

Deep co-added sky from Catalina Sky Survey images

Akshat Singhal¹,^{*} Varun Bhalerao,¹ Ashish A. Mahabal^{1,2,3}, Kaustubh Vaghmare^{1,2,4}, Santosh Jagade,² Sumeet Kulkarni,^{5,6} Ajay Vibhute,² Ajit K. Kembhavi,² Andrew J. Drake,³ S. George Djorgovski,³ Matthew J. Graham^{1,3},^{*} Ciro Donalek,^{3,7} Eric Christensen,⁸ Stephen Larson⁸ and Edward C. Beshore⁸

¹Indian Institute of Technology Bombay, Powai, Mumbai 400076, Maharashtra, India

²Inter-University Centre for Astronomy and Astrophysics (IUCAA), Post Bag 4, Ganeshkhind, Pune 411007, Maharashtra, India

³Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

⁴Persistent Systems Ltd., Pune 411016, Maharashtra, India

⁵Indian Institute of Science Education and Research, Homi Bhabha Road, Pune 411008, Maharashtra, India

⁶Department of Physics and Astronomy, The University of Mississippi, University, MS 38677, USA

⁷Virtualitics, Inc., Pasadena, CA 91101, USA

⁸Lunar and Planetary Laboratory, University of Arizona, Tucson, AZ 85721, USA

Accepted 2021 July 24. Received 2021 July 24; in original form 2020 August 4

ABSTRACT

A number of synoptic sky surveys are underway or being planned. Typically, they are done with small telescopes and relatively short exposure times. A search for transient or variable sources involves comparison with deeper baseline images, ideally obtained through the same telescope and camera. With that in mind, we have stacked images from the 0.68 m Schmidt telescope on Mt. Bigelow taken over 10 yr as part of the Catalina Sky Survey. In order to generate deep reference images for the Catalina Real-time Transient Survey (CRTS), close to 0.8 million images over 8000 fields and covering over 27 000 sq. deg have gone into the deep stack that goes up to 3 mag deeper than individual images. CRTS system does not use a filter in imaging; hence, there is no standard passband in which the optical magnitude is measured. We estimate depth by comparing these wide-band unfiltered co-added images with images in the *g* band and find that the image depth ranges from 22.0 to 24.2 across the sky, with a 200-image stack attaining an equivalent AB magnitude sensitivity of 22.8. We compared various state-of-the-art software packages for co-adding astronomical images and have used SWARP for the stacking. We describe here the details of the process adopted. This methodology may be useful in other panoramic imaging applications, and to other surveys as well. The stacked images are available through a server at Inter-University Centre for Astronomy and Astrophysics.

Key words: surveys – stars: general – transients – imaging – techniques: image processing.

1 INTRODUCTION

Time domain astronomy is a vibrant and growing field powered mainly by synoptic sky surveys, often using relatively modest diameter telescopes (for a review of different sky surveys, see e.g. Djorgovski et al. 2013). The temporal spacing in the images at a given location, along with filters, size of the mirror, total area covered, etc., determines the utility of the survey for time-varying phenomena at different time-scales and depth. When looking for rare transient events, the availability of a deeper image is essential for determining the nature of the transient sources. Ideally, such comparison images will have the same sampling and filters as the survey data themselves, and this is commonly done in all synoptic sky surveys. One such survey is the Catalina Sky Survey (Larson et al. 1998, 2001), which is unique in the combination of area coverage, depth, and the number of epochs, and the overall time baseline. All surveys, regardless of their power, will benefit from such co-addition.

The Catalina Sky Survey incorporates multiple telescopes but in this paper we focus on the 0.68 m Schmidt on Mt. Bigelow in Arizona (CSS¹). The original purpose of the survey has been to look for Near-Earth Objects (NEOs) and it has been highly successful at it.² The Catalina Real-time Transient Survey (CRTS), as the name suggests, undertakes a real-time search for transients in catalogues made from the survey data (Drake et al. 2009; Djorgovski et al. 2011; Mahabal et al. 2011, etc.). Thousands of high-significance candidates have been found by CRTS and reported to the world in real time.³ The emphasis of CRTS has been to search for transients, and develop automated methods to characterize and classify them accurately despite the sparseness of available data. Given that spectroscopic classification is not practical for all detected transients, even preliminary classifications based on the survey data are valuable, and deeper baseline images are a critical part of that, e.g. detecting stellar counterparts in their quiescent state or faint host galaxies of extragalactic transients. The most widely used standard reference

¹<http://www.lpl.arizona.edu/css/>

²See Figures at <http://neo.jpl.nasa.gov/stats/>

³<http://crts.caltech.edu>

* E-mail: crts@iucaa.in, akshats@iitb.ac.in (AS)

Table 1. CSS images taken in Follow-up, User, N, and S modes.

Field type	Code	# Fields	# Fields discarded	Max images	Total images included in co-adds
Follow-up	F	2079	15	40	12139
User	U	2113	81	61	12312
North	N	2370	1	545	528 880
South	S	1218	3	421	186 703
Other	–	11	3	12	62
Total	–	7791	103	–	740 096

image sets providing optical images for a large part of the sky are the Sloan Digital Sky Survey (SDSS; SDSS Collaboration 2017) and the Pan-STARRS survey (PS1⁴; Flewelling et al. 2016).

CSS has obtained hundreds of epochs over the years covering 27 000 sq. deg and the co-added images can be deeper than SDSS or PanSTARRS. However, CSS images are acquired without using any filter, with the wavelength response being defined by the telescope and camera system. Hence, there is no particular band in which optical magnitude is measured. We refer to these as ‘unfiltered’ images and ‘unfiltered’ magnitudes, and give comparisons with some standard filters where appropriate. On the other hand, CSS has the added advantage that the observations span several years, and it is much more likely for it to have caught multiple brightening episodes, or the occasional one in more objects (the average total effective exposure for the stacked images is 3000 s in CSS as against 54 s for the 2.5 m SDSS). Further, the area covered by SDSS DR16 is under 15 000 sq. deg. While the larger pixels of CSS – 2[″].5 pixels, which amounts to undersampling compared to 0[″].4 of SDSS – are a downside in crowded fields and can introduce issues in counterpart identification, the data set is still immensely useful for various science cases involving isolated sources.

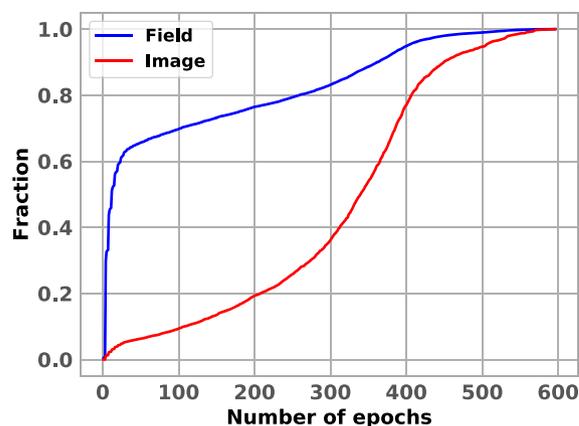
During the nightly transient detection process, CRTS uses a reference median-combined stack of at least 20 images (Drake et al. 2009). A deeper stack will lead to a deeper comparison image, and allow us to rule out fainter but consistent sources when looking for transients that brighten by several magnitudes. With its large pixels, CSS co-adds are not ideal for de-blending objects close to each other on the sky, but they are still useful at higher Galactic latitudes where we do not reach the confusion limit despite reaching fainter levels by combining multiple images.

Past co-adding efforts involving astronomical surveys include Deep Sky⁵ that used 10–100 images from Palomar-QUEST (Djorgovski et al. 2008), SN Factory (Aldering et al. 2002), and Near Earth Asteroid Tracking (Pravdo et al. 1999); Palomar Transient Factory (PTF; Law et al. 2009), Intermediate PTF (Cao, Nugent & Kasliwal 2016), Zwicky Transient Facility (ZTF; Bellm), the DESI Legacy Imaging Surveys (Dey et al. 2019), and a few others using SDSS and *Wide-Field Infrared Survey Explorer* (WISE) images (Annis et al. 2014; Lang 2014; Meisner, Lang & Schlegel 2017a; Meisner et al. 2017b).

In this paper, we describe the co-addition of CSS images to produce deep stacks. We evaluated several co-addition software packages like MONTAGE,⁶ SWARP,⁷ and I-CORE⁸ to see which one best suits our

Table 2. Comparison of statistics of the number of epochs of all observed fields for the 2003–2012 period, and the subset selected for co-addition. This subset was created by rejecting some images as discussed in Section 2. Each category has two columns: ‘All’ for all fields observed in CSS, and ‘NS’ for the N and S fields that make up the regular survey grid. All data are from telescope the 68 cm Catalina telescope.

Statistic	Total images		Co-added subset	
	All	NS	All	NS
Min	1	1	1	1
Max	596	596	545	545
Mean	101	212	95	199
Stdev	146	156	138	145
Median	12	217	8	205

**Figure 1.** The solid blue line shows the cumulative number observed fields as a function of increasing number of epochs. As defined in the text, *number of epochs* is the number of times a field is observed. Well over half the fields were observed fewer than 30 times each, majority of which are User (U) and Follow-up (F) fields. The solid red line indicates the cumulative fraction of total images covered by fields below a given number of epochs.

needs in terms of handling the image sizes and numbers as well as artefacts and sporadic World Coordinate System (WCS) issues. In Section 2, we describe the data, followed by comparison of various stacking methods in Section 3. We describe final implementation of the co-addition in Section 4, and summarize and discuss future steps in Section 5.

2 DATA

To look for NEOs, CSS uses a fixed sequence: four 30-s images spaced by about 10 min each. CSS marches from one field to another for 10 min, and then comes back to the first field for its next image. These field locations are fixed within pointing errors. The field is

⁴<https://panstarrs.stsci.edu/>

⁵<https://c3.lbl.gov/nugent/deepsky.html>

⁶<http://montage.ipac.caltech.edu/>

⁷<https://www.astromatic.net/software/swarp>

⁸<http://web.ipac.caltech.edu/staff/fmasci/home/icore.html>

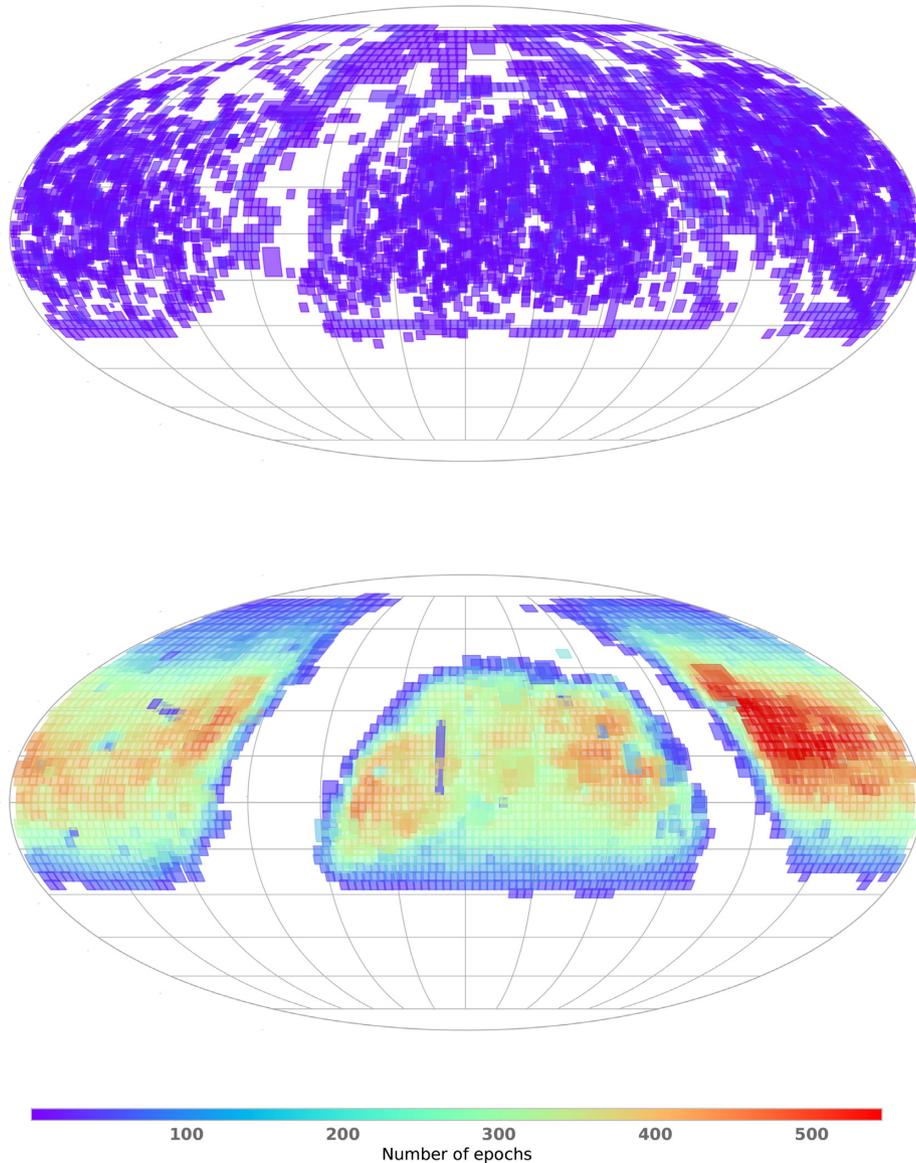


Figure 2. Distribution of fields over the sky, colour coded to show the NoE (number of images): top: shows 4871 fields with NoE between 1 and 30 and bottom: shows 2801 fields with NoE between 31 and 545.

identified with a keyword in the header. In the data, we identified 7894 distinct field IDs. These include the North (N) and South (S) fields that are part of an all-sky grid in northern and southern celestial sphere and are used for regular observations, whereas User (U) and Follow-up (F) fields are occasional and sporadic, defined by either a user for some off-the-grid observation or used for follow-up targets, respectively (Table 1). The U and F fields are typically scanned exactly thrice (in the standard four-image mode), leading to a median of twelve visits per field. This is in stark contrast with the N and S fields, which were imaged a median of 217 times each between 2003 and 2012, the period covered here (Table 2). Large variations are also present in the number of visits for N and S fields – for instance, parts close to the plane of the Galaxy are scanned infrequently. Adding the offset U and F fields with the N and S fields during stacking would have led to unnecessary depth non-uniformity. Hence, in this study, we have treated all four field types on par with each other, and created separate co-added stacks for each field ID.

A small fraction of images have incorrect WCS in their headers. Such images were identified and excluded from co-addition. Each image is $171 \text{ arcmin} \times 171 \text{ arcmin}$, and 4110×4096 pixels with a pixel size of $2''.5 \times 2''.5$. The uncompressed size of each image is 33 MB, and, 5–7 MB when H-compressed (White & Percival 1994), the standard way CSS stores images. The total compressed size for all images for the 2003–2012 period is about 5 TB.

Out of 7894 fields, 4973 were observed fewer than 30 times. We use the word *number of epochs* (NoEs) to describe how many times a field was imaged. There are 2282 fields with NoE 4, 837 fields with NoE 8, 409 fields with NoE 12, and so on. All the low-NoE fields together contain only a small fraction of the total number of images (4973 fields – 63 per cent of the total – have a NoE of 30 or less, and the image contribution is ~ 5 per cent of the total. See Fig. 1).

The high-NoE fields are relatively well tiled as compared to the low-NoE ones that are sporadically distributed, and with differing overlap (Fig. 2).

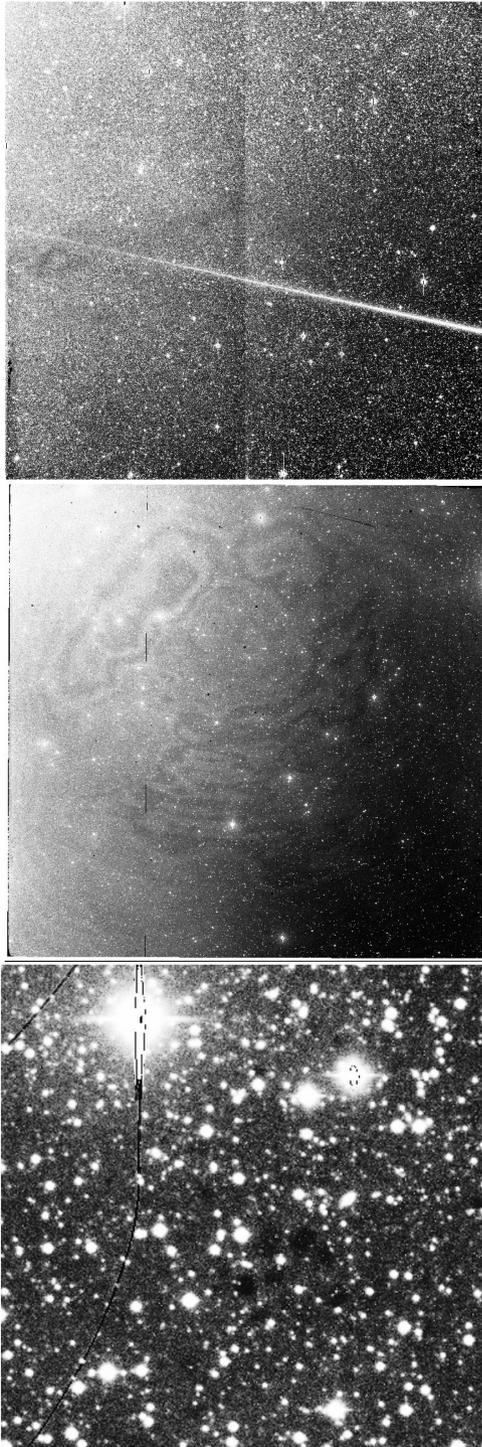


Figure 3. Examples of images in CRTS repository with artefacts likely to adversely affect the quality of final co-added images. The artefacts are: Top: A satellite trail, and discontinuous background variation across the chip. Middle: Fringes, primarily caused by red/infrared light. Bottom: Round holes due to dust and dirt. The knowledge of the location of these spots was part of the transient detection pipeline. They were constant over long periods, but required to be updated periodically. Diffraction spikes and charge ‘bleeding’ are seen for bright stars. Most of these artefacts are masked and not included in final images.

CSS images show the usual artefacts present in typical imaging data, including satellite trails, CCD fringing, bleeding, spikes for bright sources, and poorer image quality for edge pixels (Fig. 3). In particular, we note that calibrations for the CSS data (e.g. flat-field images) were infrequent, which adds some uncertainty to photometry in individual images. However, we used the images that were flat-fielded during the routine runs as in our experience; reprocessing the flat-fields would not have resulted in a substantial improvement in the quality of the data.

3 COMPARISON OF STACKING METHODS

Many software tools, such as MONTAGE (Berriman et al. 2003), SWARP (Bertin & Arnouts 1996; Bertin 2010), and I-CORE (Masci 2013a, b), are available for co-adding images. We discuss some of these tools with an emphasis on MONTAGE and SWARP. While these two tools use different algorithms for co-addition, we found that the final products were of comparable quality, resulting in an increased depth of up to 3 mag for our deepest stacks, reaching a typical sensitivity of $m_g \sim 23$. This is roughly consistent with the central limit theorem and the depth increases as \sqrt{N} increases, where N is number of co-added images.

3.1 MONTAGE

MONTAGE (Berriman et al. 2003)⁹ is a WCS enabled open-source mosaicking software suite for FITS images. MONTAGE co-adds images using ‘forward pixel mapping’, where each pixel from an input image is projected on to an output image in Cartesian space such that flux and astrometric information is preserved. A direct implementation of this approach would lead to inclusion of backgrounds in the final co-adds, causing spatial non-uniformity based on exposure and also increasing the overall noise. MONTAGE provides tools to remove the background in each image before co-adding. MONTAGE also offers a variety of statistical methods (mean, median, weighting by number of images contributing to a pixel, etc.) for combining the input images to create the output.

Rather than a single package, MONTAGE is built as a set of tools that can be used in various combinations to create the final co-added image. We used the default statistic, i.e. mean, for co-addition and compared two work flows with different methods of background estimation (Table 3). In ‘MontageA’, we use `mDiffFitExec` for comparison of each pair of input images to estimate and remove the background. Time taken to co-add the images using this method scales as the square of the number of input images and is unwieldy for large image stacks. In ‘MontageB’, the background model is created to achieve the best global fit by a least-squares method. This method is faster than MontageA, as it scales nearly linearly with the number of images.

The time complexity of the algorithm used by MONTAGE goes as $\mathcal{O}(n^p)$, where $p > 1$, and we can use multiple CPU cores in parallel to make co-adds faster. To do that, we divide N images into m groups of $k = N/m$ images each, and co-add each group separately to get m temporary mosaics, which are then combined to get the final mosaic. Even the k images can be subdivided beforehand if k is large. Appropriate weight files are used to keep track of

⁹<http://montage.ipac.caltech.edu/>

Table 3. Workflow for the MontageA and MontageB methods discussed in Section 3.1. At the right, we give a short description of each MONTAGE command. More details are available at the montage web site <http://montage.ipac.caltech.edu/docs/index2.html>.

MontageA	MontageB	Description
mImgtbl	mImgtbl	Collect data from headers that comply with the FITS standard and create a table. It reports a count of images that fail that check.
mMakeHdr	mMakeHdr	From a list of images to be mosaicked together, generate a FITS header that best describes the output image.
mProjExec	mProjExec	Reproject the image.
mDiffFitExec	↓	Compare images in pairs to determine overlaps.
↓	mImgtbl	As described above.
mBgModel	mBgModel	Determine a set of corrections to apply to each image in order to achieve the ‘best’ global fit.
mBgExec	↓	Remove background from FITS images.
↓	mFlattenExec	Reject outlier pixels to fit a plane to an image, then subtract that plane from the image.
mAdd	mAdd	Co-add the reprojected images.

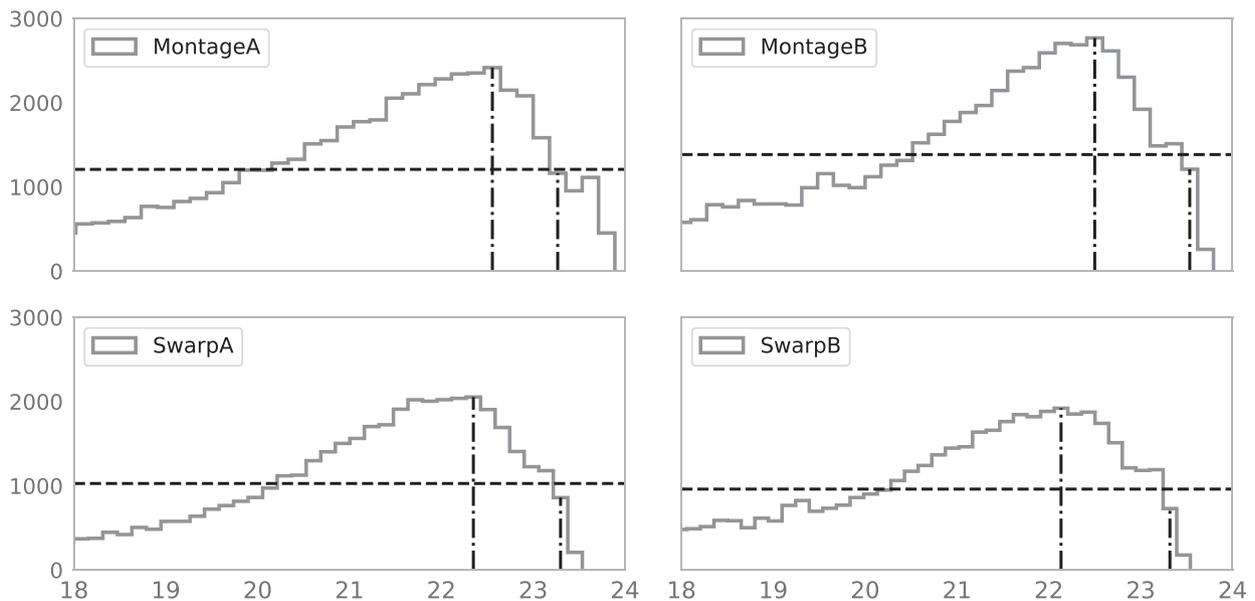


Figure 4. Histograms of source magnitudes (x-axis) in the final co-added images for one set, using the four co-addition methods. We used SEXTRACTOR isophotal magnitudes (MAG_{ISO}) and added an arbitrary but constant zero-point of 28 to the instrumental magnitudes. Vertical dot-dashed lines show the mode magnitude (m_{mode}) and the limiting magnitude (m_{lim}). The horizontal dashed lines mark half the number of sources as the histogram mode, and m_{lim} can be read off as the first bin below the rightmost intersection of this line with the histogram.

the number of times each pixel has been co-added in the co-added image. We tested this method to a limited extent, and found that the resultant co-adds were comparable to the co-adds if done all at once.

3.2 SWARP

SWARP (Bertin & Arnouts 1996; Bertin 2010) is part of the `astromatic.net`¹⁰ suite of packages and works with other tools like SEXTRACTOR. Unlike MONTAGE, SWARP uses inverse mapping that is much faster than forward mapping of an arbitrarily large image, but arguably slightly less accurate. In this method, WCS information is used to obtain an outline of the image to be output, and then input pixel values are interpolated using LANCZOS3 (Bertin 2010)

method to write out the output pixels. The main advantage SWARP has is its robust image artefact removal (Gruen, Seitz & Bernstein 2014). SWARP also allows image co-addition using a variety of statistics including average, median, weighted by various parameters, sigma clipping, etc. SWARP creates a background estimate over a grid following the SEXTRACTOR algorithm (Bertin & Arnouts 1996) and subtracts it before co-adding the images. After various trial runs, we selected two approaches for detailed study. In the first approach (hereafter ‘SwarpA’), we follow the method described in Gruen et al. (2014), using the local standard deviation (σ) and the local median (μ) to reject outlier pixels with values $>4\sigma + 0.3\mu$. In the second approach, we simply use the average of input pixels as the final output value (hereafter ‘SwarpB’). Lastly, it is straightforward to provide a bad pixel mask to SWARP for ignoring certain pixels from input images. This proved to be a very useful feature, as discussed in Section 3.5.

¹⁰<http://www.astromatic.net/software/swarp>

Table 4. Comparison of the results of 50 images added together, using four different algorithms, two each from MONTAGE and SWARP. An arbitrary but constant zero-point of 28 was used. m_{lim} is the limiting magnitude. The FWHM reported in pixels is the average for all stars. The last column shows the time taken by the co-addition to complete.

Method	Mode m_{mode}	Mag m_{lim}	FWHM pixels	Sources (3.0σ)	Time (s)
MontageA	22.5	23.3	3.6	52 519	4079
MontageB	22.5	23.5	3.4	58 498	2528
SwarpA	22.3	23.3	2.9	44 648	598
SwarpB	22.1	23.3	3.2	50 422	262

Table 5. Fraction of SExtractor-detected sources common to various methods. The numbers in each column denote the fraction of sources for that method detected in co-adds created by other methods. For instance, 86 per cent of sources in the MontageB co-add were also detected in the MontageA co-add.

Method → Comparison ↓	MontageA	MontageB	SwarpA	SwarpB
MontageA	1.00	0.86	0.70	0.70
MontageB	0.77	1.00	0.65	0.66
SwarpA	0.82	0.85	1.00	0.83
SwarpB	0.73	0.77	0.73	1.00

Table 6. Pairwise difference in magnitude for all the common sources (‘All source’) and for only those sources that follow a linear relation (‘Subset’). Ideally, difference in magnitude of all common sources should have been constant, but that was not the case.

	MontageA	MontageB	SwarpA
All source			
MontageB	0.02 ± 0.24	–	–
SwarpA	-0.01 ± 0.22	-0.02 ± 0.23	–
SwarpB	-0.06 ± 0.24	-0.07 ± 0.21	-0.05 ± 0.22
Subset			
MontageB	0.01 ± 0.25	–	–
SwarpA	-0.01 ± 0.15	0.01 ± 0.12	–
SwarpB	-0.05 ± 0.20	-0.04 ± 0.15	-0.06 ± 0.12

3.3 Other methods

In addition to MONTAGE and SWARP, we also considered some other software tools for image co-addition. A particular case is I-CORE,¹¹ a generic FITS co-addition and mosaicking software originally written for the *WISE* images (Masci 2013a, b). A key feature of I-CORE is to utilize small dithers between images to enhance the resolution of the final co-add. However, this requires good-quality flat-field images, which were not available by default for the CSS data.

Some methods were not at a mature enough stage of development for our work. For instance, IP2¹² is a software created primarily for image subtraction (Hartung et al. 2012; Hartung & Shukla 2013) but the team is creating an image co-addition version as well. Similarly, the Large Synoptic Survey Telescope (LSST; Ivezić et al. 2008) is developing an image co-addition pipeline tuned to future LSST images. This pipeline will be configurable to other data sets, and has been tested with data from the Dark Energy

¹¹<http://web.ipac.caltech.edu/staff/fmasci/home/icore.html>

¹²<http://iccs.lbl.gov/research/isaac/IP2.html>

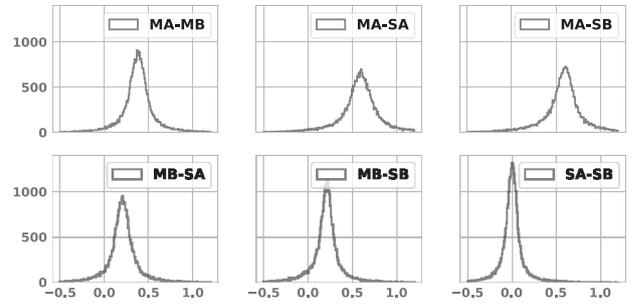


Figure 5. Distribution of pairwise FWHM difference. The X-axis represents ΔFWHM in pixels and the Y-axis represents count. The difference in FWHM is in the same order as in the legend. For example, the first figure in the top row represents pixel-FWHM of MontageA–MontageB. The mean and standard deviation values are given in Table 8.

Table 7. Mean and standard deviation of FWHM (in pixels) of common sources detected in co-added images using different methods for a few image sets. As the standard deviations are comparable to the mean, we calculate pairwise differences for the methods and report them in Table 8.

	MontageA	MontageB	SwarpA	SwarpB
Set 1	3.7 ± 1.8	3.4 ± 1.7	3.0 ± 1.4	3.3 ± 2.2
Set 2	4.0 ± 2.2	4.1 ± 2.6	3.5 ± 1.8	3.7 ± 2.5
Set 3	3.5 ± 1.8	3.2 ± 1.9	2.8 ± 1.4	3.0 ± 1.8

Table 8. Pairwise difference of FWHM (in pixels) between different methods, for sources that are common in all methods. Each value gives the difference between the column name and the row name: For instance, the FWHM of sources in MontageB is 0.21 ± 0.23 pixels larger than FWHM measured in SwarpA co-adds. The distributions of such differences are plotted in Fig. 5.

	MontageB	MontageA	SwarpB
SwarpA	$+0.21 \pm 0.23$	$+0.58 \pm 0.27$	$+0.01 \pm 0.21$
SwarpB	$+0.20 \pm 0.23$	$+0.58 \pm 0.27$	–
MontageA	-0.38 ± 0.21	–	–

Camera (Flaugher et al. 2015). The Pan-STARRS collaboration (and references therein; Flewelling et al. 2016) utilizes CosmoDM (Desai et al. 2015) for their image co-addition. Goldstein et al. (2019) are developing a method for ZTF (Bellm E.). All these software would need significant work for adapting them for the task of CSS image co-addition. As the primary focus of our research was to create deep all-sky images, it was decided that testing all possible methods would yield a limited benefit relative to the invested effort. We also note that after the completion of our work, an optimal image co-addition technique has been published by Zackay & Ofek (2017a, b). Here, they use matched filtering and take individual PSFs into consideration without degrading them. They apply it to simulated data and limited data from PTF (Law et al. 2009; Rau et al. 2009).

Other methods like Drizzle (Fruchter & Hook 2002) have been used to improve image resolution through co-addition by shrinking pixels and remapping. Given the size of the survey, our main concern was quick co-addition at native scale of hundreds of images each over thousands of pointings and we did not investigate relative benefits of adopting methods like Drizzle.

Table 9. Key parameters for SwarpB configurations.

Keyword	Value	Notes
WEIGHT_THRESHOLD	0.5	Mask pixels with value <0.5 will be ignored in the co-add. The actual mask is binary, with zeros for bad pixels and ones for good pixels
COMBINE_TYPE	CLIPPED	Use the clipping algorithm as per Gruen et al. (2014)
CLIP_AMPFRAC	0.3	Default value recommended by Gruen et al. (2014)
CLIP_SIGMA	4.0	Default value recommended by Gruen et al. (2014)
SUBTRACT_BACK	Y	Removes background, eliminating most background-related artefacts from co-added image
DELETE_TMPFILES	Y	Deletes the temporary files created during the co-add process.
IMAGEOUT_NAME	[Field_ID]_clip_mask_SBG_Y.fits	The suffix signifies the major algorithm properties: sigma clipping, masking input images, and background subtraction
WEIGHTOUT_NAME	[Field_ID]_clip_mask_SBG_Y.weight.fits	–
WEIGHT_IMAGE	Mask_S.fits	The name of the mask file shown in Fig. 11

Table 10. Keywords added in the final CSS image co-adds.

Keyword	Value
AUTHOR	akshat.singhal014@gmail.com
COADDN	number of input images (number of images co-added)
DATET	Date and time of creation of file
INSTI	‘IUCAA’
EPOCH	2000
DETECTOR	‘IMAGER LABS’
FILTER	‘NONE’
TELESCOP	‘CATALINA SCHMIDT 003’
OBSERVAT	‘UA Bigelow Station’
CRTSDATA	Total number of images in CRTS repository for this field

3.4 Quantitative analysis

We arbitrarily selected 50 images each from a few fields as the test sample for comparing the co-addition methods MontageA, MontageB, SwarpA, and SwarpB. We used SExtractor to identify all sources in the image, with the ‘DETECT_THRESHOLD’ parameter set to 3.0.¹³ We wish to compare the full width at half-maximum (FWHM) of sources and the depth of the co-added images across the four methods.

To measure depth, we need to establish a reproducible definition of the limiting magnitude. Here, we add a constant but arbitrary value of 28 to convert instrumental magnitudes to apparent magnitudes. As will be seen later, this is a reasonable value. Typical histograms of magnitudes of all sources detected in an image (Fig. 4) show increasing number of stars at fainter magnitudes up to a certain point (m_{mode} , the mode of the distribution), and then a decline to zero. The exact drop to zero is very likely affected by spurious sources in the image; on the other hand, m_{mode} was seen to be clearly brighter than the faintest detected sources. As a result, we define the limiting magnitude m_{lim} as the magnitude bin fainter than the mode, which has half as many sources as in the mode bin. This is graphically shown in Fig. 4. Armed with this definition, we can compare depths of the co-adds created by various methods by comparing the limiting magnitudes as show in Table 4, or by the total number of sources detected in an image. The latter method is based on the reasonable assumption that bright sources will be detected in all images, and any differences in source counts are likely caused by the presence of fainter sources.

¹³Pixels that are 3σ above the local background are considered part of a star.

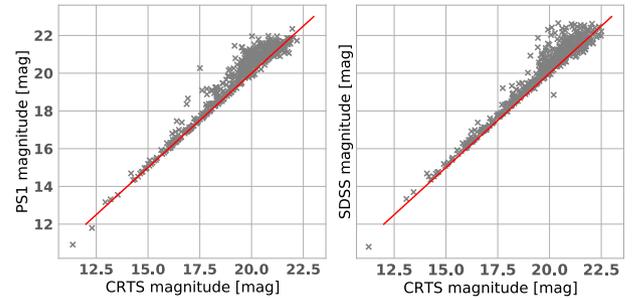


Figure 6. Constant zero-point offset added to the SExtractor *MAG_ISO* of the CRTS sources such that the magnitudes of the PS1 and SDSS (left to right) catalogue roughly match the common sources. The scatter plot shows comparison of SDSS and PS1 magnitudes with offset SExtractor magnitude superimposed with slope 1 linear line.

We find that in the four methods compared, MontageB consistently has the faintest m_{lim} , while SwarpB has the poorest. It is non-trivial to account for spurious sources detected in the image, which may include image artefacts (diffraction spikes, unmasked bad pixels, cosmic rays, etc.) and random noise. In this context, it is not surprising that both MONTAGE methods lead to more sources in the final co-added image as they do not reject any pixels. This is thus a trade-off between using all data, potentially boosting the signal-to-noise ratio for brighter sources, and retaining all noise, degrading data quality at the faint end.

To better understand the variation in total number of sources, we cross-matched the source catalogues created by each of the four methods using a matching radius of 2 arcsec (Table 5). We find that SwarpA is most effective in finding sources detected by other methods. In other words, given a source catalogue created by any algorithm, SwarpA has the highest overall reproducibility for finding sources in that catalogue. This, along with the fact that SwarpA co-adds seem to have the least number of sources in them, suggests that SwarpA is more effective at suppressing spurious sources.

CRTS images are unfiltered, which makes it difficult to compare the measured magnitudes from co-added images against a standard data base for examining the quality of photometry. However, we can undertake pairwise comparison of photometry from co-adds created by each method. For instance, we cross-match the photometric catalogues from MontageA and MontageB, and for each common source we calculate the difference in magnitudes obtained from both co-adds. We find that for these two methods, the difference is 0.02 mag with a standard deviation of 0.24 mag. Similar comparisons

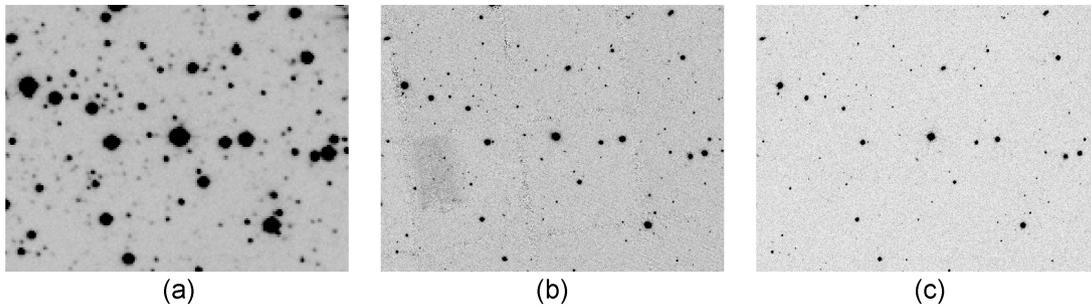


Figure 7. A small $\sim 3'.8 \times 3'.0$ postage stamp around $\alpha = 8^h 41^m 23^s$, $\delta = +17^\circ 55' 33''$ (130.34622° , 17.926073°) with intensities z -scaled to compare three different sky surveys. From left to right: (a) co-added image from CRTS, (b) corresponding cut-out from PS1 in the g band, and (c) SDSS survey image in the g band. The limiting magnitude comparisons were done on a larger area centred on these coordinates (Section 3.6).

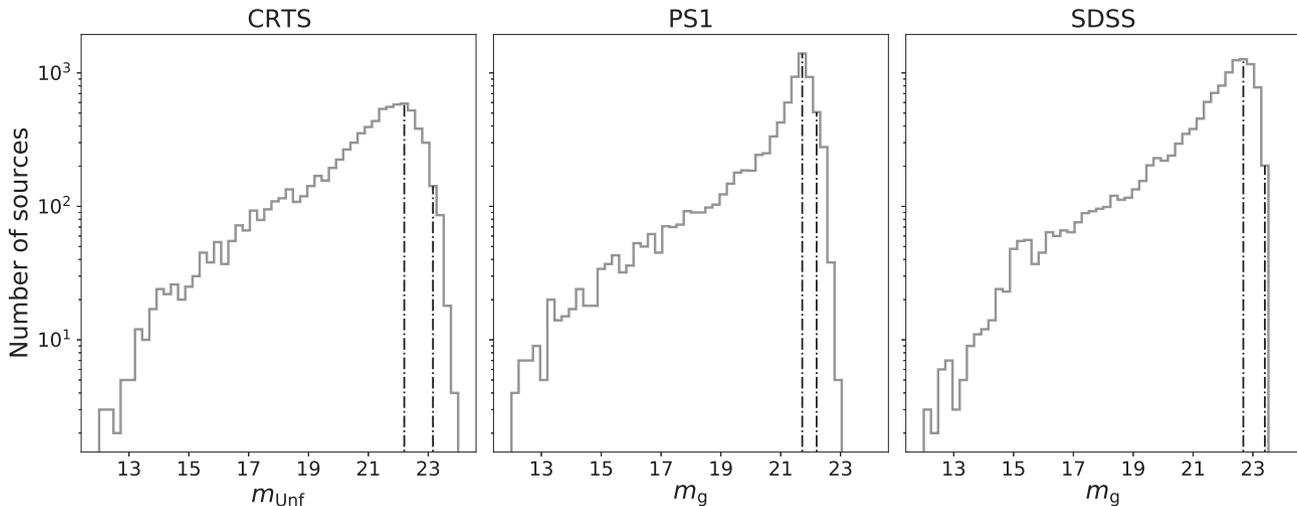


Figure 8. Histograms of magnitudes of sources detected in various surveys in a 30 arcsec region centred at $\alpha = 8^h 41^m 23^s$, $\delta = +17^\circ 55' 33''$ (Section 3.6). The vertical dash-dotted lines show the mode and the halfway point as discussed in Fig. 4. From left to right: unfiltered zero corrected CRTS, g -band sources in PanSTARRS, and g -band sources with $\Delta m_g < 0.2$ in SDSS. Numbers of sources within the histogram for CRTS, PS1, and SDSS are 7774, 8004, and 11 852, respectively.

for all methods are given in the top part of Table 6. We see that the scatter in magnitudes is higher for the brightest and faintest sources. As a result, we refine the comparison by plotting magnitudes measured from two methods against each other, and using a subset corresponding to a linear region of the plot. Values for these subsets are given in the lower part of Table 6.

Lastly, we compared the FWHMs of sources in the final co-adds by different methods as shown in Table 7. To eliminate spurious sources and a few genuine extended sources, we restricted this comparison only to sources with $\text{FWHM} < 20$ pixels. We find that SwarpA produces sharper images of stars as compared to the other three methods (Table 4). This difference is seen more clearly by pairwise comparisons of the FWHM in Table 8, analogous to the magnitude comparisons above. Fig. 5 shows histograms of the differences in FWHM for all objects, showcasing the sharper PSF obtained by SwarpA. However, it may seem contradictory with the fact that Table 4 shows a difference of 0.3 in average FWHM but considering that it shows average FWHM of all the sources including the artefacts and Table 6 represents the difference in FWHM among only the common source, the discrepancy vanishes.

In the end, we conclude that MontageA and MontageB produce slightly deeper co-adds, while SwarpA produces sharper co-adds. The key difference was that SWARP methods were modular and fast.

Although we have discussed the possibility of using multiple cores in parallel for single co-adds to reduce the time, it was found that it is far more efficient to co-add one field per core.

3.5 Comparison and method selection

Three of the four methods, which used direct averaging for creating the final co-add, give similar results. However, co-adds produced by the SwarpA clipped mean algorithm are significantly more robust to the presence of artefacts present in individual images (Fig. 3). Due to small dithers between exposures of the same field, the border regions of co-adds typically have fewer exposures. This reduces the effectiveness of the clipping procedure. However, background subtraction does a reasonable job of cleaning up these image areas.

Apart from cosmic rays, satellite trails, etc., an important artefact in images was spots due to dirt and dust. These regions consistently showed lower counts than neighbouring pixels (for example, see Fig. 3, bottom panel). These spots were present at the same pixel location over long periods. As the usual imaging procedure involved small dithers between exposures, we can address these problems by creating a bad pixel mask and excluding those from the final co-add. We created such a mask by co-adding 50 exposures of randomly selected fields. We converted this into a binary mask by applying

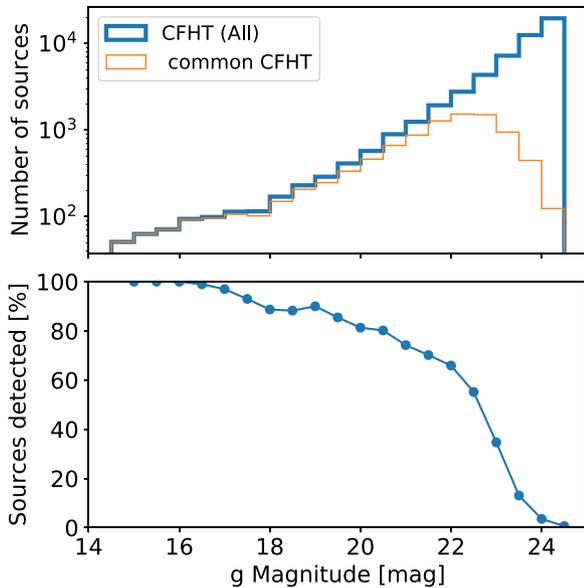


Figure 9. Top: Comparing the complete CFHT data set for the Cosmos field (blue line) with the subset of stars that were also found in our CRTS catalogue (orange line). The CFHT data are truncated at 24.5 mag, as there are no CRTS sources fainter than this value. Bottom: The fraction of CFHT sources that were detected in CRTS, calculated in 0.25 mag bins. We reach 50 per cent coverage at about 22.6 mag for this field, created by co-adding 382 images.

a simple sigma-clipping algorithm: All pixels in the mask with values 2σ below the mean were marked as ‘bad’ and assigned a Boolean value 0, while all other pixels were assigned the value 1. This was used as the weight mask in the final SwarpA co-addition process. Owing to some individual exposures having four extra pixel columns, the mask – a superset of all input images – was larger than typical images to be processed. This issue was tackled by trimming two columns from each side of the mask before using it during co-addition.

To conclude, the four methods tested were similar in many respects, but we found that robust artefact removal was crucial for producing reliable co-added images. This could have been achieved by source code modifications in MONTAGE, but was readily available in SWARP. After analysing various parameters in the test sets – e.g. time taken to co-add, modularity of the method, and quality of the co-adds – we decided to use SwarpA to co-add the whole set of images.

3.6 Comparing with other surveys

Once the preferred co-addition method is identified, we compared the CRTS co-added images with other all sky surveys,¹⁴ such as Pan-STARRS1 (PS1) and SDSS. We arbitrarily choose a region where all three surveys have an image with a reasonable overlap with each other. We arbitrarily picked the images to roughly centre around RA: 8h41m23s and Dec.: +17d54m21s. We added constant zero-point offset to the extracted sources in CRTS images as shown in Fig. 6, using cut-out of co-added images as in Fig. 7 (left to right) from CRTS, from the PS1 images, and SDSS survey image, respectively.

For limiting magnitude comparison, similar to what is being described in Section 3.5, we extracted the CRTS sources and their

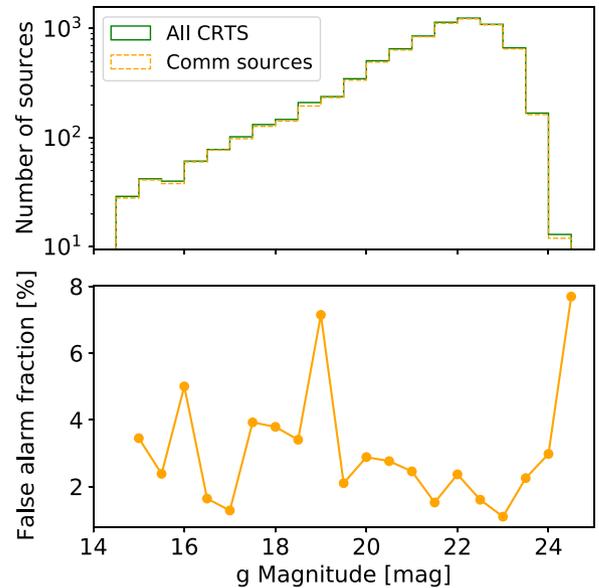


Figure 10. Evaluating the validity of the CRTS catalogue. Top: The green line shows a histogram of all sources found in the CRTS image, and the dashed orange line shows those CRTS sources that are also present in the CFHT catalogue. We see a good overlap between the two. Bottom: The fraction of sources detected in the CRTS image that were not present in the CFHT image, evaluated in 0.5 mag bins. We see that the false positive rate is at a few per cent level, independent of magnitude.

magnitudes using the SExtractor and for PS1¹⁵ and SDSS,¹⁶ catalogue cone searches were used, within 30 arcmin around the mentioned RA, Dec. in the *g* band. We see that CRTS is deeper than PS1 (Chambers et al. 2016), where the stacked images have a limiting magnitude of 23.3. Comparison with SDSS (Lundgren et al. 2015) is trickier: The limiting magnitude of SDSS DR16¹⁷ is 23.13, but the data base includes sources detected with very low significance. To approximate the 5σ sensitivity, we have created a histogram only of those sources with a *g*-band photometric error $\Delta m_g < 0.2$. In order to compare the magnitudes of all three surveys, we would need to compute the zero-point constant to the SExtractor MAG_ISO value computed for the CRTS sources, by cross-matching the common sources as shown in Fig. 6. Fig. 8 compares the histograms of limited magnitude for CRTS, PS1, and SDSS, respectively, of the sources within the region around the images in Fig. 7.

We note that CRTS images are unfiltered, with photometry corresponding loosely to Johnson’s *V* band – clearly distinct from *g*-band PS1/SDSS magnitudes for our comparisons. We take this step instead of calculating filter transformations to match the CRTS band, as these comparisons are only meant to be indicative of the overall depth and sensitivity of CRTS images.

From the number of sources and limiting magnitudes of images in this analysis, we can conclude that CRTS co-add images can reveal fainter sources than PS1 and SDSS in many parts of the sky. However, as Fig. 7 shows, the PSF of CRTS co-adds is a factor of 2 wider than the others: with the average FWHM being $3''.1$ (computed from SExtractor), as compared to $1''.39$ for PS1 (Magnier et al. 2016), and $1''.32$ for SDSS (Ross et al. 2011). This trade-off between deeper

¹⁵<https://catalogs.mast.stsci.edu/panstarrs/>

¹⁶<http://skyserver.sdss.org/dr16/en/home.aspx>

¹⁷https://www.sdss.org/dr16/imaging/other_info/

¹⁴ZTF data were not publicly available when this study was started.

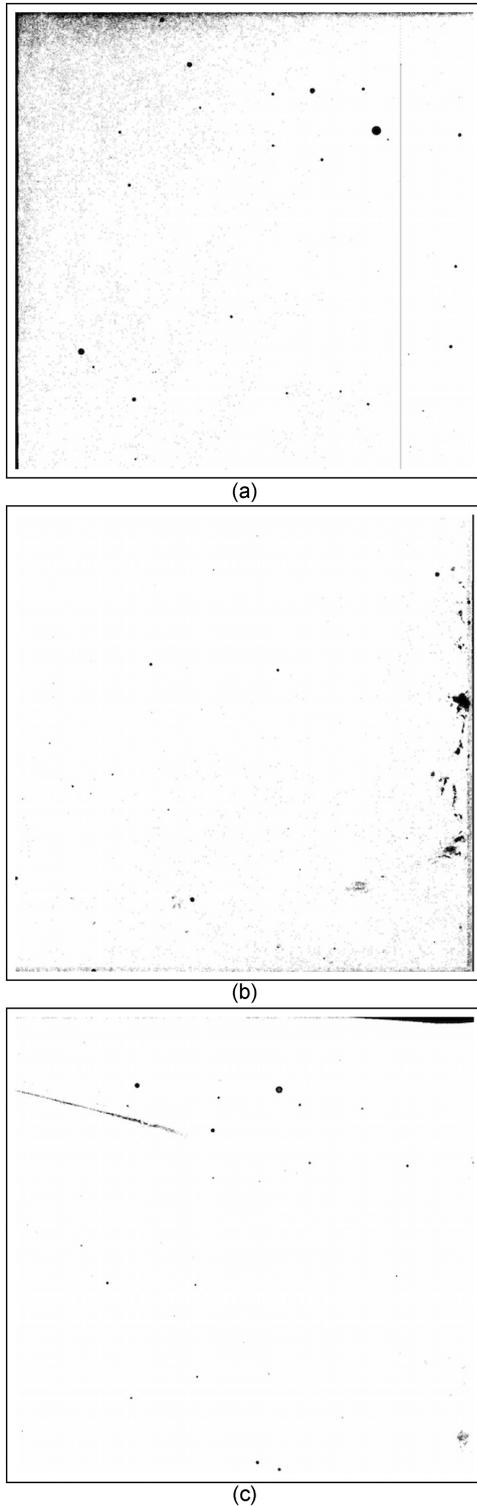


Figure 11. A combined image was produced by taking the median of 50 random images. This image was used as a mask, with zero-weights given to the spots during co-adding. The sections shown here are approximately a tenth of the original images. (Example of mask produced during the process – it is a binary image, with black pixels being masked away.)

and sharper images will have to be resolved by the end user based on their requirements in image subtractions and transient searches.

Due to the differences in aperture size, total exposure, filter used, etc., the total number of sources is often not the best metric

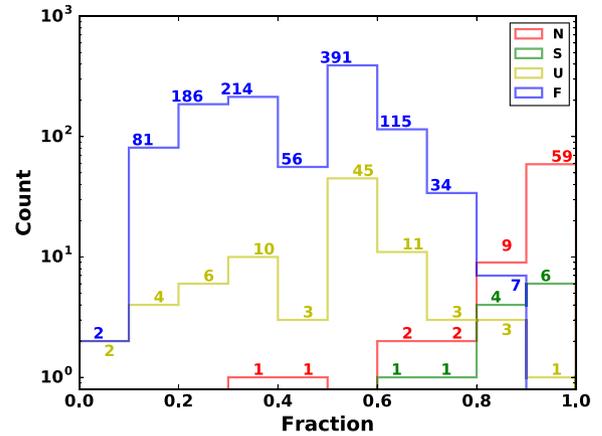


Figure 12. Of the total 7894 fields, 1260 fields were identified to have a high scatter in pointing directions across individual images. The co-adds for these fields were created by excluding the outlying fields from the final stack. The histograms here show the number of fields as a function of the fraction of total exposures retained in the final stack. For instance, 59 N fields with high scatter retained >90 per cent images in the final stack, 9 fields retained 80–90 per cent, and so on. *U* and *F* fields followed moving objects resulting in larger scatter and fewer of those fields got included in the final co-adds. Co-added images resulting from individual images with large angular separation result in large contiguous rectangular regions containing both the images. The process is computationally resource hungry with only an incremental advantage and hence such outliers were ignored.

when comparing a survey with a deeper survey. However, it is still important to do such a comparison if only to understand the fraction of sources that may be lost to noise, or blending, and to understand the usefulness of a survey for detecting, say, transients. We do so in the next section.

3.7 Comparison with a deeper catalogue

Another way to investigate the sensitivity of our co-added images and catalogue is to compare it with a significantly deeper catalogue. This allows unambiguous discrimination between true and spurious sources, as well as provides a high-quality photometric reference. For this purpose, we pick one of the deep fields in the CHFT Legacy Survey (Gwyn 2012). This survey includes ultra-deep images of the Cosmos field, obtained with the 1 sq. deg MegaCam. The catalogue created from these CFHT images reaches 50 per cent completeness at $m_g = 27.9$ (AB magnitude): significantly deeper than our typical images.

This CFHT ‘D2’ field, centred at $\alpha = 10^h00^m28^s$, $\delta = 02^\circ12'30''$, overlaps with the N01054 CRTS field that was created by co-adding 382 survey images. We pick a 20 arcmin \times 20 arcmin overlapping region from these fields for our analysis. To avoid problems with spurious sources near bright stars, we excluded a $3'.3$ region around every star brighter than $m_g = 14.5$. To get robust cross-matching, we picked a match radius of 4 arcsec: larger than diagonal of our $2''.5$ pixels. We follow a method similar to Section 3.6 and derive a zero-point offset of 27.3 mag.

Since the CRTS images are much shallower than the CFHT catalogue, we assume that any source present in CFHT is a true source. First, we find the number CFHT sources that are present in our CRTS catalogue. The top panel of Fig. 9 compares the total number of sources found in the CFHT catalogue (blue line) with the subset of these sources that was also detected in our catalogue (orange line). We

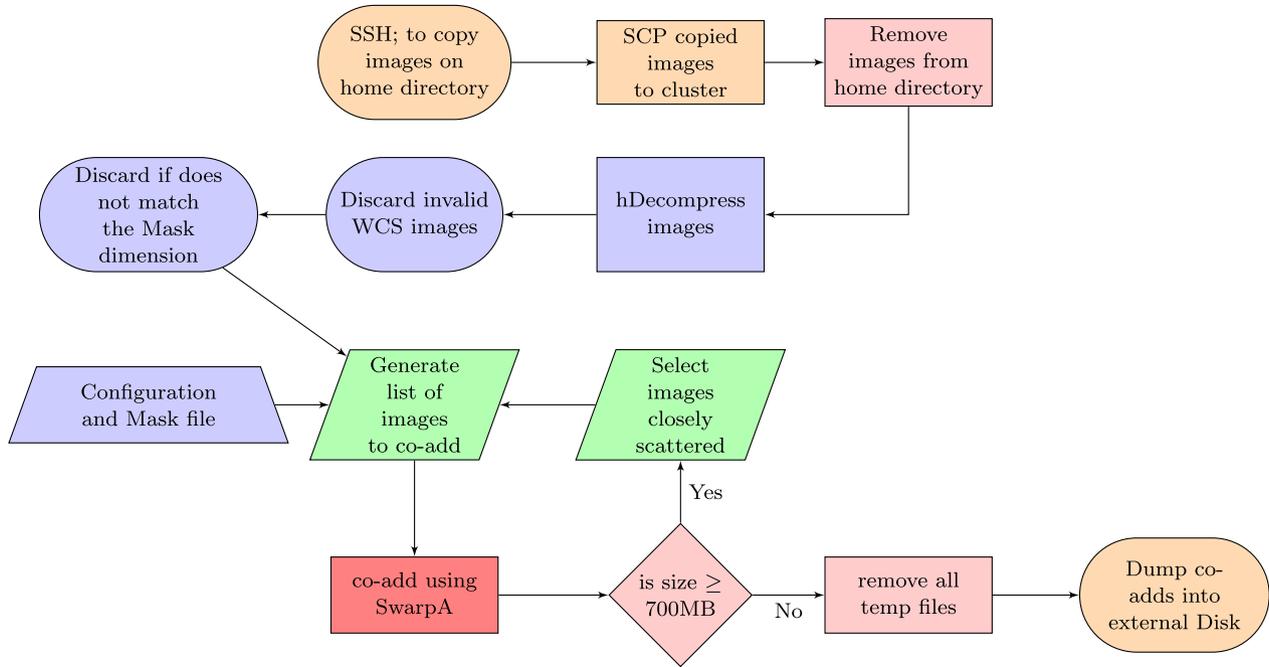


Figure 13. Flow chart of the work flow used to execute the co-addition. Colours of the nodes represent the following. Orange represents steps involving remote computers, i.e. not part of the cluster where co-addition takes place. Pink represents steps carried out only due to memory constraints; hence pink and orange processes are not mandatory and can be avoided if performing in a local directory. Violet and green are mandatory steps, with the difference being green steps may require correction. Lastly, red presents the process where core algorithm takes place.

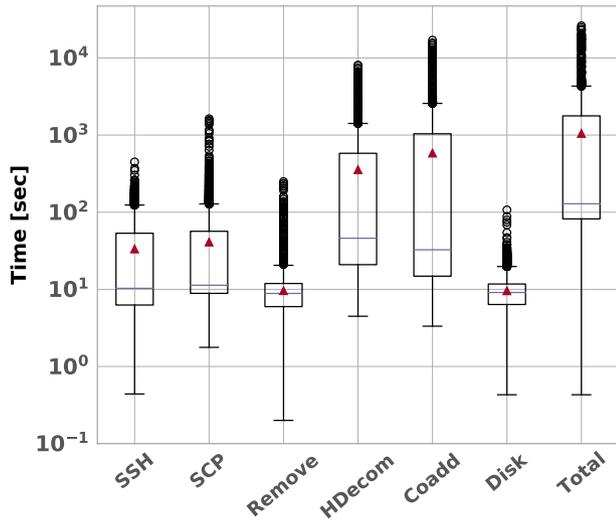


Figure 14. Box plot of time taken (in seconds) by each of the co-addition stages for different fields. In some cases, the contributors to the dispersion are mainly the number of images in a field (and the resulting sizes, operations, etc.), but another factor is also whatever else the machine may be handling at the time. For each column, the red line and the red square represent median and mean values, respectively. The height of the box indicates the interquartile range (first to third quartile), and the outliers are shown by the ‘o’ markers. In general, the data decompression took the same amount of time as the co-addition itself, and typically over an order of magnitude more than the other tasks.

see that the catalogue is nearly complete till ~ 22 mag. This completeness is better demonstrated in the lower panel, which shows the percentage of CFHT sources recovered in CRTS images. Our image of this field has 50 per cent completeness at about 22.6 mag. We note that

this depth varies over the sky based on the number of images that were co-added.

Having checked the completeness of the catalogue, we now turn our attention to validity. The upper panel of Fig. 10 compares the histogram of all sources in the CRTS catalogue (blue line) with those that are also present in the CHFT catalogue (orange line) – and the curves nearly overlap. The lower panel quantifies any discrepancies in terms of a false alarm fraction: We see that this fraction remains at a few per cent level across all magnitudes. A plausible explanation of poor coverage at some magnitude might be poor PSF, which blends in the nearby objects to the SEXTRACTOR algorithm.

4 IMPLEMENTATION

We now discuss the co-addition process in detail, including the workflow (Section 4.1), hardware used for processing (Section 4.2), processing time (Section 4.3), and accessing the final products (Section 4.4).

4.1 Workflow

The set-up at IUCAA is such that images are stored at one place, and were to be processed at another. Hence, the first step was to copy all images to a staging area on the source machine, scp them to the co-add server into a folder for a given field, and then clear the staging area. We H-decompressed the images, verified the presence of a valid WCS in them, and ensured pixel dimensions as indicated in the header were equal to that of the image mask. Fig. 11 shows portion of the image of bad pixel mask which plays an important role in removing pixel based artefacts hence, it is important for the dimension of the mask matches that of the image as well. In the few cases where these conditions were not met, the images were excluded from the final co-add. Another reason to exclude an image from final

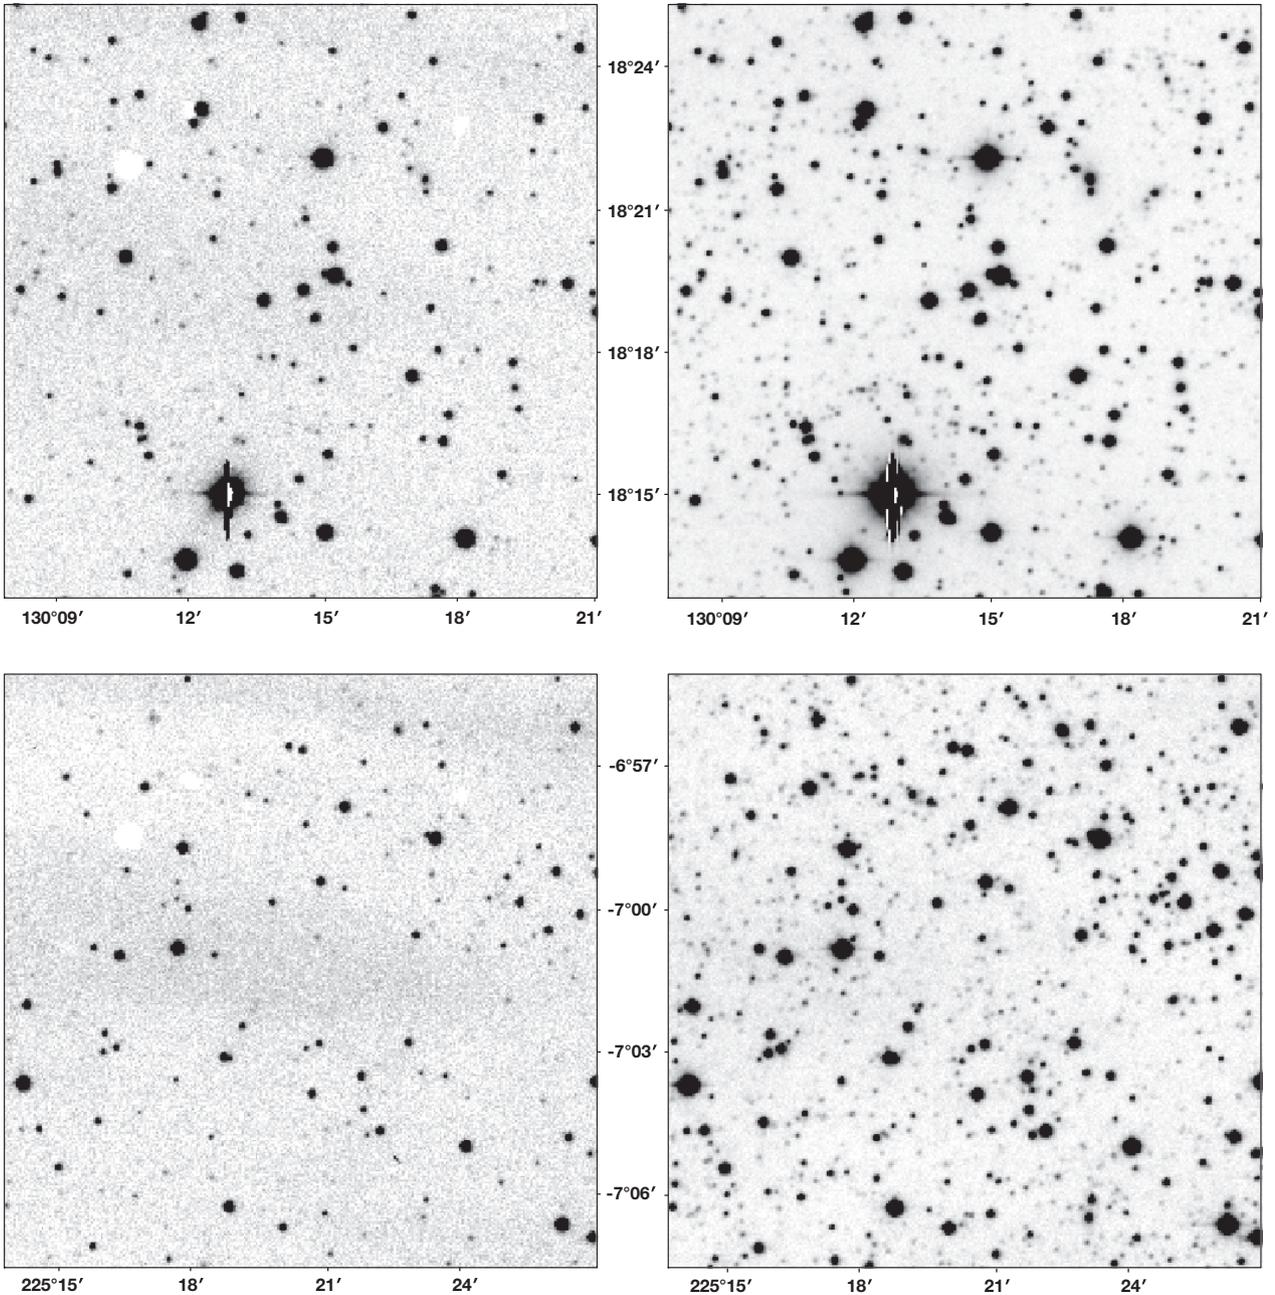


Figure 15. Top row: A section of field N18045 ~ 12 arcmin on a side around RA, Dec. 08h40m57.9354s +18d19m02.4383s. Top left: You see a single-epoch image, and top right: you see the same section for the co-added image. This field had only 439 images, out of which 497 have gone in to the co-added image. Bottom row: A section of field S07080 ~ 12 arcmin on a side around RA, Dec. 15h01m21.1668s -07d01m17.2091s. Bottom left: You see a single-epoch image, and bottom right: you see the same section for the co-added image. This field had only 268 images, out of which 299 have gone in to the co-added image. Note that defects in the single image are not seen in the co-added image.

co-add was if an individual images has a large angular separation (scattered) with respect to the image at the median position. Not excluding these images puts a heavy burden on the computational resources with little or no incremental advantage. Fig. 12 shows a histogram of fraction of images rejected for different fields. A text file containing names of all valid images was generated for each field, and SWARP configuration file was also generated. Most parameters were left at default values, with notable exceptions listed in Table 9. The master script then calls SWARP for co-addition.

SWARP then reads all input images along with their weight maps. In our case, these weight maps are simply binary masks intended to eliminate known CCD artefacts. SWARP uses this information to build a background map and subtract it from the images. All input images are then re-sampled and projected on to an appropriate subsection of the output frame rejecting the masked pixels. The images were co-added with sigma-clipping mean to reject outliers. Since all input images had the same weight, no weight files were used for this step. However, if hierarchical co-addition is performed

(for instance, Section 3.1), then appropriate weight files generated in each intermediate step should be supplied to the consecutive steps.

The master script deletes temporary files generated by SWARP, retaining only the final co-add and the corresponding weight file. The master script logs all operations for debugging and statistical purposes. It also adds appropriate FITS header keywords to the final co-added image, describing the processing involved (Table 10). The work flow is explained as flow chart in Fig. 13.

Typical output images are 70–100 megabytes in size, as the area varies a little. In some cases, we found that the final image was extremely large – this happened when one of the input images had a large on-sky offset from the typical images in the input stack. Such an offset may arise either from a wrong (but semantically valid) WCS in the image or by incorrect labelling of the field ID for an input image. As a simple cut for identifying such images, we flagged cases where the co-add was larger than 150 MB. This was noticed for 450 fields out of 7894 (5.7 per cent). Over time, we will remove them using visual inspection, a time-consuming task deferred for the future. Another future task is to exclude from co-addition images with poor seeing and/or high background sky. This is expected to marginally reduce the number of blends seen in the current co-adds.

4.2 Hardware/computing infrastructure

The CSS images are processed into catalogues by Lunar and Planetary Laboratory (LPL) in Tucson, Arizona to look for NEOs and then passed on to Caltech to look for transients. At Caltech, we obtain 5 arcmin \times 5 arcmin cut-outs centred around the found transients. The full images of the CSS are transferred to the Inter-University Center for Astronomy and Astrophysics (IUCAA), in Pune, India and served in the form of 5 arcmin \times 5 arcmin cut-outs¹⁸ through a web interface that uses a PYTHON/Flask backend. The web application serving these cut-outs is hosted on an HP Blade server with dedicated storage class server for the imaging data and a dedicated server for the photometric catalogue and the imaging metadata. Originally, we planned to use these cut-outs for co-adding, but the small dithering between revisits meant that each cut-out would have non-uniform edges and varying SNR. Initial experiments also revealed that process would take up to 50 times longer as a result of the larger number of images.

We used ‘Perseus’, a 64-node cluster at the High Performance Computing Center at IUCAA, for our stacking experiments. Each node of Perseus has two CPUs with eight cores each, each node having 128 GB RAM. Ideally, it should have been possible to execute all co-addition jobs in parallel as the number of cores, 1024, is of the same order as the number of fields, viz. 7894. However, moving the images to Perseus in real time turned out to be a bottleneck. We discovered that at most 30 processes could be run concurrently. A dedicated storage could have been attached to Perseus but it was not possible at the time of this work as the cluster was also being shared by other processing jobs. In the end, we typically used 20–25 cores at a time for the co-additions. The total time required was $\mathcal{O}(2$ weeks).

4.3 Performance/time taken

We now discuss the time taken by different processes. The times taken to SSH, SCP the images and transferring the co-adds to the disc were the fastest compared to H-decompress and co-addition processes. Fig. 14 shows a box plot of time taken per co-add by the

individual processes and total time from remotely accessing images to transferring co-adds to an external hard disc.

Once we settled on the method and the set-up and debugging was done, the total time for the co-adds was 10 clock-days. Of the 7894 fields we started with, 7791 fields were successfully co-added.

4.4 Final stack

The total size of the co-added images is 1.5 TB, and the images cover an area of ~ 27 000 sq. deg. Users can access the images from <http://crts.iucaa.in>. Users who find these data useful in their research are requested to acknowledge it as ‘*This research makes use of data from the Catalina Sky Surveys, LPL, UA and Catalina Realtime Transient Survey, Caltech, with value additions and served by IUCAA through a web application currently hosted at http://crts.iucaa.in*’. Fig. 15 shows the comparison of single image versus stacked methodology as discussed.

5 CONCLUDING COMMENTS

We have carried out the co-addition of CSS images to produce very deep co-adds with up to ~ 3 mag more depth than a single image. CRTS also includes the Mount Lemmon Survey (MLS) and Siding Spring Survey surveys. We plan to co-add these images as well to provide coverage in the South, and greater depth. CSS and MLS cameras are currently using larger, 10k \times 10k chips. In the near future, we will be co-adding images from these cameras as well.

A comparison with additional combining software will provide an insight about what should be used to continually improve reference images over large areas in the era of ZTF (Bellm E. and LSST (LSST Science Collaboration 2009).

The cut-outs will be served along with the images for individual epochs already being served from the CRTS public server at IUCAA.¹⁹ The set comprising the deep image in its entirety is 1.5 TB.

In order to realize the full potential of the stacked images, a catalogue needs to be obtained. Achieving this, however, is non-trivial since the depth of the co-adds varies from field to field. Until such a catalogue is produced, the users will have to rely on local calibration or use the deep images for detection purposes.

ACKNOWLEDGEMENTS

We thank, Bruce Berriman, John Good, Frank Masci, Steve Hartung, and Simon Krugoff for their guidance. We also thank CSS, CRTS, IUCAA, Indo-US Science and Technology Forum (IUSSTF), and National Science Foundation (NSF) for the excellent support.

The Catalina Sky Survey is funded by a grant from the National Aeronautics and Space Administration’s Near-Earth Object Observations programme. The CRTS web service at IUCAA was partly developed under the *Virtual Observatory* program funded by the Ministry of Communication and Information Technology of the Government of India and the *Data Driven Initiatives* program funded by the National Knowledge Network.

This work is based in part on data products produced at Terapix available at the Canadian Astronomy Data Centre as part of the Canada–France–Hawaii Telescope Legacy Survey, a collaborative project of NRC and CNRS.

Ashish A. Mahabal (AAM), Andrew J. Drake, Ciro Donalek, Matthew J. Graham (MJG), and S. George Djorgovski (SGD) were

¹⁸<http://crts.iucaa.in>

¹⁹<http://crts.iucaa.in>

supported in part by the NSF grants AST-0909182, AST-1313422, AST-1413600, and AST-1518308, and the Ajax Foundation. AAM, MJG, and SGD were also supported by IUSSTF. We also thank CLASSACT: Indo-US centre for astronomical object and feature characterization and classification.

This research made use of MONTAGE. It is funded by the National Science Foundation under Grant Number ACI-1440620, and was previously funded by the National Aeronautics and Space Administration's (NASA) Earth Science Technology Office, Computation Technologies Project, under Cooperative Agreement Number NCC5-626 between NASA and the California Institute of Technology. This work was done during the tenure of Akshat Singhal and Varun Bhalerao in Inter-University Centre for Astronomy and Astrophysics.

DATA AVAILABILITY

All the final co-add images and weight files are available at <http://crts.iucaa.in>. Additional material such as configuration files, scripts for processing bulk images, log files, etc. is available at <https://ddi.iucaa.in/static/crts/>. The raw data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Albareti F. D. (SDSS Collaboration), et al. (SDSS Collaboration), 2017 *ApJS*, 233, 25
- Aldering G. et al., 2002, in Tyson J. A., Wolff S., eds, *Proc. SPIE Conf. Ser. Vol. 4836, Survey and Other Telescope Technologies and Discoveries*. SPIE, Bellingham, p. 61
- Annis J. et al., 2014, *ApJ*, 794, 120
- Bellm J., 2014, in Wozniak P. R., Graham M. J., Mahabal A. A., Seaman R., eds, *The Zwicky Transient Facility, Proceedings of the Third Hot-wiring the Transient Universe Workshop*, Santa Fe, New Mexico, USA, p. 27
- Berriman G. B., Good J. C., Curkendall D. W., Jacob J. C., Katz D. S., Prince T. A., Williams R., 2003, in Payne H. E., Jędrzejewski R. I., Hook R. N., eds, *ASP Conf. Ser. Vol. 295, Astronomical Data Analysis Software and Systems XII*. Astron. Soc. Pac., San Francisco, p. 343
- Bertin E., 2010, *Astrophysics Source Code Library*, record ascl:1010.068
- Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
- Cao Y., Nugent P. E., Kasliwal M. M., 2016, *PASP*, 128, 114502
- Chambers K. C. et al., 2016, preprint ([arXiv:1612.05560](https://arxiv.org/abs/1612.05560))
- Desai S., Mohr J. J., Henderson R., Kimmel M., Paech K., Wetzstein M., 2015, *J. Instrum.*, 10, C06014
- Dey A. et al., 2019, *AJ*, 157, 168
- Djorgovski S. G. et al., 2008, *Astron. Nachr.*, 329, 263
- Djorgovski S. G. et al., 2011, preprint ([arXiv:1102.5004](https://arxiv.org/abs/1102.5004))
- Djorgovski S. G., Mahabal A., Drake A., Graham M., Donalek C., 2013, in Oswalt T. D., Bond H. E., eds, *Sky Surveys, Planets, Stars and Stellar Systems, Vol. 2*, Springer Science+Business Media, Dordrecht, p. 223
- Drake A. J. et al., 2009, *ApJ*, 696, 870
- Flaugher B. et al., 2015, *AJ*, 150, 150
- Flewelling H. A. et al., 2016, *ApJS*, 251, 7
- Fruchter A. S., Hook R. N., 2002, *PASP*, 114, 144
- Goldstein D. A. et al., 2019, *ApJ*, 881, L7
- Gruen D., Seitz S., Bernstein G. M., 2014, *PASP*, 126, 158
- Gwyn S. D. J., 2012, *AJ*, 143, 38
- Hartung S., Shukla H., 2013, *American Astronomical Society Meeting Abstracts #221, #240.06*
- Hartung S., Shukla H., Miller J. P., Pennypacker C., 2012, preprint ([arXiv:1209.5823](https://arxiv.org/abs/1209.5823))
- Ivezic Z. et al., 2008, *ApJ*, 673, 111
- Lang D., 2014, *AJ*, 147, 108
- Larson S. M., Hergenrother C., Whiteley R., Kelly C., Hill R., McNaught R., 2001, in Isobe S., Asakuro Y., eds, *Upgrading the Catalina Sky Survey and Southern Survey*, Int. Workshop Collaboration Coord. NEO Obs. Orbital Comput, Okayama, Japan, p. 13
- Larson S., Brownlee J., Hergenrother C., Spahr T., 1998, *Bull. Am. Astron. Soc.*, 30, 1037
- Law N. M. et al., 2009, *PASP*, 121, 1395
- LSST Science Collaboration et al. 2009, preprint ([arXiv:0912.0201](https://arxiv.org/abs/0912.0201))
- Lundgren B. et al., 2015, *PASP*, 127, 776
- Magnier E. A. et al., 2016, *ApJS*, 251, 5
- Mahabal A. A. et al., 2011, *Bull. Astron. Soc. India*, 39, 387
- Masci F., 2013, preprint ([arXiv:1301.2718](https://arxiv.org/abs/1301.2718))
- Masci F., 2013b, *Astrophysics Source Code Library*, record ascl:1302.010
- Meisner A. M., Bromley B. C., Nugent P. E., Schlegel D. J., Kenyon S. J., Schlafly E. F., Dawson K. S., 2017b, *AJ*, 153, 65
- Meisner A. M., Lang D., Schlegel D. J., 2017a, *AJ*, 153, 38
- Pravdo S. H. et al., 1999, *AJ*, 117, 1616
- Rau A. et al., 2009, *PASP*, 121, 1334
- Ross A. J. et al., 2011, *MNRAS*, 417, 1350
- SDSS Collaboration, 2017, *ApJS*, 233, 25
- White R. L., Percival J. W., 1994, in Stepp L. M., ed., *Proc. SPIE Conf. Ser. Vol. 2199, Advanced Technology Optical Telescopes V*. SPIE, Bellingham, p. 703
- Zackay B., Ofek E. O., 2017a, *ApJ*, 836, 187
- Zackay B., Ofek E. O., 2017b, *ApJ*, 836, 188

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