The Physics of the Accelerating Universe Survey: narrow-band image photometry

S. Serrano [1,2,3]★ E. Gaztañaga,[1,2,4] F. J. Castander,[1,2] M. Eriksen,[5]† R. Casas,[1,2] D. Navarro-Gironés,[1,2] A. Alarcon [6] A. Bauer,[1,2] L. Cabayol,[5] J. Carretero,[5]† E. Fernandez,[5] C. Neissner,[5]† P. Renard [1,2,7] P. Tallada-Crespí,[5]† N. Tonello [8]† I. Sevilla-Noarbe,[9] M. Crocce [1,2] J. García-Bellido [10] H. Hildebrandt,[11] H. Hoekstra [12] B. Joachimi,[13] R. Miquel,[5] C. Padilla,[5] E. Sánchez[9] and J. de Vicente [9]

[1]*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, E-08193 Barcelona, Spain*
[2]*Institut d'Estudis Espacials de Catalunya (IEEC), E-08193 Barcelona, Spain*
[3]*Satlantis, University Science Park, Sede Bld, E-48940 Leioa-Bilbao, Spain*
[4]*Institute of Cosmology & Gravitation, University of Portsmouth, Dennis Sciama Building, Burnaby Road, Portsmouth PO1 3FX, UK*
[5]*Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, E-08193 Bellaterra (Barcelona), Spain*
[6]*HEP Division, Argonne National Laboratory, Lemont, IL 60439, USA*
[7]*Department of Astronomy, Tsinghua University, Beijing 100084, China*
[8]*Barcelona Supercomputing Center (BSC), Plaça Eusebi Güell 1-3, E-08034 Barcelona, Spain*
[9]*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, E-28040, Spain*
[10]*Instituto de Fisica Teorica UAM/CSIC, Universidad Autonoma de Madrid, E-28049 Madrid, Spain*
[11]*Ruhr-University Bochum Astronomical Institute (AIRUB), Faculty of Physics and Astronomy, Ruhr University Bochum, German Centre for Cosmological Lensing, D-44780 Bochum, Germany*
[12]*Leiden Observatory, Leiden University, Niels Bohrweg 2, NL-2333 CA, Leiden, the Netherlands*
[13]*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

ABSTRACT

PAUCam is an innovative optical narrow-band imager mounted at the *William Herschel Telescope* built for the Physics of the Accelerating Universe Survey (PAUS). Its set of 40 filters results in images that are complex to calibrate, with specific instrumental signatures that cannot be processed with traditional data reduction techniques. In this paper, we present two pipelines developed by the PAUS data management team with the objective of producing science-ready catalogues from the uncalibrated raw images. The NIGHTLY pipeline takes care of entire image processing, with bespoke algorithms for photometric calibration and scatter-light correction. The Multi-Epoch and Multi-Band Analysis pipeline performs forced photometry over a reference catalogue to optimize the photometric redshift (photo-z) performance. We verify against spectroscopic observations that the current approach delivers an inter-band photometric calibration of 0.8 per cent across the 40 narrow-band set. The large volume of data produced every night and the rapid survey strategy feedback constraints require operating both pipelines in the Port d'Informació Cientifica data centre with intense parallelization. While alternative algorithms for further improvements in photo-z performance are under investigation, the image calibration and photometry presented in this work already enable state-of-the-art photo-z down to $i_{AB} = 23.0$.

Key words: instrumentation: detectors – methods: data analysis – techniques: image processing – techniques: photometric – large-scale structure of Universe – cosmology: observations.

1 INTRODUCTION

Current cosmological studies have been increasing their volume and complexity of data up to a point that traditional analysis methods are not practical anymore. In the past, the data obtained by an astronomer could be processed at the telescope itself or by a personal computer at the research institute. Today, the massive volume and complex

analysis require processing in a data centre, or even in a grid of computing centres. In 2000, when the Sloan Digital Sky Survey (SDSS; York et al. 2000) started their observations, the volume of data produced in a single week was larger than all previous data collected in the history of astronomy. The whole SDSS data will be negligible compared to the 60 PB of raw observations that will be produced by the Vera Rubin Observatory (Ivezić et al. 2019).

Fortunately, the increased volume of observations came with new and more powerful computing technologies that enabled its necessary analysis. The era of big data arrived in time providing the data management tools we need. The data reduction techniques used in previous surveys cannot be simply applied, and new scalable

★ E-mail: serrano@satlantis.com
† Also at Port d'Informació Científica (PIC), Campus UAB, C. Albareda s/n, E-08193 Bellaterra (Cerdanyola del Vallès), Spain.

algorithms had to be designed that meet the more strict needs of today's studies. Higher level languages such as Python enable fast and versatile program development not possible with older programming languages like FORTRAN or C. Even standard astronomy libraries such as IRAF (Tody 1986) that have been used for decades are becoming obsolete with more flexible astronomical Python libraries like Astropy (Astropy Collaboration et al. 2018) or Photutils (Bradley et al. 2020). The combination of parallel processing in high throughput computing data centres with these new advanced libraries changed the paradigm in which data reduction pipeline are being built.

In this paper, we describe the particular image processing and analysis required for a large-scale narrow-band cosmology survey, the Physics of the Accelerating Universe Survey (PAUS). To achieve the scientific goals of the survey, we built PAUCam (Padilla et al. 2019), a large field of view camera equipped with a large set of narrow-band filters that enable low-resolution spectra for all the sources in the sky. This massive camera was mounted in the prime focus of the *William Herschel Telescope* (*WHT*). Its 4.2-m diameter mirror allows observing fainter objects through the reduced transmission of the particular narrow-band filters of PAUS. Both camera and data reduction system are designed to optimize the photometric redshift (photo-z) precision, delivering complete and homogeneous galaxy catalogues down to a total magnitude of i_{AB} or $I_{auto}(AB) = 23.0$.[1]

This paper describes the imaging data set we are dealing with (Section 2), the raw instrumental detrending (Section 3), the photometric (Section 5) and astrometric calibration (Section 4), the particular forced photometry process (Section 6) that enables the science ready catalogues, and the validation process (Section 7). In Section 8, we present a validation of our pipeline using photo-z estimation, which has a direct impact on the science applications already published or plan in separate papers. In Appendix A, we present the operation and technical performance of the pipeline, including the data flow and orchestration. Appendix B gives details of the data base schema. Appendix C compares some example reference spectra to PAUS data, while Appendix D provides the list of flags used. We emphasize in the specific challenges of processing narrow-band images that prevent us from using generic software. We also describe the validation procedures and intensive operations at the computing centre. This paper is part of a series of scientific analysis which have been already published (Stothert et al. 2018; Eriksen et al. 2019, 2020; Tonello et al. 2019; Cabayol-Garcia et al. 2020; Alarcon et al. 2021; Cabayol et al. 2021, 2023; Johnston et al. 2021; Renard et al. 2021, 2022; Soo et al. 2021; Tortorelli et al. 2021; van den Busch et al. 2022) or are in preparation, which all used the data reductions presented here.

2 IMAGING DATA

Here, we describe the image data used in this work. First in Section 2.1 we describe the uncalibrated raw data that comes directly from the camera and secondly, in Section 2.2 we define the reduced data

---

[1]Note that we have extended our science analysis to $i_{AB} = 23.0$, or even to $i_{AB} = 24.0$ in some deep fields, as we have verified that our pipeline produces interesting science results for those magnitudes. But note that for the standard exposure times and number of exposures used in the wide fields, the SNR is usually below 3 for narrow bands with $i_{AB} < 22.5$. Because we use forced photometry over galaxies selected with deeper exposures in BB filters, even SNR < 3 fluxes are useful to measure photo-z (or constrain spectral energy distributions) with the 40 narrow bands, e.g. see Fig. 26.

produced by the PAUS data management system (PAUdm) after the instrumental detrending process that will be detailed in Section 3.

2.1 PAUCam raw exposures

PAUCam is an 18-detector imager camera (Padilla et al. 2019) with 40 narrow-band filters, covering the range from 450 to 850 nm in steps of 10 nm. After passing the optical system, the incident photon fluxes in the charge-coupled device (CCD) detectors are converted into photoelectrons which are then stored in the individual CCD pixel potential wells. During the read-out process of the detectors, the charge in each pixel is passed sequentially (clocked) to an on-chip amplifier that converts the charge into a voltage and amplifies this voltage before sending it to an analogue-to-digital converter (ADC). To reduce read-out time, PAUCam CCDs have four output amplifiers, one for each read-out region consisting of 4096 rows and 512 columns.

The data produced by the mosaic array is packed into multi-extension FITS (MEF) files (Wells, Greisen & Harten 1981) containing the pixel imaging data for the various types of frames, with associated metadata information in the header of each extension. Separate header-data units are created for each amplifier, resulting into a 72-extension MEF file of about 670 MiB.

During the afternoon, bias and dome flats are observed for calibration purposes. In the twilight, when the sky is too bright for science observations, a high-altitude standard star is observed with each detector-filter set for photometric calibration and total system throughput calibrations. Once the sky is dark enough, the main scientific field observations are taken.

The PAUCam raw data are organized in observation sets that group exposures from a single night. Typically for an observing night, the calibration frames (bias and dome flat-fields) and the science exposures are stored in separate observation sets. Observation sets allow also separating observations between PAUS and community observations in shared nights.

The three main types of raw exposures are as follows:

(i) **Bias frames**: Zero exposure time images for electronic pattern calibration.

(ii) **Flat-fields frames**: Images of a screen with a flat illumination at the dome for total throughput calibration.

(iii) **Science frames**: The scientific exposures with the target sources on sky. An example of a raw science exposure is shown in Fig. 1.

There are also other types of exposures [e.g. stacked focus or photon transfer curve (PTC) sequences] but these are not part of the regular data reduction process. So far we have observed 240 nights at the *WHT*, producing a total of 68 000 raw exposures.

2.2 PAUdm reduced images

After the image processing that will be described in the following sections, five more types of images are created, which are as follows:

(i) **Master bias**: Stack of corrected bias images electronic bias calibration (described in Section 3.2).

(ii) **Master flat**: Stack of processed flat-field frames for throughput variations compensation (described in Section 3.4.1).

(iii) **Reduced science**: The instrumentally calibrated science images (described in Section 3).

(iv) **Reduced mask**: The associated mask image with flag values per pixel (described in Section 3.6).
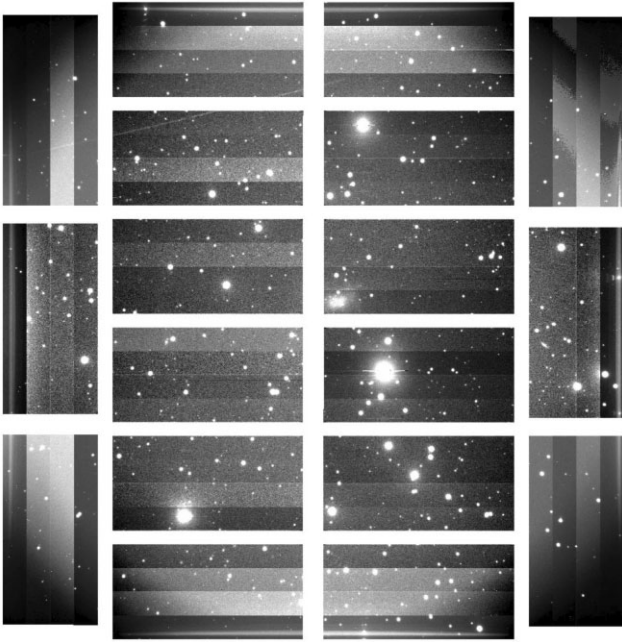
Figure 1. A sky exposure of the full 18-detector PAUCam mosaic. As a raw image, all instrumental signatures are present, and the 72 amplifiers can be identified as well as the vignetting from the *WHT* prime focus corrector.

(v) **Reduced weight**: The associated weight image (described in Section 3.6).

All reduced images contain only eight extensions, one per detector on the eight most central and illuminated CCDs. Additionally, the point spread function (PSF) models and astrometry catalogues are stored in the archive, together with other associated data that are recorded in the PAUS data base such as image zero-points (ZPs) and single-epoch detections.

3 INSTRUMENTAL DETRENDING

Raw images contain significant instrumental signatures that need to be neutralized or masked before carrying out any photometric measurement. This section describes the correction flow for all raw science exposures from PAUCam.

3.1 Gain calibration and PTC analysis

During the assembly of the PAUCam instrument (Casas et al. 2012), we performed the first PTC analysis with the method described in Janesick (2001) to estimate the gain for each of the four amplifiers in all detectors. Once on the mountain, with PAUCam mounted in the prime focus of the *WHT* under a stable installation, we repeated the test as the environmental conditions were different (low temperature, pressure, and humidity) and with the interaction with other devices and a different grounding. We decided to repeat this test on each observing run since 2015.

In order to process the PTC, we need to obtain dome flats with a broad-band (BB) filter, usually *g* or *r*, and obtain pairs of images with scaled exposure times, from 1 to 30 s, covering the range between very low counts until saturation. Due to the strong vignetting caused by the *WHT* prime focus optical system, PAUCam allocated narrow bands only on the eight central detectors. The external detectors use BB filters for calibration and guiding purposes. For this work, we will
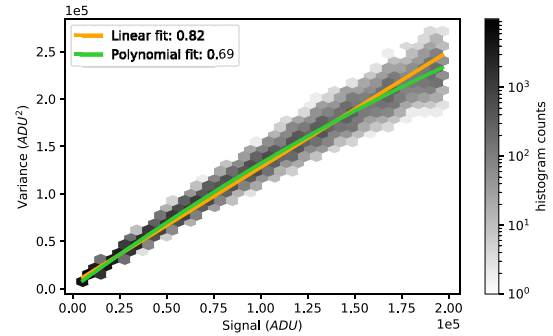
Figure 2. The relation between the mean flux value and its variance used to infer the gain as described in the PTC process for the first amplifier of CCD 1 in a test during the 19A period. The gain is measured as the inverse of the first-order component of the polynomial fit (green line). A linear fit over the whole flux range (yellow line) cannot be used to estimate the gain as it delivers a biased value due to the non-linear response of the detector.

focus on the central detectors with narrow-band filters. An analysis of the gain for the full focal plane (the 18 detector set) is described in Padilla et al. (2019). Each $2K \times 4K$ Hamamatsu CCD of PAUCam has four outputs, and each one must be analysed independently for the PTC analysis, so we analyse 32 regions from the eight central CCDs.

To compute the PTC we subtract the median overscan value for each image and average the pairs of images with the same exposure time to remove possible patterns and reduce the noise. A random choice of small squares of 100×100 pixels in the subtracted and averaged images are used to determine the mean signal [in analogue-to-digital units (ADUs)] and the variance (in ADU^2). The PTC represents the variance as a function of the signal as shown in Fig. 2. The gain can be evaluated by the fit between the variance and the signal. The inverse gain can be estimated as the slope between the variance and its signal, in units of e^- per ADU. As the detectors do not respond linearly in the high end of the flux range, as shown in Fig. 2, we have measured the slope in the fully linear regime of the detector range (<20.000 ADU) where the polynomial and linear fit estimate the same value.

Since the first installation of PAUCam in the prime focus of the *WHT* in 2015 we ran 16 PTC tests. The value of the gain varies from amplifier to amplifier but has remained stable in time across the various observing runs.

3.2 Overscan, gain, bias, and flat-field correction

The electronic read-out process adds an artificial pedestal signal across the whole image, biasing the values of each pixel by a constant number. In order to estimate this bias value that slightly varies from one amplifier to each other, we use the overscan section where only the bias signal is present. Then we estimate the value of the bias in each amplifier, computing the median row by row in the overscan region. To correct for low-frequency oscillations in the vertical read-out, an eight-row Gaussian filtering is applied, allowing a more accurate subtraction of the electronic bias as the single row median did not deliver accurate statistics of the varying bias value. The gain estimated from the PTC analysis is applied to each amplifier array, and ADUs are transformed to electrons in this early stage of the process. To compensate for read-out patterns that are present in all exposures, a master bias frame is produced, combining around 10 individual zero-exposure bias frames. As we have identified residual patterns in the first two or three exposures after a full read-out

system restart (typically once per day), the individual bias frames are analysed and those with abnormal levels of noise are removed from the median average. Finally, a master flat is produced from the individual dome-flat exposures that are taken every afternoon, prior to the night-time observations. We perform a median average of at least five individual flat-field exposures, to reduce noise fluctuations and cosmic ray hits. Scatter-light residuals found in the raw flat images needed to be removed, as this signature in the flat images is an additive component of the light, while the flat-field needs to contain only multiplicative factors of the main optical path. The process of removing the scatter-light will be described in Section 3.4.1, together with the correction of the scientific sky images. Once a clean master flat has been produced, we use it to divide the science exposures from the same night, flattening the response across the field of view.

3.3 Cross-talk calibration and correction

Since the four amplifiers are read in parallel it is possible that current is induced through magnetic fields between the different channels read at the same time, an effect that is commonly known as cross-talk (Freyhammer et al. 2001). The charges from the same row are read simultaneously in all four amplifiers of each detector and in all 18 detectors. Amplifiers 1 and 3 read the pixels from right to left while amplifiers 2 and 4 read in the opposite direction, due to its disposition in the detector. This allows us to identify unambiguously the pixels that are read simultaneously.

Another relevant instrumental effect occurs when too many electrons in a given pixel cause electrons to overflow to the nearby pixel wells in the same column. The average capacity of these detectors is $\sim 210.000\,\mathrm{e}^-$. This effect is known as saturation bleeding, producing an elongated shape around very bright stars. If the pixel signal is above 18 bit (current depth of PAUCam ADCs), it will saturate at $N^{\mathrm{sat}} = 2^{18} - 1$ (or 262 143) ADUs. With an average gain of 0.7, the final saturation value is limited by the ADC conversion limit and not by the full well. Saturation starts approximately on a $G_{\mathrm{AB}} \sim 10$ stars, affecting approximately 200 pixels for a $G_{\mathrm{AB}} \sim 9$ star due to bleeding. In the top panel of Fig. 3, one can see ghost images parallel to the elongated star image due to the cross-talk effect. Even though we can only see the ghost image from the saturated pixels, the cross-talk effect happens at all levels but is only visible when the signal is strong enough to stand out from the noise in the background.

To remove the cross-talk signal we need to estimate the induction ratios $r_{\mathrm{xy\_ij}}$ between each pair of amplifiers from detector x amplifier y to detector i amplifier j. Therefore, the signal of each amplifier would be the sum of its direct integration signal I^{int} plus the signal of all the remaining amplifiers in the mosaic being read out at the same time scaled by the induction ratio

$$I_{\mathrm{xy}}^{\mathrm{tot}} = I_{\mathrm{xy}}^{\mathrm{int}} + \sum_{i}^{18} \sum_{j}^{4} I_{\mathrm{ij}}^{\mathrm{int}} r_{\mathrm{xy\_ij}} \tag{1}$$

for a given set of detector x and amplifier y.

Although cross-talk analyses were made at the facility lab in Barcelona, the cross-talk ratios may vary once the camera is mounted in the telescope and thus we had to estimate them with sky data obtained when the instrument was mounted in the prime focus. In order to estimate the induction ratios between amplifier i and j of exposure z, we measured the average image background level in the target amplifier bg_i and subtracted it from the average level in the mirrored positions of the saturated pixels f_j, looking for any change in flux with respect to the rest of the image background due to cross-talk
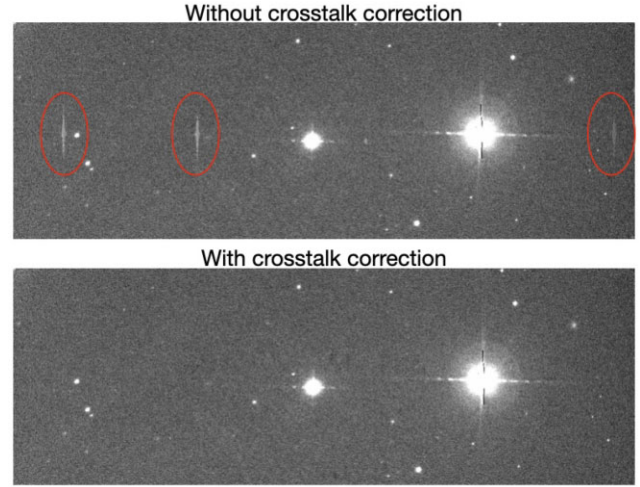
Figure 3. Top panel: An image crop showing an 8.5 magnitude star that saturates the detector producing bleeding and a visible cross-talk signal on the mirrored positions of the remaining three amplifiers. Bottom panel: The same crop with cross-talk correction enabled. The three ghost signals from the saturated star mirrored on the other amplifiers are completely removed after applying the corresponding ratios and subtraction.

such as

$$r_{\mathrm{ij}}^{z} = \mathrm{median}\{f_j(x,\,y) - bg_i\}. \tag{2}$$

where x, y are saturated pixels in amplifier image j of exposure z. If the mirrored position contains sources or is not flat enough, we discard the measurement. We combine the i, j ratio measured in all available images, weighting by the number of saturated pixels available in the each ratio as

$$r_{\mathrm{ij}} = \frac{\sum_z r_{\mathrm{ij}}^{z} N_{\mathrm{sat}}^{z}}{\sum_z N_{\mathrm{sat}}^{z}}. \tag{3}$$

We compute the ratios for the eight central detectors (where narrow-band filters are located) against the remaining 18 detectors and four amplifiers, resulting into a total of 2272 pairs. To remove the cross-talk signal, each amplifier needs to subtract the flux of the other 71 amplifiers by the corresponding ratio. Fig. 3 illustrates how the ghost produced by the bright star on the right is now removed after applying the calibrated cross-talk ratios.

To verify the correct implementation of the cross-talk calibration, we have measured the same cross-talk ratios over a data set that has been already cross-talk corrected. Fig. 4 illustrates the matrix of cross-talk ratios. It can be seen how the intra-detector cross-talk ratios are significantly higher than the inter-detector ones, as expected due to further distance of wiring in the electronics. We measure the cross-talk ratios in a large data set with more than 430 000 images, and apply those ratios to a different data set (as applying it over the same data set would give a zero residual by construction). We can see how the cross-talk ratios reduce from ~ 0.04 per cent to less than 0.002 per cent once the cross-talk calibration has been applied.

3.4 Scatter-light correction

The large narrow-band set of PAUS required engineering an efficient design to include all 46 filters, as described in Padilla et al. (2019). The filter exchange system has a total of 14 filter tray slots. Only five trays are used for narrow bands. There are also standard *ugrizY* BB filters. These are larger and they cover the full field of view each. So
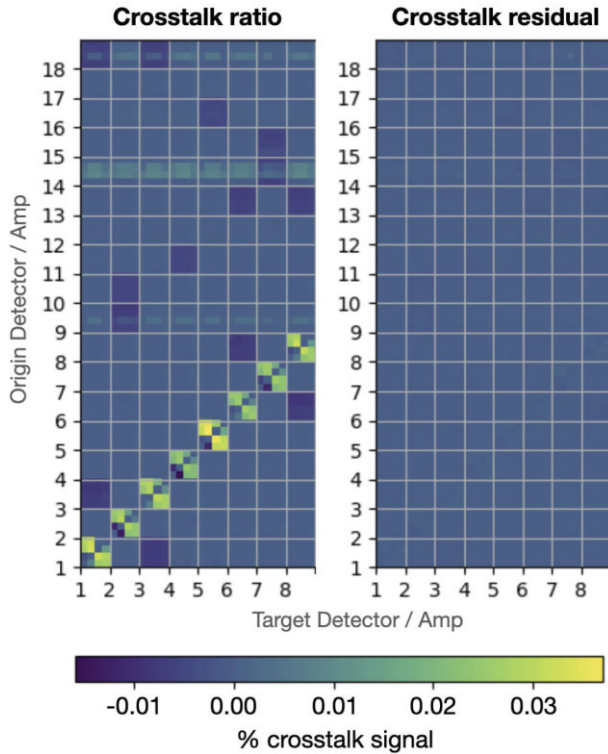
Figure 4. Left-hand panel: The cross-talk induction ratios used to apply the correction. Each box contains the 4x4 amplifiers of each detector against another detector. Right-hand panel: The residual cross-talk ratios measured on a different data set after cross-talk correction.

Figure 5. Top panel: A typical reduced science frame without scatter-light correction. Bottom panel: The flat-field science image preserving the flux from the stars and small galaxies, but affecting the larger extended sources and haloes.

one tray per BB filter. Note nevertheless that for the science fields in the PAUS we chose regions of the sky which already have deep BB *ugrizY* observations (from CFHTLenS, COSMOS, or KiDS) so we do not need to repeat those BB observations with PAUCam.

The final solution for narrow bands was to distribute multiple small filters, each covering a single detector in filter trays, such that with five trays and eight narrow bands in each covering the central most illuminated detectors, we could cover the entire wavelength range. This design caused a side effect in the quality of the images, as significant reflections from the lateral edge of the filter and the filter trays themselves created localized scatter-light at the edge of each detector image (see the top panels of Figs 5 and 6). This effect was identified quickly and the camera was reopened in mid-2016 to redesign the filter trays and minimize the problem. After the camera intervention, the background caused by scatter-light was reduced by a factor of 4 as it can be seen in Fig. 7. The scatter-light was not fully eliminated and specific processing (described in the following sections) had to be developed to correct the areas affected by this issue, without compromising the light from the sources we need to measure.

3.4.1 Implementation in flat-fields

Flat-fields can be divided into two main frequency components across the focal plane: a low-pass band due to vignetting and a high-ass pixel due to pixel variations (i.e. dust, dead, and hot pixels). Fortunately the scatter-light is in between with a mid-size frequency variation, so we are able to isolate and correct for it. First we use the BB flat-fields made of a single large filter, that does not contain scatter-light, to construct a vignette image, mostly caused by the prime
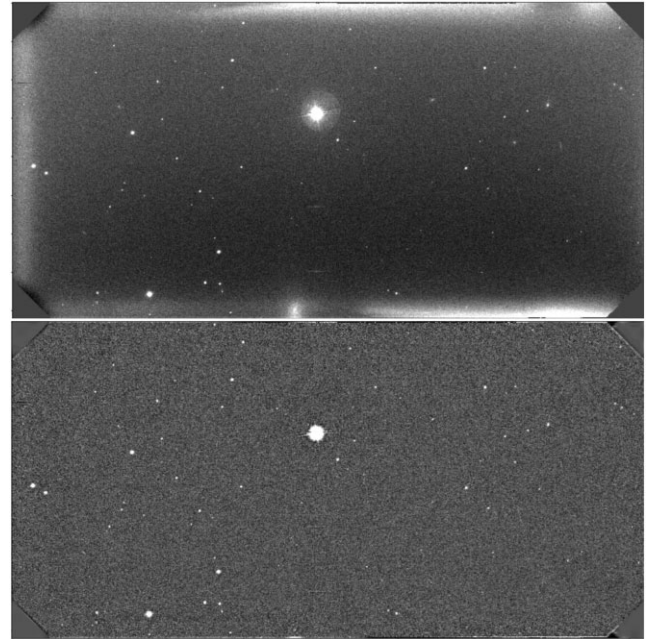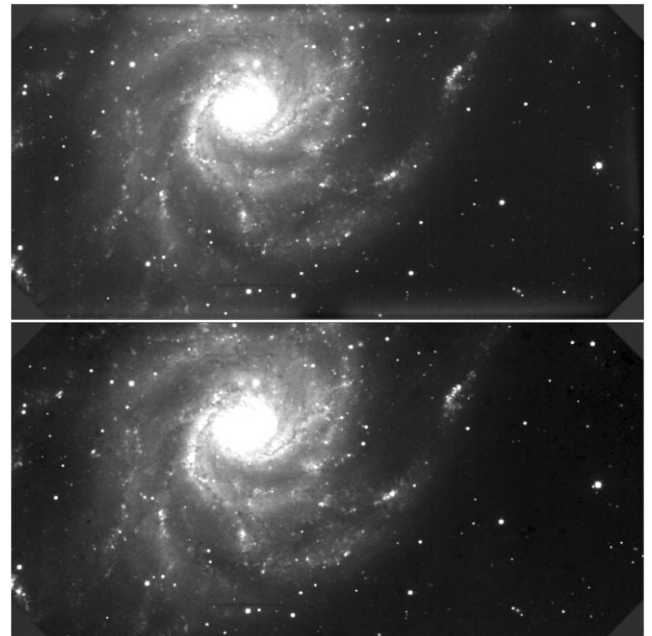
Figure 6. Top panel: A rare science frame with a very extended source (M101) without scatter-light correction. Bottom panel: A sky-flat corrected science image preserving the light from extended sources.

focus corrector optics. With a low-pass filtering we could isolate the vignette component and dismiss the high-frequency variations. Dividing the narrow-band flats with the vignette profile flattens the image and leaves the scatter-light as the lowest component in the image, ensuring that high-pass filtering would leave the scatter-light component only, which can then be subtracted from the flat. The resulting correction can be seen in Fig. 8.
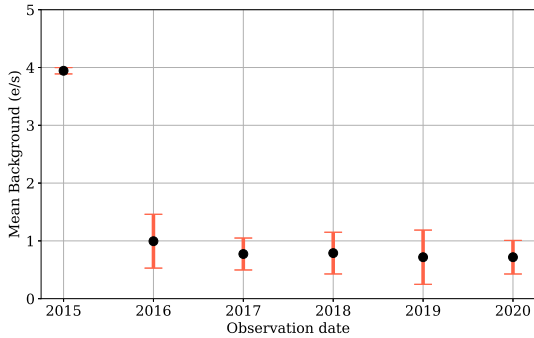
Figure 7. The mean background in the reduced PAUS images. With the original filter tray design, a substantial fraction of the light was reflected at the edges of the filter pieces and filter tray, causing severe levels of scatter-light in the science images. An intervention was carried out in mid-2016 and the mean background was reduced from $\sim 4\ e^-$/s to less than $1\ e^-$/s, as can be seen after 16B observations and beyond.
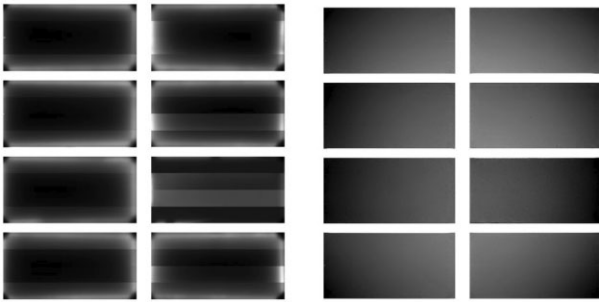

Figure 8. Left-hand panel: The raw flat-field exposure with severe scatter-light signal. Right-hand panel: The scatter-light corrected flat-fields preserving both the low frequency component of the vignetting and the high frequency of the pixel-to-pixel variations.

3.4.2 *Implementation in science images*

In the case of science images, the process is more complicated as large extended sources (i.e. large galaxies or nearby objects) may have a similar spatial frequency as the scatter-light. In the cases where the target sources are distant and small galaxies, a low-pass filtering, sigma-clipping the sources, isolates the scatter-light, without affecting the photometry of the small sources. This is similar to background subtraction techniques, with the difference that the codes used take into account the preferable direction of the scatter-light across the edges of the detector-filter system.

To preserve the flux from extended sources, a sky-flat correction was the most effective solution. To produce the sky flats, a large set of images were combined following a median average stack. We identified that around 50 images were enough to provide an accurate model of the background, including scatter-light. A combination of multiplicative factors (residuals from the dome flat) and additive components (mostly scatter-light) are present in the sky flats and therefore need to be separated to subtract and divide the images to detrend the effects properly. The complication of this method is that it requires multiple epochs to detrend an image, and the images stacked need to have similar background levels, which might not always be possible with shorter observations. However, when there are enough exposures from the same filter tray under the similar sky conditions, this method provides a more accurate illumination correction than the traditional dome flats.
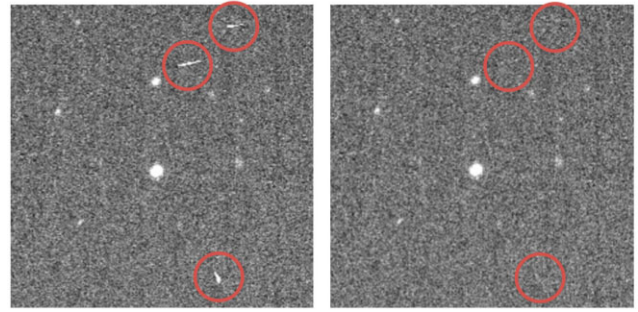
Figure 9. Left-hand panel: A crop of an image without cosmic ray masking. Right-hand panel: The same image crop after the Laplacian filter algorithm, successfully identifying, masking, and interpolating the cosmic ray hits in a single-epoch image. Red circles highlight the location of cosmics before and after detection and masking.

3.5 Cosmic ray detection, rejection, and masking

The fully depleted detectors are also very sensitive to cosmic rays impacting the silicon sensitive area and leaving a trace that will impact the photometry if it overlaps with the target source. The typical exposure times for PAUS are between 2 and 3 min, causing a reduced amount of cosmic hits in the exposure. However the camera is capable of performing long exposures thanks to its integrated auto-guiding system, and in these cases having a proper cosmic ray identification is critical. In any case, we perform a cosmic ray detection, rejection, and masking, following a Laplacian filtering algorithm known as L.A.Cosmic (van Dokkum 2001). The main idea is to take advantage of the specific sharp profile of the cosmic ray hits, due to the fact that the particle is not blurred by the atmosphere PSF as the rest of the photons in the image and can be highlighted with a Laplacian filtering. For the noise model described in the algorithm, we provide the measurement of the read-out noise taken in the overscan regions of the raw image. Fig. 9 illustrate how cosmic ray hits over a PAUS image disappear after being identified and masked with neighbouring information.

3.6 Image masking

In order to keep track of the history of each pixel, we attach a mask to each exposure image, of the same size as the original. The pixels in the mask are mapped into a bitmap where each bit corresponds to a certain issue. Bits 1–5 are left for SEXTRACTOR (Bertin & Arnouts 1996) flags, and will specify when a detected source is crowded, merged, under a halo, truncated or deblended. Bits 6–13 are set at the pixel or image level in the mask image to specify for cross-talk-corrected, detector cosmetics (dead or hot pixels), saturated pixel, rejected cosmic ray, highly vignetted area, close to edge, and very high distortion. Other bits are left for later photometry such as pixels or sources affected by scatter-light background, very high extinction, discordant measurements, astrometry issues, and noisy background. The full list of flags can be found in Appendix D.

To classify the cosmetics in the detector, a flattening process over the master flats is applied and pixels with less than 40 per cent of the normalized flux are added to the mask and flagged as bad pixels. For each science frame, we also produce the associated weight map, built from the master flat-field, with a very low weight value the pixels with mask values larger than 0. In this way, we assign a higher weight for pixels that have a better system response and neglecting the masked pixels.
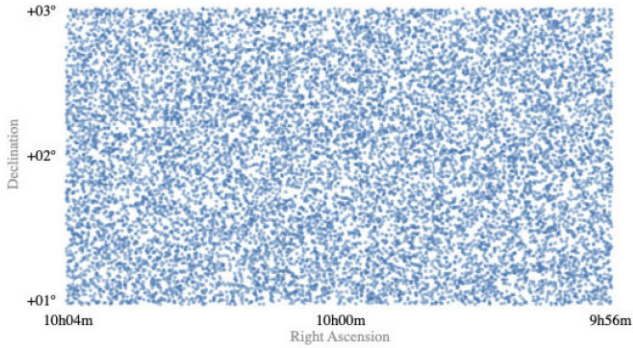
Figure 10. A uniform and dense star coverage in the COSMOS field from *Gaia* DR2 enabling an accurate reference for astrometric calibration.
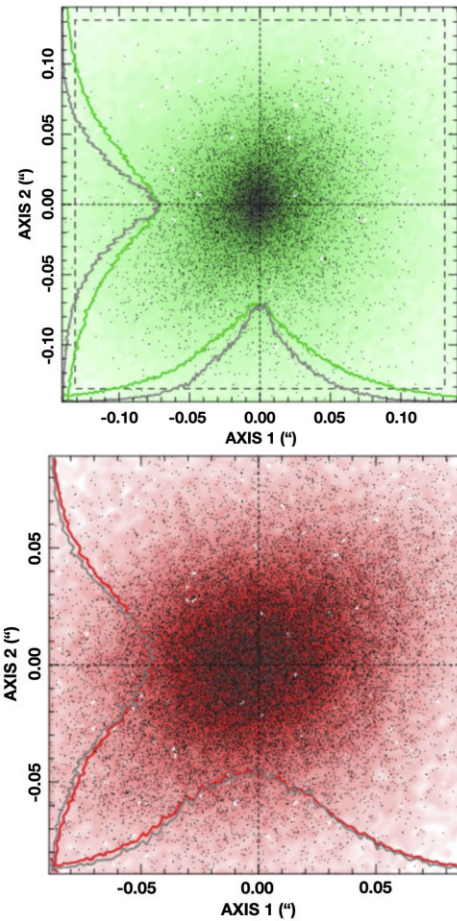
4 ASTROMETRIC CALIBRATION

Once the image has been cleaned and detrended for all known instrumental effects described in Section 3, we can proceed to correct for telescope pointing inaccuracies and optical distortions. There are three key elements in this process; the reference catalogue (Section 4.1), the calibration of the world coordinate system (WCS; Section 4.2), and the astrometry matching and correction of the single-epoch images (Section 4.3).

4.1 Gaia reference catalogue

The reference astrometric catalogue is the publicly available *Gaia* Data Release 2 (DR2) (Gaia Collaboration 2018). This data set provides the most accurate stellar catalogue, complete between $12 < G < 17$ and with a limiting magnitude of $G = 21$. It also includes proper motions, which can be critical to find an accurate astrometric solution in some situations. Nevertheless the observation period of PAUS and *Gaia* overlap in time, so accurate proper motions corrections are not as critical as with observations that are far apart in time. The whole *Gaia* DR2 was ingested into the PAUS data base, enabling a high-quality astrometric reference for any PAUCam observation in the sky.

4.2 WCS calibration

The raw mosaic exposures come with a base WCS (Calabretta & Greisen 2002) in its header that approaches the plate scale of PAUCam at the *WHT*. However, this default WCS is not enough to calibrate single-epoch exposures and we need to compute a more accurate WCS solution for the implemented focal plane geometry. Furthermore, the solution for each detector position, rotation and scale needs to be defined independently as the base WCS comes from the mechanical layouts and not from precise measurements of the built focal plane. For this purpose we use SCAMP (Bertin 2006) in a particular configuration (MOSAIC_TYPE LOOSE), giving freedom to each detector to move, scale, and rotate around the focal plane for a perfect match between the overlapping dithered exposures. We provide SCAMP all overlapping measurements in the PAUS reference sky location, the 2 deg^2 COSMOS field (Capak et al. 2007), where additional dithers were observed by the survey for validation purposes. This allows us to compute a precise solution with enough stars across the focal plane, accurately determining the position, rotation, and scale of the entire detector mosaic. The homogeneity of the reference astrometric catalogue used in the WCS


Figure 11. The position residuals corresponding to the difference between the measured position of each source after the astrometric calibration against its reference in the two coordinate axes (1: right ascension, 2: declination). The plot highlights the axis histogram distributions (colour: all sources, grey: high signal-to-noise ratio (SNR) sources). Top panel (green): The astrometric residuals between internal overlapping sources. Bottom panel (red): The astrometric residuals against the reference catalogue.

calibration can be seen in Fig. 10. Fig. 11 shows the measured errors in both coordinates when calibrating the WCS with 100 overlapping exposures with respect to the reference *Gaia* catalogue and internally to PAU. From this validation analysis we estimate an absolute astrometric accuracy of ~40-mas root mean square (RMS) and an internal consistency of ~30-mas RMS. This process only needs to be computed once (unless the astrometric reference catalogue is updated), producing a calibrated WCS of the PAUCam focal plane that is used during the nominal image processing.

4.3 Single-epoch astrometry and PSF modelling

At this stage, we perform the first extraction of sources in the single-epoch image, using SEXTRACTOR in a configuration mode specific for astrometry, with extended centroid measurements of windowed positions and a vignette matrix for each source.

To compute each exposure astrometric solution, we use again SCAMP but in a configuration that takes into account the pre-calibrated WCS (MOSAIC_TYPE SAME_CRVAL), helping the global mosaic solution and increasing the performance even with
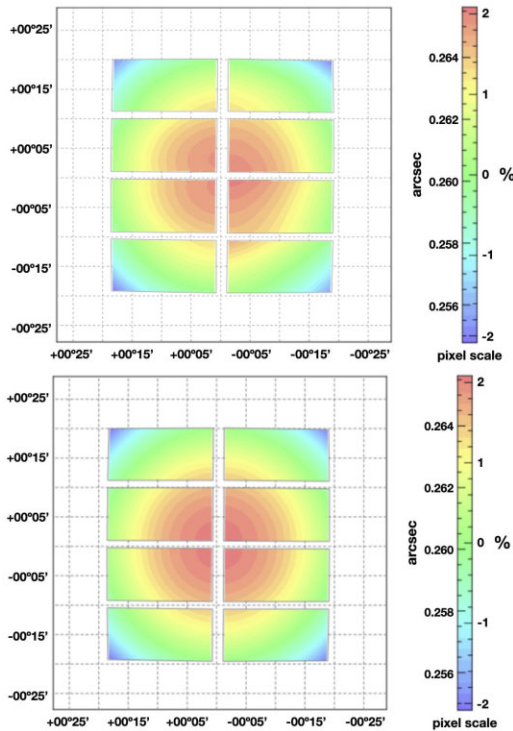
Figure 12. Top panel: The spatial variation of the pixel scale found by SCAMP in a single-epoch exposure. Bottom panel: The spatial variation of the pixel scale found by SCAMP with 100 overlapping exposures.

individual exposures. It provides an updated WCS header for the FITS file that includes the offset and distortion correction under a 3 degree polynomial fit. Even though SCAMP works best with overlapping exposures, we perform the astrometry solution on a single-epoch basis, as it already delivers 50-mas RMS of astrometric accuracy, below from what it is required to perform forced photometry ($<$100 mas). We have simulated a 50-mas aperture position error over an average galaxy size ($1''22$) and an average PSF size [$1''2$ full width at half-maximum (FWHM)], resulting into a tolerable 0.02-per cent error in the flux. This is only possible when using the base WCS calibration of the focal plane explained previously, serving as a first guess reference for the single-epoch corrections. Even though there are different filters over each CCD in the detector array, the global mosaic solution performs well and no chromatic distortions remain from SCAMP's solutions. Fig. 12 illustrates the estimated pixel scale variations across the focal plane, for the WCS calibration run with 100 overlapping exposures and for the individual single-epoch solutions. The difference between the two analyses is minor, indicating a good solution even when using a single exposure. The focal plane exhibits a strong distortion pattern with variations of 4 per cent in the pixel scale between the centre and the edges of the eight central detectors.

Using the same astrometric catalogue with additional centroiding and profile information, we model the PSF across the focal plane with PSFEx (Bertin 2011), which delivers a variable PSF model that can be reconstructed at any position in the focal plane. When there are not enough stars in a single detector to properly model the variations, the average PSF FWHM value is given. This information is key to obtain accurate aperture photometry as described in Section 6.3. The current aperture scaling uses the average PSF FWHM at each detector.

5 PHOTOMETRIC CALIBRATION

In this section, we briefly describe a key step in the data reduction process: the photometric calibration of the narrow-band images. The detailed description of the process is explained in Castander et al. (2022). Note that our calibration is not absolute but relative to SDSS photometry, which has been shown to be accurate to 1 per cent in the bands (*g*, *r*, and *i*) overlapping PAUS wavelength coverage (Padmanabhan et al. 2008).

5.1 Overview

In a photometric night, the atmospheric extinction correlates linearly with airmass. Traditionally the calibration of astronomical images is performed taking observations of photometric standards (Landolt 1992) observed at different values of the airmass to compute the extinction coefficient. After the extinction model of the night is computed, it can be applied to the rest of the science exposures. Due to the non-standard filter set of PAUS and the variety of observing conditions in the survey, we had to design a particular process that allows us to calibrate the fluxes in each image observed.

The approach presented in Castander et al. (2022) is to infer stellar templates from the SDSS BB data and compute synthetic narrow-band photometry from the stellar templates. Synthetic photometry is then compared to the PAUS measurements of the same stars to obtain a ZP for each star. Even though the stellar templates inferred from the SDSS BBs for a particular star can be wrong, combining ZPs from multiple stars cancels errors of individual star ZPs and delivers accurate narrow-band calibration.

As we compute the photometric ZP for every PAUS image independently and directly to the already calibrated SDSS synthetic photometry, we could observe under all sky conditions, even in non-photometric nights. This has been essential to maximize the use of available time in the observatory which is always limited and precious.

5.2 Implementation

The photometric calibration code is implemented inside the NIGHTLY pipeline. It has two main steps: the star photometry and the ZP calibration. For the first step, we run SEXTRACTOR on the instrumentally detrended and astrometrically calibrated images. We choose calibration stars that are moderately bright, comprising magnitudes between 14 and 19, that typically deliver a signal-to-noise ratio higher than 20. For such bright stars we do not need to optimize the aperture with complex and PSF-dependent methods that could be sensitive to the observation conditions or optical distortions in the focal plane. Instead we use a constant large aperture ($\sim$4$''$ radius) that gathers all the light from the star independently from the image PSF, ensuring that the truncated flux left outside the aperture is negligible even in the worst conditions tolerated by the survey. From simulations we verified that for the average survey PSF of FWHM $1''1$, the loss in flux is 0.03 per cent and in the worst case of $1''8$ seeing, the loss in flux is below 0.5 per cent. We tested various configurations of aperture sizes, background modelling and scatter-light correction, and the method using aperture photometry with a large aperture size was the most reliable across the different observing conditions. Once the photometry is processed, we perform a spatial matching with the SDSS DR12 catalogue (Alam et al. 2015), as these are the stars of

interest for the photometric calibration of the narrow-band images due to their accurate $\sim 1-$per cent relative calibration. We make use of the standard flags described in SEXTRACTOR documentation[2] and only clean measurements (flag = 0) are used to determine the star ZPs.

The second step is to compute the ZP for each star, as well as the combined ZP for each detector image by combining the individual star ZPs. For this process, we provide to the calibration code the narrow-band fluxes measured in the uncalibrated image, attached to the ID of the reference SDSS catalogue with its BB photometry. The calibration code computes the synthetic narrow-band fluxes for all the stellar templates and compares them to the observed uncalibrated fluxes of the selected stars in each image. It then computes the ZP for each particular set of star-templates. Finally, or each narrow-band image, the code returns a ZP per star, weighted by the χ^2 value of the stellar template fit and the synthetic BB fluxes of the stellar templates used.

The image ZP is computed by the median of all the stars available. The median is more robust than a SNR-weighted average as it reduces the weight of underestimated errors in the brightest stars that would otherwise dominate and possibly bias the final measurement. Weighting all the stars equally produces a more uniform spatial sampling of the ZP throughout the detector than inverse-variance weighting that typically determines the global ZP with just the brightest stars only sampling the detector in a few points. Even though we just use the image ZP in the calibration of the scientific catalogue, we store in the data base all the data for stellar photometry, individual star ZPs, and the combined image ZPs for validation purposes. In Section 6, we detail how to use the ZPs computed to obtain calibrated fluxes and propagate the corresponding error.

6 FORCED PHOTOMETRY

The most important aspect that maximizes photo-z accuracy is preserving the colours of its measured bands. The narrow-band images of PAUS deliver a low SNR (<3) at the target magnitude of $i_{AB} = 23$ and therefore we need to perform the photometry based on external reference catalogues. With the information from the reference catalogues, we can compute forced photometry for the PAUS images, measuring the same fraction of light in each band. This would not be possible with the PAUS images themselves for the faintest objects, as the shape could not be properly estimated on sources with such low SNR. From the reference catalogue we establish a consistent location, shape, and scale of each object and define an aperture that preserves the flux fraction at all wavelengths. The only constraint for this technique is to have good astrometry accuracy as the positions of the apertures for the sources are defined blindly (no object detection or centroiding is done in the PAUS images). As we defined in Section 4, we have a consistent astrometry at the subpixel level, enough for the purpose of forced photometry. As PAUS images have different seeing and PSF sizes, it is also important to model the PSF such that apertures are scaled accordingly for a constant flux fraction. This process is described in Section 6.3.

6.1 Reference catalogues

Performing a forced photometry technique requires a reference catalogue that overlaps with the images observed by PAUS. There are two main aspects that the reference catalogues must have: the necessary

parameters to perform the forced photometry accurately and the complementary galaxy lensing measurements that, in combination with the outstanding redshift accuracy of PAUS, deliver a unique scientific spot. Moreover, the reference catalogue must be complete down to the magnitude limit of PAUS such that the final combined catalogue has no target selection that could bias the cosmological measurements.

The selected PAUS fields are the Canada-France-Hawaii-Telescope Lensing Survey (CFHTLenS) fields W1, W3, and W4, the GAMA G09 field over the Kilo-Degree Survey (KiDS) North field, and COSMOS. In the CFHTLenS catalogue (Heymans et al. 2012), we have combined state-of-the-art reduction with THELI (Erben et al. 2013), shear measurement with LENSFIT (Miller et al. 2013), and photo-z measurements with PSF-matched photometry (Hildebrandt et al. 2012). In the case of the KiDS, we use its latest release DR4 (Kuijken et al. 2019), also with outstanding cosmological lensing measurements (Kuijken et al. 2015). And finally, in COSMOS we built a merged catalogue from Laigle et al. (2016) and the Zurich Structure & Morphology catalogue[3] for the accurate shape information. COSMOS has been our main validation sample and, with so many multiwavelength observations, it provided interesting photo-z tests with more than 70 bands (PAUS + COSMOS).

Furthermore, all these catalogues contain the information to perform the forced photometry measurements such as sky coordinates, a star-galaxy classification, a reference i_{AB} magnitude, the scale of the source deconvolved from its observed PSF, the axis ratio to estimate its intrinsic ellipticity, the position angle, and its Sérsic index (or an equivalent parameter that allows us to infer the Sérsic profile).

6.2 Background modelling

Accurate background subtraction is a key step to achieve precise photometry. The fluxes of the sources we need to measure sit on top of a floor of counts produced by the brightness of the sky, plus residuals of electronic bias and scatter-light. The latter is particularly important due to the configuration of the filter trays in PAUCam, causing significantly more light to be reflected and scattered in the edges of the narrow-band filter glass (see Section 3.4). An additional complication of the scatter-light is that it can produce a non-homogeneous background that is much harder to estimate and subtract. Underestimating the background will have a stronger impact on faint sources with few counts, as it will add a bias that scales with the size of the aperture.

For all reasons stated above, a careful background estimation needed to be implemented. Amongst the different background subtraction options, we designed an annulus around each source where we wanted to perform forced photometry. This provides an accurate estimate that takes into account large-scale variations of the scatter-light. The annulus had to be placed at a close distance to pick up smaller scale variations, but not too close to introduce flux from the source itself. Taking into account the typical size of sources we aim to measure ($1''.22$), we set a fixed limit for the inner annulus at 30 pixels from the centre of the target source. To get enough pixel statistics the outer annulus was set at 45 pixels. All bad pixels that fall into the background annulus are removed from the median statistics. A sample of an annulus with real pixels can be seen in Fig. 13. Additionally, to avoid blending sources to affect the estimate, we perform a sigma clipping in the remaining pixels

---

[2]https://sextractor.readthedocs.io/en/latest/Flagging.html

[3]Zurich COSMOS catalogue: https://irsa.ipac.caltech.edu/data/COSMOS/gator_docs/cosmos_morph_zurich_colDescriptions.html
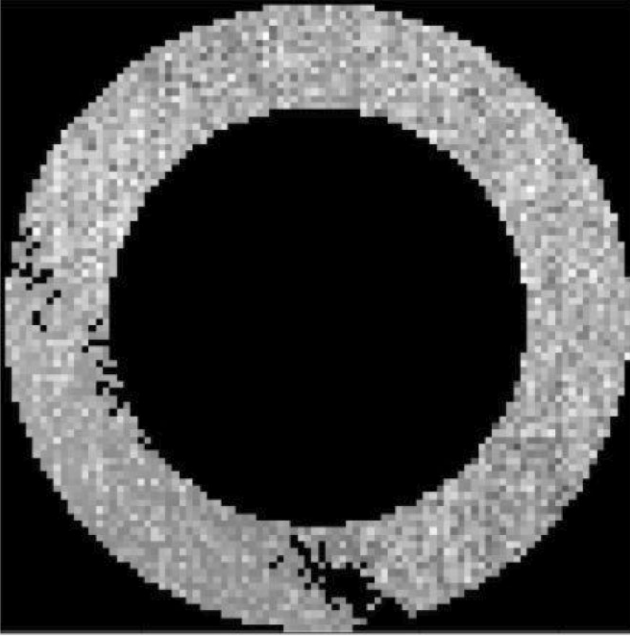
Figure 13. The annulus used to estimate the background of a particular source. It can be seen how pixels present in the image mask are discarded from the annulus and will not enter into the sigma clipping statistics.

of the annulus, leaving only background free pixels in the average calculation. The median of the available pixels provides our estimate of the background in each source. This average value is multiplied by the area of the aperture and subtracted from the measurement in the main photometry process. The standard deviation and number of pixels used in the estimate is kept as it is used in the flux error estimate.

An alternative neural network method to estimate and subtract the background of PAUCam narrow-band images proposed in Cabayol-Garcia et al. (2020) seems to deliver very accurate modelling of the background. It is not yet implemented in production and is planned to be integrated in the main pipeline for future releases.

6.3 Aperture scaling

We correct here for the impact of an elliptical PSF on an elliptical aperture for a galaxy. For simplicity, we present the results in circular coordinates ($a = b = r_0$), but these results can be extended to elliptical apertures by just rescaling the coordinate (x, y) units using elliptical coordinates:

$$x' \rightarrow x = r_0 q^{1/2} \cos \theta; \quad y' \rightarrow y = r_0 q^{-1/2} \sin \theta, \quad (4)$$

where $b \equiv qa$ is the smaller axis of the ellipses and θ the polar angle. This scaling is tested later on in Fig. 16 below. The goal is to measure the fluxes in an elliptical aperture that corresponds to the same fraction of the total light after taking into account the effects of the PSF at the time and sky location when the image was taken.

6.3.1 Sérsic profiles and aperture fluxes

We assume a Sérsic circular profile of slope n and scale r_0 for the surface brightness distribution:

$$I(r) = I(0) \exp\left[-(r/r_0)^{1/n}\right]. \quad (5)$$

The total luminosity in an aperture radius $r = A$ is

$$L(A) = 2\pi \int_0^{r=A} dr \, r \, I(r). \quad (6)$$

For a Sérsic profile, we can relate r_0 to the effective radius r_{50}, defined as the aperture which contains half the total light ($L(r_{50}) = L(\infty)/2$): $r_0 = [(0.86n - 0.142)\ln 10]^{-n} \, r_{50}$.

6.3.2 Convolved profiles

The observed surface brightness profile $I_o(r)$ will result from the convolution of the intrinsic $I(r)$ and the PSF kernel W_{PSF}. This is a two-dimensional (2D) convolution, so even when the image is originally symmetric, $I(\vec{x}) = I(x)$, the convolved image I_o might not be symmetric:

$$I_o(\vec{r}) = \int_{\text{Image}} d\vec{x} \, I(x) \, W_{\text{PSF}}(\vec{r} - \vec{x}). \quad (7)$$

For a circular PSF: $W_{\text{PSF}}(\vec{r} - \vec{x}) = W_{\text{PSF}}(|\vec{r} - \vec{x}|)$ we have:

$$I_o(r) = \int_0^\infty dx \, x \, I(x) \int_0^{2\pi} d\theta \, W_{\text{PSF}}\left(\sqrt{r^2 + x^2 - 2xr \cos\theta}\right). \quad (8)$$

6.3.3 Moffatt PSF

We will use a radial Moffatt PSF profile

$$W_{\text{PSF}}(r) = \frac{\beta - 1}{\pi \alpha}[1 + (r/\alpha)^2]^{-\beta} \quad (9)$$

which has a FWHM $= 2\alpha\sqrt{2^{1/\beta} - 1}$.

6.3.4 Effect of seeing on aperture

To obtain the same aperture flux $L(A)$ in equation (6) with a different PSF we need to change the aperture from $r = A$ to $r = A_o$ to account for the change in $I(r)$ to $I_0(r)$. This can be achieved by solving

$$L(A) = L_o(A_o) \equiv 2\pi \int_0^{r=A_o} dr \, r \, I_o(r), \quad (10)$$

where I_o is the convolved profile in equation (8).

6.3.5 Stellar apertures

A point source convolved with a Gaussian profile of width σ gives a Gaussian $I_o(r)$ and an aperture flux

$$L(A) = \sqrt{2\pi}\sigma[1 - e^{-0.5A^2/\sigma^2}] = \sqrt{2\pi}\sigma[1 - e^{-4A_F^2 \ln 2}], \quad (11)$$

where $A_F \equiv \frac{A}{2\sigma\sqrt{2\ln 2}}$ in units of the FWHM. The fraction of light is then $\frac{L(A)}{L(\infty)} = 1 - e^{-4A_F^2 \ln 2}$ which for $A_F = 1/2$ gives 50 per cent of light (FWHM). For a Moffatt profile the fraction is only a bit smaller.

6.3.6 Implementation

Computing the convolution integrals for the 2D Sérsic profile with a Moffatt PSF model for each measurement (120 per galaxy on average) is prohibitive as this operation may take several seconds when it is not optimized. We take advantage of the very optimal implementation of the fraction of light radius calculation in GALSIM (Rowe et al. 2015). Even when GALSIM claims this is only accurate down to a few per cent, with these smooth profiles we achieve precision better than 1 per cent in most cases. This is precise enough
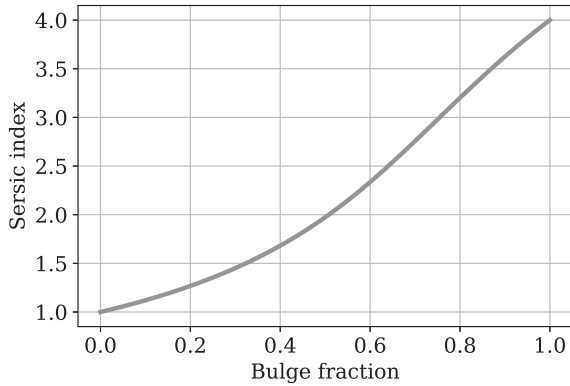
Figure 14. The relation between bulge fraction and Sérsic index assumed for CFHTLenS and KiDS reference catalogues.
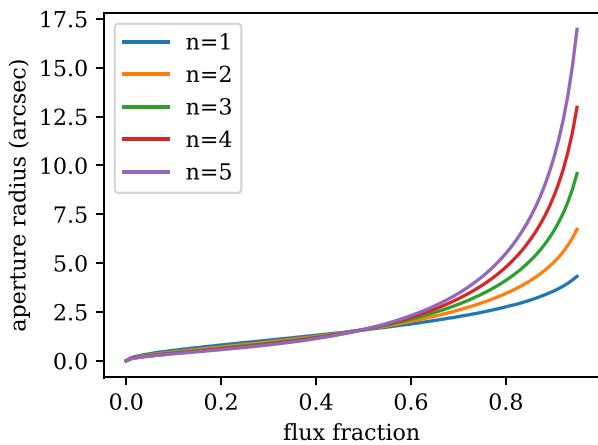
Figure 15. The aperture size radius as a function of flux fraction at various Sérsic indexes. The sample source is a galaxy with an effective radius $r50$ of $1''\!.5$ and a Moffatt PSF with a FWHM of $0''\!.7$ and a beta parameter of 4.75.

as our PSF model and the estimated effective radius will dominate the error budget and will determine the final precision we can reach in estimating the aperture.

The Sérsic profile is generated with the effective radius of the reference catalogue and the Sérsic index estimate. In case the Sérsic index is not available in the reference catalogue, we assign an index of 4 to elliptical galaxies and an index of 1 to spiral galaxies. In the case of CFHTLenS and KiDS, we derive the Sérsic index from the bulge fraction as shown in Fig. 14 (see Kannawadi et al. 2019).

The PSF is modelled as a Moffatt profile with a $\beta \sim 4.75$ (Trujillo et al. 2001) and the FWHM measured in PSFEx for each detector image. The GALSIM function calculateHLR (calculate half light radius) is applied over the convolved object and it allows to compute either the half light radius or a specific flux fraction. Typically we run the photometry over a variety of flux fractions, from 0.5 to 0.9. On average, the signal-to-noise is maximized around 0.625, even if the optimal SNR depends on each source. An example of how aperture size relates to flux fraction can be seen in Fig. 15.

As the r50 is defined as the effective radius on the major axis, we apply the same procedure for the minor axis, multiplying the r50 by the axis ratio defined in the reference catalogue.

We verified the aperture scaling method, including equation (4), with a simple error-free simulation, rendering the models described previously (Sérsic profile convolved with a Moffatt PSF) at a reference flux and performing aperture photometry. For multiple
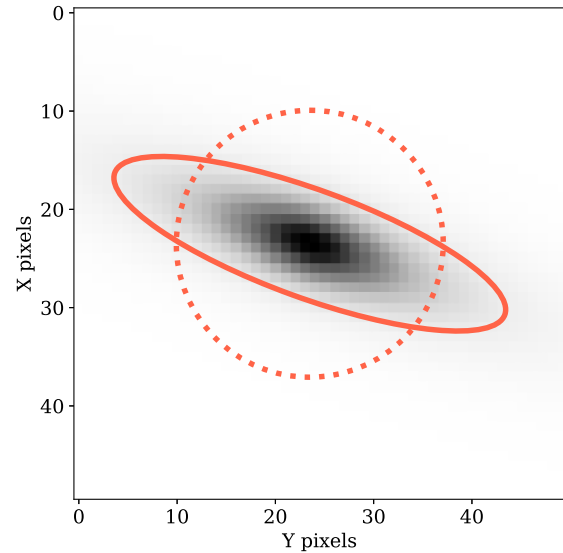
Figure 16. A simulation of an elliptical Sérsic profile of $1''\!.9$ and b/a = 0.2 convolved with a Moffatt profile equivalent to a PSF of FWHM = $1''\!.2$. Dashed line: A circular aperture containing 62.5 per cent of light accurate to 0.3 per cent. Solid line: An elliptical aperture at the same flux fraction with equivalent accuracy but with 60 per cent less area, improving the SNR of the measurement.

combinations of PSF FWHM, galaxy scales, and Sérsic indexes, the reconstructed flux was always accurate within less than 1 per cent. Fig. 16 shows an example of a convolved galaxy with a circular and elliptical aperture at 62.5 per cent of light. Both cases estimate and recover accurately the flux fraction but the elliptical aperture delivers a higher SNR.

6.4 Flux measurement and error estimate

In the forced photometry process, we blindly place the aperture in the image, without any additional centroiding. As the BB reference catalogue is significantly deeper than the narrow-band PAUS images, we can rely on good and complete positions and shapes from the reference catalogue. At this stage the narrow-band images successfully passed all astrometry quality cuts and therefore, we can use the calibrated single-epoch WCS solution to determine the pixel positions of the source with subpixel precision.

After computing both major and minor aperture radius from the previous step, we can perform aperture photometry with elliptical apertures. For that we make use of the PHOTUTILS library (Bradley et al. 2020), an affiliated package of ASTROPY (Astropy Collaboration et al. 2018). The aperture flux $f_{\rm tot}$ is measured on the reduced image, still containing the component of the background. The bitwise information of all pixels inside the aperture is propagated to the flux measurement flag.

At the same location, we estimate the background with an annulus as described in Section 6.2. The flagged pixels in the annulus do not enter into the combined flagging of the source, they are simply ignored in the statistics of the background calculation. The background modelling function estimates the median background flux $\overline{B}$, standard deviation $\sigma_{\rm bg}$, and number of valid samples in the annulus $N_{\rm bg}$.

The background-subtracted flux of the source $f_{\rm src}$ is therefore

$$f_{\rm src} = f_{\rm tot} - \overline{B} \cdot N_{\rm ap}, \tag{12}$$

where N_{ap} is the area of the aperture in pixels, taking into account the fractional overlap of the aperture and each pixel region.

The model of the error in the measurement σ_{tot} is composed by the source flux error (σ_{src}), the noise in the background (σ_{bg}), and the uncertainty in the estimation of the background.

The source flux error is assumed to be a Poisson process, and it is estimated as the variance of the flux inside an aperture. As we perform the measurement over single-epoch images, we assume there are no pixel-to-pixel correlations as it would appear on a remapped image. Therefore, the flux error of the source in electrons is

$$\sigma_{src}(e^-) = \sqrt{g \cdot f_{raw}}, \tag{13}$$

where g is the amplifier gain in the detector read-out system and f_{raw} is the raw detector counts of the source itself. As the measured units are flux rate in e^-/s, we need to scale the variance and the error such as

$$\sigma_{src}(e^-/s) = \frac{\sqrt{f_{src} t_{exp}}}{t_{exp}} = \sqrt{\frac{f_{src}}{t_{exp}}}. \tag{14}$$

Finally, we compute the total measurement uncertainty by adding in quadrature the flux error, the independent background noise, and the background estimate error in the annulus such as

$$\sigma_{tot} = \sqrt{\frac{f_{src}}{t_{exp}} + \left(N_{ap} + k\frac{N_{ap}^2}{N_{bg}}\right)\sigma_{bg}^2}, \tag{15}$$

where $k = \pi/2$ is the efficiency correction[4] for the median we used in our background estimation method.

6.5 Flux co-addition

So far all measurements were made at the single-epoch level. However, PAUS observes each area on the sky multiple times for every filter. On average we perform three passes in the main fields (W1, W2, and W3) and five times in our calibration field (COSMOS; Laigle et al. 2016). Observing the same area multiple times has some advantages, such as covering the gaps between detectors, increasing the signal-to-noise, rejecting outliers, reducing the density of cosmic rays, and increasing the dynamic range of the sources, as shorter exposure times will allow brighter stars not to saturate. This is done at the expense of an increased volume of data and a slower observing rate due to a constant read-out time of 20 s.

Most surveys combine their multiple layers at the image level, which is convenient for cosmic ray rejection (a median average almost completely removes all cosmic hits) but the stacking requires to resample the images, causing correlated noise which is complicated to model. Due to our objective of measuring very low signal-to-noise galaxies, we decided to stack the measurements at the catalogue level, performing all image measurements on individual exposures at their original pixel sampling. The combined measurements are called coadd fluxes.

Before combining the single-epoch aperture measurements, they need to be corrected to a standard system so all fluxes are consistent. Light rate from the same source may vary due to particular observing conditions like variations in the atmospheric extinction on non-photometric nights, different telescope elevation resulting into different observing airmass, or any other effect that varies the transmission with time. For this purpose, we have calibrated each

[4]https://wise2.ipac.caltech.edu/staff/fmasci/ApPhotUncert.pdf

image and assigned a multiplicative factor ZP and its corresponding calibration error σ_{zp}.

The calibrated single-epoch flux is simply defined as

$$f_{cal} = f_{src} \cdot ZP \tag{16}$$

and its calibrated error σ_{cal}, assuming non-linear error propagation with independent and not negligible variances, we derive

$$\sigma_{cal} = \sqrt{\sigma_{src}^2 \sigma_{zp}^2 + \sigma_{src}^2 ZP^2 + f_{src}^2 \sigma_{zp}^2}. \tag{17}$$

To avoid too small numbers that could require special data types, we added a magnitude offset of 26 in the calculation of the ZP and thus, one could derive the AB magnitude from the PAUS calibrated flux as

$$m_{AB} = -2.5 \log_{10}(f_{cal}) + 26. \tag{18}$$

However, with narrow-band photometry, we are often dealing with sources that have flux close to zero and magnitudes are inconvenient at that level. All the processing and archive of source brightness is done at the flux level.

Now that we have a calibrated flux and its associated error for the individual measurements, we can proceed to combine all the repeated measurements of the same source and band into a coadd flux and error using an inverse-variance weighted average such as

$$f_{coadd} = \frac{\sum f_{cal_i} \sigma_{cal_i}^{-2}}{\sum \sigma_{cal_i}^{-2}}, \tag{19}$$

where only non-flagged sources (Section 6.6) will enter into the combined measurements.

Assuming that the overlapping measurements are independent we estimate the coadd error as

$$\sigma_{coadd}^2 = \frac{1}{\sum_N \sigma_{cal_i}^{-2}}, \tag{20}$$

where N is the number of unflagged measurements to be combined.

Additionally we compute the reduced chi-square χ^2 as a measurement of consistency for the multiple measures:

$$\chi_{coadd}^2 = \sum_i^N \frac{(f_{cal_i} - f_{coadd})^2}{\sigma_{cal_i}^2}/(N-1). \tag{21}$$

All three forced photometry coadd parameters are stored in the data base for further processing and quality analysis.

6.6 Flagging

Throughout the whole processing of an image from its original raw state, we identify any possible cause that may affect the confidence of its value. To track each possible cause of problems, we use flags both in the NIGHTLY processing and in the Multi-Epoch and Multi-Band Analysis (MEMBA) pipeline. In the image calibration process of the NIGHTLY pipeline, we track the flags at the pixel level. Thus, we created a mask image where each pixel contains the flag values of the corresponding pixel in the science image. In order to track all possible flag combinations in a single value, we have mapped each flag condition to a bit in the value of the pixel, allowing for 16 different flags in a 16-bit depth image. In PAUdm, we have defined the following image-type flags:

(i) Cosmetics: Pixels not responding correctly to light, either hot ones that deliver constant high values or dead pixels that do not react to light inputs. Dust or imperfections in the detector mosaic or filter may appear here too.

(ii) Saturated: Pixels with so much flux that reached the ADC limit (18-bit in the case of PAUCam).

(iii) Cosmic Rays: Pixels identified as cosmic rays in the Laplacian filtering algorithm (van Dokkum 2001). Even though pixel values are interpolated from the neighbouring ones and cosmic rays may seem to have disappeared, the mask will keep track and the pixel value will not be used for science.

(iv) Vignetted: Areas in the focal plane with low transmission due to optical vignetting. The default value is set to 40 per cent.

(v) Cross-talk: Pixels contaminated by a strong signal of cross-talk from a related amplifier or detector (Section 3.3).

In contrast, in the MEMBA pipeline, we perform photometry of sources and flagging will take place at the catalogue level for each source measurement. The flags in the image that overlap the aperture are propagated to the measurement flag. Additionally we have defined the following catalogue-type flags:

(i) Edge: Source too close to the edge (<80 pixels) or partially out of the image array.

(ii) Distortion: Source in an area with strong optical distortion (>50 arcmin from the focal plane centre) or with an elongated PSF such that flux ratio in the aperture scaling may be inaccurate.

(iii) Scatter-light: Source with intense and spatially dependent scatter-light that could compromise the background subtraction. We estimated the presence of scatter-light in the background with two methods: variance ratio and ellipticity ratio. In the first method, we simply compared the variance in the annulus around each source compared to a global variance in the background of the whole image. If the ratio was above a certain threshold (typically 5 per cent), we flag the source. This method was effective to flag sources in scatter-light areas but was not efficient as it was overflagging sources that were simply on noisier areas. In the second method, we compute the ellipticity of the background image using the second-order brightness moments of the sigma-clipped stamp around each source, defined as

$$q_{xx} = \sum_{xy} I(x, y)(x - \overline{x})^2 \Delta x \Delta y, \tag{22}$$

$$q_{xy} = \sum_{xy} I(x, y)(x - \overline{x})(y - \overline{y}) \Delta x \Delta y, \tag{23}$$

$$q_{yy} = \sum_{xy} I(x, y)(y - \overline{y})^2 \Delta x \Delta y \tag{24}$$

and from these quadrupole moments we build the ellipticity of the stamp as

$$\epsilon = \frac{q_{xx} - q_{yy} - 2iq_{xy}}{q_{xx} + q_{yy} + 2\sqrt{q_{xx}q_{yy} - q_{xy}^2}}, \tag{25}$$

where real component measures deviations from circle along axes and imaginary component along the main diagonals (Bridle et al. 2009). Effectively this measurement estimates gradient variations in the background pixels around the target sources.
To obtain a reference ellipticity value, we compute the median ellipticity of the whole image scanning the detector in steps of 25 by 25 pixels and we flag those measurements with a background ellipticity larger than 10x the median of the image. This has proved to be an efficient method to track scatter-light residuals, flagging only sources with non-reliable background subtraction.

As described in Section 6.5, we skip all measurements that contain any flag inside its aperture for the combination of the coadd measurement. Approximately 9 per cent of measurements are flagged and therefore do not enter into the final coadd average. The full list of flags and its bit mapping value can be found in Appendix D.
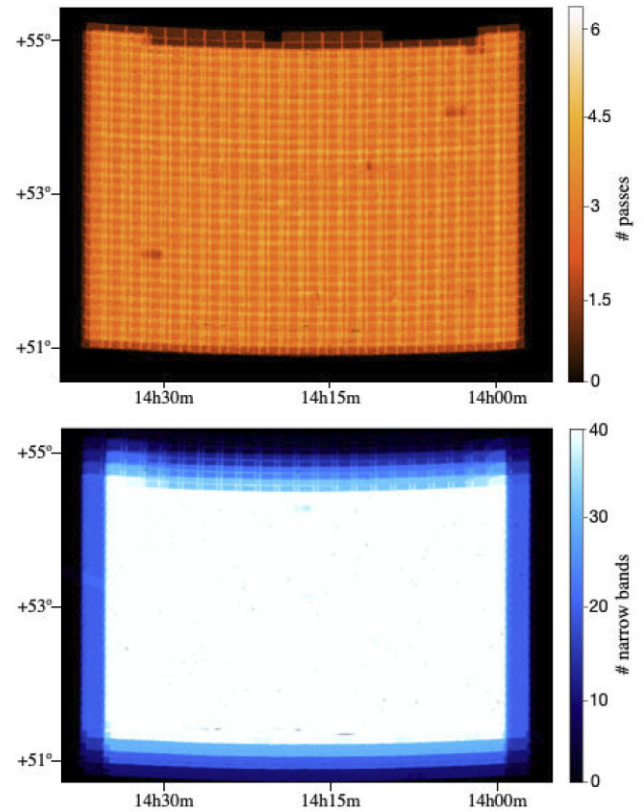
Figure 17. The PAUS mask in the 25 deg^2 W3 field. Top panel: The exposure mask of narrow-band NB655 indicating the number of overlapping observations of this filter on the sky. Bottom panel: The bands mask indicating the number of narrow bands with at least one observation in each part of the sky.

6.7 Survey mask

The particular layout of the narrow-band filters in the camera trays and the fact that CCD detectors are separated, leave gaps in the focal plane and result in a non-homogeneous coverage of the sky for each pass band. Additionally there are telescope pointing errors that result in a more inhomogeneous sky coverage. In large cosmology surveys that intend to identify statistical correlations of galaxy positions and densities, it is mandatory to accurately identify how the survey has tiled the sky with its thousands of exposures.

For this purpose, we have built a survey mask with two levels of information. First we generate the *exposure mask*, where we define for each filter how many times we observed each area in the sky for an entire field with a resolution of 5 arcsec. The mask is built taking into account variations in the system response from the flat-field and flagged pixels. This will create a complex mask that introduces effects like vignetting, bright stars that saturate, or corners in the detector not visible due to mechanical pieces in the optical path.

The second level of mask is the *bands mask* and it is built from the combination of all 40 exposure masks for each band. It represents the same area in the sky as the exposure masks but contains the number of bands available in each location of the sky with one or more effective units of exposure. An example of how both levels of mask can be seen in Fig. 17.

The survey masks are associated with MEMBA runs, as the resulting coadd catalogue is built with a set of images and this same set is the one used to build the masks. As we are performing forced photometry from an external catalogue, the selection of sources will

not depend on PAUS observations and we will need access to the survey mask provided by the external survey. However, the final survey mask must be created as an intersection between the external mask and the PAUS mask.

We used SWARP (Bertin et al. 2002) from Astromatic[5] to remap and build image stacks. We process the science images and their corresponding flat-field maps as the weight map of an entire field with SWARP. We obtain the exposure mask as the resulting combined weight map provided by SWARP in 'MAP_WEIGHT' mode. As a by-product we obtain the science-stacked image even though this is not used as part of the main science processing.

7 QUALITY ASSURANCE AND VALIDATION

One of the major challenges in PAUdm is the volume and complexity of the data to be processed and analysed. Contrary to spectroscopic surveys that can obtain a spectrum in a single observation (or stacking a reduced number of individual spectrums), each galaxy in PAUS is composed by more than 120 measurements. Furthermore a single problematic image can impact thousands of galaxies, altering its 40 narrow-band spectrum and causing catastrophic outliers in its photo-z determination.

Tuning an algorithm to process a small data set is simple, as one can manually verify the correct behaviour of the processing and its output result. However, with such large and complex data set, we need to build automatic control systems that verify that a particular code or configuration worked for the entire volume of data and raise an alarm or discard the data that did not meet some specific requirement.

7.1 Quality controls

At the time of writing this publication, the PAUdm system has processed more than 7 million images of PAU, and the number is increasing with additional reprocessing of data. The complexity and volume of this set requires an automated data quality control system to ensure that the data products meet the expected requirements. Although the most imaginative and cautious developer will miss the variety of circumstances that data from an observatory can contain. From closed petals of the main mirror, to dust from the Sahara desert, or even vapour condensation in the entry window of the camera in extreme weather conditions – these are some of the unpredictable conditions that we must catch to reject bad quality exposures and request observations to be repeated.

With this particular aim, we built a quality control system associated with the NIGHTLY pipeline. We define the following quality control tests with the corresponding tolerance limits in each metric to classify an image as valid:

(i) Readnoise: Check that electronic readnoise is under specification. This is measured in the overscan region of each amplifier. The default limit is set to 20 e⁻.

(ii) Flat-field level: Check that the flat-field image is illuminated in the correct range of values. Too bright illumination could result into saturation and too faint illumination would increase the noise of the master flat-field. The default range is set between 1000 and 120 000 ADUs.

(iii) Saturation: Check that the science images do not contain too many saturated pixels. A certain amount of saturated pixels are expected due to bright stars in the field. However, too many saturated

pixels are indicative of an issue in the exposure time, electronics, or target selected. The default limit is 0.1 per cent of saturated pixels (average is ∼0.02 per cent).

(iv) Cosmic rays: Check that the cosmic ray detection algorithm does not classify too many cosmic ray pixels. Issues in the electronics or very noisy images may affect the sensitive Cosmic Ray (CR) detection algorithm and end up with too many pixels being classified as cosmics. The default limit to reject an image is 1 per cent of pixels (average is ∼0.05 per cent).

(v) Astrometry: Check that the contrast (the ratio of the amplitude of the detected peak to the amplitude of the second highest peak found in the cross-correlation) and χ^2 to the reference catalogue are good enough to ensure SCAMP found a reliable solution. This is critical as we rely on the single-epoch astrometry and images with high extinction may end up with too few stars to deliver a solution. The default limits are contrast greater than 3 and reference χ^2 below 50.

(vi) Seeing: Check that the average image PSF FWHM measured by PSFEx is below a certain value. Large PSFs reduce the signal-to-noise and limit the target sources of interest. The default limit is 1″.8.

(vii) Calibration stars: Check that there are enough stars matched with SDSS to be used for the photometric calibration. Images observed under high-extinction atmospheric conditions may reduce the number of available stars. The default limit is set to 5 stars and approximately 5 per cent of images do not meet the required value.

(viii) ZP error: Check that the estimated error in the photometric ZP is constrained. An unusually high ZP error may be due to a non-uniform response across the detector. The default limit is set to 0.2 (flagging ∼3 per cent of the images).

The quality controls processed in each job are aggregated and propagated to the parent jobs so quality issues in large processing sets with many dependencies can be tracked easily.
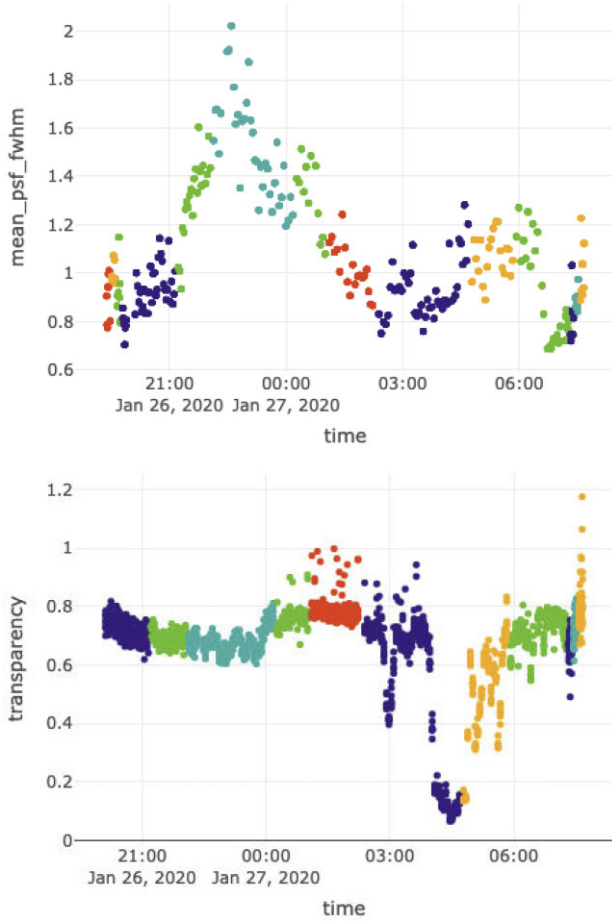
7.2 Nightly report

Periods of observation typically last for 1 to 2 weeks. During this time it is very important to provide feedback to the astronomers in the observatory in the shortest time. This was one of the key constraints in the design of the PAUdm system and we managed to process the whole night data set before the next night starts (∼8 h).

The NIGHTLY report is a web-based application that provides feedback to the astronomers on the quality of the data from the previous night so that observers can reschedule observations that did not meet a certain quality. It has also been critical to identify issues in the camera or telescope that were fixed with minimal delay for the remaining observing run. The NIGHTLY report has default quality limits necessary for the survey and generates a report file that can be ingested directly to the PAUCam control system for rescheduling targets. The application also displays statistics and evolution plots for each night, with options to adjust the metrics to be analysed and its time span, such as the ones in Fig. 18. A total of 35 parameters can be displayed to help PAUS astronomers understand the atmospheric, weather, and instrument behaviour from any previous night.

In addition to the quality checks described in Section 7.1, the NIGHTLY report displays the processing status of the main blocks in the NIGHTLY pipeline: detrending, astrometry, PSF modelling, and photometric calibration. The status in each block can be used as quality cut (i.e. repeating all observations where PSF modelling failed in any detector). Approximately 40 per cent of the exposures did not meet the image quality requirements imposed by the survey and had to be rejected due to bad weather or any other possible issues.

Figure 18. Some of the quality control plots that the nightly report produces as a result of the nightly pipeline processing. This data corresponds to one of the last nights of PAUS before the COVID-19 pandemic break. Each colour represents a different narrow-band filter tray. Top panel: Evolution of the atmospheric seeing and how it got quite high before midnight, stabilizing low before the end of the night. Bottom panel: Evolution of the transparency, where some clouds entered at around 4 am, reducing the absolute transparency to almost 10 per cent.

7.3 Comparison with SDSS and VIPERS spectra

A very reliable reference to validate PAUS measurements are high-resolution spectra from external surveys. The selection of the calibration fields was in fact driven by the overlap with data sets of carefully calibrated spectra. Our main validation field COSMOS contains ~17 000 sources with good reference spectra. Additionally our main fields contain other data sets such as VIPERS (Le Fèvre et al. 2013) in W1 or DEEP2 (Newman et al. 2013) in W3, complementing the already extensive set in COSMOS. From the calibrated spectra, we can infer the synthetic narrow band and perform a direct comparison to the PAUS measurements. Some examples are shown in Appendix C. Galaxy sources with a spectrum provide accurate spectroscopic redshift estimations which allow us to know where the expected emission and absorption lines are, providing additional confidence on a particular spectral shape.

We use spectroscopic data from SDSS and VIPERS to predict the MEMBA flux measurements. We then compare the predictions with the actual measurements to test MEMBA and the previous data calibration steps.

7.3.1 Synthetic narrow-band photometry

In order to compare spectra with PAUS narrow-band photometry, we need to generate synthetic photometry from the high-resolution spectra. This provides a high-quality reference to compare with any passband, especially valuable with the non-standard PAUS filter system. This comparison will only be possible with objects that both are observed by a spectroscopic survey and by PAU.

The first step in the process of generating the synthetic bands is to retrieve and homogenize the spectral data. In our case, we have converted all fluxes to a more common f_λ with units of erg cm^{-2} s^{-1} Å^{-1}. Generally, the data set contains the wavelengths where the spectrum is sampled, the fluxes, the noise (or inverse-variance), and a mask. Optionally SDSS also includes a measurement of the sky that allows identifying possible contamination in strong emission or absorption lines. Second, we interpolate the bandpass response $R(\lambda)$ to the sampling of the spectral data, as these two are not necessarily in the same space. Then we mask both the spectral fluxes and the passband with the flagged measurements from the spectrum mask. At this point, we can compute the integrated average flux density of the source at the specific passband in erg cm^{-2} s^{-1} Hz^{-1} such as

$$\langle F_\nu \rangle = \int \frac{f_\lambda R(\lambda) \lambda^2}{c} \mathrm{d}\lambda \tag{26}$$

and its associated integrated response

$$R_\mathrm{i} = \int R(\lambda) \mathrm{d}\lambda. \tag{27}$$

Finally we can compute the synthetic magnitudes in the AB system with the following transformation:

$$m_\mathrm{syn} = -2.5 \left(\log \frac{\langle F_\nu \rangle}{R_\mathrm{i}} \right) - 48.6. \tag{28}$$

It is also important for the statistical analysis to estimate the error of each synthetic band. As the flux in the spectrum has been weighted by the response of the transmission, we must weigh the noise in the spectrum by the relative transmission throughout the entire passband:

$$\sigma_{f_\nu}^2 = \int \frac{R(\lambda)^2 \sigma_\lambda^2 \lambda^2}{c^2 R_\mathrm{i}^2} \mathrm{d}\lambda, \tag{29}$$

where σ_λ is the noise in the high-resolution spectrum. We can approximate the magnitude error such as

$$\sigma_{m_\mathrm{syn}} \approx 1.0857 \frac{\sigma_{f_\nu}^2}{\langle F_\nu \rangle / R_\mathrm{i}}. \tag{30}$$

Following the previous procedure, we compute the photometry over all VIPERS spectra and all SDSS spectra that overlap with PAUS over the 40 PAUS narrow-band set and the SDSS and CFHT BB systems. We have flagged all measurements where the overlap between the systems response and the unmasked spectrum is below 70 per cent. An example of synthetic spectrum with SDSS over PAUS narrow bands and other BBs is shown in Fig. 19. More examples together with PAUS real observations are shown in Appendix C.

7.3.2 Recalibration of spectra

To account for remaining aperture or PSF effects in the measured spectra, we use total broad-band photometry in the corresponding reference survey (SDSS or VIPERS) to recalibrate each individual spectrum. To do this, we first estimate synthetic BBs from the spectra, $F_\mathrm{S}(BB)$, as shown in previous section. We then use the BB measured
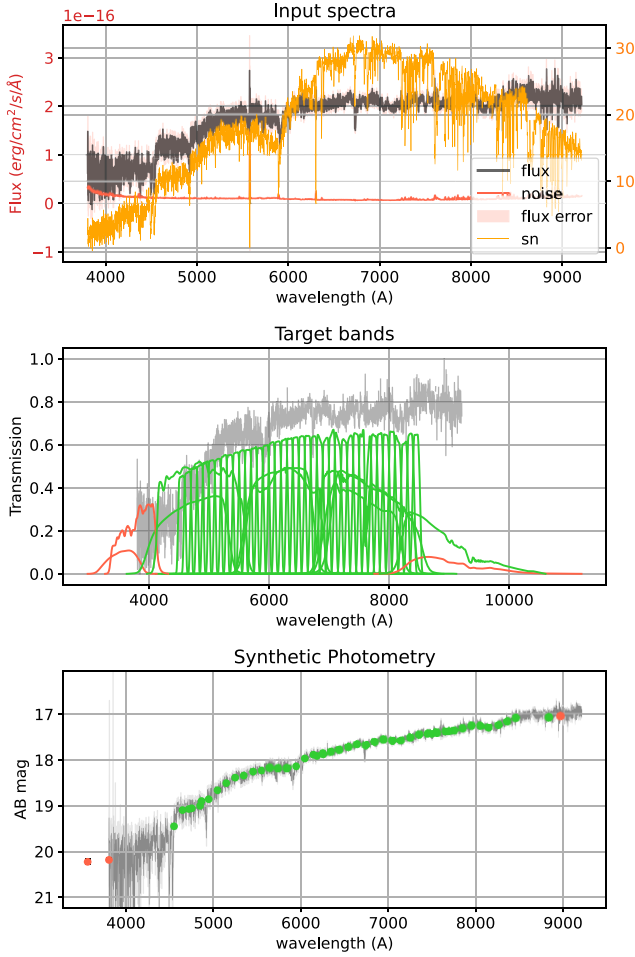
Figure 19. Top panel: A galaxy sample at $z = 0.14$ from SDSS used as input for synthetic photometry. The left-hand axis (in red) represents the flux and its error while the right-hand axis (in orange) represents the SNR. Middle panel: The target bands are the 40 narrow-band set from PAUS plus the two BB systems from SDSS and CFHT. Bottom panel: The computed synthetic photometry from the high-resolution spectrum. The bands without enough unmasked samples from the spectrum are marked in red.

flux $F_O(BB)$ to find a multiplicative ZP, ZP, which is in general different for each BB:

$$ZP(BB) = \frac{F_O(BB)}{F_S(BB)}. \tag{31}$$

We use ZP to rescale each individual spectrum. In the cases where we have 2 (or 3) BB measurements fully within the spectra wavelength coverage, we combine them using a fit to a linear (or cubic) function $ZP = ZP(\lambda)$, where λ is the mean of the bandpass response $R(\lambda)$. Each synthetic narrow-band λ_{NB} from the spectrum is rescaled by $ZP_{NB} = ZP(\lambda_{NB})$. The mean recalibration is only a 2 per cent offset with a 5 per cent scatter: $ZP = 1.02 \pm 0.05$. Similar results are found for VIPERS (see Castander et al. 2022 for more details).

7.3.3 Aperture corrections

Once the SDSS spectra are recalibrated with equation (31), we also perform aperture correction of the amplitude of each individual spectrum (S) to the PAUS measurements. This is a fit to a linear

constant $A = A(S)$

$$A(S) = \frac{\sum_i f_{PAUS}(S, i) f_{SDSS}(S, i)}{\sum_i f_{SDSS}^2(S, i)} \tag{32}$$

between PAUS raw fluxes f_{PAUS} and SDSS rescaled synthetic spectral f_{SDSS} (including the spectral recalibration). The sum is over individual PAUS measurements i in a given spectrum (S) and it uses inverse-variance weighting $w_i = 1/\sigma_i^2$, where σ_i is the joint error (from SDSS and PAUS) added in quadrature. Typically there are 200 PAUS independent measurements (corresponding to the 40 narrow-bands in the 5 dither exposures available in the COSMOS field) for each SDSS spectrum.

We study the distribution of values of A for different SDSS star calibration spectra and 42 420 independent measurements for PAUS run #955 in COSMOS. We find a mean value and scatter of $A = 0.98 \pm 0.02$, which indicates that PAUS data is in very good agreement overall with the SDSS calibration within 2 per cent overall scatter.

Fig. 20 shows a comparison of SDSS, VIPERS, and PAUS estimated narrow-band fluxes for two typical spectral energy distribution examples of galaxies (at $z = 0.4$ and $z = 0.5$) with $i_{AB} \simeq 20$. One can see the corresponding figure for stars in fig. 12 of Castander et al. (2022).

7.3.4 Colour terms

We now check for any residual differences as a function of narrow-band wavelength λ using galaxies with SDSS and VIPERS spectra. Fig. 21 shows the mean and scatter ZP difference for each narrow-band λ:

$$ZP(\lambda) = \frac{\sum f_{PAUS}(\lambda) f_{SDSS}(\lambda)}{\sum f_{SDSS}^2(\lambda)} \tag{33}$$

between PAUS raw fluxes f_{PAUS} and SDSS or VIPERS rescaled synthetic spectral f_{SDSS} (including the aperture correction A in equation 32). The sum is over all individual PAUS measurements (42 420 in total) and uses inverse-variance weighting $w = 1/\sigma^2$, where σ is the join error added in quadrature.

In the top panel, we use SDSS calibration stars which show a much smaller scatter than SDSS galaxies (middle panel). This is because we use fainter galaxies and also because the aperture correction becomes more important for extended objects.

We find a small residual colour tilt between the SDSS and PAUS narrow-band systems when using SDSS star spectra to compare

$$ZP(\lambda) = 1.05 \pm 0.04 - (0.05 \pm 0.04) \left(\frac{\lambda}{650 \text{ nm}} \right) \tag{34}$$

which is consistent with unity within errors. This correction is not applied to the data in MEMBA, but we state it here as a potential source of correction that should be applied, if needed. The scatter between the 40 bands after correcting for this linear residual slope is just 0.8 per cent. The scatter increases from 0.8 per cent to 1.1 per cent without this linear colour correction.

7.4 Forced aperture inspector

With such a large narrow-band set and with the additional overlapping exposures, each object depends on the correct reduction and calibration of hundreds of images. There are many things that can go wrong, even if all quality tests passed in the processing of an image. It is critical to identify the source of issues that may end up in outliers on PAUS spectra and that the photo-z code could misinterpret as a
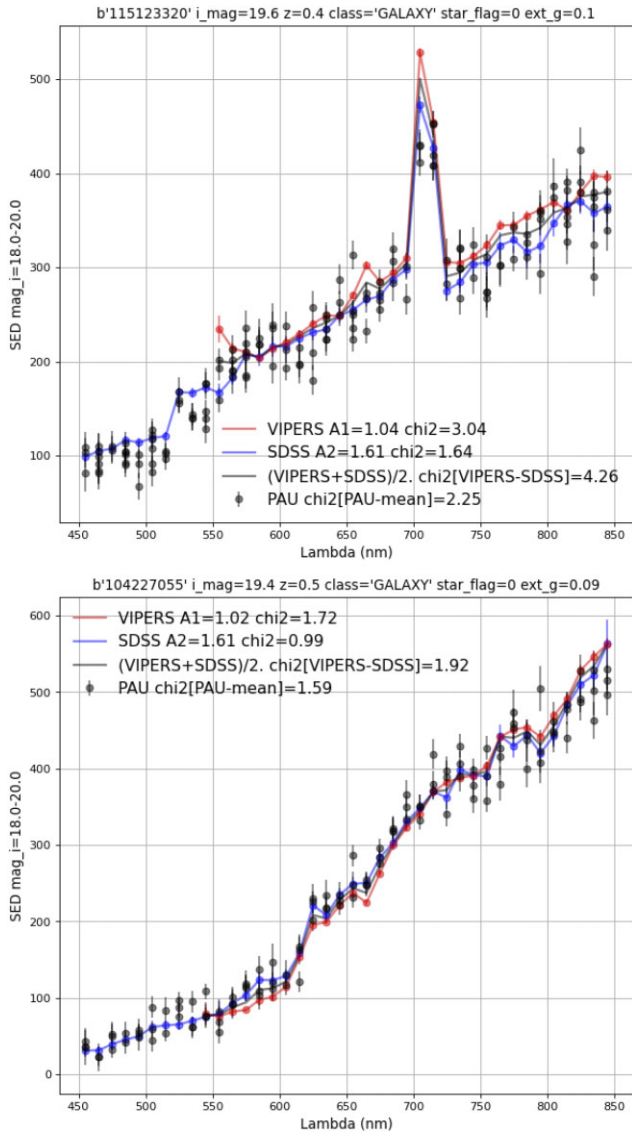
Figure 20. Example of validation spectrum for the photometric calibration in PAUS (points with error bars corresponding to five independent exposures before co-addition in Section 6.5) with synthetic narrow-band photometry from SDSS (blue) and VIPERS (red) galaxy spectrum. SDSS (VIPERS) spectra have been multiplied by A_1 (A_2), as shown in the labels, to account for possible differences in the aperture used in each observation. The 'chi2' label shows the normalized χ^2 as compared to PAUS data.

physical feature of the galaxy, resulting in a catastrophic redshift determination.

We built another quality control web application where PAUS measurements for a particular source are displayed together with the synthetic measurements from reference spectra. When the redshift is available, we also display the emission and absorption lines in the expected wavelength position. This allows us to visually inspect PAUS data and easily identify outliers or discrepancies with the reference spectrum. The application allows clicking on a particular suspicious PAUS band, displaying the measurements that contribute to that band. Clicking again on a single measurement will display the portion of that particular image together with the aperture done in MEMBA and all quality parameters associated to the image or observing conditions. Very rapidly we can identify issues from the
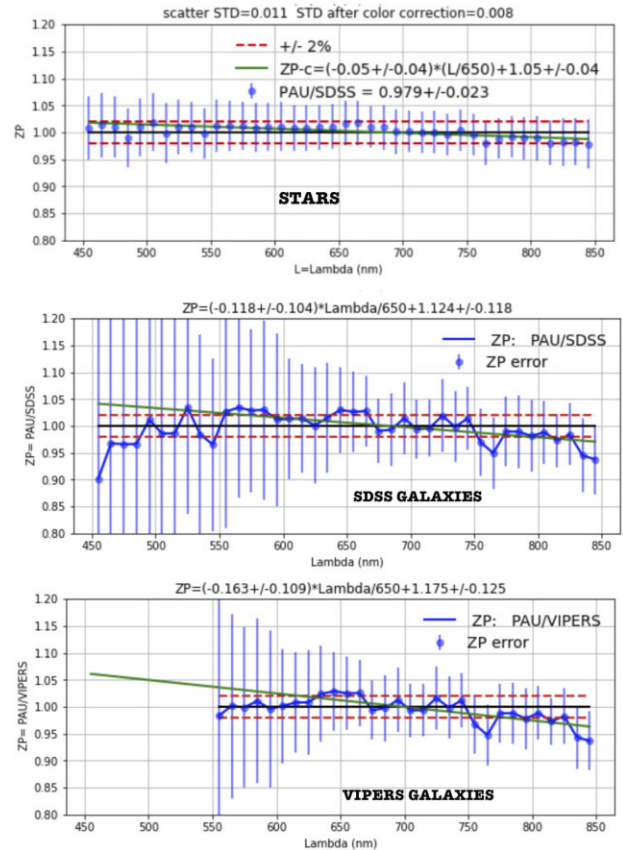

Figure 21. A validation study of the photometric calibration in PAUS (ZP = PAU/SDSS) using synthetic narrow-band photometry from SDSS stars spectra (top panel), SDSS galaxy spectra (middle panel), and VIPERS galaxy spectra (bottom panel). Blue error bars indicate the scatter in the values, which is much larger for galaxies.

final PAUS spectrum to the original image that contributed to every measurement.

As this resulted in a very powerful tool and many scientists from the PAUS collaboration contributed to it, we added a reporting system where people inspecting the sources could specify the issue that caused a particular outlier. The list of possible issues grew up as we learned more and more about the PAUS data, and we ended up with a list of 18 possible issues, such as scatter-light, blended source, cross-talk, astrometry issue, and more. Even though this is a subjective test that required some training, we could extract valid statistics and correct for multiple systematics that originally caused trouble. An example of the 40 bands of a galaxy in the forced aperture inspector can be seen in Fig. 22.

7.5 Duplicate observations test

Repeated exposures over the PAUS fields have been used to validate the final MEMBA photometry. These repeated measurements correspond to the independent exposures that are co-added in Section 6.5. We have between 3–10 independent fluxes for the same object in each narrow band. We use these catalogues to build a sample of over a million pairs of duplicate (repeated) measurements of the same object. There are 45 pairs of measurements for the same object and about 40 000 separate objects, before cuts and masking. We identified duplicates as objects with the same reference ID (which means apertures with the same position in MEMBA). We select
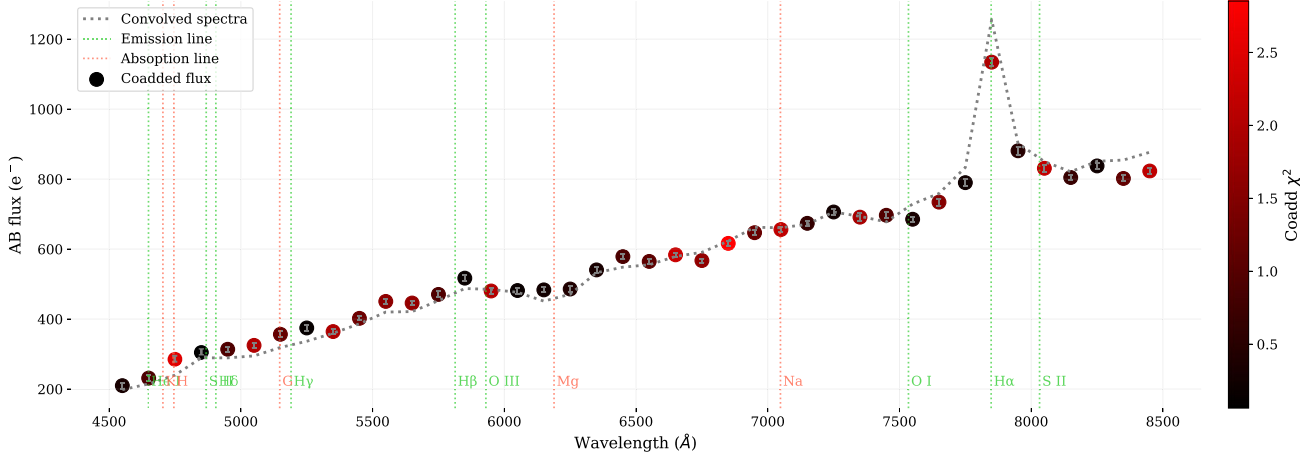
Figure 22. A bright galaxy at $z = 0.2$ from the COSMOS field. The aperture inspector displays the coadd measurements from MEMBA and overlays a synthetic narrow-band photometry from the corresponding SDSS high-resolution spectrum (when available). With the emission and absorption lines depicted, the redshift solution can be confirmed, especially on emission-line galaxies such as the example above. The redder points correspond to higher χ^2 values, caused by discrepant single-epoch fluxes, suggesting possible issues in the combined measurement.

pairs with SNR > 3. The goal is to test if the uncertainties in fluxes produced by MEMBA are consistent with repeated measurements for the same object as a function of different properties of the object and observation.

The top panels of Fig. 23 show two examples of histogram in the values of normalized flux differences:

$$df \equiv \frac{f_1 - f_2}{\sigma} = \frac{f_1 - f_2}{\sqrt{\sigma_1^2 + \sigma_2^2}} \qquad (35)$$

of duplicate measurements with fluxes f_1 and f_2. The error is just added in quadrature from the MEMBA error of f_1 and f_2. In general, the filled histograms of normalized differences follow a normal distribution (shown by dashed red lines) but the width is typically a bit larger, $\sigma_{68} \simeq 1.028$ and $\sigma_{68} \simeq 1.175$ in these two cases, than unity (black line) which is what we would expect if MEMBA errors were perfectly accurate.

The bottom panels of Fig. 23 show the width of the normalized duplicate distributions, σ_{68}, as a function of the total broad i band of the reference catalogue (W3 in CFHTLS) which we called I_{auto}. We can see that there is a strong dependence with I_{auto} which indicates that MEMBA errors are correct at the faint end but are underestimated at the bright end. Similar results are found for other PAUS fields. The machine learning background modelling prototype presented in Cabayol-Garcia et al. (2020) seems to improve this trend for brighter sources, suggesting there is a background subtraction problem in higher SNR sources.

We have seen similar tendencies in the photometry of duplicates for other BB surveys. This happens for both galaxies or stars and using aperture photometry or also SEXTRACTOR total magnitudes (e.g. in the PAUS ZP calibration with SDSS stars). We speculate that these results may come from small biases in ZP calibration values, which have spatial variations (see e.g. fig. 8 in Castander et al. 2022). These variations affect more strongly the brighter fluxes and are not accounted for by error propagation of ZPs to flux in equation (17). A ZP bias of only 1 per cent can produce up to 20 per cent mis-estimation in the flux errors of bright galaxies, which have larger fluxes and smaller relative errors.

In Fig. 24, we show $\sigma_{68}[df]$ for all other narrow-band wavelengths in PAUS. Duplicates are split in six magnitude bins: $17.5 < I_{auto} < 18.5$, $18.5 < I_{auto} < 19.5$, $19.5 < I_{auto} < 20.5$, $20.5 < I_{auto} < 21.5$,

$21.5 < I_{auto} < 22.5$, and $22.5 < I_{auto} < 23.0$ according to the galaxy total i-band magnitude. The labels show the median value of I_{auto} in each bin. The standard deviation is typically larger than $\sigma_{68} = 1$ for brighter galaxies which indicates that flux errors are underestimated in those cases. As explained above this is most likely because of ZP and PSF variations within individual CCDs, which are not accounted for in the error estimate. For scientific applications and simulations, we interpolate the values in Fig. 24 to correct for this flux-error underestimation. But note that the majority of the galaxies are faint and the most interestingly new science cases involved the faintest galaxies.

We have also looked for similar trends of $\sigma_{68}[df]$ as a function of different data reduction inputs (such as ZPs and PSF of individual CCDs, airmass, galaxy type, Sernic index, sky background, sky noise, CCD position, scatter light, exposure time, observation time, etc.) and find no significant trends. This indicates that the data reduction, the calibration, and the aperture corrections that we have applied work very well in general. We do find an error under estimation for small galaxies, with small r_{50} (half light radius), as shown in the bottom panels of Fig. 23. This is most likely caused by PSF variations within a given CCD, which are not taken into account in our data reduction.

8 PHOTO-Z PERFORMANCE

Photo-z performance of PAUS data is presented in separate papers (e.g. see Eriksen et al. 2019, 2020; Alarcon et al. 2021; Soo et al. 2021; Navarro-Gironés et al., in preparation), but it has been a key requirement for the PAUS data reduction presented here.

As an illustration of the data reduction improvements, we include some plots of photo-z performance. The accuracy of the photo-z (z_p) estimation is defined as the scatter $\sigma_{68}[dz]$ of the redshift difference $dz = (z_s - z_p)/(1 + z_s)$, as compared to a spectroscopic z_s validation sample. We also define the outlier fraction as the fraction of galaxies which has $dz > 0.02$. For more details, see Eriksen et al. 2019, 2020, Alarcon et al. 2021, and Soo et al. 2021.

Fig. 25 shows the relative improvement in the photo-z performance as a function of accumulate I_{auto} magnitude bins. We compare the ratio of photo-z performance (top accuracy and bottom outlier
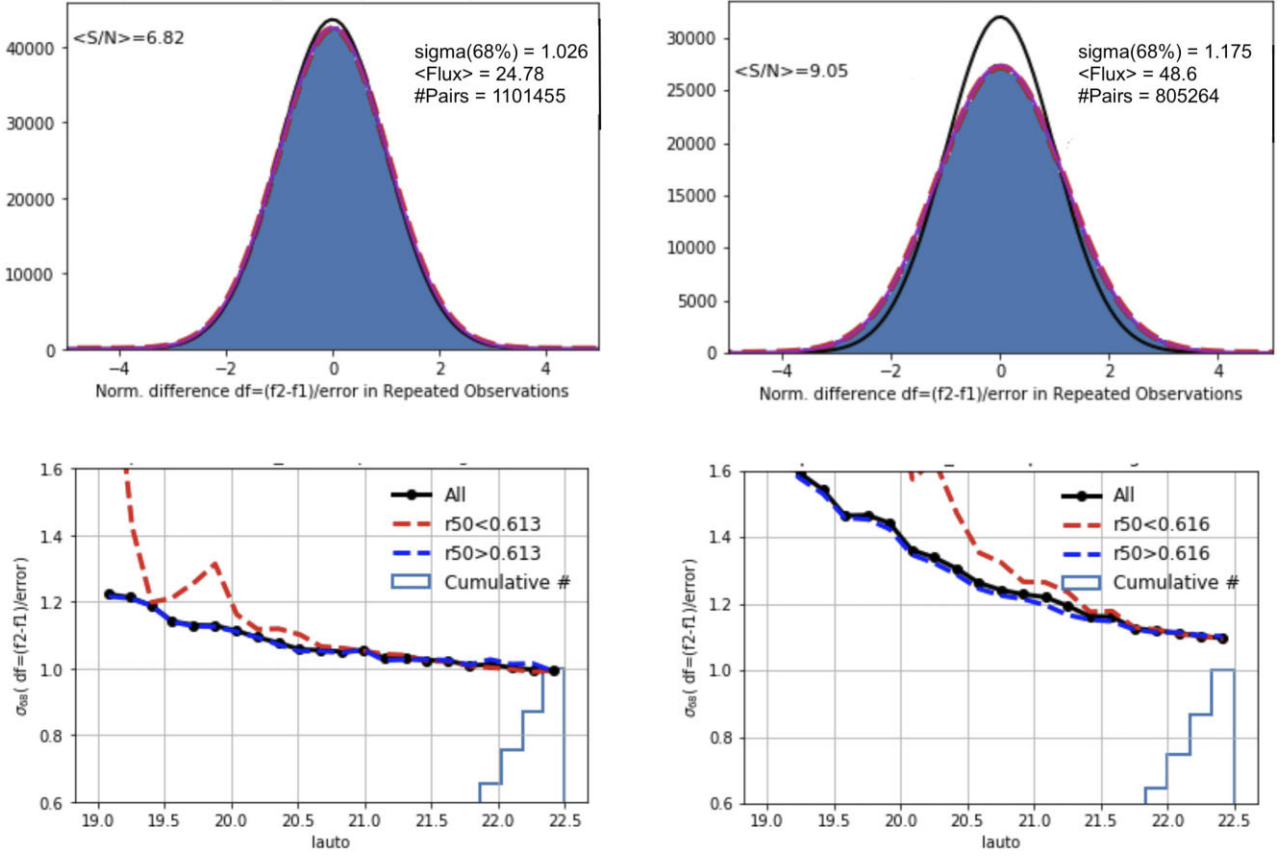
Figure 23. Top panel: Duplicates for W3 galaxies (run 941) with $NB = 575$ nm (left-hand panel) and $NB = 755$ nm (right-hand panel). Histograms of values of $df \equiv (f_2 - f_1)/\sigma$. There are $\simeq 1.1 \cdot 10^6$ pairs with mean SNR of 6.82 and mean flux of 24.78 for $NB = 575$ nm. Bottom panel: σ_{68} as a function of the I_{auto} total i-band magnitude. Continuous lines correspond to all duplicates. Dashed red line shows values with smaller half light radius $r_{50} < 0.616$ arcsec. The histogram (below the lines) shows the relative number of duplicates which increases sharply at the faint end and dominates the statistics for a full population.
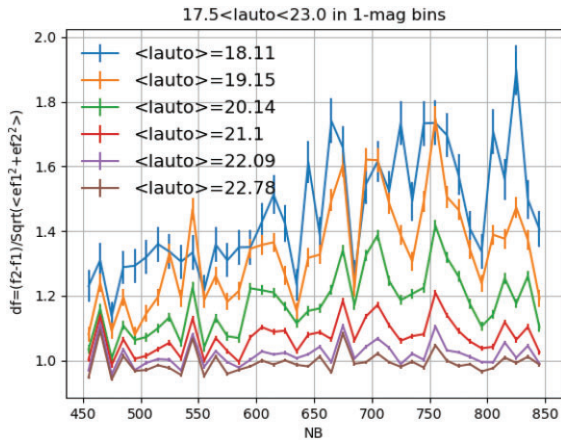
Figure 24. Normalized error (corresponding to the width σ_{68} of the histograms of Fig. 23 for the latest run 1012 of MEMBA in W3) in $df \equiv (f_2 - f_1)/\sigma$ as a function of narrow-band wavelength. Each line corresponds to the $\sigma_{68}[df]$ values for duplicates of galaxies in six separate magnitude bins within $17.5 < I_{\text{auto}}(AB) < 23.0$ in steps of $\Delta I_{\text{auto}}(AB) = 1$ (the label shows the median in the bin). Error bars correspond to 2σ estimates based on the total number of pairs in each bin.
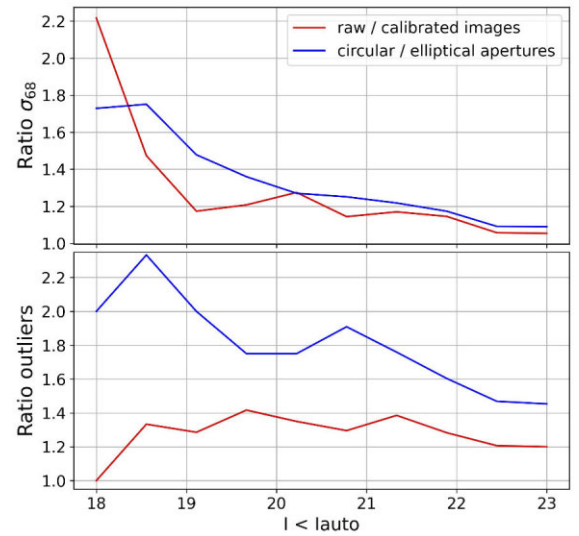
Figure 25. Top panel: Ratio of the photo-z performance, given by σ_{68} in the scatter of spectroscopic to photo-z in a COSMOS validation sample, as a function of total cumulative i-band magnitude $I < I_{\text{auto}}$. The red line compares the ratio from raw images compared to that of calibrated images. The blue line compares the ratio using circular apertures instead of elliptical ones. Bottom panel: The corresponding ratios for the outlier fraction.
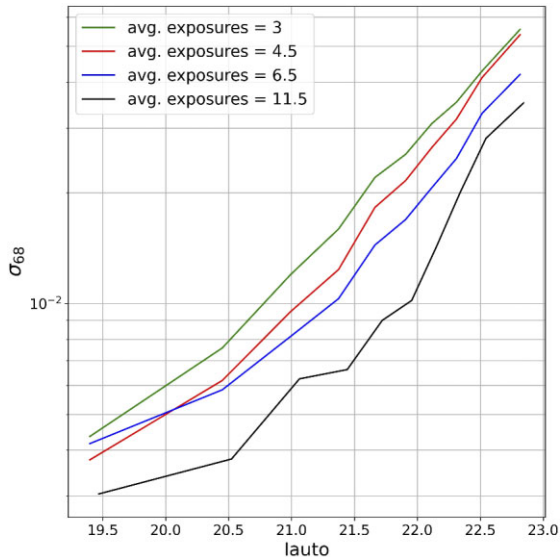
Figure 26. Photo-z accuracy ($\sigma_{68}[dz]$) in I_{auto} magnitude bins for the COSMOS field using different number of repeated exposures. The bottom line has twice the SNR in the measured fluxes than the top line. This corresponds to a shift of one magnitude in sensitivity, in good agreement with the results shown.

fraction) using raw versus calibrated images and also the ratio of using circular versus elliptical apertures. In both cases, we see a significant improvement of the photo-z performance, which validates these steps of the data reduction.

Fig. 26 shows that the accuracy of $\sigma_{68}[dz]$ increases with the number of exposures that we use in the MEMBA coadds. The difference between the top and bottom lines is about a factor 4 in the number of exposures. This translates into an improvement of a factor of 2 in SNR of the narrow-band flux measurements, which is equivalent to a shift of one magnitude (galaxies that are one magnitude brighter have about 2 times larger SNR). The change in $\sigma_{68}[dz]$ shown in Fig. 26 is caused by the degradation in the SNR of the narrow-band measurements as we consider fainter galaxies. The difference in $\sigma_{68}[dz]$ between the top and bottom line roughly correspond to a shift of about one magnitude in the curves, as expected from the corresponding increase in the SNR. This agreement validates our method for adding different exposures, as explained in Section 6.5. It also allows us to predict what type of photo-z performance we can expect for surveys with different exposures (or collecting power) using our pipeline.

9 KNOWN LIMITATIONS AND FURTHER WORK

Since the first light of PAUCam we have been constantly improving both NIGHTLY and MEMBA algorithms, continuously improving our understanding of the systematics and the instrument behaviour. Even though we reached outstanding photometry (see Section 7.3) and photo-z accuracy (see Section 8), there are still steps in the overall process that can be improved. However, in this paper, we refer only to those algorithms that were implemented in the main pipeline and that delivered scientific results in already published papers. We plan to improve and present further algorithms with the corresponding upgrade in scientific results once these are stabilized and validated.

Regarding the flux calibration step in the NIGHTLY pipeline, we know the current dome flat-fields do not reproduce the sky

illumination accurately. Therefore, the detector response after the flat-field calibration is not homogeneous. A possible solution is processing sky flats and leaving the dome flats for small-scale pixel variations and not to correct the large-scale vignetting and illumination patterns. This has the difficulty of the scatter-light as an additive component which complicates the sky-flat processing. The solution proposed in Section 3.4.2 has only been applied to particular studies involving extended objects such as the M101 of Fig. 6. Once we verify the fluxes of the main target galaxies (smaller and fainter down to $i_{AB} < 23$) are preserved, we will implement the sky-flat scatter-light subtraction to the main processing.

The background modelling implemented in MEMBA is a simple but reliable method. Due to the complex varying background in PAUS images resulting from the flat-field and scatter-light residuals, a more complex background estimator that understands trends of the background can provide significant benefits. This is especially important for low-SN sources where small residuals of the background can bias the measurement. A machine learning technique has been studied and published in Cabayol-Garcia et al. (2020) and will soon enter into the main processing of PAUdm. There is also an innovative method developed for PAUS (Cabayol et al. 2021) that provides flux estimates from neural network algorithms, improving the accuracy and increasing the signal-to-noise of the measurements.

The areas in the focal plane with most distortion are currently flagged. There is a possibility to increase the survey efficiency by accurately modelling the PSF at the edges of the focal plane and adapt the apertures to include a larger area of the unvignetted mosaic.

Modelling the PSF and computing the astrometric solution is done at the single-epoch level, independently for each exposure, for simplicity and because it delivers good enough precision. However, more stable and accurate solutions can be obtained by computing astrometry and PSF models with multiple overlapping exposures, at the expense of complicating the NIGHTLY processing adding dependency between single-epoch reductions. The astrometric precision would only improve marginally at the subpixel level, which could have some benefits for particular applications. But a more accurate and stable varying PSF modelling could improve the photometry and increase the area efficiency, contrary to the simple model we currently use.

Similar to the previous point, the photometric calibration delivers ZPs independently for each detector image. There are algorithms such as the Übercalibration that compute a global solution from the overlapping measurements between images, homogenizing the calibration and ensuring a flat response across wide area in the sky. However, as we calibrate against SDSS stars and those have been globally calibrated in a similar process, we would expect this improvement to be minor for PAU, assuming a good match with SDSS photometry.

10 SUMMARY AND CONCLUSIONS

The PAUdm system described here has been able to provide the most accurate photo-z catalogues available down to $i_{AB} < 23$, dealing with very particular aspects of the narrow-band photometry and specifics of the instrument. Custom algorithms were designed to calibrate narrow bands down to 1-per cent accuracy. Subtle systematic effects had to be modelled and corrected such as the specialized processing to deal with scatter-light residuals caused by the unusual filter tray disposition of PAUCam. The technical implementation presented has also been challenging due to the large volume and complex data set for this survey. It has been key to orchestrate the processing and metadata around a powerful data base with flexibility to modify

and extend the processing as needed and allowing very complex analysis that enforced the scientific exploitation of the data. Although some of the adopted solutions are bound to the infrastructure of the data centre, it can be adapted to different surveys or hardware configurations with similar volumes of data (below 100 TB and 10^{10} data base entries).

The current PAUdm implementation has some limitations that we are currently working to improve. Even under these limitations, the photometric catalogues published by PAUS deliver the most precise photo-z down to $i_{AB} < 23$. PAUS data were available in the Early Data Release (Eriksen et al. 2019), in the PAUS+COSMOS photo-z catalogue (Alarcon et al. 2021) and in Soo et al. (2021). Fig. 21 shows that our pipeline reduction delivers an inter-band photometric calibration of 0.8 per cent across the 40 narrow-band set. Fig. 26 shows how the stacking (or co-addition) of independent PAU measurements results in a photo-z accuracy that scales as expected with the number of exposures. This provides a direct validation of the whole pipeline presented in this paper. It also indicates that the pipeline is optimal as it saturates the performance that can be achieved for a given signal-to-noise in the input data. These measurements open new windows in various astronomy scientific areas as published in Stothert et al. (2018), Tortorelli et al. (2021), Johnston et al. (2021), Renard et al. (2021, 2022), van den Busch et al. (2022), or Cabayol et al. (2023) with its large-scale imaging survey of narrow-band photometry where each pixel is a low-resolution spectrum.

ACKNOWLEDGEMENTS

DATA AVAILABILITY

The data underlying this article are available in the PAUS website, under the Data Releases section, at https://pausurvey.org. The available catalogues include the Early Data Release (EDR) catalogue and the PAUS+COSMOS photo-z catalogue. The EDR corresponds to PAUS data obtained in the COSMOS field, described and used in Eriksen et al. (2019). The PAUS+COSMOS photo-z catalogue contains accurate and precise photometric redshifts in the ACS footprint from the COSMOS field for objects with iAB < 23 combining all 40 PAUS bands with 26 BBs from the COSMOS2015 catalogue (Alarcon et al. 2021). Further data releases are expected with the photometry presented in this article.

REFERENCES

Alam S. et al., 2015, ApJS, 219, 12
Alarcon A. et al., 2021, MNRAS, 501, 6103
Astropy Collaboration et al., 2018, AJ, 156, 123
Bertin E., 2006, in Gabriel C., Arviset C., Ponz D., Enrique S., eds, ASP Conf. Ser. Vol. 351, Astronomical Data Analysis Software and Systems XV. Astron. Soc. Pac., San Francisco, p. 112
Bertin E., 2011, in Evans I. N., Accomazzi A., Mink D. J., Rots A. H., eds, ASP Conf. Ser. Vol. 442, Astronomical Data Analysis Software and Systems XX. Astron. Soc. Pac., San Francisco, p. 435
Bertin E., Arnouts S., 1996, A&AS, 117, 393
Bertin E., Mellier Y., Radovich M., Missonnier G., Didelon P., Morin B., 2002, in Bohlender D. A., Durand D., Handley T. H., eds, ASP Conf. Ser. Vol. 281, Astronomical Data Analysis Software and Systems XI. Astron. Soc. Pac., San Francisco, p. 228
Bradley L. et al., 2020, astropy/photutils: 1.0.0. Zenodo. https://doi.org/10.5281/zenodo.4044744
Bridle S. et al., 2009, Ann. Appl. Stat., 3, 6
Cabayol-Garcia L. et al., 2020, MNRAS, 491, 5392
Cabayol L. et al., 2021, MNRAS, 506, 4048
Cabayol L. et al., 2023, A&A, 671, A153
Calabretta M. R., Greisen E. W., 2002, A&A, 395, 1077
Capak P. et al., 2007, ApJS, 172, 99
Casas R. et al., 2012, in Holland A. D., Beletic J. W., eds, Proc. SPIE Conf. Ser. Vol. 8453, High Energy, Optical, and Infrared Detectors for Astronomy V. SPIE, Bellingham, p. 845326
Castander F. J., Serrano S., Eriksen M., Gaztanaga E., Casas R., 2022, AJ, 157, 246
Erben T. et al., 2013, MNRAS, 433, 2545
Eriksen M. et al., 2019, MNRAS, 484, 4200
Eriksen M. et al., 2020, MNRAS, 497, 4565
Freyhammer L. M., Andersen M. I., Arentoft T., Sterken C., Nørregaard P., 2001, Exp. Astron., 12, 147
Gaia Collaboration, 2018, A&A, 616, A1
Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, ApJ, 622, 759
Heymans C. et al., 2012, MNRAS, 427, 146
Hildebrandt H. et al., 2012, MNRAS, 421, 2355
Ivezić Ž. et al., 2019, ApJ, 873, 111
Janesick J. R., 2001, Scientific Charge-Coupled Devices. SPIE, Bellingham
Johnston H. et al., 2021, A&A, 646, A147
Kannawadi A. et al., 2019, A&A, 624, A92
Kuijken K. et al., 2015, MNRAS, 454, 3500
Kuijken K. et al., 2019, A&A, 625, A2
Laigle C. et al., 2016, ApJS, 224, 24
Landolt A. U., 1992, AJ, 104, 340
Le Fèvre O. et al., 2013, A&A, 559, A14
Miller L. et al., 2013, MNRAS, 429, 2858
Newman J. A. et al., 2013, ApJS, 208, 5
Padilla C. et al., 2019, AJ, 157, 246
Padmanabhan N. et al., 2008, ApJ, 674, 1217
Renard P. et al., 2021, MNRAS, 501, 3883
Renard P. et al., 2022, MNRAS, 515, 146
Rowe B. T. P. et al., 2015, Astron. Comput., 10, 121
Soo J. Y. H. et al., 2021, MNRAS, 503, 4118
Stothert L. et al., 2018, MNRAS, 481, 4221

Tody D., 1986, in Crawford D. L., ed., SPIE Conf. Ser. Vol. 627, Instrumentation in Astronomy VI. SPIE, Bellingham, p. 733

Tonello N. et al., 2019, Astron. Comput., 27, 171

Tortorelli L. et al., 2021, J. Cosmol. Astropart. Phys., 2021, 013

Trujillo I., Aguerri J. A. L., Cepa J., Gutiérrez C. M., 2001, MNRAS, 328, 977

van Dokkum P. G., 2001, PASP, 113, 1420

van den Busch J. L. et al., 2022, A&A, 664, A170

Wells D. C., Greisen E. W., Harten R. H., 1981, A&AS, 44, 363

York D. G. et al., 2000, AJ, 120, 1579

APPENDIX A: OPERATION AND TECHNICAL PERFORMANCE

The PAUdm system has been designed to operate in the infrastructure at the Port d'Informaciò Científica (PIC). In this section, we describe the technical aspects of the main scientific pipelines and the tools required to operate the pipeline under the available infrastructure. The solution presented below was not the original design as both the infrastructure and the project required changes since the beginning of the operation. A more technical and infrastructure-oriented description of PAUdm is described in Tonello et al. (2019). This section also includes updated data and pipeline flows from the ones presented in the technical paper.

A1 Archive

The PAUS camera produces ∼300 GB of raw data per observing night. These are mostly FITS files that contain the exposure images and additional metadata in its header. The data are processed in the main pipelines where more subproducts are generated, multiplying the volumes of raw data. The actual size of the raw archive (until 20A observations) is 42 TB and it is safely archived as a two-copy tape and a third copy on disc. The processed data are significantly larger due to the increased bit depth and the various subproducts generated per exposure (science, mask, weights, PSF models, etc.). For the reduced data, we have a single copy on disc, except for published releases where we include an additional copy on tape. The raw archive tree organizes the data in observation sets, following the schema in the PAUCam temporary archive at the observatory. The reduced tree

adds an additional layer to account for multiple reprocessing of the same data, that we call productions.

Even though most of the access to the archive system is provided by the nodes in the computing farm, we wanted to make available both raw and reduced data to the PAUS collaboration. For that purpose, we set up a WebDAV server that allows web access to the entire archive in a user-friendly format.

A2 Data base

With such a large and complex data set where millions of galaxies, measurements, and images are related, it has been key to set up a relational data base that tracks all the information and metadata of the survey. We have chosen a PostgreSQL data base running on a powerful 12-core 96 GB server in a twin configuration for reliability and performance. The data base can be accessed by the pipelines via an object relational mapper (ORM) for better integration and reliability under the pipeline environment. It is also accessed by the different web applications such as the NIGHTLY report and the forced aperture inspector described in previous sections. As PAUS is a large collaboration, we make available via the PAUdm website a dynamic view of the data base to browse the schema and perform simple queries. Additionally, for development and validation purposes, it can be accessed via Python notebooks under a read-only role, dumping queries directly into data frames with all the flexibility and potential that these objects provide.

The data base model was designed to allow for reprocessing of data at any level, tracked by the `production` table. Careful constraints were set on each table to ensure unique entries under each production set. We defined four main pipelines: (i) The pixel simulation is where we produce survey and pixel image simulation for development and assessment of performance. (ii) The NIGHTLY pipeline (Section A4.1) is where the main image reduction and calibration happens. The NIGHTLY can process input productions from the pixel simulation or real observations from PAUCam. (iii) The MEMBA pipeline (Section A4.2) is where we perform forced photometry over the NIGHTLY images. (iv) And finally the photo-z pipeline, a wrapper to BCNz (Eriksen et al. 2019) where we estimate photo-z from MEMBA photometry. Each pipeline can be processed independently, allowing processing multiple times a given input production. For instance, one
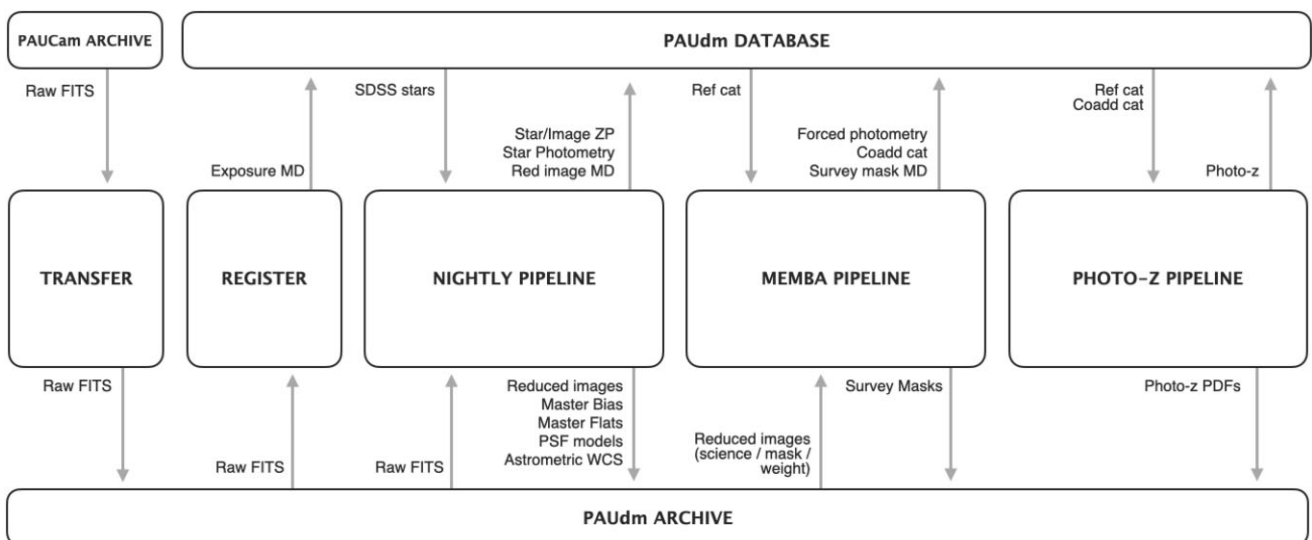
Figure A1. A simplified schema of the PAUS data flow, where interactions among pipeline, storage, and data base are shown.
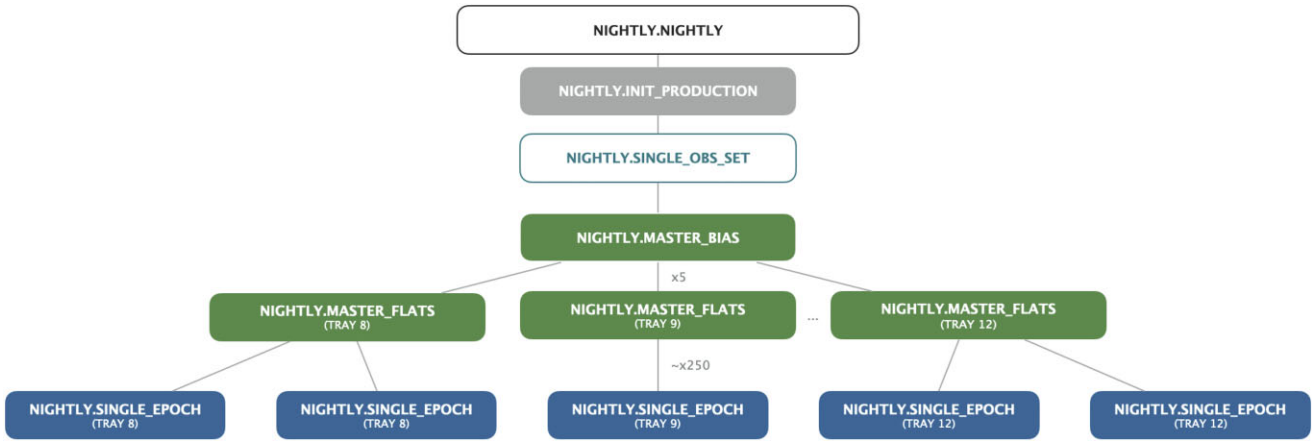
Figure A2. The dependency chart of the NIGHTLY pipeline. Empty boxes define jobs creating subjobs while filled boxes refer to jobs processing data at the computing farm. The dependency at the filter level starts with the master flat-fields and the single epochs are related with its own masters.
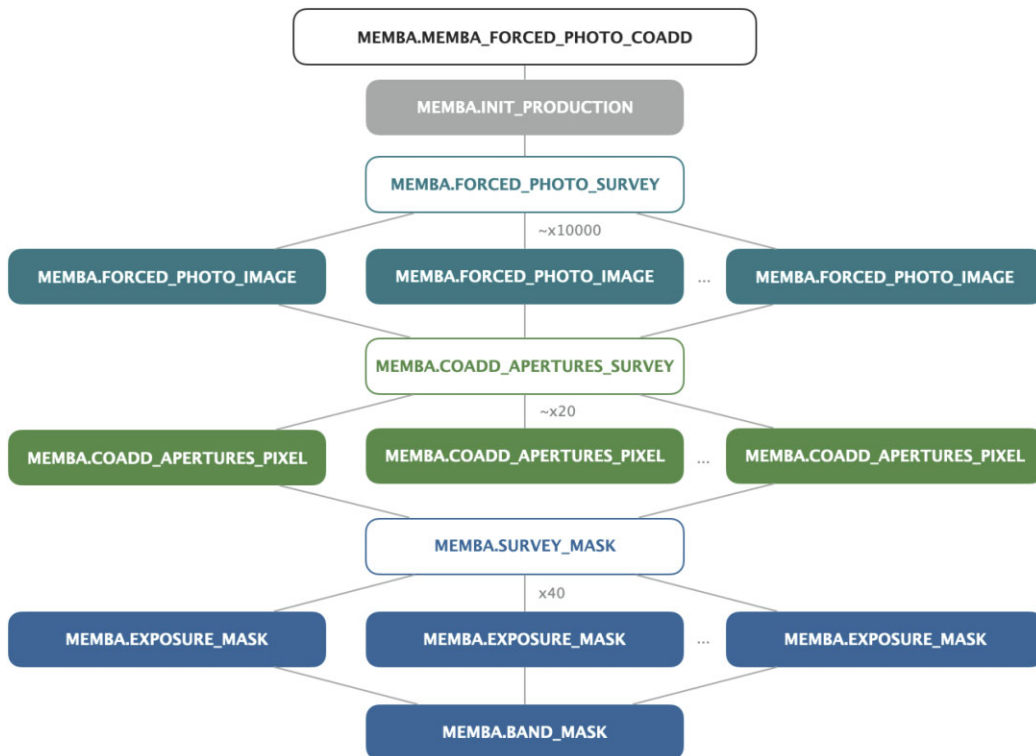
Figure A3. The dependency chart of the MEMBA pipeline. Empty boxes define jobs creating subjobs while filled boxes refer to jobs processing data at the computing farm. The parallelization takes place at three levels: forced photometry, coadds, and survey masks.

can process different aperture photometry in MEMBA under various configurations with the same set of image reduction made by the NIGHTLY pipeline.

In addition to the data-related tables, we have a survey strategy data base, synchronized with the one at the observatory, containing all information about fields, exposure status, and observation progress. These tables are used throughout the night to schedule targets that need to be observed by PAUCam. It is also necessary in the regular PAUdm processing to select the exposures associated to a survey field and that have been classified as valid exposures.

Finally, the data base have all data operation tables with information about job configuration, status, and their dependencies. The use

of these is described in the next subsection (Section A3). Additionally the whole data base schema with the most relevant columns in each table can be found in Appendix B.

A3 Processing

The high volume of data and its complex analysis requires processing the pipelines in a data centre with enough computing power such as at the PIC. This is a high-throughput computing data centre and therefore we had to split each pipeline into smaller tasks than can be processed independently with limited consumption of memory and CPU time. Consequently, a pipeline may result in hundreds or even

thousands of jobs, with its configuration and inter-dependencies that need to be launched, monitored, and operated.

For that purpose we designed a job orchestration tool named Brownthrower (BT) that gives us the flexibility to operate the pipelines at PIC. With BT we can create jobs, add dependencies between them so certain jobs do not begin processing until others are complete, share configuration between jobs, and monitor the status of a pipeline and its subjobs. To process these jobs, we submit to the computing farm via HTCondor a set of pilot jobs that are continuously grabbing free jobs to be processed (in status queued and without pending dependencies) until all jobs are being processed.

A JUPYTER LAB web service was recently set at PIC, running over actual nodes from the computing farm (with memory up to 32 GB) or even GPU nodes. This service has been of great help to develop and deploy new algorithms, as well as for validation and quick test purposes.

A4 Data flow and orchestration

The data flow in PAUdm is divided between the main PAUdm archive (where large files are stored) and the data base (where metadata and information that requires complex selections are uploaded), orchestrated by the different pipelines. A summary of the data flow is depicted in Fig. A1. In the first place, the raw data are transferred from the observatory on La Palma to the archive set in Barcelona at PIC (detailed in Tonello et al. 2019). Immediately, the exposure metadata are registered in the data base. Next, the NIGHTLY pipeline begins its image calibration and archives the clean images, its PSF models, and the astrometric solutions in WCS. Photometric measurements and their calibration ZPs are uploaded to the data base. Once enough sky area has been processed by the NIGHTLY, MEMBA can begin the galaxy photometry and reingest its coadd catalogue once all measurements have been done. MEMBA is also in charge of producing survey masks and thus, it stores them in the archive. Finally, the photo-z pipeline obtains MEMBA's measurements and computes the photo-z for each galaxy. The photo-z values and estimated errors enter the data base while the large redshift probability distribution files are archived in the storage.

A4.1 Nightly pipeline

This is the pipeline in charge of the image calibration, as described in Section 3. It begins with a set of raw exposures, including flat and bias calibration images. It ends up with the science exposures astrometrically and photometrically calibrated, ready to perform the flux measurements.

The NIGHTLY pipeline has two main steps. First, we have the production of master bias and master flats and secondly the single-epoch reduction of science images. During observation periods, we process the pipeline in batches of observation sets. Typically an observation set contains the exposures from a single night. However, as PAUCam is also available in the *WHT* as a community instrument, it allows having more than one observation set per night in cases where observations belong to different surveys. When we operate in observing mode, the NIGHTLY pipeline tree starts with the master bias, then with the master flats of each filter tray, and associated to each tray the corresponding sky images. This means that the image calibration would not start until the jobs of master bias and master

flat are successfully completed. The pipeline tree can be seen in Fig. A2.

There is a second mode to operate the NIGHTLY pipeline meant to process entire fields from multiple nights. This is used when we improve the NIGHTLY code and want to reprocess a given subset. First we process all calibration frames (master bias and master flats) from the nights with science exposures that we plan to process. At a second stage, once all the calibration frames are available, we analyse all the science images in parallel.

The pipeline has evolved significantly since the beginning as we had to deal with heterogeneous data sets such as very cloudy skies, observation sets without calibration images, saturated flats, etc. A major effort had to be done to automatically detect any possible situation (detailed in Section 7.1) and either correct for it or classify the faulty data as invalid.

To allow for the 8-h rapid feedback during an observation period, we launch 50 BT pilots to the computing farm that process all images in time. Reprocessing entire fields involves a much greater set of images and therefore we increase the number of pilots to 100. Master bias and master flat jobs are processed in 10 min while single-epoch exposure reduction can take up to 20 min per job.

The NIGHTLY pilots require intense I/O access to the archive system to retrieve raw data and ingest reduced images and for this reason we do not allow more than 100 NIGHTLY pilots to run in parallel.

A4.2 MEMBA pipeline

The MEMBA pipeline is intended to perform the photometry over the reduced images across the different bands and overlapping exposures from different epochs. It can be divided into three main steps: the forced photometry, the coadd catalogue, and the production of survey masks. The pipeline tree can be seen in Fig. A3.

Each forced photometry job takes care of running the photometry of a single detector image. It will load the reference catalogue, the corresponding reduced image, and mask and will upload the measurements to the data base once completed.

The coadd jobs are divided in different areas in the sky. We select the areas by HEALPIX pixels (Górski et al. 2005) of *N* side 128. This approach allows us to limit the load of each job, retrieving only the overlapping measurements of a reduced area, defined by the pixel size. Co-adding a larger field would only increase the number of jobs but those will always be constrained in memory and processing time.

The dependency of jobs is set between the forced photometry tasks and the coadds, as the latter requires the photometry to be complete at all bands and layers prior to the combination. On the contrary, mask jobs can be processed independently as these do not require inputs from the forced photometry or coadd tasks.

As MEMBA jobs do not have such intense access to the archive and interact mostly with the data base, we can increase the number of parallel pilots up to 200. The CPU time of a MEMBA run is dominated by the forced photometry jobs. The largest fields are made of $>30\,000$ images. Each job lasts approximately 15 min, resulting in a total process time of 1 d for every 10 deg^2.

APPENDIX B: DATA BASE SCHEMA

In this appendix, we describe the PAUS data base schema. In Table B1, we group different tables in each pipeline, which are given by different colours in Fig. B1.

Table B1. Data base schema.

| Pipeline	Data base	Description
Survey strategy		
	ss_target	Survey strategy targets from observations
NIGHTLY		
	detection	Detections on the image after the NIGHTLY data reduction
	image	List of images associated to the mosaics (CCD and single amplifier images)
	mosaic	List of mosaic exposure images (raw and reduced)
	obs_set	List of observation sets registered in the data base
	obs_set_project	Projects associated to observation sets
	project	List of projects associated to PAUCam observations
	quality_control	Quality control entries
Cross-talk		
	cross-talk_diff	Cross-talk differences between raw images
	cross-talk_ratio	Cross-talk ratios between amplifiers
Photometric Calibration		
	image_zp	Image ZP measurements for each photometry-calibration method
	phot_method	Method used during photometry for image calibration
	phot_zp	Photometric ZPs per CCD
	star_photometry	Photometry for each star in the NIGHTLY photometry
	star_zp	ZP measurements for each star in the NIGHTLY photometry
MEMBA		
	forced_aperture	Single-epoch forced aperture photometry
	forced_aperture_coadd	Coadd forced aperture fluxes per band
	forced_aperture_report	Reports for forced aperture inspector
	mask_image	List of mask images (band and field)
	memba_ref_cat	Reference catalogue used in each MEMBA production
Photo-z		
	photoz_bcnz	Photo-z from BCNz code
External		
	cfhtlens	External CFHTLenS catalogue for forced photometry
	cosmos	zCOSMOS (DR3). Redshifts for forced photometry and validation
	deep2	The DEEP2 DR4 redshift catalogue
	gaia_dr2	*Gaia* DR2 stellar catalogue
	kids	KiDS KV450-G9 reference catalogue
	sdss_spec	SDSS Spectra catalogue
	sdss_spec_photo	SDSS DR12 (Spec_Photo view). Spectrum for forced photometry and validation
	sdss_star	SDSS DR12 (Star view). Stars for simulation and calibration
Synthetic photometry		
	match_to_spec	Match table between forced aperture catalogues and spectra catalogues
	spec_conv	Convolved fluxes derived from spectra in external surveys (SDSS, COSMOS, and DEEP2)
	synth_sdss	Synthetic photometry over SDSS spectra (over COSMOS and W1)
	synth_vipers	Synthetic photometry over VIPERS spectra (over W1)
BT (operation tables)		
	dependency	Tracks the dependency between BT jobs
	job	Tracks the list of BT computing jobs (operation table)
	tag	Contains tags for BT jobs (operation table)
Production		
	production	Tracks the different processing production runs for all pipelines

Figure B1. PAUS data base schema. Inside each table with show some example columns. Tables are grouped under the same colour.

## APPENDIX C: SYNTHETIC SPECTRA AGAINST PAUS PHOTOMETRY EXAMPLES

In this section, we include some interesting reference synthetic spectra samples (Figs C1–C4) against the PAUS narrow-band measurements after all the processing described in this study. This highlights only a star, a galaxy, and a Quasi-Stellar Object (QSO) but it illustrate the possibilities of PAUS and the validation with this synthetic reference method.
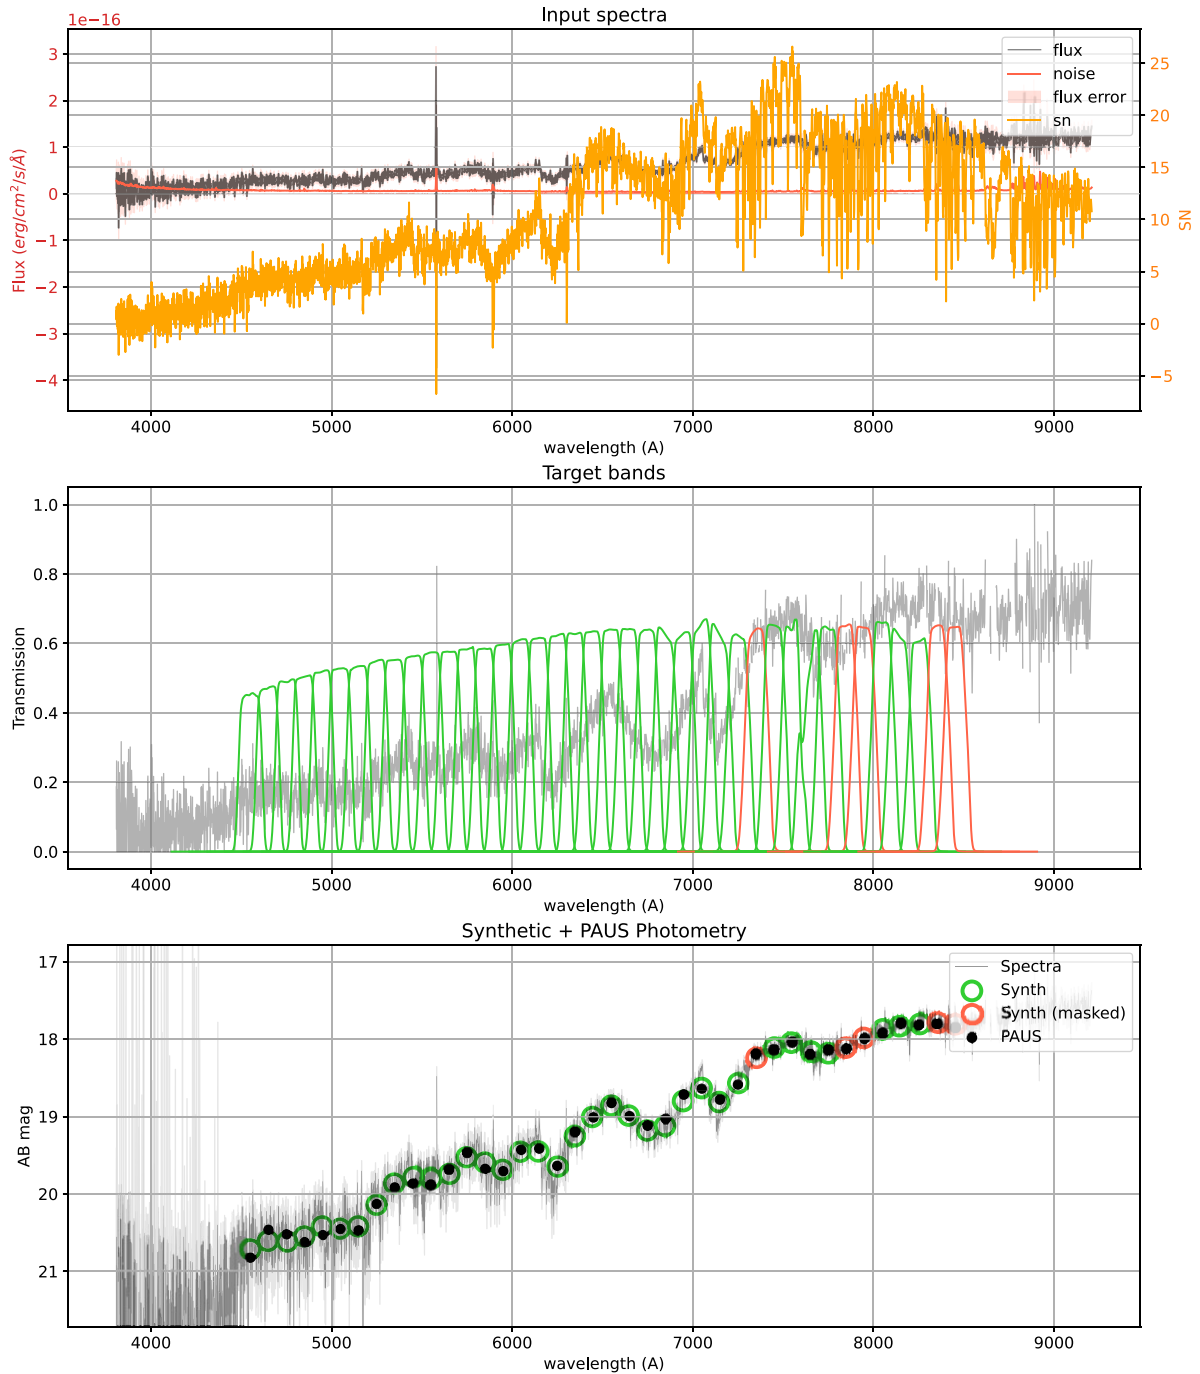
Figure C1. A M3 star observed by PAUS (COSMOS object 79081) with reference SDSS spectrum (Plate 501 - MJD 52235 - FiberID 461 - SNR 8.95).
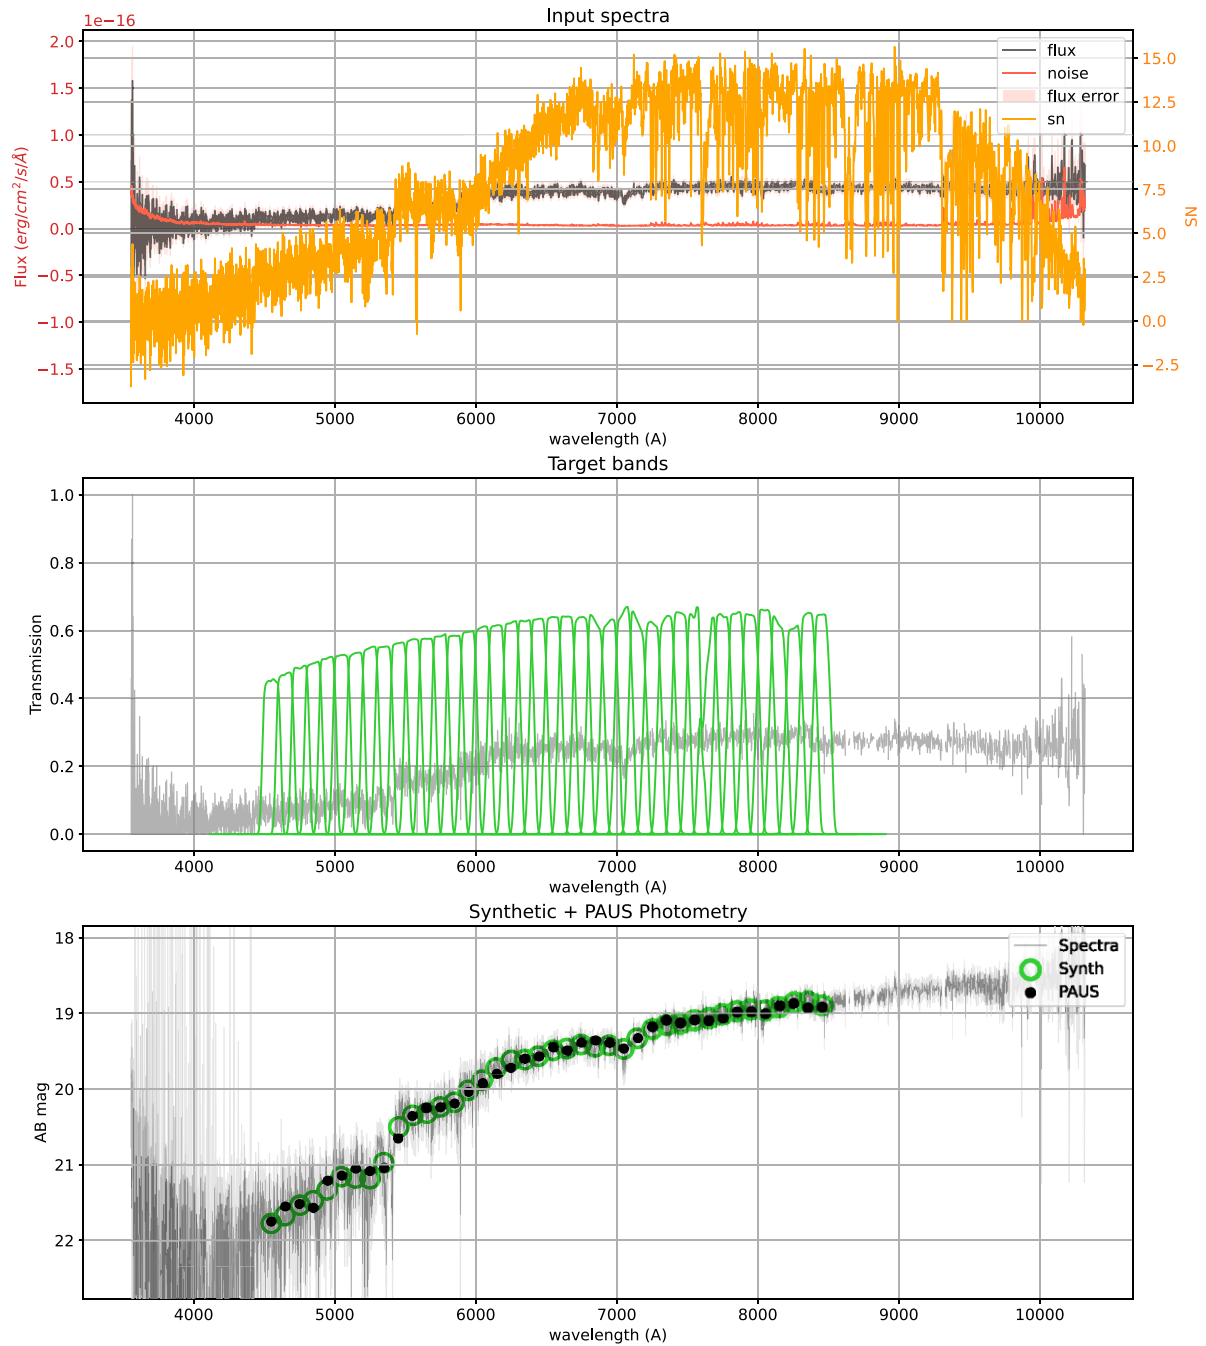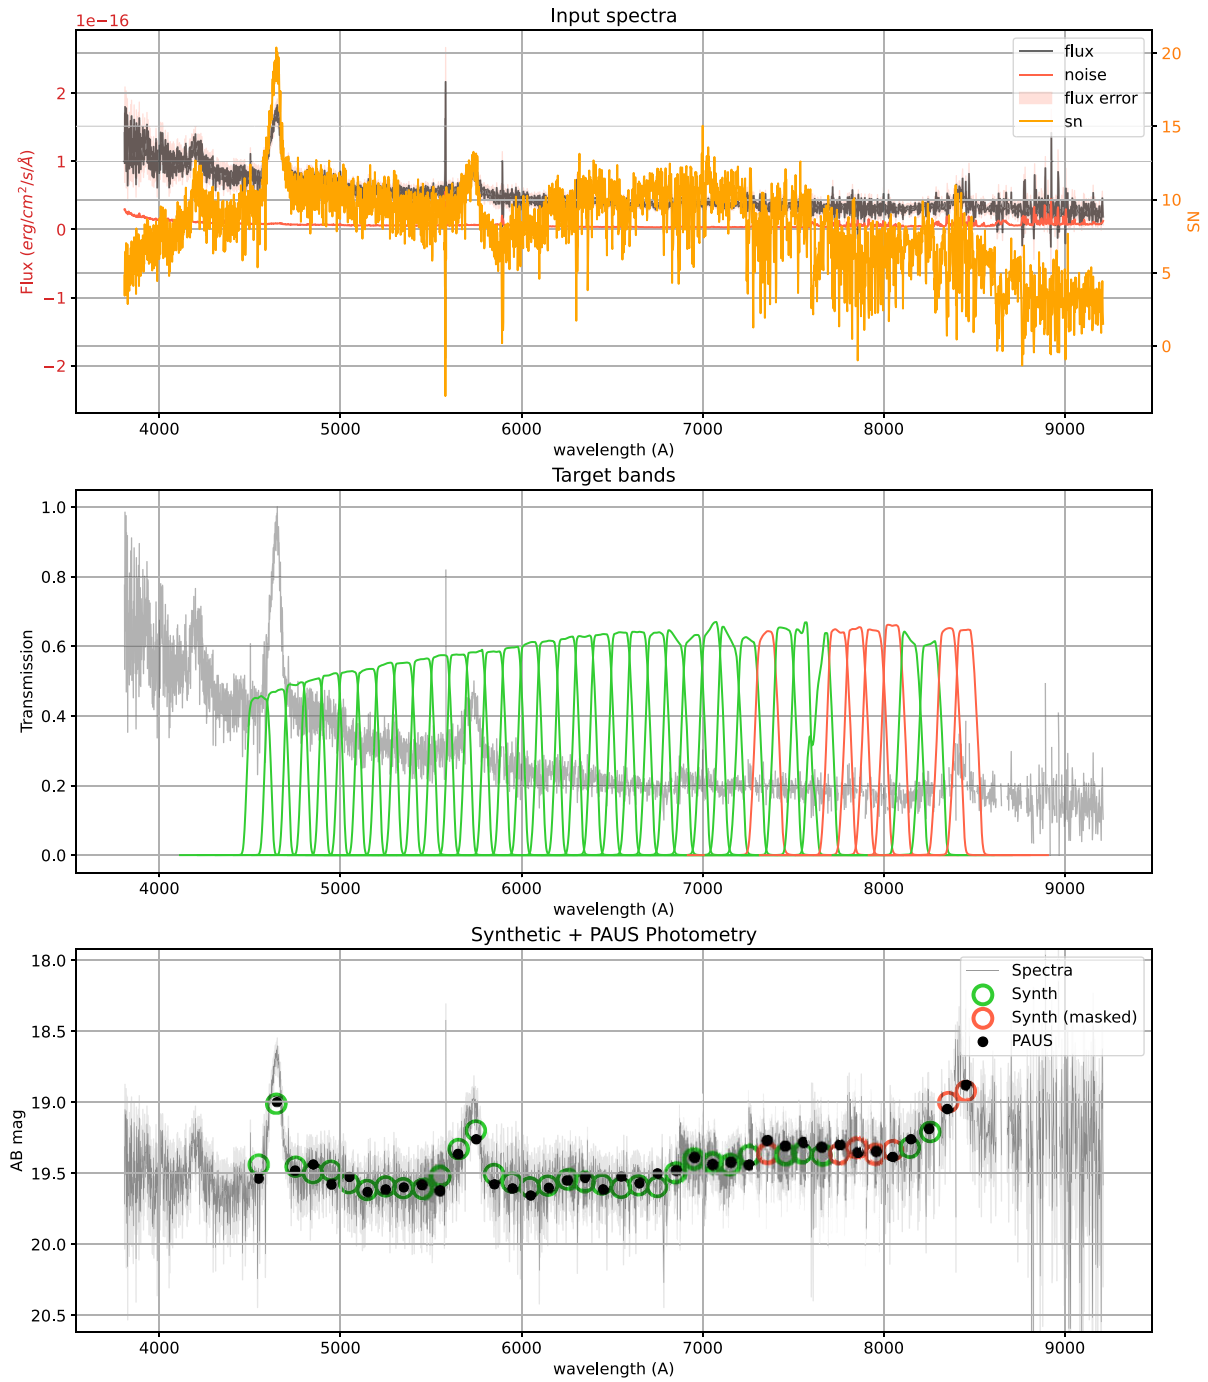
Figure C2. A red galaxy observed by PAUS (COSMOS object 3956) with reference SDSS spectrum (Plate 4737 - MJD 55630 - FiberID 58 - SNR 8.37) at redshift 0.362.
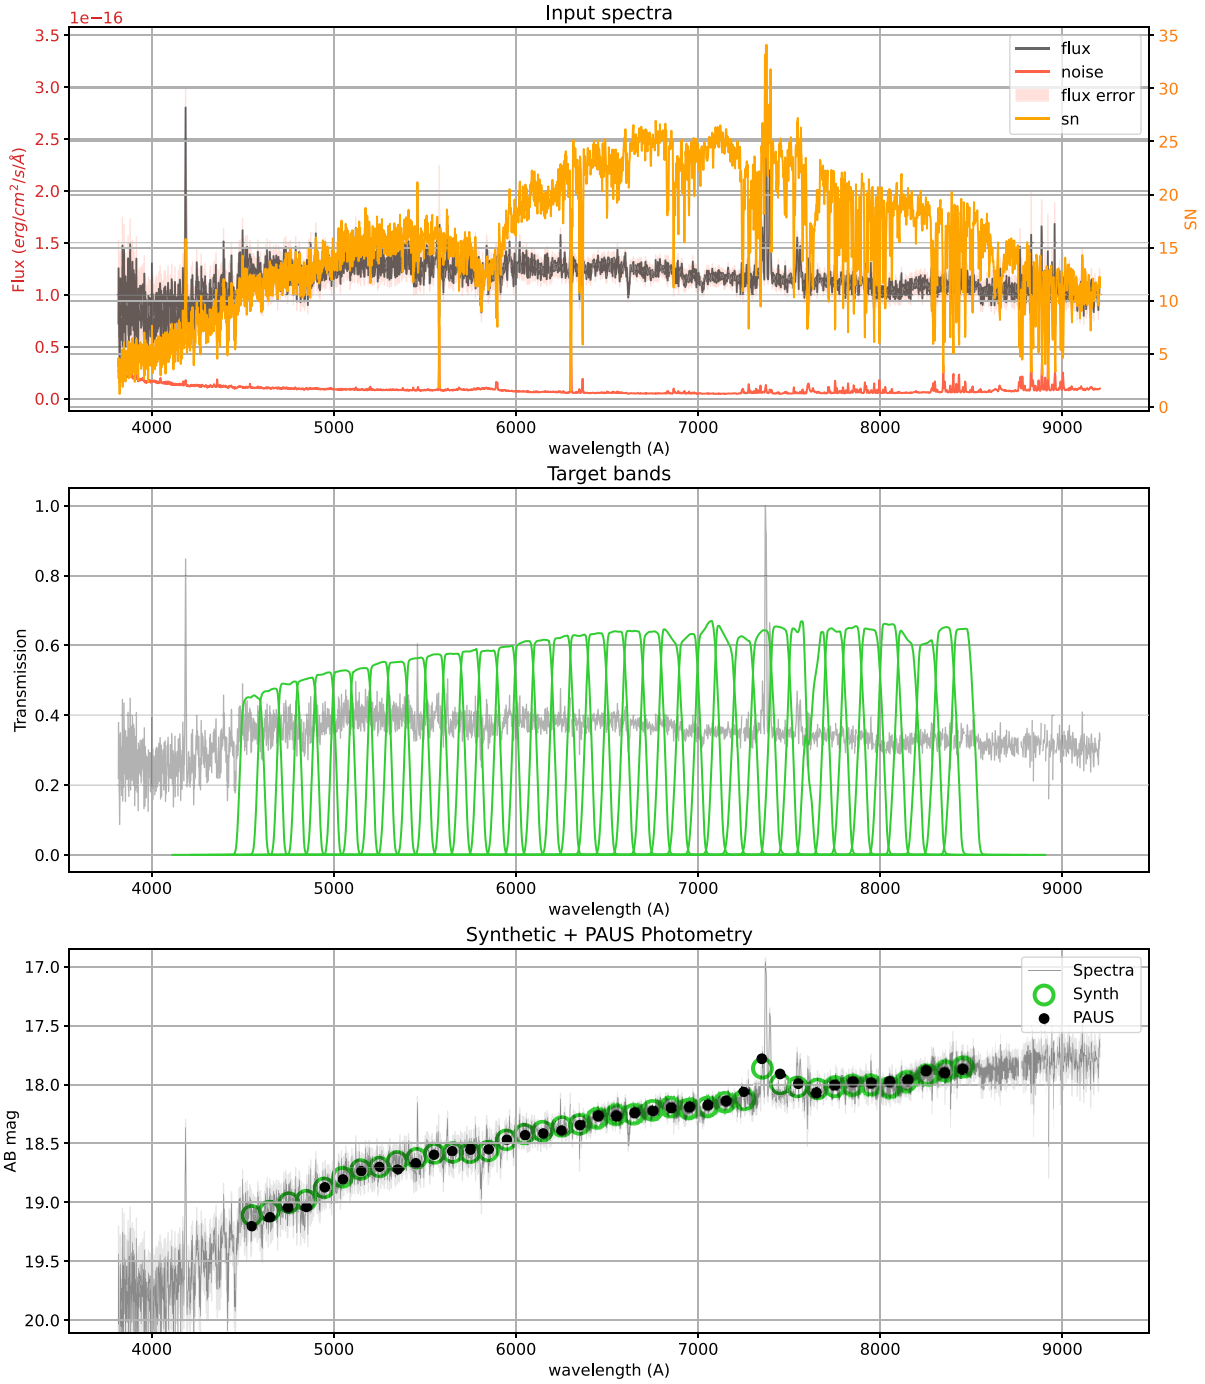
Figure C3. An Hα star-forming galaxy observed by PAUS (COSMOS object 67024) with reference SDSS spectrum (Plate 500 - MJD 51994 - FiberID 587 - SNR 19.41) at redshift 0.122.

Figure C4. A QSO observed by PAUS (COSMOS object 80935) with reference SDSS spectrum (Plate 501 - MJD 52235 - FiberID 462 - SNR 9.31) at redshift 2.00.

APPENDIX D: FLAGGING

In this section, we describe in Table D1 the list of flags that any source can have throughout the whole data processing of the PAUdm system. Each flag is assigned to a bit such that with a single integer we can obtain the unique list of flags affecting each source.

Table D1. The list of flags used across the PAUdm system. The table specifies the flag reason, the assigned value and bit, the software/pipeline origin, and the level at which the flag is defined.

Flag	Value (bit)	Origin	Level
Crowded	1 (1)	SExtractor	Source
Merged	2 (2)	SExtractor	Source
Halo	4 (3)	SExtractor	Source
Truncated	8 (4)	SExtractor	Source
Deblended	16 (5)	SExtractor	Source
Cross-talk	32 (6)	Nightly mask	Pixel
Scatter-light	64 (7)	Nightly mask	Pixel
Extinction	128 (8)	Nightly photometry	Image
ZP	256 (9)	Nightly photometry	Image
Cosmetics	512 (10)	Nightly mask	Pixel
Saturated	1024 (11)	Nightly mask	Pixel
Cosmics	2048 (12)	Nightly mask	Pixel
Vignetted	4096 (13)	Nightly mask	Pixel
Discordant	8192 (14)	MEMBA photometry	Source
Edge	16 384 (15)	MEMBA photometry	Source
Distortion	32 768 (16)	MEMBA photometry	Source
Noisy	65 536 (17)	MEMBA photometry	Source
Astrometry	131 072 (18)	MEMBA photometry	Source

This paper has been typeset from a TeX/LaTeX file prepared by the author.