Generating Lightcurves from Forced PSF-fit Photometry on ZTF Difference Images

Frank Masci, Russ Laher, David Shupe & Ben Rusholme

2022-04-06, v2.3

Outline

1 Introduction

This document describes how to submit forced photometry requests to operate on archived difference images from the ZTF survey. These difference images would have been initially generated by IPAC's real-time production pipeline. We also cover the fundamentals of generating publication-quality lightcurves from the raw output, as well as caveats, warnings, and suggestions for optimizing them, for example by combining measurements in order to place tighter constraints or recover signal below the single-exposure limits.

Please familiarize yourself with the advisories in §2 and §8. Also, be prepared to write a script to analyze the raw table output and plot your lightcurves. Stay tuned for more advisories as we learn them, primarily from you.

If your publication uses products from this service, we recommend referring to this as: "The ZTF forced-photometry service…" and citing the following:

> *Masci, F. J., Laher, R. R., Rusholme, B., et al. 2018, The Zwicky Transient Facility: Data Processing, Products, and Archive, PASP, 131, 995.*

Also, if your publication used this service, please include the following in your acknowledgments:

2 Before you submit your request

(a) You will first need to register by email requesting access to ztf@ipac.caltech.edu. You will be notified with a personal password. **Note:** the email address you use to register will be the one you use to submit requests and where notifications that products are ready are received. **Please indicate in your request for registration if you are affiliated with the ZTF Partnership or ZTF Project at Caltech. This includes students and postdocs. If so, it is strongly advised that you register with your institutional email address.**

(b) Ensure your target of interest has been vetted by other means, i.e., there is independent confirmation that it is real or potentially "interesting". This goes without saying, right? Basically, the ZTF forced-photometry service is not a vetting service for disseminating candidates discovered from the individual epochal images. If the images covering your position(s) are publicly available, we recommend examining them. Images can be retrieved from the archive portal: https://irsa.ipac.caltech.edu/Missions/ztf.html.

(c) Ensure your supplied R.A., Dec. position is accurate and reflects what would be measured directly off IPAC's astrometrically-calibrated science images in the archive. The position based on a single discovery epoch may not be accurate enough. Images can retrieved from the archive portal: https://irsa.ipac.caltech.edu/Missions/ztf.html.

(d) The maximum number of individual (distinct) target positions you can submit at any given time is 100. You will be notified if this quota is exceeded. You will be able to submit more jobs once the number in your queue has dropped below 100. We recommend waiting until you have been notified that all 100 requests are complete and ready for download before submitting more.

(e) Ensure you request a sufficient number of future and/or historical epochs after/prior to the Julian Date (JD) defining your transient event or characteristic behavior. We suggest ≥ 30 additional epochs. The purpose is to enable a possible baseline correction (Section 10) and/or uncertainty rescaling (Section 11). This will also allow you to combine the single-epoch measurements in order to go deeper (Section 14). Ensuring additional future/historical epochs is usually unnecessary for continuous (a)periodic variables. It is assumed you understand your targets well enough to judiciously select their timespans.

3 How to submit your requests

An individual submission consists of supplying a *single* R.A., Dec. position, a time-range of interest, and your credentials. You can submit requests using the graphical user interface (GUI) form:

https://ztfweb.ipac.caltech.edu/cgi-bin/requestForcedPhotometry.cgi

The **username, password** for initiating the GUI service is: ztffps, dontgocrazy!

These credentials are fixed and should not be confused with your personal password, which you should have received by registering – see Section 2, note (a). The requestor's email address and password you supply on the form are precisely those you registered with.

Alternatively, you can submit requests using the *wget* command-line utility. Below is a synopsis, where quantities in the angled brackets denote user inputs. Ensure there are no spaces inside the "quoted text".

```
wget --http-user=ztffps --http-passwd= dontgocrazy! -O log.txt
"https://ztfweb.ipac.caltech.edu/cgi-
bin/requestForcedPhotometry.cgi?ra=<RA>&dec=<DEC>&jdstart=<JD_START>&jdend=<JD
_END>&email=<EMAIL_ADDRESS>&userpass=<USERPASS>"
```

where:
  RA = Right Ascension of your target position [decimal degrees; J2000]
  DEC = Declination of your target position [decimal degrees; J2000]
  JD_START = Julian Date corresponding to earliest epoch of single exposure [decimal days]
  JD_END = Julian Date corresponding to latest epoch of single exposure [decimal days]
  EMAIL_ADDRESS = The email address you registered with (see Section 2, note a)
  USERPASS = The password you were provided upon registering (see Section 2, note a)

The Julian Date parameters are optional, if not supplied will default to:
  JD_START to the beginning of survey 2458194.5 (2018-03-17Z00:00:00)
  JD_END to current JD
Thus, a request specifying no Julian Dates will automatically return all possible data. This will be significant for high-cadence fields, please use sparingly/appropriately!

Here is an example (to be executed as a single line on your command-line prompt):

```
wget --http-user=ztffps --http-passwd=dontgocrazy! -O log.txt
"https://ztfweb.ipac.caltech.edu/cgi-bin/requestForcedPhotometry.cgi?
ra=280.8058788&dec=45.2077645&jdstart=2458231.891227&jdend=2458345.025359&
email=xxxx&userpass=xxxx"
```

In case you're interested, this example returns a subset of measurements covering a previously discovered Type Ia Superluminous Supernova.

4 Checking the status or history of your submissions

You can check the status or history of your submissions using the utility:

https://ztfweb.ipac.caltech.edu/cgi-bin/getForcedPhotometryRequests.cgi

The **username, password** for initiating this service is also: ztffps, dontgocrazy!

Due to possible excessive database loads, the job queue is refreshed every hour. It is refreshed on the hour according to wall clock. Therefore, the job status you see could be up to one hour old. Furthermore, only requests submitted within the last month can be accessed.

5 What image data is included in requests?

Generally, forced photometry is performed on all epochal images that are:

(i)   From the public portion of the ZTF Survey, up to and including the *most recent* night, and,
(ii) From the private (partnership) surveys that are *now* in the public domain.

Additionally, if you indicated you were affiliated with the ZTF Partnership or Project at Caltech when you registered (see Section 2 note a), the moderators will have consulted with project management to determine specific access privileges. If access was approved, measurements will also be performed on epochs associated with the *unreleased* private data (from either the partnership and/or Caltech-allocated programs).

6 How will I be notified when products are ready?

Users will be notified of completed jobs via email. The address used is the one you registered with (see Section 2 note a). This email will contain URLs pointing to the location of your forced-photometry products. You will be able to download your products from these URLs by supplying **username, password** = ztffps, dontgocrazy!

Others will not be able to access your lightcurves unless you supply them with the precise URLs. They will not be able to navigate through the URL directory tree.

You will receive a separate email for each submitted request.

7   Raw output table: what does it mean?

The output is a table in ASCII format. This consists of metadata and forced-photometry measurements for all observation epochs that you were authorized to access (Section 5). A header describing the table columns is provided.

An example output table (with the first two rows of data) is shown on pages 7 – 8.

This output table does not represent a stand-alone lightcurve that you can immediately plot. It contains difference-image forced photometry measurements on the requested input position and JD range covering **all fields, CCD-quadrants, and filters** (g, r, or i) for which difference-images exist in the archive and which you are authorized to access. Before generating a lightcurve (Sections 12 and 13), you will need to analyze the contents of this file and extract the rows (epochs) for the specific *field ID* and *filter* of interest, including any quality metrics you wish to filter on.

Multiple *field IDs* may cover your target and each *field ID* (with specific CCD-quadrant therein) will have its own reference image, which may have been created from an entirely different set of observation epochs. You will need to treat each *field* for a given *filter* separately to account for possible contamination from a transient's flux (see Section 10 for details).

Ignoring all the quality metrics you may want to threshold, the minimum columns for generating a forced PSF-fit-photometry differential lightcurve **for a given ZTF *field* and *filter*** are: *field, filter, jd, zpdiff, forcediffimflux,* and *forcediffimfluxunc.*

```
# ---------------------------------------------------------------------
# Generated by forcedphotometry_trim.pl, v3.3 on 2020-11-30 at 11:34:50
#
# Requested input R.A. = 265.42935 degrees
# Requested input Dec. = 67.961877 degrees
# Requested JD start = 2458198 days
# Requested JD end = 2458268 days
# Authenticated program IDs = 0, 1, 2, 3
#
# Column definitions:
# index = sequential counter
# field = ZTF Field identifier
# ccdid = CCD identifier (1..16)
# qid = Quadrant (CCD-amplifier) identifier (0..3)
# filter = Filter string: ZTF_g, ZTF_r, ZTF_i
# pid = processed image Operations DB identifier
# infobitssci = processing summary/QA bits for sci image
# sciinpseeing = Effective FWHM of sci image [pixels]
# scibckgnd = Background level in sci image [DN]
# scisigpix = Robust sigma per pixel in sci image [DN]
# zpmaginpsci = Photometric zeropoint for sci image [mag]
# zpmaginpsciunc = 1-sigma uncertainty in zpmaginpsci [mag]
# zpmaginpscirms = RMS (deviation from average) in difference between instrumental mags
#                  and PS1 calibrators [mag]
# clrcoeff = Linear color coefficient from calibration; for ZTF_g,r,i, PS1 color used
#            is g-r, g-r, r-i respectively
# clrcoeffunc = 1-sigma uncertainty in clrcoeff
# ncalmatches = Number of PS1 calibrators used in initial calibration of sci image
# exptime = Integration time for sci image [sec]
# adpctdif1 = Full sci image astrometric RMS along R.A. with respect to Gaia1 [arcsec]
# adpctdif2 = Full sci image astrometric RMS along Dec. with respect to Gaia1 [arcsec]
# diffmaglim = Magnitude limit in difference image [mag]
# zpdiff = Photometric zeropoint for difference image [mag]
# programid = Program identifier [0=engineering; 1=public; 2=private; 3=Caltech time]
# jd = Julian Date at start of exposure [days]
# rfid = Rerence image Operations DB identifier
# forcediffimflux = Forced difference image PSF-fit flux [DN]
# forcediffimfluxunc = 1-sigma uncertainty in forcediffimflux [DN]
# forcediffimsnr = Signal-to-noise ratio for forcediffimflux measurement
# forcediffimchisq = Reduced chi-square in PSF-fit
# forcediffimfluxap = Forced difference image aperture flux using a fixed 9-pixel
#                     diameter aperture [DN]
# forcediffimfluxuncap = 1-sigma uncertainty in forcediffimfluxap [DN]
# forcediffimsnrap = Signal-to-noise ratio for forcediffimfluxap measurement
# aperturecorr = Aperture (curve-of-growth) correction factor that was applied to
#                forcediffimfluxap measurement
# dnearestrefsrc = Distance to nearest ref image source if found within
#                  5 arcsec [arcsec]
# nearestrefmag = Magnitude of nearest ref image source [mag]
# nearestrefmagunc = 1-sigma uncertainty in nearestrefmag [mag]
# nearestrefchi = Chi parameter for nearestrefmag measurement (ratio: RMS in
#                 PSF-fit residuals / expected RMS from priors)
# nearestrefsharp = Sharp parameter for nearestrefmag measurement
#                   (~0 => PSF like; >>0 => extended; <<0 => pixel spike or hard edge)
# refjdstart = JD of earliest sci image used for ref image [days]
# refjdend = JD of latest sci image used for ref image [days]
# procstatus = Per-epoch processing status codes (0 => no warnings); if non-zero, see
#              accompanying log file and document link below
#
# Documentation: https://zwicky.tf/vgt
#
# ---------------------------------------------------------------------
# Order of columns below:
 index, field, ccdid, qid, filter, pid, infobitssci, sciinpseeing, scibckgnd,
scisigpix, zpmaginpsci, zpmaginpsciunc, zpmaginpscirms, clrcoeff, clrcoeffunc,
ncalmatches, exptime, adpctdif1, adpctdif2, diffmaglim, zpdiff, programid, jd, rfid,
forcediffimflux, forcediffimfluxunc, forcediffimsnr, forcediffimchisq,
```

forcediffimfluxap, forcediffimfluxuncap, forcediffimsnrap, aperturecorr,
dnearestrefsrc, nearestrefmag, nearestrefmagunc, nearestrefchi, nearestrefsharp,
refjdstart, refjdend, procstatus
#
 0 848 7 3 ZTF_r 444364252715 0 2.3975 188.193 6.63524 26.2349 4.7425e-06 0.0217544
0.0915153 7.56753e-06 669 30. 0.047876 0.046107 20.2909 26.2349 1 2458198.8642593
848120227 -250.415274640521 27.3875326806812 -10.1709778472871 2.17363468110553 -
83.6123400317302 48.9406316255288 -1.71995070673677 1.03625409860041 0.340755063179189
16.9430002212524 0.0509999990463257 4.34899997711182 0.243000000715256 2458198.864259
2458258.948056 0
 1 848 7 3 ZTF_r 444365232715 0 2.4255 188.429 6.64423 26.2327 5.8564e-06 0.0241068
0.0960336 9.3264e-06 669 30. 0.051224 0.047296 20.2558 26.2327 2 2458198.8652315
848120227 -242.213408467973 29.1486336595707 -9.24346455530012 1.77261561156291 -
59.802634347219 51.3060597609554 -1.17345609487057 1.03709824026656 0.340755063179189
16.9430002212524 0.0509999990463257 4.34899997711182 0.243000000715256 2458198.864259
2458258.948056 0

...

8 Advisories and cautionary notes

(a) Forced photometry requests will be processed at a slower rate between the hours of sunset to sunrise (Pacific Time). This is due to the limited resources available during nightly pipeline operations.

(b) The output table contains difference-image forced photometry measurements on the requested input position and JD range that covers **all** fields, CCD-quadrants, and filters (*g*, *r*, or *i*) for which difference-images exist in the archive and for which you are authorized to access.

(c) The forced photometry service only operates on difference images that already exist in the archive. The earliest observation JD in your returned table reflects the earliest difference image that exists, for the given field, CCD-quadrant, and filter. This is contingent on when the corresponding reference image was created, and not on the earliest science image that exists. The reference-image creation timestamp is provided in the output table: *refcreated.*

(d) Requested JD_START earlier than the beginning of the ZTF survey, or JD_END in the future, will be reset to default values (2458194.5, current JD respectively).

(e) We recommend using the PSF-fit photometry measurements for all lightcurves (*forcediffimflux* and *forcediffimfluxunc*). Curve-of-growth-corrected aperture photometry measurements are also provided (*forcediffimfluxap* and *forcediffimfluxuncap*). The latter are noisier and serve more as a sanity check for the PSF-fit photometry.

(f) All columns in the output table prefixed with "*forced*…" refer to **differential** flux measurements. These are relative to the flux level in the archived reference image. We advise examining the time range of the input data used to generate the initial reference image (*refjdstart* and *refjdend*) for possible contamination by the actual transient flux. To be sure, we also advise visually examining the reference image. If transient flux is present (and you will see this in the time series of the historical and/or future differential fluxes), you will need to perform a baseline correction (see Section 10).

(g) If your target is a variable star or some other (a)periodically varying source, the reference image will likely contain a time-averaged signal from the source. In this case, you may be interested in generating an *absolute* photometric lightcurve by combining the differential fluxes with the magnitude of the nearest reference-image catalog source: *nearestrefmag*, but first checking its distance from your target: *dnearestrefsrc.* See Section 13 for details.

(h) Beware of bad pixels impacting the individual epochal measurements. This will be indicated by a *procstatus* containing the value 56. See Section 9 for code definitions.

(i) A significant fraction of the ZTF data is non-photometric, that is, was acquired through intermittent cloud cover and/or was contaminated by the moon. This would have impacted the photometric calibration. It is advised to trend the calibration and noise metrics provided in the output table: *zpmaginpsci*, *zpmaginpscirms*, and/or *scisigpix*. Outliers can be identified in these trends and omitted from your final lightcurve. We find the *scisigpix* metric to be the most useful.

(j) The *jd* and *obsdate* (UT) timestamps pertain to the **start** of an exposure (nominally 30s).

(k) All astrometry pertains to equinox J2000 and all archived products were astrometrically calibrated with respect to Gaia DR1.

9 Exit and warning codes (the *procstatus* column)

Overall exit and warning codes from execution of your request are provided in the log file that accompanies the output table. Processing status codes are also provided on a per-epoch basis. This is the *procstatus* column in the output table. These codes are defined below.

- 0   => Successful execution

- 56   => One or more epochs have photometry measurements that may
        be impacted by bad (including NaN'd) pixels

- 57   => One or more epochs had no reference image catalog source
        falling with 5 arcsec

- 58   => One or more epochs had a reference image PSF-catalog that
        does not exist in the archive

- 59   => One or more epochs may have suspect photometric uncertainties
        due to early creation date of difference image in production

- 60   => One or more epochs had upsampled diff-image PSF dimensions
        that were not odd integers

- 61   => One or more epochs had diff-image cutouts that were off the
        image or too close to an edge

- 62   => Requested start JD was before official survey start date
        [3/17/18] and was reset to 2018-03-17T00:00:00.0 UT

- 63   => No records (epochs) returned by database query

- 64   => Catastrophic error (see log output)

- 65   => Requested end JD is before official survey start date [3/17/18]

- 255 => Database connection or query execution error (see log output)

10 Baseline correction using the raw *flux* lightcurve of a suspect transient

For a given *field* and *filter* in your output table, it is recommended you examine a plot of differential flux *forcediffimflux* [DN] versus time [JD] to determine if there is any residual offset in the "baseline". The baseline here refers to the level defined by either historical or future epochs, where the transient signal relative to the reference image is stationary in time, relative to the background noise. For example, if the reference image used to generate the difference images was contaminated by the transient's flux, your baseline will be < 0. The reference image may also have been affected by some other systematic in its generation or calibration, and therefore the baseline could be > 0.

Requesting a special time range for regenerating a reference-image for the purpose of ensuring a zero baseline in the differential photometry is superfluous. You will *always* need to check and correct for a possible non-zero baseline *for transients,* due to possible systematics, and hence obtain the best photometric accuracy. This *is not necessary for (a)periodic variables* that may fluctuate about a constant long-term flux-level. The reason is that even if the reference image was not a true (unbiased) time-average of all the input epochs (at random lightcurve phases), any residual baseline will need to be retained when converting from differential to *absolute* photometry using the reference image source catalog flux (see Section 13).

The baseline correction needs to use a sufficient number of future and/or historical measurements. We suggest ≥ 30 epochs. The more the better since epochs may be affected by low transparency (cloud cover) and/or moon glow, – see note (i) in Section 8. It is assumed you have some prior knowledge of the behavior of your transient to include a sufficient number of before and/or after epochs before submitting your request.

Figure 1 shows an example of a transient. The flux versus time plots in *g* and *r* are on the top row. The *r*-filter flux plot on the right shows a *future* baseline of ∼ –90 DN (dashed red line), as inferred from the differential fluxes at epochs with relative JD < 60 days where the transient reaches quiescence. A trimmed average or median of the measurements in this "stationary" period guided by visual examination of the flux lightcurve is recommended.

Following an estimate of the baseline level, you will need to *subtract* this from all the differential flux measurements (*forcediffimflux*) before converting them to calibrated magnitudes and upper limits (see Section 12).
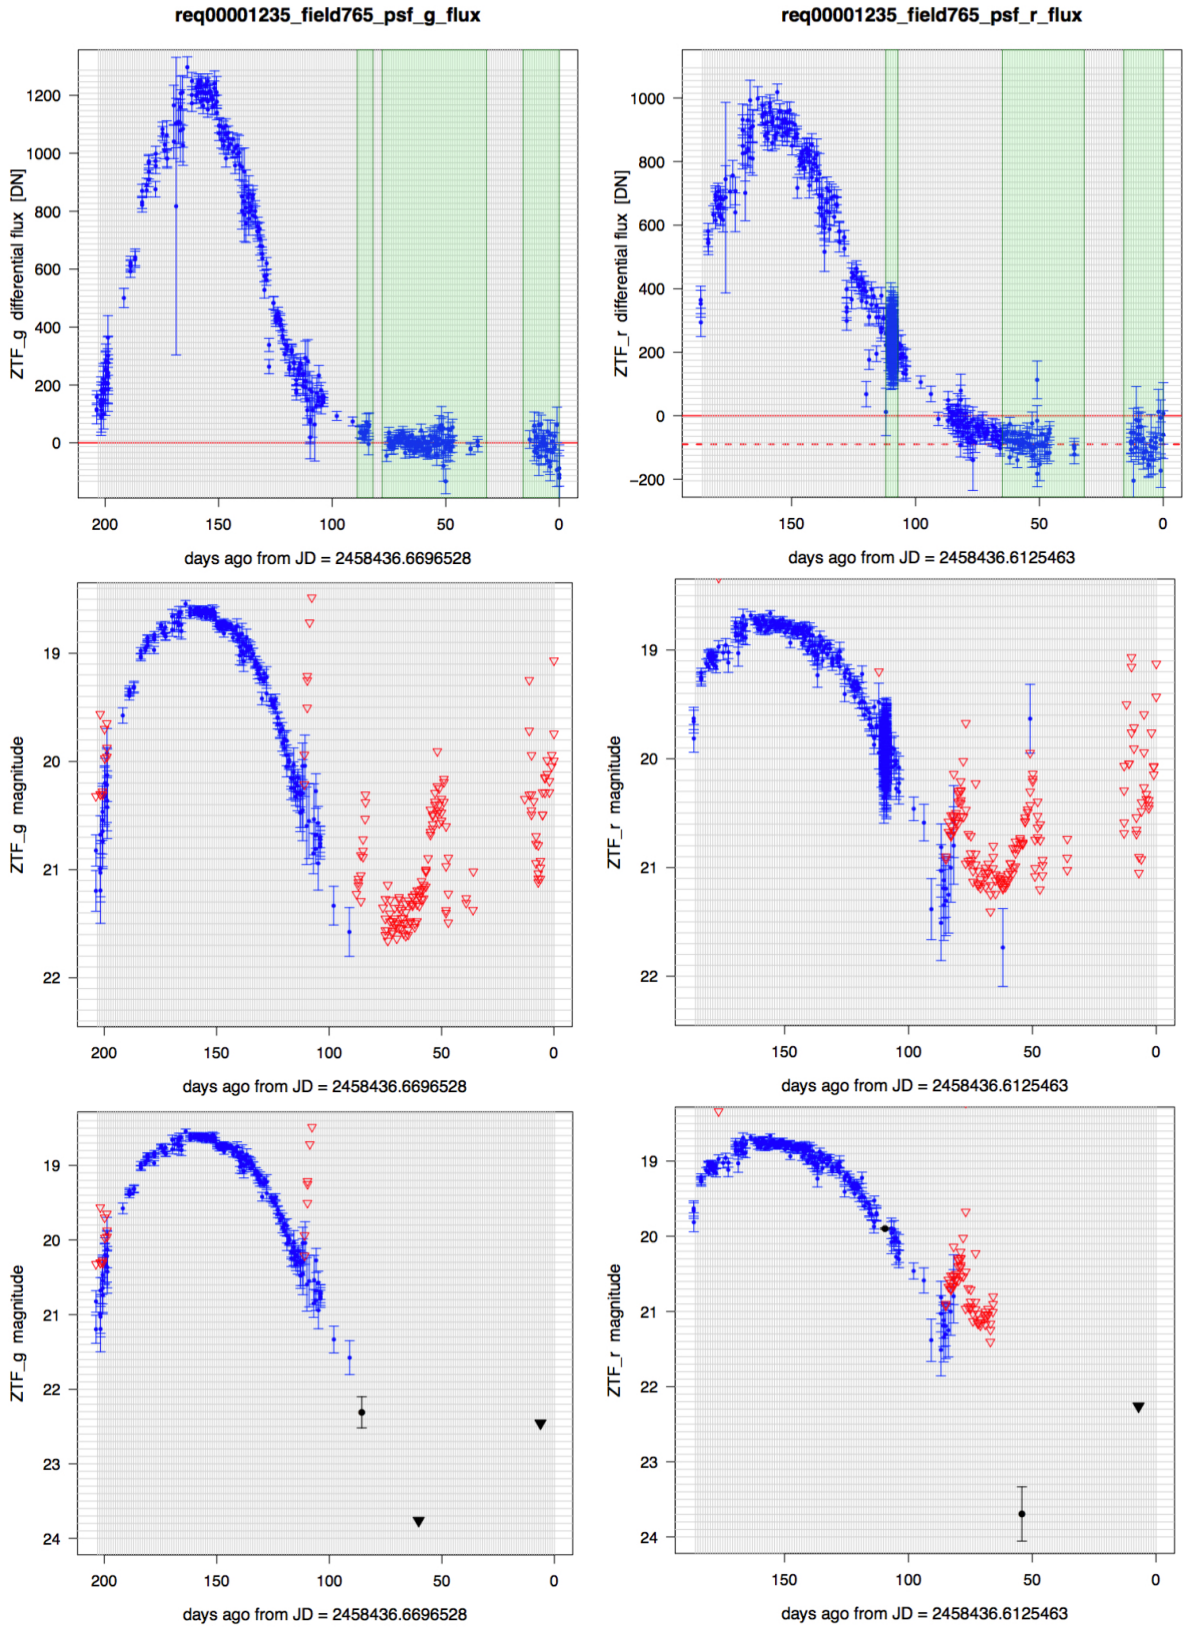
Figure 1: top to bottom row: (i) *g* **&** *r* **lightcurves in flux space; (ii) in calibrated magnitudes where red triangles denote single-epoch upper limits; (iii) also in magnitudes but combining the measurements in the shaded windows shown in the flux plots (black symbols).**

11 Validating photometric uncertainties and possible rescaling

Despite valiant efforts to ensure the 1-σ uncertainties in the PSF-fit fluxes (*forcediffimfluxunc*) are correct, it is advised to check these are still plausible, at least with respect to the photometric repeatability expected at a given magnitude, for the *field* and *filter* being considered, i.e., in the frequentist sense.

The measurement uncertainties are based on propagating a semi-empirical model of the statistical (random) noise expected in the detector pixels through the difference- (and reference-) image pipelines. This model does not account for possible systematics, for example, errors in PSF estimation, photometric zero-point and color-term calibrations, flatfielding errors, astrometric calibrations (which determine ultimate PSF-placement), and/or user-error in the supplied target position.

First, we advise checking the distribution of the PSF-fit reduced χ^2 values (*forcediffimchisq*) for a given *field* and *filter*, following any quality filtering (see note (i) in Section 8). This can also be examined as a function of flux (*forcediffimflux*). The average, *<forcediffimchisq>*, should be ≈ 1. If not, we advise multiplying the raw *forcediffimfluxunc* values by √*<forcediffimchisq>.* Any flux-dependent *forcediffimchisq* could mean the [e-/DN] gain-factor for estimating the Poisson-noise component is not properly tuned or the PSF shape and hence flux scaling is not correct. Ensuring *<forcediffimchisq>* ≈ 1 only means the pixel uncertainties are consistent, in a statistical sense, with the PSF-fit residuals *per epoch*. It will not catch and correct for possible systematics across the epochs (see below).

Following the *forcediffimchisq* check and possible rescaling, you can stop here or continue with another check to account for possible temporal systematics. This can be done by computing a robust RMS (e.g., via 0.5×[84th – 16th percentile]) in the signal-to-noise (S/N) ratios (*forcediffimflux/forcediffimfluxunc* values) of the "stationary-flux" epochs used to infer your baseline level (Section 10). These ratios would be computed following any baseline correction to the *forcediffimflux* values and any prior scaling of the *forcediffimfluxunc.* If the uncertainties are plausible, this RMS should be ≈ 1. If not, the *forcediffimfluxunc* values can be multiplied by this RMS. This will ensure the resulting distribution in S/N will have zero-mean and unit variance. Beware of epochs with bad data (see note (i) in Section 8). This assumes the noise is ≈ stationary during these historical/future epochs and they sample, to first order, the same underlying background and conditions as in epochs where the transient was detected. Note that this check alone will not validate the Poisson-noise component in epochs when the transient has high signal-to-noise.

A final coarse check (for validating specifically the Poisson-noise-dominated regime) can be performed by comparing against plots of prior-calibrated photometric-repeatability (stack-RMS) versus magnitude[1]. This will require you to first convert the fluxes to calibrated magnitudes (see Section 12 or 13). **Beware:** these priors will be *field* and *filter* dependent. A future version of the forced photometry software will perform an internal check against these priors with automatic rescaling of the uncertainties if necessary.

---

[1] For example, see Figures 9 and 10 in https://zwicky.tf/3o9

12   Generating *differential*-photometry lightcurves with non-detections

It is assumed you have corrected your differential fluxes (*forcediffimflux*) for any non-zero baseline (Section 10) and validated and possibly rescaled the uncertainties (*forcediffimfluxunc*; Section 11). If the uncertainties are over- or underestimated, your significance levels and hence derived upper limits will be wrong. I.e., "5-sigma" only has meaning if "sigma" is correct, that is, it captures all noise sources that may have corrupted the measurements.

The pseudo-code for obtaining calibrated magnitudes is below. This uses the following two parameters:

SNT – the signal-to-noise threshold for declaring a measurement a "non-detection" so it can be assigned an upper-limit.
SNU – the actual signal-to-noise ratio to use when computing a "SNU-sigma" upper-limit.

We recommend SNT = 3 and SNU = 5. For details on why these were picked, see:
http://web.ipac.caltech.edu/staff/fmasci/home/mystats/UpperLimits_FM2011.pdf

```
if( (forcediffimflux / forcediffimfluxunc) > SNT )
   # we have a "confident" detection, compute and plot mag with error bar:
   mag = zpdiff — 2.5*log₁₀[forcediffimflux]
   σ_mag = 1.0857* forcediffimfluxunc / forcediffimflux
else
   # compute flux upper limit and plot as arrow:
   mag = zpdiff — 2.5*log₁₀[SNU*forcediffimfluxunc]
```

Example lightcurves in *g* and *r* following this construction are shown in the second row of Figure 1.

13   Generating *absolute*-photometry lightcurves for variable sources

Transients are generally defined as sudden events where signal suddenly appears out of the noise and then disappears. The goal is to analyze the energy released during the explosive phase, in excess of any pre-existing signal. For continuous (a)periodic variables or transients associated with pre-quiescent sources with detectable signal, one may want to characterize the *absolute* (or total) signal versus time. Here, you will need an estimate of the source flux from the reference-image that was used to generate the difference images on which the forced photometry was performed.

Photometry and quality metrics for the nearest reference image source are given by columns prefixed by "*nearestref*…" in the output table. These are from the archived PSF-fit reference image catalog. You will need to threshold *dnearestrefsrc* to ensure a reference source exists close to your target position. We recommend *dnearestrefsrc* <~ 1 arcsec. The *nearestrefchi* and *nearestrefsharp* metrics can be used to check if the reference source is extended or PSF-like. If significantly extended, the PSF-fit photometry (*nearestrefmag*) will be erroneous. You will want to ensure it is a point source.

Below are the steps for obtaining an absolute photometric light curve in calibrated magnitudes (modulated by variations inferred from differential photometry).

(a) Compute the reference source flux and its uncertainty on the same calibration (zero-point) as the differential flux at each epoch:

nearestrefflux = 10**[0.4 * (*zpdiff* – *nearestrefmag*)];
nearestreffluxunc = *nearestrefmagunc* * *nearestrefflux* / 1.0857,
where *zpdiff*, *nearestrefmag*, and *nearestrefmagunc* are given in the output table.

(b) Compute the total flux, its uncertainty, and signal-to-noise ratio at each epoch:

$Flux_{tot}$ = *forcediffimflux* + *nearestrefflux*;
$Fluxunc_{tot}$ = $\sqrt{(\,forcediffimfluxunc^2 - nearestreffluxunc^2\,)}$;
SNR_{tot} = $Flux_{tot}$ / $Fluxunc_{tot}$.

The negative sign in the square-root follows from the (implicit) assumption that the reference flux was measured from the same reference image used to generate the "science – reference" differential flux. Hence, the noise is anti-correlated. If this is not true or if the argument in the square-root is ≤ 0 for whatever reason, we suggest adding the variances instead. This will be a conservative estimate.

(c) Using the outputs from (b) and the SNT, SNU parameters defined in Section 12, here's the pseudo-code for generating an absolute-photometry lightcurve:

```
if( SNR_tot > SNT )
   # we have a "confident" detection, compute and plot mag with error bar:
   mag = zpdiff — 2.5*log₁₀[Flux_tot]
   σ_mag = 1.0857 / SNR_tot
else
   # compute flux upper limit and plot as arrow:
   mag = zpdiff — 2.5*log₁₀[SNU*Fluxunc_tot]
```

This logic reduces to the differential-lightcurve generation logic in Section 12 when *nearestrefflux* = 0 and $Fluxunc_{tot}$ = *forcediffimfluxunc*.

The above construction to compute calibrated magnitudes assumes your source has zero color in the Pan-STARRS1 photometric system (g_{PS1} – r_{PS1} = 0 for g and r flux measurements; r_{PS1} – i_{PS1} = 0 for i flux measurements). If you have prior information on the color or spectrum of your source in these bandpasses across epochs, the magnitude computations above can include the additive term: + *clrcoeff**(g_{PS1} – r_{PS1}) for g and r measurements, or + *clrcoeff**(r_{PS1} – i_{PS1}) for i-filter measurements, where *clrcoeff* is in the output table.

14 Going deeper: combining single-epoch measurements

If you're feeling ambitious and want to make your lightcurve measurements more statistically significant and/or tighten your upper-limits, you can attempt to combine the flux measurements within carefully selected time-windows using some optimal method. This procedure assumes your single-epoch differential fluxes have been corrected for any non-zero baseline (in the case of transients; see Section 10) and uncertainties have been validated and rescaled if necessary (Section 11).

First, you will need to place the fluxes on the same (fiducial) photometric zeropoint (call this "*zpavg*"). You can pick any value for zpavg. A value that's close to the epoch-dependent zeropoints (*zpdiff$_i$* values) will suffice. You will then rescale the input fluxes and uncertainties (*forceddiffimflux, forceddiffimfluxunc*) as follows:

$$forceddiffimflux\_new_i = forceddiffimflux_i * 10^{0.4(zpavg - zpdiff_i)}$$

$$forceddiffimfluxunc\_new_i = forceddiffimfluxunc_i * 10^{0.4(zpavg - zpdiff_i)}$$

One method for combining the measurements is to assume the underlying source signal *is stationary* within a time-window and collapse the rescaled single-epoch fluxes (*forcediffimflux_new$_i$*) therein using an inverse-variance weighted average:

$$Flux = \frac{\sum_i w_i forceddiffimflux\_new_i}{\sum_i w_i},$$

where

$$w_i = \frac{1}{forceddiffimfluxunc\_new_i^2}.$$

The uncertainty in *Flux* is:

$$Fluxunc = \left[\sum_i w_i\right]^{-1/2}.$$

The last line is effectively $\approx forceddiffimfluxunc\_new/\sqrt{n}$ if one assumes the *n* windowed single-epoch flux measurements have approximately the same uncertainty equal to some constant value or RMS. Therefore, the improvement in signal-to-noise assuming uncorrelated inputs is no more than a factor of $\approx \sqrt{n}$. The important assumption is that the underlying source signal is constant within the window. It may appear constant within measurement error, but when many measurements are available, you can attempt to bin them in various ways to tease out possible hidden trends in the signal. A moving (weighted) average may also work, using a properly tuned window size. There's also a large collection of methods on local-polynomial regression fitting.

You can also attempt to collapse the measurements within windows by fitting a prior model of flux versus time, i.e., if you have prior (or contextual) knowledge from other higher signal-to-noise data. The advantages of combining measurements in source-space across epochs as opposed to co-adding entire images in time-ordered slices are: (i) speed and (ii) greater flexibility when combining the data.

Having computed the window-collapsed fluxes (*Flux* values), you can convert them to calibrated magnitudes with upper-limits using the formalism in Section 12 (for transients) or Section 13 (for absolute photometry on variables). The only difference is that you will need to use the *zpavg* value assumed when rescaling the input fluxes (see above). The top row in Figure 1 shows example windows (shaded green) where the single-epoch flux measurements are combined. These were combined using a weighted average (as above) and the calibrated photometry with upper limits are shown in the bottom row.

15 Update history

Date	Version	Changes
2022-05-06	2.3	Contact ztf@ipac.caltech.edu JD defaults to from beginning of survey to present time