Zubercal

Version 1.2.1, Feb 2025

A.J. Drake, F.J. Masci, R.M. Smith, N. Regnault, M. Rigault, P. Rosnet, B. Racine, F. Feinstein, A. Wold, J. Purdum, R.R. Laher, B. Rusholme, T.X. Chen, R. Paladini, M.J. Graham, S.R. Kulkarni and members of the ZTF Calibration Working Group

Table of Contents

1. What is Zubercal?

Zubercal (ZTF-ubercal) is a new set of photometry based on a detailed recalibration of ZTF science image-based PSF photometry. This set includes all ZTF detections from data release 20 with matches to PS1 sources (< 1.2") and colour information. Zubercal photometry is calibrated to the PS1 photometric system.

2. Why was ZTF PSF photometry recalibrated?

The ZTF project aims to provide photometry calibrated to 1% (or better), to optimize the scientific return for current and future time-domain projects. However, it was discovered that the accuracy of the zero-point (ZP) from the original calibration process varies depending on the number of calibrator sources within a field. Sparse fields have poorer calibration than dense fields. Thus, increasing the number of calibrators, by calibrating multiple fields simultaneously, could potentially reduce these ZP uncertainties.

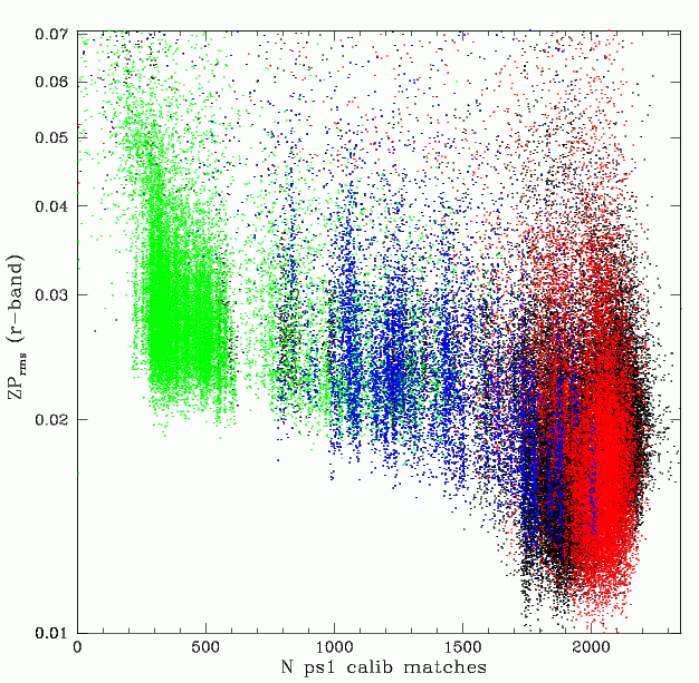

Figure 1. Uncertainty in ZP vs number of PS1 calibrators. Black points (|b| < 10 deg), red (10 < |b| < 20), green (|b| > 20), blue (field 539, l=44.3, b=3.6 deg).

To demonstrate the importance of the number of calibrators, in Figure 1, we plot ZP uncertainties vs the number of PS1 calibrators across several fields. The selections show how the varying stellar density affects the uncertainties in the current pipeline photometry.

The calibration equations used in the current calibration of ZTF photometry to PS1 are as follows:

$$m_{cal} = m_{inst} + ZP_f + c_f\left(m_1^{PS1} - m_2^{PS1}\right),$$

thus

$$g^{PS1} - g = ZP_g + c_g(g^{PS1} - r^{PS1})$$
$$r^{PS1} - r = ZP_r + c_r(g^{PS1} - r^{PS1})$$
$$i^{PS1} - i = ZP_i + c_i(r^{PS1} - i^{PS1})$$

Thus, the current ZTF calibration process includes only the instrumental magnitude, a single ZP and a colour term. Although calibration is easily obtained by applying a zero-point and a colour term (Masci et al. 2018, S3.5.2), this does not fully account for variations in the telescope/hardware and observing conditions over time. Additionally, differences between ZTF and PS1 filter transmission functions complicate transformation between ZTF and PS1. Colour terms are required for both the filter systems and the colour-dependent airmass term. Thus, a more complete equation is required to account for the two separate dependencies.

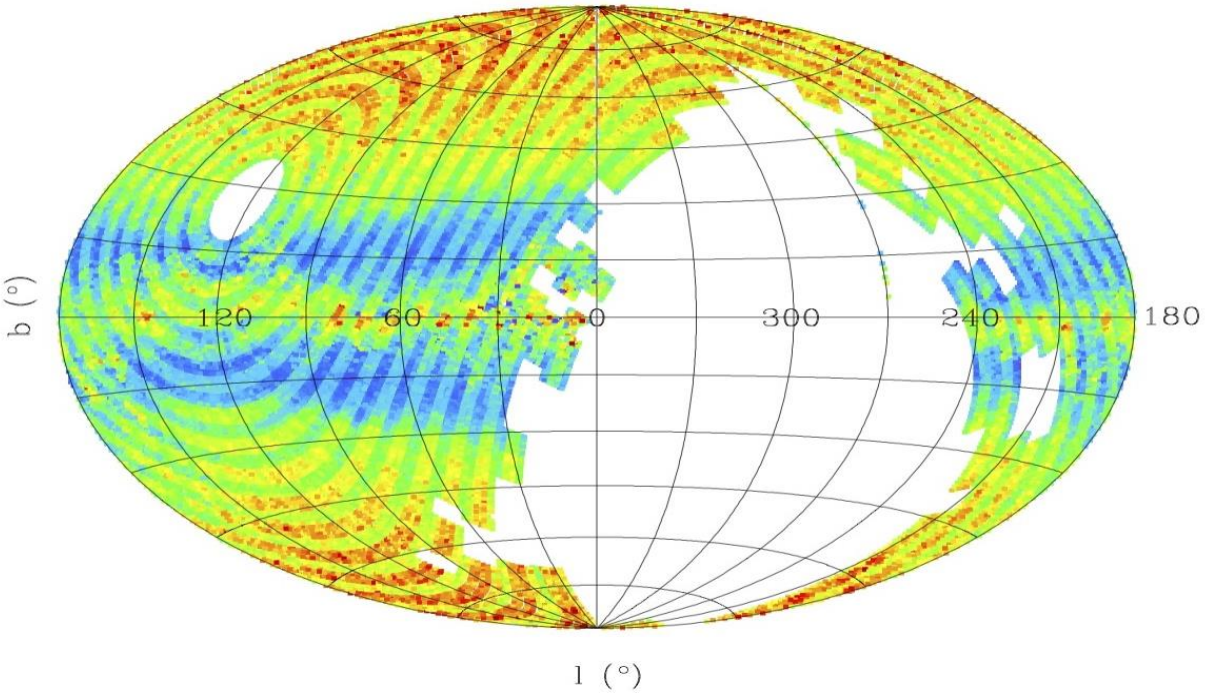

Figure 2. The variation in ZTF r-band colour coefficients across the sky. Median ZTF colour coefficients range from 0.07 (blue pixels) to 0.13 (red pixels). The banding structure is due to varying response of the ZTF CCDs.

Furthermore, since the PS1 and ZTF systems filters differ, the amount of dust induced extinction must also vary. In Figure 2, we plot the median colour coefficient vs Galactic coords. The observed structure is due to the combination of differences in reddening as well as variations in the coatings of the ZTF CCDs. The central eight ZTF CCDs have two layers of anti-reflection coating while the outer CCDs have one. The differences in the transmission functions of these CCDs cause the stripes. The ZTF CCDs also come from different batches with varying sensitivity (e.g., quantum efficiency, QE).

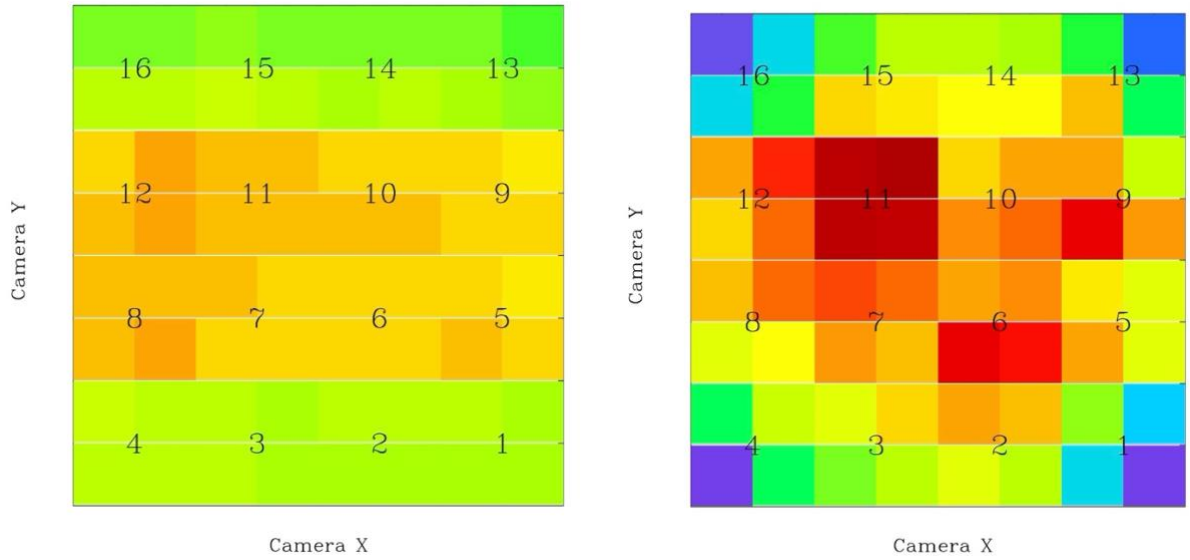

Figure 3. Variations in sensitivity across the 16 CCDs of the ZTF camera. Left: Median differences in colour coeffs with ZTF CCDs (-0.014 (green) to 0.014 (orange)). Right: variations in median ZP (0.2 (red) to -0.2 mags (blue))

To demonstrate these differences, in Figure 3 we show the variations in the median r-band colour coefficients and ZPs for each of the 16 ZTF CCD cameras and 64 quadrants (each camera is read via four quadrants). We also show the difference in average ZP. Here we see that CCD 11 has an average ZP that is 0.4 mags higher than the outer corners. However, this difference is exaggerated by vignetting which causes the corner CCDs/quadrants to receive less light.

2.1 Systematics in ZTF PSF photometry.

An investigation of the initial ZTF photometry revealed complex systematic photometric residuals of up to 5%. These systematics include spatial, chromatic, temporal, and source brightness-based dependencies.

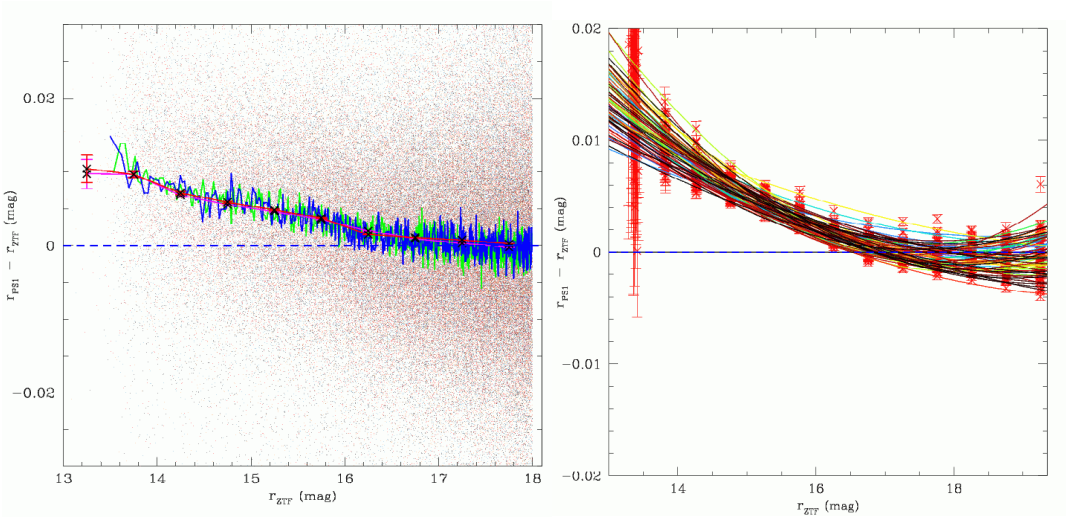

Figure 4. ZTF photometric magnitude biases. Left: differences between ZTF and PS1 magnitudes vs source brightness for one field. Right: quadratic fits of the magnitude biases for all ZTF fields.

The first clear bias we discovered in ZTF photometry was a source brightness bias. This showed a slight (1-2%) difference between bright stars observed in PS1 and ZTF. This may be related to the brighter-fatter effect (e.g. Astier & Regnault 2022). In Figure 4, we show the scale of the biases present within the original ZTF photometry. However, plot suggests that the effect varies with the ZTF field. It is also possible that this bias could be due to uncertainties in PS1 rather than ZTF. Nevertheless, corrections are applied to all ZTF processing to remove this field-dependent magnitude bias effect.

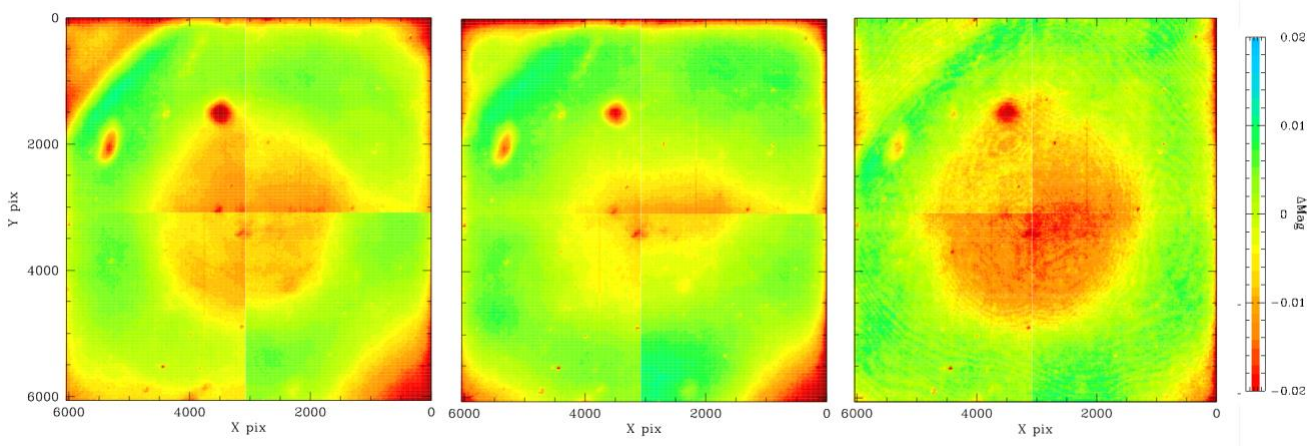

Figure 5. Spatial systematics in g, r and i-band for ZTF CCD 16.

The second type of bias discovered within ZTF photometry was spatial offsets between calibrated ZTF magnitudes and PS1 magnitudes. In Figure 5, we present the spatial photometric residuals determined by calibrating all the ZTF-I photometry to PS1. Here, residuals from across each CCD are spatially binned into groups of super-pixels. The features on the edges of the CCDs are due to glints off lens holders. The round regions are due to dust spots on the CCD. The large circular shape in each CCD is due to the radially varying thickness profile of the CCD. The fingerprint-like feature in the *i*-band residual map is due to fringes caused by internal interference. The differences in the spatial structure between the filters also demonstrate a colour dependence in the residuals.

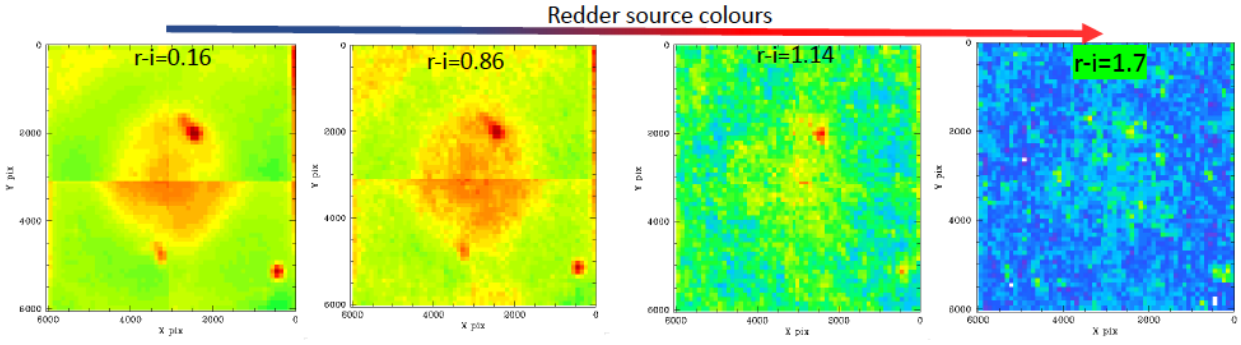

Figure 6. Spatial systematics in CCD6 g-band for sources divided by g-i colour. Each image represents sources in a 0.14 mag r-i colour bin from 0.16 to 1.7. The colours represent residual values from 0.05 (dark red) to -0.04 mags (purple).

Subsequent analysis led to the discovery that the residual structure varies with the colour of the sources involved even within a given ZTF filter. In Figure 6, we show this variation for objects with differing g-i colours. This shows that one needs to correct not only for location but also the colour of a source. In fact, the source colour dependence varies between dust spots, CCD thickness variations and edge glints.

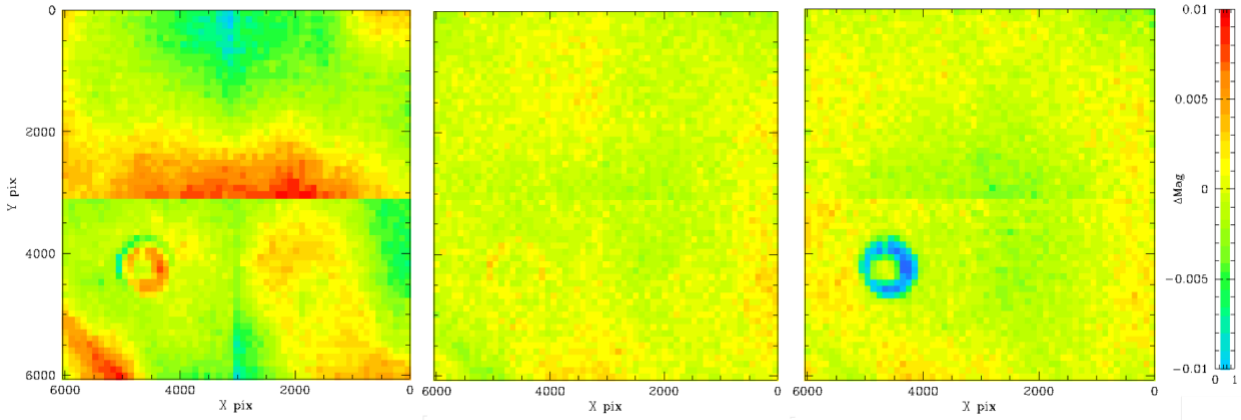

Figure 7. Variation in spatial systematics in CCD6 g-band for sources over time. 2018-06-16, 2019-04-17, 2020-09-28**.**

Since the sensitivity of instruments change over time, it is expected that there will be changes in the spatial residuals over time. In Figure 7, we show the difference between the spatial systematics measured on specific dates and the median systematic structure. As expected, the spatial structure also varies with time. In this example, we can see that a ring due to a dust spot was removed in the later observations.

2.2 Improving ZTF Coverage.

ZTF data release photometry is only available from IRSA for the areas of ZTF fields covered by reference images. Due to pointing variations, areas on the edges of images often fall outside of the reference frame region and are *not included in current data releases*. Additionally, in the *i*-band, ZTF reference frames cover only 70% of the area observed by ZTF (see DR20 release notes). Zubercal does not require a ZTF reference frame. Thus, it includes more photometry on the edges of ZTF images (where there are CCD chip gaps), or areas without any reference frames.

2.3 Complete Source Calibration.

Currently, all ZTF PSF photometry publicly available from **IRSA interfaces are only partially calibrated**. Astronomers downloading ZTF photometry must perform positional matches between ZTF and PS1 to obtain the PS1 g,r,i photometry required to calibrate each ZTF data point. Additionally, since the current calibration is performed on a field + quadrant (1/64th of the focal plane) basis, variations in telescope pointing and field grid lead to multiple (four or more) separately labelled lightcurves for the same object (for each photometric band). Zubercal provides a single fully calibrated lightcurve for each ZTF filter for all PS1 sources detected.

2.4 Improved Photometric Uncertainties.

Accurate uncertainties are necessary to assess the significance of data. However, the photometric uncertainties provided by current ZTF pipeline are based on fits to magnitude vs object uncertainties combined across all observations. The ZTF pipeline calibration assumes the same uncertainty for **all detections of a given apparent magnitude within a quadrant**. Since the photometric uncertainties vary with instrumental noise and source flux (due to varying cloud, seeing, airmass, etc.), the uncertainties in the regular ZTF data releases can be *underestimates* or *overestimates*. In contrast, uncertainties in Zubercal are estimated based on the measurements from individual observations.

2.5 Improved Calibration for Red Sources.

As noted above, the current ZTF pipeline calibration does not account for the physical throughput differences between the ZTF and PS1 filter systems. Correcting for these differences requires a more complex calibration method. For example, colour degeneracies of stars within the PS1 system cause ZTF r-band measurements of red sources (g-r > 1.2) to be very poorly calibrated within the current photometry.

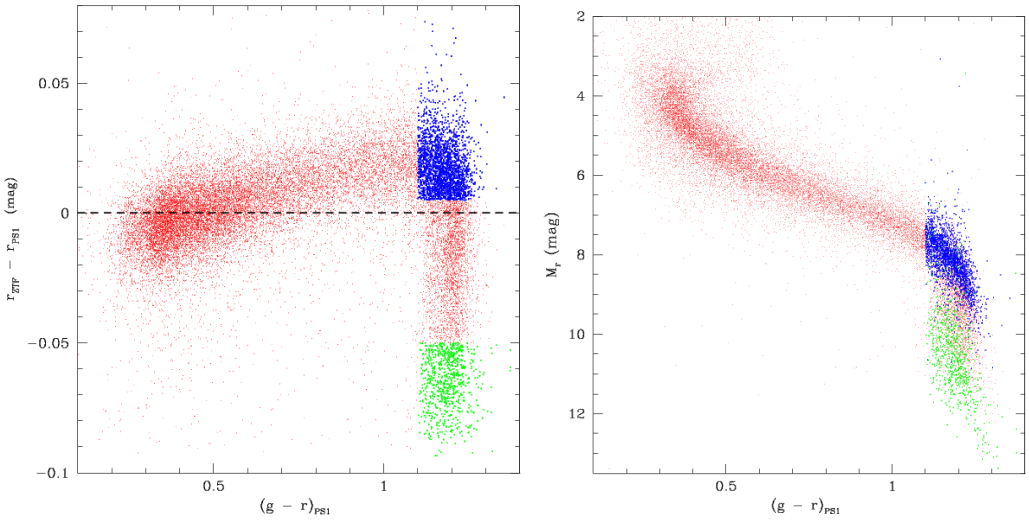

Figure 8. An example of the systematic bias in calibrated ZTF r-band photometry relative to PS1 magnitudes. Stars with varying absolute magnitude presented in the right panel have the same colour range in PS1, but varying ZTF g-r colours, resulting in systematic differences.

In Figure 8, we present the photometric residuals for ZTF r-band photometry calibrated to PS1 using the PS1 colours. Here we select two sets of stars with g-r > 1.2 that exhibit significant systematic differences between ZTF and PS1. Green stars with *delta mag* > 0.05 mags and blue with *delta mag* < -0.01. These objects have the same PS1 g-r colour range but different absolute magnitudes based on Gaia distances. The g-r colours vary due to differences between the PS1 and ZTF filter sets. This PS1 colour degeneracy means that the r-band ZTF photometry for **red sources (g-r > 1.2) must be calibrated using different PS1 colour combinations** (e.g., g-i or r-i).

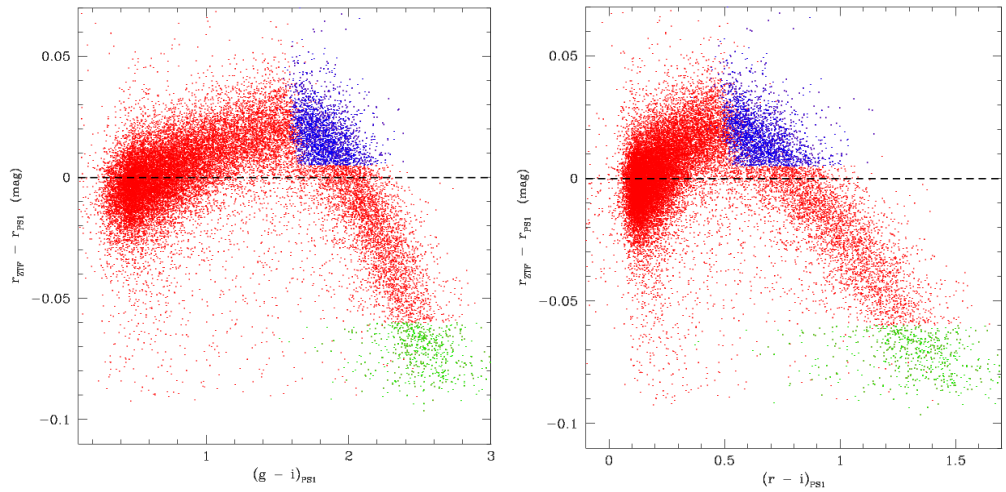

Figure 9. Variations in current ZTF r-band photometric systematics from Figure 8 when plotted with the PS1(g-i) and (r-i) colours of the stars**.**

In Figure 9, we show the same star selections as the prior plot vs their PS1 g-i and r-i colours. We find that calibration with PS1 g-i or r-i colours removes the g-r calibration degeneracy. However, we also note that the colour terms required for the calibration are clearly non-linear.
As noted above, differences between ZTF and PS1 also make it necessary to consider differences in reddening. This is particularly important in the high-extinction regions near the Galactic plane.

2.6 Photometric biases due to incomplete readout: *the pocket effect*

Throughout the ZTF survey, several changes have been made to improve the ZTF image quality by changing the waveform used to read the charge from the 16 ZTF CCDs. A change made on 2019-10-22 had the side effect of causing the charge to be incompletely read. A small amount of charge remains in each pixel *(a pocket effect)*. This scale of this effect varies between the CCDs but is particularly evident when the background level is low.

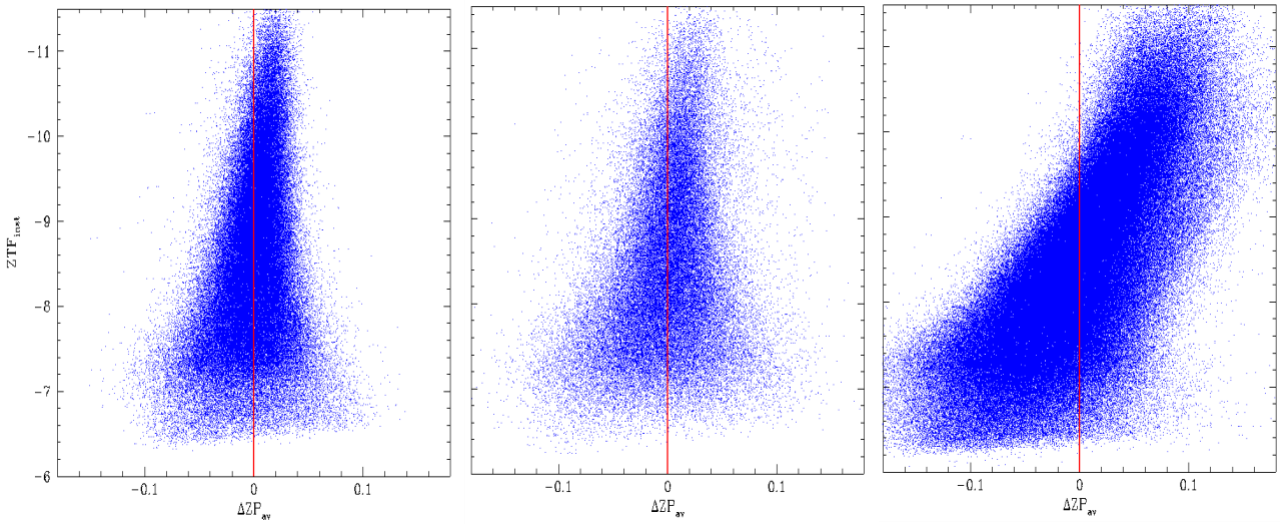

Figure 10. Examples of the variation in the readout linearity for ZTF CCD-6 photometry on three different nights. Large slope corresponds to lower background flux. Left: 2019-03-30, middle 2020-03-06, right 2020-02-24.

To illustrate the scale of the pocket effect on the photometry, in Figure 10, we present photometry from a single ZTF field observed on three nights. In the left plot, we see that photometry is moderately well calibrated using the current system, with only a slight magnitude bias (as noted above). However, on the other two nights, a low background level produces a strong pocket effect. This produces a trend in brightness that cannot be calibrated with a constant ZP and colour term.

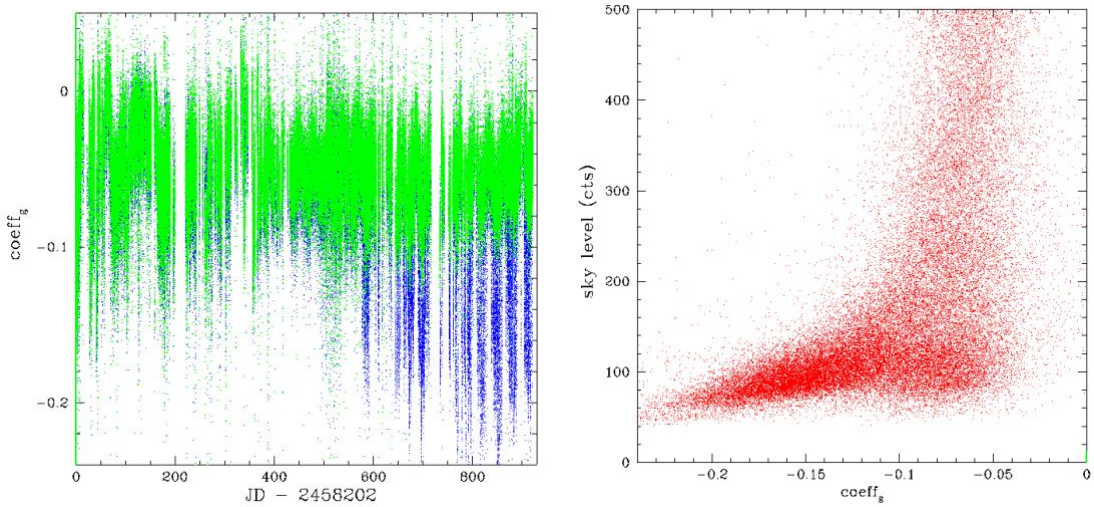

Figure 11. Values of ZTF calibration colour coefficients. Left: variation in pipeline ZTF calibration coefficients over time. Green values from CCD-7 and blue from CCD-6. Right: variations in the calibration colour coefficients with sky background level.

In Figure 11, we demonstrate when the problematic photometry began. Here, we show that the range of colour coefficients suddenly increased CCD-6 when the CCD readout waveforms were changed in Oct 2019. In the right plot, we see that a spur of low colour coefficients occurs when the sky level in the images is low. This pocket of unread charge disappears when the pixels have sufficient counts from the sky background to fill it. Zubercal calibration includes the sky level within the fit to reduce this effect. Nevertheless, a far better solution would have been to adjust the waveforms to mitigate the problem.

2.7 Correcting ZTF Photometry for Transparency.

All ground-based images are affected by atmospheric extinction. Optical observations are mainly affected by Rayleigh scattering, ozone, aerosols, and water vapour. The most obvious effect of the atmosphere in images is the decrease in transmission with an increase in airmass. Atmospheric transmission can be far more complicated in the presence of particulates from ash or water vapour in clouds.

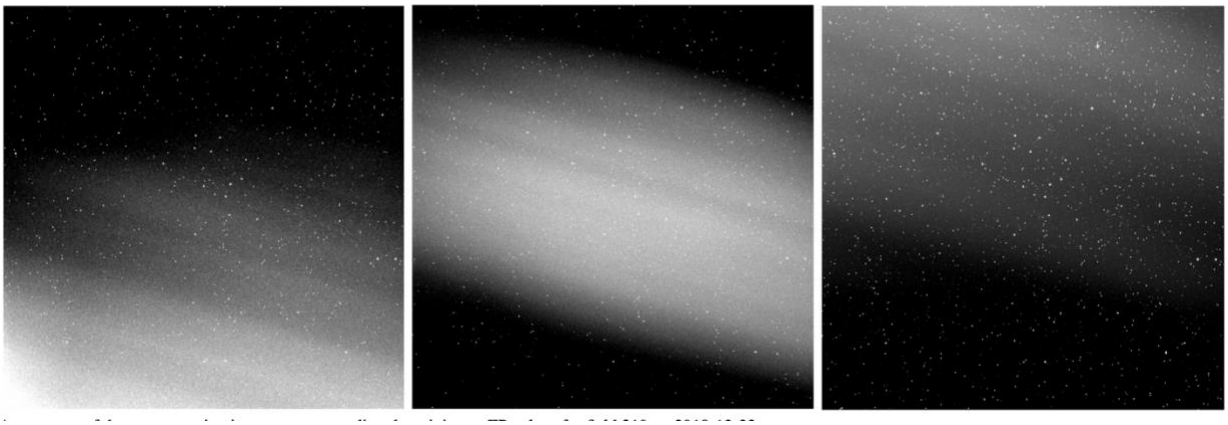

Figure 12. Variation in background flux between three ZTF frames taken over two minutes**.**

Photometric calibration is complicated by the fact that the extent and structure of features, such as clouds, can vary over short periods. In Figure 12, we show an example of how much cloud cover can change within two minutes. These images show that clouds can create varying absorption structures on spatial scales less than a single ZTF quadrant. Such variations necessitate spatial corrections to the individual zero points used in the pipeline calibration process.

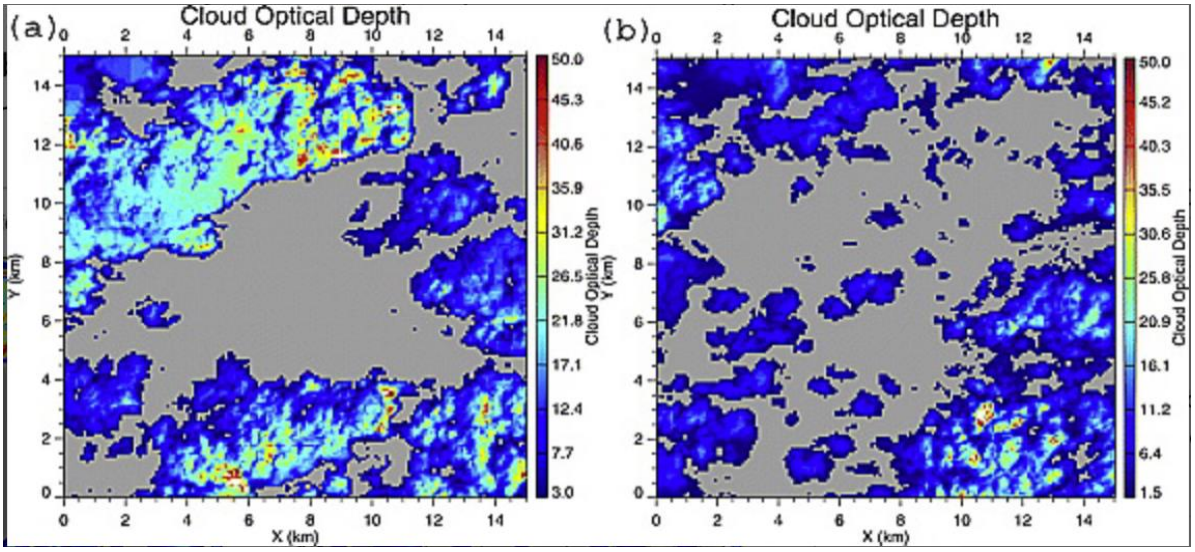

Figure 13. The optical depth of thick (av. depth 14) cumulus clouds vs thin cumulus clouds (av. depth 7) from Wen et al. (2008).

In Figure 13, we show examples of the optical depth of two different types of clouds. The extinction caused by clouds can cause significant variations in attenuation across an image. From this plot, we see that variations in optical depth of up to a factor of ~40 can occur across 0.5km. Assuming the average Palomar wind speed of 11km/hr, such changes can occur in ~3 mins. In these extreme cases, ZTF photometry would vary by **4 magnitudes in ZP over ~5 ZTF exposures**.

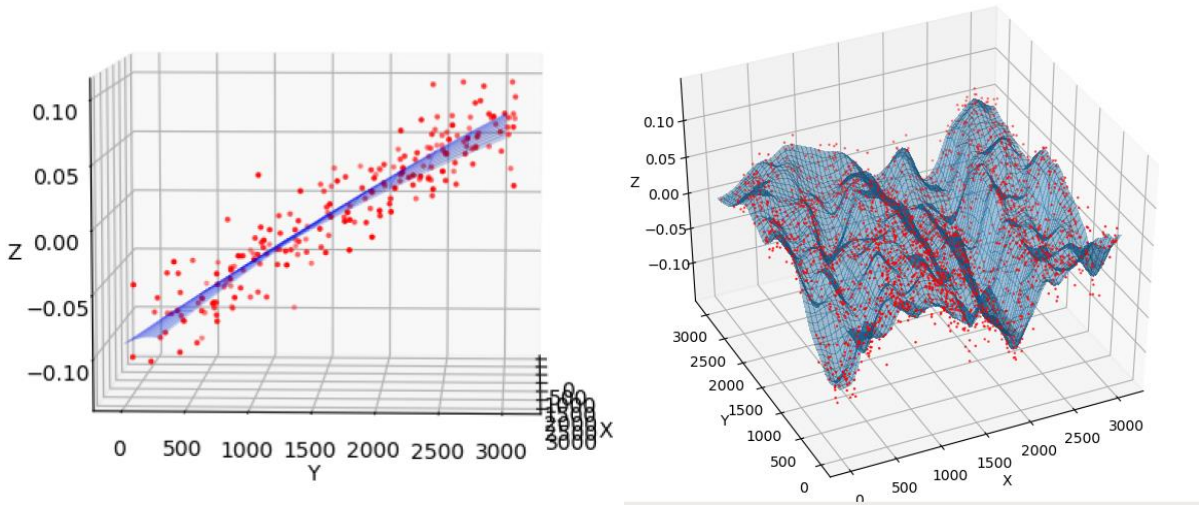

Figure 14. Two examples of ZTF spatial photometric residuals due to attenuation by clouds**.**

By using hundreds of bright ZTF stars calibrated to PS1, we investigated the structure of the cloud attenuation with ZTF images. In Figure 14, we show two examples of the spatial variations in photometric residuals due to clouds. The left plot shows a simple situation where the amount of extinction across an image simply increases from one side to the other. In the right plot, we show an example where the amount of attenuation is very complex. To correct such complex spatial attenuation due to clouds, we fit a series of polynomial surfaces and a radial basis function to every ZTF frame. We then pick the best fit based on the Akaike information criterion (AIC) and the reduced Chi^2 value of the fits. The complexity of these Zubercal surface fits is limited by the number of bright sources available within each image. Thus, heavily extincted images may still have poor corrections for spatial variations in the amount of attenuation.

3. The Zubercal Process.

The first step in the Zubercal process consisted of the bulk calibration of all bright stars (mag < 18) in ZTF-I PSF photometric catalogs using the original calibration equations. Next, all of the calibrated photometry were combined and binned to create high-resolution maps of spatial systematics for each of the 16 ZTF CCDs, in each of the three ZTF filters. The super-pixel bins in each map are the averages of the sigma-clipped photometric residuals at that image location. Maps with varying spatial scales were investigated. The final super-pixel scale was chosen to have sufficient resolution and signal-to-noise to correct for variations at the milli-magnitude level (such as radial fringing residuals due to varying CCD thickness variations). These maps are used to correct for the average of systematic effects, such as reflections along image edges. However, since such effects can vary with time, additional corrections are required.

The second step in the Zubercal process was to determine the spatial variations in the ZTF camera's chromatic response. The ZTF camera exhibits strong chromatic effects (due to CCD coating variations), radial (thickness-based) response variations and the differing response of dust spot regions.

As with all telescopes, the Palomar 48in used by ZTF exhibits temporal changes in throughput over time. The presence of dust accumulating on optics causes temporal ZP changes. Additionally, there have been changes to the CCD gain, and changes and patterns in the readout noise (due to interference), as well as changes to the readout to remove problems due to readout saturation and related artifacts. To mitigate these issues, we fit the temporal variations in spatial response.

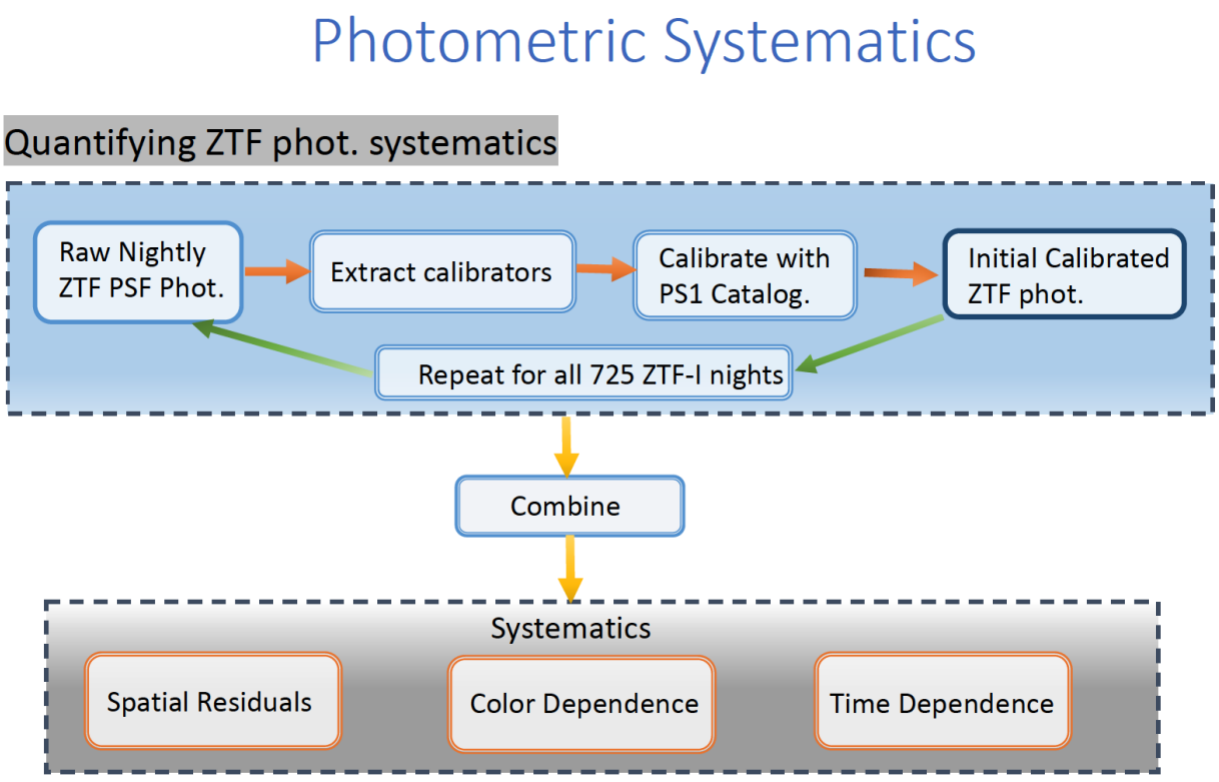

Figure 15. Steps for quantifying the systematics in ZTF pipeline processing photometry.

In Figure 15, we show how all of the ZTF-I data photometry was combined to determine each of the systematic effects within the current pipeline photometry.

3.1 The Zubercal Calibration Equation.

Once the spatial structure has been removed from the input catalogs, the next step is to perform a high-order fit including additional fit terms that account for the dependence of ZTF photometry on PSF shape, airmass, Galactic extinction, and sky level. This process is performed by simultaneously calibrating all PSF photometry taken within an individual quadrant (1/64th full FoV) and filter on a single night. This combined calibration allows us to increase the complexity of the calibration fit, taking advantage of the improved calibration possible with a large number calibrator stars. The calibration is limited to individual image quadrants, rather than calibrating the full field or CCD, to avoid the systematic ZP variations between adjacent CCD quadrants caused by readout gain variations and the flatfield processing (that is performed at the quadrant level).

$$m_{cal} = m_{inst} + ZP_f + c_1(m_1 - m_2) + c_2 X + c_3 m_{inst}^2 + c_4 \mathrm{chi} + c_5 \mathrm{sharp} + c_6 E[i] + c_7 S[i]_g$$

Airmass

Non-linearity

PSF shape dependence

Extinction [frame]

Skylevel dependence [frame]

This fit is performed iteratively to include ZP offsets between exposures. Once estimates of ZP offsets are determined, sigma rejection criteria are applied, and the photometric data is refit.

Corrections are applied to the calibrated data to remove the spatial chromatic variations and to correct the temporal variation in the spatial residuals.

The photometric residuals in each quadrant and observation are fit with a sequence of extinction surface models where the best fit is selected based on the AIC as noted above. These fits remove much of the temporally and spatially varying atmospheric extinction.

The final step is to determine the photometric uncertainties. We determine this by fitting the distribution of photometric residuals between the calibrated ZTF photometry and PS1 for each observation. We adopt uncertainties from the current calibration when the uncertainties from individual fits are unreliable. As with the current ZTF pipeline calibration, we do not keep the individual DAOPhot uncertainties. This means that the uncertainties on individual measurements will not be accurate in all cases.

4. Zubercal Photometric Performance.

To test the accuracy of the final Zubercal calibration, we have compared the Zubercal results with the current calibration in several ways.

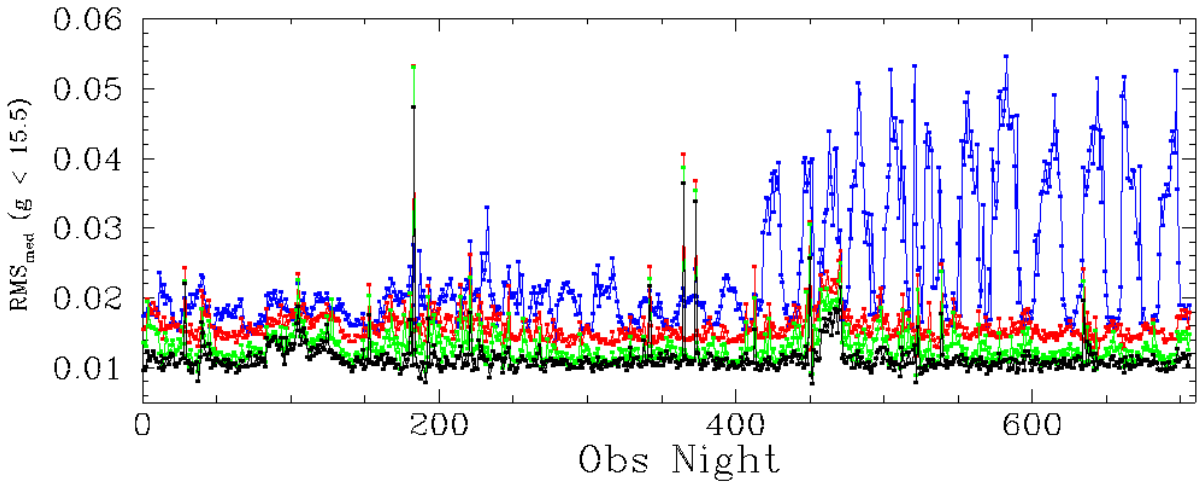

Figure 16. Median g-band RMS of photometric residuals by observation night for CCD 13 for bright stars. Each line gives a different calibration stage. The **blue** points give the current ZTF pipeline calibration with colour coefficient but no other corrections. The **red** points give the initial Zubercal calibration. The **green** points give the values when corrected for spatial structure and colour offsets. The **black** points show the residuals when corrected for fits to the transparency structure.

In Figure 16, we show the effect the Zubercal corrections have on the resulting scatter for ZTF-I photometry. The scatter in the current pipeline calibration (blue line) shows a distinct increase in scatter after observation night 400 (Oct 2019). This is due to the change in the CCD readout waveforms that created the *pocket effect*. As noted above, the affect increases when the sky background is low and exhibits a trend that depends on the lunar phase. Zubercal removes much of the temporal structure and reduces the difference between PS1 and ZTF measurements to ~1% for bright sources.

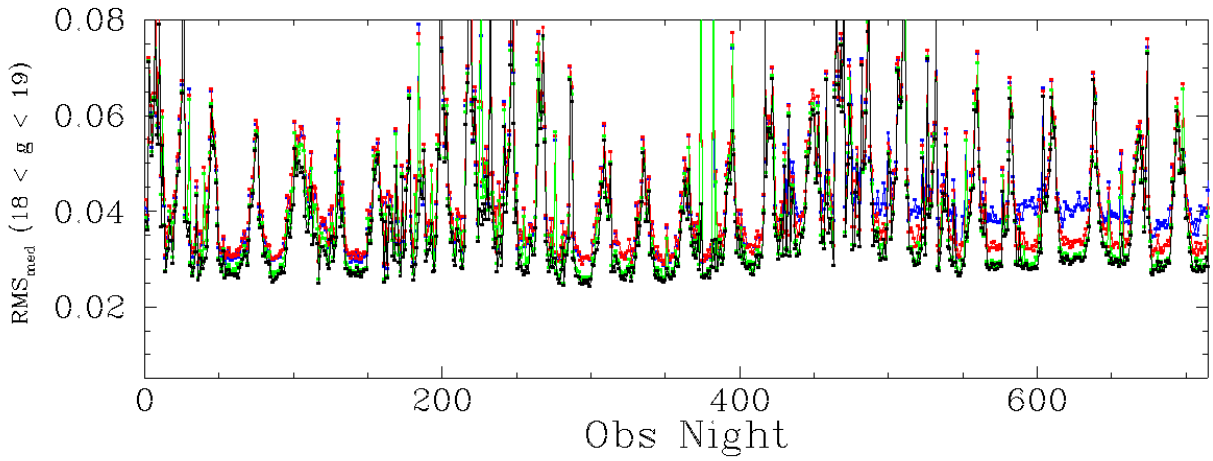

Figure 17. Median g-band RMS of photometric residuals by observation night for CCD 13 for faint stars (18 < mag < 19). Lines and points as defined in plot above.

In Figure 17, we show the temporal trend in photometric residuals for faint sources. In this case, there is only a slight reduction in scatter because the photometric uncertainties are dominated by Poisson noise from the source and sky and the read noise from the CCD pixels. Nevertheless, a significant improvement is seen for the later data where the pocket effect is still present.

4.1 Zubercal Absolute Calibration.

The accuracy of the determination of fundamental astronomical parameters, such as the distances to objects, depends on how well their absolute brightness is measured. Under the assumption that *PS1 is a well-calibrated photometric system*, differences between Zubercal calibrations and PS1 measurements should determine the absolute accuracy of the absolute calibration.

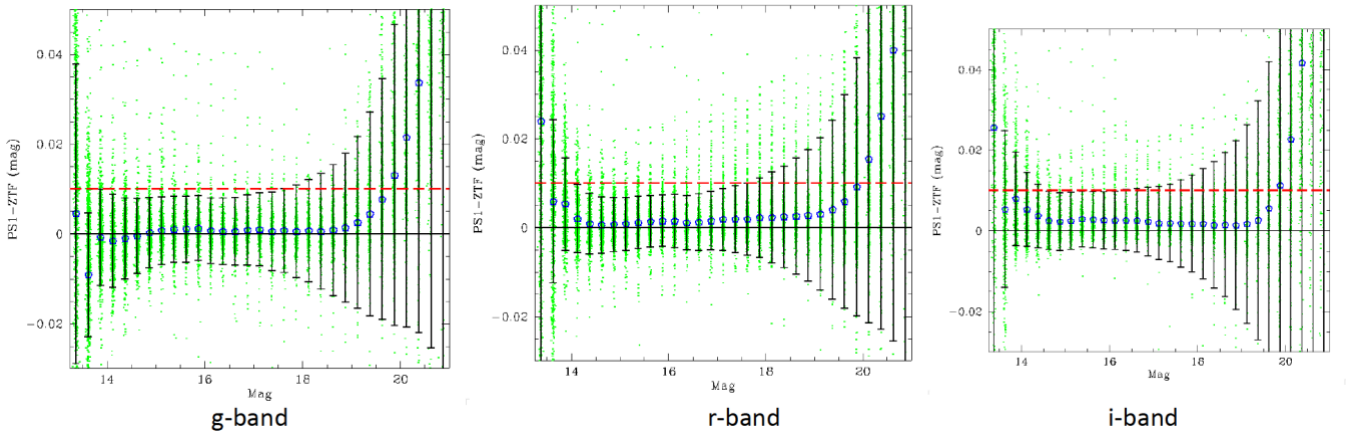

Figure 18. Systematic differences between ZTF and PS1 magnitudes (CCD-10, quad-1).

In Figure 18, we plot the difference between Zubercal magnitudes and PS1 as a function of brightness. Each data point gives the median value for a ZTF field for one CCD quadrant. This plot shows that, in the range between 14th and 19th mag, *systematic differences between PS1 and Zubercal mags are* ***< 0.005 mags***. Thus, moderately bright sources should provide accurate magnitudes. At faint magnitudes, completeness biases begin to dominate. A systematic trend occurs as faint sources are detected when they are brightest due to statistical scatter. Additionally, PS1 stars are only secondary standards. Thus, there is an indirect link to primary standards that can lead to additional uncertainties.

4.2 Zubercal Photometric Scatter.

In cases where absolute calibration is unnecessary, such as detecting the presence of variability, it is usually still important to understand the relative calibration uncertainties. The scale of these uncertainties can be determined by measuring the scatter of the measurements on calibrator sources.

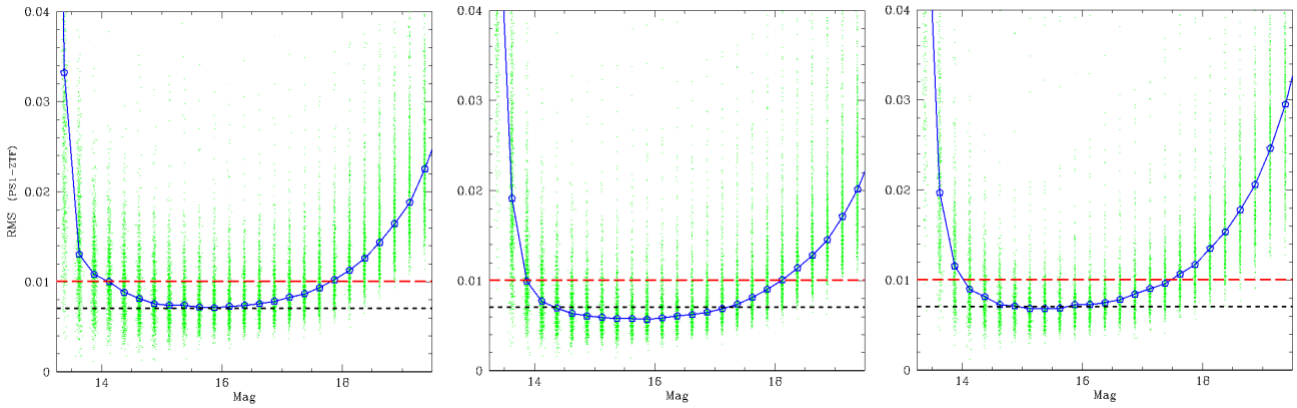

Figure 19. Zubercal RMS vs magnitude for g,r, and i-band, respectively (CCD11, quadrant 2). **Blue** pentagons show the median RMS for all fields and each magnitude bin.

In Figure 19, we plot the scatter as a function of magnitude for all measurements in one quadrant for the three ZTF filters. Each point provides the average scatter within a different ZTF field and magnitude bin. The plot shows that between 14th and 18th mag the Zubercal photometric uncertainties in most fields are typically **less than 1%**. The least scatter in this quadrant occurs at ~16th magnitude in r-band where the median difference between ZTF and PS1 is ~0.006 mags. At bright magnitudes, the saturation of PS1 and ZTF sources leads to large scatter while for faint sources noise from the image readout and background increases the scatter.

To compare the accuracy of Zubercal with the current calibration, we also investigate the RMS scatter of bright sources within the dataset on a frame basis. In Figure 20, we present histograms of the binned values of the median RMS in each frame for bright stars measured in ZTF g and r-band observations. Median frame RMS residuals peak at 1% for bright stars in Zubercal g and r-bands, while the current calibration peak is 1.5%. The maximum RMS in Zubercal is < 0.1 mags, while they are much larger in the current calibrations. These plots also show us that there are far fewer frames where the median RMS is 0.04 mags. Such frames are most effectively removed by the surface fits that are applied to remove variations in transparency across the images (due to clouds, etc.).

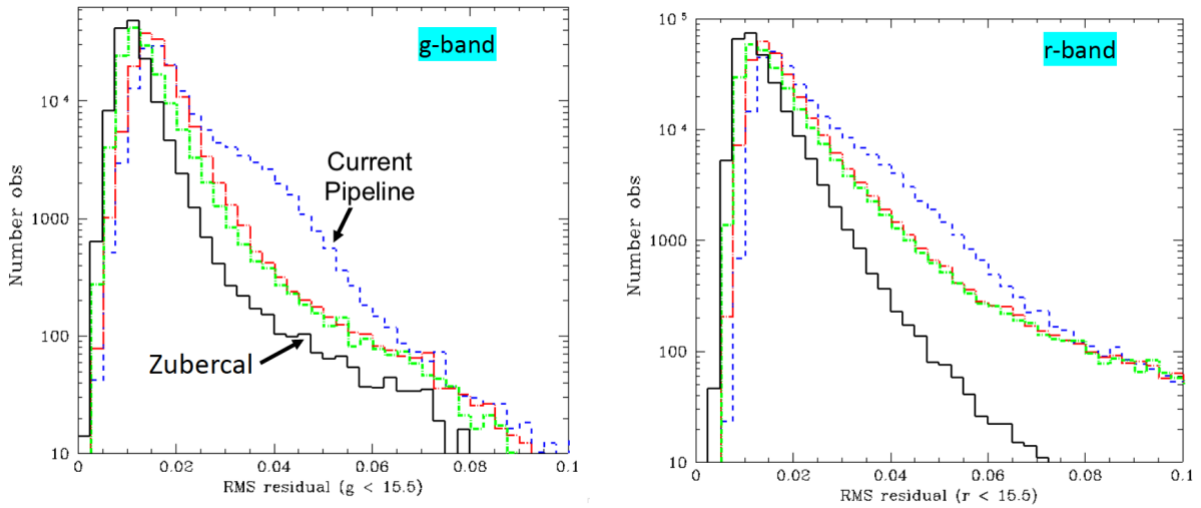

Figure 20. RMS scatter for Zubercal calibration of bright stars (mag < 15.5) at various stages

compared to the current ZTF calibration. Left: scatter for g-band photometry. Right: scatter for r-band photometry. **Blue** line: current ZTF calibration. **Green** line: calibration using Zubercal equation fitting. **Red** line: calibration with Zubercal equation, spatial residual maps and temporal corrections. **Black** line: Final Zubercal scatter after including residual maps and surface fits.

5. Zubercal Data Storage.

Zubercal was designed as a new and improved calibration that would complement and link to the existing ZTF photometry on a detection level. Specifically, Zubercal would provide fully calibrated magnitudes for each ZTF measurement with a matching PS1 source. However, the services have not been implemented. Thus, to maintain the utility of Zubercal as a data product, it was necessary to combine the photometry into a single data set that could be accessed without the existing IRSA database infrastructure.

Due to the size of the data and the need to obtain flexible and rapid access to it, it was decided to spatially partition the sky into tiles of half-degree width. This spatial scale varies from ZTF image (quadrant) sizes (~0.84 deg wide). Zubercal tiles are ~1/3rd of the area of a ZTF image.

For storage space and access speed, we chose Apache's Parquet format with snappy compression. This open-source format provides columnar compression so that repeated entries, such as object IDs, are saved with a minimal footprint using run-length encoding (RLE). Additionally, object IDs can be indexed as in a regular database for quick retrieval without long scans. Zubercal parquet files will be updated with new data on a different schedule than regular ZTF data releases.

5.1 Zubercal Data Partitioning.

Since every ZTF detection within Zubercal is matched to a PS1 ID, and the PS1 IDs are derived from RA and Dec, one can uniquely assign all Zubercal measurements to a spatial grid based on the coordinates associated with a PS1 ID. The formula for creating PS1 IDs from positions is given by the Flewelling et al. 2020 as:

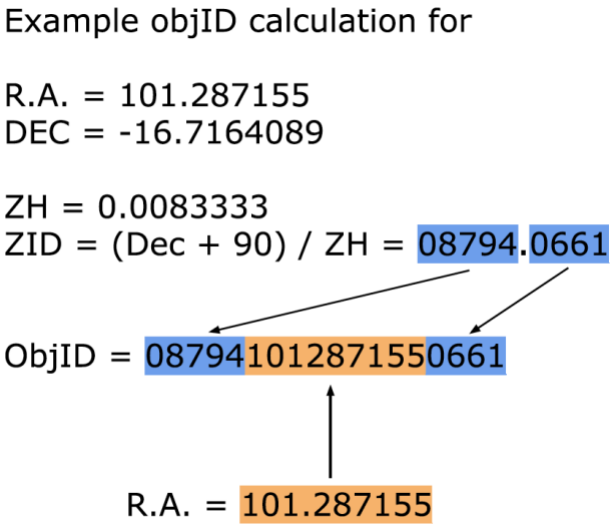

To create a spatial grid, Zubercal parquet photometry files are contained in directories that are divided by half-degree of RA. These directories are named with the combination of field "F" + a four-digit RA number. This number is thus: round[RA(deg),0.5]*10. For example, objects in the range 20 < RA < 20.5 deg are stored in directory F0200, and objects within 120.5 < RA < 121 deg are found under F1205. The number of directories was selected to approximately optimize the directory file depth.

Each object's RA is taken from the four digits in the PS1 ID that encodes it (i.e. ps1id[5:9]). As with the directories, Zubercal filenames use the *Right Ascension* digits rounded down to the nearest 0.5 deg. The filenames also include the four-digit *Declination* part of PS1 ID where this is encoded as above (i.e. Dec = float(int(PS1ID[:5])*0.0083333)-90). The final part is the ZTF filter (e.g. g,r or i). As with the Zubercal RAs, the Decs encoded within parquet filenames are multiplied by 10 and then rounded down. The rounding is to the nearest 5 for Dec < 60 deg, or the nearest 10 when Dec > 60 deg.

For example, a Zubercal parquet file named "*ztf_1310_1400_i.parquet*" will contain ZTF i-band photometry within the ranges: 131 < RA < 131.5 deg and 26.6662 < Dec < 27.0829 [i.e. (1400*0.083333)-90 < Dec < (1405*0.083333)-90].

The approximate RA and Dec limits of each parquet file and filter are provided in the release directory files: Zubercal_g.txt, Zubercal_r.txt, Zubercal_i.txt.

5.2 Zubercal Parquet File Parameters.

Each Zubercal parquet file contains ZTF PSF photometry in the following format:

Quantity	Parameter name	Data Type
PS1 ID	objectid	64 bit int
MJD (tdb)	mjd	64 bit float
mag	mag	32 bit float
mag error*10000	magerr	16 bit uint
ZTF field ID	fieldid	16 bit uint
RC ID	rcidin	16 bit uint
Infobits	info	32 bit int
Phot flag	flag	8 bit int
RA	objra	32 bit float
Dec	objdec	32 bit float

The following parameter definitions and selections are applied.

MJD_(tdb)=Barycentric Dynamical Time (TDB) in MJD format, via astropy.time.
The coordinate localized times are included for higher accuracy.
PS1 ID, closest matching PS1 object. Limit 1.2".
RA and **Dec** are the coordinates of the ZTF detection.
RC ID is the ZTF readout channel number.
ZTF field ID is the ZTF field grid ID number.
Infobits is the IPAC pipeline Infobits flag (includes bit 26).
mag is the Zubercal-based PS1-calibrated magnitude.
mag error the detection uncertainty multiplied by 10000 [to store as a 16 bit uint vs a 32 bit float].
Phot flag is a Zubercal bitwise photometry error flag.

The photometry flag values have the following definitions:
0 = no error

1 = offset spatial correction map value exceeds 0.2 mag limit
2 = colour correction spatial map value exceeds limit (g=0.05, r & i=0.02)
4 = time-based ZP magnitude correction exceeds 0.03 mag limit.
8 = Quadrant-based surface fit correction exceeds 0.1 mag limit.
16 = No correction for best-fit surface applied.
32 = Mag uncertainties from fit < 0.0015 mags. Previous uncertainty value adopted.

Note: No PS1 RAs and Decs are stored within the parquet files. Thus, any query based on coordinates must first be matched with PS1 data. However, individual detection RAs and Decs are retained. These are kept enabling future matching with other catalogs.

6. Zubercal Services and Bulk Downloads.

The current ZTF Zubercal DR20 dataset is ~10TB and is stored in *parquet* format files as noted above. These files are divided into **0.5 x 0.5 deg** spatial tiles over the ZTF sky coverage area (i.e. 0 < RA < 24 hrs, Dec > -30 deg).

Queries for individual objects can be carried out via the Zubercal Release pages. These services allow searches for matching Zubercal lightcurves by Coordinates, Name or PS1 ID, or for Periodicity within the lightcurves. *This service may be discontinued in the near future*.

The **bulk** parquet **format** data files are currently available at:
http://atua.caltech.edu/ZTF/Fields/

Future bulk data releases of the Zubercal data files will be via IRSA.

7. Zubercal Photometry Caveats.

The Zubercal photometric calibration is usually better than the current ZTF photometry available via the IRSA interface, but not always.

In most cases, transient objects detected by ZTF, such as supernovae, are unlikely to have any nearby PS1 match. Such sources cannot be directly calibrated to PS1 and are thus not contained in this dataset. Additionally, any sources detected in ZTF that are missing from the PS1 catalogs will not be included.

Zubercal photometry does not contain all of the error flags that are available with pipeline processing detections (e.g. the *catflags* descriptor is not included) since Zubercal was designed to be included with the regular calibration. The additional quality flags can be obtained by querying the source location via the IRSA interface.

The IRSA interface often returns different numbers of points than Zubercal. This is because the current pipeline uses different data flagging and does not combine detections of the same source detected in different fields/quadrants. Additionally, a small number (~0.1%) of the Zubercal r-band database files were corrupted during processing.

8. Possible Future Improvements to Zubercal.

The Zubercal photometry is based on sources detected in ZTF images at a 5-sigma threshold. The use of a lower threshold would increase the number of detections but also result in many false detections due to image artifacts. However, PanSTARRS and ZTF reference image catalogs provide much deeper source catalogs. Thus, it is possible to perform forced photometry at the location of known sources in ZTF images without introducing any false sources. Work on creating deeper forced photometry catalogs is underway.

9. Where can I get more details about Zubercal?

Additional details of the Zubercal analysis can currently be found at:
http://atua.caltech.edu/ZTF/Web/Calib.html

Appendix

The Zubercal Nightly Processing Pipeline

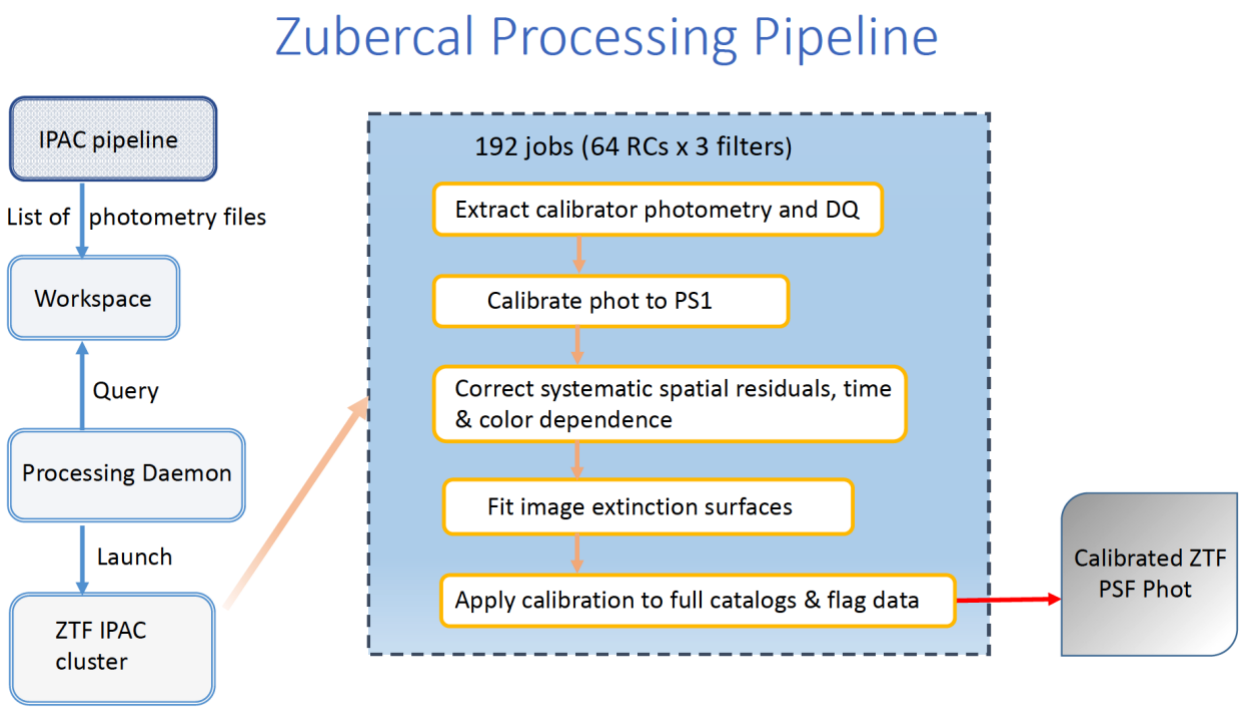

Figure 21. The Zubercal processing pipeline.

The Zubercal processing pipeline runs continuously on the ZTF data processing cluster with a daemon that automatically submits jobs to a *slurm* data queue at the end of each night. This pipeline matches all ZTF detections to PS1 objects and produce ascii *match* catalogs including the calibrated photometry, following the steps in Figure 21. Sources that are detected in ZTF images that have no PS1 match are retained in separate (*unmatched*) ascii files.

Zubercal Data Release Processing.

Prior to a new data release, all of the metadata required to create new parquet entries is combined into lists (MJDs, Infobit, etc). This process typically takes one day.

Zubercal parquet files are created or updated from the individual Zubercal photometry catalogs by reading every new catalog within a given ZTF image quadrant and filter. This data is combined and appended to the matching spatial parquet tiles. Since the parquet tile areas are much smaller than ZTF quadrants, and do not align with the ZTF field grid, it is typical for between 4 and 16 parquet files to be updated for each ZTF field/RC/filter combination. As with the nightly Zubercal processing pipeline, jobs are posted to the ZTF cluster where they must be monitored to avoid any

possible delaying of regular ZTF nightly processing jobs. This process typically takes around two weeks.