



# Sub-band Image Reconstruction Using Differential Chromatic Refraction

Matthias A. Lee<sup>1</sup> , Tamás Budavári<sup>1,2,3</sup>, Ian S. Sullivan<sup>4</sup>, and Andrew J. Connolly<sup>4</sup>

<sup>1</sup>Department of Computer Science, The Johns Hopkins University, Baltimore, MD, USA

<sup>2</sup>Department of Applied Mathematics and Statistics, The Johns Hopkins University, USA

<sup>3</sup>Department of Physics and Astronomy, The Johns Hopkins University, USA

<sup>4</sup>Department of Astronomy, University of Washington, Seattle, WA, USA

Received 2018 June 20; revised 2019 March 24; accepted 2019 March 25; published 2019 April 22

## Abstract

Refraction by the atmosphere causes the positions of sources to depend on the airmass through which an observation was taken. This shift is dependent on the underlying spectral energy of the source and the filter or bandpass through which it is observed. Wavelength-dependent refraction within a single passband is often referred to as differential chromatic refraction (DCR). With a new generation of astronomical surveys undertaking repeated observations of the same part of the sky over a range of different airmasses and parallactic angles, DCR should be a detectable and measurable astrometric signal. In this paper we introduce a novel procedure that takes this astrometric signal and uses it to infer the underlying spectral energy distribution of a source; we solve for multiple latent images at specific wavelengths via a generalized deconvolution procedure built on robust statistics. We demonstrate the utility of such an approach for estimating a partially deconvolved image, at higher spectral resolution than the input images, for surveys such as the Large Synoptic Survey Telescope.

**Key words:** astrometry – catalogs – galaxies: statistics – methods: statistical – surveys

## 1. Introduction

Modern surveys observe the same region of the sky hundreds or thousands of times over the course of their lifetime. Due to changes in the atmosphere, airmass, and parallactic angle of these observations the image quality and effective throughput of these exposures can vary with time. Time-domain and multicolor studies need to consider these images jointly and, therefore, require advanced methods to combine the data to maximize their information content. Image stacking is one of these techniques and is traditionally accomplished by either coadding images or selecting a small subset of the highest quality data (often referred to as lucky imaging). Coadding can include convolving images to a common resolution before addition (Lucy & Hook 1992), combining images within a limited range of image qualities (Annis et al. 2014), stacking all images while weighting by signal-to-noise (Szalay et al. 1999) or by using a matched-filter approach (Zackay & Ofek 2017). Lucky imaging, in contrast, selects the exposures with the highest image quality but at the cost of a reduction in signal-to-noise, and hence depth of the observations. In each of these approaches we sacrifice information (e.g., spatial resolution or depth) for computational expediency or to optimize some aspect about the properties of the images (e.g., image quality).

In Lee et al. (2017), we introduced the concept of learning an underlying model that represents the sky from a sequence of images with differing image qualities and depths. In this approach, an image  $y_t$  at epoch  $t$  is a simple convolution of the point-spread function (PSF)  $f_t$  and a latent model image  $x$  plus some (uncorrelated) normal noise  $\epsilon_t$ ,

$$y_t = f_t * x + \epsilon_t \quad \text{with} \quad \epsilon_{t,i} \sim \mathcal{N}(0, \sigma_{t,i}^2), \quad (1)$$

where  $\sigma_{t,i}$  is the estimated standard deviation of pixel  $i$ , i.e.,  $\sigma_t^2$  is the variance map of the  $y_t$  image. Since the convolution is a linear operation, in the simplified limit of one-dimensional vectors, the estimated  $\hat{y}_t$  image can be written as a matrix

multiplication

$$\hat{y}_t = f_t * \hat{x} = F_t \hat{x}, \quad (2)$$

where  $\hat{x}$  is the estimated latent image, without loss of generality. Given a set of  $\{y_t\}$  exposures and their known  $\{f_t\}$  PSFs we have shown that we can solve for the model image  $x$  in an incremental way by iterating over the observations, considering them one by one, and updating the latent model using the data from the current epoch,  $t$ . Omitting the  $t$  index for clarity, the multiplicative update formula is

$$\hat{x}^{(\text{new})} = \hat{x} \odot \frac{F^T W y}{F^T W \hat{y}}, \quad (3)$$

where the dividing bar and the  $\odot$  sign are elementwise division and multiplication operations, and the  $\hat{y} = F \hat{x}$  is the predicted image (Lee & Budavári 2015; Lee et al. 2017). The  $W$  matrix is a combination of the variance map and the binary masks for the exposure including censored areas, e.g., saturated pixels or bad camera columns, and robust weights that arise as a modified minimization of the dispersion of the  $r_t = y_t - \hat{y}_t$  residuals in the form of a function  $\rho(t)$  with

$$\min_{\hat{x}} \sum_{t,i} \rho\left(\frac{r_{t,i}}{\sigma_{t,i}}\right) \quad (4)$$

rather than the traditional

$$\min_{\hat{x}} \frac{1}{2} \sum_{t,i} \left(\frac{r_{t,i}}{\sigma_{t,i}}\right)^2. \quad (5)$$

In the classical limit of  $\rho(t) = t^2/2$ , these two optimization problems are identical, however, for other (more) robust  $\rho$ -functions, they significantly differ. The typical choice for robust optimization is for  $\rho$  to be quadratic for small residuals, but in case of large deviations  $\rho$  becomes a constant value, say 1. In this way the contribution of outliers to the cost function is

limited (Maronna et al. 2006). Such robust minimization essentially corresponds to a maximum likelihood estimation using a likelihood function with a longer tail than the Gaussian to accommodate outliers. Our standard solution is an iterative process where we re-weight the quadratic terms with a weight derived from the robust  $\rho$ -function,  $\mathbb{W}(t) = \rho'(t)/t$ . In the classical case of  $\rho(t) = t^2/2$ , the weight is a constant of 1, but for cases with significant outliers we use a Cauchy distribution with  $\mathbb{W}(t) = 1/(1 + t^2/c^2)$  to down weight these outliers; see Maronna et al. (2006) and Lee et al. (2017).

In this paper, we expand upon these ideas to learn not just the latent image but also to infer the spectral properties of sources from astrometric shifts introduced through differential chromatic refraction (DCR). In Section 2 we describe our new method capable of extracting color information from sequences of monochromatic images. We apply this technique to simulated images in Section 3 where we study the limits of the approach in realistic settings. The results and potential applications of this approach are discussed in Section 4.

## 2. Differential Chromatic Refraction

Astronomical sources emit light across the continuum of the electromagnetic spectrum,  $S(\lambda)$ , with varying intensity. Observed broadband magnitudes or fluxes  $Y$  are an integral over the filter and instrument characteristics  $r(\lambda)$  through which the source is observed (including the detector's quantum efficiency, camera rotator, etc.).

Considering that the CCD counts photons, the equation is given by

$$Y = \frac{\int d\lambda \lambda r(\lambda) S(\lambda)}{c \int d\lambda r(\lambda)/\lambda}, \quad (6)$$

where  $c$  is the speed of light. By substituting for a specific filter

$$R(\lambda) = \frac{\lambda r(\lambda)}{c \int d\lambda r(\lambda)/\lambda}, \quad (7)$$


the expression for the flux is significantly simplified,

$$Y = \int d\lambda R(\lambda) S(\lambda). \quad (8)$$

The above equation applies to each pixel in a  $y_t$  exposure. The latent image is, however, complicated by the convolution from the PSF and the refraction of the atmosphere. DCR is due to the refraction of light as it passes through the atmosphere similar to light passing through a prism. The differential component of DCR refers to the change in the PSF across the bandpass of a given filter. It results in a positional shift and distortion of a source that depends on the source's spectral energy distribution (Meyers & Burchat 2015). This shift and distortion is a function of wavelength,  $f_t(\lambda)$ , and the altitude and azimuth of the telescope for each epoch,  $t$ , of the observation. In image coordinates, the effect will also depend on the orientation of the camera rotator. The magnitude of the refraction depends on the temperature and humidity of the atmosphere Filippenko (1982).

Instead of a latent image  $x$ , we introduce the density image  $\xi(\lambda)$ , and formulate a more general model as

$$y_t = \int d\lambda R(\lambda) [f_t(\lambda) * \xi(\lambda)] + \epsilon_t, \quad (9)$$



**Figure 1.** The DCR effect is wavelength-dependent with lower wavelengths exhibiting a significantly larger astrometric offset. In this figure, we show the astrometric offset for a monochromatic source if we change the wavelength of the source by 10 nm (assuming a zenith angle of  $50^\circ$ ) on the center location of a PSF. This offset is measured in units of LSST pixels (0.2 arcsec). For reference we also show the LSST bandpass throughputs.

where again the  $*$  sign denotes the spatial convolution. In the limit of negligible dependence of the PSF on the wavelength, we get back the previous model

$$y_t = f_t * \left[ \int d\lambda R(\lambda) \xi(\lambda) \right] + \epsilon_t = f_t * x + \epsilon_t. \quad (10)$$

In Figure 1 we illustrate the strength of the astrometric shift as a function of wavelength. The black solid line plotted over the throughput curves shows the relative shift in pixels (0.2 arcsec) for a 10 nm difference in observed wavelength. For example the differential shift at a zenith angle of  $50^\circ$  across the extremes of the  $u$  band,  $\lambda_1 = 320$  nm and  $\lambda_2 = 400$  nm, is over 6 pixels, similarly over the  $g$  band,  $\lambda_1 = 400$  nm and  $\lambda_2 = 560$  nm, it is approximately 5.5 pixels. Of course in reality, the total apparent shift is dependent on the spectral energy distribution (SED) of the source (Kaczmarszik et al. 2009; Peters et al. 2015). The elongation of the PSF arises because the PSF is simply the weighted summation of all monochromatic PSFs between  $\lambda_1$  and  $\lambda_2$ .

To account for the wavelength dependence of DCR we, therefore, consider that a latent image is comprised of a number of sub-band images each comprising a limited wavelength interval within a given photometric passband. Given the spectral dependence of the PSF, if we can measure an astrometric offset or variation in the PSF as a function of airmass or parallactic angle (telescope pointing), we can in principle infer the underlying SED. Within these narrower spectral ranges, we assume a constant PSF and SED (though the technique can account for wavelength-dependent PSFs even within these sub-bands). We split the filter's wavelength range  $[\Lambda_0, \Lambda_K]$  at preset  $\{\Lambda_k : k = 1 \dots K - 1\}$  values, and define the response function of these sub-bands as

$$R_k(\lambda) = \begin{cases} R(\lambda) & \text{if } \lambda \in [\Lambda_{k-1}, \Lambda_k] \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

for all  $k$  from 1 to  $K$ . Naturally,

$$y_t = \sum_{k=1}^K \int d\lambda R_k(\lambda) [f_t(\lambda) * \xi(\lambda)] + \epsilon_t \quad (12)$$

is equivalent to Equation (9), and its terms can be approximated similarly to that in Equation (10), which yields

$$\hat{y}_t = \sum_k f_{t,k} * \hat{x}_k = \sum_k F_{t,k} \hat{x}_k, \quad (13)$$

where each  $x_k$  is a latent image in the given wavelength range of sub-band  $k$ , and  $F_{t,k}$  is the linear operator that corresponds to the convolution with the PSF  $f_{t,k}$  at the appropriate wavelengths.

Introducing a tall  $\hat{x}$  vector that contains all these  $\hat{x}_k$  (one-dimensional) images

$$\hat{x} = \begin{bmatrix} \hat{x}_1 \\ \hat{x}_2 \\ \vdots \end{bmatrix} \text{ and } F_t = [F_{t,1} \ F_{t,2} \ \dots] \quad (14)$$

as the horizontal concatenation of the convolution matrices, we can verify that the equation

$$y_t = F_t \hat{x} + \epsilon_t \quad (15)$$

is equivalent to Equation (13). Considering that Equation (15) is formally the same as the previous model, see Equations (1) and (2), the iterative updates of Equation (3) directly apply to our new wavelength-dependent model.

Furthermore, considering that the transpose of  $F$  (omitting the  $t$  index) is

$$F^T = \begin{bmatrix} F_1^T \\ F_2^T \\ \vdots \end{bmatrix}, \quad (16)$$

we see that the updates for the individual latent images at different  $\lambda_k$  wavelengths take the familiar forms of

$$\begin{aligned} \hat{x}_1^{(\text{new})} &= \hat{x}_1 \odot \frac{F_1^T W y}{F_1^T W \hat{y}}, \\ \hat{x}_2^{(\text{new})} &= \hat{x}_2 \odot \frac{F_2^T W y}{F_2^T W \hat{y}}, \\ &\vdots \end{aligned} \quad (17)$$

and so on. The important difference is that now the common  $\hat{y}$  term, the predicted image, is the sum of the constituents wavelength-dependent image,

$$\hat{y} = \sum_k F_k \hat{x}_k. \quad (18)$$

Note that in the limit of a single latent image, we get back the original equations and the corresponding iterative algorithm, hence this is a more general approach. With this formalism we can use our robust iterative deconvolution method with only minor modifications. This iterative method is comprised of the following steps. We begin by initializing our models,  $[\hat{x}_1 \dots \hat{x}_k]$ , to be uniform in value such that the sum of all pixels is 1.0, then start by selecting a random observation,  $y_t$ , for which we estimate the appropriate PSF, which we then use in conjunction with Equation (17) to update our models. We repeat this

procedure until we have reached convergence. Note that we reuse observations, once we have exhausted our unique supply.

### 3. Recovery from Synthetic Exposures

To study the properties and limitations of this method we analyze synthetic images with only two components. This essentially corresponds to objects with spectra that only contain two discrete wavelengths. In other words, we attempt to solve for two images at preset wavelengths from a set of exposures that combine these wavelengths into single passband observations.

#### 3.1. Simulated Images

Our synthetic exposures are created by observing a latent image through a PSF generated at one of two preset wavelengths. Our latent images are a simple grid of point sources where the intensity of the individual sub-bands as well as the total flux of the source vary across the image. This simulates sources with varying brightness or color. We will refer to the sub-band images as  $\hat{x}_1$  and  $\hat{x}_2$ , corresponding to  $\lambda_1$  and  $\lambda_2$ , respectively. While our simulations are of point sources, we note that this technique works equally well for resolved and even blended sources.


In order to choose realistic wavelengths, we take the Large Synoptic Survey Telescope (LSST)  $g$  band and select the effective wavelengths of the sub-bands produced by halving the  $g$  band into two equal passbands between the wavelengths,

$$\lambda_1 = 443.7 \text{ nm and } \lambda_2 = 513.5 \text{ nm.} \quad (19)$$

To generate a synthetic observation, we choose a zenith and azimuth angle, then generate a PSF for each wavelength at those angles and convolve each latent image with their respective PSFs. We use the GALSIM library (Rowe et al. 2015) to generate Kolmogorov PSFs with a FWHM of 1.0 arcsec, the same settings used by the StarFast Simulator (Sullivan 2016). Finally, we add the two resulting images together into our observation. It is important to correctly choose the altitude and airmass of the observations, as these control the strength of the DCR effect at given wavelengths. The zenith angle controls the dispersion and the azimuth controls rotation or direction of the refraction. As the DCR effect is most strongly dependent on the zenith angle, we want to ensure we select a realistic set of zenith angles. To do so we randomly sample the expected distribution of zenith angles, generated based on LSST's OpSim (Delgado et al. 2014) and derived from Sebag & Vogiatzis (2014). Figure 2 shows the probability density as a function of the zenith angle in degrees. As expected, there is only a small fraction of the images taken at large angles (due to high airmass and telescope limitations). Typical zenith angles fall between  $40^\circ$  and  $50^\circ$ . For the simulation, we sample zenith angles from this distribution. For azimuthal angles, we randomly sample from a uniform distribution of angles ranging from  $0^\circ$  to  $360^\circ$ .

We introduce shot noise for both the background and sources into the synthetic exposures and set the source fluxes such that their extracted measurements have realistic uncertainties. Using the StarFast Simulator of Sullivan (2016), we generate 200 exposures with zenith and azimuth angles, chosen as described above, using the default configuration. We run tests on subsets of these 200 observations (5, 15, and 50) to simulate how well we can recover the DCR corrected images in the early years of the survey. Additionally, when noise is introduced, we generate






**Figure 2.** Distribution of zenith angles sampled in our simulations, see Sebag & Vogiatzis (2014).

50 random realizations of the noise to quantify its effect. Below we analyze these and present our results. Figure 3 illustrates a random realization of these synthetic exposures. The sources in  $\hat{x}_1$  vary from bottom to top in flux and are approximately 0, 1, 3, 5, 7, and 9 signal-to-noise ratio (S/N),  $\hat{x}_2$  has the same S/N sampling but varying from left to right.


### 3.2. Recovered Sub-band Fluxes

Our procedure is illustrated in Figure 4 where the recovered sub-band fluxes are plotted for all synthetic stars as a function of the iteration assuming 50 observations. For this study we measure the fluxes using simple aperture photometry. The left and right panels show the fluxes in  $\sigma$  units at  $\lambda_1$  and  $\lambda_2$ , respectively. By design the true fluxes are 1, 3, 5, 7, and  $9\sigma$ , which the algorithm quickly can recover despite the (intentionally) bad initialization in this illustration.

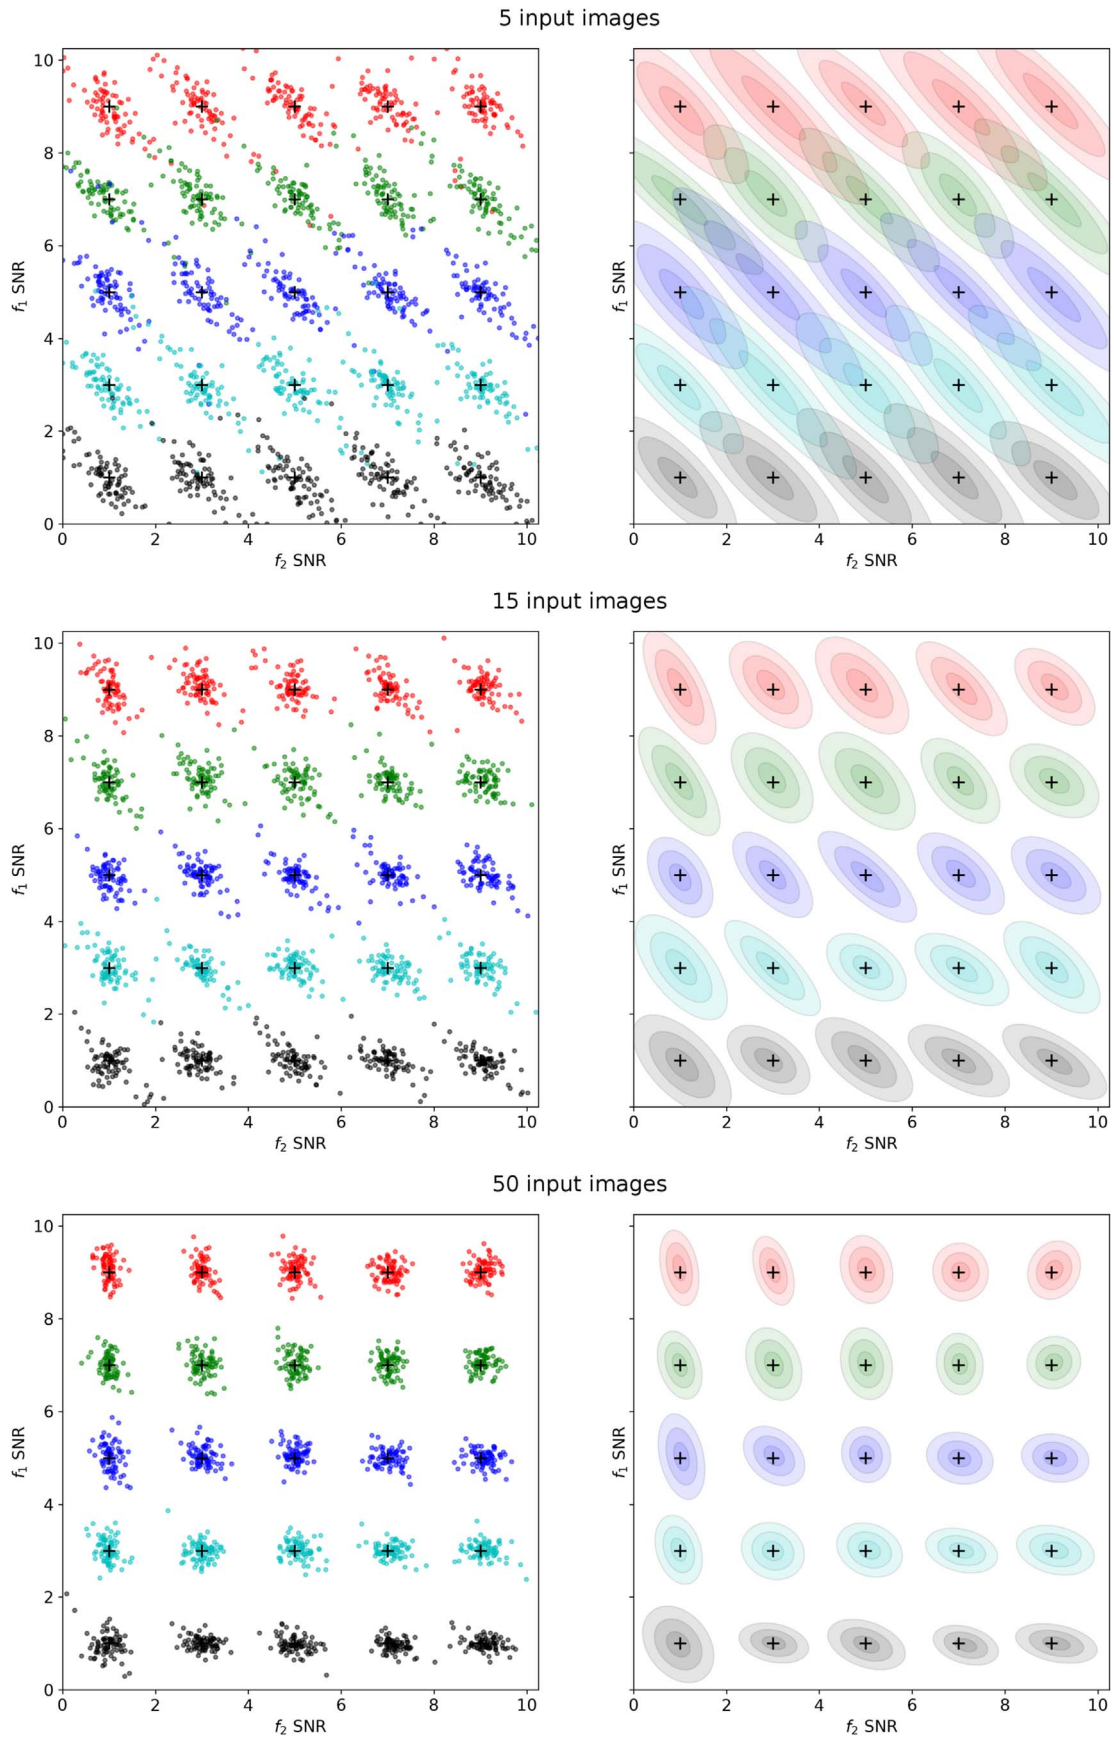
Due to the noise in the observed images, we see errors in the final recovered fluxes. To study the uncertainty of the



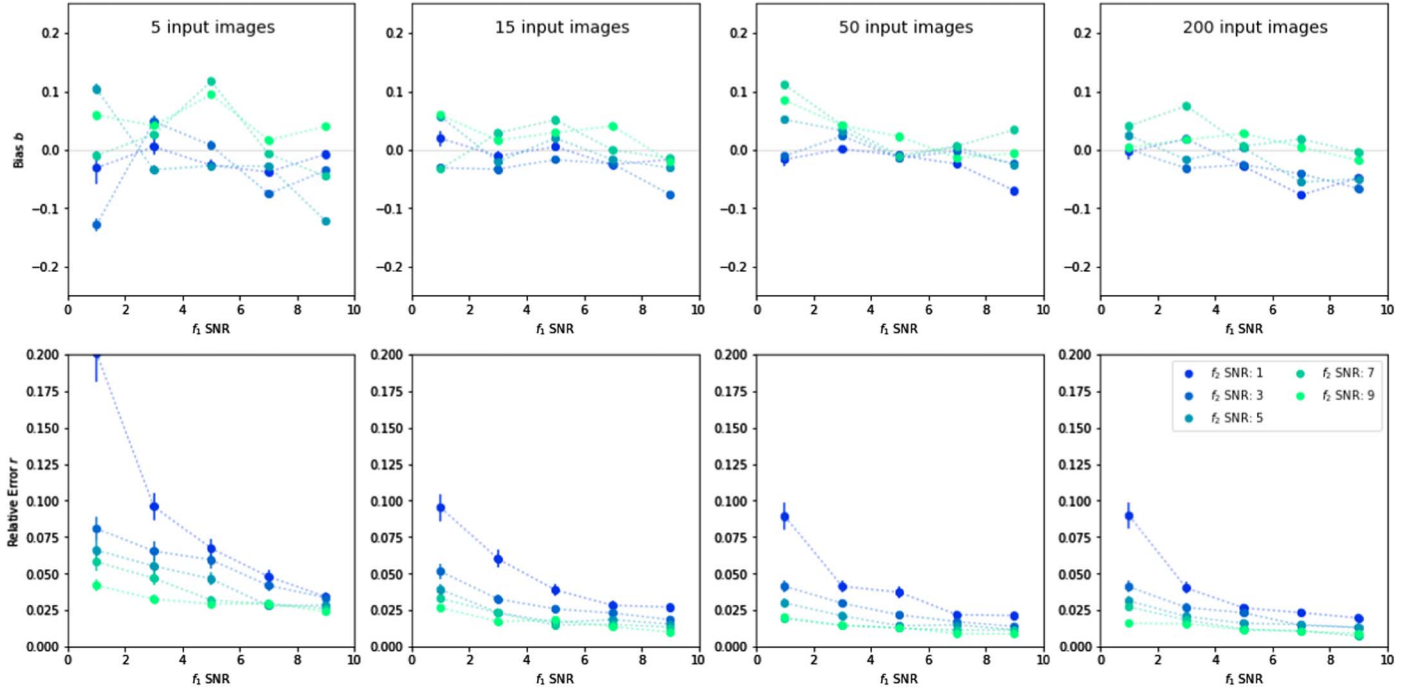
**Figure 3.** Synthetic exposures sampling source signal-to-noise and color. Differences between these images are slight and difficult to see by eye. Left: zenith angle  $9^{\circ}85'$  and azimuth angle  $147^{\circ}$ , source fluxes at 443.7 and 513.5 nm. Center: zenith angle  $25^{\circ}04'$  and azimuth angle  $334^{\circ}$ , source fluxes at 443.7 and 513.5 nm. Right: zenith angle  $50^{\circ}0'$  and azimuth angle  $77^{\circ}$ , source fluxes at 443.7 and 513.5 nm.



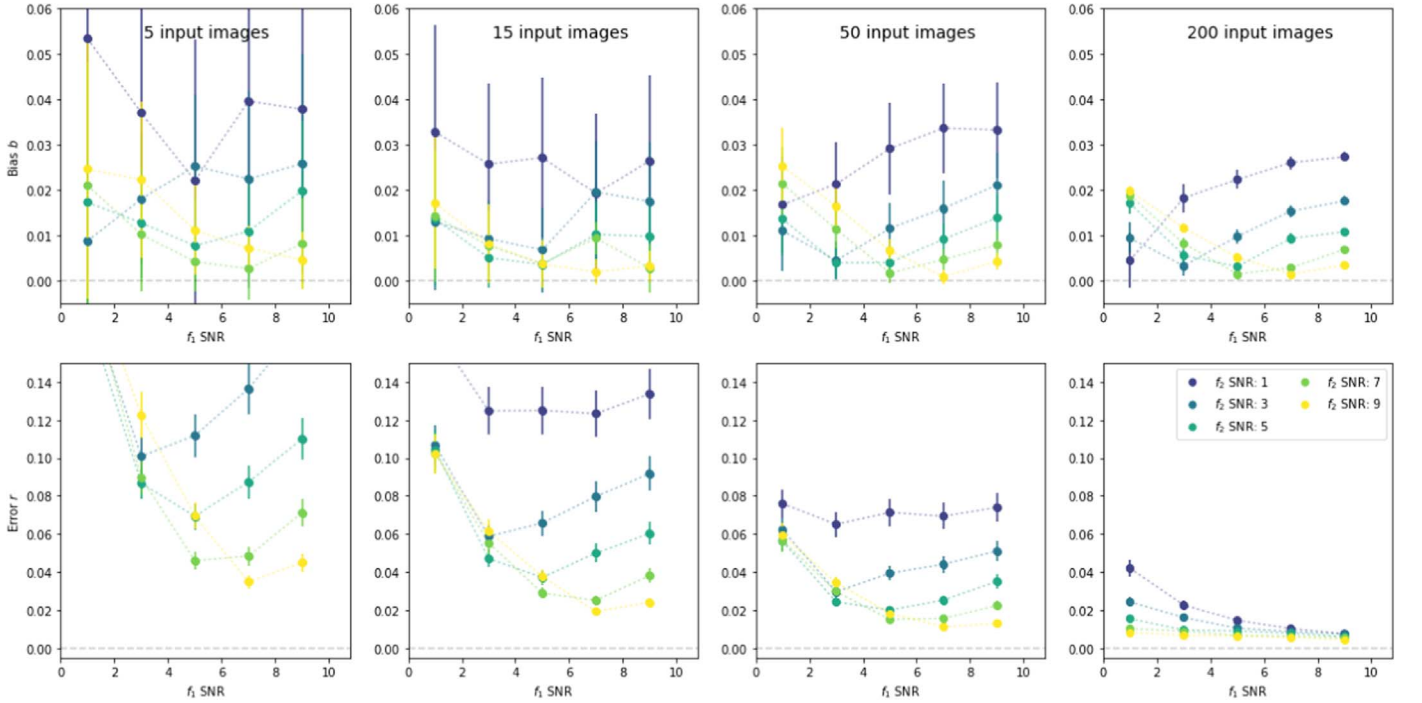
**Figure 4.** Recovered fluxes of objects as a function of iteration. Left: flux 1 measured in units of  $\sigma$  as a function of iteration. Right: flux 2 measured in units of  $\sigma$  as a function of iteration. The true flux values in units of  $\sigma$  for both sets of sources are 1.0, 3.0, 5.0, 7.0, and 9.0, indicated by the black horizontal lines. The flux 1 portion is represented by the color of the line and the flux 2 portion is represented by the line style. For example, the red dashed line indicates a flux 1 of  $9.0\sigma$  and a flux 2 of  $7.0\sigma$ . Reconstruction based on a set of 50 observations. Each observation has a different realization of zenith angle and azimuth.



**Figure 5.** Left: scatter plot of the recovered fluxes across multiple realizations; the different groups of points correspond to the different sources in Figure 3. Right: covariance distribution for each flux 1–flux 2 pairing, ellipses drawn at  $1\sigma$ ,  $2\sigma$ , and  $3\sigma$ . Each row corresponds to a different number of input images: first row, 5 images; second row, 15 images; third row, 50 images.



**Figure 6.** Relative bias and uncertainty of the recovered fluxes. The flux 1 component is plotted along the  $x$ -axis and the flux 2 component is indicated by the color, see the legend. The bias and the error are reported as relative to the expected flux. The shown error bars indicate the standard error.



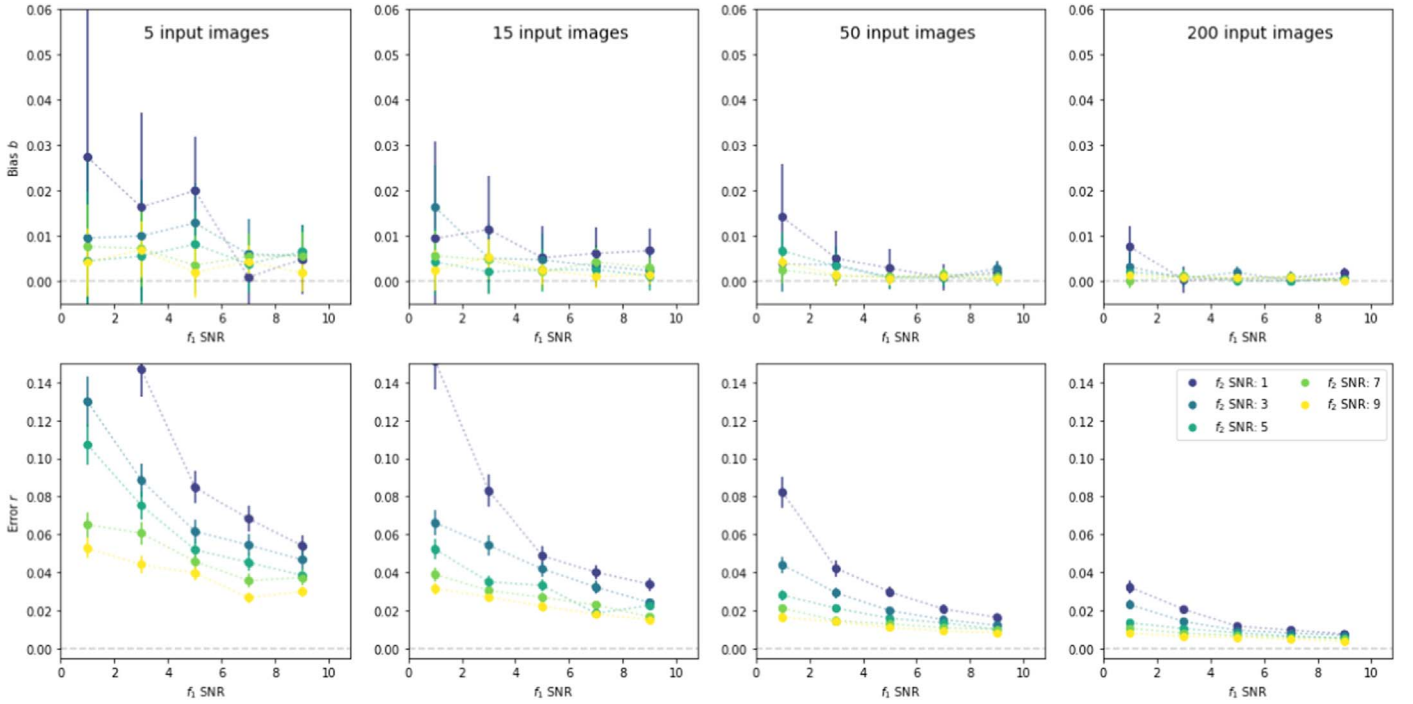
**Figure 7.** Bias (top row) and uncertainty (bottom row) of the positions of sources measured in pixels extracted from the coadded images without a correction for DCR. The flux 1 component is indicated along the  $x$ -axis of each panel and flux 2 component is represented by the color, see the legend. The shown error bars indicate the standard error.

reconstruction, we create 50 realizations and show the scatter plot of the recovered fluxes in Figure 5. The left panels show the resulting realization and right panels summarize these by visualizing the sample covariance matrices. The top, middle, and bottom panels correspond to scenarios with 5, 15, and 50 observations, respectively. For the small number of observations we see a strong anticorrelation between the two fluxes, as expected their sum is better constrained than the individual sub-


bands. As we increase the number of observations, the error ellipses shrink.

In Figure 6 we study the reconstruction as a function of the true flux. The top row shows the average bias in the flux and the bottom illustrates the sample variance; both of these are calculated to be relative to the expected flux. Beyond the initial step from 5 to 15 input images, we see the bias stay relatively constant with the increasing number of images, while the





**Figure 8.** Bias (top row) and uncertainty (bottom row) of the positions of sources measured in pixels extracted from  $x_1$  and  $x_2$  after the iterative image recovery. The flux 1 component is indicated along the x-axis of each panel and the flux 2 component is represented by the color, see the legend. The shown error bars indicate the standard error.



**Figure 9.** Left: two typical sample observations containing a combination of three sub-band images plotted as an RGB image. Center: the corresponding simulated broadband  $g$ -band observations. Right: result of our deconvolution process, resolving the corresponding RGB colors of the left column, using only images similar to the center column as an input. We show a deconvolution result not only resolving a higher color resolution than the observed images, but also significantly deblur the image and correct for the positional error caused by DCR.

variance decreases with the signal-to-noise ratio of the sources and the number of exposures.

### 3.3. Improved Astrometry

Beyond evaluating the photometry, we also study the astrometric accuracy of the new reconstructions. For the comparison, we created coadded images by averaging the input exposures. Given that the simulation did not include pointing uncertainties, the simulated exposures were not registered explicitly. Also, our study was restricted to monochromatic observations and does not make use of

observed color information that multicolor surveys would have access to. In Figure 7 we plot the quality of the astrometry estimated for each source as measured by the length of the average of the offset vector from the true coordinates to the observed centers of mass. The top panels show this bias, and the bottom panels illustrate their relative errors. We see that the estimated error improves with increasing number of images, but the offset shows significant shifts even for the largest signal-to-noise scenarios. These offsets are dependent on the sub-band color as expected from shifts due to DCR.

The reconstructed sub-band images can be combined to derive the positions of the sources and their astrometric uncertainty. In Figure 8 the bias and error plots are shown in the same cases as the coadds. The difference is day and night: the recovered images provide improved astrometry as we increase the signal-to-noise ratio in a sub-band.

In particular, the result shows little bias across the board, with an average positional bias below 0.01 pixels. Even more so, in the 50 and 200 input image case, where the average bias is below 0.005 pixels. This is largely due to the fact that the deconvolution process deblurs the source, allowing for a more accurate position finding (center of mass). The error is very similar between the two methods in the case of 200 exposures, but quickly diverges for smaller numbers of input images. For 15 images, the average error across all sources for the coadd is 0.08 pixels, where as for the deconvolution result it is only half of that.

## 4. Discussion and Summary

Deconvolutions such as these are incredibly computationally complex, requiring multiple fast fourier transform (FFT) per input image for every loop over the set of available images. Graphics processing units (GPUs) or similar massively parallel co-processors are, therefore, required in order to tackle these

types workloads within reasonable timeframes. We therefore spent a considerable effort ensuring we optimize the computations of our research code to take full advantage of our GPUs. While precise timings vary depending on the available hardware and the runtime settings, most of our test runs for this article concluded in less than 60 minutes. We make extensive use Nvidia’s cuFFT library as well as a variety of custom kernels, allowing us to keep the majority of our computation on the GPU while minimizing slow transfers to and from the GPU.

With this work, we show that small extension to our latent image recovery technique (Lee et al. 2017) enables us to account for the effects of DCR in broadband images and to use the effect of refraction to recover the underlying SED of a source (see also Sullivan 2018). Essentially we can use atmospheric refraction as a low-resolution spectrograph to improve the spectral resolution of broadband imaging. Utilizing the LSST filter response curves for the  $u$  and  $g$  bands (where the DCR effect is strongest) we can improve astrometric and photometric accuracy over standard coaddition techniques that do not incorporate corrections for DCR. From the analysis of the astrometric and photometric performance we show that the improvements in uncertainties in the recovered positions and fluxes are consistent with the increase in signal-to-noise gained by combining multiple images. For intermediate signal-to-noise sources (i.e.,  $S/N \sim 6$ ) accounting for DCR in the latent image reduces the bias in the astrometric solutions by a factor of five.

While the main focus in this paper has been on resolving two sub-bands, this technique is easily extensible to more sub-bands. In Figure 9, we use our method to resolve three sub-bands given a realistic set of images produced using LSST’s Starfast Simulator (Sullivan 2016). In the left column we show two example simulated three-band LSST images where the  $g$  band was subdivided into three equally wide wavelength regions. The color-dependent position of the sources and the extension of the PSF along the parallactic direction are clearly visible. The simulated observations (where we have a single broadband image) for the  $g$ -band filter are shown in the center column. The DCR-induced distortion to the PSF is clearly present within these images. The right panel shows the result of processing 200  $g$ -band observations, similar to those in the center column. This latent image resolves the sub-band images (shown as an inverted RGB) as well as partially deconvolves the underlying image. Considering that the method is able to handle PSFs and pixel sizes of various kinds simultaneously, we expect this method to work across instruments and telescopes as well.

The current model of the latent image of the sky assumes constant brightness across the observations and does not accommodate possible variability of sources explicitly. The robust statistical treatment of our approach will detect the varying flux of these objects and will essentially down weight their observation when they are significantly different from the typical brightness. Variability is a key topic that we plan to address in the future.



A caveat to the current study and its conclusions is that the observations of a given astronomical object will not generally

have that same distribution as the expected distribution used derived for the whole ensemble. For example, observations of the celestial pole will always occur at the same zenith angle for a given telescope. For a given object (really, a given decl.), the zenith angles and parallactic angles available are also highly correlated, occupying a one-dimensional subspace of the two-dimensional manifold. These effect should be included for detailed analyses for specific experiments.

Our results are consistent with the findings of Sullivan (2018)<sup>5</sup> and demonstrate the potential for such an approach to learn the underlying model for the sky. We note that this first in a series of papers and does not include the uncertainties introduced due to PSF estimation and errors in the underlying astrometry. We will address these issues in a forthcoming paper together with the expected spectral resolution that can be achieved by surveys such as the LSST. Given the speed of the GPU implementation we expect that one year of LSST survey data for approximately 50 arcmin<sup>2</sup> of sky could be processed in less than 30 minutes on a single GPU.

T.B. acknowledges partial support from NSF Grant AST-1412566 and IIS-1447639, as well as NASA via the awards NNG16PJ23C and STScI-49721 under NAS5-26555. A.J.C. acknowledges partial support from DOE award DE-SC0011635, NSF awards AST-1409547 and OAC-1739419, and from the DIRAC Institute. I.S. acknowledges support from the NSF and the DIRAC Institute. We thank the anonymous referee for the careful review and useful suggestions.

## ORCID iDs

Matthias A. Lee  <https://orcid.org/0000-0001-5337-1583>  
 Andrew J. Connolly  <https://orcid.org/0000-0001-5576-8189>

## References

- Annis, J., Soares-Santos, M., Strauss, M. A., et al. 2014, *ApJ*, **794**, 120
- Delgado, F., Saha, A., Chandrasekharan, S., et al. 1982, *Proc. SPIE*, **9150**, 915015
- Filippenko, A. V. 1982, *PASP*, **94**, 715
- Kaczmareczik, M. C., Richards, G. T., Mehta, S. S., & Schlegel, D. J. 2009, *AJ*, **138**, 19
- Lee, M., Budavari, T., White, R., & Gulian, C. 2017, *A&C*, **21**, 15
- Lee, M. A., & Budavári, T. 2015, in ASP Conf. Ser. 495, *Astronomical Data Analysis Software and Systems XXIV*, ed. A. R. Taylor & E. Rosolowsky (San Francisco, CA: ASP), 261
- Lucy, L., & Hook, R. 1992, adass I in ASP Conf. Ser. 25, *Astronomical Data Analysis Software and Systems I*, ed. D. M. Worrall, C. Biemesderfer, & J. Barnes (San Francisco, CA: ASP), 277
- Maronna, R., Martin, D., & Yohai, V. 2006, *Robust Statistics* (Chichester: Wiley)
- Meyers, J. E., & Burchat, P. R. 2015, *ApJ*, **807**, 182
- Peters, C. M., Richards, G. T., Myers, A. D., et al. 2015, *ApJ*, **811**, 95
- Rowe, R. B., Jarvis, M., Mandelbaum, R., et al. 2015, *A&C*, **10**, 121
- Sebag, J., & Vogiatzis, K. 2014, *Proc. SPIE*, **9150**, 91500R
- Sullivan, I. 2016, *Starfast—A Fast Simulation Building Tool for Testing Algorithms*, <https://dmtn-012.lsst.io/>
- Sullivan, I. 2018, *Dmtn-037: Dcr-matched Template Generation*, Zenodo, <https://doi.org/10.5281/zenodo.1492936>
- Szalay, A. S., Connolly, A. J., & Szokoly, G. P. 1999, *AJ*, **117**, 68
- Zackay, B., & Ofek, E. O. 2017, *ApJ*, **836**, 187

<sup>5</sup> LSST report at <https://dmtn-037.lsst.io>.