The ZTF Science Data System (ZSDS) Advisories & Cautionary Notes

Version 1.3, June 10, 2020

Contents

1 INTRODUCTION 2

2 ADVISORIES & CAUTIONARY NOTES 2

2.1 General 2

2.2 Raw Image Files 3

2.3 Single-Exposure (Epochal) Image products 3

2.4 Single-Exposure (Epochal) Photometry Catalogs 4

2.5 Reference Images (co-adds) and their Source Catalogs 8

2.6 Lightcurves ingested from Source-Matchfile products 9

2.7 Objects Database Table to support Lightcurve Queries 11

2.8 Alert Packets 12

1 INTRODUCTION

Below we summarize the known caveats and warnings to be aware of when using products from the ZSDS. Advisories on how to identify and omit bad-quality data from archive queries are also given. These notes are a work in progress and will be revised as issues become known or are resolved. If you come across new issues or tips worth sharing, please send them to: ztfdatasystem@gmail.com

Some points to be aware of:

- The notes below are a reproduction of the **Cautionary Notes in Section 13** of the unwieldy ZSDS Explanatory Supplement: *https://zwicky.tf/ykv*
  These notes are reproduced here for easier access.

- Unless otherwise stated, section numbers referenced below pertain to those in the ZSDS Explanatory Supplement (*https://zwicky.tf/ykv*).

- The notes below should be used in conjunction with a science product of interest. These products are summarized in the documentation accompanying a specific ZTF Public Data Release (see *https://www.ztf.caltech.edu*).

2 ADVISORIES & CAUTIONARY NOTES

2.1 General

1. The ZSDS archive contains products with basic *automated* data quality assurance applied during pipeline processing. These are products where astrometric and photometric calibration was possible according to thresholds defined by metrics in the `INFOBITS` string (§10.4). This yields potentially usable products, however, "bad" data is still possible. This may be data affected by clouds, low atmospheric transparency, moon contamination, bright source scattering artifacts, unmasked aircraft and artificial satellites, transient detector behavior, and/or inaccurate telescope tracking. The archive metadata contains flags for identifying and omitting bad quality data. For images and accompanying source catalog files derived therefrom, the appropriate metric is `INFOBITS`. For lightcurve measurements (on a per-epoch basis), the metric is *catflags*. These are described further below in §2.4 and §2.6 respectively.

2. The date/timestamps attached to all file-based products (i.e., values for the FITS header keywords: `OBSJD, OBSMJD, SHUTOPEN, DATE-OBS`) pertain to the *start* of the exposure, not the middle. The timestamps attached to the lightcurve photometry measurements however (*hjd, and mjd*) pertain to the *middle* of each exposure, i.e., with 0.5*`EXPTIME` added to the start observation times.

3. The start of an exposure, as defined by the start of photon integration (keywords in note 2 above) pertains to the approximate N-S centerline of the focal plane, i.e., parallel to the shutter blades. Due to the finite time it takes for the shutter to open and close, the actual total integration time depends on a source's position in the focal plane. The effective exposure time is smaller by at most 60 millisec at the E-W edges of the focal plane relative to its center. This decrease towards the edges translates to a relative reduction in photometric throughput of <~ 0.2%. This spatial variation is implicitly caught by the photometric zeropoint calibration at the CCD-quadrant level.

4. The `EXPTIME` (exposure or integration time) keyword value is always 30 seconds for public (MSIP) exposures. This is also the nominal integration time for private and Caltech-time observations, however, for some science programs, it can be 60, 90, 300, or 600 sec.

5. The root filenames and identifiers assigned to the single epoch-based file products are described in §7.1. The fractional time of day represented in a filename (i.e., the *dddddd* in ztf_<YYYYMMDD*dddddd*>…) is UT-based. E.g., 176377 => 0.176377 day => 04h:13m:58.97s UT. It's important to note that these times are typically one second earlier than the start-observation timestamps in the FITS headers (e.g., `OBSJD`). This difference is due to the asynchronous FITS-keyword/telemetry computation and image write time at the P48.

2.2 Raw Image Files

1. The raw CCD data are packaged as Rice-(*fpack*-)compressed multi-extension FITS files (see §6.2). There is one file per CCD, with 16 CCD files per exposure of the focal plane. There are nine FITS extensions: the first extension is a global header; the next four contain the CCD photon-sensitive quadrant pixel data (or readout-channels); the last four contain the over-scan (bias) strips for each respective CCD-quadrant.

2. The photon-sensitive CCD-quadrant pixel data in the raw CCD files *have not* been over-scan corrected using their corresponding over-scan data. There are no other raw intermediate products in the ZSDS where the raw CCD data have been over-scan (bias) corrected.

2.3 Single-Exposure (Epochal) Image products

1. Beware of unmasked artifacts in all images, i.e., not flagged in the accompanying bit-mask images (§10.3). The detection and masking of optical and detector artifacts was performed on a best effort basis. There is no guarantee that all flavors of ghosts, cross-talk, bright-star CCD-bleed artifacts and bad-pixels were tagged. Their presence will have also contaminated the photometry catalogs (§2.4 below), but not necessarily alert packets generated from difference-images (§9) since for these, other *source*-based metrics were used to mitigate artifacts following detection. The machine-learned reliability score also provides a metric for filtering alerts triggered from artifacts.

2. Photometric calibration accuracy (estimates of zeropoints and color coefficients with respect to PS1; §6.5.2) depends on many variables, e.g., seeing, airmass, non-uniform transparency due to intermittent/partial cloud cover, background level, CCD/detector-noise, scattering from the moon and other bright objects, and level of source confusion. There is one estimate of `MAGZP` and `CLRCOEFF` *per CCD-quadrant*. Spatial variations in transparency at the intra-quadrant level will lead to errant and often unusable science products. This includes difference images and alerts derived/triggered therefrom. See note 2 in §2.4 below for instructions on how to identify and omit bad quality data from your queries.

3. In the Public Data Releases, a difference image (file suffix **scimrefdiffimg.fits.fz*) will only accompany a single-exposure science image in the archive if a reference image corresponding to the same field/CCD-quadrant/filter *existed* on/before a single-exposure image was processed. Reference image generation continued throughout the survey during 2018 and image differencing (with alert generation) could only be triggered for fields/CCD-quadrants/filters where a survey-quality reference image was present in the archive. There was no attempt to generate difference images for historical epochs. We expect future data releases to be more complete in terms of the difference images across epochs.

2.4 Single-Exposure (Epochal) Photometry Catalogs

1. Only non-saturated sources are detected and measured, both during PSF-fitting and aperture photometry. The saturation limits lie within ~ 12.5 to 13.2 mag for all filters and CCD-quadrants. The limits are fuzzy due to (i) electronic detector-gain variations and (ii) a given star will fluctuate between being saturated in its core when the seeing is good and not saturated when the seeing is bad. The limit of 12.5 mag is typically for quadrants near the center of the focal plane and with a derived image quality of ~ 2.0 arcsec.

2. The ZTF Observing System is a fully robotic system. One of the drawbacks is that it cannot determine in advance when conditions are non-photometric or the atmospheric transparency is very low, for example, when there are intermittent clouds partially covering the FOV during an exposure. This severely affected the accuracy of the derived photometric calibration solutions. There is one estimate of `MAGZP` and `CLRCOEFF` *per CCD-quadrant.* Therefore, spatial variations in transparency at the intra-quadrant level will lead to errant and often unusable science products.

   "Bad-quality" or generally unusable CCD-quadrant-based images from individual epochs (including accompanying source catalog files) can be omitted when querying the archive by thresholding the `INFOBITS` value in the archive metadata. If `INFOBITS` for an image has value < 33554432 (i.e., does not contain bit 25), the image and catalog data are probably usable.

   Data flagged as bad/unusable were determined beforehand by examining metrics from processed CCD-quadrants for 42 survey fields covering different parts of the sky and observing conditions, and then deriving thresholds. The criteria used to set the bad-data quality flag (specifically bit 25 in the INFOBITS string) for all image/catalog file-based products in the Public Data Releases are:

```
if fid=1 (g filter):
 magzp > 26.7 - 0.2*airmass OR
 magzprms > 0.06 OR
 nps1matches < 80 OR
 magzp < zp_thres[rcid] - 0.2*airmass

if fid=2 (r filter):
 magzp > 26.65 - 0.15*airmass OR
 magzprms > 0.05 OR
 nps1matches < 120 OR
 magzp < zp_thres[rcid] - 0.15*airmass

if fid=3 (i filter):
 magzp > 26.0 - 0.07*airmass OR
 magzprms > 0.06 OR
 nps1matches < 100 OR
 magzp < zp_thres[rcid] - 0.07*airmass
```

`zp_thres[rcid]` means the threshold is CCD-quadrant dependent, where *rcid* = 0…63.

The `zp_thres[rcid]` values are listed in
http://web.ipac.caltech.edu/staff/fmasci/ztf/zp_thresholds_quadID.txt

To demonstrate these criteria, Figure 1 shows an example exposure with all 64 CCD-quadrants impacted by clouds that also happen to be reflecting moonlight onto the focal plane. This exposure is on galactic plane field 806 and was observed on 2018-12-21 UT. The photometry catalogs from these quadrants are not usable. Figure 2 shows trends in *magzp* (photometric zeropoint), *magzprms* (RMS in *magzp*-fit residuals), and *nps1matches* (number of PS1 calibrators) for the same survey field. The scatter at an individual observation JD represents the dispersion across all 64 CCD-quadrants of an exposure. There is also intra-night scatter at this resolution. The *magzp* values shown here were corrected for airmass variations (to a fixed airmass of 1) and for possible variable integration times (to the nominal 30 sec) using *magzp* + $2.5\log_{10}$(AIRMASS) – $2.5\log_{10}$(EXPTIME/30). The bad exposure in Figure 26 is at relative *obsJD* ~ 174 in Figure 2 and is indicated by a huge scatter in the metrics. The red horizontal lines are thresholds derived from the above flagging criteria. The usable *magzp* range is conservative in the sense that the tightest "good-data" *magzp* range across all CCD-quadrants is shown.

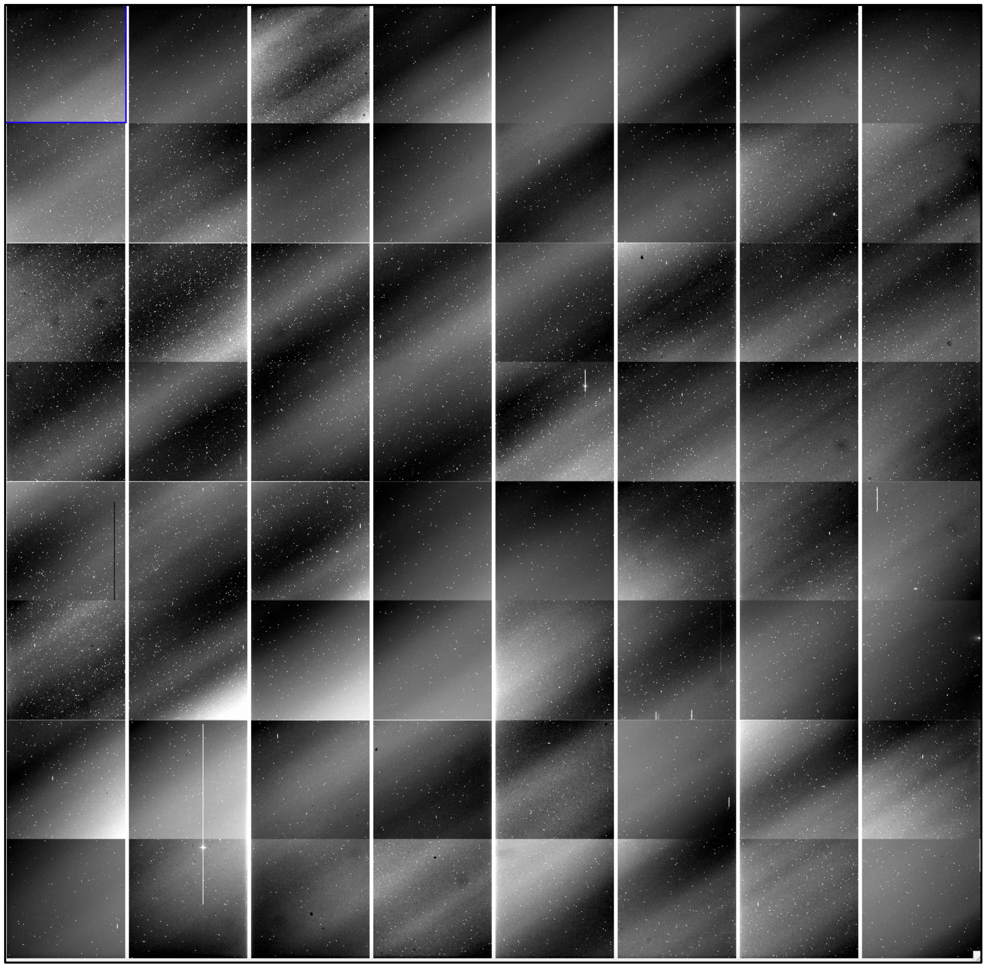

Figure 1: example exposure (one epoch on survey field 806) covered by intermittent clouds that are also reflecting moonlight onto the focal plane. This exposure resides at relative *obsJD* ~ 174 in Figure 2. Sky north is up and east is to the right.

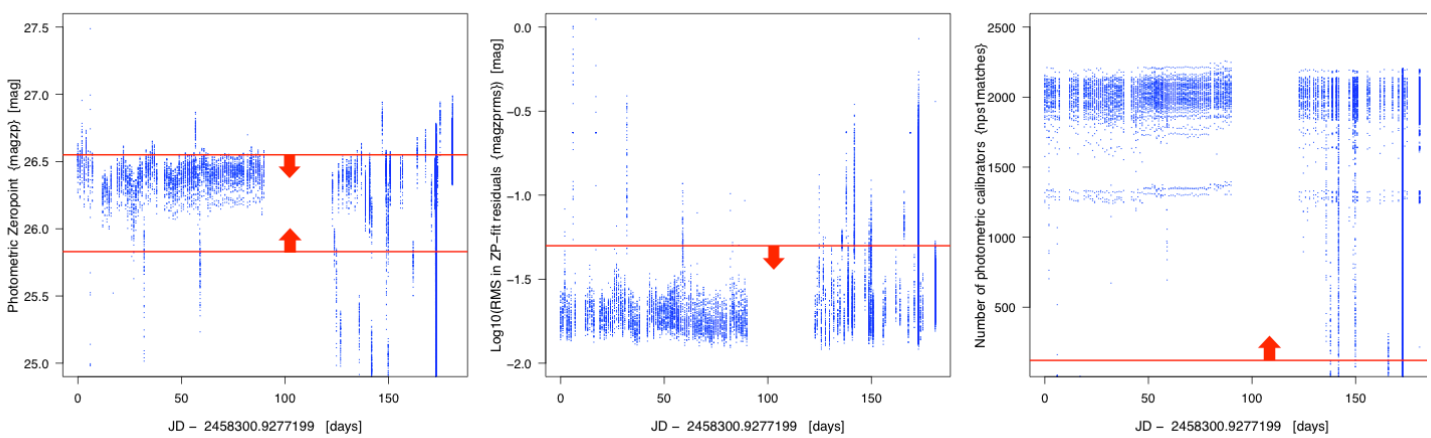

Figure 2: trends in *magzp*, *magzprms*, and *nps1matches* per CCD quadrant in *r*-filter exposures for survey field 806 from 2018-07-01 to 2018-12-29. Red lines with arrows indicate approximate thresholds for "good" (usable) image data. Data outside these limits are flagged in the archive metadata. See cautionary note (2) in Section 3.4.

3. Varying atmospheric transparency on the scale of an individual CCD-quadrant can have a significant impact on the quality of the *i-filter* fringe corrections. Fringe corrections were calibrated from high throughput images acquired under photometric conditions. The presence of intermittent cloud cover, depending on its spatial variation, can lead to erroneous corrections and unusable images.

4. The `INFOBITS` value stored in the archive metadata database (used to flag "bad-quality" data; see note in red text in note 2 above) is not always consistent with that present in the FITS header of a single-epoch image. `INFOBITS` in the FITS headers do not record bit 25, – the bad-data quality flag. This bit was assigned following archival analysis. All other bits (see §10.4 for definitions) are consistent between FITS headers and records in the archive metadata database.

5. The aperture and PSF-fit based single-epoch catalog files (file suffixes **sexcat.fits* and **psfcat.fits* respectively) were constructed independent of each other. I.e., the catalogs are not merged according to position. The PSF-fit catalogs usually contain more sources due to de-blending during PSF-fitting, particularly in crowded regions.

6. Comparisons between calibrated ZTF photometry and PS1 on a CCD-quadrant-basis show biases in the PSF-fit fluxes of up to 0.02 mag for predominately bright sources with *g, r, i*, <~ 15.5 mag (§12.4). These biases are not yet understood. They are also field-dependent and whether there was significant contamination by the moon.

7. Similar photometric biases are seen in the aperture-based catalogs with respect to PS1 (§12.4). As expected, biases in the aperture photometry are typically larger in high source-density fields across the full flux range due to contamination from source crowding. PSF-fitting is more immune to such contamination.

8. Photometric uncertainties in the PSF-fit based single-epoch catalog files (file suffix **psfcat.fits*) *can be* overestimated by factors of up to 3 relative to estimates of the RMS in photometric repeatability using non-variable (noisy) source populations. This overestimation is a function of source flux and is generally larger for brighter sources. These uncertainties have been corrected in the final lightcurve photometry following source-matching across the single epoch catalog files (see note 3 in §2.6 below).

9. Bad photometric quality can also arise from inaccurate over-scan (bias) subtractions during processing of CCD-quadrant images. The presence of bright/saturating stars near or on an over-scan boundary can lead to anomalous bias-level estimates as a function of pixel position and therefore impact the entire image and photometry derived therefrom. See note 10 in §2.6 below for a method to avoid selecting lightcurve data that may have been impacted by this effect. A fix was implemented to mitigate this effect on February 21, 2020, following the DR3 data cutoff of Dec 31, 2019.

2.5 Reference Images (co-adds) and their Source Catalogs

1. Reference images are co-adds of individual epochal images on a CCD-quadrant/filter basis that have been selected to satisfy the quality criteria outlined in §6.7 (step 1 therein). Their depth varies from 15 (minimum) to 40 single-epoch images (maximum).

2. Epochal images from **all** survey programs (public, private, and Caltech-time) that satisfy the quality criteria (§6.7) are used to construct reference images. This means there could be a mixture of input integration times (30, 60, 90, 300 or 600 sec). The total effective integration time is indicated by the `TOTEXPT` keyword in the reference image FITS headers.

3. Only non-saturated sources are detected and measured for inclusion in the reference image photometry catalogs, both during PSF-fitting and aperture photometry. The saturation limits lie within ~ 12.5 to 13.2 mag for all filters and CCD-quadrants, similar to those in the single epoch images (§2.4 below).

4. Photometric calibration solutions are not explicitly derived for reference image products. The single-epoch photometric solutions are propagated to a reference image product by rescaling the input single-epoch images (pixel values) to a common photometric zeropoint before co-addition (see §6.7). Currently, the fixed reference image magnitude zeropoints (`MAGZP`) for *g, R, i* filters are `MAGZP` = 26.325, 26.275, 25.660 respectively. I.e., the reference image products are not recalibrated against any external catalog.

5. The following four metrics residing in the reference-image FITS headers are derived from a median of the corresponding input single-epoch image values: `MAGZPUNC, NMATCHES, CLRCOEFF, and CLRCOUNC`. The `MAGZPRMS` metric for reference-image products represents a robust RMS (based on the MAD) of the input single-epoch `MAGZP` values.

6. The aperture and PSF-fit based reference image catalog products (file suffixes **refsexcat.fits* and **refpsfcat.fits* respectively) were constructed independent of each other. I.e., the catalogs are not merged according to position. The PSF-fit catalogs usually contain more sources due to de-blending during PSF-fitting, particularly in crowded regions.

7. Comparisons between calibrated ZTF photometry and PS1 on a CCD-quadrant-basis show biases in the PSF-fit fluxes of up to 0.02 mag for predominately bright sources with *g, r, i*, <~ 15.5 mag. These biases are not yet understood. They are also field-dependent.

8. Similar photometric biases are seen in the aperture-based catalogs with respect to PS1. As expected, biases in the aperture photometry are typically larger in high source-density fields across the full flux range due to contamination from source crowding. PSF-fitting is more immune to such contamination.

2.6 Lightcurves ingested from Source-Matchfile products

Access to the lightcurve query service was described in §7.6. Table columns and metrics returned by this service are defined in §18. *Matchfile* (lightcurve) generation was described in §6.8.

1. Lightcurves are generated per field, CCD-quadrant, and filter. Due to overlapping fields/quadrants on the sky-tiling grid, it is likely the query service will return multiple lightcurves for the same object. There is no merging of same-object lightcurves corresponding to the same filter. You will need to splice them after retrieving them. Furthermore, the different fields/quadrants will likely have different photometric calibration solutions with respect to PS1. The final lightcurve photometry however is calibrated on the same system.

2. Lightcurves are constructed using seed detections from the deeper (co-add) Reference Image PSF-fit catalogs. These "seed sources" are loaded into the searchable Objects Table (§2.7 below) where their photometry is ~ 2 to 2.5 magnitudes deeper than the single-exposure catalogs. Therefore, a majority of faint sources in the Objects Table (fainter than the single-exposure sensitivity limits) will not have lightcurves. Futhermore, objects with photometry hovering near the single-exposure detection limits will have incomplete temporal-coverage and therefore more sparsely sampled lightcurves.

3. The primary lightcurve photometry is represented by column *mag* with 1σ uncertainty: ± *magerr*. The *mag* values already have their epoch-dependent photometric zeropoints applied (column: *magzp*). By default, all lightcurve photometry is in the native ZTF photometric system. No color corrections are applied (or rather, zero color in the *AB* system is assumed). To transform your photometry onto the *AB* (PS1) system, you will require knowledge of the source color in that system. The exact color used during initial photometric calibration against PS1 is defined by column *pcolor*, the actual color coefficient to enable transformation to the PS1 system is given by column *clrcoeff*, and the procedure for applying color corrections was described in §10.1.1.

4. The magnitude uncertainties *magerr* include empirically-derived corrections so they are overall consistent with the temporal RMS inferred from the photometric repeatability in stationary (noisy) lightcurves at the given *mag*. For details, see §6.8.1.

5. Analogous to the flagging of "bad-quality" images and associated catalogs (§2.4 above), bad or generally unusable observation epochs in lightcurves can be omitted by thresholding the *catflags* column. If *catflags* for an image has value < 32768 (i.e., does not contain bit 15), the photometry at that epoch is probably usable. This flagging removes epochs based on their overall image/calibration quality. The *catflags* value also encodes possible source-specific issues, for example, bad-pixel information propagated from the input catalog *flags* defined in §10.6. If you demand perfectly clean extractions, we advise specifying *catflags* = 0 instead when querying lightcurve epochs.

6. Related to the previous note, the number "good" observation epochs in a lightcurve, i.e., with *catflags* = 0 is given by the *ngoodobs* metric. The total number of epochs is *nobs*. Analogous to these counters is the actual number of epochs in a public release (usually a subset of all epochs). These numbers are given by *nobsrel* (all epochs, regardless of quality) or *ngoodobsrel* (with *catflags* = 0).

7. As of April 8, 2019, relative photometric refinement is turned off during lightcurve generation. Analyses have shown that this post-processing step can give erratic ZP corrections, usually caused by occasional bad (unflagged) data going into the solver, which was difficult to track.

8. A weak but significant periodic signal with period ~ 28 to 29 days is seen in the lightcurves of faint sources with g or r >~ 20.5 mag (close to the 5σ detection single exposure limit). This no doubt is due to background contamination by the moon and is a function of its distance. The local background estimation during PSF-fit photometry is relatively robust against linearly-varying background gradients, but can be biased if the background has a large variation or is significantly non-linear at high spatial frequencies. Furthermore, if you are applying color corrections to obtain photometry in the *AB* system (note 2 above), contamination by the moon may have biased the color of your source due to residuals from inconsistent background subtraction across filters.

9. Lightcurve photometry and all associated products are not replaced if there is any reprocessing or replacement of the reference image and/or single-epoch catalog products in the archive. Publicly-released lightcurve products are frozen at release time and never updated. Any updates to the inputs will propagate into future public releases of the lightcurve database.

10. To additionally avoid possibly corrupt lightcurve measurements (e.g., due to unmasked ghosts or inaccurate over-scan subtractions as described in note 9 of §2.4), one can threshold on the *chi* metric, which is available on a per-epoch basis, but only from the lightcurve GUI service (not lightcurve tarballs). A recommended threshold for eliminating bad photometric measurements is *chi* > 4. This would be in addition to any *catflags* filtering (note 5 above).

2.7 Objects Database Table to support Lightcurve Queries

1. As described in §7.6, access to the lightcurve database first entails performing a spatial-search (combined with other optional filters) on the ZTF Objects Table. These are all the objects detected and extracted (with PSF-fit photometry) from the reference images, i.e., those that reside in the CCD-quadrant-based reference image catalog files (§2.5 above). The cone search for a single query is currently limited to a radius of 10 arcmin.

2. Currently, the ZTF Objects Table does not contain flux transients, that is, single epoch detections with no counterpart in a reference image. Only sources detected in the reference images are represented. The reason is that a large fraction of the transients appear to be associated with image artifacts, detections falling off the overlapping reference image footprint, and blended detections with ambiguous (non-unique) counterparts in the deeper, better quality reference images. The alert packets (§9) are the primary products from which transients should be mined.

3. The Objects Table also contains a selection of collapsed-lightcurve metrics. Source data spanning all observation epochs from **all** survey programs (public, private, and Caltech-time) are used to compute these metrics. The total number of epochs used to compute these metrics is given by *nobs*. Publicly-released lighturves however contain a subset of the measurements are their number is indicated by *nobsrel* which is typically < *nobs*. Any lightcurve-collapsed metric can be included and filtered as part of your spatial query.

4. Photometry and metrics in the Object Table are not replaced if there is any reprocessing or replacement of the reference image and/or single-epoch catalog products in the archive following initial source matching to generate lightcurves. Publicly-released lightcurve products are frozen at release time and never updated. Updates to the inputs will propagate into future public releases of the lightcurve database.

2.8 Alert Packets

The contents and structure of the alert packets were described in Section 9. Below are the known issues and behaviors to be aware of when using these products.

1. All packets have been (lightly) filtered to remove obvious false-positives using the metrics therein (see §9.1). Therefore, some of the metrics will have limits set by the thresholds used in this filtering. Be aware that not all packets correspond to reliable events. Low-level instrumental calibration residuals and anomalies in the image-subtraction process still lurk.

2. The differential photometry (*magpsf* or *magap*) can sometimes exhibit significant intra-night scatter (larger than the quoted errors) when constructing lightcurves from either the photometric history in each packet, or by merging packets corresponding to the same *objectId*. We believe this is due to systematics that need to be corrected upstream. We appreciate such cases be reported to the Data System team.

3. If the number of bad pixels in a 5x5 pixel region (*nbad*) is non-zero, the *magpsf* is more likely to be erroneous. In such cases, we force *magpsf* = *magap*. The *magap* is likely to yield a more accurate measurement due to internal bad-pixel interpolation used during the aperture photometry step. This will be revised when the issue in (2) is addressed.

4. When merging packets to generate lightcurves for a given *objectId* (from both detections and associated photometric histories, including upper limits), it is possible for some packets to have upper-limits in their histories at a specific observation epoch and other packets to have bona fide detections at the *same* epoch. This is due to the finite matching radius (1.5 arcsec) used to (i) assign the initial *objectId* and (ii) construct the photometric history in each individual packet. That is, each alert packet's history is constructed using a positional match to the *detected* event that triggered it, independent of previous alerts. Astrometric and centroiding error can cause positional discrepancies across packets. We advise only reporting photometric upper limits at a given epoch in a lightcurve if detections from other packets don't already exist.

5. Due to astrometric and centroiding error, alerts from the same astrophysical object may not receive the same *objectId* if their measured positions hover around the 1.5 arcsec matching radius used for *objectId* assignment or reuse across epochs. The important point is that the *nearest* previously assigned *objectId* (if found) is reused. It is therefore possible for two very close alerts (detected at different epochs) being assigned different *objectIds* according to this rule (see §9.4 for details).

6. Multiple alerts detected within 1.5 arcsec of each other on the *same image* (observation epoch) may lead to all/some being assigned the same *objectId*. This would be the *objectId* of a previously named (historic) alert if found within 1.5 arcsec (see §9.4 for details). If no previously named alert exist, each alert would acquire a unique *objectId* for the first time. Matching to future alerts (if they reoccur close to these earlier alerts) will then lead to ambiguities in the assignment of *objectId,* again due to astrometric and centroiding error.

These ambiguities are more likely to occur in high source density regions, e.g., the galactic plane.

7. The PSF-fit differential photometry estimates (metrics *magpsf* and *sigmapsf* in the alert packets) cannot easily be reconstructed using the available image products (difference image and accompanying PSF from the archive). These estimates are computed from other internally generated products from ZOGY (see §6.6). The intent here was to reduce photometric scatter early in the project (see note 2 above) without disrupting existing image products from which all other metrics are computed and used to train and update the *RealBogus* classifiers.

8. Beware when interpreting the lightcurves constructed from merging the differential photometry from multiple alert packets for an *objectId*. There are instances where an object is covered by different overlapping survey fields and detector quadrants. As a reminder, each field/quadrant is treated independently in the ZSDS, where each has its own reference image, possibly constructed from different input-image epochs. There is no guarantee that the reference images are not contaminated (by different amounts) by the transient or variable source flux sought. This means that the differential photometry can show significant scatter if packets originating from different fields/quadrants (with different "reference baselines") are combined.

9. It can be a challenge to ascertain the reliability of a single alert using its packet metadata alone, including the supplied *rb* or *drb* scores. Unless the alert is associated with a moving object (as indicated in the metadata), one recommendation is to perform a spatial cross-match with earlier or later alerts (if available) and construct a lightcurve. Temporal correlation is a powerful way to determine if a collection of alerts are real (astrophysical) or not.

10. The *sgscore* (star/galaxy separation score) metrics attached to the nearest PS1-catalog associations have at times found to be unreliable, most notably for events from variable stars. This is due to issues with the original PS1 photometry catalogs since these were used to train the machine-learned star/galaxy classifiers outside the ZSDS.

11. All alert packets generated per CCD-quadrant subtraction image are archived at the ZSDS and stored in gzipped tar files (generic suffix: *_alerts.tar.gz*). These can be accessed from IRSA using the same GUI and API tools for querying the single epoch CCD-quadrant image products (§7.6).

12. The reliability of an alert or its photometric accuracy strongly depends on the calibration quality of the CCD-quadrant used during image subtraction. Photometric calibration accuracy (estimates of zeropoints and color coefficients with respect to PS1; §6.5.2) depends on many variables, e.g., seeing, airmass, non-uniform transparency due to intermittent/partial cloud cover, background level, CCD/detector-noise, scattering from the moon, and level of source confusion. There is one estimate of `MAGZP` and `CLRCOEFF` *per CCD-quadrant*. Spatial variations in transparency at the intra-quadrant level will lead to

unreliable and often unusable alerts extracted from the quadrant subtraction image. See note 2 in §2.4 above for a method on how to identify and omit all alerts triggered from a suspect badly-calibrated CCD-quadrant image.

13. The alert packets can also be combined to construct absolute photometric (DC) lightcurves, for example, for (a)periodic variable sources. Such lightcurves can be computed by combining the nearest reference image catalog magnitude (*magnr*), differential magnitude (*magpsf*), and *isdiffpos* (positive or negative difference image detection) as follows:
    $mag_{abs} = -2.5\log_{10}(10^{[-0.4magnr]} + sign10^{[-0.4magpsf]})$,
    where
    $sign = 1$ if $isdiffpos$ = '*t*' or $sign = -1$ if $isdiffpos$ = '*f*'.
    Before using the nearest reference image source magnitude (*magnr*), you will need to ensure the source is close enough to be considered an association (e.g., *distnr* <~ 1.5 arcsec). It is also advised you check the other associated metrics (*chinr* and/or *sharpnr*) to ensure it is a point source. These metrics are defined in §10.6. We recommend 0.5 <~ *chinr* <~ 1.5 and/or –0.5 <~ *sharpnr* <~ 0.5.

14. NaN'd pixel values in the alert packet image cutouts are an indicator of bad (unusable) pixels. These NaNs were propagated from upstream processing and are also present in the archived difference image products. Bad (fatal) pixels identified in upstream processing are intentionally forced to NaN prior to differencing. Specifically, the bad pixels set to NaN are those tagged with bits 0,2,3,5,8,9,10,12 in the science images and are defined in §10.3. Various convolutions and interpolations during the PSF-matching process will have spatially-expanded these bad-pixel regions. NaNs were the simplest way to propagate and track this "badness".

15. Alert packets are generated close to realtime typically within 20 minutes of a camera exposure, using filtered events extracted from difference images. Image differencing requires the presence of a reference image. Therefore, alerts can only be generated if a reference image corresponding to the same field/CCD-quadrant/filter *existed* on/before a single-exposure image was processed. Reference image generation continued throughout the survey and therefore there are regions of sky where no alerts could be generated during the course of the survey.

16. A note of caution on the *ndethist* metric in alert packets. This metric is not related to the contents of the *prv_candidates* block in a packet. This metric enumerates every associated detection (within 1.5 arcsec of the target alert's position) from an *internal* candidates database down to S/N ~ 3. Only historical detections falling on the same survey field and CCD-quadrant from which the target alert was extracted are counted. The S/N here is shown as an approximation (~ 3) since it is based on a point source match-filtered estimate (peak S/N estimate, a proxy for the photometric S/N). This should not be confused with associations in the *prv_candidates* block which contain detections down to a *photometric* S/N cut of exactly 5 (precisely 1.0857/*sigmapsf*).