WP3.7: Software to Output Metrics That Keep Track of Improvements to the Pipeline Sky Subtraction

Deliverable 3.7.3

Project Acronym	LUSC-B
Project Title	UK Involvement in the Legacy Survey of Space and Time
Document Number	LUSC-B-19

Submission date	31/August/2021
Version	
Status	Final Draft
Author(s) inc. institutional affiliation	Aaron Watkins (Hertfordshire), Chris Collins (LJMU), Sugata Kaviraj (Hertfordshire)
Reviewer(s)	John Stott (Lancaster), Ivan Baldry (LJMU)

Dissemination level	
Public	*Public*

Version History

Version	**Date**	**Comments, Changes, Status**	**Authors, Contributors, Reviewers**
0.1	31/AUG/21	First Draft	Watkins, Collins, Kaviraj, Stott, Baldry
0.2	1/OCT/21	Second Draft	Watkins, Collins, Kaviraj, Stott, Baldry
1.0	1/OCT/21	Accepted Draft	Watkins, Collins, Kaviraj, Stott, Baldry

Table of Contents

List of Figures

1. Executive Summary

Low surface brightness science contains much of the potential discovery space for LSST, in galaxies, solar system, and Milky Way research. Unfortunately, we have previously demonstrated that the current LSST data reduction pipeline's sky subtraction routine over-subtracts flux in the outskirts of extended objects, making low surface brightness science with LSST a potentially very difficult enterprise. Changes to the pipeline are thus required to ensure that LSST's scope is not limited in this way.

As we develop an alternative sky-subtraction routine to resolve this issue, it is necessary to track improvements using quantitative metrics. This in turn requires software that can measure and output these metrics alongside regular runs of the pipeline. This report details the current state of such software, which we are now making available to the LSST:UK community. We also describe future plans for the software, to address its current slow speed and to more easily integrate some version of it into the LSST pipeline itself for use by the data management team.

2. Introduction

As we have established in previous reports [1, 2], the current LSST data reduction pipeline fails to preserve flux in the presence of extended objects, like galaxies. The current implementation of the sky subtraction algorithm overestimates the flux in the local night sky adjacent to such objects in a manner that scales with the objects' brightnesses. We demonstrate this in Figure 1, which shows the total flux loss sustained by model galaxies post-sky-subtraction, normalized by the model area—the positive trend toward brighter magnitudes indicates that the amount of sky being subtracted scales with the model brightnesses. We also demonstrated that most of this flux is being lost at low surface brightness (LSB, below $\mu \approx 26.5$ mag arcsec^{-2}), which we show in Figure 2. Visually, the trend shown in Figure 2 appears as dark over-subtraction rings around extended objects, akin to those seen in HSC PDR1 [6] though occurring at much fainter surface brightness. This problem appears to affect all extended objects, but is most impactful in the LSB regime. This poses a serious problem for the potential scope of LSST science: in galaxies science in particular, much of the potential discovery space available with LSST lies in the LSB regime, including dwarf galaxies (e.g., [3, 7]), tidal debris (e.g., [5]), and intragroup or intracluster light (e.g., [4]). LSB science in other sub-fields will, of course, also be affected, including the study of comet tails, Galactic cirrus, and Galactic or external emission line regions like supernova remnants.

Some improvements to the sky subtraction algorithm thus are merited, as well as a means to track said improvements. The former will be the ultimate goal of WP 3.7, while the latter is the subject of this report.

2.1. Summary of Metrics

Sky over-subtraction is a problem of flux loss. A complete understanding of sky over-subtraction's impact on different science cases thus requires relative and absolute measures—both global and local—of this flux loss. We have thus settled on three basic flux-loss metrics for assessing the LSST pipeline sky-subtraction's impact, which we briefly summarize here:

- Change in total magnitude: Δm
- Total flux loss per unit area: ΔF/A
- Median isophote among all models at which N% loss in surface brightness occurs: $\langle \mu_{\Delta\mu=0.n} \rangle$

In the latter, N refers to the percentage loss we wish to probe, while n is the corresponding decimal fraction in magnitudes. Δm assesses the relative flux loss of an object in astronomically relevant units, to more easily understand the flux loss's impact on magnitude-derived quantities such as spectral energy distributions (SEDs). ΔF/A assesses how much linear flux is being lost per pixel in each object's vicinity, which in turn helps estimate the object's influence on the amount of sky being subtracted locally to it (in the best-case scenario, the object would have no influence on this). Finally, $\langle \mu_{\Delta\mu=0.n} \rangle$ assesses at what surface brightness, on average among all objects tested, this flux loss is actually occurring, as well as what typical losses are experienced at different surface brightnesses. Photometric accuracy near any given survey's limiting surface brightness is typically of order a few tenths of a magnitude, even for azimuthally averaged radial profiles (see, e.g., Figure 10 of [9]); we will thus track this third metric primarily using the 10% flux loss isophote, to ensure that any systematic over-subtraction is pushed to surface brightnesses well-below the expected photometric accuracy at the LSST 10-year depth.

Specific details on how we derive these metrics can be found in our previous Work Package

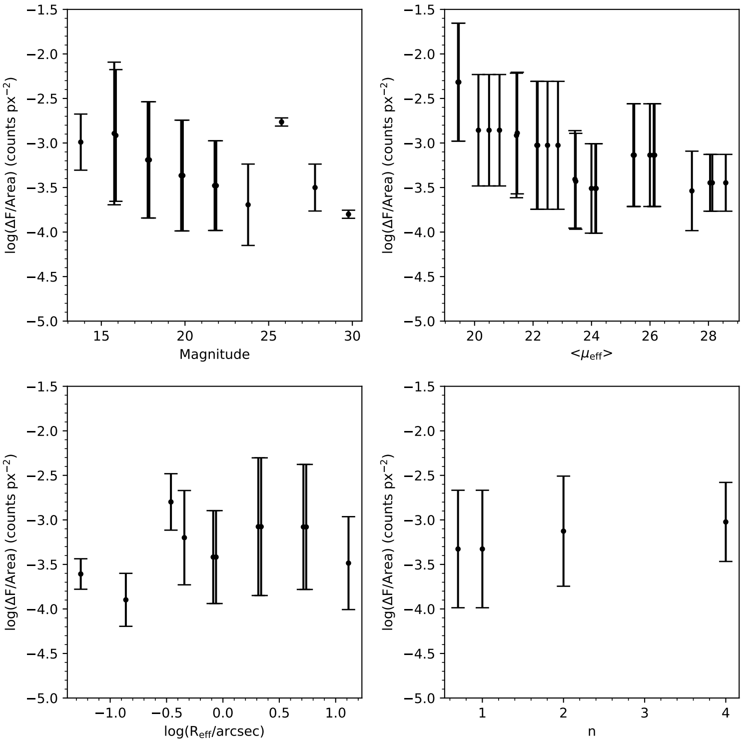

Figure 1: Total flux loss (in HSC coadd flux units, derived as $10^{-0.4(m-27.0)}$ where m is the model magnitude) among single Sérsic index model galaxies injected into the pipeline just prior to sky subtraction, normalized by the model stamp areas in pixels, as a function of four parameters: model magnitude, mean surface brightness within the effective radius, effective radius, and Sérsic index. The full focal plane sky solution—the sky subtraction procedure we will be continuing to test and develop using these metrics—should not be influenced by astrophysical sources, but here we see a correlation between model magnitude and total flux being subtracted as sky.

report [2].

We track each of these metrics by doing comparative surface photometry of model galaxies, as explained in detail in WP 3.7.2 [2]. This report describes the first version of the software used to do this photometry and output those metrics, which is now being made available to the LSST:UK community. This report does *not* describe the final version of this software, which will ultimately be implemented into the LSST pipeline itself to automatically track these metrics as development of an alternative sky subtraction algorithm continues.

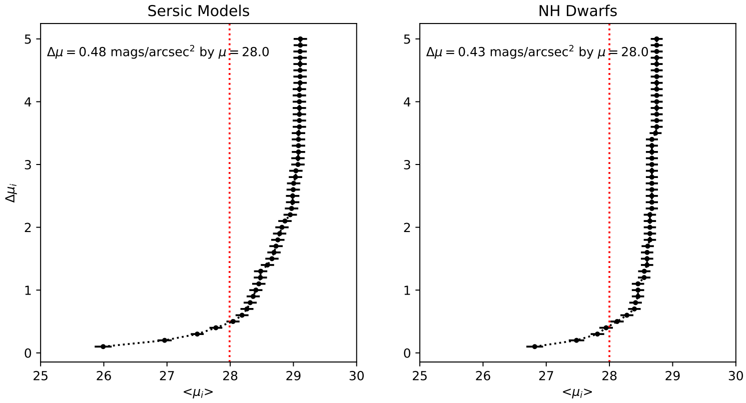

Figure 2: Change in surface brightness as a function of the median surface brightness, among all models, at which that much surface brightness loss occurs. Each point is a step in 0.1 mag arcsec^{-2} in $\Delta\mu$; for example, the first point shows the median surface brightness among all models at which a $\sim$ 10% flux loss occurs ($\mu \sim$ 26–27 mag arcsec^{-2}). The **left panel** shows the trend for single Sérsic index models, while the **right panel** shows the trend for simulated dwarf galaxies from New Horizon [8]. In both sets of models, we see the over-subtraction shows asymptotic behavior with surface brightness, such that the lowest surface brightness isophotes are losing magnitudes of total flux—this is reminiscent of the over-subtraction problem present in HSC PDR1.

3. Deliverable 3.7.3

This software was developed in the Rubin Science Platform (RSP), and requires access to that platform's Notebook aspect to run. As such, all of the software is written in Python. The software relies on image retrieval via the Butler (as of this writing, the second generation; this will be upgraded to the third generation as soon as new model injections are made available in the third generation Butler format) and performs photometry directly on these images, the outputs of which are stored as Python dictionaries in `pickle` files. Diagnostic plots and metric measurements are then made from these `pickle` files. The following subsections describe these three broad aspects of the software. The software itself can be found here: `https://github.com/lsst-uk/sky-estimation-WP3.7/tree/master/measureMetrics`

3.1. Image retrieval

Retrieving image cutouts located at a specific celestial coordinate using the RSP is unfortunately a non-trivial task. To do our photometry, we needed two kinds of images: visit-level (or pre-sky-subtraction) cutouts of models, to measure the baseline photometry; and coadd-level (or post-sky-subtraction) cutouts of models, to compare against the former to measure the sky-subtraction's impact.

The former case is required to ensure that models are being injected properly, by comparing model photometry pre-sky-subtraction to their input catalogue parameters. An added benefit of this step is the direct measurement of model radial profiles as they appear post-PSF-convolution, an effect that is difficult to estimate theoretically without fully understanding the specifics of this convolution as performed by the LSST pipeline. This step requires retrieval of individual CCD exposures ("visits", in the LSST parlance), meaning cutout sizes here are limited by the CCD size (2048 × 4176 px), specifically the shortest axis (2048 px), as our model stamps are all square. The visit and CCD identifiers are, rather inconveniently, associated with the `lsst.afw.image.exposure` objects returned by the Butler, so in order to find a visit in which a given model is located far enough from the CCD boundaries to not be truncated, one must loop through every visit associated with the patch on which the model is injected until such an exposure is located. The returned model injections have each been convolved with the PSFs specific to those retrieved visits, which are subtly different from the coadd PSF, but this effect is mild and limited to the model cores, not the LSB outskirts that our analysis is primarily interested in.

The latter case is required to understand the sky subtraction's impact on the models, by comparison with either the model catalogue parameters or the pre-sky-subtraction photometry provided by the above case. The coadd image retrieval is based on patches ($\sim 4000 \times 4000$ px), hence cutout size is less of an issue. It is only a problem if the requested cutout bounding box exceeds the patch boundaries, but this problem is surmountable through retrieval of images from adjacent patches and some stitching together of images. We designed our current model catalogue so that no models are injected between patches, but we include such a "robust" image retrieval function, written and provided to us by Markus Dirnberger (private communication), because it is a manifestly useful function.

3.2. Surface photometry

Because one of our metrics relies on knowing where the post-sky-subtraction surface brightness profile deviates from its pre-sky-subtraction counterpart, and because the LSST pipeline only outputs course-resolution (logarithmically spaced) integrated fluxes within circular apertures, we do our own surface photometry on the images we retrieve by the means discussed in the previous section. We do this in the same manner in which it has been done for decades: by measuring the mean (and median) flux values within growing consecutive elliptical annuli of fixed width, out to some maximum elliptical semi-major axis (SMA) length. We also measure the total integrated flux contained within growing consecutive elliptical apertures out to that same maximum SMA—the curve of growth. The largest value of the curve of growth is roughly equivalent to the object't total flux. The module is general purpose: we demonstrate this in Figure 3 using a barred spiral galaxy (IC 1010) found in HSC tract 9615. Using 1 px-wide bins on this 1000×1000 px stamp, out to a maximum SMA of 500 px, takes on average 1.16 s (measured using Python's `timeit` module).

We provide the option to include a mask value for this photometry as well, such that all pixels in the input image array with that value are ignored when measuring the mean (and median) surface brightnesses. This mask only applies to the derived surface brightness profiles, and not the curve of growth; to properly estimate the unmasked flux, the latter would require some estimate of the missing light beneath the mask, which is unnecessary for the purposes of this work package at the moment. Given this limitation, the general user is advised that the curves of growth produced by this code are only approximate, and so should not be used for any rigorous analysis.

3.3. Deriving the metrics

To derive our metrics, we do surface photometry on each injected model in the manner described in WP3.7.2 [2], storing the output photometry as Python dictionaries in `pickle` files. We then use the profiles stored in these `pickle` files to derive our various metrics, all of which are measured simply through subtraction of the relevant pre-sky-subtracted value from its coupled post-sky-subtracted value (for example, Δm is simply the difference between the pre- and post-sky-subtracted integrated magnitudes, pulled from the curves of growth). The usage of `pickle` files to store photometry is critical because of the time it takes to run image retrieval and surface

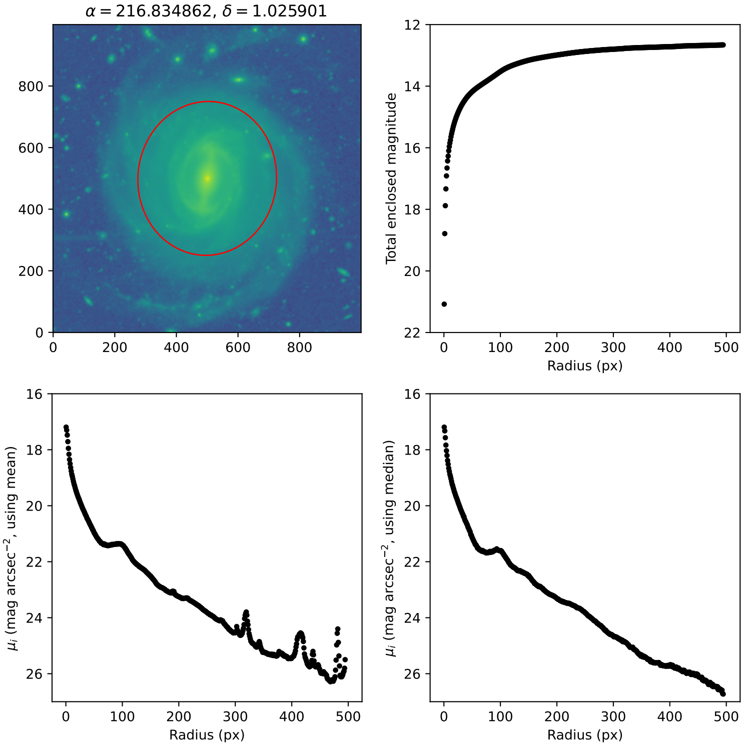

Figure 3: Demonstration of the general usage capability of our surface photometry module, used on a large barred spiral (IC 1010) found by eye in HSC tract 9615. The galaxy image is shown in the **top left** panel, along with its approximate central coordinates and an ellipse with the rough by-eye estimated parameters used for the photometry. The **top right** panel shows the curve of growth. The **bottom left** and **bottom right** panels show the surface brightness profiles made using the mean and median values in each isophote, respectively. We applied no masking, as this is simply a demonstration.

photometry on every model. Doing photometry on each model takes only a fraction of a second on average (excluding image retrieval, which takes several seconds), but this scales non-linearly with model size—as demonstrated above, a 500 px radius model requires $\sim$ 1 sec, but a 2000 px model requires $\sim$ 90 sec. Both catalogues combined have more than 2000 models total, meaning it takes several hours to analyze all models in all photometric bands. Later versions of this software may address the slow speed, possibly by rewriting some of the code in a compiled language. A factor of 100–1000 would be necessary if we wanted to implement this package directly into the pipeline.

Given this, and also to produce a continual independent check on our own photometry, in the near future we will write complementary software to produce approximate versions of these metrics using the pipeline output catalogues. These are limited in depth by image background noise, may suffer from deblending issues and source confusion, and in their surface photometry have courser resolution (due to logarithmic spacing) and smaller maximum radii (70 pixels) than we might desire, but should provide more easily trackable versions of the metrics to be integrated directly into the pipeline itself, at least until we increase the speed of our own software to an acceptable level.

We provide the current model catalogues and photometry `pickle` files in the module directory. We also provide a number of visualization tools. Usage of each of these modules is also demonstrated in accompanying Jupyter Notebooks: the `CataloguePhotometry.ipynb` notebook can be run to reproduce the included `pickle` files, while the `Overview.ipynb` notebook shows a complete summary of all pertinent results from the version of the catalogue available at the time of writing. This includes our two primary results, which we display in this report in Figures 1 and 2, which show that the pipeline over-subtracts flux significantly starting at surface brightnesses below $\mu \sim 26$, in a manner that scales with the brightnesses of our injected models.

References

[1] Report on optimal metrics for measuring the impact of the LSST pipeline sky subtraction on low-surface-brightness flux at different spatial scales, Project Deliverable D3.7.1

[2] Report on Mock Testing Results, Project Deliverable D3.7.2

[3] Dalcanton J. J., Spergel D. N., Gunn J. E., Schmidt M., Schneider D. P., 1997, AJ, 114, 635

[4] Burke C., Collins C. A., Stott J. P., Hilton M., 2012, MNRAS, 425, 2058

[5] Kaviraj, S., 2014, MNRAS, 440, 2944

[6] Aihara, H., et al., PASJ 71, 114

[7] Martin G., et al., 2019, MNRAS, 485, 796

[8] Dubois, Y., et al., 2021, AA, 651, 109

[9] Trujillo, I., Fliri, J., 2016, ApJ, 823, 123

Annex A Acknowledgements

We are grateful to the members of the DM pipeline team who have worked with us throughout this project: Lee Kelvin, Yusra AlSayyad, Robert Lupton, Sophie Reed, and Dan Taranu. Without them, this project will not be possible.

We also thank Johan Knapen for his continued support of this project and valuable feedback.