












Removing Atmospheric Fringes from Zwicky Transient Facility *i*-band Images using Principal Component Analysis

Michael S. Medford^{1,2} , Peter Nugent² , Danny Goldstein² , Frank J. Masci³ , Igor Andreoni⁴ , Ron Beck³,
Michael W. Coughlin⁵ , Dmitry A. Duev⁴ , Ashish A. Mahabal^{4,6} , and Reed L. Riddle⁷ 

¹Department of Astronomy, University of California, Berkeley, CA 94720, USA

²Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA 94720, USA

³IPAC, California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125, USA

⁴Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

⁵School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

⁶Center for Data Driven Discovery, California Institute of Technology, Pasadena, CA 91125, USA

⁷Caltech Optical Observatories, California Institute of Technology, Pasadena, CA 91125, USA

Received 2021 February 16; accepted 2021 May 5; published 2021 June 24

Abstract

The Zwicky Transient Facility is a time-domain optical survey that has substantially increased our ability to observe and construct massive catalogs of astronomical objects by use of its 47 square degree camera that can observe in multiple filters. However the telescope’s *i*-band filter suffers from significant atmospheric fringes that reduce photometric precision, especially for faint sources and in multi-epoch co-additions. Here we present a method for constructing models of these atmospheric fringes using Principal Component Analysis that can be used to identify and remove these artifacts from contaminated images. In addition, we present the Uniform Background Indicator as a quantitative measurement of the reduced correlated background noise and photometric error present after removing fringes. We conclude by evaluating the effect of our method on measuring faint sources through the injection and recovery of artificial stars in both single-image epochs and co-additions. Our method for constructing atmospheric fringe models and applying those models to produce cleaned images is available for public download in the open source Python package *fringeZ* (<https://github.com/MichaelMedford/fringeZ>).

Key words: Algorithms – Computational methods – Open source software – Principal component analysis – Publicly available software – Astronomical optics – CCD photometry

1. Introduction

Large scale synoptic surveys have produced an abundance of optical images at scales previously unseen. Surveys such as the Palomar Transient Factory (Law et al. 2009), Catalina Real-Time Transient Survey (Drake et al. 2009), Zwicky Transient Facility (Masci et al. 2018; Bellm et al. 2019b; Graham et al. 2019) and others have revolutionized our understanding of the universe by generating massive transient data sets. These data sets require advances in computational processing techniques. The Vera C. Rubin Observatory will generate 20 terabytes of data per night, totaling over 500 petabytes of imaging data over the 10 yr of the survey (Ivezic et al. 2019). Such data flows require that the reduction of raw images into calibrated science images must be systematic and require minimal human intervention.

One significant source of noise in long wavelength optical imaging is atmospheric emission lines. These emission lines are produced by highly non-thermal atomic and molecular transitions (primarily O II and OH) and are influenced by the temperature and density of the upper and lower atmosphere as well as the current solar activity. Thus the strength of these

lines can vary throughout the night and is proportional to the airmass through which the telescope is pointed. Fringe patterns appear when photons at these wavelengths fall onto thin charge-coupled devices (CCDs) due to the self-interference caused from light reflecting off of the back of the imaging instrument before it is absorbed by the CCD itself (Bernstein et al. 2017). Thick CCDs rarely see this effect save at the longest wavelengths. For thin CCDs, these fringe patterns can introduce significant noise into *i*-band and *z*-band images, rendering photometry and image subtraction ineffective at faint magnitudes without a calibration correction that successfully removes them. These fringes also appear in interference image spectrometers where attempts have been made to remove them using wavelet transformations (Ren et al. 2017).

Principal component analysis (PCA) is a statistical method for reducing data to a set of orthogonal components by finding vectors of minimal variance within the data (Jolliffe & Cadima 2016). PCA has been shown to be an effective method for modeling and removing atmospheric fringes from red optical and near-IR images. These methods work by building a set of orthogonal component images, constructed from a large

representative sample, which can approximate the fringes of a single image through linear combination. The dot product of individual images against these orthogonal components results in eigenvalues that, when used as weights to these orthogonal images, generate bias images of the atmospheric fringes. Subtracting these bias images removes the fringe noise and improves the photometric precision of astrophysical source measurements.

Previous PCA methods for atmospheric fringe subtraction constructed orthogonal images by down-sampling fringed images into a lower resolution before re-parameterization of the data. Bernstein et al. (2017) compress an image into a sparse set of features before attempting reduction into orthogonal components. This method has strong results but lacks the ability to resolve contributions from individual atmospheric lines due to its compression. PCA performed on a per-pixel basis is more computationally expensive but has the power to capture these individual fringe effects, resulting in more accurate photometry.

The computational resources dedicated to the reduction and calibration of astronomical images has grown over the decades to keep up with increasing data flows. The development of highly optimized PCA algorithms, along with additional computational resources, have now made it possible to model fringes via per-pixel PCA analysis. Here we present the implementation of such a method on full resolution Zwicky Transient Facility *i*-band data. In Section 2, we outline the Zwicky Transient Facility instrument and data set. In Section 3, we present our method for implementing per-pixel PCA atmospheric fringe modeling and removal, as well as the Uniform Background Indicator as a quantitative measurement of correlated background noise. In Section 4, we analyze the results of applying our method, including increased photometric precision on faint sources and the ability to detect otherwise undetectable sources in multi-epoch co-additions. We discuss and conclude in Section 5.

2. The Zwicky Transient Facility Instrument

The Zwicky Transient Facility (ZTF) is an optical time-domain survey that has been operating on the 48 inch Samuel Oschin Telescope at Palomar Observatory since 2018 March (Bellm et al. 2019b; Graham et al. 2019). ZTF’s camera covers 47 square degrees in a single exposure, enabling coverage of the entire visible Northern sky every few nights in ZTF *g*-band, *r*-band and *i*-band filters with an average point-spread function (PSF) full width at half maximum of $2''.0$ on a plate scale of $1''.01 \text{ pixel}^{-1}$. The ZTF camera is divided into 16 CCDs, each covered with an anti-reflective (AR) coating, with each CCD split into four separate readout channels for a total of 64 readout channels (Dekany et al. 2020). Surveys with the telescope over its first several years of operations (Bellm et al. 2019a) take standard 30 s exposures that achieve median 5σ

limiting magnitudes of $r \approx 21.0$, $g \approx 21.25$, and $i \approx 20.0$ (Bellm et al. 2019b). The Infrared Processing and Analysis Center (IPAC) reduces the data for the ZTF survey, including producing a real-time alert stream triggered by transient event detections on difference images (Patterson et al. 2018). In addition to these alerts, the ZTF collaboration routinely produces public data releases that contain, among other data products, lightcurves assembled from single image PSF photometry for every star in the northern sky that appears in a deep co-added reference image (Masci et al. 2018). Reference images are ideally constructed from 40 individual exposures, although weather and visibility produces variable results for different areas of the sky. ZTF’s observing time is split between public observations (40%) (funded by the National Science Foundation’s Mid-Scale Innovations Program or MSIP), partnership observations (40%) (which are held in a proprietary period for collaboration members of the survey), and Caltech observations (20%). Starting in 2020, ZTF began operating its second phase as ZTF-II with 50% of observation time funded by a MSIP survey that images the entire visible sky every other night. The ZTF *i*-band filter is used exclusively for partnership observations, such as the ZTF Ultra Deep Survey (ZUDS), before being released to the public after the partnership’s proprietary period. Due to the decreased limiting magnitude of the *i*-band relative to the *g*- and *r*-band, standard exposures in the *i*-band are set to 90 s and increase the filter’s median limiting magnitude to a comparable value. Atmospheric fringes appear exclusively in the *i*-band and therefore this work uses partnership data not yet released to the public.

3. Method for Removing Atmospheric Fringes

Our method for removing atmospheric fringes was a two step process. First the PCA eigenvectors for a single readout-channel were extracted from a large set of images with significant fringing. Second each image was processed through this model to generate a unique bias image that, when subtracted from the original image, removed the atmospheric fringes. Figure 1 outlines a visual representation of the steps in this method. An explanation of the data products and corresponding symbols used throughout this section can be found in Table 1.

3.1. Generating PCA Fringe Models

ZTF’s reduction pipeline was designed to process each of the camera’s 64 readout-channels separately. We therefore began by gathering images of a common readout-channel together for model generation, with the goal of generating 64 fringe models. Training images were selected by gathering all *i*-band images between 2019 April 1 and 2020 April 1 and removing images with a limiting magnitude less than 19. This cut on limiting magnitude removed cloudy images from our sample that would fail to have a photometrically accurate measurement of the

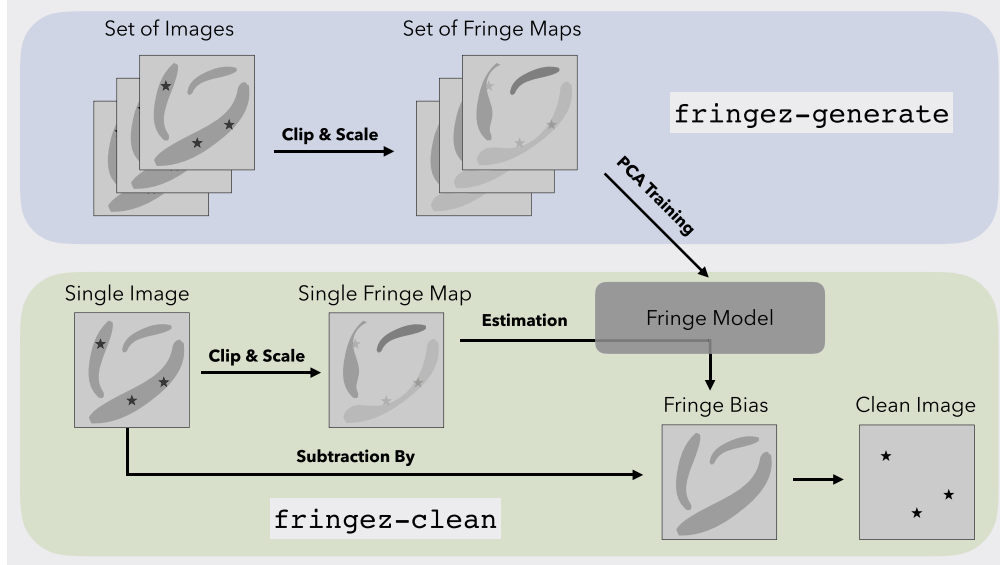


Figure 1. The process for removing atmospheric fringes was done in two steps. A fringe model for each readout channel was constructed by running fringe maps through PCA training, where fringe maps are a clipped and scaled version of single-epoch images. This step was only performed once per readout channel. Cleaning every single-epoch image thereafter was performed by using the eigenvectors in the fringe model to generate a fringe bias image for each fringed image. This fringe bias was subtracted from the fringed image to create a clean image. Functions for generating fringe models and cleaning fringed images can be found in the open source `fringe-z` (<https://github.com/MichaelMedford/fringe-z>) package under the executables `fringe-z-generate` and `fringe-z-clean` respectively.

Table 1
`fringe-z` Math Symbols

Description	Symbol	Dimensions
Fringed Image	$\mathbf{I}_{\text{fringe}}$	$[N_{\text{rows}} \times N_{\text{cols}}]$
Flattened Fringed Image	\vec{I}_{fringe}	$[1 \times N_{\text{pixels}}]$
Fringe Map	\vec{F}_{map}	$[1 \times N_{\text{pixels}}]$
Fringe Map Eigen-values	$\vec{\lambda}_{\text{map}}$	$[1 \times N_{\text{comp}}]$
Fringe Bias	\vec{F}_{bias}	$[1 \times N_{\text{pixels}}]$
Average of Training Set	$\vec{\mu}_{\text{model}}$	$[1 \times N_{\text{pixels}}]$
Model Eigen-images	$\mathbf{v}_{\text{model}}$	$[N_{\text{comp}} \times N_{\text{pixels}}]$
Model Explained Variance	\vec{R}_{model}	$[1 \times N_{\text{pixels}}]$

night sky and thus its atmospheric lines. By including all images within these dates we ensured a representative sample of atmospheric conditions and airmasses which are correlated with the strength of atmospheric fringes. There were 550,365 *i*-band images included in total, with each fringe model trained on between 8062 and 8898 images.

We began by selecting all fringed images ($\mathbf{I}_{\text{fringe}}$) for a single readout channel. The location and strength of atmospheric fringes are independent of the astrophysical sources within an image. Therefore we needed to remove sources from each image prior to training in order to reduce the confusion in the PCA variance reduction process. Our preferred method for identifying pixels containing astrophysical sources was to use the quality masks produced by IPAC that labels pixels containing a source in each image’s source catalog. In cases where this pixel mask could not

be obtained, pixels containing sources were identified by calculating the median absolute deviation for the image and flagging pixels that were 5 standard deviations above or below the image’s median value. This process did not clip the atmospheric fringes themselves, which were at the level of the background noise and therefore significantly below the level of these thresholds. All pixels flagged as containing sources were replaced with the value of the image’s global median, removing stars from our images without changing the value of the vast majority of pixels. A possible improvement to our method could be to replace the pixel instead with an estimate of the local background. We choose the global median because the local background can be difficult to estimate due to the fringes themselves and in particular in regions where the transverse length scale of the fringes approaches the pixel scale of the image. Each image was then scaled so that they could be used for training together with other images taken at different airmasses and limiting magnitudes. We named each training image a fringe map (\vec{F}_{map}), as it traces the relative location and strength of only the atmospheric fringes (Figure 1). Fringe maps were created by subtracting the median from each image, dividing the result by the image’s median absolute deviation, and flattening the image into a $[1 \times N_{\text{pixels}}]$ array:

$$\vec{F}_{\text{map}} = \frac{\mathbf{I}_{\text{fringe}} - \vec{I}_{\text{fringe}}}{\text{MAD}(\mathbf{I}_{\text{fringe}})} \Big|_{\text{flattened}}. \quad (1)$$

(Because PCA is a linear method that calculates coefficients irrespective of the basis in which the data is represented,

performing our analysis in one-dimension has no effect on the results.) Each pixel in the fringe map was treated as a separate variable in our PCA analysis, resulting in 9,461,760 variables (3080 rows by 3072 columns) for each model. The eigenvalues of the feature vector were determined using the randomized Singular Value Decomposition method (Halko et al. 2011; Martinsson et al. 2011) in the `scikit-learn` Python package (Pedregosa et al. 2011). Each one-dimensional eigenvector was reconstructed into the original image shape to create an eigen-image. The final fringe model was a set of eigen-images that captured orthogonal contributions to the training set variance in each pixel. We also recorded each component’s explained variance (\vec{R}_{model}), or the amount of training set variance pointed in each eigenvector’s direction, for later use in the processing pipeline. We then repeated this process to construct a unique fringe model for each of the camera’s 64 readout channels.

3.2. Removing Fringes from Single-epoch Images

Fringes were removed from images by applying the fringe model in these steps:

$$\vec{\lambda}_{\text{map}} = \frac{(\vec{F}_{\text{map}} - \vec{\mu}_{\text{model}}) \cdot \mathbf{v}_{\text{model}}^T}{\sqrt{\vec{R}_{\text{model}}}}, \quad (2)$$

$$\vec{F}_{\text{bias}} = [[\vec{\lambda}_{\text{map}} \cdot (\sqrt{\vec{R}_{\text{model}}} \mathbf{v}_{\text{model}})] + \vec{\mu}_{\text{model}}] \cdot \text{MAD}(\vec{I}_{\text{fringe}}), \quad (3)$$

$$\vec{I}_{\text{clean}} = \vec{I}_{\text{fringe}} - \vec{F}_{\text{bias}}. \quad (4)$$

First, a fringe map was regenerated for each fringed image in our sample. The dot product of this fringe map (less the average value of each pixel in the training set $\vec{\mu}_{\text{model}}$) and the model’s eigen-images was calculated, and then divided by the square root of that component’s explained variance (Equation (2)). This produced the eigenvalues ($\vec{\lambda}_{\text{map}}$) for each fringe map. These eigenvalues were used as coefficients to linearly combine re-scaled eigen-images and the average values of the training sets were added back in (Equation (3)). This image, after being re-scaled by the median absolute deviation of the original fringed image, was named a fringe bias (\vec{F}_{bias}). When this fringe bias was subtracted from the original fringed image it produced a clean image (\vec{I}_{clean}) with significantly less fringing (Equation (4)).

An example of each of the images in this process are shown in Figure 2. This figure visually demonstrates the results of our method. The original image is a 90 s *i*-band exposure that contains a representative amount of atmospheric fringes. The fringe map is almost entirely devoid of individual sources, although the source pixel identification method does struggle around particularly bright stars. The fringe bias generated from processing the fringe map through the fringe model successfully identifies the location

and strength of each atmospheric fringe. The clean image has nearly all of the atmospheric fringing pattern removed while retaining nearly all of the astrophysical sources.

The arithmetic average of the fringe maps used to train the fringe model for each readout-channel is shown in Figure 3, with each readout-channel placed on a common gray-scale. The pattern that clearly emerges is coherent on the level not just of the readout channel, but also over the full CCD. The etching process for creating a thinned CCD uses a circular buffer that removes layers of the chip to create a uniformly thick device. The average of the training fringe maps identifies the residual thickness variations for each CCD. Work is ongoing to use these psuedo-measurements of the thickness to improve the ZTF data quality pipeline (R. Dekany 2021, private communication). The inner 32 readout channels ($16 \leq \text{rcid} < 48$) and the outer 32 readout channels ($0 \leq \text{rcid} < 16$, $48 \leq \text{rcid} < 64$) have distinctly different amounts of atmospheric fringing present in their images. The inner readout channels have two layers of AR coating while the top and bottom rows of CCDs have a single layer coating. This causes the outer readout channels to have a higher reflectivity at the longer wavelengths where fringing occurs and thus a larger contrast of the fringing pattern.

The eigen-images for readout channel 13 are shown in Figure 4 as a representative example of a fringe model. Each pixel in the image was trained as a separate variable and yet the fringe patterns across pixels remains coherent after reconstruction into the original image dimensions. This confirms that the components contain correlated eigenvectors for the different features. In addition to the atmospheric fringes, the PCA training captured large scale variations in the background that remain after flat-fielding. These backgrounds dominate the higher order components. This indicates that some global scale variations in flux remains after the execution of the ZTF flat-fielding pipeline that could possibly be improved by implementing a PCA method (Bernstein et al. 2017). We note that PCA requires setting the number of components, or number of eigenvectors, for the reduction algorithm. The fractional explained variance for all of the components are also shown in Figure 4. On average across the 64 readout channels, the first component captured 64.6% of the pixel variance while the sixth component captured only 5.0% of the pixel variance. In total six PCA components reconstructed 95.0% of the variance seen in our training sample. Adding additional components failed to significantly increase the amount of fractional explained variance as the previous component. We therefore determined that six components was sufficient to capture the variation due to atmospheric fringes.

Successfully training the fringe models required simultaneously holding large numbers of training images in computer memory in order to execute singular value decomposition. The large number of images that trained each of our 64 fringe models ranged in memory from 285 gigabytes to 314 gigabytes. We trained our models on the Cori Haswell nodes

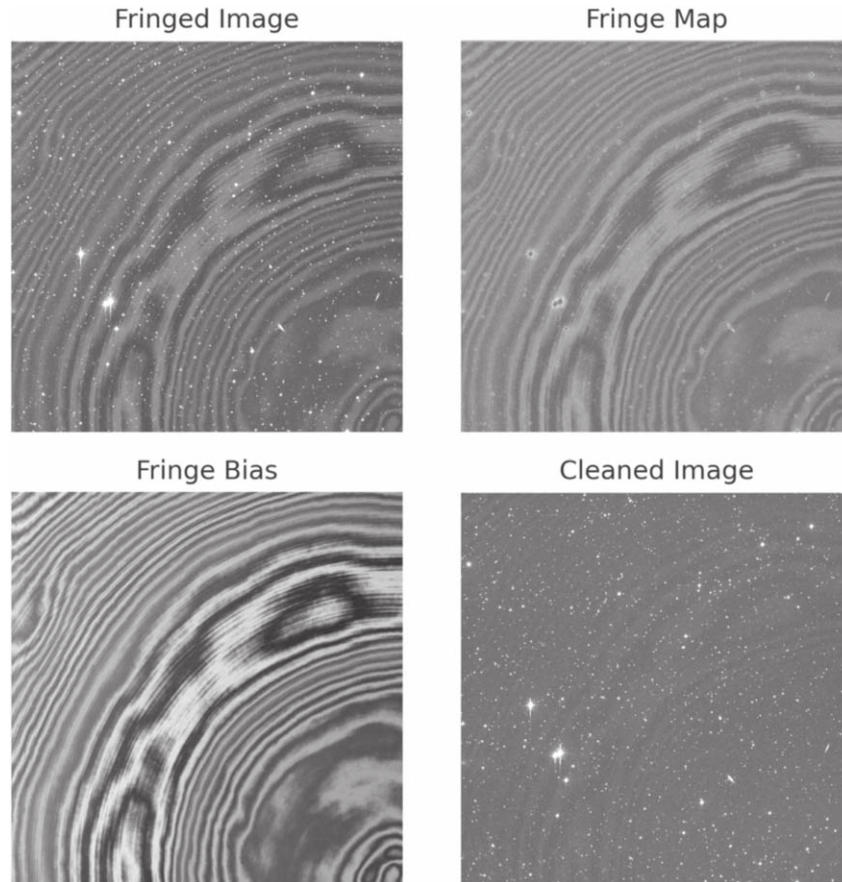


Figure 2. An example of an *i*-band image in the various steps of our fringe removal method. Fringed images were clipped and scaled to produce a fringe map. The fringe map was processed through a fringe model to generate a fringe bias. This bias image was subtracted from the original fringed image to create a cleaned image. While a fringe model only needs to be constructed once, each single-epoch image has a unique fringe bias determined by the linear combination of eigen-images in the fringe model that best reconstructs the fringe map. This ensures that the fringe model can successfully remove fringes arising from a variety of airmasses and seeing conditions.

at the NERSC Supercomputer located at the Lawrence Berkeley National Laboratory. The Haswell nodes each have 128 gigabytes of RAM, making it impossible to include such a large number of images in our training sets without modifying our method. To adapt to cases where access to computational resources are limited such as this one, our training method was modified by down-sampling the training set. Training images were sorted by limiting magnitude and split into 300 sub-stacks of approximately equal depth. In order to combine each sub-stack into a single training image we calculated the sub-stack’s median. While taking the median can distort the correlations between pixels that PCA is attempting to solve for, the median is also more robust than the mean to removing outliers that have contaminated the signal from a different distribution than the one attempting to be measured. Smooth background gradients remain on some ZTF images that can confuse the PCA reduction algorithm if not removed as outliers. We therefore settled on taking the median instead of the mean to combine the sub-stack of fringe maps into a single training

image. An alternative approach could be to entirely remove the images with smooth gradients through a different method and then combine the remaining images in the sub-stacks using the mean. These 300 training images totaled approximately 11 GB and were therefore able to be fed simultaneously into our training process. Comparing the fringe models that result from training on down-sampled fringe maps produced comparable results to training on individual fringe maps for smaller sample sizes.

The code used for generating fringe models from ZTF *i*-band images and cleaning fringed images has been released in the open-source package `fringeZ` (Medford 2021). The `fringeZ` package includes command line executables and Python functions for downloading the fringe models computed as a result of this work, producing fringe bias images, and cleaning fringed images. This package is available for download and installation at <https://github.com/MichaelMedford/fringeZ>, and all fringe models are available for download at <https://portal.nersc.gov/project/ptf/iband>. In 2019 October, `fringeZ` was implemented into the ZTF

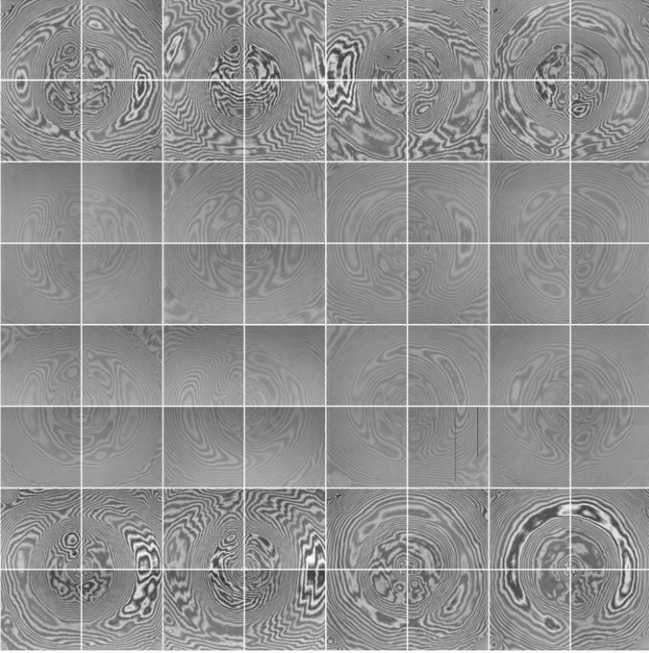


Figure 3. The average of the training fringe maps for each readout channel’s fringe model on the ZTF camera, all placed onto a common gray-scale. These circular patterns roughly trace the thickness variations in the CCD and show the circular pattern resulting from flattening the device. The inner 32 readout channels show significantly less atmospheric fringing due to an additional layer of anti-reflective coating, as compared to the outer 32 readout channels.

IPAC data reduction pipeline with fringe models built on 500 training images per readout channel. All *i*-band images taken previous to this date were also reprocessed by IPAC to remove atmospheric fringes. In 2020 November, fringe models were updated to include the 550,365 *i*-band images described and investigated in this paper. This implementation of `fringe2` and the current version of fringe models will continue to be a part of the IPAC data reduction pipeline in ZTF-II. The current implementation of `fringe2` can also be generalized to other instruments by simply altering the FITS extensions and header keywords that are currently chosen for ZTF images.

4. Measuring Improved Photometric Precision

4.1. The Uniform Background Indicator

In order to assess the effectiveness of our de-fringing method, we created a procedure for measuring correlated background noise. In this method, aperture photometry is taken at random locations on the background of an image. At the location of each aperture the flux (F_{bkg}) and its associated measurement error (σ_{bkg}) are measured, resulting in a set of aperture fluxes ($\{F_{\text{bkg}}\}$) and a set of errors ($\{\sigma_{\text{bkg}}\}$). We then compare the standard deviation of background fluxes ($\text{std}(\{F_{\text{bkg}}\})$) to the average error on the background flux measurements ($\text{median}(\{\sigma_{\text{bkg}}\})$). We call

the ratio of these two values the Uniform Background Indicator:

$$\Psi = \frac{\text{std}(\{F_{\text{bkg}}\})}{\text{median}(\{\sigma_{\text{bkg}}\})}. \quad (5)$$

For an image of uniform Gaussian noise, $\Psi \approx 1$ as the distribution in background fluxes will be equal to the average flux error across all measurements. Here there exists no global variance which is not captured by the local error term. For an image with correlated background noise, $\Psi > 1$ because the different background values sampled by the apertures will introduce additional variance into the numerator of Equation (5) not captured by the local error terms in the denominator. The flux errors in Equation (5) must be calculated locally and be in the same units as the flux measurement with respect to aperture area. The aperture size should be chosen to be approximately the PSF scale of the instrument so that UBI measurements will indicate what fringe signal (or any other source of correlated background noise) a particular instrument would observe.

We validated using the UBI as a measurement of correlated background noise through a controlled experiment. We generated images containing only Gaussian noise that were the same size as ZTF images. The error on each pixel value was set as the root-mean-square of the median-subtracted pixel values in a 10 pixel by 10 pixel box around that pixel. The size of this square had to be smaller than the spatial scale of the correlated background noise so that it contained minimal variation in the background value. We then laid 50,000 circular apertures onto the image and calculated the aperture flux and aperture flux error of each measurement. The UBI as outlined in Equation (5) was evaluated five times on each image producing an average UBI as well as an error on that measurement.

In Figure 5, we show the results of this experiment for varying aperture sizes. For all Gaussian images and aperture sizes we found $\Psi \approx 1$, confirming our interpretation that this value indicates no correlated background noise on an image. The location and scale of the Gaussian noise was also found to have no effect on this measurement. Figure 5 shows three example images with correlated background noise and the UBI values that they produce, covering the range of Ψ that we detected throughout this analysis. These example images show the correlation between a quantitatively larger Ψ and a qualitative increase in the appearance of atmospheric fringes.

Having verified that the UBI is a valid indicator of correlated background noise, we next measured the effect of removing atmospheric fringes on the UBI. A sample of *g*-band, *r*-band and *i*-band images was created by downloading one random image for each filter, readout-channel and field observed in the ZUDS survey from the week of 2020 February 1 to 2020 February 8 with a limiting magnitude greater than 19. This sampling method ensured a representative sample of airmasses and limiting magnitude for ZTF observations, while removing images of extremely low quality from the sample. Our final

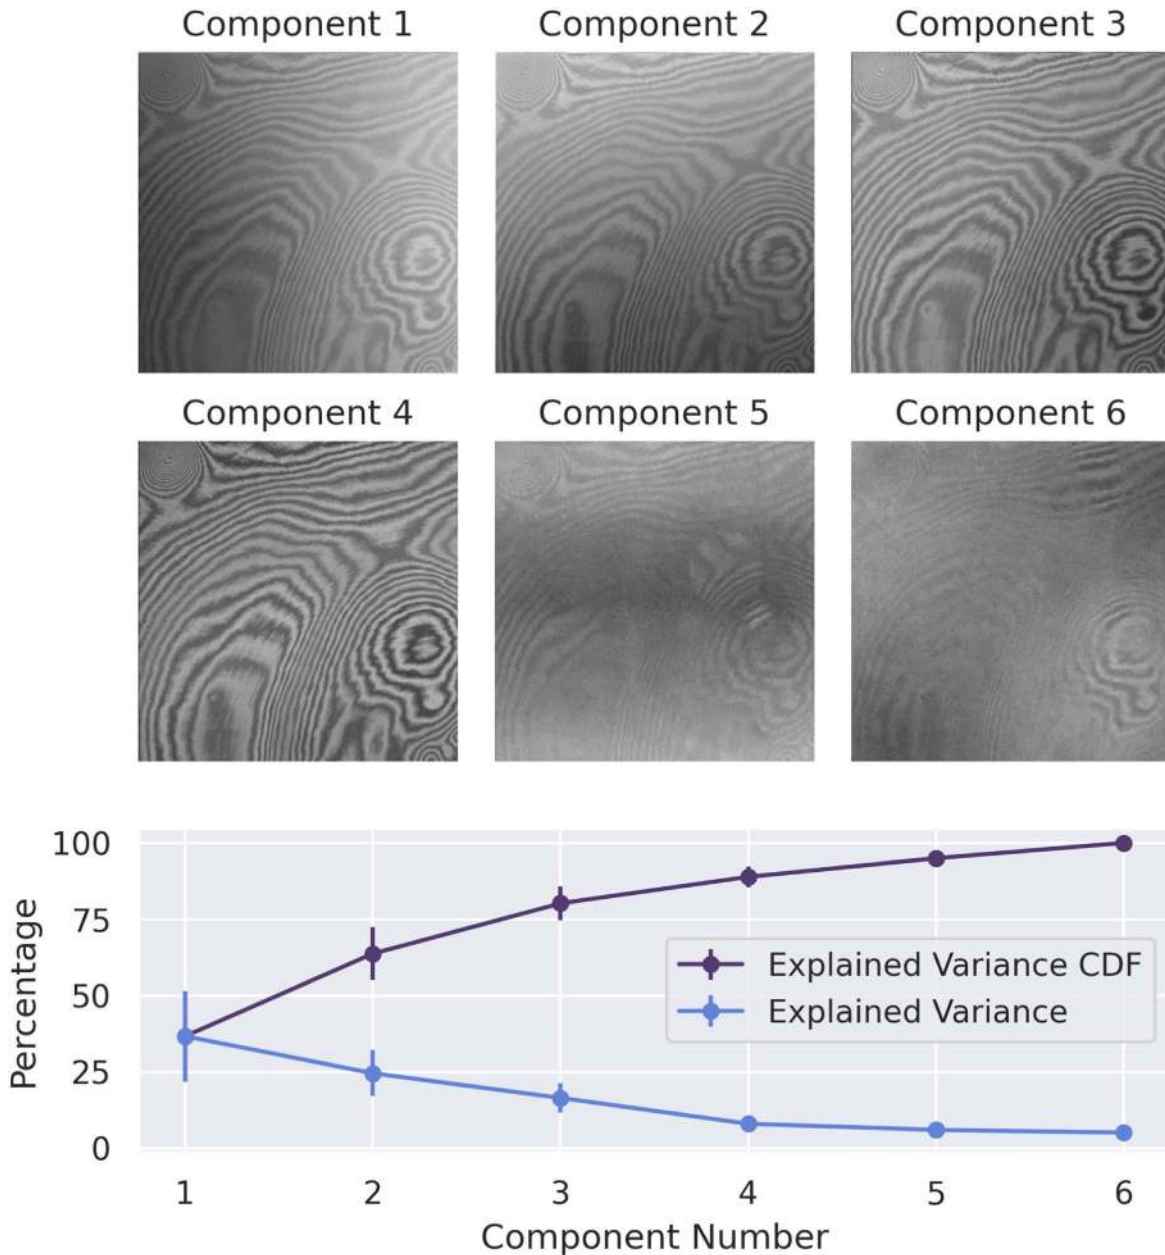


Figure 4. Each readout channel has a distinct fringe model with six eigen-images constructed from the reduction of thousands of training fringe maps. Here is an example of the eigen-images from readout channel 13 (top). While the fringe models were trained on a one-dimensional array of pixels, the two-dimensional fringes remain intact. The fringe pattern of the readout channel is clearly evident with slight variations in position and strength among the first four eigen-images. However the last two components (and to some extent the first four) contain smooth global gradients that would ideally be removed by flat-fielding. The fractional explained variance for each component across all 64 fringe models (bottom) show that the fifth and sixth components captured far less variance in the training sample. The first component captured 64.6% of the pixel variance while the sixth component captured only 5.0% of the pixel variance. In total six PCA components reconstructs 95.0% of the variance seen in our training sample. We therefore chose to have six components in our models.

sample contained 1408 *g*-band, 1407 *r*-band, and 1472 *i*-band images across a range of stellar densities. Each of the *i*-band images was cleaned using our method described in Section 3.2. The UBI was then calculated for each of the fringed images, as well as the cleaned images.

The distribution of the UBI for all of the images in our sample is shown in Figure 6, with $\Psi \geq 1$ as expected. The *g*-band and *r*-band images each have a nearly normal UBI distribution with medians of $\Psi = 1.13$ and $\Psi = 1.16$, indicating a small but measurable amount of correlated background noise

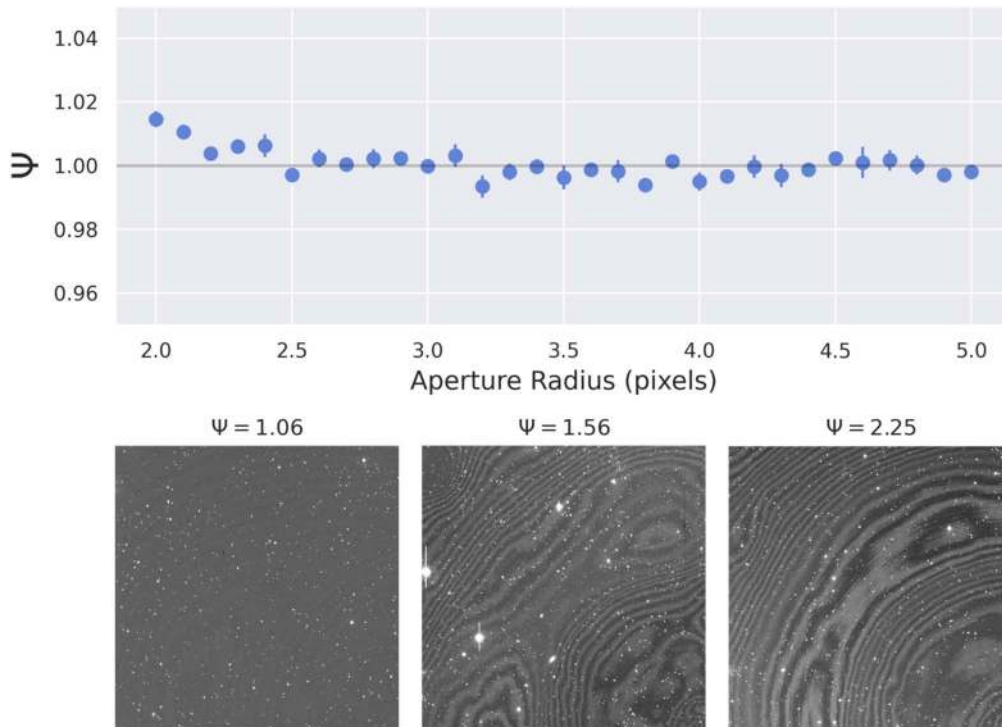


Figure 5. The Uniform Background Indicator is $\Psi \approx 1$ for a range of aperture sizes on images of Gaussian noise (top). This validates interpreting $\Psi \approx 1$ as measuring no correlated background noise in an image. Three example *i*-band images and their Ψ values are shown (bottom). More prominent atmospheric fringes correlates with an increase in the value of Ψ .

remaining after flat-fielding. The *i*-band images contain far more correlated background noise and are notably bimodal, with one population averaging $\Psi = 1.33$ and a second wider distribution averaging a larger $\Psi = 1.91$. This split is caused by the location of the readout channel on the image plane. The inner 32 readout channels ($16 \leq \text{rcid} < 48$) with their additional layer of AR coating are less susceptible to atmospheric fringing. The outer 32 readout channels ($0 \leq \text{rcid} < 16$, $48 \leq \text{rcid} < 64$) experience significantly more fringing due to a lack of an additional layer of AR coating. Less than 20% of the images have a $\Psi < 1.72$ and the 20% most fringed images have a $\Psi \geq 2.20$. The cleaned *i*-band images show indistinguishable behavior between the inner and outer readout channels forming a single distribution averaging $\Psi = 1.15$. This population appears similar to the *g*-band and *r*-band populations, indicating successful removal of atmospheric fringes. However the *g*-band and *r*-band images have significantly longer tails with an 80th percentile of $\Psi = 1.28$ and $\Psi = 1.32$ respectively, compared to an 80th percentile for cleaned *i*-band images of $\Psi = 1.20$. This indicates that there is correlated background noise occurring in *g*-band and *r*-band images that PCA analysis could potentially model and remove. It is clear that the process of removing atmospheric fringes significantly reduces correlated background noise from *i*-band images.

It is reasonable to predict that the UBI will increase with airmass, as the presence of atmospheric fringes is caused from the column of atmosphere through which images are taken. This is found to be the case in Figure 7, where collecting the images into airmass bins and calculating the median UBI shows a trend toward larger UBI for larger airmass in *i*-band images. There is a slight increase in the *g*-band and *r*-band images as well, although the effect is relatively weak. The cleaned *i*-band images have comparable UBI values to the *g*-band and *r*-band images, again confirming the effectiveness of our method.

4.2. Photometric Error Due to Fringes

While the UBI is a useful measurement of the presence of correlated background noise, we sought to quantitatively measure the improvement in photometric precision that resulted from removing atmospheric fringes. First we measured the photometric error caused by atmospheric fringes relative to a reference catalog. Next we measured how de-fringing removes this error by injecting faint artificial sources into single images. Last we measured how de-fringing recovers lost signals by injecting extremely faint sources into images before combining them through co-addition. For each of these experiments we will demonstrate how our method significantly improves the

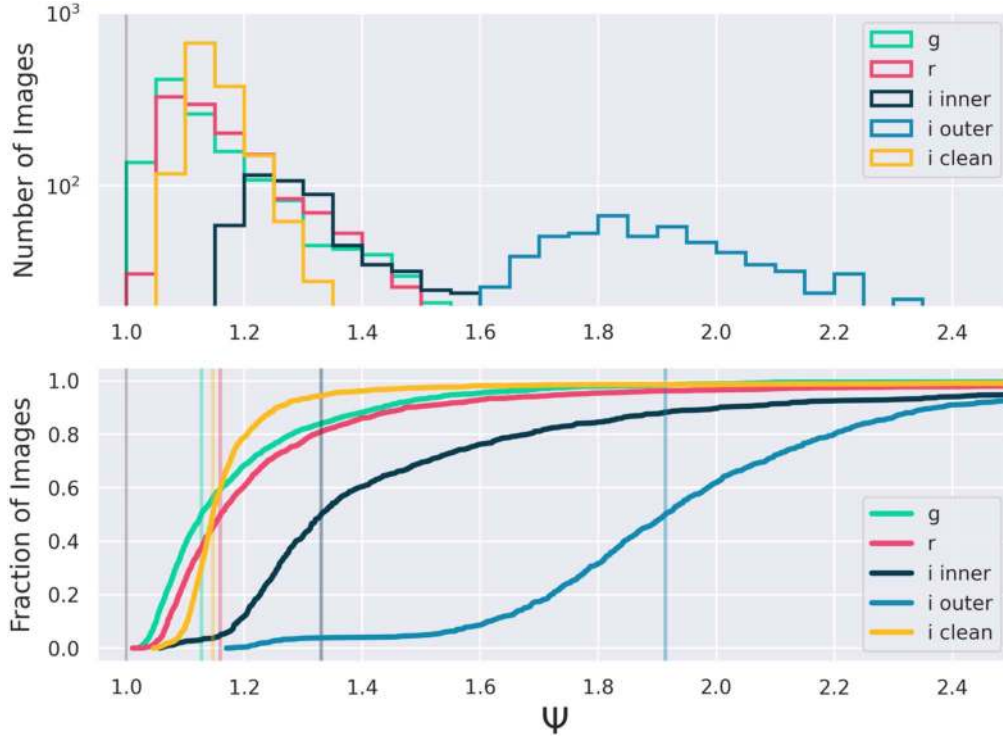


Figure 6. Our sample of *g*-band, *r*-band and *i*-band images show distinctively different distributions in their Uniform Background Indicator (Ψ) as shown in both a histogram (top) and cumulative distribution function (bottom). The *g*-band (green) and *r*-band (red) average $\Psi \approx 1.15$ indicating a small but measurable amount of correlated background noise after flat-fielding. The *i*-band is split between two populations. The images taken on the inner 32 readout channels (light blue) are only moderately affected by atmospheric fringing, averaging $\Psi = 1.33$. However the outer 32 readout channels (dark blue) are significantly affected by these fringes, with an average of $\Psi = 1.91$ and less than 20% of the images with $\Psi < 1.72$. The population of cleaned *i*-band images (yellow) is monomodal and has a median value similar to the *g*-band and *r*-band of $\Psi = 1.15$. ZTF *i*-band images processed with our method show similar amounts of correlated background noise as is present in *g*-band and *r*-band images.

photometric precision of faint source measurements affected by atmospheric fringes.

First we measured the photometric error caused by atmospheric fringes. For each of the *i*-band images in our sample, we used SExtractor (Bertin & Arnouts 1996) to generate an instrumental photometric aperture catalog of astrophysical sources. We then cross-matched the original and cleaned images with Pan-STARRS1 (PS1) (Chambers et al. 2016) *i*-band catalogs downloaded from the VizieR database using `astroquery` (Ginsburg et al. 2019) to create cross-matched catalogs. A zero-point was calculated for each image using matching sources with PS1 *i*-band magnitudes less than 17. This zero-point was used to transform the instrumental photometric catalogs into ZTF *i*-band magnitudes comparable with PS1 *i*-band magnitudes, as well as to calculate a 5σ limiting magnitude for each image. After removing images with a limiting magnitude less than 21 to ensure reasonable measurement of faint sources, our final sample size was 738 images. For each pair of fringed and cleaned images, the cross-matched catalogs were used to find the variance in the

difference between ZTF and PS1 magnitudes for stars within a magnitude bin. The photometric error caused by fringes was then calculated as:

$$\sigma_{\text{mag}} = \text{std}(m_{\text{PS1}} - m_{\text{ZTF}}) \quad (6)$$

$$\sigma_{\text{fringe}} = \sqrt{\sigma_{\text{mag,fringed}}^2 - \sigma_{\text{mag,cleaned}}^2} \quad (7)$$

The photometric error we measured due to fringing is shown in Figure 8, plotted separately for different faint magnitude bins against the UBI of the fringed images before cleaning. Images with larger UBI values have larger amounts of additional magnitude scatter relative to the PanSTARRS catalog than those images with smaller UBI values. Those images with $\Psi \leq 1.3$ appear to have marginal improvement to their photometric precision due to removing fringes. However, images with $\Psi > 1.3$ show significant photometric errors due to fringing that our method removes. Images with $\Psi = 1.75$ have a 0.21 mag error on 19.5 mag stars, 0.35 mag error on 20.5 mag stars and a 0.43 mag error on 21.5 mag stars. Images with $\Psi = 2.0$ have a 0.28 mag error on 19.5 mag

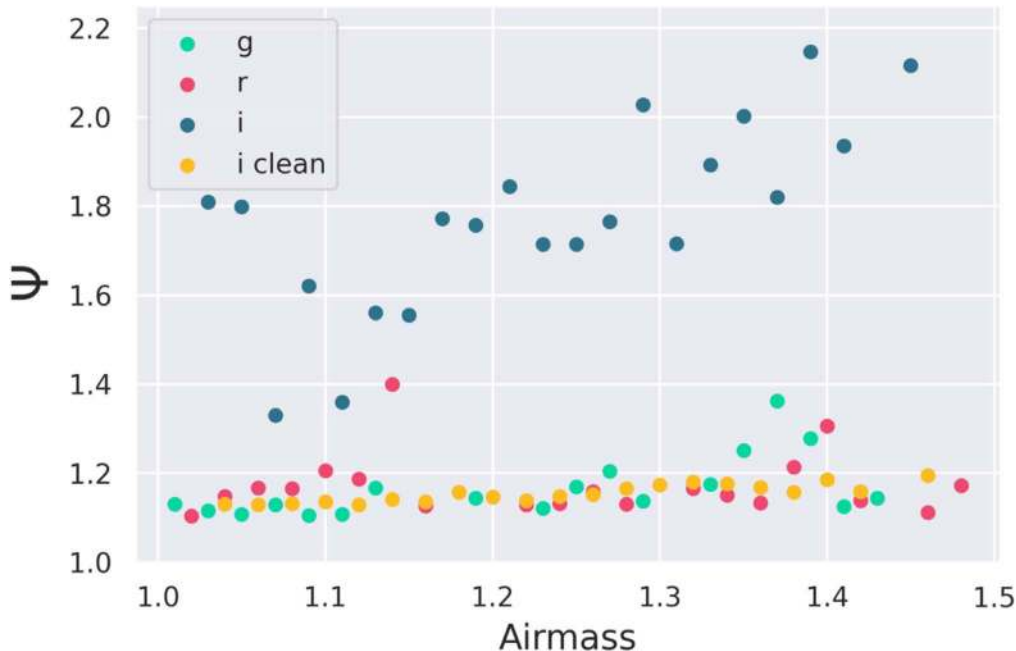


Figure 7. The Uniform Background Indicator (Ψ) correlates with airmass for *i*-band images with significant amounts of atmospheric fringes. Observing sources through the additional column of atmosphere increases the exposure to stimulated emission of atmospheric lines that causes the emergence of fringes. Cleaning the *i*-band images removes this correlation and produces a relationship with airmass indistinguishable from the *g*-band and *r*-band populations.

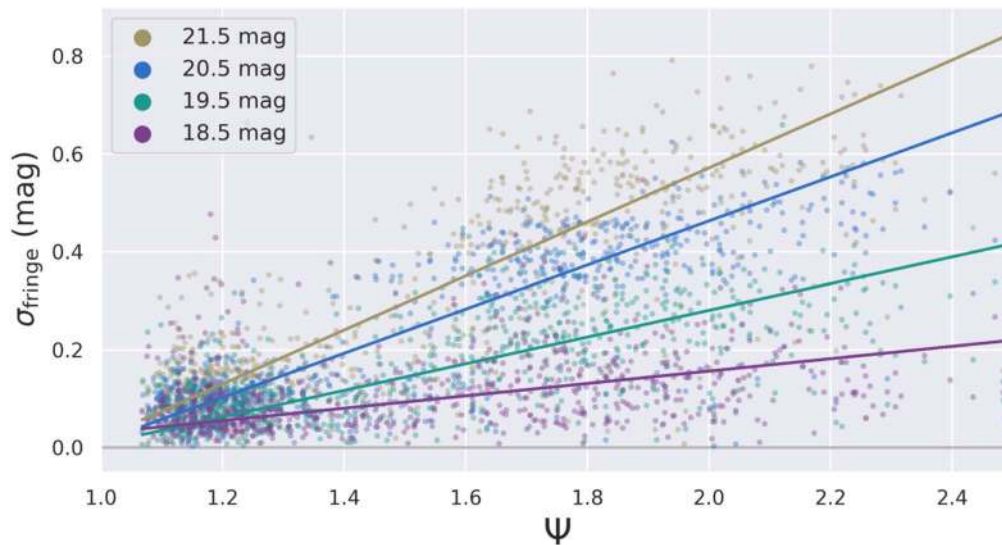


Figure 8. Measurement of the photometric error on faint sources due to fringing on 738 high-quality *i*-band images as shown against the Uniform Background Indicator (Ψ) of the fringed image before cleaning. This error is calculated by comparing the variance in the difference between ZTF and PS1 magnitudes before and after removing fringes. Larger UBI values correlate with larger amounts of photometric error, getting as large as 0.46 mag for 20.5 mag sources and 0.57 mag for 21.5 mag sources at $\Psi = 2.0$. Atmospheric fringes add a significant systematic error to the photometry of faint sources that our method is able to remove.

stars, 0.46 mag error on 20.5 mag stars and a 0.57 mag error on 21.5 mag stars. The worst 10% of *i*-band images in the outer readout channels have $\Psi \geq 2.4$, resulting in errors as large as 0.21 mag for 18.5 mag stars and up to 0.39 mag for 19.5 mag stars. Images with $\Psi \geq 2.3$ have very few sources

fainter than 19.5 mag in the figure, indicating that these images fail to pass the limiting magnitude cut. Our method enables the recovery of the faint end of the luminosity function that would otherwise be unobservable in ZTF *i*-band images.

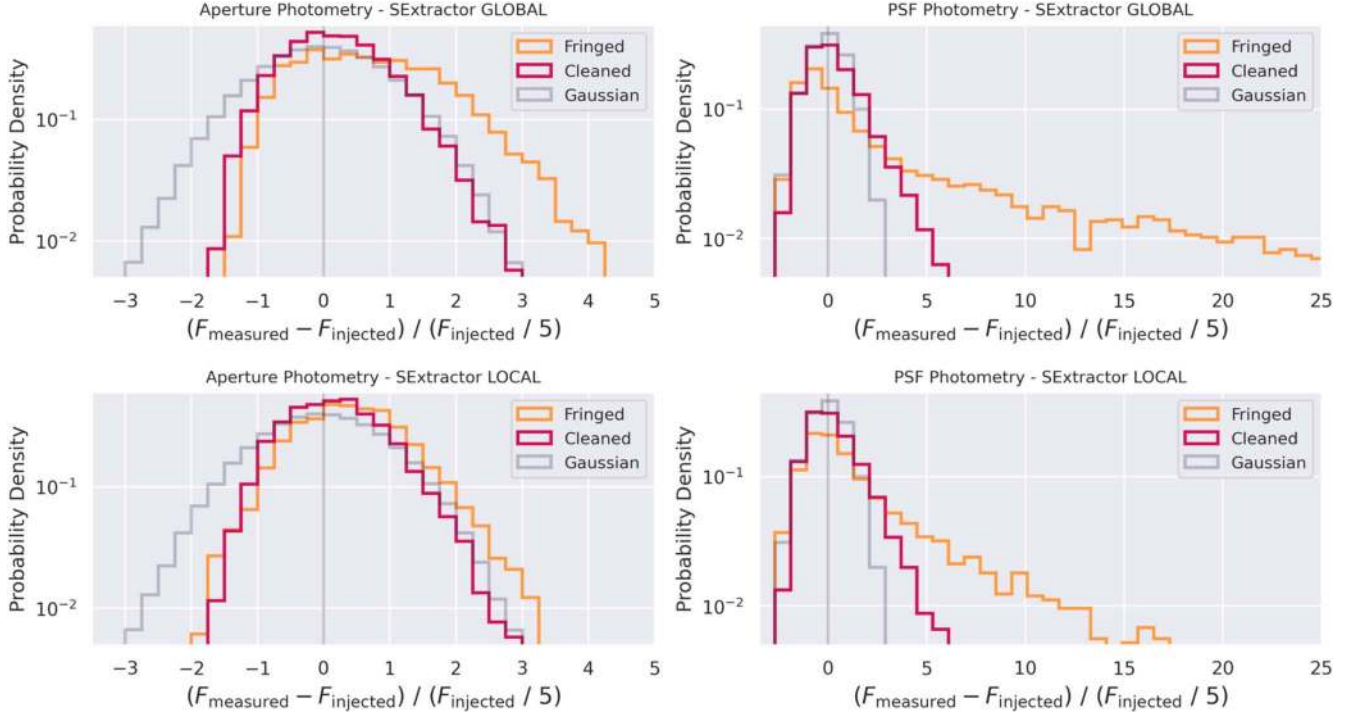


Figure 9. The measured flux of 5000 fake 5σ sources injected into an i -band image on the fringed images (orange) and after cleaning (red) using the `SExtractor GLOBAL` background setting (top), `LOCAL` background setting (bottom), aperture photometry (left) and PSF photometry (right). On each sub-figure a Gaussian of 5000 sources is drawn (gray) as a visual guide. In all cases the cleaning method significantly increases the accuracy and precision of the recovered flux, particularly for PSF photometry where attempting photometry on images with fringes can often result in overestimating the brightness of the source. Aperture photometry also overestimates the flux resulting in a deficit of lower flux detections than would be statistically expected, although to a lesser degree. If a method for removing fringes cannot be applied, it is best to use the `SExtractor LOCAL` background setting and a catalog of aperture photometry to most accurately measure the true magnitude of faint sources.

4.3. Effects on Fake Sources: Single Epoch

Another way to measure the effect of atmospheric fringes on photometric precision of faint sources is through the injection and recovery of fake sources. We selected an i -band image with $\Psi = 1.78$ as a representative image. The image’s zero-point and 5σ limiting magnitude were calculated using a cross-match to PanSTARRS1 sources as described above. A PSF for the image was derived using `psfex` (Bertin 2011) and executed using wrappers written in the `galsim` (Rowe et al. 2015) Python package. 100 sources with magnitudes equal to the 5σ limiting magnitude were injected into the original image using this PSF model at random locations, including Poisson noise. The image was then cleaned using the `fringe2` package. Photometric catalogs using both aperture and PSF photometry were calculated on the original and cleaned images using `SExtractor` with a 3 pixel 1σ detection threshold. Measurements in the catalogs at the location of the injected fake sources were then recovered for comparison to the true modeled flux. We also calculated the ideal aperture corrected flux of the injected fake sources by calculating the median ratio of aperture to PSF fluxes of high signal-to-noise astrophysical sources from the catalogs, for both fringed and cleaned images.

We then repeated this process 50 times for a total of 5000 fake sources injected to form our sample. We also duplicated this experiment with the `SExtractor BACKGROUND` parameter set to `LOCAL` and `GLOBAL` to test the effect of forcing the source identification algorithm to attempt to characterize local variations in the background noise. Figure 9 shows the results of this experiment. The distributions show the fractional offset of the measured flux to the injected flux scaled by the theoretical error on a 5σ source. Plots are drawn in log-scale to highlight the long tail of overestimated measurements for images that have not been cleaned, with a Gaussian distribution drawn in gray as a reference.

We first note some observations about the quality of recovered fake sources on the images which have not been cleaned. Setting the `SExtractor BACKGROUND` parameter to `LOCAL` improves photometric measurements on these fringed images. `GLOBAL` computes the background flux across the entire image and underestimates the amount of background underneath a source sitting on a bright fringe. These sources will have overestimated brightness and form the long tail of measurements with larger than expected fluxes. `LOCAL` computes the background flux with a rectangular annulus around the source that prevents attributing the

additional brightness to the source but instead to the background. Switching to a LOCAL background therefore reduces this tail. Aperture photometry performs better than the PSF photometry on these fringed images with a reduced tail for both SExtractor settings. This prominent tail may be due to a bias in the PSF model caused by including stars that fall on fringes, inflating the wings of the PSF model. When that model is applied to stars that fall on particularly bright fringes the increased background may be included in the measured flux, resulting in an overestimate of the flux. Aperture photometry does not suffer from this potential model bias and therefore has a less prominent tail of overestimated fluxes.

The sources from the cleaned catalogs are more accurately and precisely measured under all conditions. There is a slight tail of overestimated fluxes when performing PSF photometry but it is far closer to a Gaussian distribution than without cleaning. For both LOCAL and GLOBAL backgrounds, the distribution of recovered sources in the cleaned catalogs very nearly resembles a Gaussian distribution. There exists only a few sources with measured fluxes exceeding the injected flux causing a deficit of fainter sources. There still remains a noticeable deficit of sources at the faint end of the brightness distribution, indicating that our method is not removing all additional flux in the image background. Aperture photometry remains the best way to evaluate sources even where atmospheric fringes have been removed. Failure to clean images containing atmospheric fringes results in a systematic overestimation of the flux of faint sources. Applying our method enables near-ideal recovery of 5σ sources for a variety of measurement methods.

4.4. Effects on Fake Sources: Multi Epoch

It is particularly difficult to overcome the effects of fringes when combining multiple images of the same field to recover sources fainter than a single image's limiting magnitude. Changing atmospheric conditions and various observational airmasses will alter the strength of the fringes, while dithering and sky motion will place astrophysical sources onto slightly different pixels for each exposures. Failure to remove fringes will contribute significant excess flux to a co-addition of multiple images. We demonstrate here this effect quantitatively by injecting extremely faint sources into individual images and attempting to measure them in a co-added image.

100 *i*-band images of the same field were zero-pointed using the previously outlined method. Sources in each image were measured and the signal-to-noise versus magnitude was fit to extrapolate the 0.5σ magnitude of all images. Each image was injected with 100 sources at this 0.5σ magnitude at a common list of coordinates in R.A. and decl. These sources would be expected to appear as 5σ sources after combining the 100 *i*-band images because signal-to-noise increases as the square root of the total exposure time. All images were then cleaned

with the *fringe* package. SCAMP (Bertin 2006) was run to find astrometric projection parameters for each of the images such that they could be transformed onto a common reference frame. All fringed and cleaned images were then combined into two separate co-additions using SWarp (Bertin et al. 2002), produced with a median combination filter and background subtraction. Aperture photometry catalogs were then generated on the final co-additions using SExtractor with the LOCAL background setting. Sources at the locations of the injected signals were recovered and their signal-to-noise measured as the ratio of their aperture flux to their locally determined aperture flux error. We repeated this process 50 times for a total of 5000 fake sources distributed over 50 fringed and 50 cleaned co-additions.

Figure 10 shows the quantitative and qualitative photometric improvements to these recovered sources. Sources observed on fringed co-additions peak at a signal to noise of 3, pushing them below the typical 5σ observable threshold. These sources also have a long tail of excessive flux stretching as high as 10σ due to the additional flux caused by atmospheric fringes. Sources observed on cleaned co-additions are much closer to their theoretical distribution. These sources peak at exactly 5σ and the large majority of sources are within the normal signal-to-noise range of 4–6. There still exists a tail of excessive fluxes, demonstrating that our method is failing to remove all excessive correlated background from the individual exposures. However our method clearly produces an observed flux distribution that is much closer to a Gaussian distribution.

The largest difference between the two populations appears in their yields. Of the 5000 sources injected into each of the fringed and cleaned images, 1161 sources were recovered from cleaned co-additions and only 156 sources were recovered from fringed co-additions. Over 95% of the 0.5σ sources injected into the fringed images failed to be recovered after co-addition. The order of magnitude increase in the number of sources recovered by our method enables ZTF *i*-band surveys to recover faint sources that would otherwise have been extremely unlikely to observe. We note that forced photometry at the known locations of the fake sources may have increased the recovered yields for both populations but would not be an accurate representation of the observation process undertaken for unknown sources.

Our analysis demonstrates the power of our method to remove photometric error due to atmospheric fringes and enable the recovery of faint sources in both single images and co-additions that would otherwise have been undetectable. Failure to implement a method for removing atmospheric fringes greatly reduces the effectiveness of the *i*-band filter for observing any sources fainter than 18th magnitude. Our method increases the photometric precision of the *i*-band to that of *g*-band and *r*-band images that do not suffer from atmospheric fringes.

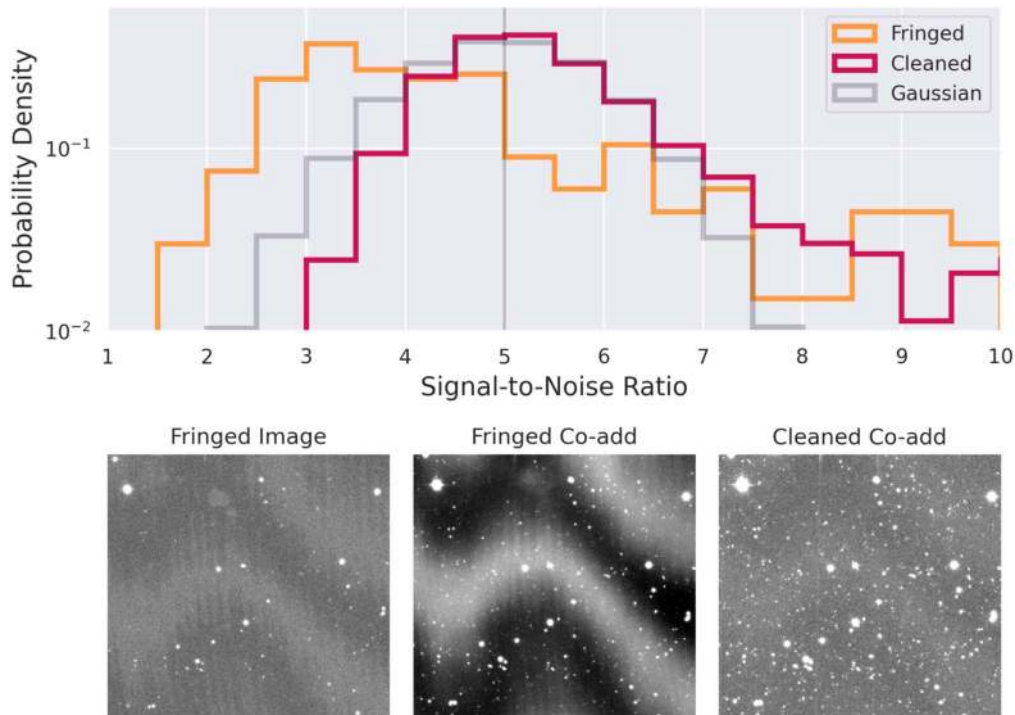


Figure 10. The signal-to-noise distribution of 5000 0.5σ sources injected into 100 images after a median co-addition (top) for images with atmospheric fringes (orange) and after cleaning (red). The photometric catalogs were generated with a `SExtractor LOCAL` background and aperture photometry. 1161 sources were recovered in the cleaned images and only 156 sources were recovered in the fringed images, demonstrating the necessity to clean *i*-band images in any attempt to find faint sources after multi-epoch co-addition. Those sources that were recovered in the fringed images are most likely to be detected at an artificially low signal-to-noise, with a long tail of higher signal-to-noise due to falling on positive fringes. The cleaned images have a distribution much closer to Gaussian, with a small deficit at the low signal-to-noise end that also appears as a tail at the higher end. The photometric improvement in multi-epoch co-additions can be clearly seen in as a smoothly varying background and the presence of faint sources after cleaning has been applied (bottom).

5. Discussion

Our method has several benefits beyond improving the photometric quality of individual *i*-band images. Image references are constructed through the combination of as many as 40 epochs taken at a single field. As shown in Section 4.4, removing fringes significantly improves the photometric precision on a multi-epoch co-addition. Application of our method also greatly improves the quality of the ZTF alert stream. The alert stream packets are generated on difference images created from subtracting individual epochs from a multi-epoch co-addition reference (after appropriate scaling). Correctly removing fringes makes visible the faint end of the luminosity function that would otherwise not be observable in the *i*-band alert stream.

Future work on removing atmospheric fringes using PCA could improve upon our method in several ways. Training on CCDs instead of readout channels could produce better eigen-images by including more correlated pixels in the PCA feature identification. Also, different methods of calculating the local flux error such as calculating the root-mean-square error on a smaller (or larger) square around each pixel could produce a less noisy UBI measurement. Lastly, investigating the

eigenvalues for fringe bias images generated on different readout channels from the same exposure could reveal correlations that could be used to perturb the fringe bias into a better fit for each readout channel. Treating each readout channel as entirely independent, while convenient and a natural fit for the IPAC processing pipeline, may leave out valuable information that could improve our method.

The eigen-images shown in Figure 4, as well as the eigen-images of many of the other fringe models, show significant smooth variations that are not being removed by the current flat-fielding pipeline. This indicates that the fringe bias images include not only atmospheric fringes, but residual global gradients. Future work could be done on exploring the application of this PCA method to supplement or even replace the current flat-fielding pipeline on not only the *i*-band images, but *g*-band and *r*-band images as well. The Uniform Background Indicator can be used as a quantitative measurement to compare how well a PCA method, as compared to more classical flat-fielding methods, generates astronomical images with normal backgrounds.

This work is based on observations obtained with the Samuel Oschin Telescope 48 inch and the 60 inch Telescope at the

Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

We acknowledge support from the University of California Office of the President for the UC Laboratory Fees Research Program In-Residence Graduate Fellowship (Grant ID: LGF-19-600357). This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We acknowledge support from the DOE under grant DE-AC02-05CH11231, Analytical Modeling for Extreme-Scale Computing Environments. M.W.C. acknowledges support from the National Science Foundation with grant No. PHY-2010970.

ORCID iDs

Michael S. Medford,  <https://orcid.org/0000-0002-7226-0659>

Peter Nugent  <https://orcid.org/0000-0002-3389-0586>

Danny Goldstein  <https://orcid.org/0000-0003-3461-8661>

Frank J. Masci  <https://orcid.org/0000-0002-8532-9395>

Igor Andreoni  <https://orcid.org/0000-0002-8977-1498>

Michael W. Coughlin  <https://orcid.org/0000-0002-8262-2924>

Dmitry A. Duev  <https://orcid.org/0000-0001-5060-8733>

Ashish A. Mahabal,  <https://orcid.org/0000-0003-2242-0244>

Reed L. Riddle  <https://orcid.org/0000-0002-0387-370X>

References

- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019a, *PASP*, **131**, 068003
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019b, *PASP*, **131**, 018002
- Bernstein, G. M., Abbott, T. M. C., Desai, S., et al. 2017, *PASP*, **129**, 114502
- Bertin, E. 2006, in ASP Conf. Ser. 351, *Astronomical Data Analysis Software and Systems XV*, ed. C. Gabriel et al. (San Francisco, CA: ASP), 112
- Bertin, E. 2011, in ASP Conf. Ser. 442, *Astronomical Data Analysis Software and Systems XX*, ed. I. N. Evans et al. (San Francisco, CA: ASP), 435
- Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393
- Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in ASP Conf. Ser. 281, *Astronomical Data Analysis Software and Systems XI*, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco, CA: ASP), 228
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
- Dekany, R., Smith, R. M., Riddle, R., et al. 2020, *PASP*, **132**, 038001
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, **696**, 870
- Ginsburg, A., Sipőcz, B. M., Brasseur, C. E., et al. 2019, *AJ*, **157**, 98
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, **131**, 078001
- Halko, N., Martinsson, P., & Tropp, J. 2011, *SIAMR*, **53**, 217
- Jolliffe, I. T., & Cadima, J. 2016, *RSPTA*, **374**, 20150202
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, **121**, 1395
- Martinsson, P.-G., Rokhlin, V., & Tygert, M. 2011, *Appl. Comput. Harmon. Anal.*, **30**, 47
- Masci, F. J., Laher, R. R., Rusholme, B., et al. 2018, *PASP*, **131**, 018003
- Michael, M. 2021, fringeZ: Correct Atmospheric Fringes in ZTF i-Band Images, Zenodo, doi:10.5281/zenodo.3404081
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2018, *PASP*, **131**, 018001
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Ren, W., Cao, Q., Wu, D., et al. 2017, *OExpr*, **25**, 17039
- Rowe, B. T. P., Jarvis, M., Mandelbaum, R., et al. 2015, *A&C*, **10**, 121