

**OPEN ACCESS**

Design and Operation of the ATLAS Transient Science Server

K. W. Smith¹ , S. J. Smartt¹ , D. R. Young¹, J. L. Tonry², L. Denneau², H. Flewelling², A. N. Heinze², H. J. Weiland², B. Stalder³, A. Rest^{4,5}, C. W. Stubbs⁶, J. P. Anderson^{7,8}, T.-W. Chen⁹, P. Clark¹, A. Do², F. Förster^{8,10}, M. Fulton¹, J. Gillanders¹, O. R. McBrien¹, D. O'Neill¹, S. Srivastav¹, and D. E. Wright¹¹

¹ Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast, BT7 1NN, UK

² Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA

³ Vera C. Rubin Observatory Project Office, 950 N Cherry Ave, Tucson, AZ 85719, USA

⁴ Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA

⁵ Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

⁶ Department of Physics, Harvard University, Cambridge, MA 02138, USA

⁷ European Southern Observatory, Alonso de Córdova 3107, Casilla 19 Santiago, Chile

⁸ Millennium Institute of Astrophysics, Chile

⁹ The Oskar Klein Centre, Department of Astronomy, Stockholm University, AlbaNova, SE-10691 Stockholm, Sweden

¹⁰ Center for Mathematical Modeling, Universidad de Chile, Chile

¹¹ School of Physics and Astronomy, 116 Church Street SE, University of Minnesota, Minneapolis, MN 55455, USA

Received 2020 March 16; accepted 2020 May 15; published 2020 June 23

Abstract

The Asteroid Terrestrial impact Last Alert System (ATLAS) system consists of two 0.5 m Schmidt telescopes with cameras covering 29 square degrees at plate scale of 1.86 arcsec per pixel. Working in tandem, the telescopes routinely survey the whole sky visible from Hawaii (above $\delta > -50^\circ$) every two nights, exposing four times per night, typically reaching $o < 19$ magnitude per exposure when the moon is illuminated and $c < 19.5$ magnitude per exposure in dark skies. Construction is underway of two further units to be sited in Chile and South Africa which will result in an all-sky daily cadence from 2021. Initially designed for detecting potentially hazardous near earth objects, the ATLAS data enable a range of astrophysical time domain science. To extract transients from the data stream requires a computing system to process the data, assimilate detections in time and space and associate them with known astrophysical sources. Here we describe the hardware and software infrastructure to produce a stream of clean, real, astrophysical transients in real time. This involves machine learning and boosted decision tree algorithms to identify extragalactic and Galactic transients. Typically we detect 10–15 supernova candidates per night which we immediately announce publicly. The ATLAS discoveries not only enable rapid follow-up of interesting sources but will provide complete statistical samples within the local volume of 100 Mpc. A simple comparison of the detected supernova rate within 100 Mpc, with no corrections for completeness, is already significantly higher (factor 1.5 to 2) than the current accepted rates.

Key words: surveys – minor planets – asteroids – supernovae – variable stars

Online material: color figures

1. Introduction

The use of ground-based wide-field telescopes equipped with large format digital cameras has opened a window on the time varying universe over the last decade. The late 1990s and first decade of the twentieth century saw relatively high redshift surveys for type Ia supernovae (e.g., Schmidt et al. 1998; Astier et al. 2006) and low redshift surveys focused on galaxy targeted searches (Leaman et al. 2011; Li et al. 2011b) The former typically employed what were then regarded as wide-field CCD

camera mosaics on 2–4 m telescopes (of order 0.5 to 1° diameter field-of-view; Perlmutter et al. 1997) with carefully selected targets for spectroscopic follow-up on the new and precious 8–10 m aperture facilities (Howell et al. 2005; Wood-Vasey et al. 2007).

Low redshift supernovae (SNe) were found primarily by narrow field, small telescopes (0.5–1 m; Pignata et al. 2009; Leaman et al. 2011) and by well equipped and experienced amateur astronomers. These bright events were subject to follow-up by many teams around the globe with moderate sized telescopes and many were observed and detected from X-rays (Immler & Lewin 2003; Immler et al. 2006; Soderberg et al. 2008) to ultraviolet (Brown et al. 2009; Bufano et al. 2009) and from the mid-infrared (Kotak et al. 2006; Fabbri et al. 2011) to

 Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

radio (Sutaria et al. 2003; Ryder et al. 2004; Stockdale et al. 2007; Soderberg et al. 2010; Margutti et al. 2017). Both the low and high redshift surveys were highly successful in their respective goals and in particular in allowing understanding of the physics of type Ia SNe to shape their use in cosmology.

The successful techniques of both these approaches have been combined in the last decade to allow wide-field, virtually all sky, surveys of the local universe with 10 to 100 sq degree cameras (Quimby 2006; Drake et al. 2009; Law et al. 2009; Baltay et al. 2013; Chambers et al. 2016; Holoien et al. 2017; Tonry et al. 2018; Bellm et al. 2019). This allows a large volume to be searched and importantly this volume samples the local universe, enabling extensive multi-wavelength monitoring of bright objects. Photon starvation from distant objects tends to raise barriers to physical understanding and development of radiative transfer simulations, shock physics models and theories of powering mechanisms (e.g., Kromer & Sim 2009; Kasen & Bildsten 2010; Rabinak & Waxman 2011). The rate of discovery of supernovae and transients within $z \lesssim 0.1$ increased dramatically and significant numbers of transients were found either in faint host galaxies or in environments far from their high mass hosts (tens of kiloparsecs). The combination of the sheer number of discoveries and different explosion environments probed in comparison to the historic surveys revealed an unexpected and remarkable diversity in types of stellar explosions and merger events occurring in the local universe. An important scientific advance was combining the discovery of transients with rapid spectroscopic follow-up to determine redshifts, association with host galaxies and chemical composition of the ejecta (e.g., Quimby et al. 2011; Inserra et al. 2013; Gal-Yam et al. 2014). Furthermore, as these low redshift transients were relatively bright then multi-wavelength, time domain follow-up has unveiled new explosion physics. And remarkably for the first time, a particle messenger other than a photon provided an alternative alert that triggered exact identification of the source with optical-near-infrared telescopes: the joint discovery of GW170817, GRB170817A and AT2017gfo (Abbott et al. 2017a, 2017c, 2017b).¹²

A critical piece of a wide-field imaging survey and discovery facility is the computing software and hardware that processes the data to produce a stream of clean, real, astrophysical variables and transients. This comprises two parts, the first being the detector detrending, source identification and astrometric and photometric calibration (e.g., Waters et al. 2016). The principles of such processing are based on well known methods applied to narrow-field telescopes, although wide-field camera calibration has its own challenges (Magnier et al. 2016a, 2016c). The now ubiquitous use of difference

imaging has provided a method to identify and quantify variables and transient sources (Alard & Lupton 1998; Becker 2015; Zackay et al. 2016). However at that point the night’s catalogs of objects still requires further processing to provide a useful information stream that can be used scientifically, and immediately, by astronomers. This means filtering the bogus detections that inevitably exist on the difference images of a CCD, associating the detections together and assimilating them into distinct objects and providing basic classifications (Brink et al. 2013; Wright et al. 2015; Goldstein et al. 2015; Narayan et al. 2018; Mahabal et al. 2019). Real astrophysical detections on a difference image comprise satellite trails, minor planets, fast moving near earth asteroids (point sources and trails), variable and erupting stars in the Milky Way and Local Group galaxies, novae, active galactic nuclei (AGN), tidal disruption events and extragalactic supernovae, stellar mergers and afterglows of gamma-ray-bursts. Typically most wide-field survey and discovery projects have built a complex software system that provides a means of extracting and classifying the sources and feeding them to spectroscopic follow-up programmes (Smartt et al. 2015; Blagorodnova et al. 2018; Graham et al. 2019). Classification of the lightcurves to filter and select objects of interest has had some success (Lochner et al. 2016; Muthukrishna et al. 2019).

We will use the term “transient science server” for this system of software and computing hardware. The back end of these systems which focused on the assimilation of individual detections into object lightcurves and classification of the objects from their host object (star or galaxy) have recently been termed “brokers.” The Zwicky Transient Facility (Bellm et al. 2019) issues a public alert stream of transient and variable sources detected on their difference images (Patterson et al. 2019), with a view to testing methods for the upcoming Rubin Observatory Legacy Survey of Space and Time (LSST). These alert packets are ingested and processed by a series of “brokers” to provide added value such as context (variable star and AGN identification, host galaxy, redshift, association with known transients, variables or outbursting source) and lightcurve classification. Currently four systems ingest these data and provide different functionality to users, Lasair¹³ (Smith et al. 2019b) ANTARES¹⁴ (Narayan et al. 2018), ALerCE¹⁵ (Förster 2019, white paper 2020, in preparation) and MARS.¹⁶

In this paper we describe how the ATLAS Transient Science Server functions and operates on a daily basis. The ATLAS system was defined in Tonry et al. (2018) and the first major data release (a variable star catalog) was described in Heinze et al. (2018). ATLAS is designed, funded, and operated using

¹² While multi-messenger astrophysics did begin with solar and supernova neutrinos, neither of these photon emitting objects were discovered due to the directional information of their neutrino detections. AT2017gfo was discovered by astronomers directly following an alert from a non-photon detector.

¹³ <https://lasair.roe.ac.uk/>

¹⁴ <https://antares.noao.edu/>

¹⁵ <https://alerce.online/>

¹⁶ <https://Mars.lco.global/>

NASA funds to find dangerous asteroids which are routinely detected and submitted to the Minor Planet Center (Denneau, L. et al. 2020, in preparation). The nature of the search for moving objects yields detections of variable and transient objects for free. To date, around 25 science papers have been published on transient and variable objects from the ATLAS data processing. This includes the discovery of AT2018cow (Prentice et al. 2018) a candidate for a pulsational pair instability supernova (Chen et al. 2018) a possible white dwarf—neutron star merger (McBrien et al. 2019) and the lowest luminosity supernova ever discovered (Srivastav et al. 2020). ATLAS has covered the gravitational wave skymaps of a number of LIGO Virgo events, often covering a complete fraction of the map and discovering candidates rapidly. For example we discovered a fast fading transient in the skymap of GW170104, which turned out to be the afterglow of a gamma-ray burst (GRB170105A) but was discovered without the high energy trigger (Stalder et al. 2017). After describing the system in Sections 2–5, we briefly outline the scientific return in Section 6 and improvements for the future in Section 7.

2. The ATLAS Telescopes and Survey Description

ATLAS currently consists of two identical telescopes which are a variant of a “Wright Schmidt” design. They consist of a 0.5 m Schmidt corrector, a 0.65 m spherical primary mirror, and a three lens field corrector sitting in front of the filter changer and camera. They are situated at the Haleakala Observatory (HKO) on Maui and the Mauna Loa Observatory (MLO) on the Big Island of Hawaii. The first telescope (on HKO) has been fully operational since 2015, and the second (on MLO) was commissioned at the end of January 2017. The telescopes use two main filters in normal survey operations. One is “cyan” (we refer to it as c) which is used during dark time and is roughly equivalent to a composite PS1 $g + r$ (420–650 nm). The other is “orange” (or o), used during bright time and roughly equivalent to $r + i$ (560–820 nm). The Haleakala unit hereafter referred to as ATLAS-HKO has both the c and o filters and the the Mauna Loa unit (ATLAS-MLO) has only the o filter.

The cameras on each telescope (which we call an “ACAM”) each contain a STA 1600 $10,560 \times 10,560$ pixel single CCD detector. The pixel scale is 1.86 arcsec per pixel, and excluding the borders of the CCDs (a 30 pixel boundary) the net field of view is 5.4×5.4 , equating to 29 square degrees. In normal survey operations, the telescope schedules are jointly planned to tile the sky with 30 second exposures. As described in Tonry et al. (2018), the overhead from shutter closure to beginning of the next frame is 10 s. This is determined by the readtime of the chip (which is ~ 9 s) since the ATLAS telescope mounts are fast and reliable. As described in Tonry et al. (2018) for telescope movements smaller than 45° it takes 6.5 ± 0.8 s to

slew and resume tracking. Hence a total duration (expose plus overheads) of 40 s per frame means that each telescope can take a *maximum* of about 1000 exposures per night during the longest night duration of 11.5 hr. In practice, between 800 and 1000 frames per telescope per night are routinely taken (see Figure 1), with 900 frames covering 26,100 square degrees. This means each of the telescopes could comfortably cover the entire visible sky from Hawaii ($\sim 24\,500$ deg²: north celestial pole down to about -45 degrees decl.) every night. The entire funding for ATLAS comes from NASA to find dangerous asteroids; hence operations are designed to optimize the detectability of moving objects. It is mandatory to collect 4 views of each field every night over a span of less than 1 hour in order to reach a manageable false alarm rate. This facilitates identification of moving objects, linkage of the tracklets, orbit estimation and submission to the Minor Planet Center. Hence the total sky footprint covered by each telescope, 4 times each night, is around 6500 square degrees. The repeat field coverage also enables intra night (and multi night) difference stacks to be created, allowing ATLAS to detect slow moving and static transients at $m > 20$ mag across the whole sky, including the galactic plane.

The telescopes are scheduled to scan decl. strips of the sky and, weather permitting, each will return to the same footprint area every 4 nights. Hence with two units working in tandem, the whole sky visible from Hawaii is covered every 2 nights with a 4×30 sec observing sequence. A typical 2 day footprint is illustrated in Figure 5 of Tonry et al. (2018). The southern pointing limit is a boresight of $\delta = -47.5^\circ$, and the southern egde of these exposures reach $\delta = -50^\circ$. The survey strategy has been defined to optimize for NEO cadence (e.g., Vereš et al. 2009; Tonry 2011; Tonry et al. 2018), but a 2-day, all sky cadence is an excellent survey base for early discovery and lightcurve classification of stationary transients.

3. ATLAS Data Processing and Object Detection

The data processing steps to detrend, process and to photometrically and astrometrically analyze the images are described in Tonry et al. (2018), along with all the timings for each step (see Table 3 in that paper). Each 30 s frame is processed on summit computers and it takes 310 s to apply the first initial step of detector detrending and photometric and astrometric calibration. It requires a further 1280 s to produce a matching reference image (from the ATLAS wallpaper), subtract it from the input image, detect and classify the sources and write out a file containing the detections with $S/N > 5\sigma$. Hence a total of 27 minutes is required from the time of shutter opening for an exposure to the final processed file containing a list of detections. The processing runs multi-threaded, and not in serial, hence 27 minutes is the typical delay time from observing to initially quantifying detections on the cameras.

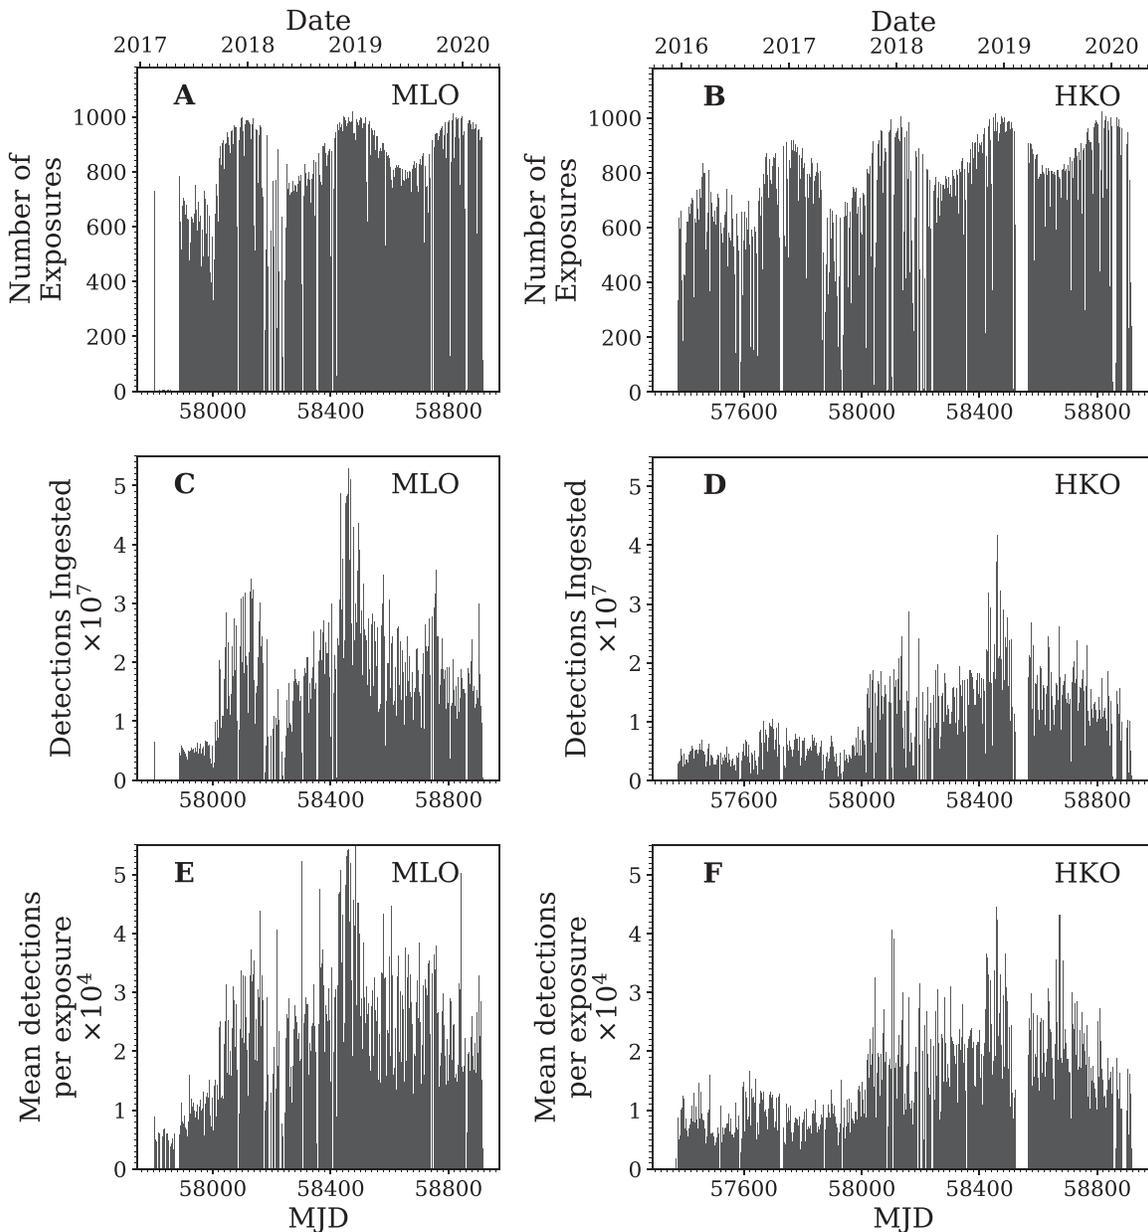


Figure 1. A, C, E: Number of exposures, number of detections ingested and average number of detections per exposure for MLO. The detections are the number of sources at 5σ or greater that are present on the difference images (in the `.ddc` files). B, D, F: The same for HKO. The gap in the HKO data was caused by an ice storm on Haleakala. See also Figure 7. The rise in the average number of objects per exposure around MJD = 58000 is because of a change of detection file schema and improvements in source extraction.

We have the capability of reducing the elapsed time by factor of ~ 2 through deployment of multi-threading on the summit computers, but currently do not do this since the primary science goal of asteroid detection has a latency of ~ 30 mins due to the 4 visits.

A key part of transient object detection is subtraction of a reference sky frame from the input image to allow for rapid and automated detection of variables, moving objects and transients (Tomaney & Crotts 1996; Alard & Lupton 1998). In

ATLAS we refer to this reference sky as the ATLAS wallpaper, and we currently rebuild this wallpaper about once per year. We store the wallpaper in files of square degree size, and routines allow any ATLAS pointing on the sky to trigger a reconstruction of a wallpaper reference image of the appropriate size and alignment. We use a modified version of the image subtraction algorithm HOTPANTS (Becker 2015) to subtract the matched, convolved reference image from the target frame.

Table 1The Contents of a `.ddc` File for a Difference Image Contains All Detections (5σ Significance or Greater), Each with the Following Measured Attributes

<code>.ddc</code> file entry	meaning
RA	astrometrically calibrated R.A. in decimal degrees of centroid
Dec	astrometrically calibrated decl. in decimal degrees of centroid
mag	AB magnitude from PSF fit
dmag	error in AB mag
x	x pixel on CCD of centroid
y	y pixel on CCD of centroid
major	PSF FWHM major axis [pixels]
minor	PSF FWHM minor axis [pixels]
phi	angle of major axis, counter-clockwise from x-axis [degrees]
det	integer indicating fixed, free, trailed, streak, or negative flux analysis
chi/N	reduced χ^2 of the PSF fit
pvr	vartest probability of source being a variable star
ptr	vartest probability of source being a moving object or astrophysical transient
pno	vartest probability of source being just a noise fluctuation
pbn	vartest probability of source being a vertically extended electronic “burn” artifact
pcr	vartest probability of source being a cosmic ray hit
pxt	vartest probability of source being an electronic crosstalk artifact
psc	vartest probability of source being a star subtraction residual “star scar”
pmv	vartest probability of source being a moving object (subset of ptr)
pkn	vartest probability of source being a known minor planet
dup	integer indicating whether this PSF measurement should be kept or culled
WPflx	forced flux measured in the wallpaper (reference sky) at the position of the transient
dflx	error in WPflx

The ATLAS reduction pipeline has a custom built point-spread-function fitting routine that runs on the difference images to produce flux measurements of all sources that are detected at 5σ or more above the background noise. This routine (written in C) is called `tphot`, and is based on the algorithms discussed in Tonry (2011) and Sonnett et al. (2013). This produces one headed ASCII table per image (internally appended with `.ddc`), and the contents are summarized in Table 1. Some of these quantities are used in an initial filtering process to remove the bogus sources on the difference images (see Section 4.3).

On a typical night, a single ATLAS telescope produces about 900 catalog files (one per exposure). Each catalog file contains on average about 20,000 detections, so about 10–20 million detections are ingested per night from the two telescopes combined. The vast majority of these sources are not astrophysically real objects and a series of quality control procedures and filters, together with machine learning algorithms are run to reject the false positives and leave real transient objects. Figure 1 illustrates the number of detections ingested from these `.ddc` files per night and example breakdowns of their classifications are in Table 2.

4. The ATLAS Transient Science Server

4.1. Object Definition and Spatial Indexing

The essential functions of a hardware and software system to deliver real transients from a synoptic wide-field survey such as ATLAS includes amalgamating detections into objects,

removing false positives, assimilating lightcurves, and associating the newly discovered objects with known catalogs. We have built the ATLAS Transient Science Server to provide all these functions. The Server is currently based on individual detections from single ATLAS exposures, but in future we intend to combine the information from the nightly difference stacks to allow us to detect transient objects earlier.

The ASCII catalog files of objects detected on the difference images (the `.ddc` files described in Section 2) are transferred to a machine at Queen’s University Belfast (QUB). This is a 28 core server with 128 GB RAM (also used as the ingestion machine). The files are transferred (using parallel download threads) in the background every 30 minutes between 04:00 UT and 22:00 UT (corresponding to the earliest start of a Hawaiian observing night at 18:00 and cleaning up after the night is done at 12:00 next day). Just before an ingest cycle begins, an extra check is done to make sure we have the very latest data set. The reduced, calibrated images and their associated subtracted frames are also transferred (also in parallel) to QUB during the Hawaiian night. Typical end to end data rates from the University of Hawaii Manoa campus to the QUB machine are 40 MBytes s^{-1} . Since individual download threads are throttled by the network infrastructure, using 8 parallel threads ensures we use the maximum available bandwidth. Each of the 8 threads achieves 5 MBytes s^{-1} and using more than 8 threads brings very little gain in download speed.

The `.ddc` file contains a header comprising 37 key-value pairs describing metadata of the exposure, such as the telescope

Table 2
Example Nightly Ingest Numbers for Two ATLAS Telescopes on the Night of MJD = 58816 (2019 November 29)

Stage	Number	(2 detection Number)
Detections on the difference images ($\geq 5\sigma$ in the .ddc files)	3.1×10^7	
Single detection objects (probably noise, edge, streaks and 1.6×10^5 known movers)	2.1×10^7	
Objects with 3 (2) or more spatially coincident detections	6.2×10^5	(2.6×10^6)
Objects with 3 (2) or more spatially coincident, good quality detections	2.2×10^4	(3.6×10^5)
Objects remaining after stars removed	9450	(68,542)
Objects remaining after AGN removed	9300	(68,068)
Objects remaining after untagged known slow movers removed (e.g., comets)	9280	(68,008)
Objects remaining after ML and $RB \geq 0.2$ applied	340	(8341)
Good objects after human scanning	40	

Note. Good quality detections are defined as having $pvr > 50$ or $ptr > 50$ and not a known or possible mover, must be positive flux and not crosstalk. The detection must also not be within 100 pixels of the edge of the detector. Note that the number of “good” objects for this night is higher than average. This was possibly because of dark time, the slightly longer winter night, and a full night of good quality difference images. The figures assuming 2 spatially coincident detections rather than 3 are also shown in brackets.

pointing coordinates, exposure MJD, exposure time, filter, PSF FWHM, zero-point and 5σ limiting magnitude. This is followed by a headed table with each entry corresponding to a single detection (see Table 1). The metadata and detections are automatically read by a C++ ingester (generated by python code which reads the .ddc structure) at QUB which connects to a local MySQL database. Detections are first checked via cone searching (see below) against previously ingested objects. If a detection is within $3''6$ (twice the ATLAS pixel size) of an existing object, it is ingested and formally associated with that object, becoming another point on the lightcurve. If there are no objects within the specified proximity of the ingested detection, a new object is created and the detection is associated with the new object. To speed up the ingest, multiple ingester processes are run in parallel using a python wrapper (to utilize all the cores of the ingestion machine).

Hence all records in the .ddc files are preserved in the database. The critical values ingested are R.A., decl., calculated HTM ID (see below), mag and dmag. Other detection properties are also ingested to facilitate post ingest cutting.

The detection data are spatially indexed using Hierarchical Triangular Mesh (Szalay et al. 2007). Since the vast majority of queries are cone searches of the order of a few arcsec in size, HTM Level 16 ($2 \times 4^{(16+1)}$ triangles, hence mean triangle area = 15.6 square arcsec) is used. The HTM value, calculated for every detection’s R.A. and decl. upon ingest, is an integer and this column is indexed inside the relational database. The C++ HTM library provides a cone searching API which returns the ranges of triangles that overlap a cone on the sky. This is used to create an SQL WHERE clause. All objects in the database that have those triangle IDs are returned to the requesting code. The cone search code then just needs to eliminate objects outside the requested radius (i.e., the edges of the resultant jagged cone are smoothed).

4.2. Selection Criteria, Machine Learning and Classification

After ingest, selection criteria and algorithms are applied to select real astrophysical transients from the detection stream.

Multiple detections: since the standard quad set of 30 s images are separated across a 1 hr window of observation, spatial coincidence of the sources is used to reject some moving objects and other non-astrophysical contaminants. Three or more good quality, co-spatial detections are required to define an object and all three must be on the same night. They must also not be within 100 pixels of the edge of the detector. There is no selection applied to the coherence of the magnitudes of three detections, and we do find rapidly evolving transients (over the course of 1 hour), all of which we have identified as M-dwarf flares apart from one GRB afterglow (Stalder et al. 2017).

The main benefit of reducing the triggering threshold to 2 detections rather than 3 would be that we may discover and detect objects earlier. Given the 2 day cadence of ATLAS, objects with only 2 detections on a single night and no other 5σ detections are either real transients that just rise to the sensitivity limit in that window, earlier detections of objects that continue to rise (and are subsequently detected with ≥ 3 detections afterwards) or spurious artefacts. We carefully checked 50 real objects from 2020 April, and found that 40% could potentially have been discovered earlier had we used a 2 detection minimum criterion rather than 3. The mean time difference was a gain of 4 days. However while that means 40% of objects have the potential to have been discovered earlier, it does not mean they would definitely have been promoted. The early detections would need both a high RB score (see below) and to be confidently judged by human scanners as real. Our human scanners will often wait for an extra night of data if the images are ambiguous. This gives a sense of the maximum gain if we were to go to a criterion of

2 images. It is possible, but the penalty is a large increase in false positives. For example, the figures shown in brackets in Table 2 illustrate that human scanners would have been presented with over 24 times more objects. Seven objects submitted to the IAU TNS by other surveys (PS1 and ZTF), all fainter than 19 mag at peak, would have been triggered with a 2 detection threshold by ATLAS, all with a high RB factor, but *after* the other surveys had already reported them.

Source classification: a pixel-based classification of candidate source detections in the ATLAS difference images is performed using a program called `vartest`, written in C (Heinze, A. N. et al. 2020, in preparation), which runs in Hawaii after the difference image detections have been extracted. `vartest` attempts to classify sources into several categories of real and spurious objects. It is not a machine learning code, but rather uses a logical cascade of manually tuned thresholds and analytical probability estimators applied to various pixel-based metrics. These metrics include centroid offset and flux ratio between the original and difference images; elongation of the PSF; pixel-coordinate association with bright stars; ratio of nearby bright and dark pixels; and many others. The classification is done probabilistically, with $p = 0$ being the lowest probability and $p = 999$ being the highest. The categories of real objects are known asteroids (`pkn`), unknown asteroids or transients (`ptr`), and variable stars (`pvr`). Objects with nonzero values of `ptr` are also assigned a probability of being a fast-moving (i.e., trailed) detection (`pmv`). The categories of spurious detections are noise (`pno`), electronic column artifact or “burn” (`pbn`), cosmic ray (`pcr`), electronic crosstalk (`pxt`), and star subtraction residual or “star scar” (`psc`). These are summarized in Table 1. The sum of the values for a single detection (excluding `pmv` and `pkn`, which are subsets of `ptr`) will not exceed 999. `vartest` also identifies duplicate detections in its input file, and attempts to indicate the “best” detection (stored in the `.ddc dup` column).

From the `vartest` set of parameters we define potential “good quality” detections as those with values `pvr` > 50 or `ptr` > 50. If either of these two criteria are not met, the detection does not count toward the 3 minimum good quality detections within a night required to define an object. Additionally, the detection must NOT belong to a known mover (`pkn` < 500) or a possible mover (`pmv` < 500) or be marked as crosstalk (`pxt` < 500). The value of `det` must not indicate negative flux and the value of `dup` must not indicate a deprecated duplicate value. Figure 2 shows histograms of the vector values for the full (dark) night of MJD = 58816 (2019 November 29). The data rows used in the plots are only those which comprised objects of 3 or more distinct detections. The dotted line on the `ptr` and `pvr` plot shows where the quality cut is currently set.

Object recentering: if an object passes the above criteria, the detection closest to the mean coordinates of the object is identified as the representative detection for that object. This ensures that subsequent detections are associated with the

correct objects. (The code is periodically rerun to make sure the object is properly centered.)

Ephemeris checker: an initial cross-match with known minor planets is done in Hawaii (within `vartest`) using a local copy of the Minor Planet Center (MPC) orbital elements and a 9.3 arcsecond (5 pixel) association radius. This produces an entry in the `.ddc`, labeled as `pkn` which is the probability that the detection is that of a known asteroid. However, this occasionally misses some asteroids, and does not currently label any comets and thus we implement a further check. Three times weekly, at QUB, the Minor Planet Center (MPC) orbital element database is downloaded and converted into an XEphem database file. At the same time, the comet elements, provided by the MPC in XEphem format already, are appended to the converted downloaded file. After each batch of data is ingested, the latest detection of an object that passes previous cuts is passed to code which locally generates the solar system object ephemeris from the orbital element data based on the MJD of the exposure. The detections are grouped by MJD to minimize the number of times an ephemeris needs to be generated. A cone search is then run against the ephemeris to identify any solar system object (asteroid or comet) that may lie within a specified search radius (currently 10 arcsec). If it lies within the search radius, the object is tagged as a probable known mover and labeled with the solar system object ID. The ephemeris check is done offline without interrogating the MPC servers (apart from the periodic data downloads) to minimize MPC load. Additionally, moons of the giant planets are not currently recorded in the MPC database of orbital elements and therefore fail to be found in this basic check. For this reason, the positions of all the satellites of the giant planets have been generated¹⁷ for every hour for the next 10 yr and these are stored in a local database table at QUB. A simple cone search is run against this for the most recent detection of each object, with an additional temporal check of any matched object. If a matched planetary satellite is within 20 arcsec and half a day of the timestamp of the detection, the object is labeled as a probable mover.

Although every effort is made to reject moving objects from the flow of data to human scanners, two serendipitous solar system discoveries have been made with the ATLAS Transient Science Server. We discovered unusual activity of the asteroid (6478) Gault (Kleyna et al. 2019; Smith et al. 2019a), arising from a comet-like “tail” which provided a locus of detections offset from the asteroid itself. We also discovered a distant previously unknown long period comet C/2019 K7 (Smith) (Heinze 2019a, 2019b), at 4.5 au perihelion which was therefore slow moving and made its way through the Transient Science Server as a possible stationary transient.

¹⁷ via Giorgini, JD and JPL solar system Dynamics Group, NASA/JPL Horizons On-Line Ephemeris System, <http://ssd.jpl.nasa.gov/?horizons>, data retrieved 2019 February 21.

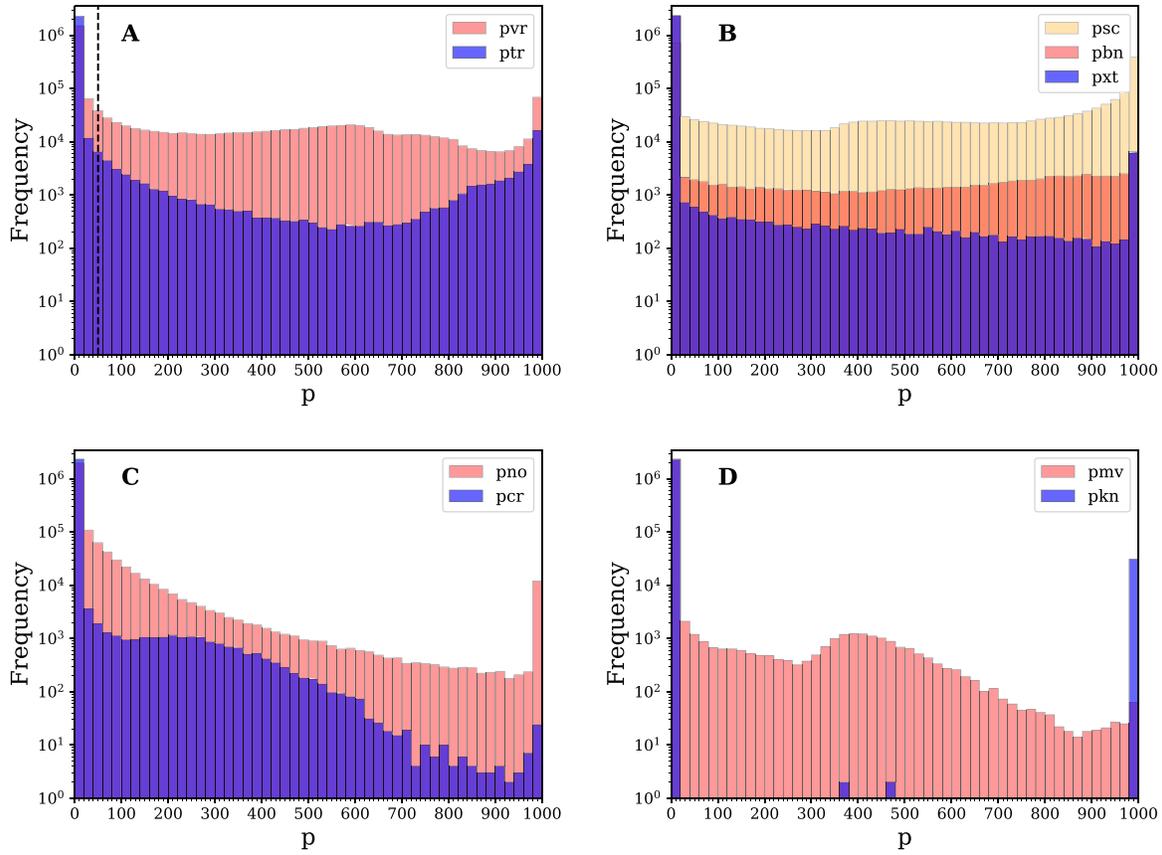


Figure 2. Histograms of the *vartest* detection values for night of MJD = 58816 (2019 November 29). The output of *vartest* produces values for each parameter $0 < p < 999$. Of the 31 million detections that night, only those that form objects comprising at least three distinct intra night detections (2.4 million detections, 602,000 objects) are shown. (A): histogram of the values of p_{vr} and p_{tr} . The dotted line is the current threshold above which a detection is regarded as good quality for the purpose of detecting variable objects and transients. (B): histogram of the values of p_{sc} , p_{bn} and p_{xt} . (C): histogram of the values of p_{no} and p_{cr} . (D): the values of p_{kn} and p_{mv} are treated separately and the histogram of values is plotted. The value of p_{kn} is normally 999 or 0 (i.e., a detection is associated with a known object or it is not), whereas the value of p_{mv} will span the entire range. Of 160,000 known minor planet detections that night, 31,000 detections were of movers slow enough to be aggregated into 4800 objects comprising 3 or more detections. (A color version of this figure is available in the online journal.)

Sherlock: once an object is defined, a boosted decision tree algorithm (internally known as *Sherlock*) mines a library of historical and on-going astronomical survey data and attempts to predict the nature of the object based on the resulting crossmatched associations found. One of the main purposes of this is to identify variable stars, since they make up about 50% of the objects (see Table 2), and to associate candidate extragalactic sources with potential host galaxies. The full details of this general purpose algorithm and its implementation will be presented in an upcoming paper (Young, D. R. et al. 2020, in preparation), and we give an outline of the algorithm as implemented within the ATLAS Transient Science Server here.

The library of catalogs contains data sets from many all-sky surveys such as the major Gaia DR1 and DR2 (Gaia Collaboration et al. 2016, 2018), the Pan-STARRS1 Science Consortium surveys (Chambers et al. 2016; Flewelling et al. 2016; Magnier et al. 2016a, 2016c, 2016b) and the catalog of probabilistic classifications of unresolved point sources by

Tachibana & Miller (2018) which is based on the Pan-STARRS1 survey data. The SDSS DR12 PhotoObjAll Table, SDSS DR12 SpecObjAll Table (Alam et al. 2015) usefully contains both reliable star-galaxy separation and photometric redshifts which are useful in transient source classification. Extensive catalogs with lesser spatial resolution or color information that we use are the GSC v2.3 (Lasker et al. 2008) and 2MASS catalogs (Skrutskie et al. 2006). *Sherlock* employs many smaller source-specific catalogs such as Million Quasars Catalog v5.2 (Flesch 2019), Veron-Cett AGN Catalogue v13 (Véron-Cetty & Véron 2010), Downes Catalog of CVs (Downes et al. 2001), Ritter Cataclysmic Binaries Catalog v7.21 (Ritter & Kolb 2003). For spectroscopic redshifts we use the GLADE Galaxy Catalogue v2.3 (Dálya et al. 2018) and the NED-D Galaxy Catalogue v13.1.¹⁸ *Sherlock* also has the ability to remotely query the NASA/IPAC Extragalactic

¹⁸ <https://ned.ipac.caltech.edu/Library/Distances/>

Database, caching results locally to speed up future searches targeting the same region of sky, and in this way we have built up an almost complete local copy of the NED catalog. More catalogs are continually being added to the library as they are published and become publicly available.

At a base-level of matching Sherlock distinguishes between transient objects *synonymous* with (the same as, or very closely linked, to) and those it deems as merely *associated* with the cataloged source. The resulting classifications are tagged as *synonyms* and *associations*, with synonyms providing intrinsically more secure transient nature predictions than associations. For example, an object arising from a variable star flux variation would be labeled as *synonymous* with its host star since it would be astrometrically coincident (assuming no proper motion) with the cataloged source. Whereas an extragalactic supernova would typically be *associated* with its host galaxy—offset from the core, but close enough to be physically associated. Depending on the underpinning characteristics of the source, there are 7 types of predicted-nature classifications that Sherlock will assign to a transient:

1. Variable Star (VS) if the transient lies within the synonym radius of a cataloged point-source,
2. Cataclysmic Variable (CV) if the transient lies within the synonym radius of a cataloged CV,
3. Bright Star (BS) if the transient is not matched against the synonym radius of a star but is associated within the magnitude-dependent association radius,
4. Active Galactic Nucleus (AGN) if the transient falls within the synonym radius of cataloged AGN or QSO.
5. Nuclear Transient (NT) if the transient falls within the synonym radius of the core of a resolved galaxy,
6. Supernova (SN) if the transient is not classified as an NT but is found within the magnitude-, morphology- or distance-dependent association radius of a galaxy, or
7. Orphan if the transient fails to be matched against any cataloged source.

For ATLAS the synonym radius is set at $1.5''$. This is the crossmatch-radius used to assign predictions of VS, CV, AGN and NT. The process of attempting to associate a transient with a cataloged galaxy is relatively nuanced compared with other crossmatches as there are often a variety of data assigned to the galaxy that help to greater inform the decision to associate the transient with the galaxy or not. The location of the core of each galaxy is recorded so we will always be able to calculate the angular separation between the transient and the galaxy. However we may also have measurements of the galaxy morphology including the angular size of its semimajor axis. For ATLAS we reject associations if a transient is separated more than 2.4 times the semimajor axis from the galaxy, if the semimajor axis measurement is available for a galaxy. We may also have a distance measurement or redshift for the galaxy enabling us to convert angular separations between transients

and galaxies to (projected) physical-distance separations. If a transient is found more than 50 Kpc from a galaxy core the association is rejected.

Once each transient has a set of independently crossmatched synonyms and associations, we need to self-crossmatch these and select the most likely classification. The details of this will be presented in a future paper (Young, D. R. et al. 2020, in preparation). Finally the last step is to calculate some value added parameters for the transients, such as absolute peak magnitude if a distance can be assigned from a matched cataloged source, and the predicted nature of each transient is presented to the user along with the lightcurve and other information (see Figure 5).

We have constructed a multi-billion row database which contains all these catalogs. It currently consumes about 4.5 TB and sits on a separate, similarly specified machine to that of the ATLAS database. It will grow significantly as new catalogs are added (e.g., Pan-STARRS 3π DR2, VST and VISTA surveys, future Gaia releases etc.).

Cross-matching with other surveys: the flagged objects are also cross-matched against a database of transient objects discovered by other surveys and tagged if they are already known. We compile a local database from the IAU TNS and also cross-match against objects from ZTF (Bellm et al. 2019), which are contained in the Lasair broker and are likely to be supernova or extragalactic transients of some sort (Smith et al. 2019b). Hence users know almost instantaneously if an object has been discovered by another survey and registered in the TNS as a discovery, or if it has multiple detections in the ZTF public stream. This crossmatch runs at the end of every ingestion cycle (at least 4 times a day) with a match radius of $3''0$.

We also trawl “The Astronomer’s Telegram”¹⁹ text reports (ATels) for information connected to ATLAS transients. There are two distinct searches performed against the human-written, plain-text content found in the title and body of each ATel; a name-search and a coordinate-search. As ATLAS and all other major transient search projects have consistent transient naming schemes, we have found it possible to compile a set of regular-expressions that match most of the transient names reported in the ATels. All matched names are recorded in a separate table to the original database table hosting the ATel titles and content, with ATel number as the relational key between tables. Furthermore a long and complex regular expression is employed to extract out all sets of equatorial coordinates found within ATels. This expression has to accommodate for the plethora of formats of sexagesimal coordinates and also for the surprisingly common human-errors found in those reports. The coordinates are homogenised into decimal degree format and recorded in another database table which is then indexed using the familiar HTM indexing scheme. Whenever a new ATel is

¹⁹ <http://www.astronomerstelegram.org>

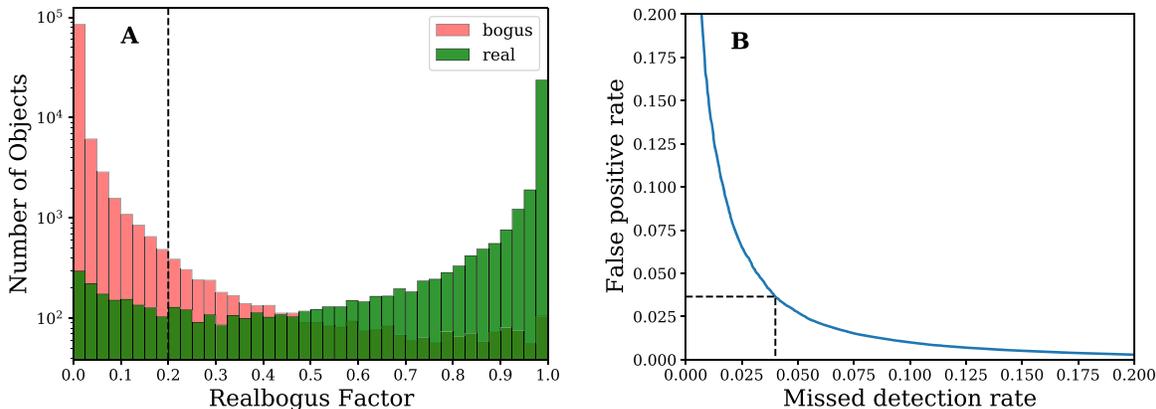


Figure 3. (A) Histogram showing the classifier performance against the training set (for the Haleakala (HKO) classifier). The dotted line shows the RB factor threshold which corresponds to a 4% missed detection rate of real objects. The real/bogus value threshold is $RB = 0.2$, below which objects are regarded as bogus. (Note the logarithmic vertical scale.) (B) Detection vs Error tradeoff. The dotted line shows the false positive rate (3.65%) at the set missed detection rate of 4%, based on a test subset of the training set. This is for an $RB = 0.2$ threshold for real/bogus selection.

(A color version of this figure is available in the online journal.)

released, transient names and coordinates can then be matched against ATLAS transient names, pseudonyms from other surveys and coordinates. The discovery of a new ATLAS transient also triggers a name and coordinate-search for matches in past ATels. The resulting ATel matches are reported on the ATLAS transient pages providing direct web-links to the associated ATel reports. However use of “The Astronomer’s Telegram” for registering, announcing and classifying extragalactic transients and supernovae is dwindling, being replaced by a more robust and accessible database system in the IAU TNS. The latter allows full photometric and spectroscopic registration in an automated way, combined with a well maintained database and links to WISEREP (Yaron & Gal-Yam 2012). It also allows reports to be written and released, as AstroNotes (Gal-Yam et al. 2019), all of which are indexed on the NASA Astrophysics Data System and are citable, hence providing all the functionality of ATels but with added data curation and searching value.

Machine learning image recognition: for the objects that remain after the stars are removed, we construct a triplet of “postage stamps.” The target image, reference and difference images centered on the detections are produced of size 6.2 arcmin (200×200 pixels). A machine learning algorithm, trained to recognise “good” subtractions (i.e., PSF-like objects in the center of the difference stamp) is run on the multiple (at least three) difference images associated with each object. A median score between 0 (bogus) and 1 (real) is applied to each object. All objects with a score above a chosen threshold are passed to human scanners. The training sets were built from the 20×20 pixel core of difference image stamps centered on 200,000 known asteroids (real) from data sampled over six separate nights and a random selection of the 20×20 cores of difference image stamps of 600,000 objects that were not known variables (bogus). Two training sets were built, one for

each telescope, both of which have similar performance, though the HKO data set was built using three cyan and three orange nights, whereas the MLO classifier only uses orange nights (because it only has the one filter). Figure 3 shows the performance of the classifier, displaying a plot similar to a receiver operator characteristic (ROC) curve and a hypothesis distribution. We began using the machine learning algorithm as described in Wright et al. (2017) and Wright (2015) (random forest and sparse filtered neural network) but now use a Keras (Chollet et al. 2015) convolutional neural network with a Tensorflow (Abadi et al. 2015) backend. Convolutional neural networks (CNNs) were applied successfully for transient discovery and detection in image data by Cabrera-Vives et al. (2017), showing excellent promise for use in high data volume surveys (see also Reyes et al. 2018). Our Keras/Tensorflow CNN, which is a modification of an example model used to recognise the MNIST database of handwriting digits, consists of three 2D convolutional layers with 16, 32 and 64 filters and a kernel size of 2×2 pixels. Each convolutional layer is followed by a maxpooling layer that pools over 2×2 pixel regions. The final pooling layer feeds into a dropout layer (to prevent overfitting) with a dropout rate of 0.3, which is followed by a Dense (fully connected) layer with 500 units and another dropout layer with a dropout rate of 0.4. The output layer is a Softmax classifier that produces two outputs; the probability the detection is real and the probability the detection is bogus. All units other than those in the output layer are rectified linear units (relu). “Same” padding (Dumoulin & Visin 2016) is used throughout and the model is trained with the categorical cross entropy loss (Hastie et al. 2009) and optimized with adam (Kingma & Ba 2014).

We are of course free to set the threshold for the RB factor to any value to optimize either completeness or purity. In reality we run our daily transient searches by selecting objects with

$RB \geq 0.2$, in the knowledge that this results in 96% completeness. As can be seen in the example numbers in Table 2, this reduces the objects that humans will scan by a factor 20, from about 9000 to 300, leaving quite a manageable number per day. The classifier is periodically retrained as the training sets are refined. This is necessary because of changing conditions (e.g., PSF improvements over time). Selection of representative “bogus” data is also improving as human scanners log more examples where machine classification has failed. We foresee improvements in the training sample by using all the images of all objects we now consider as real stationary transients. This may improve the performance of the classifier on highly varying spatial backgrounds (which leave noise residuals in the difference images). With close to 8000 real objects in our database and at least 20 images per object this will provide a rich and useful training set. Employing an unsupervised method could allow clusters of real and bogus objects to emerge and further improve performance.

Forced photometry: for each potential extragalactic transient, or orphan transient, that has minimum of 3 detections at 5σ significance, and a machine learning score above our chosen threshold (i.e., has passed the above filters), we carry out forced photometry on historic images up to 30 days before initial discovery at the position of the transient (using the rolling average of its R.A. and decl. in the unforced .ddc files). This normally results in one or more pre-discovery epochs with multiple $2-4\sigma$ detections on the individual 30s frames. We calibrate and measure the forced photometry in flux (μ Jansky), since conversion between AB magnitudes is trivial, allowing low significance positive values, zero values, and negative values to provide real meaning (which is lost when converting to Pogson magnitudes). This is invaluable in constraining the explosion epoch for nearby objects and providing constraining early data for all transients. This information is presented to the user along with all the above information, as explained in the next section (see Figure 5 for an example). The forced photometry is again PSF based, with t_{phot} used in the same way as for unforced PSF flux measurements, but simply with the centroid position fixed and the PSF modeled from the input image before subtraction. Forced photometry is also done against a stack of the four intra-night difference images for an object, which can help to tie down the discovery epoch for an object.

4.3. User Interaction and Object Selection

The user is presented (via the web interface) with a triplet of postage stamps: the target image (typically 4 per night), the reference image used as the template for subtraction, and the difference image. The scanners inspect these and use a web form to decide whether or not they are indeed real or bogus transients that have a high machine learning score. Figure 4 shows examples of the triplet of images presented to the users. As well as these triplet images, the user is presented with a range of other

visualisation plots and tools which include the forced photometry, Aladin-lite (Bonnarel et al. 2000; Boch & Fernique 2014) plugin showing the Pan-STARRS1 3π archive image (if above $\delta > -30^\circ$, otherwise it defaults to the DSS), Galactic coordinates, a summary of the Sherlock cross-matching results (along with host galaxy information). Figure 5 shows examples of the panels displayed to users. At this point, human scanners decide whether to log the object in the ATLAS “Good Objects” list, reserve judgement and save it to a holding list (called “Possible Objects”) to wait for further data, or to archive it. The object remains in the database even if archived, and can be retrieved. The Transient Name Server (TNS) python API²⁰ is used to submit objects promoted to the TNS as soon as objects have been identified as “Good” by the scanners. All future detections from ATLAS that are spatially associated with an object get added to the lightcurve, and forced photometry runs on all objects in the “Good Objects” and “Possible Objects” list once a day, and the lightcurve plots are updated. A summary of nightly ingests are given in Table 2.

GW Event Tagging: software has been written (in python) to download the latest LIGO-Virgo GW event Healpix maps and calculate which ATLAS objects lie within the 90% contour region of a GW event. Objects are tagged with the event ID, event MJD and contour value (in 10% divisions) if they are within -10 to $+21$ days of the GW event. E.g. “Possible GW Events Association: GW190425 (MJD 58598.3458914-60% contour).” The ATLAS web interface has an additional filter facility that can be used to search for named GW events and display lists of objects that lie within the 90% region. This facility can be used to prioritise the scanning of relevant objects after a GW event. This was used to discover the afterglow of GRB 170105A which was in the skymap of the binary black hole merger GW170104 (Stalder et al. 2017).

4.4. Hardware

The processing of data is done by two rack servers—a 28 core server with 128 GB RAM for the ATLAS data processing and another identical but separate machine to host the database. The servers are linux based (CentOS7) and are equipped with 26 TB (RAID6) of spinning disk (for postage stamp storage) and 16 TB of NVME SSD in a RAID0 configuration (for database storage, to reduce I/O latency). The database SSDs are backed up to spinning disk twice a week, so loss of an SSD, while destructive to the array, would not result in permanent loss of data.

ATLAS postage stamps are stored on a 26 TB RAID6 spinning disk array, local to one of the above servers, while the full sized reduced ATLAS images (input, reference and difference) are stored on a bank of 4 RAID6 70 TB dedicated network attached storage arrays.

The webserver requires relatively little CPU or memory, so is a separate virtual machine (running on shared university

²⁰ <https://wis-tns.weizmann.ac.il/content/tns-getting-started>

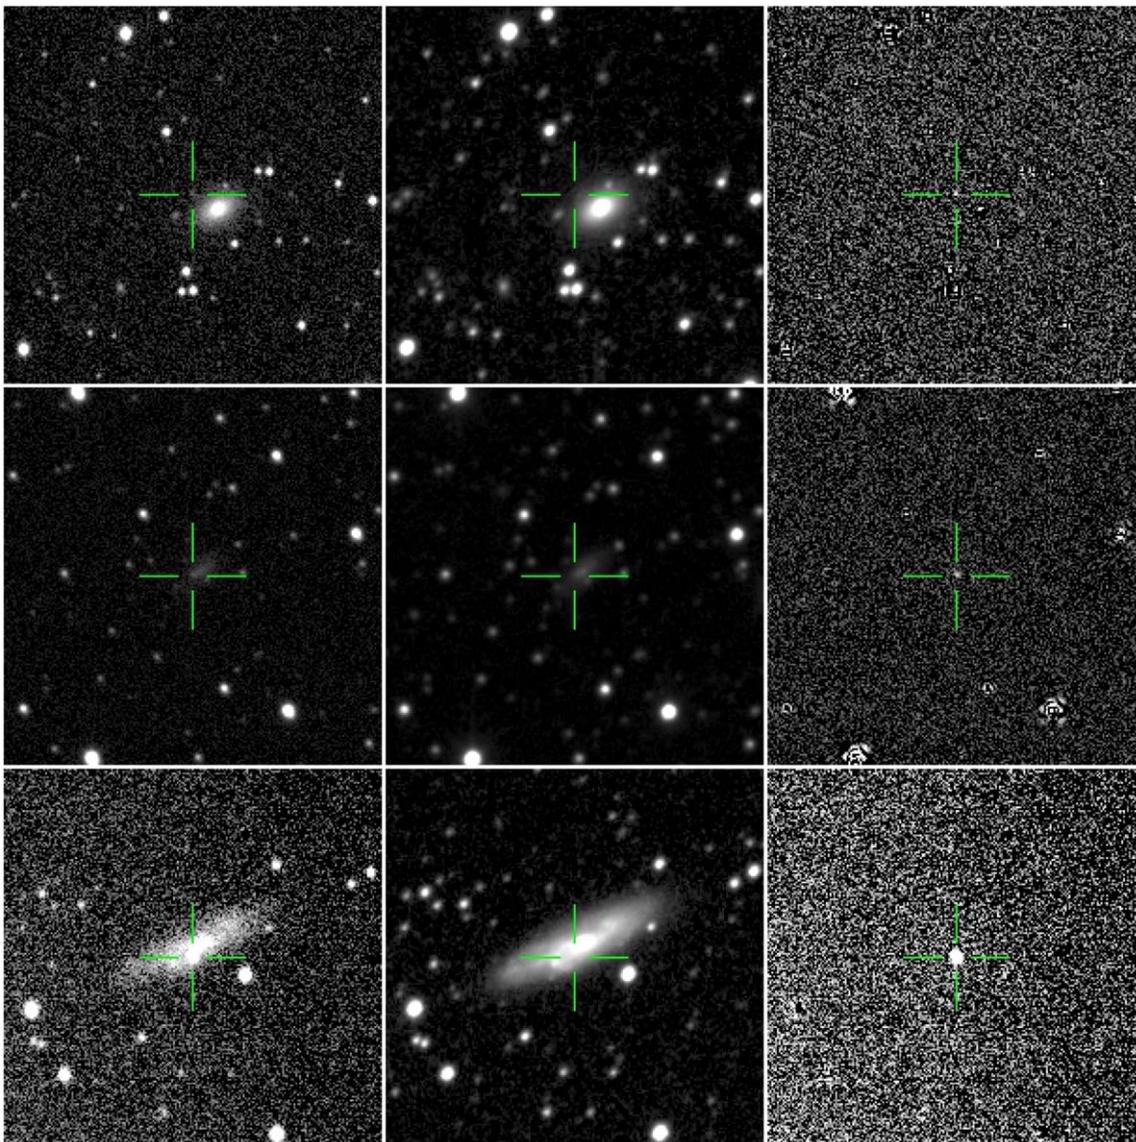


Figure 4. Three examples of triplet images, target (left), reference (middle) and difference (right) of objects that have passed the automated filters, are associated with galaxies with known redshifts and are presented to human scanners for final selection. They are confirmed supernovae SN2019shb, SN2019tua and SN2019vsa (respectively from top to bottom), the top one has a tensorflow real-bogus factor $RB = 0.71$ and the other two have $RB = 1$.

(A color version of this figure is available in the online journal.)

hardware) with an allocation of 4 cores and 8 GB RAM. Local webservice data is stored on a small SSD with postage stamp images accessed via NFS from the above 26 TB array.

Raw ATLAS data (currently about 280TB) is backed up directly from Hawaii to the QUB petabyte scale HPC cluster. Periodic ATLAS database dumps and postage stamps are also stored there.

4.5. Software

Most of the software is written in Python, with C++ (ingester) and javascript (dynamic lightcurve rendering) components. The ingester, post ingest cutting, context classification and external

survey crossmatching are marshalled by bash scripts called via CRON.

Ingester: the ingester is written in C++ (with the template functions generated by python code) and dependent on the C++ HTM library. The code is spawned using a Python Multi-processing wrapper.

Post Ingest Cutting and Crossmatching: the post ingest processing, context classification (*Sherlock*) and external crossmatching code (including minor planet rejection) is written in python, with extensive utilisation of Astropy (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018). Most of

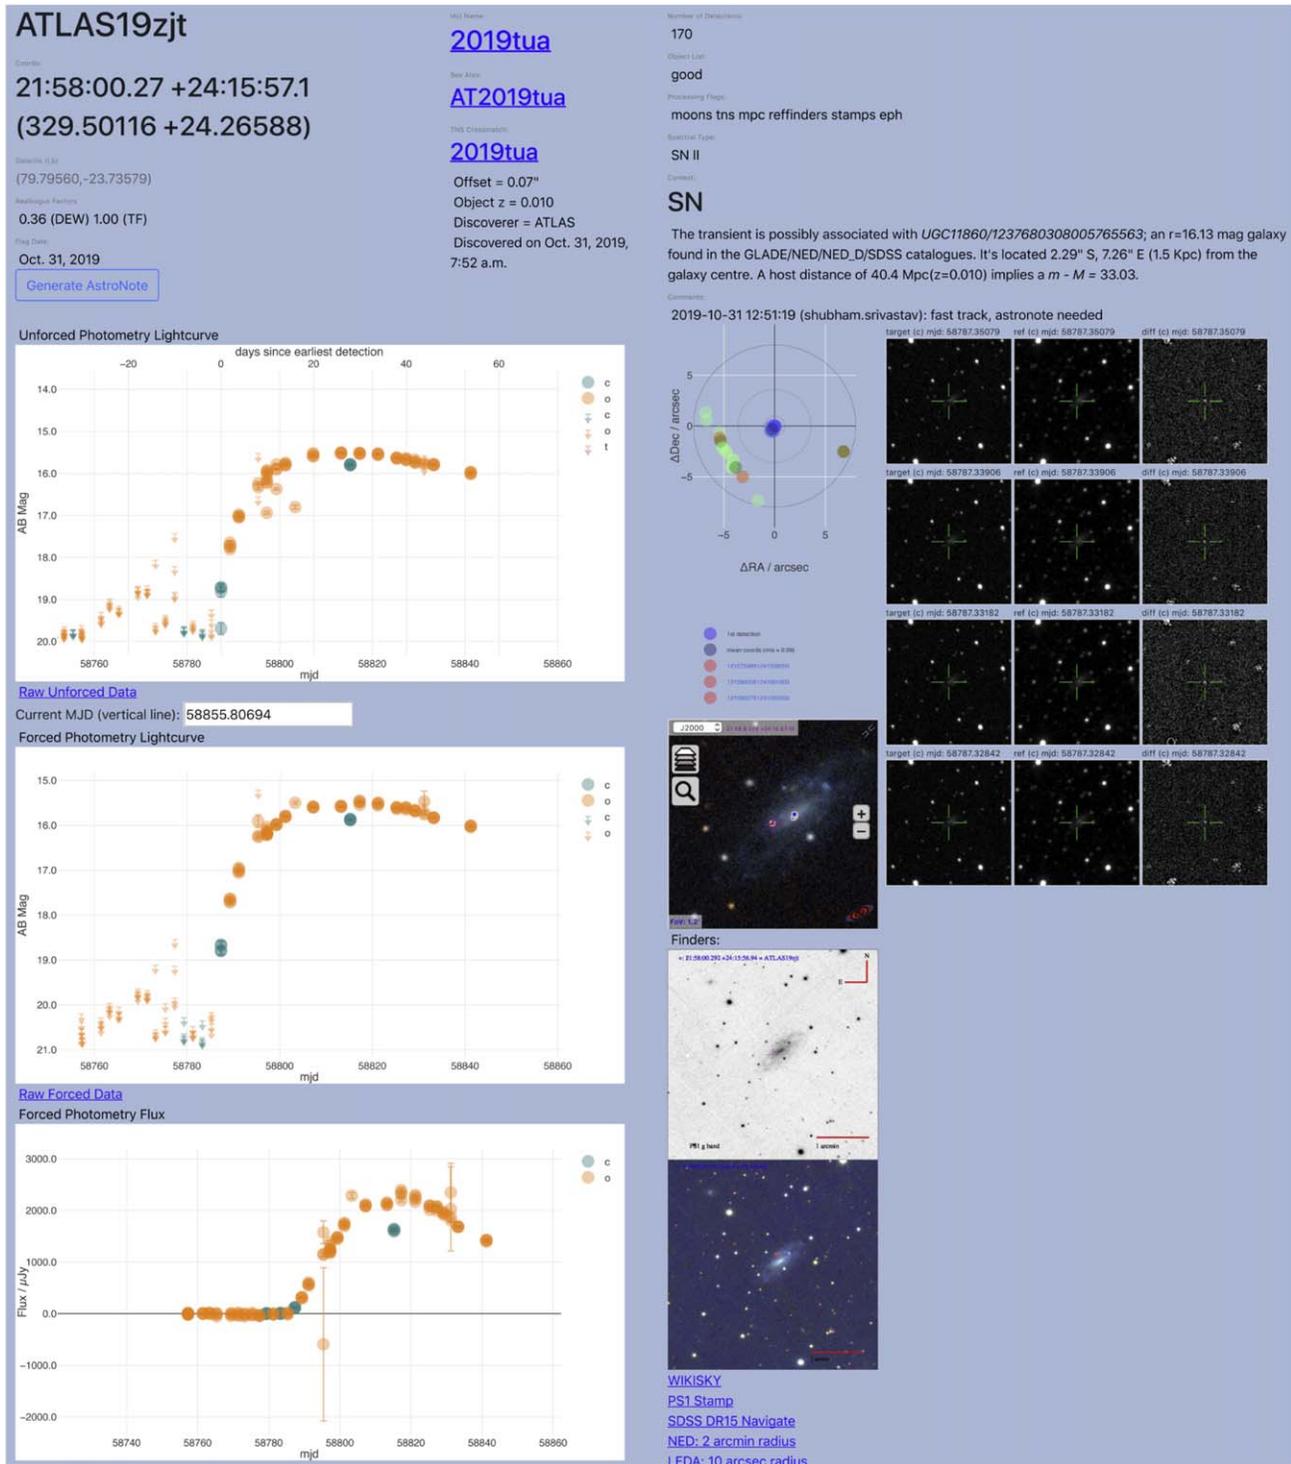


Figure 5. Candidate object page showing unforced, forced and forced photometry in flux space (μ Jansky units) lightcurves. (Non-detections in the magnitude plots are shown as arrows, which display the limiting magnitude (5σ for unforced, 3σ for forced.) The most recent four rows of images on the day it was flagged are also displayed, along with a scatter plot of detections (which also shows a passing asteroid). An Aladin-Lite widget is shown displaying Pan-STARRS imagery, and underneath this are shown two finders from Pan-STARRS data in both PS1 g-band and color (PS1 gri). Text information on the host galaxy is shown. A crossmatch with the Transient Name Server is also displayed, and the object spectral type is shown to be a type II supernova. (A color version of this figure is available in the online journal.)

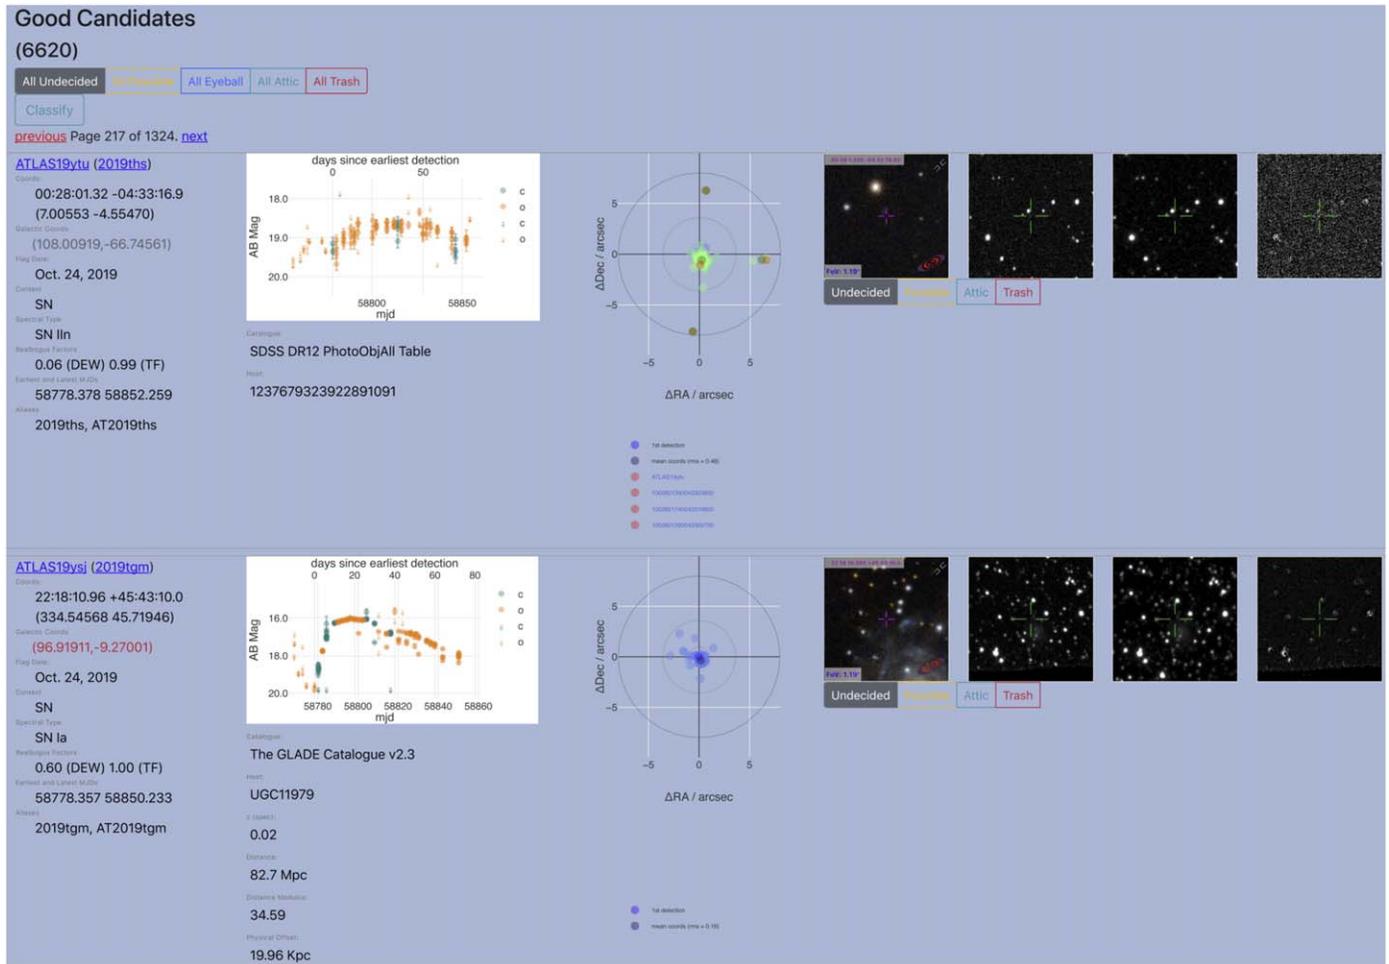


Figure 6. A quickview page showing multiple objects in the “good” list. Summary information including lightcurves, scatter plots and Aladin-Lite widgets are presented, along with a representative triplet of postage stamps. Objects can be moved one-at-a-time using buttons under the stamps, or all at once using “select all” buttons at the head and foot of the page. Lists of objects can be filtered by (e.g.,) context type, or GW event ID, and ordered by (e.g.,) RB factor. Note that the galactic coordinates are highlighted in red if the object is within galactic latitude $|b| < 15$. (A color version of this figure is available in the online journal.)

the code is called via the Python Multiprocessing module to make maximum use of our CPUs.

Database: the database technology is MySQL. The “back-end” storage engine chosen for most tables is the non-transactional MyISAM engine. Although this is becoming obsolete in future, with users strongly being advised to move to InnoDB (transactional storage), the MyISAM engine is still the preferred solution for ATLAS for the time being. Its main advantage is that the database files are a third of the size of InnoDB database files. Given that the database is currently 6 TB in size, we would require 18 TB of SSD storage to run the InnoDB version just to stand still. Another advantage is that the database files can be copied directly from one machine to another without having to apply transactional logs (assuming the source and destination machines are running the same version of MySQL). Additionally, the way the database is

currently designed (with primary keys of objects represented by their location in the sky) does not lend itself well to the InnoDB model of database row insertion (though this would be relatively easy, albeit time consuming, to rectify by the creation of a new running counter primary key). One final advantage is MyISAM tables can be ordered by a specified index. One property of HTMs is that spatially nearby triangles are physically close to each other in the database (because they are numerically close). Hence if a table is ordered by the HTM index, this reduces lookup time. The main disadvantage of MyISAM is that the whole table is locked when a single row is updated (forcing many parallel updates to be serialised internally). However, parallel inserts are allowed, and this property is used by the ingester.

The ATLAS database, which has been running since 2015 December and contains 20 billion detections (5.7 billion

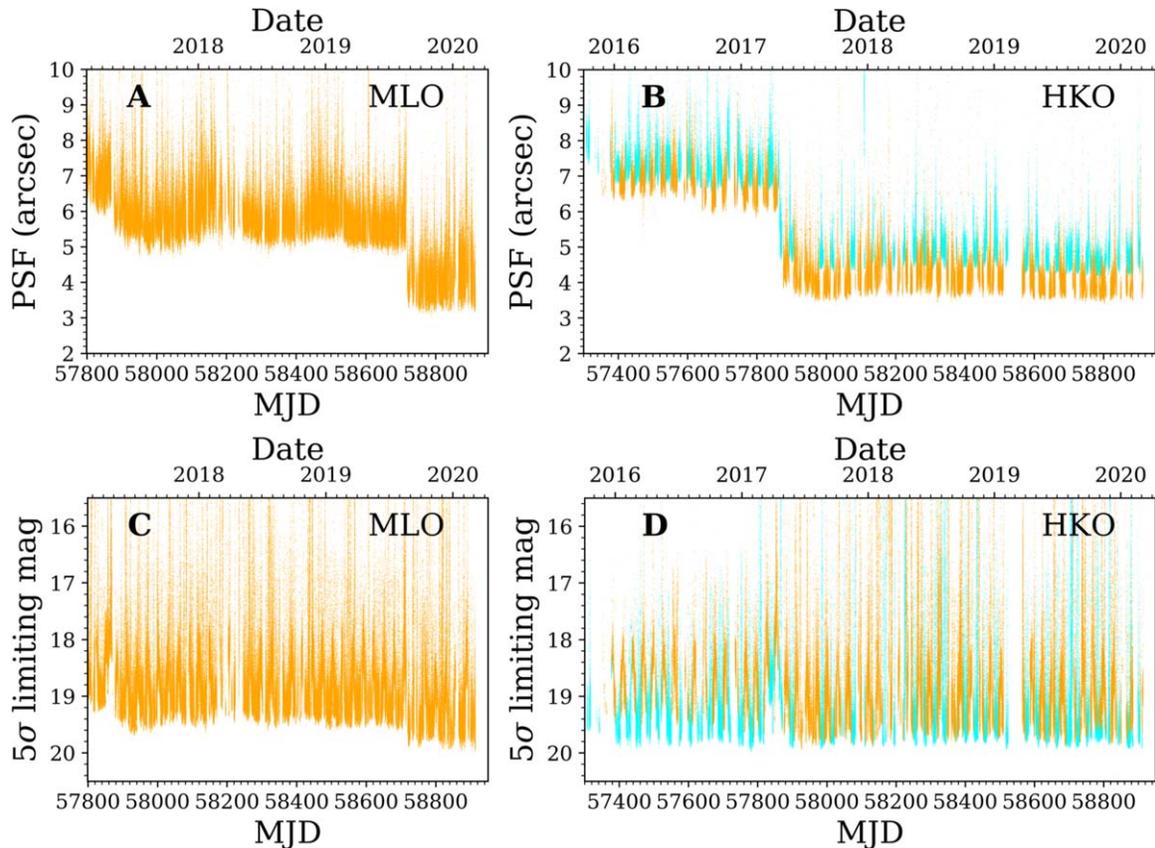


Figure 7. Image quality monitoring for ATLAS HKO and MLO for their full history. Cyan points represent measurements in the *c* filter and orange in the *o* filter. Recall that Mauna Loa only has had an *o* filter in place to the present date. Note that the plots show limits *per exposure*. (A) and (B): PSF measurements—the recorded FWHM (in arcseconds) measured on the detector. (C) and (D): The 5σ limiting magnitude, which is a function of zero-point variation (extinction), PSF variation and sky brightness.

(A color version of this figure is available in the online journal.)

objects) currently occupies about 6.2 TB (as of 2020 March 12). The size will likely double with the introduction of the southern telescopes in Chile and South Africa. The ATLAS postage stamps currently occupy 36 TB (mostly junk or asteroid data, which is kept for future machine learning), but extraction of the stamps also requires download of the original input, reference and difference images—which currently accounts for 0.2 TB per day per telescope. Hence a single telescope database is growing by 0.75 TB per year, stamps are growing by 4.5 TB per year and image data (if download continues at the current rate) will require 60 TB per year per telescope. Several times this space for backups (twice weekly and monthly) on spinning disk is also reserved.

An additional unsolved problem is the database storage of the photometric time series images on all the individual reduced images (not the difference images) to provide variable star and AGN lightcurves. Heinze et al. (2018) presented the first catalog of variable stars from ATLAS, which is 4.7 million stars, selected from 142 million stars that each had 100 photometric points or more. A more ambitious project would

be to create an ATLAS billion star database which would store 1000 points on the lightcurve per year for a billion stars. The size of this database would potentially be at least 100 TB per year, assuming just 100 bytes per detection and the database would need to ingest 64,000 detections per second, assuming 50% utilisation. The current ATLAS Transient Science Server infrastructure is not yet designed to cope with a database of this size and a different architecture and design would be required to serve all variables.

Web Interface: a web interface (written in python using the Django Web Framework,²¹ Bootstrap²² and Plotly²³) has been built atop the database to allow the human scanners to visually inspect objects that have a reliable machine learning score. The web interface is designed to be “responsive” hence will adjust how pages are rendered based on whether they are being viewed via a mobile or desktop browser. Individual objects are

²¹ <https://www.djangoproject.com>

²² <https://getbootstrap.com>

²³ <https://plot.ly/javascript/>

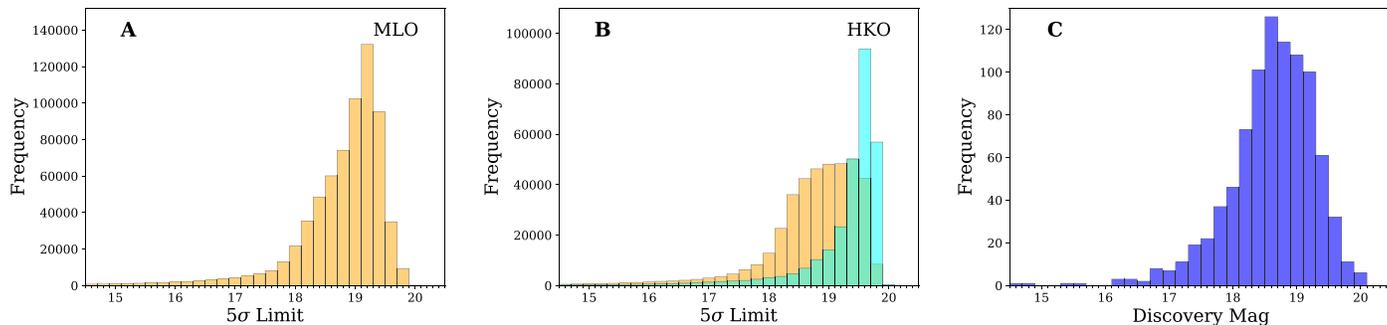


Figure 8. (A) Histogram of 5σ limiting magnitudes for MLO in the o band, with the same data as in Figure 7. (B) Same as in (A), but for HKO with the c and o bands plotted together. (C) The discovery magnitude distribution for all ATLAS discovered transients reported to the TNS that have been classified as supernovae. (A color version of this figure is available in the online journal.)

displayed to the user utilizing as much of the space as possible (in desktop orientation) to display the maximum amount of information without having to excessively scroll down (Figure 5). Users can be presented with lists of objects in either simple tabular format or with full rendered lightcurves, scatter plots, Aladin-Lite context, summary information and the most recent image triplet (Figure 6). The latter allows a user to see groups of 50 (by default) objects at a time (adjustable), allowing them to archive multiple objects or retain for promotion with a single click. The lists may be sorted by (e.g.) RB factor, latest magnitude and/or filtered to display only subsets of objects (e.g., objects tagged as SN or NT, or objects associated with a specific LIGO/Virgo event).

5. On-sky Performance and Monitoring

As with any survey facility, the sensitivity of an ATLAS image is dependent on atmospheric transparency, sky background (moon phase) and the image quality or PSF as measured on the detector. We monitor all of these parameters with measurements made on every on sky science, reduced science frame. The values are written into the reduced image FITS header to allow on-sky performance monitoring and visualisation of long term trends. (Most of this information is also propagated to the .ddc files and hence stored in the database.) Plots of the temporal variation are given in Figure 7 for the full ATLAS history. Points to note on these plots are:

1. The significant improvement in PSF around MJD 57864 (2017 April 21) is due to the replacement of the Schmidt correctors on both telescopes. The improvement is initially less obvious on ATLAS-MLO for the reasons stated below.
2. The throughput at ATLAS-HKO was observed to roll off by 0.2 mag between 58420 (2018 October 29) and 58605 (2019 May 2). This was due to a faulty dehumidifier which emitted an oily residue that gathered on the CCD window. The window was cleaned on 58606 and immediately the

throughput returned to normal. The faulty dehumidifier has been replaced.

3. The large gap at MJD = 58525 (2019 February 11) to 58565 (2019 March 23) was due to an ice storm on Haleakala, which enforced a long period of dome closure. The storm lasted just a few days but disabled electrical service to the Haleakala summit for 42 days. The gap is also visible in Figure 1.
4. The ATLAS-MLO sensitivity saw a slow, gradual improvement (due to improving the PSF through focus stability) until MJD = 58715 (2019 August 20). Until that point, the PSF measured on the ATLAS-MLO detector was significantly worse than on ATLAS-HKO (5.5 arcsec median versus 4.0 arcsec median on HKO). A new STA 1600 detector was installed on MJD = 58715 which immediately improved the stellar image widths to a median of 3.8 arcsec, even better than ATLAS-HKO. This was due to charge diffusion on the old CCD which degraded the FWHM (caused by a manufacturing defect). This effect corresponded to a blurring of approximately 2 pixels FWHM, in quadrature.
5. The 5σ limiting magnitude displays the expected modulation due to lunar phase and the variation in sky background.

The histogram of the 5σ limiting magnitudes for each telescope are given in Figure 8 for the calendar year 2019. This includes all sky conditions and all atmospheric transparencies. As can be seen in both Figures 8 and 7 there is a long tail toward poor quality images with limiting magnitudes $m_{5\sigma} < 16$, which we deem not to be of science quality. These typically make up 2% of the total science images. The median limiting magnitudes for each of the systems (from 2019) is $o < 19.0$ mag (ATLAS-HKO) $c < 19.6$ mag (ATLAS-HKO) and $o < 19.0$ mag (ATLAS-MLO). Although ATLAS-MLO has historically had poorer PSF measurements, and a similar zero-point (i.e., similar throughput), the fact we employ it in both dark and bright time compensates since ATLAS-HKO is preferentially used when the moon is above the horizon. Since

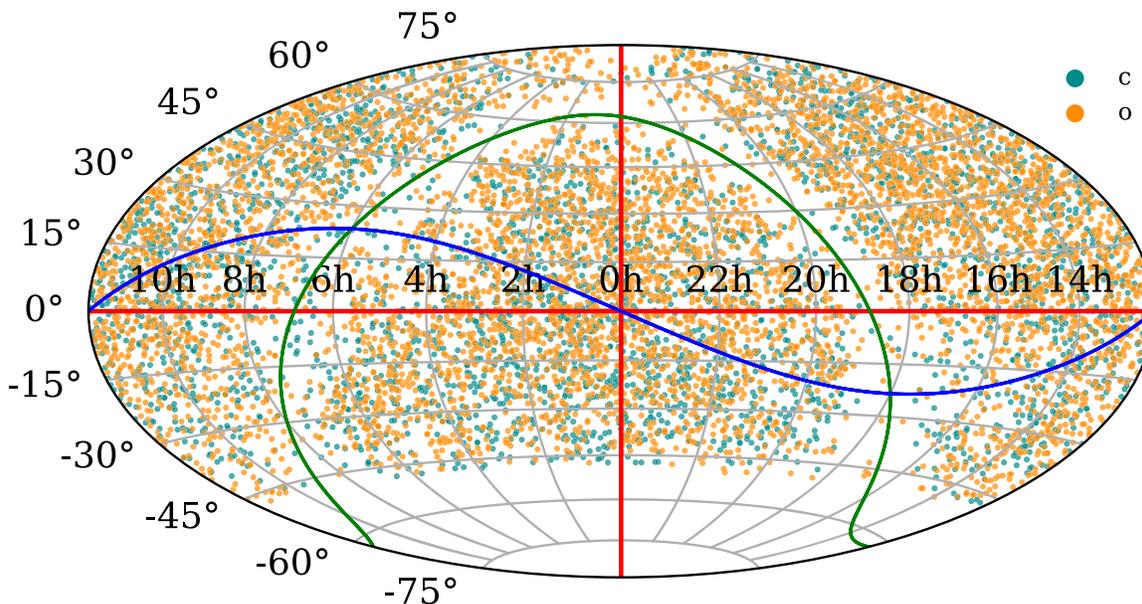


Figure 9. Hammer projection of the ATLAS transients discovered from the start of the survey until 2020 March 15. Cyan and orange points represent the discovery filter of the object. Green and blue lines are the galactic plane and ecliptic respectively. (A color version of this figure is available in the online journal.)

the new CCD was installed on ATLAS-MLO on MJD = 58715, the median has improved to $\sigma < 19.3$ mag.

6. Summary of Results and Discoveries

As described in Section 4.3, real astrophysical transient sources are selected by the ATLAS Transient Server system, with the final step being human scanning of all the candidate sources which result from the computational processing steps. Table 2 indicates that we typically typically have 300–400 objects per day to present for human vetting which results in a typical mean discovery rate of 10 to 15 real extragalactic transients per day. All these real sources are automatically registered on the International Astronomical Union’s Transient Name Server (IAU TNS),²⁴ and a catalog of our discoveries is easily recovered from the TNS web interface or through simple scripts (e.g., Kulkarni 2020).

To date (2020 March 15), we have registered 5282 transients as primary discoveries, and have detected 10,332 in total in the “Good” list in our Transient Science Server database (the difference being those discovered by other surveys, mostly iPTF, ZTF, Gaia, ASASSN and Pan-STARRS). (See Figure 9.) We aim to register extragalactic transients that are not known AGN on the TNS. Hence, as far as possible we remove objects that are coincident with known AGN and CVs (see Section 4) and fast declining transients (more than 0.5 mag within the 1 hr quad) that are generally associated with faint red objects and are hence M-dwarf flares (Martínez et al. 2019). While we save

these in our database as real objects, our intention is not to let these slip into the TNS reporting stream.

We also do not report objects that are almost certain CVs, which means a very fast rise of several magnitudes within 1–3 days (typically a gradient of $\text{grad} < -0.5 \text{ mag day}^{-1}$) from last non-detection and which have no obvious host object visible to g , r or $i \simeq 23$ mag in the Pan-STARRS 3π archive imaging. While this behavior is almost exclusively a sign of CVs and they dominate the sky density of very fast rising transients, it is important not to be over zealous with this filter, since a number of rare, fast, blue extra-galactic transients inhabit this parameter space. In addition, if extragalactic transients are more than a projected distance of $R_p \simeq 50$ kpc from their host, then association with the correct host and redshift becomes confusion limited. The discovery of AT2018cow showed how a fast blue transient with a similar rise as a CV could turn out to be a luminous and unusual transient at 60 Mpc, which radiated from the hard X-ray to the radio (Prentice et al. 2018; Smartt et al. 2018; Margutti et al. 2019).

The selection of transients from detections on the subtractions (in the .ddc files) to final registration on the TNS depends on several steps, each of which are not guaranteed 100% completeness. This is also true of other surveys, so a straight comparison of discoveries does not yield a meaningful result. As a test of completeness of our processing, we extracted transients discovered by two other surveys during the 6 month period 58620 (2019 May 17) to 58787 (2019 November 16), after the HKO ice-storm and the ATLAS-HKO CCD window fogging issue.

²⁴ <https://wis-tns.weizmann.ac.il/>

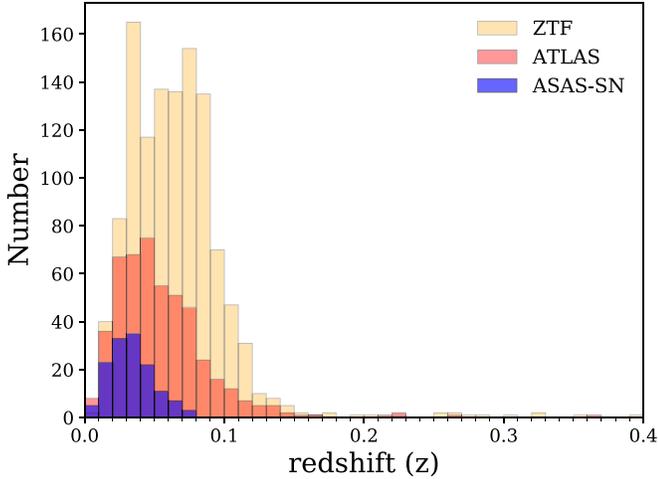


Figure 10. Redshift distribution of all ATLAS discovered and detected extragalactic transients reported to the TNS in 2019 that have a spectroscopic redshift. For comparison, the redshift distribution of spectroscopically confirmed supernovae reported to TNS in 2019 by ZTF and ASAS-SN are also shown. It should be noted that ZTF have dedicated spectroscopic followup facilities.

(A color version of this figure is available in the online journal.)

The ASAS-SN project (Holoiën et al. 2019) registered 88 objects as discoveries on the TNS above $\delta = -50^\circ$, and we searched for these in our data. 13 of these were discovered at RAs with large hour angles, and in an area of sky that ATLAS had either stopped observing for the season or had yet to start. We could not confirm the existence of 12 of the 88 (and they are as yet unconfirmed on the TNS), suggesting they are not real. This left 53 transients that should be in the ATLAS database and registered as objects. All were detected on the ATLAS difference images (multiple times at $S/N > 5$), and all 53 objects were in the ATLAS Transient database. Three objects did not make the initial filter of ≥ 3 , good quality detections, at greater than 5σ , on any one night and two had low machine learning RB scores. A 4% missed detection rate from the machine learning is consistent with the levels we set for selection from the ROC curve discussed in Section 4.3. One further object was rejected as being associated with a compact galaxy labeled as stellar in the PS1 catalog (therefore Sherlock classified it as a VS). Therefore our computer filtering from end to end appears susceptible to a 10% incompleteness. A further 8 objects were not promoted as real sources by human scanners, as they judged them not to be reliable astrophysical transients on the first night that they were detected. This indicates the level of uncertainty introduced when scanners have a small number of detections on which to judge real or bogus (typically introducing a 15% hit on efficiency). This human scanning incompleteness usually arises from low significance detections on the first night of detection, and the objects are archived, while subsequent nights will produce much more reliable and secure detections. (See improvements

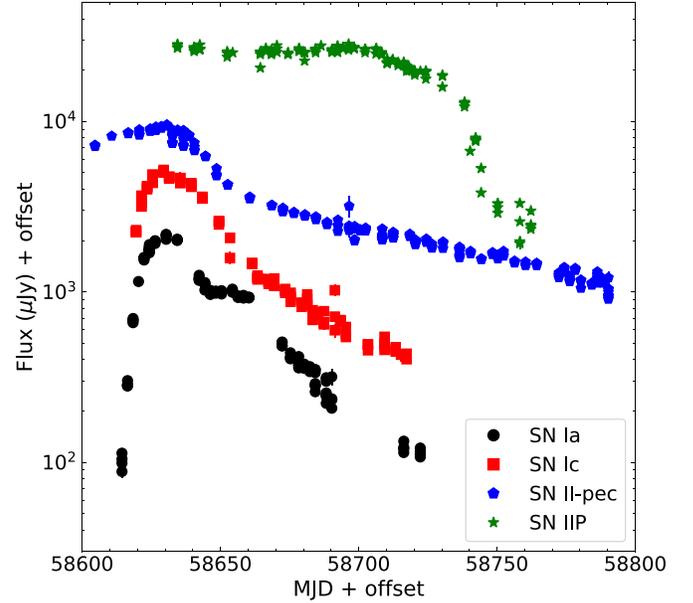


Figure 11. Well-sampled *o*-band ATLAS light curves of 4 nearby SNe of types Ia, Ic, IIP and II-peculiar. The “shoulder” or secondary peak in the light curve of the SN Ia ~ 30 days after the primary maximum is seen clearly, whereas the stripped envelope SNe (SESNe: Ib, Ic and IIb) do not exhibit this feature. The light curves were shifted along the *x* and *y* axes for clarity.

(A color version of this figure is available in the online journal.)

in Sections 7 item 3.) A search for ZTF objects with a peak $r < 19$ mag and at least 4 detections in *Lasair*, and a cross-correlation with ATLAS objects showed similar results. The end to end comparison showed we had approximately 80% efficiency with the losses split roughly equally between automated computer filtering (basic cuts and machine learning) and human decisions.

In Figure 10 we illustrate the redshift distribution of the spectroscopically classified ATLAS transients in 2019. We include only the objects that were formally discovered by ATLAS (which we define as first report to the TNS) and exclude those detected by ATLAS but were discovered by other surveys. (Spectroscopically confirmed objects reported to TNS by ZTF and ASAS-SN in 2019 are also shown for comparison.) At a redshift of $z = 0.15$, a discovery magnitude of $m \sim 19.5$ corresponds to $M \sim -19.6$ mag, or a typical Ia at peak. Those SNe beyond $z = 0.15$ are significantly brighter than the bulk of the normal population. Of the 15 SNe found at $z > 0.15$, 7 are super-luminous SNe (SLSNe), whereas the remainder include bright Ia (Ia 91T-like or Ia-CSM), II and Ic broad-line events. Figure 11 shows example light curves of different classes, indicating that in the *o*-filter, the near-infrared secondary peak in normal type Ia SNe is clearly detected, allowing SNe Ia to be distinguished from SNe Ib/c. This will assist SN rate calculations for the supernovae for which we do not have a spectroscopic classification.

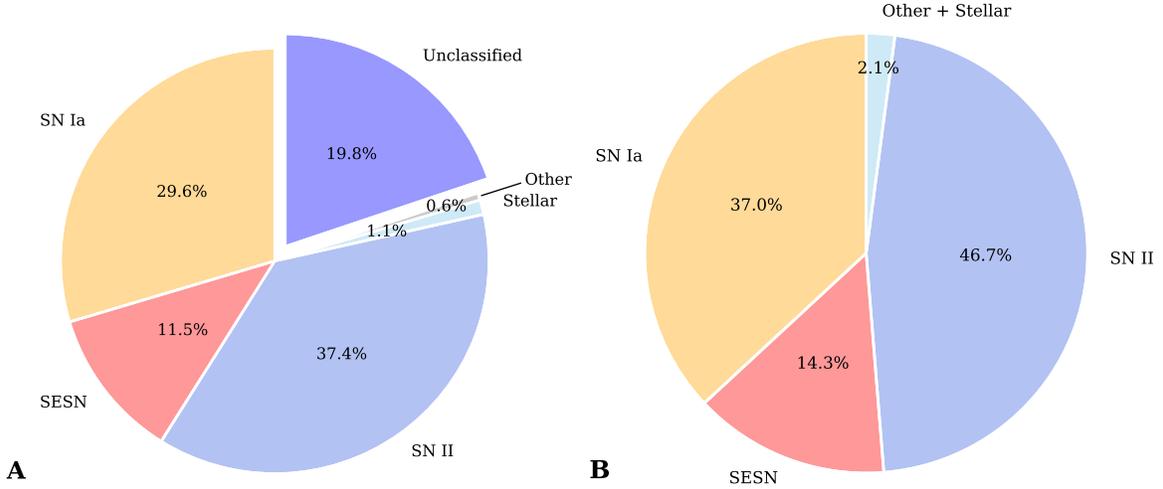


Figure 12. Pie charts of the spectroscopic classifications of ATLAS detected transients associated with galaxies with distances within 100 Mpc. (A): All ATLAS transients. (B): Subset of spectroscopically classified transients. SESN represents the family of stripped envelope SNe that includes types Ib/c, IIb and Ic broad-line. Stellar transients include Galactic CVs and variable stars as well as extragalactic novae, whereas the “Other” category includes a TDE, ILRTs and LBVs. A total of 540 transients were included, out of which 107 transients are without a spectroscopic classification.

(A color version of this figure is available in the online journal.)

6.1. Spectroscopic Types within the Local 100 Mpc Volume

An interesting scientific survey that ATLAS will contribute to is measuring complete, volumetric rates of transients within the local universe. At a distance of 100 Mpc ($\mu = 35$), we are sensitive to transients of absolute magnitude of $M \lesssim -16$ (in regions of low Milky Way extinction) and is an upper limit to the distance within which we will carry out this survey. It is also approximately the upper distance limit to which we would be sensitive to faint, fast fading kilonova type objects like AT2017gfo (which peaked at $M_{r,i} = -16$ mag, within the first 24 hr; see the early combined r and i photometric evolution in Andreoni et al. 2017; Arcavi et al. 2017; Coulter et al. 2017; Soares-Santos et al. 2017; Smartt et al. 2017; Valenti 2017).

A total number of 540 transients, associated with host galaxies within 100 Mpc, were detected by ATLAS and reported to the TNS between 2017 September 22 (MJD 58018.5) and 2019 December 03 (MJD 58820.2). These galaxies had either distances or spectroscopic redshifts (in our three primary catalogs: NED, GLADE and SDSS) and the association radius was set at a sky projected distance of 50 kpc in *Sherlock*. These include transients discovered by other surveys, but all were independently detected by ATLAS. Out of these, 107 transients (19.8%) do not have a reported classification on TNS. Of the 433 transients with a reported spectroscopic classification, 6 (1.1%) were stellar sources, i.e., foreground Galactic CVs and variable stars or extragalactic novae. Tidal disruption events (TDEs), intermediate luminosity red transients (ILRTs) and luminous blue variable outbursts (LBVs) together constitute a small fraction (0.6%) of the classified transients. The breakdown of the classifications is

shown in Figure 12 and we emphasize that this should not be taken as a true measure of relative rates since we have carried out no correction for completeness as a function of absolute magnitude of each type. It is purely illustrative of the numbers of transients detected which are associated with galaxies within this distance limit and shows that about 1 object per day is found which is plausibly within 100 Mpc. As an illustrative comparison we can compare our results to the Lick Observatory Supernova Search sample (LOSS; Li et al. 2011b, 2011a), which was a fully corrected, volume limited survey of cataloged galaxies within 60–80 Mpc, from a sample of 180 SNe in total. The relative rate for core collapse SNe (CCSNe) in our sample (433 objects after removing unclassified transients) is 61.0%, whereas SNe Ia account for 37.0%. In comparison, the relative rate of CCSNe and SNe Ia was 60.6% and 37.7% respectively, in the Lick Observatory Supernova Search sample (LOSS; Li et al. 2011a).

Within this 100 Mpc volume, we can estimate the simple “observed” volumetric rates for supernovae. Assuming that ATLAS covers the whole sky down to $\delta \sim -50^\circ$ (88% of 4π) and that ± 60 degrees in RA from the Sun is unobservable and lost (33% of 4π), then ATLAS covers approximately 60% of the full 4π sky with a 2 day cadence. The number of type Ia and CCSNe that we have detected within 100 Mpc implies that the lower limit to the rate of supernova is $R_{Ia} = 5.3 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for SNe Ia and $R_{CC} = 9.0 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ for CCSNe. This is based on all ATLAS detections, not just on discoveries and these volumetric rates represent a hard lower limit, since correction factors involving sensitivity and sky coverage have not been taken into account. When we calculate such corrections they will only have the effect of significantly increasing the rates since we are not 100% efficient in detecting

SNe within this volume. The above is based on assuming the 540 transients contain 2% of non-supernovae (which we know is the fraction of the 433 classified transients), and that the rest are supernovae with a relative rate the same as in the 433 classified events. One may worry about a bias introduced by the selection of transients for spectroscopic follow-up. False positives such as mistaken CVs, outbursting stars or non-astrophysical artefacts would lie in this un-classified list simply because those with spectroscopic resources considered them not likely to be real supernovae. We have visually inspected all lightcurves of these unclassified 107 objects. There are 6 which are likely either CVs, stars or poor subtractions leading to false positives. All the other 101 have clear supernova like lightcurves and are associated with 100 Mpc galaxies with no clear candidate for a background galaxy observed to Pan-STARRS1 3π depths. Many of the objects are toward the fainter end of the distribution, $m \gtrsim 18.5$ mag and some have incomplete lightcurves due to seasonal and or weather constraints. However they are all very likely SNe of some type and a future paper will address classification of these from their lightcurves, with a particular focus on the “shoulder” feature in the o -band lightcurve which distinguishes type Ia SNe from the others. The 6 possible impostor objects do not influence the above calculation.

It is interesting to note that these observed rates, with no corrections are already higher than those previously published at low redshift. For type Ia SNe, Li et al. (2011a) report a rate of $3 \pm 0.6 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ within 80 Mpc for targeted galaxy searches and Frohmaier et al. (2019) report $\approx 2\text{--}3 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ within $z \lesssim 0.09$ (400 Mpc). For CCSNe, LOSS find a rate of $7.1 \pm 1.6 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$ Li et al. (2011a), which is higher (but not seriously discrepant, considering the uncertainties) to that from Cappellaro et al. (1999) of $4 \pm 2 \times 10^4 \text{ Gpc}^{-3} \text{ yr}^{-1}$. Our estimate of the type Ia rate is nearly a factor 2 higher than currently accepted, which is quite a difference and the discrepancy would only increase when a proper efficiency and survey simulation calculation is carried out. The CCSN rate is less discrepant, at 30% higher than the LOSS Li et al. (2011a) value. However, the completeness correction for CCNSE is likely to be much larger than for type Ia, due to their fainter luminosities (Li et al. 2011b, 2011a), and we are certainly far from complete for a $M_o \simeq -16$ mag core-collapse SN at 100 Mpc (since, $o = \mu + M_o + A_o \gtrsim 19$ mag). Therefore it is possible that some of the SNe which are classified as type Ia on the TNS are actually type Ic (or Ib), since this mis-identification is often made (Clocchiatti et al. 2000). This would decrease the Ia rate and increase the CCSN rate, although the total SN rate would be preserved, and that is already 50% higher than the published values. The ATLAS SN counts give a total SN rate of $1.4 \times 10^5 \text{ yr}^{-1} \text{ Gpc}^{-1}$, compared to the LOSS total of $1.0 \times 10^5 \text{ yr}^{-1} \text{ Gpc}^{-1}$. Another possibility is that both the type Ia and CCSN rate really are significantly higher than measured in LOSS and indeed all other low redshift

surveys. This could be due to a higher specific SN rate in lower mass galaxies, meaning targeted galaxy surveys will be biased against the high rates of production in low mass galaxies. This type of effect has been suggested by ASAS-SN for type Ia SNe (see the ASASSN results in an all-sky survey Brown et al. 2019). If the core-collapse SN rate is significantly higher than the LOSS rates, and if such a systematic underestimate were also present at higher redshifts, then it may close the discrepancy between the cosmic core-collapse rate and the star formation rate from other indicators as noted by Horiuchi et al. (2011). Although Botticella et al. (2012) find that very local SN rates and the traditional star formation rate indicators ($H\alpha$ and FUV fluxes) are discrepant in the opposite sense. To address these important problems, we will undertake further careful calculation of the local SN rates in future papers. At face value, these simple estimates indicate that the accepted local universe supernova rates could be too low by a factor 1.5 to 2.

7. Conclusions and Future Improvements

The ATLAS system now routinely surveys the whole sky visible from Hawaii (above $\delta > -50^\circ$) every two nights with the twin telescope system. This results in 10–15 extragalactic transients discovered or detected (where the latter means discovered first by other surveys) every night with a 5σ detection limit of $o < 19$ mag. All discoveries are made public promptly, however the spectroscopic classification completeness is only about 25%. The processing at QUB, during day time while the ATLAS telescopes are observing produces rapid discoveries now routinely during the Hawaiian night. Computer algorithms for filtering detections focus on real-bogus classifications, removal of moving objects and stars, and a combination of boosted decision trees and machine learning provides an efficient method to present users with a short list of real extragalactic transients. However human vetting is still essential in this last step to ensure an optimal balance between false positive rate and missed detection rate. The ATLAS discoveries not only enable rapid follow-up of interesting sources (e.g., AT2018cow Prentice et al. 2018) but will provide complete statistical samples within the local volume of 100 Mpc.

Looking forward, we foresee the following improvements to the system

1. Addition of two more units in the southern hemisphere at the South African Astronomical Observatory, and El Sauce in Chile. These will allow the full sky to be covered every night with our quad sequence and transients to be discovered routinely within 12–24 hr. An assessment of the data flow speed to QUB and what extra hardware is required is essential. It will be prudent to review if the QUB Transient Science Server should continue to be the main processing node or if Chilean/South African resources (such as the ALerCE

architecture) can improve the speed and reliability and if there should be multiple sources of data release for the transients.

2. Process the co-added 4×30 s exposures each night as one, deeper 120 s exposure. We now make a co-added nightly stack of the difference images, which results in a single image deeper by 0.75 magnitudes. Implementing our 5σ detection process on these images should enable us to reach $o < 20$ mag and $c < 20.3$ mag routinely. We are experimenting with the real-bogus balance and how to robustly select objects with one detection on nightly stack. During image subtraction, it would be worthwhile to pursue improvements to the convolution calculation, optimized for our sampling and point-spread-function. This should help decrease the false positive rate on the co-added difference images.
3. Development of code that will resurrect any object that has been rejected by human based on low signal-to-noise detections if new detections arise that are more significant.
4. Application of a machine learning algorithms that will classify objects primarily by their lightcurve (e.g., RAPID Muthukrishna et al. 2019). However combining this with host object information (galaxy redshift, offset from the galaxy center, morphology, and color) should lead to improvements in the prediction of supernova/transient type from an incomplete lightcurve.
5. Rapid release of transients as Kafka alerts, as is currently done with ZTF, and combining automatically with ZTF to provide extra time resolution.
6. Improvements to the input galaxy catalog, with more spectroscopic redshifts and any reliable photometric redshift measurements (e.g., Takada et al. 2014; DESI Collaboration et al. 2016; de Jong et al. 2019).
7. Further, closer ties to follow-up programmes to increase spectroscopic classification rates and early data.
8. Creation of a public user interface that provides access to all ATLAS lightcurve data and all information that is currently proprietary to the science teams. This is a significant effort to provide both the computing and software resources to support a large, public user base.
9. Creation of a large database of variability of every object in the sky brighter than $o \sim 19$ mag. And to provide public access as in the previous point.

This work has made use of data from the Asteroid Terrestrial-impact Last Alert System (ATLAS) project. ATLAS is primarily funded to search for near earth asteroids through NASA grants NN12AR55G, 80NSSC18K0284, and 80NSSC18K1575 (under the guidance of Lindley Johnson and Kelly Fast); byproducts of the NEO search include images and catalogs from the survey area. The ATLAS science products have been made possible through the contributions of the University of

Hawaii Institute for Astronomy, the Queen’s University Belfast, the Space Telescope Science Institute, and the South African Astronomical Observatory. We acknowledge support for transient science exploitation from the EU FP7/2007–2013 ERC Grant agreement No. [291222], STFC Grants ST/P000312/1, ST/N002520/1, ST/S006109/1 and support from the QUB Kelvin HPC cluster, and the QUB International Engagement Fund. TWC acknowledges EU Funding under Marie Skłodowska-Curie grant agreement No. 842471.

ORCID iDs

K. W. Smith  <https://orcid.org/0000-0001-9535-3199>

S. J. Smartt  <https://orcid.org/0000-0002-8229-1731>

References

- Abadi, M., Agarwal, A., Barham, P., et al. 2015, TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, [tensorflow.org](https://arxiv.org/abs/1603.04467)
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *PhRvL*, **119**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *ApJL*, **848**, L12
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, *ApJL*, **848**, L13
- Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, *ApJS*, **219**, 12
- Alard, C., & Lupton, R. H. 1998, *ApJ*, **503**, 325
- Andreoni, I., Ackley, K., Cooke, J., et al. 2017, *PASA*, **34**, e069
- Arcavi, I., Hosseinzadeh, G., Howell, D. A., et al. 2017, *Natur*, **551**, 64
- Astier, P., Guy, J., Regnault, N., et al. 2006, *A&A*, **447**, 31
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Baltay, C., Rabinowitz, D., Hadjijska, E., et al. 2013, *PASP*, **125**, 683
- Becker, A. 2015. HOTPANTS: High Order Transform of PSF ANd Template Subtraction, Astrophysics Source Code Library, [ascl:1504.004](https://www.ascl.net/1504.004)
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002
- Blagorodnova, N., Neill, J. D., Walters, R., et al. 2018, *PASP*, **130**, 035003
- Boch, T., & Fernique, P. 2014, in ASP Conf. Ser. 485, Aladin Lite: Embed your Sky in the Browser, ed. N. Manset & P. Forshay (San Francisco, CA: ASP), 277
- Bonnarel, F., Fernique, P., Bienaymé, O., et al. 2000, *A&AS*, **143**, 33
- Botticella, M. T., Smartt, S. J., Kennicutt, R. C., et al. 2012, *A&A*, **537**, A132
- Brink, H., Richards, J. W., Poznanski, D., et al. 2013, *MNRAS*, **435**, 1047
- Brown, J. S., Stanek, K. Z., Holoien, T. W. S., et al. 2019, *MNRAS*, **484**, 3785
- Brown, P. J., Holland, S. T., Immler, S., et al. 2009, *AJ*, **137**, 4517
- Bufano, F., Immler, S., Turatto, M., et al. 2009, *ApJ*, **700**, 1456
- Cabrera-Vives, G., Reyes, I., Förster, F., Estévez, P. A., & Maureira, J.-C. 2017, *ApJ*, **836**, 97
- Cappellaro, E., Evans, R., & Turatto, M. 1999, *A&A*, **351**, 459
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, [arXiv:1612.05560](https://arxiv.org/abs/1612.05560)
- Chen, T. W., Inserra, C., Fraser, M., et al. 2018, *ApJL*, **867**, L31
- Chollet, F., et al. 2015, Keras, <https://github.com/fchollet/keras>
- Clocchiatti, A., Phillips, M. M., Suntzeff, N. B., et al. 2000, *ApJ*, **529**, 661
- Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, *Sci*, **358**, 1556
- Dályá, G., Galgóczi, G., Dobos, L., et al. 2018, *MNRAS*, **479**, 2374
- de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, *Msngr*, **175**, 3
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036)
- Downes, R. A., Webbink, R. F., Shara, M. M., et al. 2001, *PASP*, **113**, 764
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, **696**, 870
- Dumoulin, V., & Visin, F. 2016, [arXiv:1603.07285](https://arxiv.org/abs/1603.07285)
- Fabrizi, J., Otsuka, M., Barlow, M. J., et al. 2011, *MNRAS*, **418**, 1285
- Flesch, E. W. 2019, [arXiv:1912.05614](https://arxiv.org/abs/1912.05614)
- Flewelling, H. A., Magnier, E. A., Chambers, K. C., et al. 2016, [arXiv:1612.05243](https://arxiv.org/abs/1612.05243)
- Förster, F. 2019, Transient Name Server Discovery Report 2019-1403, 1
- Frohmaier, C., Sullivan, M., Nugent, P. E., et al. 2019, *MNRAS*, **486**, 2308
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2016, *A&A*, **595**, A2
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1

- Gal-Yam, A., Arcavi, I., Ofek, E. O., et al. 2014, *Natur*, 509, 471
- Gal-Yam, A., Yaron, O., Sass, A., et al. 2019, Transient Name Server AstroNote, 1, 1
- Goldstein, D. A., D'Andrea, C. B., Fischer, J. A., et al. 2015, *AJ*, 150, 82
- Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, 131, 078001
- Hastie, T., Tibshirani, R., & Friedman, J. 2009, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Berlin: Springer)
- Heinze, A. 2019a, CBET, 4645
- Heinze, A. N., Tonry, J. L., Denneau, L., et al. 2018, *AJ*, 156, 241
- Heinze, A. 2019b, CBET, 4648
- Holoien, T. W. S., Brown, J. S., Valley, P. J., et al. 2019, *MNRAS*, 484, 1899
- Holoien, T. W. S., Stanek, K. Z., Kochanek, C. S., et al. 2017, *MNRAS*, 464, 2672
- Horiuchi, S., Beacom, J. F., Kochanek, C. S., et al. 2011, *ApJ*, 738, 154
- Howell, D. A., Sullivan, M., Perrett, K., et al. 2005, *ApJ*, 634, 1190
- Immler, S., Brown, P. J., Milne, P., et al. 2006, *ApJL*, 648, L119
- Immler, S., & Lewin, W. H. G. 2003, in *X-Ray Supernovae*, ed. K. Weiler, Vol. 598 (Berlin: Springer), 91
- Inserra, C., Smartt, S. J., Jerkstrand, A., et al. 2013, *ApJ*, 770, 128
- Kasen, D., & Bildsten, L. 2010, *ApJ*, 717, 245
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kleyna, J. T., Hainaut, O. R., Meech, K. J., et al. 2019, *ApJL*, 874, L20
- Kotak, R., Meikle, P., Pozzo, M., et al. 2006, *ApJL*, 651, L117
- Kromer, M., & Sim, S. A. 2009, *MNRAS*, 398, 1809
- Kulkarni, S. R. 2020, Transient Name Server AstroNote, 1, 1
- Lasker, B. M., Lattanzi, M. G., McLean, B. J., et al. 2008, *AJ*, 136, 735
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, 121, 1395
- Leaman, J., Li, W., Chornock, R., & Filippenko, A. V. 2011, *MNRAS*, 412, 1419
- Li, W., Chornock, R., Leaman, J., et al. 2011a, *MNRAS*, 412, 1473
- Li, W., Leaman, J., Chornock, R., et al. 2011b, *MNRAS*, 412, 1441
- Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. 2016, *ApJS*, 225, 31
- Magnier, E. A., Chambers, K. C., Flewelling, H. A., et al. 2016a, arXiv:1612.05240
- Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., et al. 2016b, arXiv:1612.05242
- Magnier, E. A., Sweeney, W. E., Chambers, K. C., et al. 2016c, arXiv:1612.05244
- Mahabal, A., Rebbapragada, U., Walters, R., et al. 2019, *PASP*, 131, 038002
- Margutti, R., Kamble, A., Milisavljevic, D., et al. 2017, *ApJ*, 835, 140
- Margutti, R., Metzger, B. D., Chornock, R., et al. 2019, *ApJ*, 872, 18
- Martínez, R., Lopez, L. A., Shappee, B. J., et al. 2019, *ApJ*, 892, 144
- McBrien, O. R., Smartt, S. J., Chen, T.-W., et al. 2019, *ApJL*, 885, L23
- Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., & Hložek, R. 2019, *PASP*, 131, 118002
- Narayan, G., Zaidi, T., Soraisam, M. D., et al. 2018, *ApJS*, 236, 9
- Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, *PASP*, 131, 018001
- Perlmutter, S., Gabi, S., Goldhaber, G., et al. 1997, *ApJ*, 483, 565
- Pignata, G., Maza, J., Antezana, R., et al. 2009, in *AIP Conf. Ser.* 1111, ed. G. Giobbi et al. (Melville, NY: AIP), 551
- Prentice, S. J., Maguire, K., Smartt, S. J., et al. 2018, *ApJL*, 865, L3
- Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, 156, 123
- Quimby, R. M. 2006, PhD Thesis, The University of Texas at Austin
- Quimby, R. M., Kulkarni, S. R., Kasliwal, M. M., et al. 2011, *Natur*, 474, 487
- Rabinak, I., & Waxman, E. 2011, *ApJ*, 728, 63
- Reyes, E., Estévez, P. A., Reyes, I., et al. 2018, in 2018 Int. Joint Conf. on Neural Networks (IJCNN) (Piscataway, NJ: IEEE), doi:10.1109/IJCNN.2018.8489627
- Ritter, H., & Kolb, U. 2003, *A&A*, 404, 301
- Ryder, S. D., Sadler, E. M., Subrahmanyan, R., et al. 2004, *MNRAS*, 349, 1093
- Schmidt, B. P., Suntzeff, N. B., Phillips, M. M., et al. 1998, *ApJ*, 507, 46
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
- Smartt, S. J., Chen, T. W., Jerkstrand, A., et al. 2017, *Natur*, 551, 75
- Smartt, S. J., Clark, P., Smith, K. W., et al. 2018, *ATel*, 11727, 1
- Smartt, S. J., Valenti, S., Fraser, M., et al. 2015, *A&A*, 579, A40
- Smith, K. W., Denneau, L., Vincent, J. B., & Weryk, R. 2019a, CBET, 4594
- Smith, K. W., Williams, R. D., Young, D. R., et al. 2019b, *RNAAS*, 3, 26
- Soares-Santos, M., Holz, D. E., Annis, J., et al. 2017, *ApJL*, 848, L16
- Soderberg, A. M., Berger, E., Page, K. L., et al. 2008, *Natur*, 453, 469
- Soderberg, A. M., Chakraborti, S., Pignata, G., et al. 2010, *Natur*, 463, 513
- Sonnett, S., Meech, K., Jedicke, R., et al. 2013, *PASP*, 125, 456
- Srivastav, S., Smartt, S. J., Leloudas, G., et al. 2020, *ApJ*, 892, L24
- Stalder, B., Tonry, J., Smartt, S. J., et al. 2017, *ApJ*, 850, 149
- Stockdale, C. J., Williams, C. L., Weiler, K. W., et al. 2007, *ApJ*, 671, 689
- Sutaria, F. K., Chandra, P., Bhatnagar, S., & Ray, A. 2003, *A&A*, 397, 1011
- Szalay, A. S., Gray, J., Fekete, G., et al. 2007, *Indexing the Sphere with the Hierarchical Triangular Mesh*, Tech. Rep. MSR-TR-2005-123
- Tachibana, Y., & Miller, A. A. 2018, *PASP*, 130, 128001
- Takada, M., Ellis, R. S., Chiba, M., et al. 2014, *PASJ*, 66, R1
- Tomaney, A. B., & Crofts, A. P. S. 1996, *AJ*, 112, 2872
- Tonry, J. L. 2011, *PASP*, 123, 58
- Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, *PASP*, 130, 064505
- Valenti, S. 2017, David, Sand, J., et al. *ApJL*, 848, L24
- Vereš, P., Jedicke, R., Wainscoat, R., et al. 2009, *Icarus*, 203, 472
- Véron-Cetty, M. P., & Véron, P. 2010, *A&A*, 518, A10
- Waters, C. Z., Magnier, E. A., Price, P. A., et al. 2016, arXiv:1612.05245
- Wood-Vasey, W. M., Miknaitis, G., Stubbs, C. W., et al. 2007, *ApJ*, 666, 694
- Wright, D. 2015, PhD thesis, Queen's Univ. Belfast
- Wright, D. E., Lintott, C. J., Smartt, S. J., et al. 2017, *MNRAS*, 472, 1315
- Wright, D. E., Smartt, S. J., Smith, K. W., et al. 2015, *MNRAS*, 449, 451
- Yaron, O., & Gal-Yam, A. 2012, *PASP*, 124, 668
- Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, *ApJ*, 830, 27