region ‘e’ that was effectively exposed for a total of 18 hrs. It is only by superposition of all input data that such a depth could be reached.

Most deep maps are similar to the one from Figure 11, but some are more complex. In the following we sketch the most complex situations we have encountered so far. Figure 12 shows a mapping like in the previous figure, with an additional central pointing. For this deep datacube, the exposure map becomes a bit complex. It can be derived by compiling the individual FOV plots.

In Figure 13 and Figure 14 we illustrate another configuration with a 3x3 grid and an additional deep exposure. With a total of 275 input files we were unable to process all of them into a single deep map, but we could come close to the ideal solution with 5 deep datacubes for pointings UDF-03, 06, 07, 08, and 09, plus one deep datacube combining UDF-02, 04, 05, and 10, and a final one combining UDF-01 and 10. Note that in this exceptional case we have used the photons from pointing UDF-10 twice, a situation which is so far unique within the MUSE-DEEP release.

[18] PI J. Richard, programme 094.A-0115A and following.

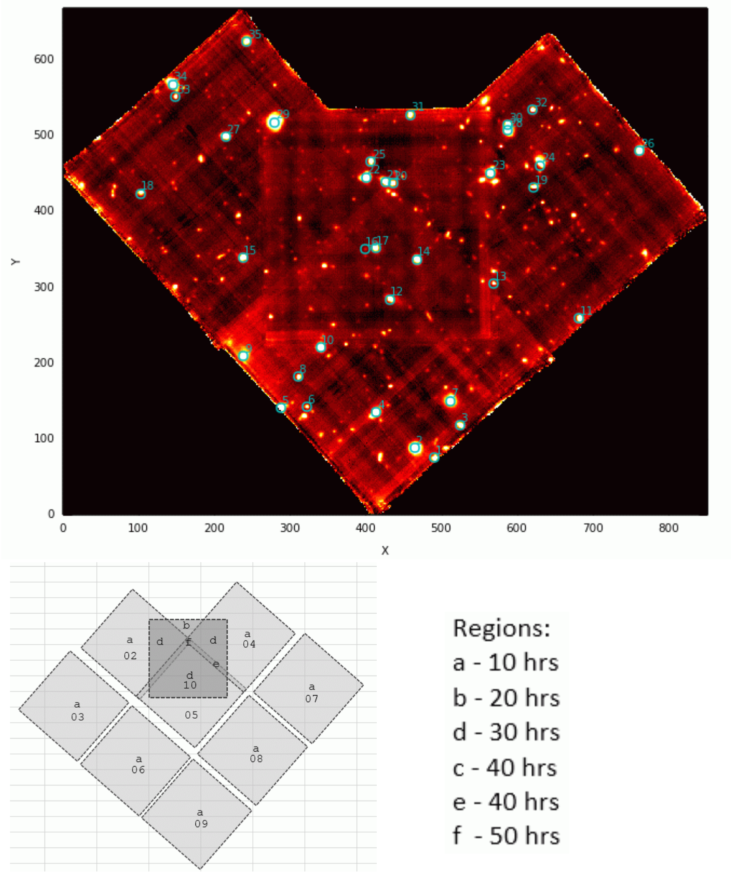


Figure 13. Complex 3+1 mapping of UDF pointings[19].

The entire map has 9+1 pointings (a 3x3 grid and a central pointing, see sketch at bottom). Altogether this would amount to the combination of 9*25 + 1*50 = 275 exposures which is more than a factor of 2 beyond our capacity. We have decided to process 5 of the pointings (UDF-03, 06, 07, 08, 09) into separate deep datacubes. Then we have combined UDF-02, 04, 05 and 10 into the one displayed at left (with 123 input files). It has the optimal depth everywhere except for the small region 'b' which lacks the contribution from pointing UDF-01. In total there are 7 deep datacubes, which are close to the theoretical optimum. The regions in the exposure map are defined at left. To optimize the region 'b' we have created another deep datacube, see the next figure.

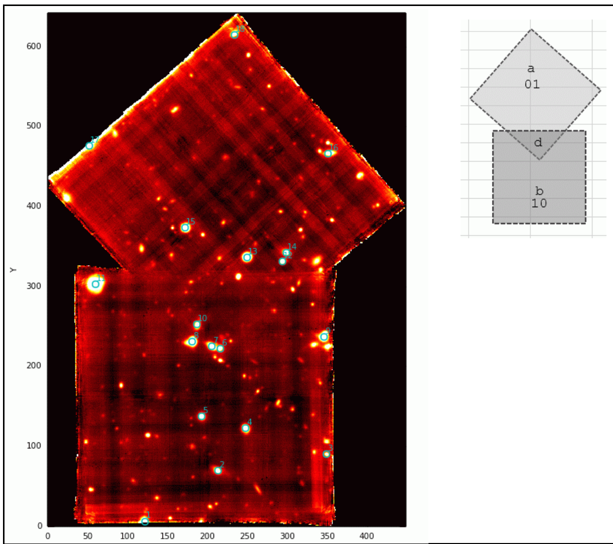

Figure 14. This is the final pointing UDF-01/10, with N=74 input files. It provides the same depth for the small region 'd' of the central field UDF-10 as its other parts in the previous datacube.

## Known features and issues

### Issues

**General.** Files known from the MUSE release to have issues like guiding errors, derotator problems etc., have not been selected for the MUSE-DEEP datacubes.

[19] PI R. Bacon, programme 094.A-0289B and following.

**Misalignment.** While all deep datacubes have been checked visually for misalignment (at the IMAGE_FOV level), and while there are also automatic checks, there is a non-zero chance that cases of misalignment have been overlooked. In Figure 15 below is an example how subtle this effect can be. The double sources might be overlooked easily, and the ring-like signature becomes visible only with the upper brightness threshold set rather low. **We strongly recommend visual checks for this issue.** This is particularly true for the NFM data which often have only one bright source in the center.

You may want to point your viewer (e.g. ds9) to a wavelength where (redshifted) emission lines become visible, which might give a brighter signal than continuum sources, and also choose the dynamic range appropriately.

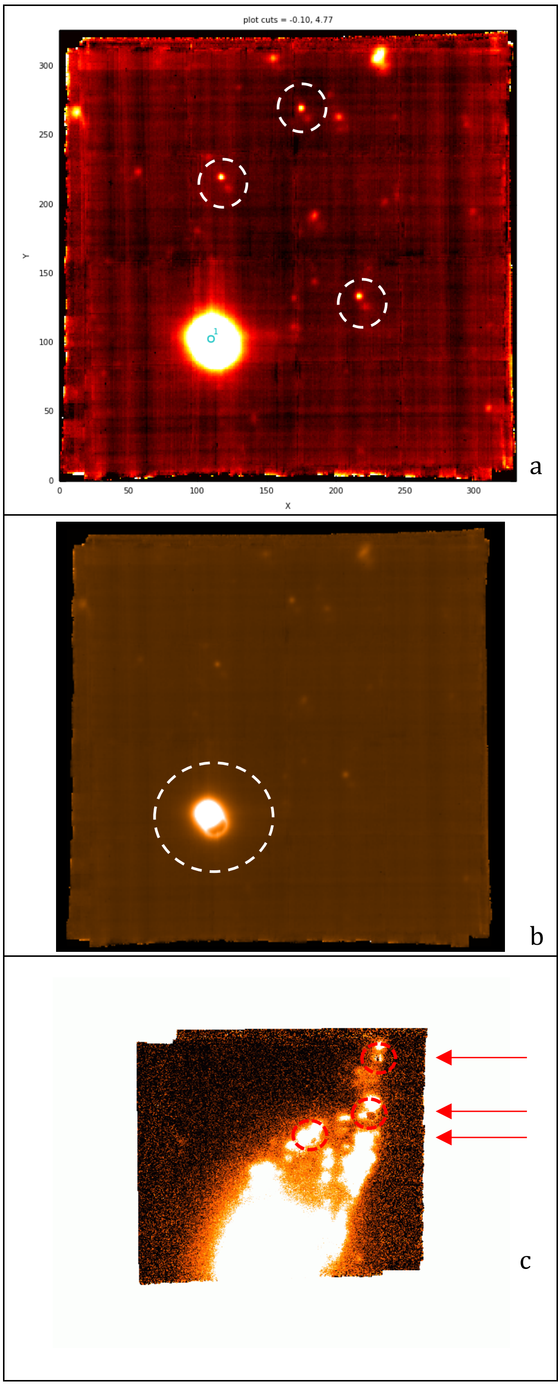


Figure 15 a-c. Examples for misalignment. Panels a and b show results for a deep combined datacube with formally well-behaved alignment correction recipe, but with subtle visual indications of misalignment. The Figure a (top) exhibits some duplicated fainter sources. Once alerted, additional evidence for misalignment comes from the very bright sources which, if displayed as IMAGE_FOV fits file with an upper threshold set to low values, exhibits the typical crater-like symptom (Figure b). This comes from the stacking procedure in `muse_exp_combine`. In this example, one out of 8 input frames is shifted against the others. The bright outlier signal is clipped (dark hole) except for its outer wings which are within the acceptance threshold giving rise to the narrow bright ring. This signature is typical for alignment issues with *one* input frame while the others are well-behaved.
In Figure c we show a different example as taken as screenshot from the display of ds9. Only when the tool is pointing to the wavelength of a strong emission line (here: 6640 A), the bright knods as indicated by arrows clearly show the displacement, as a pair of "mountain-valley" structures.

**Wiggles in AO-E data.** For AO-E data (extended mode), the current pipeline version does not correct properly for bumps and wiggles in the instrument response function. These are caused by the transmission of the Na filter that is used to block the contamination from the AO laser star. The wiggles are propagated to the science spectra. They are particularly evident in the blue part of the spectrum (see Figure 16).

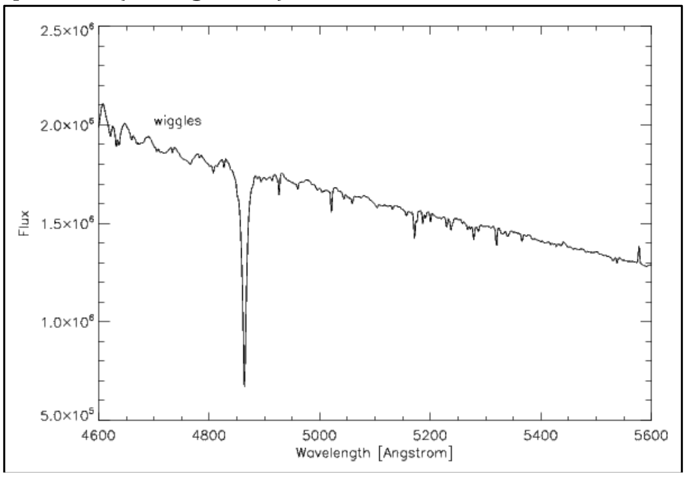


Figure 16. Wiggles in the spectra of for instrumental setup AO-E.

**Raman scattered laser lines.** The MUSE observations using the AO system are affected by the Raman-scattered light from the lasers. Its contribution is seen mainly as emission lines at 6485 A and 6827 A, which fluxes vary slowly across the field of view by about 5%. Currently, the MUSE cubes are not corrected for this effect.

**Quality issue with data obtained between February and April 2019.** Due to an issue with the telescope dome safety camera, MUSE observations taken between February 1st and April 18th, 2019, suffered from light contamination. The contamination is visible as an excess of continuum emission between 800 and 900 nm, with the peak around 860 nm (Figure 17). The rest of the spectral range is not be affected.

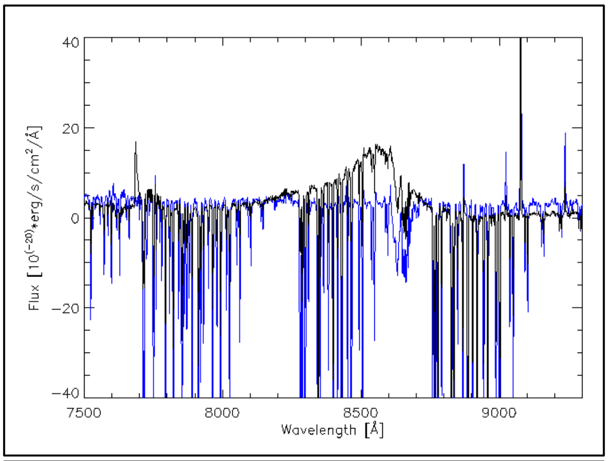


Figure 17. Example of the red part of MUSE spectra of the same program and the same pointing, where data were taken before (blue) and during (black) the affected period. The contamination is visible as a bump between 800 and 900 nm peaking at 860 nm.

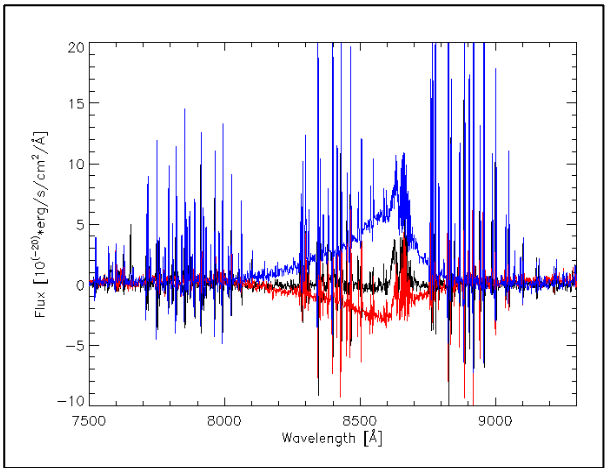


Figure 18. Examples of the contamination residuals in the red part of spectra of different sky regions (darkest regions of the field of view) of a single data product. Note the variable contamination residuals between 800 and 900 nm.

The analysis of data products from that period has shown that contamination might vary from exposure to exposure. This is propagated to the final data product. In some cases, when the sky background is estimated on the same science frame, the pipeline successfully removes contamination from the final data cube. However, when the sky background is created from a dedicated SKY pointing, there are often residuals of contamination in the final products, presumably due to the variability of the sky.

Moreover, the residuals of contamination in MUSE data products vary across the field of view (Figure 18). This is probably due to variability in the light source itself and to possible reflections within the telescope dome.

Users are urged to check the MUSE-DEEP IDPs which have contributions from that period of time with caution. Any residuals of the contamination should be taken into account for the analysis of the scientific signal.

**Quality issue with data obtained between June 15 and June 19, 2022.**

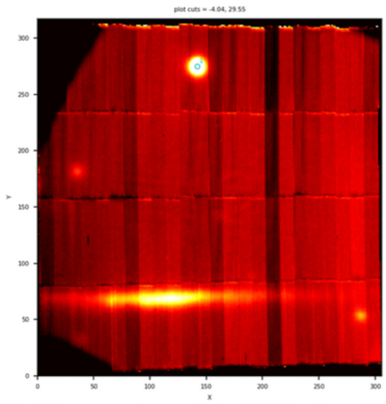

Figure 19: Example of the vignetted field-of-view due to camera beam splitter in the optical path.

The MUSE on-sky: science, standard star and twilight flat observations, acquired in that period of time, show vignetting (see example). It was caused by the GALACSI commissioning camera beam splitter inserted in the optical path. In the WFM, this element defines a square 82"x82" within which the light that is passed to MUSE is decreased by approximately 50% at 500 nm to 75% at 900 nm. The light modification may vary across the field-of-view. The vignetted master calibrations were used to process vignetted science exposure.

Users are urged to consider the MUSE-DEEP IDPs from that period of time with caution.

## Features

**Background variations.** For input files with extra SKY pointings, taken under non-photometric conditions, the individual background may show fluctuations, because the SKY pointing was taken under different photometric conditions than the OBJECT pointing. This likely broadens the background peak in the histogram and limits the reachable ABMAGlim values. Nevertheless the deep combined data show better SNR in the sources.

**Crowded fields with varying background.** Occasionally crowded field data, although processed without any SKY subtraction, show an artificially high background that is due to a pipeline issue that is not solved. Figure 20 shows an example. If analysed with aperture photometry these artefacts should be irrelevant.

**Saturation.** Check carefully the saturation flag #6. If 1, then at least one of the input files has more than 300 saturated pixels. In a deep datacube it might be difficult to tell which spaxels got affected. The QC plots of the individual datacubes might give further information about the level of saturation. If saturated pixels occurred, be very cautious with the analysis.

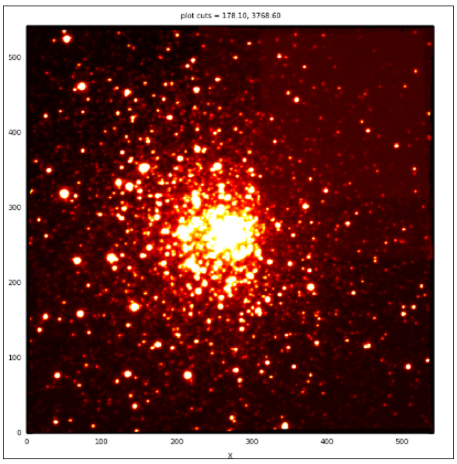

Figure 20. Deep mapping, with the upper right quadrant having a higher background level than the others. This is due to an unsolved pipeline issue.

**Flux scale inaccuracies.** In rare cases it turned out upon deep combination of OBs that the input candidates had a strongly deviating flux scale, sometimes by more than a factor 10, caused by using an inappropriate flux standard star measurement. The combination of such data might lead to unwanted and unexpected results, e.g. a bad alignment (because the alignment correction algorithm uses noise criteria to identify candidate sources). If discovered, we have tried to fix the issue by choosing another standard star for the flux calibration, or we have rejected the product cube. Nevertheless there might be cases that escaped our attention. The signature of this issue is unusual skyline residual patterns, misalignments, and strongly different flux scales.

**Deep datacubes with mixed AO and NOAO data.** Since August 2017 some MUSE data are taken in laser-assisted AO mode with ground-layer correction. The wavelength range between about 580 and 596 nm (N range), or between 576 and 601 nm (E range) is suppressed (flux set to zero) in the pixel-tables of these data, due to laser-induced sodium lines. A few deep datacubes contain mixed NOAO and AO input data. This is justified because in general these programmes are designed to have matching seeing constraints. Note that in those cases there is no spectral gap (it is filled with NOAO data) but the SNR is lower across this range, and there might be steps in the spectral fluxes because of the different number of combined spectra within and outside the sodium range.

**Transients.** Satellite trails and other transients (like minor planets) get diluted over the deep combination of OBs. The user should check for faint linear structures in the cubes. The user should also be aware that we have effectively destroyed any time variability information in the data. For time domain analysis, the user should always check the OB-combined and the single datacubes.

**Multiple run IDs.** Some of the deep datacubes are combined from OBs obtained in multiple runs. The file headers contain the key `PROG_ID` which either lists the run ID (if unique), or is filled with '`MULTI`' and then is followed by additional keys `PROGIDi` listing all contributing run IDs.

**OB IDs.** All participating OBs are listed in the headers as `OBIDi`.

**Provenance and access rights.** All participating raw files are listed under `PROVi`. The access rights are derived under the rule that a deep datacube is public only if *all input data are public*. If a datacube is not yet public and all input files belong to the same run ID, the datacube is accessible to the PI of that run only. If a datacube is not yet public and the input files belong to different runs (PROG_ID = '`MULTI`'), the whole datacube is not accessible, even to the PI(s).

**OB grades and OB comments.** The OB grades and comments (if available) are not stored in the headers but in the associated text file with name *r.MUSE…dpc.log* where this information is found at the end.

**Combined datacubes with more than 130 input files.** Our processing scheme for the last combination step (muse_exp_combine) drizzles all pixel-tables generated in the previous steps onto the common spatial grid of the final datacube. This algorithm preserves the spatial resolution in the best possible way, but it is memory-intensive. With 2 TB of memory available for the MUSE-DEEP processing, we are limited to about 130 input files. In those (very few) cases when this number is exceeded, we have achieved the final solution as spectral "sub-cubes", by processing several wavelength bands separately. We have then stitched the sub-cubes together, into the final deep cube with the full wavelength range. Find the list of stitched deep datacubes under the URL http://www.eso.org/qc/PHOENIX/MUSE_DEEP/MUSE_DEEP_stitched_cubes.txt.

## Tips and tricks

**Post-pipeline removal of sky lines.** See the MUSE release description.

**Analysis software package.** See the MUSE release description.

**Working with pipeline log files.** See the MUSE release description. In addition, the log file has a section 3 at the end ("Selection file for this combined datacube"). It lists the products of the selected input OBs, with the OB IDs, OB names, the user-defined ambient constraints for the seeing ("AMBI_REQ"), the OB grades and the OB comments. The listed pipeline product names are the names of the COMBINED datacubes that are also available as MUSE datacubes. Finally, all raw file IDs used for the deep datacube are listed.

# Data Format

## Files Types

The primary MUSE-DEEP product is the 3D datacube:

| ORIGFILE names starting with | Product category HIERARCH.ESO. PRO.CATG | Format | How many input files? | Description |
|---|---|---|---|---|
| MU_SCBD | DATACUBE_DEEP | 3D spectro-image | N>1 | combined datacube from OBJECT observations in at least 2 OBs |

Each product has two ancillary FITS files (only one if processed before 2019-11-20):

| ORIGFILE names starting with | Product category HIERARCH.ESO.PRO.CATG | Format | How many input files? | Description |
|---|---|---|---|---|
| MU_SIMD | IMAGE_FOV_DEEP | 2D image | N>1 | collapsed white-light image of combined FOV |
| MU_SXPM | EXPOSURE_MAP (since 2019-11-20) | 2D image | N>1 | imaged in the same format as IMAGE_FOV_DEEP, with the total exposure time as pixel values |

Furthermore the following non-FITS files are delivered with each datacube:

| ORIGFILE names starting with | ASSOCn (listed in the header of the main product) | Format | How many? | Description |
|---|---|---|---|---|
| r.MUSE...dpc.png | ANCILLARY. PREVIEW | png file | 1 | See Figure 3. |
| r.MUSE...pst1.png | ANCILLARY. PREVIEW | png file | *N* | One for each input file; see Figure 4 |
| r.MUSE...dpc.log | ANCILLARY. README | text file | 1 | all recipe processing logs for the deep datacube |

The following naming convention applies to the ORIGFILE product: e.g. the name `MU_SCBD_1117772_2015-04-12T00:56:47.087_WFM-NOAO-E_OBJ.fits` has the components:

| ORIGFILE component: | MU | SCBD | 1117772 | 2015-04-12T00:56:47.087 | WFM-NOAO-E_OBJ.fits |
|---|---|---|---|---|---|
| refers to ... | MUSE | product type (S stands for science, CB for cube, D for deep) | first OB ID | timestamp of first raw file | setup string:<br>wide-field mode, no AO, extended wavelength range;<br>DPR.TYPE=OBJECT (always) |

The ancillary files have the following ORIGFILE names:

Table 5. Naming conventions of ANCILLARY files

| type | example | rule |
|---|---|---|
| ANCILLARY.README | r.MUSE.2015-04-12T00:56:47.087_dpc.log | Technical filename of the main fits file, with extension 'log' instead of 'fits' |
| ANCILLARY.PREVIEW | r.MUSE.2015-04-12T00:56:47.087_dpc.png | same name, with extension 'png' instead of 'log' |
| ANCILLARY.PREVIEW (*N*) | r.MUSE.2015-04-12T00:56:47.087_pst1.png, etc. | names of all individual exposures |

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

## File structure

The MUSE-DEEP datacube product has two 3D image extensions:

- 3D datacube with 2 spatial dimensions and 1 wavelength axis, with flux-calibrated spatial pixels; the EXTNAME key is 'DATA'.
- 3D datacube with the variance, EXTNAME is 'STAT'.

### File size

The typical size of a deep datacube is 3-5 GB if it was collected from one pointing only (with small jitter offsets). The size grows in proportion to the number of non-overlapping pixels. The larger values apply to datacubes with orientations inclined with respect to the RA/DEC grid.

Processing a deep datacube from 30 input files takes about 500 GB of memory. That amount scales with the number of files. Our current deepest datacube is made from 124 input files (requiring 2 TB memory).

## 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/42 .*

All users are kindly reminded to notify Mrs. Grothkopf (esodata at eso.org) upon acceptance or publication of a paper based on ESO data, including bibliographic references (title, authors, journal, volume, year, page numbers) and the observing programme ID(s) of the data used in the paper.