

The gravitational-wave optical transient observer (GOTO) data pipeline and workflow for transient discovery

J. D. Lyman¹,¹★ D. O'Neill^{1,2} T. Killestein¹ D. Jarvis³ A. Kumar^{1,4} K. Ulaczyk¹ K. Ackley¹ P. Chote¹ M. J. Dyer^{3,5} M. Pursiainen¹ D. Steeghs¹ B. Godson¹ M. Magee¹ J. R. Mullaney³ B. Warwick¹ S. Belkin⁶ D. K. Galloway^{6,7} G. Ramsay⁸ V. S. Dhillon^{3,9} P. O'Brien¹⁰ K. Noysena¹¹ R. Kotak¹² R. P. Breton¹³ L. K. Nuttall¹⁴ B. Gompertz² D. Pollacco¹ J. Casares^{9,15} D. L. Coppejans¹ R. A. J. Eyles-Ferris¹⁰ O. Graur^{14,16} L. Kelsey¹⁷ M. R. Kennedy¹⁸ A. Levan^{1,19} S. Littlefair³ S. Mandhai¹³ D. Mata Sánchez^{9,15} S. Mattila^{12,20} J. McCormac¹ S. Moran¹⁰ C. Phillips¹ K. Pu⁶ A. Sahu¹ M. Shrestha⁶ E. Stanway¹ R. L. C. Starling¹⁰ L. Vincetti^{8,21,22} E. Wickens¹⁴ and K. Wiersema²³

Affiliations are listed at the end of the paper

Accepted 2026 June 10. Received 2026 June 10; in original form 2026 February 27

ABSTRACT

Wide-field and high-cadence sky surveys are the first step in the chain of discovery and characterization of astrophysical transients such as supernovae, kilonovae, and tidal disruption events, each linked to the varied demise of stellar systems. The gravitational-wave optical transient observer (GOTO) is a telescope array of thirty-two 40 cm unit telescopes split over two almost antipodal sites. It performs a regular time-domain sky-survey in the optical to ~ 20 mag in addition to immediate scheduling of follow-up observations at the locations of external multiwavelength and multimessenger triggers. To facilitate the timely recovery of optical counterparts to these triggers, as well as the presence of serendipitous discoveries of astrophysical transients in the regular sky-survey, a low-latency data pipeline and workflow was developed. The implementation of this workflow is described herein and the quality of GOTO data delivered by it assessed, alongside its performance for prompt transient recovery. Utilizing difference image analysis to identify candidate discoveries, the process is typically complete ~ 7 min after shutter close on the telescope. We further describe later processing of these candidates – both automated and human-in-the-loop – including reporting to the wider community and the triggering of more detailed observations, with a focus on immediate, intranight characterization. The workflow is meeting the needs of GOTO to promptly discover, report, and characterize infant transients. Nevertheless, areas for further development and improvements are also highlighted.

Key words: software: data analysis – data methods – transients.

1 INTRODUCTION

Very regular patrol of the visible night-sky ($\lesssim 3$ d cadence) is now possible from a given site to significant depths ($m_{\text{opt}} \sim 18 - 22$ mag) and with atmosphere-limited spatial resolution. This is largely thanks to major breakthroughs in the manufacture and affordability of large format charge-coupled device (CCD) and complementary metal-oxide-semiconductor (CMOS) sensors, as well as the affordability of high-quality optical tube assemblies (OTA) being available at scale. Surveys built around this model are a rich resource for time-domain astrophysics.

The need for wide-field and high-cadence monitoring of the sky has been amplified following the emergence of immediate, autonomous alerting in the field of multiwavelength and multimessenger time-domain astrophysics (exemplified by B. P. Abbott et al. 2017). This can necessitate rapid coverage of large localization areas – delivered in low-latency by high-energy or gravitational-wave observatories – to pinpoint optical counterparts. Beyond these targeted searches for counterparts, within the broad scope of time-domain astrophysics an ever-widening diversity of optical transients have come to the fore. This outcome is due in no small part to the advent of recent and on-going large-scale synoptic time-domain facilities such as the All-Sky Automated Survey for SuperNovae (C. S. Kochanek

* E-mail: j.d.lyman@warwick.ac.uk

et al. 2017), the Asteroid Terrestrial-impact Last Alert System (ATLAS; J. L. Tonry et al. 2018b), BlackGEM (P. J. Groot et al. 2024), the Large Array Survey Telescope (E. O. Ofek et al. 2023), the Panoramic Survey Telescope and Rapid Response System (K. C. Chambers et al. 2016), the Wide Field Survey Telescope (T. Wang et al. 2023), and the Zwicky Transient Facility (ZTF; E. C. Bellm et al. 2019). The newly uncovered behaviour of these transients, along with unprecedented populations of more common-place variants, is continually challenging our understanding of the demise of stars: the explosive endpoints of stellar evolution. The prospects for further discovery are encouraging with upcoming facilities including the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; Ž. Ivezić et al. 2019), the La Silla Schmidt Southern Survey and the *Nancy Grace Roman Space Telescope* (D. Spergel et al. 2015) set to probe new regions of luminosity, time-scale, and cosmological distance for transients, coupled with overwhelming numbers of discoveries.

The gravitational-wave optical transient observer (GOTO) is a recently commissioned two-node telescope array, building on a prototype design (D. Steeghs et al. 2022), with a principal science focus of obtaining rapid optical electromagnetic constraints for a variety of multiwavelength and multimessenger events. The wide-field of view and robotic operations of the constituent GOTO nodes enables quick and automated reaction to external triggers, even with coarse localization. Between observation campaigns triggered by these external triggers, a routine all-sky survey prioritizing more even coverage of the whole sky is performed. GOTO acts complementarily to the ongoing surveys listed previously and adds crucial longitudinal coverage (Section 2) to the global effort for discovering electromagnetic transients in their infancy, during which follow-up observations can be at their most diagnostic for understanding the nature of their explosions and progenitors. Furthermore, as GOTO’s primary science aim is reacting to multimessenger and multiwavelength, they occupy a consequently dominant portion of the observing mode and strategy. This is in some contrast to other wide-field facilities, which typically require more tensioning of their observing priorities with other science goals surrounding time-domain astrophysics. GOTO is therefore able to react to alerts not routinely covered by other surveys promptly, through a combination of relaxed triggering criteria to follow more marginal GW alerts (B. P. Gompertz et al. 2020; Ackley, Dyer et al., in preparation), and coverage of relatively high-rate and large localization triggers such as *Fermi* Gamma-ray Burst Monitor alerts (C. Meegan et al. 2009) with location uncertainties of tens of degrees, as seen in S. Belkin et al. (2024) and A. Kumar et al. (2025).

With wide-field sky-surveys acting as transient discovery machines, an emphasis on rapid dissemination of new discoveries to the community has permitted expedited characterization of these discoveries during their crucial infant stages. This systematic excavation of the short time-scale parameter space has revealed previously unseen behaviour – both intrinsically rapidly-evolving transients, and short-lived phenomena in the infant stages of more typical supernova.

This paper describes the immediate data processing and transient discovery workflow operations for GOTO facility: from shutter close of the telescope to announcement to the community and further characterization by other facilities.

2 GOTO OVERVIEW

The GOTO project design, hardware, general performance characteristics, and observing strategies are fully described in Steeghs et al. (in preparation) (see also M. J. Dyer et al. 2024). Here we briefly overview the facility, and in Section 2.1 detail the observing strategy and survey modes conducive to transient discovery.

GOTO collectively describes a telescope array composed of a large number of small-aperture unit-telescopes (UTs) co-located on a small number of large mounts. There are two nodes – one in La Palma, Canary Islands (28°45′36″ N, 17°52′45″ W) and one in Siding Spring, Australia (31°16′24″ S, 149°03′51″ E). Each node comprises two mounts, with eight UTs on each mount, for a total of 32 UTs. The individual field of view (FoV) of each UT is a little over 5 deg², and they are arranged into a 4 × 2 rectangular grid of minimally overlapping FoVs per mount, giving a mount FoV of ∼ 45 deg². UTs are equipped with ∼50 megapixel (∼ 8000 × 6000 pixels) FLI ML50100 CCD sensors with 1.3 arcsec per pixel plate-scale. The UTs each have a filter wheel containing a Baader LRGB filter set.

Each of GOTO’s mounts (and the parent domes) are operated autonomously and independently by separate instances of the control software, G-TECS (GOTO Telescope Control System; M. J. Dyer et al. 2020), whilst a central scheduler dictates the optimal next observation(s) for the facility as a whole. Data from each UT are first written to files on local discs at each node, whereupon they are immediately transferred to a central data centre in the UK for processing and long-term storage.

A table of specific relevant terms and their meaning in the context of discussing GOTO operations and pipeline in subsequent sections is provided in Table 1 for reference. Terms with long-established meaning (e.g. image types like *bias*, *science*) are omitted.

2.1 Observing strategy

GOTO generally observes within a fixed grid on the sky, with the grid size dictated by the footprint of a single GOTO mount. These pointings are referred to as *tiles* and the celestial sphere requires 1024 *tiles* to be fully covered (M. J. Dyer 2020).¹ The use of a fixed grid of positions simplifies the difference image analysis steps and additionally allows us to treat the individual mounts of GOTO as a single telescope since the provenance of an exposure is only relative for the data reduction steps. The telescope scheduler operates on this assumption, such that a given *tile* is valid for all mounts and, assuming it is the next highest ranking *tile* to be observed, is picked-up on a first-come-first-served basis by the mounts at a given site.

¹<https://github.com/GOTO-OBS/goto-tile>

Table 1. Definitions of terms used to refer to data types and products.

Term	Description
tile	A pointing location within a fixed grid covering all-sky, which is defined by the footprint of an individual GOTO array (Section 2.1).
event	An external trigger causing specific tile observations to be made (Section 2.1).
single	A reduced image from an individual exposure (Section 4.5).
set	A stack of single made from multiple exposures taken at a single pointing (Section 4.6).
template	A historical set image taken with a larger number of exposures, and identified as being of good image quality.
difference	A set image after subtraction of a corresponding template (Section 4.7.2).
super	A large stack, typically monthly, of bias, dark, or flat calibration frames used for the low-latency processing (Section 4.4.2).
columntrap	A specific type of super calibration image made of a large number of single images with low sky-background that identifies steps in background level along CCD columns (Section 4.4.3).

Observations of tile pointings during regular operations are made in two modes, namely the fall-back ‘Sky-survey’, and ‘Follow-up’. The former is GOTO’s default observation strategy when no higher ranking observation types are available. Specific tile pointings are ranked higher based on a combination of when they were last observed, and their current airmass, but otherwise provide an unbiased survey covering the whole sky (for full details, see M. J. Dyer et al. 2018). Within the ‘Sky-survey’, groups of 4×45 s exposures are taken per pointing and have a unique running integer value (setnum) assigned to them. These exposures are individually reduced and analysed into single images (Section 4.5), before being combined into a set image for increased depth (Section 4.6).

GOTO continually listens to VOEvent streams of various multimessenger and multiwavelength alerting facilities (M. J. Dyer et al. 2020)² – events in the GOTO scheduler parlance – which are formulated into observing campaigns using pre-built strategies depending on the origin and exact nature of the alert. These ‘Follow-up’ observations will trump the ‘Sky-survey’ and insert a series of higher ranking tile pointings, typically with a higher cadence. Indeed different alerts can also override other follow-up, depending on the science ranking of the alert. ‘Follow-up’ usually comprises 4×90 s exposures per pointing, which are reduced and analysed analogously to the ‘Sky-survey’ pointings.

Additionally, the GOTO project began taking template images for each tile from early 2023 when the full hardware implementation was complete. Over > 99 per cent of template images currently used had been assigned as such by mid 2024. These template images are composed of 12×90 s set exposures, significantly deeper than the usual all-sky survey 4×45 s or GRB/GW follow-up 4×90 s stacks, but otherwise processed identically. The use of the template images in difference imaging is detailed in Section 4.7.2.

Observations taken as part of transient ‘Sky-survey’ and ‘Follow-up’ are taken in the wider L filter ($\sim 400 - 700$ nm). This filter is currently used exclusively during transient-related observations as it offers the benefit of higher throughput, allowing GOTO to achieve more sensitive image-depths, at the cost of colour information, which is typically delegated to external facilities to acquire. This choice also simplified template image build-up (since they only need to be acquired in L), allowing GOTO to be online and effective for transient discovery faster.

The longitudinal distribution of the GOTO-nodes means that there is no time when both sites are making night-sky observations. However, their latitudinal positions means a region of sky with declination $-20 < \delta < 20$ is accessible at relatively low-airmass from both sites. This region can consequently benefit from added cadence and indeed real-time follow-up campaigns with external facilities aimed at identifying infant transients have exploited this, such as the GOTO Fast Analysis and Spectroscopy of Transients (GOTO-FAST) survey (B. Godson et al. 2023; Godson et al., in preparation).

3 HARDWARE IMPLEMENTATION

All processing is performed on a GOTO-dedicated data centre at the University of Warwick, UK, made from a mix of computing and high-capacity storage servers. Some of the computing servers are dedicated to other aspects of GOTO operations. Each computing server used by the pipeline has 32 physical CPU cores (virtually split to provide 64 independent logical processing) and 256 GB of memory, which allows multiple processes to concurrently perform the typically memory-intensive and sometimes computationally intensive tasks associated with reduction and analysis of GOTO images. Following the telescope control software writing exposures to disc at each site, the data are immediately transferred to this Warwick server for processing (Section 4.1). The decision to transfer data immediately instead of processing at the observatory sites was made to ameliorate logistical challenges with performing maintenance tasks on the servers, since no dedicated local staff effort is present on site. Cloud-based hosting was not considered due to concerns over cost, potential volatility of data in the absence of secure long-term funding, and the desire to benefit from working with local scientific computing expertise at Warwick.

²<https://github.com/GOTO-OBS/gtecs-alert>

A series of six computing servers act as *workers* for the tasks of the pipeline, with two of these also sharing duties of scheduling the tasks and one providing the webserver for pipeline monitoring (see Section 4.2). The marshall web framework (Section 5), and its own task worker, run on a separate, dedicated server.

Two servers host the data base cluster, which serve data bases for the pipeline and marshall, alongside other ancillary data bases (e.g. contextual information). One server is the primary, with the other a standby which is populated via PostgreSQL streaming replication. The standby is a constantly up-to-date read-only copy of the primary data base and is the entry-point for user and marshall queries that need only to fetch data, thus lightening the load on the primary, to which the pipeline intensively reads and writes records. The standby is also capable of being immediately promoted to primary, should the primary server undergo failure.

Data are stored locally to the cluster on a rolling retention policy, typically six months' worth of processed image data. This time-window is more than sufficient to ensure required access for further processing, forced photometry etc. is expedient due to high-throughput network links between servers in the server. Raw and processed data are additionally archived shortly after processing to a distinct PB-scale storage server managed by the Scientific Computing Research Technology Platform of the University of Warwick.

4 PIPELINE IMPLEMENTATION

The pipeline is intended for rapid transient identification and dissemination. Precise photometric calibration is not emphasized and deferred to future data products from GOTO (Jarvis et al., in preparation). Following the tradition of overly grandiose naming for software packages, we dubbed the data reduction and analysis pipeline KADMILOS during its development, and will refer to it as such hereafter.

The data processing of astronomical data from a facility like GOTO largely falls under the Extract-Transform-Load (ETL) model prevalent in many other data handling systems (P. Vassiliadis 2009). Succinctly, data from one or more sources is input, undergoes transformation (in this case requiring significant amounts of data reduction and analysis), before storing distilled outputs of the data in such a way that they can be disseminated to users or further processing tools. Practically, components of the pipeline are responsible for taking raw image data from each GOTO UT, and outputting final image reductions, photometric measurements, and new difference image sources. We here detail various aspects of the implementation of KADMILOS, including the associated solutions for transferring data, monitoring the processing, and storing the output. An overview schematic of the data and software flow is shown in Fig. 1.

4.1 Transferring raw data from the observatories

As the computing infrastructure is held at the University of Warwick, raw data in the form of FITS files³ must be copied immediately from each site upon writing before being picked up by KADMILOS. A small package RAWTRANSFER,⁴ builds around `ssh` data transfer, for this purpose. Additional safeguards are in place to ensure the uniqueness of file transfers, error handling for invalid FITS files, retries of transfers in the case of, e.g. temporary network down-time, and a data base record of each successful transfer. A copy of RAWTRANSFER runs locally at each site and a monitoring thread listens for filesystem events from the Linux kernel *inotify* subsystem indicating new a raw data file was written by the telescope software G-TECS. The file is then queue-scheduled for transfer to Warwick, with multiple worker threads processing the queue in sequence. In the absence of network issues, this queue processing is effectively immediate.

Using a PostgreSQL data base for the ground-truth of transfers allows for the creation of triggers (functions executed when a specific operation occurs, e.g. insertion of a new row into a table) which themselves emit a NOTIFY message. External services can listen to these NOTIFY messages. KADMILOS does exactly this for the purpose of initiating the pipeline processing cascade for the transferred file. The 50th and 95th percentile timings for FITS checking and transfer from the La Palma site to Warwick are $P_{50} = 9$ and $P_{95} = 31$ s post-file-writing. The La Palma node benefits from a dedicated level 2 VLAN connection to Warwick. For the Siding Spring site, we are at the mercy of a dynamic route over a public network between gateway machines at Warwick and Australia National University, reflected in slower transfer timings: $P_{50} = 21$ and $P_{95} = 207$ s. The typical size of a raw science frame transferred is comparatively small at 42 MB since the integer data extension can be stored efficiently.

4.2 Apache airflow: orchestrating KADMILOS

A primary concern with deciding on the pipeline implementation was a robustness of tasks execution on different schedules, the ability for multiple work-flow patterns to be decided on-the-fly, the means to easily re-process (or partially re-process) batches of data, and a monitoring system to identify issues as they may arise.

Apache Airflow⁵ (Airflow) well suited these requirements, with the added bonuses of being an open source project implemented in Python, the language also used for KADMILOS. Airflow is particularly adept at defining ETL pipelines, and offers easy horizontal scalability that is attractive for a project like GOTO where we require immediate processing of significant chunks of data from 16 UTs at $\sim$ minute intervals almost continually during good weather at both sites.

³https://fits.gsfc.nasa.gov/fits_standard.html

⁴<https://github.com/GOTO-OBS/rawtransfer>

⁵<https://airflow.apache.org/>




Figure 1. A schematic showing the data flow and software implementation of the GOTO transient discovery workflow. Full details are given in the respective sections. Succinctly, data acquired and written by the GOTO mounts to a local file system at the respective observatory (see M. J. Dyer 2020). A RAWTRANSFER service (Section 4.1) at each site monitors these local file systems, packages the new FITS files as they are created and transfers them to a large capacity file system at the University of Warwick, writing the details of the transfer to a record in the transfer data base, also held in Warwick. The insertion of the record emits a NOTIFY signal which tells the Airflow (Section 4.2) scheduler that a new raw insertion DAG (Section 4.2.1) should be scheduled to process the file. The tasks of the KADMILOS pipeline (e.g. Section 4.5) are performed by a cluster of Celery workers, scheduled by Airflow to follow the respective DAG's logic. As part of these tasks, the workers write processed data and results to a central filesystem and data base. These central data resources are pulled by the GOTO marshal (Section 5), where additional metadata such as contextual information is fetched, and source associations are made. The result source data is then visualized, vetted, and reported by collaboration members via a webserver.

Workflows in Airflow are designed as directed acyclic graphs (DAGs) of tasks, where each task can have a conditional dependence on one or more upstream tasks. A specific processing instance of a workflow is referred to as a DAG run, which comprises one or more task instances. Airflow comprises four main components, a metadata data base, one or more schedulers, one or more workers, and a web server. Extensive documentation is available on the Airflow webpages, and so we only briefly summarize the components here.

The metadata data base is a relational data base created by Airflow to allow it to monitor the processing environment and its own actions, as well as keep a record of all historical actions such as initiating DAG runs, and the outcome of task instances. We implement it using PostgreSQL for our deployment, which allows for multiple schedulers to be run concurrently.

The scheduler makes primary use of the metadata data base, inserting and updating records to journal and monitor the status of all task instance and DAG runs – deciding what needs to be executed next on a regular interval. For example, it identifies which tasks are currently due for processing and passes those to the workers via a task queuing system, marks downstream tasks of any recently completed tasks as requiring to be processed, and updates the status of DAG runs based on outcomes of their child tasks. Additionally the scheduler checks the outputs of the workers to update whether tasks are currently being processed, have successfully completed, have failed etc. A scheduler instance naturally runs on a cycle of operations, each taking finite time. To improve the prompt scheduling of tasks and increase their throughput to be picked up by the workers, we run two scheduler instances on separate computer servers. The use of PostgreSQL as the metadata data base allows for record-level locking and removes collision and deadlock potential from having multiple schedulers access it pseudo-concurrently. The scheduler is also responsible for parsing the files that define the workflows (or DAGs), such that their order of task dependence and scheduling requirements can be stored in the metadata data base.

The workers are the processes that process tasks defined in a workflow (or DAG). Storing KADMILOS on a shared filesystem in the data-centre allows each worker server to see identical code for the parsing and execution of DAGs and their constituent tasks, avoiding any conflicts that may arise during updates to either. Airflow, by default, makes use of Celery⁶ as an executor of tasks. Celery is a distributed task queue, which takes care of ensuring the worker processes pick up, execute, and report back on tasks. Many worker processes can be spawned on a given server, with distribution to multiple servers also readily supported. The workers update the metadata data base once tasks are completed (and if they failed, succeeded etc.), which the scheduler will query on its next iteration to determine the next course of action for each DAG.

Lastly, the Airflow web server provides an informative UI for interrogating the status of the scheduler, workers, and DAG/task processing. It also provides a user interface to the Airflow API, allowing the user to manipulate the processing of DAG runs and tasks, especially useful for re-running previously failed tasks due to a transient issue (e.g. network or filesystem). The web server is routinely checked by operations teams to assess the health of the system, and DAG run failures are propagated to the GOTO internal communication platform (Slack) for assessment.

4.2.1 DAG details

KADMILOS uses both triggered DAGs and scheduled DAGs. The latter run at specific intervals or at a given time of day. For the former either an external process, or another DAG, is responsible for initiating a DAG run.

The primary triggered DAGs are named *raw_insert* and *single_and_set_creation*. Upon insertion of a record in the RAWTRANSFER data base (Section 4.1), due to the successful transfer of a FITS file, a notification is issued with a payload including the file path of the newly transferred file. This payload is picked up by a dedicated listener (running as a system service) which triggers the *raw_insert* DAG using the payload. The cascade of KADMILOS processing for that file then ensues. For *science* frames the *single_and_set_creation* DAG is then triggered after some rudimentary quality checking of the data. This latter DAG takes the raw data through the steps detailed in Sections 4.4 to 4.7.

The primary DAG running on a schedule is the creation of *super* calibration frames at each site: *super_calibration_la_palma*, *super_calibration_siding_spring* (Section 4.4.2).

The above are the DAGs most pertinent to the transient workflow in GOTO, although additional DAGs for housekeeping as well as deprecated workflows (e.g. forced photometry) are visible in the public repository.⁷

DAGs are written to be intentionally idempotent. This naturally facilitates the ability to re-process (or ‘back-fill’) data. (Non-astrophysical) transient issues such as an inaccessible data base connection, or compromised data integrity due to a hardware issue, cause DAGs to fail. The web UI or command-line Interface to Airflow allows the GOTO operations team to promptly identify and rerun the failed DAGs once the issue is corrected. Due to the splitting of DAGs into individual component tasks, it is also typically the case that only a sub-set of the processing for each failed DAG needs to be attempted again.

4.2.2 DAG versioning

Given the slowly changing dimensions of GOTO’s configuration (e.g. replacement of cameras) we employ a DAG versioning system based on the notion of a ‘hardware version’ for GOTO. Once a change is made to the system that alters the expected flow of processing and/or invalidates calibration files for a given UT, we generate a new hardware configuration file to describe it, which automatically authors a new version for those DAGs that rely on knowledge of the hardware set-up (i.e. *super* calibration creation, Section 4.4.2), and appends the hardware version to the name of the DAG. An example of a hardware configuration file can be found in the frozen public repository of KADMILOS.⁸ We note that the release of Airflow 3.0 in April 2025 allows for DAG versioning natively in the Airflow ecosystem, and is a planned upgrade to KADMILOS.

⁶<https://docs.celeryq.dev/en/stable/>

⁷<https://github.com/GOTO-OBS/kadmilos-public/tree/main/kadmilos/dags>

⁸https://github.com/GOTO-OBS/kadmilos-public/blob/main/kadmilos/example.hardware_config_2020_0.yml

As part of this hardware versioning, a camera configuration file is held to store immutable properties of the cameras being used, such as the expected shape of the data, and the pixel scale when mounted on a GOTO UT. This configuration is used to validate raw data and in some processing steps.

4.2.3 Overall impressions

Airflow is not typically employed or designed for very high-throughput, low-latency processing as we are performing with GOTO. Indeed the tool is specifically advertised as such, citing concepts like a non-focus on minimizing lags for scheduling tasks (and their subsequent pick-up by workers). Although undoubtedly additional performance could be gleaned from a concerted effort in tweaking of the Airflow parameters and better DAG design (in addition to time savings through optimization of the tasks themselves), the time-scales for processing have met our original aims for the project, and so effort has not yet been directed to this aim. Considering the large and typically continuous wave of DAG run scheduling, in terms of overall throughput, Airflow has performed admirably in the face of our demands. Currently there have been over 13 million DAG runs and 130 million tasks instances fully documented in its metadata data base and counting.⁹ The orchestration by Airflow rarely encounters issues, except those externally enforced, and scheduling and execution of tasks maintain the same performance level as it did initially.

4.3 PostgreSQL: persistent data products storage

Alongside the Airflow metadata data base, the choice for data products from KADMILOS is also a PostgreSQL data base, in the same cluster. The data base contains schemas holding: raw, calibration, and image data; photometry; scheduler events and surveys; externally derived catalogues.

For photometry tables, Q3C spatial indices (Quad Tree Cube; S. Kposov & O. Bartunov 2006, 2019) are added on the RA and Declination columns to allow for performant coordinate-based querying (e.g. cone or polygon searches). A copy of the ATLAS-REFCAT2 catalogue (J. L. Tonry et al. 2018a), used for photometric calibration (Section 4.5.5), is also held locally in a table of the data base, for which Q3C spatial indexes were also generated. The photometry tables are by far the largest, to the point where partitioning significantly helps with their management and querying. These partitions are done on the `date_mid` timestamp value, with each photometry table holding one calendar month of records. Cone searches of $\lesssim 1$ arcmin take ~ 0.1 s to search 1 yr of single photometry partitions – $\mathcal{O}(10^{11})$ records.

4.4 Basic raw data processing

The following subsections describe the CCD reduction applied to every new GOTO *science* frame in low-latency, after ingestion into KADMILOS.

Other image types of calibration frames (*bias*, *dark*, *flat*) are ignored during immediate low-latency processing as they are not needed individually. Instead, these frames are batch processed into *super* calibration frames (Section 4.4.2), for use in the reduction of future *science* frames.

4.4.1 Pre-reduction corrections

The edges of the CCD arrays – ~ 10 pixels along each edge (M. J. Dyer 2020; D. Steeghs et al. 2022) – are typically compromised due to detector edge-effects and increased noise level. These regions are therefore trimmed from all frames. The median value in an overscan region (removed during prior trimming) is then subtracted from all remaining pixel values, to account for the hardware bias level. The data are then multiplied by a gain value (e^-/count) for the specific camera channel, to convert the pixel values from counts to electrons. Gain values for every camera were determined for each of the two readout channels of the sensors, with the mean and variance of a set of differently illuminated flat-fields used to calculate them.

4.4.2 Generation of super calibration frames

Regular CCD processing involves the removal of pixel-to-pixel pattern noise in the bias level, and any dark current contribution to the pixel values. Pixel-to-pixel sensitivity to incident photons and an overall illumination correction are needed to further standardize the pixel values across the CCD. Finally, dead, hot, or otherwise non-linear, pixels must be identified to prevent their erroneous readings propagating into the analysis. The sensors in GOTO are maintained at 253K during operation, resulting in only a modest dark current. Nevertheless, the fraction of non-linear and hot pixels is elevated above that of high-grade, cryostat enclosed CCDs. Full details of the sensor operation will appear in Steeghs et al. (in preparation).

To correct for these effects, KADMILOS generates monthly banks of *super* calibration frames. These frames are heavily stacked versions of individual calibration frames, produced per camera for those within the current hardware configuration (Section 4.2.2).

⁹In practice, using Airflow's archiving command, old metadata records are moved to separate tables from the ones frequently queried by the scheduler. We employ this as a best practice but saw no significant degradation of scheduler performance prior to this.

Images of type `bias`, `dark`, and `flat` are taken on a regular schedule each evening and morning by the GOTO control software (M. J. Dyer et al. 2020). A scheduled DAG spawns tasks to search a maximum of 3 months previous to fetch up to 80 observations of each image type, starting from the most recently acquired. If fewer than 50 of any type is found, the `super` calibration creation is skipped for that camera (and low-latency processing will fall back to the previous month's calibration images). Individual frames are combined pixel-wise with sigma-clipping of pixels. No weighting is applied to individual frames during combination. For `dark` frames, the combination is performed per unique camera–exposure-time combination. Exposure time subsets are 30, 45, 60, or 90 s as these match the possible choices used for `science` images in GOTO's observing strategy. Data are scaled within these subsets by their exact exposure time. The `flat` frame combination is performed per unique camera–UT–filter combination, and is preceded by scaling each frame by its respective inverse median.

Some crude quality assessment is performed on individual calibration frames when they are initially read, checked, and stored by KADMILOS during the `raw_insert` DAG. For `bias` and `dark` (and also `science`) image types, these checks solely identify obviously problematic data that would fail during processing, e.g. they ensure the background level is within a sensible range and the shape of the data array is as expected for the camera. For `flat` frames, however, the pixel-value distribution is further analysed as part of these checks. The skew and kurtosis values of this distribution were found to be a strong indicator of both the presence of cloud cover, and the abundance of astrophysical objects in the frames (due to extended exposure times during dark conditions). As each of these effects compromises the quality of the flat, we set empirically defined limits on these numbers. Any frame lying outside these limits is appropriately flagged during the data base ingestion stage, such that it is not considered for inclusion in any `superflat` frame stack.

The `superbias` frames are simply subtracted from the `science` data during reduction, but the `superdark` and `flat` frames additionally allow for the identification of misbehaving pixels. This task calls for the additional creation and storage of a mask alongside the parent calibration data. Hot or dead pixels are identified in the `dark` frame stack by calculating the ninety percentile range (P_{90}) of pixel-wise distribution of values in the constituent individual frames. Pixels in the `superdark` frame that have a P_{90} which is itself more than six standard deviations from the distribution of P_{90} values of all pixels are flagged due to erratic behaviour. Further, the highest 0.1 per cent of pixel values are additionally flagged as the population of hot pixels. Non-linear pixel responses are identified through ordering the individual `flat` frames by their median pixel value. The first and last quartiles of frames ordered thus are separately combined. A 'ratio' frame is calculated from a division of these stacks, which is modelled with a smoothly varying background using the same procedure as done for `science` frame reduction (Section 4.5.2). Any pixels in the background-subtracted 'ratio' frame that are more than 7σ deviated based on the local background root mean square (RMS) are flagged.

Currently, `super` calibration frame creation is a trusted process – i.e. newly created files are employed in low-latency image processing immediately after generation. Significant failures are not typically seen owing to the selection of a large number of individual calibration frames from a significant time-period (alongside quality control of `flat` frames). A more secure method, however, is envisaged whereby manual inspection of the `super` calibration frames is performed, prior to them being made active by the pipeline. The lack of poor performance by the `super` stacking routines has de-prioritized adding in this stage so far.

After applying these standard CCD corrections to `science` frames, a remaining feature was seen in the data from manual inspection during development of the pipeline. The feature manifests due to trapped charge in columns from defective pixels, which calls for an additional `super` calibration frame to be produced and used in our reduction recipe (see the next section).

4.4.3 Identifying and correcting column traps

In the transfer of charge along the readout axis (columns, in this instance) of a CCD, electrons can become 'stuck' in the semiconductor pixel wells. This trapping can cause a partial, or entire loss of signal along the remainder of the column after the misbehaving pixel, commonly referred to as column traps or charge traps. Column-wise traces of pixel values are shown in Fig. 2. To create these traces, a 700 pixel wide region of 500 randomly chosen `single` images (i.e. `science` frames that have been `bias`-, `dark`-, and `flat`-corrected) taken during February 2025 by GOTO-2 UT2 were collected. (This sensor, and the specific region highlighted in the figure, were chosen as a region of elevated column trap density to illustrate the method.) These data were then combined, pixel-wise, using a weighted (inverse background variance) mean with sigma-clipping. These frames were selected for having a low sky-background level of < 400 counts since bright sky conditions make the identification of these (often subtle) bad column behaviours difficult, and had their backgrounds removed (see Section 4.5.2) prior to combination. Manual inspection of the stacked data was used to identify the columns displayed.

The functional form of the traces seen can be variably best described as a constant, or a constant with a step function acting over 1 pixel. A series of model fitting is therefore performed per-column using LMFIT to find the optimal description of the trace and identify a pedestal level which can correct each column of the CCDs, as follows.

First, a constant model is fitted to each column using least squares. Since only a small fraction of columns display a column trap, these initial constant model fits are used to identify them under the assumption that most columns should be well-described by the constant model. To lessen the effect of surviving hot or dead pixels in the stacked image, the constant model is fitted to a median-filtered version of each row using a kernel size of five pixels. Per image, a conservative constraint indicating a poor fit is established as:

$$\chi_{\text{red, thresh}}^2 > \bar{\chi}_{\text{red, i}}^2 + 8 \times \text{med}(|\chi_{\text{red, i}}^2 - \bar{\chi}_{\text{red, i}}^2|), \quad (1)$$




Figure 2. Column-wise traces of pixel value in deep-stacked single images of a GOTO camera highlighting various column-trap behaviour. Each panel shows the trace for a visually identified bad column (dark line), alongside the trace for the neighbouring, good, column (light line). The top three panels show ‘step’ features in the traces of varying depths, with the second panel also showing some offset *above* that of the neighbouring good column. The bottom two panels show overall offsets from the good trace along the entire column length.




Figure 3. The use of a column-wise step-model to correct column charge trap defects in GOTO data. Left: A deep stack of 500 single science images prior to the correction. Middle: The step-model of the column traps following the procedure described in the text (Section 4.4.3). Right: Subtraction of the step-model from the uncorrected data showing the removal of almost all trap features. Very subtle and or small-length traps remain, visible due to aggressive visualization scaling. The read-noise alone in a given GOTO image dominates over the amplitude of these features. Axis values relate to pixel coordinates and the colour scale employs a broken log-linear scheme to enhance contrast at the low background level.

where $\chi_{\text{red},i}^2$ is the reduced chi-squared of the i th column, and an overhead tilde indicates the median of the corresponding set. A separate constraint on the level of the constant model for the i th column, C_i , is also established:

$$C_{\text{thresh}} < \tilde{C}_i - 2 \times \text{med}(|C_i - \tilde{C}_i|). \quad (2)$$

Any constant model satisfying either or both of these criteria is put forward as a candidate column trap. The first and last 4 columns of the CCD are not considered, as edge effects introduced in reduction resulted in these being typically always flagged as candidates.

A one-step model is then fit to each candidate column. The step is forced to act over an infinitesimal length, to mimic the discrete change in level from one pixel to the next. As such a function cannot be differentiated, differential evolution is used for the fitting method (R. Storn & K. Price 1997). As for the constant model, the threshold of equation (1) is applied to the χ_{red}^2 values of the one-step models. For any columns still meeting this criterion, a two-step model is fitted, allowing for two discrete jumps in the pixel value trace. Where this two-step model produces a better χ_{red}^2 value than the one-step model, the two-step model is used.

A final `columntrap` calibration frame is produced using these one- and two-step model values, and filled with zero elsewhere. The process is visualized in Fig. 3. In regular creation of `supercolumntrap` calibration frames KADMILOS uses between 120 and 200 single images since it was found stacking much beyond ~ 120 single science images gave negligible changes in the calibration output. The addition of this step model cleanly removes those columns with strongly defective charge transfer – the difference in pixel




Figure 4. Column-wise traces of pixel value in `columntrap` calibration frames showing the evolution of column-trap behaviour. Each panel shows the trace for the same columns as shown in Fig. 2, with the colour indicating the date the calibration frame was created.

level from neighbouring well-behaved columns is small ($\mathcal{O}(0.1\text{ e}^-)$) – compared to the read-noise level of a single image – $\mathcal{O}(10\text{ e}^-)$. Of the marginally defective columns that are undetected in the step model fitting (right panel Fig. 3), their effect does not contribute significantly to the uncertainty budget for photometry. These subtly defective columns are apparent in Fig. 3 only owing to the use of a stack of 500 single frames, but are undetectable in any given exposure as the read-noise dominates.

When applying the `columntrap` correction during data reduction, a final check is made: there can be scenarios where the background level of a science frame is smaller than the step in the calibration frame. To avoid imprinting negative versions of strong traps during reduction of these frames, a limit is placed on the correction value to be at most the level of the local background (calculated in the fashion of Section 4.5.2).

We found that the evolution of the traps is non-negligible, albeit slowly evolving (Fig. 4). Monthly creation of `columntrap` calibration frames has proven a good compromise to capture this evolution for use during the low-latency image processing, without unnecessarily creating them too frequently.

4.4.4 Application of super calibration frames

During low-latency processing of a science frame, the appropriate `bias`, `dark`, `flat`, and `columntrap` frames are found – i.e. those created most recently to the observation date of the science frame,¹⁰ and which match camera–UT–filter combination, as required.

The corrections are applied in the typical CCD reduction manner (e.g. S. B. Howell 1992): subtraction of the `bias`, `dark`, and `columntrap` frames, before division of the normalized `flat`. Pixels identified as hot/dead/non-linear (Section 4.4.2) are propagated in a bit-wise fashion to generate an initial bad pixel mask for the now-reduced science frame. Columns corrected by the `columntrap` calibration frame are not identified in the mask since their pixels are sufficiently well-recovered and behaving to be useful in photometric measurements. At this point the reduced science frames are ready for data processing.

4.5 Creation and science processing of single images

A single image is a processed and analysed science frame produced by a single exposure of a GOTO UT. The creation and processing of CCD-reduced data (Section 4.4.4) is described in the following subsections.

Unless stated, where values are calculated and/or visualized in this and subsequent sections, they are taken from GOTO images taken in the time period 2025 July 28 to 2026 January 28.

4.5.1 Cosmic-ray removal and bad-pixel correction

Cosmic-rays originating from the high energy particle background pose as random hits on CCD detectors causing a localized spike in signal. Through the difference image process (Section 4.7.2), these can manifest as positive detections, unnecessarily burdening the process of identifying real astrophysical transients. The sharp profile of cosmic ray detections typically allows for their robust identification in oversampled ground-based astronomical imaging, with the most popular technique – convolution of the data with a Laplacian-like kernel (P. G. van Dokkum 2001) – able to also perform well in critically sampled data. Given the pixel size of GOTO and the optical performance, under good conditions at the observatories the PSF can enter the regime of critical sampling. The algorithm

¹⁰KADMILOS further enforces these calibration frames must be historic (pre-observation), in the case of reprocessing a science frame.

used is a direct implementation of the P. G. van Dokkum (2001) method, accessed via ASTRO-SCRAPPY (C. McCully et al. 2018). It is applied iteratively, identifying and masking cosmic ray hits each time until no new pixels are identified.

The cosmic ray pixel mask is combined with the camera's bad pixel mask (comprised of non-linear and hot/dead pixels identified during the production of calibration frames; Section 4.4.4). Since one cannot recover useful information from these pixels, it is viable to simply track these pixels via a bad pixel mask and remove their influence from any measurements or photometry. However, if they are left untreated they can cause further unwanted consequences. For example, when transforming data spatially (as needed for difference imaging), interpolation of sharp features on to a new pixel grid can produce aliasing effects, spreading their impact, and they can affect segmentation of the image at the source-detection stage. We thus perform a (largely aesthetic) interpolation over the bad pixels using a 2D Gaussian kernel with a small (~ 1 pixel) sigma value, replacing their value by the value of the kernel. The replacement values are much closer representations of the expected good pixel value, diminishing their adverse effect on further processing (cf. examples given above). Their positions are nevertheless traced via the bad pixel mask, still allowing for them to be accounted for during final photometry. The mask itself is an integer array, allowing for bit-wise flagging of pixels to trace their respective provenance.

4.5.2 Background removal

Scattered-light and sky-brightness act as noise contributions to the reduced science frame data, as well as elevating the mean background pixel count above zero (the idealized case).

A typical method to remove a potentially spatially varying background is to identify the local background on a coarse mesh of points and then interpolate between those points using some smoothly evolving functional-form such as a low-order spline. In practice, the determination of the local background at the mesh points, and the mesh-points themselves, are subject to sigma-clipping to remove outliers, which may be caused by real astrophysical objects. The background calculation is performed by the SEP package (K. Barbary 2016), itself using the algorithm of E. Bertin & S. Arnouts (1996).

A two-step process is employed. First, an initial background formed from a very coarse resolution mesh is used to identify pixels $> 2.5\sigma$ above the local background. The pixels are included, along with the frame's bad pixels (Section 4.5.1), in a mask for construction of a background based on a finer resolution mesh background. This final background allows for the subtraction of the spatially varying sky level from the pixel values, as well as calculation of the spatially varying RMS of the frame, for later use in uncertainty calculation during photometry. The final background mesh is constructed with points separated by ~ 64 pixels in each direction (exact mesh size determined by the final image size to ensure a roughly integer number of mesh spacings). Although large nebulae and galaxies will have their extended emission captured (and removed) by this background-resolution, we are motivated to reduce issues for image subtraction (Section 4.7) and ensure reliable point-source photometry, rather than preserve extended object morphology.

Due to slight differences in the gain and noise properties of the two readout channels that make up each GOTO camera CCD, this process is performed separately for each channel, before being combined to produce a final background object. The presence of bright stars on the dividing line between the two readout channels can produce unrepresentative discrete steps in the background. However, the locations around bright stars are typically compromised due to scattered light from the PSF of the star in any case. Furthermore, the issues due to bright stars on the intersection of the channels occur much less frequently than issues caused by trying to fit a single, smoothly varying spline interpolation of the background to the discrete change in noise and level properties of the image.

4.5.3 Source detection

Point-like sources are identified and measured in each image to enable further calibration, required for transient discovery. The detection routine is that of SEP (K. Barbary 2016), also used for background removal (Section 4.5.2), and makes use of the bad pixel mask, and the background RMS to create a de-blended segmentation map (E. Bertin & S. Arnouts 1996) of sources with at least 7 pixels above 2.5σ of the local background. Following this initial detection pass, cuts are made to remove sources with any of the following:

- (i) a not-finite value for one or more shape parameters.
- (ii) a flag from SEP indicating the photometry is compromised.
- (iii) a full-width half-maximum (FWHM) value $\lesssim 1$ pixel.
- (iv) a centre too close to the edge of the field of view (within $1 \times$ the average FWHM of the image).
- (v) an overall SNR < 5 .

Half-flux diameter (HFD) measures the diameter of a circle centred on the source that contains half the total flux. For a Gaussian profile, it is related to the FWHM by $\text{HFD} \simeq 0.6 \times \text{FWHM}$. HFDs are calculated for each source, again using SEP. The distribution of HFD values is sigma-clipped, before taking a median to obtain a representative HFD for the image. Two aperture radii, 'small' and 'large', are then defined as 1 and $3 \times \text{HFD}$, respectively. Although the smaller aperture is nominally larger than one which will maximize signal-to-noise (SNR) for a Gaussian profile,¹¹ it ameliorates centroid issues, particularly a concern under good conditions given the relatively

¹¹<https://web.ipac.caltech.edu/staff/fmasci/home/mystats/GaussApRadius.pdf>




Figure 5. Density distributions of source HFD, calculated from medians in individual GOTO single images. Individual UT distributions are shown (thin coloured lines), along with the overall distribution per mount (thick dashed line). Density estimates were determined using a Gaussian kernel with bandwidth using the method of D. Scott (1992). The distributions were initially filtered to remove the low number of images with median HFD values of > 15 pix. The more varied performance of GOTO-1 is reflective of a more aged collimation of that mount compared to the others.

coarse pixel-scale of GOTO compared to the delivered image quality. The PSF shape is also unable to be held constant over the wide field of view of the fast UTs, guaranteeing departures from a Gaussian profile for any given image.

Photometry of each source passing the cuts above is performed in each aperture size, before a final cut to remove any sources that have $\text{SNR} < 3$ in the small aperture. At this point a detection table is constructed from all sources that pass the above cuts, which is stored in a FITS extension upon storage of the image (Section 4.5.6), as well as in the `photometry.single` table of the KADMILOS data base.

The distribution of HFDs for all sources in useful images, split by mount and UT, are shown in Fig. 5. The 50th and 95th percentiles of the HFD distribution for most UTs across the array are $P_{50} \simeq 2.6$, and $P_{95} \simeq 4.2$ pixel.

4.5.4 Astrometry

Positions of high SNR (> 20) sources in an image's detection table are used to determine a World Coordinate System (WCS) solution for each image, via the `ASTROMETRY.NET` package (D. Lang et al. 2010), using their 5000-series index files, built from *Gaia* DR2 (Gaia Collaboration 2016, 2018). With the large, distorted images from GOTO UTs, the main `solve-field` routine of `ASTROMETRY.NET` is called multiple times per image to improve the overall solution since it was found by experimentation to perform poorly in the outer reaches of the field of view otherwise.

First, the central quarter area of the image is solved without using any higher order tweaks such as simple image polynomials (SIP; D. L. Shupe et al. 2005). Initial guesses for the pixel scale, based on the camera's configuration file (Section 4.2.2), and the central Right Ascension (RA) and Declination, from the telescope pointing information in the FITS header of the image, are used. This initial solution provides an estimate for the footprint of the image on sky, which is used in turn to query the local copy of the ATLAS-REFCAT2 catalogue (Section 4.3) to find calibration sources within the image, fetching, at most, 50 000 sources in descending brightness order. Cross-matching of image sources and calibration sources is performed using their RA and Declination coordinates (respectively, from the initial WCS solution of the image, and those stored in ATLAS-REFCAT2). Any not matched within a 3 arcsec radius, or which are ambiguously matched (i.e. not one-to-one), are discarded. A simple linear photometric calibration is then applied to search for outliers in a fitted relation between instrumental GOTO magnitudes and the ATLAS-REFCAT2 magnitudes in the filter used for later photometric calibration (Section 4.5.5). This step effectively removes spurious bright-faint matches, which occur more frequently when the depth of a GOTO image is not comparable to the faintest calibration sources, such that the fraction of sources appearing in both the image and the calibration catalogue becomes diminished.




Figure 6. Heatmap of astrometric residuals of high SNR GOTO source centroids from ATLAS-REFCAT2 catalogue *Gaia*-based positions. Quantities were calculated from 5000 randomly selected images and are shown with uncertainty based on the P_{16} and P_{84} values of the distribution. KADMILOS splits the images into 25 (5×5) sub-tiles for astrometric stats, but only calculates for central and extrema sub-tile in each axis – i.e. the coloured regions of this figure. For an estimate of the total GOTO astrometric residual of detections, an additional statistical component should be considered in the low SNR regime (see the text).

The image is then divided into four sub-tiles, by splitting the image in along each axis. The process above is then repeated for each sub-tile: performing a no-tweak WCS solution and cross-matching sources to ATLAS-REFCAT2. For each pass, the parent cross-matching table is updated with new cross-matches that pass the cuts for that sub-tile’s WCS solution within the given sub-tile’s footprint.

This parent cross-match table is then used to fit SIP corrections to the spatially varying astrometric residuals of positions between the image and catalogue positions. These take the form of two third-order polynomial surfaces describing the residuals in each image axis. The fitting occurs iteratively, a maximum of 10 times. Each iteration:

- (i) Computes a SIP correction to the no-tweak WCS based on the current cross-match table.
- (ii) Calculates the median astrometric residual of sources after including this SIP correction.
- (iii) Checks the solution is converging by ensuring there are more cross-matches and a smaller median residual compared to the previous iteration.
- (iv) Updates the cross-match table for the next iteration.

The checks within (iii) are early stopping criteria meaning the iterative fitting almost always terminates before the maximum 10 iterations.

With the final SIP-tweaked WCS solution in hand, a final, tighter cross-matching radius of 1.5 arcsec, is applied to create the final cross-matched table of sources. Statistics for the median and standard deviation of residuals are calculated after splitting images into 5×5 sub-tiles, which are then stored in the FITS header and data base. A visualization of these residuals is shown in Fig. 6, indicating 0.32–0.38 arcsec residuals, depending on source position – around a quarter to a third of a GOTO pixel. Collimation procedures try to maximize the area of ‘good’ PSF coverage, which of course prioritizes central areas of the image, the typically poorer PSF performance in the corners is evident from a slightly elevated astrometric uncertainty, but this is limited to a 0.05 arcsec change across a typical image. A more detailed view of the telescope-specific astrometric residuals can be seen in Fig. A2 – the effect of different collimation performance (a maintenance task that is typically performed per telescope) is evident, although median differences are small, at the level of ~ 0.1 arcsec.

It should be stressed that this astrometric performance is based on the statistics of high SNR sources. An additional source of positional uncertainty, and one particularly relevant for discovery of faint infant transients, is the centroid of the source itself. One may estimate centroid uncertainty for a given SNR, assuming an idealized Gaussian profile as:

$$\sigma_{\text{cen}} \simeq \frac{\text{FWHM}}{2\sqrt{2\ln 2} \cdot \text{SNR}}. \quad (3)$$

For a marginal $\text{SNR} = 5$ source in a typical $\text{FWHM} = 4.3$ pixel image (Section 4.5.3 and Fig. 5), $\sigma_{\text{cen}} \simeq 0.4$ pixel $\simeq 0.5$ arcsec. This positional uncertainty is representative of that due to the source itself, and should be considered in quadrature with the local image-level astrometric residuals to better describe the absolute sky-position uncertainty of marginal sources. We find no strong dependence




Figure 7. The number of ATLAS-REFCAT2 calibration sources in 1.4 deg cone searches at central locations of GOTO tile positions, shown as a function of Galactic latitude (black circle markers, left axis). The thin black dashed lines indicate the P_{5-95} range of these values. The right axis and differently coloured marker shapes show the fraction of those total sources that survive each photometric quality cut detailed in Section 4.5.5, as well as those that survived all cuts. The shaded region for ‘Passed all cuts’ source counts similarly indicate the P_{5-95} range of values. Ranges are omitted for individual cuts for visual clarity. Survival rate of sources drops dramatically at low Galactic latitude, primarily due to a lack of isolated sources, but the survival number (total sources $\times$ passed all cuts) is relatively stable.

of σ_{cen} with field density or extended structures due to the whole-sky availability of good-quality astrometric positions from *Gaia* and the robust background subtraction and point-source extraction in our GOTO data.

4.5.5 Photometric calibration

The workhorse filter of GOTO is a wide optical filter, L (400–700 nm; see D. Steeghs et al. 2022). Although each UT houses a filter wheel and an array of filters, L is currently used exclusively during transient discovery observations and so we only discuss its calibration here. L broadly covers the more standard g and r filters used in, e.g. Sloan Digital Sky Survey (M. Fukugita et al. 1996) and Panoramic Survey Telescope and Rapid Response System (Pan-STARRS; N. Kaiser et al. 2002). As such, photometric calibration is done using these comparison filters, including a colour term given the significant width of the GOTO filter.

The detection table of an image (arising from Section 4.5.3) is filtered for sources that have: an astrometric cross-match with ATLAS-REFCAT2 (Section 4.5.4), a high SNR (> 75), and that are not within 20 pixels of the image edge. The cross-matched catalogue photometry in ATLAS-REFCAT2 for each source is then used to further filter for those sources which have:

- (i) A $g-r$ colour within 0.0–1.3 mag.
- (ii) A brightness of $11.5 < m_g < 20.5$ mag.
- (iii) No nearby confusion sources, by requiring less than 10 per cent of their flux to be contained within 7.5 arcsec of their position based on *Gaia*.
- (iv) A comparatively stable magnitude determined from the χ^2/DOF value in ATLAS-REFCAT2, from Pan-STARRS epoch photometry (J. L. Tonry et al. 2018a).
- (v) A low but non-zero proper-motion value of < 0.1 arcsec per year. Zero proper-motion values may be indicative of extragalactic sources.

A population of cone searches across the sky, designed to mimic a typical GOTO UT field of view, were used to determine the survival of ATLAS-REFCAT2 catalogue sources in the face of these criteria. The results are shown in Fig. 7, plotted as a function of Galactic latitude, the principal component affecting the survival of sources. Most cuts are relatively stable across the sky and have a high fraction of survivors, with colour and brightness cuts only dipping in the central few degrees of the Galactic plane, primarily due to extinction effects. Isolation, as expected, drops quite dramatically in the crowded low latitude fields, but this is offset by an increase in available sources. For the regions of primary interest for transient searches, namely $|l| > 10$ deg, around 55–65 per cent of sources from the catalogue survive. In practice, the requirement for a high SNR detection in the GOTO image is typically the limiting factor.

In Fig. 8 we show the number of astrometric and photometric sources used in the calibration of images. A minimum of 100 sources




Figure 8. The number of sources used in the astrometric and subsequent photometric calibration of GOTO images. Solid lines indicate the median value (binned to daily frequency), with the shaded regions showing the P_{5-95} ranges. Also plotted is the moon illumination (thin dark line, right axis), correlating with small dips in source numbers around full illumination due to an imparted lower sensitivity of GOTO images. The stable upper bound on the number of astrometric matches is a consequence of the 50 000 maximum sources fetched from ATLAS-REFCAT2 per footprint.

are required to be used in the calibration steps, in order to robustly overconstrain a spatially varying calibration.

In order to capture the spatially varying nature of the calibration across a UT field of view, a surface, $z = f(x, y)$ must be defined, where x and y are pixel coordinates of the image, and z is some spatially varying quantity of the image. This surface must be constructed from a sequence of irregularly sampled positions – i.e. positions of stars – rather than an idealized dense, regular mesh of points. Typical approaches use flavours of low-order polynomial or spline for f (as indeed is the case for astrometric residuals, e.g. D. L. Shupe et al. 2005). Options such as splines (G. Wahba 1990) offer advantages through providing a smoothed representation of these intrinsically noisy and irregularly spaced tie-points. However, splines are limited in their capability to model sharply diverging areas of the surface, or very localized phenomena (without consequently overfitting more smoothly behaving regions). A solution was found in the framework of radial basis functions (RBFs), specifically using the kernel of a thin-plate (polyharmonic) spline (e.g. A. Iske 2003) – an example application of RBFs for astronomical data calibration can be found in M. A. Mitchell et al. (2026). The surface, z , is produced through a linear combination of RBFs, each produced from the measurement and position of a star that passed the above filtering. To ease computational cost, a thinning method is applied to these stars which subsamples those in dense regions, whilst retaining isolated ones, thereby maintaining information on the spatial behaviour with fewer measurements. Practically, the number of neighbours (N) within a certain pixel radius (r) is calculated for each star. The probability of retaining a star P_{ret} is then:

$$P_{\text{ret}} = \begin{cases} 1, & \text{if } N < N_{\text{lim}} \\ N_{\text{lim}}/N, & \text{otherwise,} \end{cases} \quad (4)$$

where N_{lim} is a limit on the number of close neighbours a star can have before subsampling is enacted. In KADMILOS, $N_{\text{lim}} = 10$ and $r = 180$ pixels are used.

An RBF surface is constructed to describe the aperture correction of the image (i.e. the difference in magnitudes between the small and large aperture photometry (Section 4.5.3), using all detections that pass the above criteria. This surface is then applied to the small aperture photometry of all stars to get their large-aperture-corrected magnitudes. A visualization of the detection filtering and construction of this aperture correction surface is shown in Fig. 9. The shape of this surface highlights the collimation challenge with such fast aperture telescope design that utilizes almost the entire unvignetted field of view. The area of consistent performance is maximized during the collimation process, resulting in a ring of best performance with a very slight degradation at the centre, and the unavoidable degradation in the corners.

A second-order polynomial is then fitted to the colours $g - r$ and $g - L_{\text{inst}}$ using all filtered sources in the image, where L_{inst} is the instrumental GOTO L -band magnitude, aperture-corrected using the RBF surface calculated above. The relation is used to calculate residuals of stars – the difference in magnitudes between what is measured and that predicted from the polynomial relation. These residuals are used as tie points for another RBF surface calculated analogously to the aperture correction, which defines a delta zeropoint surface, allowing the calibration to account for differential atmospheric transparency, for example, across the image.

Results of the calibration in the form of zeropoints for GOTO-1 UTs are shown in Fig. 10, with equivalent plots for remaining mounts given in Fig. A1. Most UTs had shown overall little degradation compared to the initial installation and design sensitivity (M. J. Dyer 2020). However, GOTO-1 UT2 degraded noticeably faster than other UTs (which are all relatively consistent). A more rapid build up of dust on the secondary mirror of this UT, for unknown reasons, was found. A maintenance trip in September 2025 managed to recover close to design sensitivity of this UT following an intensive cleaning operation. Subsequent maintenance trips to La Palma (GOTO-1 and GOTO-2) and Siding Spring (GOTO-3 and GOTO-4) in late October and late November 2025, respectively, are evident by increases in the zeropoint values of UTs (Figs 10 and A1). A more detailed discussion of the zeropoint behaviour as related to the telescope hardware will appear in Steeghs et al. (in preparation). Short spikes of lowered sensitivity related to poorer ambient conditions, with the overall sky-brightness effect of the Moon evident in Figs 10 and A1 given the blue-optical wavelength coverage of the L -band.




Figure 9. A heat map of the aperture correction radial basis function (RBF) surface for a representative GOTO-1 UT-1 image. Overlaid on this heat map are scatter points of detections used to construct the surface. They are colour coded on the same scale and outlined for visual clarity. A zoom-in region is indicated, showing the underlying image data with coloured markers surrounding high SNR detections that fail the various criteria (i.e. low-SNR sources are already excluded and not marked in the figure). The criteria are described further in Fig. 7 and Section 4.5.5. Sources may be marked multiple times, indicating they failed multiple criteria. No detections failed the colour cut in this zoom-in region. Further, those culled due to the probabilistic thinning technique (see the text) are overlaid with crosses. Only detections passing all cuts and not thinned are used in the construction of the RBF surface. A comparison of the zoom-in detections with the outlying scatter points in the aperture correction heat map shows they are due to stars with close neighbours. Although the $rp > 7.5$ arcsec criterion limits the impact of egregiously crowded sources, neighbours can still affect aperture photometry of those making it through the cut. Nevertheless, the heat map shows that these individual outliers do not influence the surface, which maintains a good description of the overall image quality, and demonstrates the robustness of using RBFs in this domain.




Figure 10. Results of photometric calibration of 45 s single GOTO-1 images, split by UT (colours). Shown on the plot are daily bins of the zero-point (thick solid lines) and median detection magnitude of 5σ sources in the images (thick dotted lines). Overlaid is the moon illumination (thin dashed line), causing a reduction in image depths around full moon. Periods of no data are either due to bad weather, technical downtime, or nights when no 45 s exposures were taken (e.g. during intensive external trigger follow-up, which employs a different observing strategy.) The effect of maintenance trips to clean the primary and secondary mirrors can be seen as step changes in the zero-point performance – late August 2025 for UT2 and mid-October 2025 for the remaining UTs.




Figure 11. 2D histograms of GOTO set image depths as a function of moon illumination in the L -band. The two main observing strategies are indicated as titles above each subplot. Note that the colour scaling of bins is logarithmic.

We emphasize that the performance shown in these figures relates to 45 s single images of GOTO. In practice, both the sky-survey and follow-up modes employ stacking of single images into set images to improve sensitivity (Section 4.6 and Fig. 11).

4.5.6 Storage format of single images

After the processing of previous subsections, the single image data are normalized to one second by dividing the pixel values by the exposure time. single images are then stored as a multi-extension FITS file with the following extensions:

Extension	Description
science	Reduced, cleaned, background-subtracted, normalized single image data.
science_mask	Bit-wise integer mask of bad or otherwise affected pixels.
science_bkg_rms	Spatially varying RMS map determined from background fitting.
science_phot	Source position and photometry table, including cross-matched catalogue data.

The location of the FITS file, as well as statistics on the image quality (HFD, number of detections), astrometric and photometric fits is stored as a record in the appropriate PostgreSQL table. This record holds identifiers of the calibration frames (bias etc.) used in the processing of the single image. The two RBF objects generated for each image during photometric calibration (Section 4.5.5) are additionally serialized to bytes using DILL¹² and stored directly within the PostgreSQL record for the image. The (de)serializing of these objects is a quick process, and provides a space efficient means to store them. This allows the surfaces to quickly be loaded and used during subsequent analysis (e.g. forced photometry, Section 5.5).

¹²<https://dill.readthedocs.io/en/latest/>

4.6 Creation and science processing of set images

As detailed in Section 2.1, each GOTO pointing consists of one or more `single` images meant to be combined into a stacked `set` image. The addition of read-noise (cf. a single deep exposure) is offset by mitigation of cosmic ray artefacts, and the rejection of short-lived poor conditions or technical issues.

The `single` images to be combined into `set` images are tracked via unique combinations of UT, mount, and ‘setnum’ values – incrementing integer designating each set. Once all `single` images from a given set have been processed, the workflow to combine them is initiated.

4.6.1 Determination of relative image weights

The combination of `single` images is done in a weighted fashion, closely following the formalism of J. Annis et al. (2014) to generate relative weights for each image. Namely:

$$w_i = \frac{T_i}{\text{FWHM}_i^2}, \quad (5)$$

where w_i is the weight of the i th image, and T_i and FWHM_i^2 are relative values of the transparency (based on zeropoint) and FWHM of the image, with respect to the ‘best’ image in the set. The best image in the set then naturally has $w = 1$. Any image with $w < 0.6$ is excluded from the combination process as it was found these are typically compromised images, due to various issues such as mount tracking errors, scattered light from artificial sources, rapid degradation of ambient conditions, or similar anomalous phenomena. In the final `set` images they contribute largely noise, having minimal, sometimes negative, influence on the quality of the stacked image, even in spite of their lower weight.

4.6.2 Astrometric registration

Between the individual exposures of a pointing, the GOTO mounts dither their position by several arcsec, which can depend on the number of sub-exposures although a 4-point pattern is typical for both sky-survey and follow-up observations discussed here. This approach shifts bad pixels spatially, such that stacked pixels are almost always composed of mostly good pixels. The exception to this can be large CCD defects or defects in the optics (e.g. mirror blemishes). The dithering step introduces the requirement that the processed `single` images in a set are astrometrically aligned prior to combination. (Such alignment would most likely be required even in the absence of dithering due to, e.g. imperfect tracking, shifting airmass).

SPALIPY (J. D. Lyman 2021) is an astronomical-image alignment tool utilizing source detection centroids to calculate a differential transformation rather than alignment to an absolute frame such as a WCS (see Appendix B for more details). To transform the small shifts between `single` images of a set, an affine transformation ignoring shear terms proves sufficient. This affine transformation is determined using several hundred to several thousand source positions. The first `single` in the set is used to define the pixel coordinate system, and others are aligned to this, including their background RMS and bad pixel maps. Any failed alignments are removed from the list of `single` images to combine.

4.6.3 Weighted combination

Alongside the overall relative image weights (Section 4.6.1), the combination of the `single` images employs a further pixel-wise weighting. This weighting is the inverse variance of the background, calculated from the background RMS map stored within the `single` FITS file (Section 4.5.6). Pixel-wise combination is done using a weighted mean, with sigma-outlier clipping. Due to memory demands of holding 10–12 GOTO images (in the case of deep observations) in memory, plus their background maps, per `set` image generation, this stage is performed in sections. The stack is split into 100 chunks along the longer image axis, and the data are held in memory-mapped files on fast, local discs via NUMPY (C. R. Harris et al. 2020). The combined data array generates the `set` image.

4.6.4 Science processing

The science processing of `set` images follows that of the `single` images, as detailed in Sections 4.5.3 to 4.5.5. Although the astrometric calibration may be largely redundant, the benefit of additional depth and suppressed noise in the `set` image means a recalibration ensures the stored astrometric solution most accurately reflects the final `set` image.

The image depths of `set` images taken in the regular sky-survey mode (4×45 s) and triggered follow-up of external alerts (4×90 s) are shown in Fig. 11. Sky-survey observations deliver a typical 5σ depth of $m_L \sim 20$ mag, reducing in sensitivity by ~ 1 mag in the brightest lunar conditions ($\gtrsim 85$ per cent illumination). Follow-up of external triggers delivers $m_L \sim 20.5$ mag, approaching ~ 21 mag in better conditions (and of course likewise suffering from the same lunar impact as the sky-survey).

4.6.5 Storage format of *set* images

Although *set* images are stored following processing, they are typically not interacted with by end users, and are picked up immediately by the next stage of the pipeline, where they are written with the full suite of data and table extensions. The reader is directed to Section 4.7.3 for information on this format.

4.7 Difference image analysis

Difference image analysis (DIA) is a common technique in transient sky surveys, whereby a historical, high-quality image of the sky (ideally taken with the same telescope and optical set-up) is used to subtract constant-flux sources, leaving behind only transient, variable, and moving sources (A. C. Phillips & L. E. Davis 1995). These historical images are referred to as *template* images, with the result being a *difference* image (Table 1).

The observation strategy for *template* images is provided in Section 2.1. The noise of the *set* and *template* are, to first order, added in quadrature during DIA to give the overall noise level of the *difference* image. Using a much deeper *template* image ensures only a small noise contribution to the *difference* image, which consequently retains a depth comparable to the original *set* image.

A process of manual vetting of *templates* was done using consortium members via a private Zooniverse¹³ platform. After each vetting they were incrementally ingested into a specific PostgreSQL table, where they could be searched for by the DIA stage of the low-latency processing DAG *single_and_set_creation*.

4.7.1 Creation of *difference* images

A *template* search is initially made to find one matching the mount, UT, filter, and *tile* of an incoming *set* image. However, it was found that, visually, the impact from using a *template* image acquired on one mount at a given site to perform DIA of a *set* image acquired on the other mount was not significantly degrading the quality of the subtraction. Rather, issues with DIA mainly arise due to data quality issues, including poor conditions. Therefore, to streamline the generation of *template* images, only one *set* of valid *template* images were obtained per site (rather than per mount). As such, if a *template* is not initially found, the requirement for having the same mount is then relaxed (although originating from the same site remains enforced). A final check in the search for a valid *template* is to ensure the distance between the centre of the *set* and *template* images is ≤ 0.35 degrees. Due to distortions in GOTO images in the outer regions of the FoV, severely mismatched images on-sky prove difficult to accurately align. Due to the similar footprint across mounts, and the consistent *tile* grid used, this limit is violated only infrequently.

The on-sky overlap between the *set* and *template* is calculated using their respective astrometric footprints in order to trim the source detection catalogues for each image of those falling outside the overlap. These overlapping sources are then used to align the *template* to the pixel-coordinate system of the *set* image, including its background RMS map. As for the registration of the constituent *single* images (Section 4.6.2), SPALIPY is used to calculate and perform the alignment. In this scenario, however, complex transformations can be required, and so the full transformation capability of SPALIPY is used, to allow (third order) spline corrections to additionally warp the affine transformation to account for non-linear distortion of the images.

With the *template* data and background RMS map registered to the *set* image, the image differencing is performed using a multithreaded version of the High Order Transform of Psf AND Template Subtraction code (HOTPANTS; A. Becker 2015).¹⁴ The algorithm of HOTPANTS is itself building upon the work of C. Alard & R. H. Lupton (1998) and C. Alard (2000), through the use of a spatially varying regulated kernel which, when convolved with the *template* image, will generate a new image with a PSF and scaling matching (as closely as possible) that of the *set* image. Prior to running the subtraction, the masks of the *template* and *set* images are logical-or combined, to build a mask for the *difference* image. Owing to sub-pixel shifts in alignment, the *template* mask would nominally suffer from interpolation artefacts and incorrect propagation of its integer mask. To circumvent this, we apply a ‘nearest’ interpolation technique, which preserves the values of the initial mask, at the small expense of missing some partial pixel overlaps.

4.7.2 Analysis of *difference* images

The analysis of the *difference* image data begins with source detection, as is done for both *single* and *set* images (Section 4.5.3). Unlike in those latter cases, additional checks are here put in place to test the validity of detections owing to the potential for DIA artefacts and correlated noise spikes producing spurious detections in *difference* images generated via convolution (C. Alard & R. H. Lupton 1998; C. Alard 2000). First, forced photometry is applied to the original *set* image at the position of the *difference* candidate detections. Any candidate with a SNR < 5 in the *set* image is discarded. A further check is made for those detections found in DIA of *set* images that themselves are composed of a low number (< 5) of *single* images (this is the usual case). Although the combination

¹³<https://zooniverse.org>

¹⁴<https://github.com/Lyalpha/hotpants>

stage (Section 4.6.3) employs pixel-wise sigma-clipping, artefacts present in just one or two images (e.g. cosmic rays, satellite trails) can still trickle through to the stacked pixel data, causing a spurious ‘detection’ in the `difference` data. As such, the constituent `single` images also have forced photometry performed at the locations of difference candidate detections. For each source, the fraction of `single` images in which there is a $\text{SNR} > 2.5$ excess of flux, f_{single} , at the position of the candidate is recorded.

As the DIA stage aligns and scales the `template` image to match the `set` image, we can use the photometric calibration of the `set` image as a valid solution for the subsequent `difference` image. Due to aforementioned DIA artefacts, source extraction on subtracted data typically finds many more astrophysically bogus sources of flux than genuine, real detections. Due to the volume of detections, it is unfeasible to expect manual inspection of even a sizeable fraction. For this reason schemes to separate ‘real’ and ‘bogus’ detections have been developed, typically appealing to machine-learning models to provide candidate detections with a ‘realbogus’ score ranging from zero to one. These model implementations principally rely on high-level features extracted from the detection (e.g. J. S. Bloom et al. 2012), and/or the pixel data directly (e.g. D. A. Duev et al. 2019). The model used by GOTO is fully described in T. L. Killestein et al. (2021) and consists of a convolutional neural-network trained on GOTO pixel data using algorithmically constructed training data, supplemented by synthetic-transient insertions. The trained model is loaded by KADMILOS and used for inference on each candidate detection found in the `difference` image. The previously calculated f_{single} is now used as a modulator of the native realbogus score provided by the model. Strong, spreading cosmic ray hits in an individual `single` image, which are not fully cleaned (Section 4.5.1), can particularly mimic faint sources in the ultimate `difference` image. These are particularly tricky for a model to distinguish, but show up clearly as low f_{single} values. This modulation of the model score means one-off artefacts now only propagate to the later vetting stages when they unfortunately align with a pre-existing astrophysical object, however visualization of the `single` data during vetting makes their removal trivial. More persistent artefacts – namely residuals around bright sources – remain a challenge to fully clean at the model score stage, and so are the chief contaminant in later stages.

4.7.3 Storage format of difference images

Data from the `set`, `template`, and `difference` images used during DIA are collated and stored in a single multi-extension FITS file with the following extensions:

Extension	Description
<code>science</code>	Reduced, cleaned, background-subtracted, normalized <code>set</code> image data.
<code>science_mask</code>	Bit-wise integer mask of bad or otherwise affected pixels.
<code>science_background_rms</code>	Spatially varying RMS map determined from background fitting.
<code>science_photometry</code>	Source position and photometry table, including cross-matched catalogue data.
<code>template</code>	Aligned <code>template</code> image data.
<code>template_mask</code>	Bit-wise integer mask associated with the <code>template</code> image.
<code>template_background_rms</code>	RMS background map for the <code>template</code> image.
<code>difference</code>	<code>difference</code> image (<code>science</code> - <code>template</code>) data.
<code>difference_mask</code>	Bit-wise integer mask for the <code>difference</code> image.
<code>difference_background_rms</code>	RMS background map for the <code>difference</code> image.
<code>difference_photometry</code>	Photometry table of candidates identified in the <code>difference</code> image.

As there is no further processing required of the data, these FITS files are stored using lossy compression of the float data in image extensions to save on storage space (W. D. Pence, R. L. White & R. Seaman 2010). The `fpack`¹⁵ tiled-compression package is used (W. D. Pence, R. Seaman & R. L. White 2009) with the RICE compression algorithm. The quantization for compression is chosen sufficiently high (512) to have an unnoticeable effect on the precision of photometry required ($\ll 0.01$ mag) during immediate future use such as forced photometry.

In addition, a record is stored in the KADMILOS data base detailing the path to this FITS file, as well as high-level quantities related to the image such as number of difference image detections. Records are linked to the `set` and `template` tables, allowing one to determine the provenance of the image, back to the raw and calibration data used.

4.8 Pipeline timings

The Airflow metadata data base provides a means to interrogate the performance of processing by KADMILOS. Of principal interest is the time taken to complete the *single_and_set_creation* DAG (Section 4.2.1) – the workflow responsible for creating and analysing the various science image products. This DAG always produces `single` images, but may also produce `set` and `difference` images if the exposure being processed is the last of a pointing. The distribution of process timings for the latter case is shown in Fig. 12. The distribution is split on N_{det} , the number of source detections found in the `single` image being processed. This quantity is

¹⁵<https://heasarc.gsfc.nasa.gov/docs/software/fitsio/fpack/fpack.html>




Figure 12. Total times for GOTO images to be created and analysed by KADMILOS, based on DAG run timings in Airflow. Dashed lines indicate time taken for image processing only – i.e. the creation and analysis of a sequence of `single`, `set`, and `difference` images. Solid lines measure this processing time from shutter close of the telescope for the image, thereby also including raw data readout, transfer, and ingestion stages. The overall distribution is shown in thick black lines. These are also split by the number of source detections (N_{det}) in the image, since this is a principal variable affecting processing time. Note the plot has been cut at 25 min and a small fraction of timings extend significantly beyond this (see the text).

chiefly responsible for the spread of timings in the distribution owing to many per-source calculations (e.g. shape, photometry) being performed, whereas image-level calculations (e.g. CCD reduction, alignment) remain constant. A small but extended tail is present in the distributions, reaching to many days or weeks. The culprits for such delays are typically issues with the original processing (e.g. a new camera was installed but did not yet have a valid `superflat` calibration image), hardware/network failures (e.g. data base cluster being inaccessible), and delays/errors due to software mishaps. Fig. 12 also shows the same distributions but counting the start time as time of shutter close for the telescope, better representing the entire process from writing the initial FITS file of the raw data at the respective site and ending with transient candidate detections stored in a PostgreSQL table. It is important to note that these distributions do not include those images which failed processing for any reason.

Although the results in Fig. 12 indicate images with $N_{\text{det}} \geq 100\,000$ take significantly longer, these are not typical cases (cf. the ‘all’ distribution), and, further, are of the least interest for extragalactic transients searches since they are crowded, extinction-heavy Galactic plane pointings.

4.9 Pipeline data volumes

As mentioned in Section 3, KADMILOS runs in low-latency on locally held data in Warwick once the data are transferred from the respective sites (Section 4.1). We show in Fig. 13 the nightly counts of `single` and their composite `set` images needing to be generated based on *observed* raw data frames taken at each site. Image quality issues mean some frames may fail some processing stages, and so the final data products produced may be slightly below these maximal counts. Shown alongside counts are approximate data volumes assuming a typical 225 and 580 MB file-size on disc for `single` and differenced `set` images, respectively. The median (95th percentile) disc storage volume required is ~ 2.7 (4.9) TB per 24 h for the stored FITS data products. Owing to the availability of filesystem block-level compression on the storage devices, which is able to further compress non-data FITS extensions, this is a small ($\sim 5 - 10$ per cent) overestimate. The median (95th percentile) of raw frames observed per 24 h of ~ 7700 (13800) translates to an input data rate of ~ 0.3 (0.6) TB.

5 MARSHALL IMPLEMENTATION

A common means for visualizing data from transient sky-surveys is the concept of a ‘marshall’ system (e.g. N. M. Law et al. 2009; K. W. Smith et al. 2020), also referred to as a target and observation manager (TOM; e.g. R. A. Street et al. 2018). This is typically a web-interface that allows collaboration members access visualizations of survey outputs, and, when occasion demands, initiate further actions. Tools for target identification and prioritization enable the reporting of findings promptly to the community and/or actioning of follow-up observations for science characterization.

The role of a marshall is typically also to perform several other important tasks, including source association (linking new detections with historical ones at the same location) and generating a wide range of high-level metadata for the source and its detections. This metadata can include contextual information from cross-matching with astronomical catalogues, calculation of macroscopic light-curve properties, and application of algorithmic or machine-learned science classification (e.g. D. Muthukrishna et al. 2019; S. Gomez et al. 2020; X. Sheng et al. 2024) or prioritization (e.g. H. F. Stevance et al. 2025) models. Initially a per-survey or per-facility component to the




Figure 13. The nightly observed *single* (top) and *set* (bottom) image counts and approximate data volumes across both GOTO sites and over a representative date range. The sites are almost antipodal, so the data processing occurs sequentially for each site, meaning stacked counts are effectively 24 h windows of data processing. Horizontal bars on each panel show the 5th, 50th (median), and 95th percentile of the totals.

transient science workflow, recent implementations concentrate on aggregating source management to and from various surveys and follow-up facilities in one common interface (e.g. M. M. Kasliwal et al. 2019).

An important further component is the development of ‘broker’ systems (e.g. J. Nordin et al. 2019; A. Möller et al. 2020; T. Matheson et al. 2021; F. Förster et al. 2021; R. D. Williams et al. 2024), which take the heavy lifting of parsing data streams from multiple surveys and providing a consistent API to access and process the streams. Alongside built-in high-level classification and data-tagging tasks (e.g. P. Sánchez-Sáez et al. 2021; R. Carrasco-Davis et al. 2021), the brokers provide entry points for contributions by the community to develop tools that add highly specialized metadata or scores to candidates as they are processed (e.g. X. Sheng et al. 2024). These can help de-duplicate effort in self-contained marshall systems, although typically the models are specialized and not able to be effortlessly applied to different survey data – significant additional training of models is required at the least.

Science cases may demand some combination or subset of these architectures. For example, a spectroscopic follow-up collaboration may deploy a marshall system to ingest from several surveys – typically via a broker of choice – and coordinate their follow-up campaigns (e.g. S. J. Smartt et al. 2015). In contrast, survey teams themselves may have privileged additional information above that released via a forced photometry service or broker alert stream, and wish to conduct an in-house marshall in addition, in order to further facilitate their own science goals.

In the following subsections we describe the marshall developed for GOTO, from ingestion of new difference image candidate detections, to reporting of transients and triggering of external follow-up.

5.1 Software overview

The GOTO marshall is built using the DJANGO web-framework.¹⁶ Asynchronous task handling is done through the use of a dedicated Celery worker, similar to our Airflow implementation (Section 4.2). The back and front-ends of the framework interact with a dedicated PostgreSQL data base to retrieve and update all source information.

5.2 Ingestion of candidate detections

The processed data base of KADMILOS, specifically the `difference` image records (Section 4.7.3), is regularly polled by the GOTO marshall. Any image records found that are not already held by the marshall data base are marked for processing. The original reason for observing the image is held by the KADMILOS data base, and this dictates the thresholds for pulling detections associated with the image. Regular sky-survey observations require a `realbogus` score of > 0.7 for detections to be ingested into the marshall, whereas observations taken due to an `event` (i.e. GW, GRB, or neutrino follow-up) have this threshold relaxed to > 0.4 , since completeness over purity is desired for these cases. Records for each difference detection passing the cut for the new image are similarly fetched and marked for ingestion.

Ingestion of a new image causes a signal to add a task to the Celery queue to generate a JPEG preview of the whole image. Ingestion of new detections within the image follows a series of checks and tasks that are inserted into the Celery queue for processing in parallel or sequence, as required. A schematic overview of this process is shown in Fig. 14, and described here.

(i) Check existing detections from the same image to ensure none with the same RA, and Declination coordinates is already held by the marshall data base. Without this check, other methods of fetching photometry (see Section 5.5) could conceivably ingest duplicates.

(ii) Perform a ‘get or create’ operation for a source, to either associate the detection with an existing source, or create a new one. For this, a row-level lock is acquired on a table of HEALPIX (K. M. Górski et al. 2005) indices covering those at and surrounding the detection coordinates, which prevents race conditions for source association or creation. The following are tried in order, ending early on success and then releasing the table lock:

- (a) Return an exact match for the source at the exact RA, Declination coordinates.
- (b) Return the nearest existing source in a 4 arcsec cone-search around the coordinates.
- (c) Create a new source at the detection coordinates

(iii) Generate JPEG cutouts around the detection from the `difference` image FITS file, producing cutouts for the `set`, `template`, and `difference` data.

(iv) Query the KADMILOS data base to find the constituent `single` images used to create the `set` image and similarly generate cutouts of those.

(v) If a new source was created during ingestion of this detection:

- (a) Check if spatial-temporal coordinates are coincident with a known moving Solar System body (Section 5.3.3).
- (b) Determine source context using MORIARTY (T. Killestein 2023, Killestein et al., in preparation), which queries a range of catalogues to determine source disposition.
- (c) Fetch stamps at the source position using the cutout services of optical sky-surveys SDSS DR12 (S. Alam et al. 2015), Pan-STARRS DR2 (H. Flewelling 2018), SkyMapper DR4 (C. A. Onken et al. 2024), and Legacy Survey DR10 (A. Dey et al. 2019).
- (d) Perform automatic vetting procedures (Section 5.3).
- (e) Check if the source meets the criteria for automated follow-up (Section 5.6.1).

(vi) Update or create the static light-curve plot of the source with the new detection’s photometry.

After the above ingestion and processing steps, sources are then ready to be presented to collaboration members within the marshall, and are assessed and interacted with as deemed necessary, e.g. manual vetting and reporting (Section 5.3.4) or triggering of follow-up (Section 5.6).

5.3 Vetting processes

The marshall is chiefly a platform to vet candidate detections into respective categories or source lists (Section 5.3.1). A stream of detections filtered on `realbogus` score from a model will still be impure, containing ambiguous detections or artefacts, as well as astrophysically real sources unrelated to transients. Human sifting therefore remains the final step before community reporting of transients, although several automated processes and additional filtering ease the task.

¹⁶www.djangoproject.com




Figure 14. Schematic visualization of the processing performed upon ingestion of a new difference image source into the GOTO marshal. Fuller details on these processes are in Sections 5.2, 5.3, 5.5, and 5.6. The summing junction ($\otimes$) must have all input paths followed by a source for the flow to continue (i.e. forced photometry is only performed if a source (1) was autovetted as `Inbox` due to high SNR or multiple detections, (2) has the required context for follow-up and (3) is not matched with a minor planet or a planet's Hill Sphere. Colours are arbitrary and simply highlight the start and important end stages of the flow.).

5.3.1 Source lists

Organization of sources is primarily on source lists within the marshal: broad categorical tagging to focus human effort. The lists and short descriptions are given in Table 2.

All new ingestions enter the marshal as a `Pending` source list. The rate of entries into `Pending` is typically too numerous to be sifted manually. Sources are therefore automatically vetted based on several criteria.

5.3.2 Automatic vetting

The automated vetting decisions are shown as part of Fig. 14, and they rely on either external contextual information for the source, or the nature of its detections by GOTO.

Table 2. Description of source lists used in the GOTO marshall vetting process.

List	Description
Pending	Ambiguous or unclear sources. All new sources created initially begin here.
Inbox	Sources deemed to require manual, human vetting. This is the primary list monitored by users for transient identification.
Stream	Transient sources such as supernovae, afterglows and related objects.
Store	Persistently varying or out-bursting sources, comprising mostly Galactic sources as well as Active Galactic Nuclei.
Junk	Ephemeral detections such as image artefacts, satellite trails, and moving objects. New detections nearby these positions will not be matched to existing sources in this list, meaning future detections are treated as independent new sources.
Banish	Detections in permanently compromised sky locations, such as the wings of bright stars. New detections at these positions will be matched to existing sources in this list to prevent the need for repeated vetting.

Any source matched to a known Solar System minor body is vetted as *Junk* (see Section 5.3.3 for more details). Sources matched to a stellar source in either of the *Gaia* (Gaia Collaboration 2016) or AAVSO International Variable Star Index (C. L. Watson et al. 2006) catalogues are vetted as *Store*. A source with either a match in the Transient Name Server¹⁷ (TNS) or Living *Swift* XRT Point Source Catalogue¹⁸ (LSXPS; P. A. Evans et al. 2023) are vetted to the *Inbox* for human inspection and reporting, as necessary. Should a source lie within 1 arcmin of a galaxy catalogued by either the Galaxy List for the Advanced Detector Era+ (GLADE+; G. Dálya et al. 2022) or NASA Extragalactic Database Local Volume Sample (NED LVS; D. O. Cook et al. 2023) it is also promoted to the *inbox*.

Beyond contextual classification, a source may be automatically vetted based on the results of its detection(s). If multiple difference image detections are associated to the source, or it has a single detection with $\text{SNR} > 15$, it is sent to the *Inbox*.

5.3.3 Identification of moving-bodies

Frequent visitors in time-domain survey difference images are moving objects. Since they are spatially displaced between the epochs of template and science observation, they manifest as astrophysically real sources of flux in difference images and can confound transient searches. Although they will appear at any given position only once, and so can be effectively mitigated by awaiting another epoch of observation for the same patch of sky, in the era of rapid transient identification and characterization, it is desirable to identify objects in single epochs of detection.

High-proper motion stars are one such moving-body, although the number of such stars with angular velocities sufficient to displace more than a typical GOTO PSF width in the maximal few year epoch difference between science and template observations is small. Additionally, these objects are typically nearby, bright stars making such sources clear to identify by human eyeballers. Contextual information providers can also accommodate such objects trivially in labelling schemes given the exquisite proper motion constraints available from *Gaia* for sources down to the typical depth of GOTO. Given their low incidence rate, manual eyeballing of these sources is relied upon for vetting when they appear in the *Inbox* currently, although automated contextual labelling is being implemented.

Artificial satellites also pose as interlopers, with *JWST* a typical culprit. Such objects are typically bright, $L < 17$ mag, and unassociated with any source on sky (stars or galaxies). As the sky density of such sources is low, and they are distinctive in their brightness and environment, they can effectively be managed by manual cross-checks of satellite positions for suspicious detections.

Solar-system bodies are numerous, and move on angular speeds ranging from marginal displacement between successive epochs in GOTO to significantly streaking within a single exposure. In the latter case, these are not a major contaminant for human vetters, since the exposure stacking minimizes the presence of streaks in the combined *set* images, and the realbogus classifier GOTORB (T. L. Killestein et al. 2021) is further adept at ensuring they are not promoted as an object of interest for transient searches. In the parameter range of speeds where the bodies appear point-like within $\sim$ minute-long exposures, but move significantly between successive epochs (hours–days), there is the most difficulty in distinguishing them from genuine transients. For these cases, we rely on the ephemeris service maintained by the Minor Planet Center¹⁹ (MPC). The MPC provides orbital parameters and optical magnitude estimates for $\sim 10^5$ Solar-system minor-bodies. Given the high rate of candidates requiring checking per night – $O(10^{3.5})$ – we perform the calculation of ephemerides locally. The package PYMPC (J. D. Lyman 2025) was developed to perform this task, and its implementation is detailed in Appendix C. The output of PYMPC is used to perform the following actions on new sources based on a minor or major Solar-system spatial match:

- (i) Matches of minor bodies brighter than GOTO’s detection limit within 7.5 arcsec of the source’s position are automatically vetted as *Junk* and not promoted to users.

¹⁷<https://www.wis-tns.org/>

¹⁸<https://www.swift.ac.uk/LSXPS/>

¹⁹<https://www.minorplanetcenter.net/iau/mpc.html>

Table 3. Breakdown of automatic vetting decisions by the marshall (Section 5.3.2), and the final destination of sources sent to the Inbox – i.e. those which have been further checked by collaboration members (Section 5.3.4). Numbers are rounded, and given for an average 24 h period from sources ingested over the time-frame 2025 June 1 to July 10.

Origin	Destination	N per 24 h	Percentage
Auto vetting			
Pending	Inbox	123	1.2 per cent
Pending	Store	1610	16.3 per cent
Pending	Junk	8158	82.5 per cent
Final vetting			
Inbox	Stream	17	14.2 per cent
Inbox	Store	23	18.4 per cent
Inbox	Pending	17	13.5 per cent
Inbox	Banish	2	1.5 per cent
Inbox	Junk	64	52.4 per cent




Figure 15. The time-delay between marshall autovetting placing a source into the Inbox and the *final* manual vetting action performed on a source for sources in Table 3. Note that sources may have been vetted multiple times in this window. Examples include being sent to the *Stream*, then moved to the *Store* once a spectroscopic classification reveals the source to be of that disposition, or to *Junk/Pending* if future observations of the location raises questions about its realness. Such cases contribute the small number with a delay of > 1 weekend.

- (ii) The nearest match within 90 arcsec, if one exists, is displayed as contextual information for a source.
- (iii) A cautionary note is added if the source is within the Hill sphere of a major Solar-system body.

A source must have no nearby minor body, nor be within a major body's Hill Sphere, to proceed for checking of automatic follow-up criteria (Section 5.6.1).

5.3.4 Manual vetting and reporting

Collaboration members primarily check the Inbox, to make final vetting decisions on those identified by the autovetting decisions as worthy of consideration as real. One option is to return the source to the Pending source list, at which point it will be resurrected to the Inbox if new photometry meets the Inbox criteria again. The breakdown for destinations of Inbox sources is shown in Table 3. Typically ~ 100 objects require vetting per 24 h period, easily manageable by collaboration members. Fig. 15 shows the delay-time distribution of manual vetting.

Once a source is identified as a candidate transient and sent to the *Stream* source list, the marshall initiates a forced photometry (Section 5.5) task to gather the full history of variability at the location. The results of this forced photometry are then used to provide the discovery detection and recent non-detection when reporting the source to the TNS, which is performed via a form in the web interface, or externally via a Gamma-ray Coordination Network circular.²⁰

5.4 External trigger cross-matching

GOTO responds to different alerts at scheduler level – currently GW, GRB, and neutrino alerts from a variety of space missions and detectors. The same package is responsible for notifying the telescope scheduling software to react to a trigger is polled by the marshall to perform its own, independent cross-matching of sources and observations to triggered events.

²⁰<https://gcn.nasa.gov/>

5.4.1 The need for an independent cross-matching system

Sources discovered in images that were observed as part of a follow-up campaign of an alert have the alert's event name assigned to them by KADMILOS, which is pulled to the marshall, allowing users to filter for potential counterparts to the event. For a variety of reasons it is of interest that the marshall performs an independent cross-match of sources with potentially linked skymap localizations of events, regardless of the telescope scheduler's observing reason when it discovered a source. These are:

- (i) The time between event triggers and the release of the alert may be delayed by minutes to hours. During that delay, sources discovered during serendipitous coverage of the localization by GOTO would be missed.
- (ii) Scheduled follow-up is linked to at most one-event, whereas alerts may overlap spatiotemporally and so a given source should be identified as a potential counterpart to more than one alert.
- (iii) Particularly for GW events, both the localization and electromagnetic science interest of events can change via alerts after updated analysis by the GW detector teams. Events may not warrant explicit follow-up at the time but are interesting to cross-match later, even with only serendipitous coverage.

5.4.2 Alert ingestion and processing

New alerts are polled regularly from the GOTO scheduling data base which handles alert parsing (G-TECS; M. J. Dyer 2020; M. J. Dyer et al. 2018). If an alert relates to a new event, relevant trigger information is stored in the marshall data base and a skymap is extracted from the alert and assigned to it. The cross-matching system builds upon the Gravitational wave Electromagnetic RecoveRY code (GERRY; D. O'Neill et al. 2024) and HEALPIX-ALCHEMY (L. P. Singer et al. 2022), accessed via underlying PostgreSQL rangeset data types. These rangesets are used to store both Multi-order Coverage (MOC) skymaps (P. Fernique et al. 2014, 2019, 2022) and image HEALPIX pixel information. This approach enables rapid querying of source positions and image coverage against a skymap.

All skymaps, source positions, and image footprints are ingested into a specific cross-matching DB. In the case of GW events, distance information is also stored alongside the HEALPIX values. Whilst typical skymaps released by LVK in GW alerts are in the MOC format, which primarily allow for a large reduction in the number pixels to describe the skymap, skymaps released by other facilities in the GRB domain use flat resolution maps with ~ 100 times the number of pixels. To alleviate the storage and computational cost, any flat resolution skymap is converted to a MOC prior to ingest via a modified implementation of `bayestar_adaptive_grid` in the LIGO.SKYPACK package (L. P. Singer & L. R. Price 2016).²¹ The probability density function is first estimated by interpolating the flat skymap using pixel coordinates and pixel probability values. The repeated rounds of upsampling draw from this interpolation until the maximum resolution of the MOC matches the resolution of the original skymap. Typically the process results in a MOC skymap with ~ 3000 pixels corresponding to a >98 per cent reduction in overheads, without sacrificing the accuracy of the localization contours.

Not all alerts provide a skymap. In particular, instruments capable of arcminute localizations, such as *Swift*'s X-ray Telescope (D. N. Burrows et al. 2005) or Einstein Probe's Wide-field X-ray Telescope (W. Yuan et al. 2022), report only coordinates and a localization error radius in alert files. For these events a flat high resolution skymap is created assuming a 2D Gaussian probability distribution at the coordinates, with a spread matching the localization error given in the alert. In these cases, only pixels inside a 5-sigma contour of the Gaussian centre are ingested.

If another alert and/or skymap is released for the same event, the event values are updated and the new skymap is ingested, which is set as the currently active skymap for the purposes of cross-matching. Sources already cross-matched with a previous skymap version, but are not matched to the active skymap, are flagged with a caution.

Coordination of survey effort during GW alert follow-up by surveys like GOTO has been formalized in platforms such as the GW Treasure Map²² (S. D. Wyatt et al. 2020), with an aim to avoid duplicated effort, and maximize recovery chances (I. Keinan & I. Arcavi 2025). Although GOTO participates via submission to the GW Treasure Map platform, the scheduler is not influenced or aware of other surveys' pointings via any mechanism, and reacts in an isolated fashion to each alert.

5.4.3 Source and image cross-matching

New sources in the marshall are ingested into the cross-matching data base, ready to be cross-matched with external triggers. Alerts are initially marked as 'high' and 'low' priority, assigning, respectively, an immediate or interval cross-matching schedule by the marshall. GRB, neutrino, and high significance GW alerts with prospects for an EM counterpart are all given 'high' priority.

If there are active events designated as 'high' priority, new sources are cross-matched immediately as they are ingested. For 'low' priority alerts, source cross-matching is performed by a dedicated task at 15 min intervals, which also carries out image cross-matching and 2D localization coverage calculations, following D. O'Neill et al. (2024). If a source is cross-matched with an event skymap, the joint source-skymap properties such as contour level and absolute magnitude (in the cases where the skymap pixels have distance

²¹<https://lscsoft.docs.ligo.org/ligo.skymap/index.html>

²²<https://treasuremap.space/>

information) are calculated and stored in a junction table between the source and skymap tables. Users can then filter cross-matched sources based on these properties. Events have ‘active’ status and are cross-matched for fixed windows of time post-event: 3 and 5 d for GRB and GW events, respectively. After this, events are labelled ‘archived’ with no regular cross-match updating. Events that do not qualify for automatic cross-matching (i.e. are immediately labelled as ‘archived’ by the marshall) can have cross-matching run manually using the marshall interface.

5.4.4 Promising sources

The cross-matched sources streams are regularly parsed to identify the most promising optical counterpart candidates for both GW and GRB events. Sources in the `Stream`, `Inbox`, and `Pending` source lists that were discovered within 24 h of the trigger time are required to have multiple detections post-trigger, with no pre-trigger detections from forced photometry (Section 5.5).

Depending on the type of event, identification of a promising source follows different routes. For GW events, sources are marked as promising if they have been matched to a galaxy which has a distance estimate consistent with the GW distance estimate at the source position, and if the source has a peak absolute magnitude that is very generously consistent with a kilonova – a wide range of $-10 \geq M_L \geq -17$ mag is used, to allow for uncertain distances and photometry systematics.

For GRB events, a source is marked as promising if the light curve can be fit with a power-law decay with an index $-0.3 \geq \alpha \geq -3$, which encompasses theoretical and observational constraints on the optical behaviour of GRBs (e.g. A. Melandri et al. 2008). If a promising source is identified by either means, the source is highlighted visually in the marshall UI, automatically appears at the top of any cross-match source lists and triggers a team Slack alert for collaboration members.

5.5 Forced photometry

For the generation of sources, the marshall relies on the ingestion of photometry of automatically extracted sources (following Section 4.5.3) in difference images. This process is however incomplete for marginal detections, which may not pass the source detection criteria, nor is it able to provide non-detections for transient sources. For this, we rely on a forced photometry routine, which mimics the aperture-corrected photometry automatically extracted, and uses the same calibration products for the respective image. The positions of apertures are instead selected irrespective of the image properties, and are rather set by the source(s) one wishes to measure.

The implementation of forced photometry has been extracted from the KADMILOS infrastructure and now lives as a stand-alone package accessible via an API or a web-frontend to handle both users’ and automated requests (e.g. from the marshall), independent of the pipeline. Full details of the implementation are deferred to Jarvis et al. (in preparation).

5.6 Triggering of follow-up facilities

The marshall has the capability to trigger other facilities via user-interaction, or automatically once a source meets pre-defined criteria (Section 5.6.1). Currently liverpool telescope (I. A. Steele et al. 2004) imaging, spectroscopy (A. S. Piascik et al. 2014) and polarimetry (H. Jermak, I. A. Steele & R. J. Smith 2016), and pt5m (L. K. Hardy et al. 2015) imaging are implemented.

Users can define the observation sequence for each instrument via a web-form, which is then sent to the relevant facility (Fig. 16). Additional integration of facilities with APIs is ongoing.

5.6.1 Automatic follow-up scheduling

The GOTO Auto-Trigger (GOAT) is a system within the marshall which performs a series of checks on new sources, based on both contextual and photometric information (Fig. 14), to determine if the source is viable for automated follow-up by another facility. These checks performed by the GOAT on each new source arriving into the marshall are described in Table 4. The variables shown are tuned per-follow-up facility, to reflect the external telescopes’ capabilities.

In Fig. 17 we show a breakdown of the ultimate nature of sources that were triggered by the GOAT during 2025. Happily enough, even from a single detection, the checks select a majority of real astrophysical objects (transients, Galactic sources, and AGN), although both image quality (both data issues during hardware acquisition and artefacts from software processing) and moving objects (Solar-system objects and *JWST*) contribute contaminants.

It should be noted that this distribution includes triggers made during development of the final checking criteria, and so includes a number of triggers that would no longer be made in the current iteration. For example, early iterations did not include a check of whether a source was in a planet’s Hill sphere, which allowed for Jovian Moon triggers when they are vexatiously in conjunction with catalogued galaxies.

Fig. 18 shows that the GOAT typically inserts observing sequences into the LT observation queue around 11 min after shutter close. If they are picked up expediently by the LT scheduler, this offers a route to early characterization of new discoveries (e.g. M. R. Magee et al. 2025).

trigger LT/SPRAT for GOTO25glk

this will trigger a LT/SPRAT observation of the source - be sure!

exposure times can be estimated from the [Exposure Time Calculator](#)

proposal: 23:1904.77
the proposal to trigger the observations in must be in the format hh:mm:ss.ss

dec: -08:28:40.12
must include sign and be in the format +/-dd:mm:ss.ss

exptime: 400
exposure time in seconds for a single exposure (10-1200)

airmass: 2.0
airmass constraint

seeing: 3
seeing constraint in arcseconds

valid for nights: 2
number of nights group valid for, e.g. 5 means until next to polaris sunrise (1-30)

grating: blue
sprat grating to use

number exposures: 2
number of exposures (1-20)

sky brightness: 4
sky brightness constraint expressed in 14 mag above disk

photometric: no
photometric constraint

trigger

Figure 16. An example webform in the GOTO marshall allowing a user to trigger an LT spectroscopic observation using the SPRAT (Spectrograph for the Rapid Acquisition of Transients) instrument.

Table 4. Summary of triggering criteria checks applied to candidate transient sources. θ values refer to separations, t values are times, and m are magnitudes.

Criterion	Condition	Description
<i>Source metadata and sky position</i>		
Vetting result	<code>vetting_result</code> $\in$ {Stream, Inbox}	Exclude sources where autovetting has not indicated they are a viable transient candidate.
Galactic coordinates	$ b > b_{\min}$ OR $ \ell - 180^\circ < 90^\circ$	Avoid crowded Galactic plane regions.
Declination	$\delta_{\min} < \delta < \delta_{\max}$	Restrict to telescope-accessible declinations.
Reddening	$E(B - V) \leq E(B - V)_{\max}$	Exclude regions with high Galactic extinction.
<i>Single catalogue cross-matches</i>		
Nearest Minor planet	$\theta_{\text{MP}} \geq \theta_{\text{MP,lim}}$ OR $m_{\text{MP}} \geq m_{\text{MP,lim}}$	Avoid triggering on known minor planets.
Planet Hill sphere	No match	Reject sources coincident with planetary Hill spheres.
Bright nearby star	NOT ($m_{\text{bright}} < 9$ AND $\theta_{\text{bright}} < 25$ arcsec) NOR ($m_{\text{bright}} < 14$ AND $\theta_{\text{bright}} < 7.5$ arcsec)	Avoid diffraction spikes or bright star residuals.
Gaia match	$\theta_{\text{Gaia}} > 2.5$ arcsec UNLESS $\theta_{\text{galaxy}} < 10$ arcsec	Avoid triggering on <i>Gaia</i> stellar sources which are not located near to a catalogued galaxy (where galaxy nuclei or point-like components could exist).
<i>Moriarty contextual information</i>		
Moriarty classification	<code>moriarty_classification</code> $\notin$ {var, agn}	Avoid known stellar variable and AGN sources.
Candidate host association	$\theta_{\text{galaxy}} < \theta_{\text{galaxy,lim}}$ OR located in a galaxy cluster	Require some extragalactic context.
<i>Light-curve and photometric behaviour</i>		
Recent robust detection	$\exists i \in \mathcal{D}_{\text{det}} : t_i > t_{\text{now}} - \Delta t_{\text{lim}}$ and $\text{RB}_i \geq \text{RB}_{\text{lim}}$	Ensures there exists at least one recent, high realbogus entry in the set of detections.
Number of detections	$\text{count}\{i \in \mathcal{D}_{\text{det}} : t_i > t_{\text{now}} - \Delta t_{\text{lim}} \text{ and } \text{SNR}_i \geq \text{SNR}_{\min}\} \geq N_{\text{det,min}}$	Ensure a minimum number of recent, high SNR entries in the set of detections.
Historical detections	$\text{count}\{i \in \mathcal{D}_{\text{det}} : t_i \leq t_{\text{now}} - \Delta t_{\text{lim}}\} \leq N_{\text{hist,max}}$	Exclude sources with historical entries in the set of detections (e.g. variables, AGN, persistent difference artefacts).
Non-detections	$\forall j \in \mathcal{D}_{\text{non}} : t_j > t_{\text{first}} \Rightarrow m_j < m_{i_{\text{first}}}$	Enforces all non-detection limits after the first detection are brighter than it, to ensure a rising light curve.
Latest detection age	$t_{i_{\text{latest}}} > t_{\text{now}} - 24$ h	Ensure source has been recently observed.
Brightness	$m_{i_{\text{latest}}} < m_{\text{lim}}$	Latest magnitude must be bright enough to follow up.
Rise rate	$m \leq m_{\min}$	Require a rise brightness rate, measured in magnitudes, and relative to non-detections or earlier detections.




Figure 17. A breakdown of the nature of sources passing the GOAT checks (Table 4) for automatic triggering of spectroscopic follow-up by LT/SPRAT. The triggers include those for sources passing early implementations of the checks, making it a pessimistic view of the current performance (see Section 5.6.1).




Figure 18. Cumulative delay-time distribution for LT/SPRAT triggers submitted by the GOAT. The time is measured from GOTO shutter close until the marshall sends the request to the LT API. Percentiles describing the distribution are shown in the plot as text.

5.7 Citizen science integration

The marshall also has close integration with the Kilonova Seekers citizen science project (Killestein et al. 2024), which makes use of volunteer classifications of difference image stamps to identify various transients and outbursting sources in GOTO data (T. L. Killestein et al. 2025). The marshall regularly polls the Kilonova Seekers backend via a private API, awaiting new high confidence sources from the project. Once a source has at least eight votes cast by volunteers, with an agreement of ≥ 80 per cent that the source is astrophysically real, the marshall gains access to Kilonova Seekers information on the source. This information includes names of a random subset of volunteers who voted the candidate as real, which allows subsequent TNS reporting of the source via the marshall to automatically populate a remark in the discovery report, providing credit for the volunteers' efforts.

Additionally, the Kilonova Seekers upload pipeline bootstraps contextual information from the marshall, enabling uploads to be rebalanced dynamically to maximize the number of probable transients sent to the platform – specifically by not uploading transients already reported, and downweighting minor planets in the random subset selected.




Figure 19. The cumulative sum of TNS discovery reports by GOTO (solid line), and the subset of first-discoveries (dashed line). The dotted line shows the fraction of template images held by the pipeline – a reasonable proxy of sky coverage fraction where difference imaging could be performed. Significant milestones assisting the reporting of new transient candidates are also indicated.

5.8 Evaluation of transient recovery performance

A high-level overview of the overall performance of the transient workflow in the GOTO project is made in this section, using reporting statistics from the public TNS data base.

The TNS data base shows 4623 first discoveries by GOTO between the dates of 2019 May 8 and 2026 January 28, with a total of 11 800 reported transient candidates in that same window. The typical TNS reporting rate was ~ 2 reports per day until May 2024 when a ~ 50 per cent increase in template coverage (reaching almost complete coverage) and new pipeline features (primarily Autovetting) were deployed. These combined to provide a step change in reporting to 15–16 reports per day, which has remained consistent until early 2026. The cumulative distribution of GOTO TNS reports is visualized in Fig. 19. The median magnitude of first detection magnitudes reported to the TNS for GOTO is $m_L \simeq 19.5$ mag over 2025, broadly consistent with expectations from the sky-survey depths (Fig. 11) when considering a more robust confirmation than a single 5σ detection is required to meet the threshold of reporting an object to the community.

The overall TNS reporting rate did not significantly alter in the periods of GW-detector down-time (namely April–June 2025 and November 2025–early 2026). Outside these windows, significant time was dedicated to covering GW localizations during the O4 observing runs of the GW detectors, in favour of the all-sky survey. There is the hint of an up-tick in the first-discovery rate since late 2025, when observing strategies to revisit regions of the sky at higher cadence have been implemented.

In addition to the overall transient discovery rate, the workflow has delivered confirmed optical afterglow detections for GRBs (S. Belkin et al. 2024; Dimple et al. 2025; A. Kumar et al. 2025), including events initially localized only to degrees-scale regions. These recoveries, within hours of the high-energy triggers, exemplify the combination of wide-field observations via rapid robotic response, with low-latency data analysis and automated candidate vetting to efficiently isolate true optical counterparts from large candidate streams – a practical validation of the pipeline prioritization strategy.

6 SUMMARY AND OUTLOOK

We have detailed the baseline infrastructure for transient discovery with GOTO in the preceding sections, and its performance and feature set has met the design requirements of the transients working group within the project: robust detection and reporting of sources to the community starting just minutes after GOTO shutter close, with automatic follow-up being triggered for high priority sources. GOTO is discovering transients at a rate commensurate with prolific contemporary surveys such as ATLAS and ZTF, and adds valuable coverage to the global effort to map the changing night-sky on short time-scales.

Future upgrades for this transient workflow continue. Some key avenues for extracting additional robustness of the framework, and science exploitation of the discoveries have been identified:

- (i) Airflow 3.0 was released in April 2025, offering DAG versioning. An upgrade is planned which will help maintain provenance of GOTO outputs and a more tractable DAG history in the face of the slow-changing dimensions of GOTO hardware.
- (ii) The *Gaia* Synthetic Photometry Catalogue (Gaia Collaboration 2023) offers a route to directly compute a calibration catalogue natively in the model GOTO-L filter bandpass, with a consequently smaller colour term required to be calculated per UT (to account for filter variations). An in-development online calibration process for forced photometry employs a similar technique using synthetic photometry and time-variable colour corrections to achieve increased accuracy of GOTO photometry (Jarvis et al., in preparation).

(iii) The bank of template images could be superseded by identification deeper stacks of multiple existing set images now that GOTO has visited its entire survey footprint many times. This would prevent the depth of the template images compromising the difference image depths in the face of improved throughput of the GOTO hardware and/or longer exposure-time sub-surveys.

(iv) Software developed as part of the Vera C. Rubin LSST (Ž. Ivezić et al. 2019) include advanced Solar-system and moving body packages (S. R. Merritt et al. 2025), which can be used for improved cross-checking of discoveries against known objects.

(v) Photometry and calibration of all sources in the science frames is not required for transient discovery. The differential in processing time arising from the density of sources (Fig. 12) could be removed by trimming the number of calibration and photometered sources per image down to a level where it does not materially affect the final results. In practice, however, the most crowded images are also likely those suffering from the most foreground extinction, and so are already of the least interest for transient searches. Although complete photometry is of interest for variable star searches, bespoke workflows for crowd-optimized photometry is required for robust results, and KADMILOS remains focused on the low-latency discovery of transients.

(vi) Improvements to the contextual classifier, MORIARTY, to include deep survey catalogues, redshift catalogues from large spectroscopic surveys (e.g. DESI Collaboration 2026), and more robust galaxy host association, to minimize false associations.

(vii) Deploying new machine learning models for real-bogus classification and more granular multitask labelling, which run online to mitigate the effects of data/concept drift, and can be regularly fine-tuned with modest numbers of recent, high-quality labels.

(viii) Source vetting currently is a final step in the marshall, whereas interesting behaviour can be observed in historic sources (late rebrightenings of transients, extreme behaviour of AGN etc.) Filters designed to place human eyes back on these sources can be added to the framework of the marshall.

(ix) The current reporting method is via TNS discovery reports. In future, a model of releasing streams of alerts to which brokers can listen, akin to that employed by ZTF (M. T. Patterson et al. 2019) and LSST,²³ will further open the use of GOTO data to the wider community.

ACKNOWLEDGEMENTS

The authors' thanks go to the contributors of the open source software projects upon which this work stands. Gratitude is given to David Young for fruitful conversations and inspiration on marshall development.

We additionally thank the two reviewers for constructive feedback and suggestions to improve the clarity of the work.

The Gravitational-wave Optical Transient Observer (GOTO) project acknowledges the support of the Monash-Warwick Alliance; University of Warwick; Monash University; University of Sheffield; University of Leicester; Armagh Observatory & Planetarium; the National Astronomical Research Institute of Thailand (NARIT); Instituto de Astrofísica de Canarias (IAC); University of Portsmouth; University of Turku; University of Birmingham; and the UK Science and Technology Facilities Council (STFC, grant numbers ST/T007184/1, ST/T003103/1 and ST/Z000165/1).

JDL, DON and MP acknowledge support from a UK Research and Innovation Future Leaders Fellowship (grant references MR/T020784/1 and UKRI1062). TLK acknowledges support from a Warwick Astrophysics prize post-doctoral fellowship made possible thanks to a generous philanthropic donation. DS acknowledges support from The Science and Technology Facilities Council (STFC) via grants ST/T007184/1, ST/T003103/1, ST/T000406/1, ST/X001121/1. BW acknowledges the UKRI's STFC studentship grant funding, project reference ST/X508871/1. JC acknowledges support by the Spanish Ministry of Science via the Plan de Generación de Conocimiento through grant PID2022-143331NB-I00. DLC acknowledges support from the Science and Technology Facilities Council (STFC) grant number ST/X001121/1. LK acknowledges support for an Early Career Fellowship from the Leverhulme Trust through grant ECF-2024-054 and the Isaac Newton Trust through grant 24.08(w). DMS acknowledges support through the Ramón y Cajal grant RYC2023-044941, funded by MCIU/AEI/10.13039/501100011033 and FSE+. SMa acknowledges financial support from the Research Council of Finland project 350458. SMO is funded by Leverhulme Trust grant RPG-2023-240. AS acknowledges the Warwick Astrophysics PhD prize scholarship made possible thanks to a generous philanthropic donation. RLCS acknowledges funding from Leverhulme Trust grant RPG-2023-240.

This research has made use of data and/or services provided by the International Astronomical Union's Minor Planet Center.

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

This publication uses data generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation.

Computing facilities were provided by the Scientific Computing Research Technology Platform of the University of Warwick.

This research has made use of the NASA/IPAC Extragalactic Data base, which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology, including NED LVS (10.26132/NED8).

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III web site is <http://www.sdss3.org/>. SDSS-III is managed

²³[doi:10.71929/rubin/2586493](https://doi.org/10.71929/rubin/2586493)

by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

The national facility capability for SkyMapper has been funded through ARC LIEF grant LE130100104 from the Australian Research Council, awarded to the University of Sydney, the Australian National University, Swinburne University of Technology, the University of Queensland, the University of Western Australia, the University of Melbourne, Curtin University of Technology, Monash University and the Australian Astronomical Observatory. SkyMapper is owned and operated by The Australian National University's Research School of Astronomy and Astrophysics. The survey data were processed and provided by the SkyMapper Team at ANU. The SkyMapper node of the All-Sky Virtual Observatory (ASVO) is hosted at the National Computational Infrastructure (NCI). Development and support of the SkyMapper node of the ASVO has been funded in part by Astronomy Australia Limited (AAL) and the Australian Government through the Commonwealth's Education Investment Fund (EIF) and National Collaborative Research Infrastructure Strategy (NCRIS), particularly the National eResearch Collaboration Tools and Resources (NeCTAR) and the Australian National Data Service Projects (ANDS).

The Pan-STARRS1 Surveys (PS1) and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation Grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

The Legacy Surveys consist of three individual and complementary projects: the Dark Energy Camera Legacy Survey (DECaLS; Proposal ID #2014B-0404; PIs: David Schlegel and Arjun Dey), the Beijing-Arizona Sky Survey (BASS; NOAO Prop. ID #2015A-0801; PIs: Zhou Xu and Xiaohui Fan), and the Mayall z-band Legacy Survey (MzLS; Prop. ID #2016A-0453; PI: Arjun Dey). DECaLS, BASS and MzLS together include data obtained, respectively, at the Blanco telescope, Cerro Tololo Inter-American Observatory, NSF's NOIRLab; the Bok telescope, Steward Observatory, University of Arizona; and the Mayall telescope, Kitt Peak National Observatory, NOIRLab. Pipeline processing and analyses of the data were supported by NOIRLab and the Lawrence Berkeley National Laboratory (LBNL). The Legacy Surveys project is honored to be permitted to conduct astronomical research on Iolkam D'uag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

NOIRLab is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation. LBNL is managed by the Regents of the University of California under contract to the U.S. Department of Energy.

This project used data obtained with the Dark Energy Camera (DECam), which was constructed by the Dark Energy Survey (DES) collaboration. Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundacao Carlos Chagas Filho de Amparo, Financiadora de Estudos e Projetos, Fundacao Carlos Chagas Filho de Amparo a Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministerio da Ciencia, Tecnologia e Inovacao, the Deutsche Forschungsgemeinschaft and the Collaborating Institutions in the Dark Energy Survey. The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energeticas, Medioambientales y Tecnologicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zurich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, NSF's NOIRLab, the University of Nottingham, the Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University.

BASS is a key project of the Telescope Access Program (TAP), which has been funded by the National Astronomical Observatories of China, the Chinese Academy of Sciences (the Strategic Priority Research Program 'The Emergence of Cosmological Structures' Grant #XDB09000000), and the Special Fund for Astronomy from the Ministry of Finance. The BASS is also supported by the External Coop-

eration Program of Chinese Academy of Sciences (Grant #114A11KYSB20160057), and Chinese National Natural Science Foundation (Grant #12120101003, #11433005).

The Legacy Survey team makes use of data products from the Near-Earth Object Wide-field Infrared Survey Explorer (NEOWISE), which is a project of the Jet Propulsion Laboratory/California Institute of Technology. NEOWISE is funded by the National Aeronautics and Space Administration.

The Legacy Surveys imaging of the DESI footprint is supported by the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH1123, by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract; and by the U.S. National Science Foundation, Division of Astronomical Sciences under Contract No. AST-0950945 to NOAO.

Software: Apache Airflow,²⁴ *astropy* (Astropy Collaboration 2013, 2018, 2022), *lmfit* (M. Newville et al. 2025), *astroscrappy* (C. McCully et al. 2018), *sep* (K. Barbary 2016), *astrometrynet* (D. Lang et al. 2010), *q3c* (S. Koposov & O. Bartunov 2006), *rbf*,²⁵ *dill*.²⁶

CONFLICT OF INTEREST

Authors declare no conflict of interest.

DATA AVAILABILITY

Underlying image-level data are to be released within the scope of GOTO data releases. Prior to this, reasonable requests to the corresponding author can be made. The high-level data required to reproduce figures is available upon request. A static copy of the main KADMILOS data pipeline code as described here is held at <https://github.com/GOTO-OBS/kadmiolos-public>.

REFERENCES

- Abbott B. P. et al., 2017, *ApJ*, 848, L12
 Alam S. et al., 2015, *ApJS*, 219, 12
 Alard C., 2000, *A&AS*, 144, 363
 Alard C., Lupton R. H., 1998, *ApJ*, 503, 325
 Annis J. et al., 2014, *ApJ*, 794, 120
 Astropy Collaboration, 2013, *A&A*, 558, A33
 Astropy Collaboration, 2018, *AJ*, 156, 123
 Astropy Collaboration, 2022, *ApJ*, 935, 167
 Barbary K., 2016, *J. Open Source Softw.*, 1, 58
 Becker A., 2015, Astrophysics Source Code Library, record ascl:1504.004
 Belkin S. et al., 2024, *Res. Notes Am. Astron. Soc.*, 8, 6
 Bellm E. C. et al., 2019, *PASP*, 131, 018002
 Bertin E., Arnouts S., 1996, *A&AS*, 117, 393
 Bloom J. S. et al., 2012, *PASP*, 124, 1175
 Burrows D. N. et al., 2005, *Space Sci. Rev.*, 120, 165
 Carrasco-Davis R. et al., 2021, *AJ*, 162, 231
 Chambers K. C. et al., 2016, preprint ([arXiv:1612.05560](https://arxiv.org/abs/1612.05560))
 Cook D. O. et al., 2023, *ApJS*, 268, 14
 DESI Collaboration, 2026, *AJ*, 171, 285
 Dálya G. et al., 2022, *MNRAS*, 514, 1403
 Dey A. et al., 2019, *AJ*, 157, 168
 Dimple et al., 2025, *MNRAS*, 544, 548
 Duev D. A. et al., 2019, *MNRAS*, 489, 3582
 Dyer M. J., 2020, PhD thesis, University of Sheffield, UK
 Dyer M. J., Dhillion V. S., Littlefair S., Steeghs D., Ulaczyk K., Chote P., Galloway D., Rol E., 2018, in Proc. SPIE Conf. Ser. Vol. 10704, Observatory Operations: Strategies, Processes, and Systems VII. SPIE, Bellingham, p. 107040C
 Dyer M. J. et al., 2020, in Guzman J. C., Ibsen J., eds, Proc. SPIE Conf. Ser. Vol. 11452, Software and Cyberinfrastructure for Astronomy VI. SPIE, Bellingham, p. 114521Q
 Dyer M. J. et al., 2024, in Marshall H. K., Spyromilio J., Usuda T., eds, Proc. SPIE Conf. Ser. Vol. 13094, Ground-based and Airborne Telescopes X. SPIE, Bellingham, p. 130941X
 Evans P. A., Page K. L., Beardmore A. P., Eyles-Ferris R. A. J., Osborne J. P., Campana S., Kennea J. A., Cenko S. B., 2023, *MNRAS*, 518, 174
 Fernique P., Boch T., Donaldson T., Durand D., O'Mullane W., Reinecke M., Taylor M., 2014, *MOC - HEALPix Multi-Order Coverage map Version 1.0, IVOA Recommendation 2 June 2014*

²⁴<https://github.com/apache/airflow>

²⁵<https://github.com/treverhines/RBF>

²⁶<https://github.com/uqfoundation/dill>

- Fernique P., Boch T., Donaldson T., Durand D., O'Mullane W., Reinecke M., Taylor M., 2019, [MOC - HEALPix Multi-Order Coverage map Version 1.1, IVOA Recommendation 7 October 2019](#)
- Fernique P. et al., 2022, [MOC: Multi-Order Coverage map Version 2.0, IVOA Recommendation 27 July 2022](#)
- Flewelling H., 2018, in American Astronomical Society Meeting Abstracts #231. p. 436.01
- Förster F. et al., 2021, [AJ](#), 161, 242
- Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, [AJ](#), 111, 1748
- Gaia Collaboration, 2016, [A&A](#), 595, A1
- Gaia Collaboration, 2018, [A&A](#), 616, A1
- Gaia Collaboration, 2023, [A&A](#), 674, A33
- Godson B. et al., 2023, Transient Name Server AstroNote, 224, 1
- Gomez S., Berger E., Blanchard P. K., Hosseinzadeh G., Nicholl M., Villar V. A., Yin Y., 2020, [ApJ](#), 904, 74
- Gompertz B. P. et al., 2020, [MNRAS](#), 497, 726
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, [ApJ](#), 622, 759
- Groot P. J. et al., 2024, [PASP](#), 136, 115003
- Hardy L. K., Butterley T., Dhillion V. S., Littlefair S. P., Wilson R. W., 2015, [MNRAS](#), 454, 4316
- Harris C. R. et al., 2020, [Nature](#), 585, 357
- Howell S. B. ed. 1992, ASP Conf. Ser. Vol. 23, Astronomical CCD Observing and Reduction Techniques. Astron. Soc. Pac., San Francisco
- Iske A., 2003, Rendiconti del Seminario Matematico, 61, 247
- Ivezić Ž. et al., 2019, [ApJ](#), 873, 111
- Jermak H., Steele I. A., Smith R. J., 2016, in Evans C. J., Simard L., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI. SPIE, Bellingham, p. 99084I
- Kaiser N. et al., 2002, in Tyson J. A., Wolff S., eds, Proc. SPIE Conf. Ser. Vol. 4836, Survey and Other Telescope Technologies and Discoveries. SPIE, Bellingham, p. 154
- Kasliwal M. M. et al., 2019, [PASP](#), 131, 038003
- Keinan I., Arcavi I., 2025, [ApJ](#), 985, 142
- Killestein T., 2023, PhD thesis, University of Warwick, UK
- Killestein T. L. et al., 2021, [MNRAS](#), 503, 4838
- Killestein T. L. et al., 2024, [MNRAS](#), 533, 2113
- Killestein T. L. et al., 2025, [A&A](#), 699, A8
- Kochanek C. S. et al., 2017, [PASP](#), 129, 104502
- Koposov S., Bartunov O., 2006, in Gabriel C., Arviset C., Ponz D., Enrique S., eds, ASP Conf. Ser. Vol. 351, Astronomical Data Analysis Software and Systems XV. Astron. Soc. Pac., San Francisco, p. 735
- Koposov S., Bartunov O., 2019, Astrophysics Source Code Library, record ascl:1905.008
- Kumar A. et al., 2025, [MNRAS](#), 544, 1541
- Lang D., Hogg D. W., Mierle K., Blanton M., Roweis S., 2010, [AJ](#), 139, 1782
- Law N. M. et al., 2009, [PASP](#), 121, 1395
- Lyman J. D., 2021, Astrophysics Source Code Library, record ascl:2103.003
- Lyman J. D., 2025, Astrophysics Source Code Library, record ascl:2501.002
- Magee M. R. et al., 2025, [MNRAS](#), 543, 3731
- Matheson T. et al., 2021, [AJ](#), 161, 107
- McCully C. et al., 2018, [astropy/astrocrappy: v1.0.5 Zenodo Release](#). Zenodo
- Meegan C. et al., 2009, [ApJ](#), 702, 791
- Melandri A. et al., 2008, [ApJ](#), 686, 1209
- Merritt S. R. et al., 2025, [AJ](#), 170, 100
- Mitchell M. A. et al., 2025, [MNRAS](#), 545, staf2100
- Möller A. et al., 2020, [MNRAS](#), 501, 3272
- Muthukrishna D., Narayan G., Mandel K. S., Biswas R., Hložek R., 2019, [PASP](#), 131, 118002
- Newville M. et al., 2025, [LMFIT: Non-Linear Least-Squares Minimization and Curve-Fitting for Python](#). Zenodo
- Nordin J. et al., 2019, [A&A](#), 631, A147
- O'Neill D. et al., 2024, in Benn C. R., Chrysostomou A., Storrie-Lombardi L. J., eds, Proc. SPIE Conf. Ser. Vol. 13098, Observatory Operations: Strategies, Processes, and Systems X. SPIE, Bellingham, p. 1309818
- Ofek E. O. et al., 2023, [PASP](#), 135, 065001
- Onken C. A., Wolf C., Bessell M. S., Chang S.-W., Luvaul L. C., Tonry J. L., White M. C., Da Costa G. S., 2024, [Publ. Astron. Soc. Aust.](#), 41, e061
- Patterson M. T. et al., 2019, [PASP](#), 131, 018001
- Pence W. D., Seaman R., White R. L., 2009, in Bohlender D. A., Durand D., Dowler P., eds, ASP Conf. Ser. Vol. 411, Astronomical Data Analysis Software and Systems XVIII. Astron. Soc. Pac., San Francisco, p. 25
- Pence W. D., White R. L., Seaman R., 2010, [PASP](#), 122, 1065
- Phillips A. C., Davis L. E., 1995, in Shaw R. A., Payne H. E., Hayes J. J. E., eds, ASP Conf. Ser. Vol. 77, Astronomical Data Analysis Software and Systems IV. Astron. Soc. Pac., San Francisco, p. 297
- Piasecik A. S., Steele I. A., Bates I. A., Mottam C. J., Smith R. J., Barnsley R. M., Bolton B., 2014, in Ramsay S. K., McLean I. S., Takami H., eds, Proc. SPIE Conf. Ser. Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V. SPIE, Bellingham, p. 91478H
- Sánchez-Sáez P. et al., 2021, [AJ](#), 161, 141
- Scott D., 1992, Multivariate Density Estimation: Theory, Practice, and Visualization. A Wiley-interscience publication, Wiley, New York https://books.google.co.uk/books?id=7crCUS_F2ocC
- Sheng X. et al., 2024, [MNRAS](#), 531, 2474
- Shupe D. L., Moshir M., Li J., Makovoz D., Narron R., Hook R. N., 2005, in Shopbell P., Britton M., Ebert R., eds, ASP Conf. Ser. Vol. 347, Astronomical Data Analysis Software and Systems XIV. Astron. Soc. Pac., San Francisco, p. 491
- Singer L. P., Price L. R., 2016, [Phys. Rev. D](#), 93, 024013

- Singer L. P., Parazin B., Coughlin M. W., Bloom J. S., Crellin-Quick A., Goldstein D. A., van der Walt S., 2022, *AJ*, 163, 209
- Smartt S. J. et al., 2015, *A&A*, 579, A40
- Smith K. W. et al., 2020, *PASP*, 132, 085002
- Spergel D. et al., 2015, preprint (arXiv:1503.03757)
- Steeghs D. et al., 2022, *MNRAS*, 511, 2405
- Steele I. A. et al., 2004, in Oschmann J. M., Jr, ed., *Proc. SPIE Conf. Ser. Vol. 5489, Ground-based Telescopes*. SPIE, Bellingham, p. 679
- Stevance H. F., Smith K. W., Smartt S. J., Roberts S. J., Erasmus N., Young D. R., Clocchiatti A., 2025, *ApJ*, 990, 201
- Storn R., Price K., 1997, *J. Global Optim.*, 11, 341
- Street R. A., Bowman M., Saunders E. S., Boroson T., 2018, in Guzman J. C., Ibsen J., eds, *Proc. SPIE Conf. Ser. Vol. 10707, Software and Cyberinfrastructure for Astronomy V*. SPIE, Bellingham, 1070711
- Tonry J. L. et al., 2018a, *ApJ*, 867, 105
- Tonry J. L. et al., 2018b, *PASP*, 130, 064505
- van Dokkum P. G., 2001, *PASP*, 113, 1420
- Vassiliadis P., 2009, *Int. J. Data Warehousing Mining (IJDWM)*, 5, 1
- Wahba G. ed. 1990, *Spline Models for Observational Data*. SIAM, Philadelphia
- Wang T. et al., 2023, *Sci. China Phys. Mech. Astron.*, 66, 109512
- Watson C. L., Henden A. A., Price A., 2006, *Soc. Astron. Sci. Annu. Symp.*, 25, 47
- Williams R. D., Francis G. P., Lawrence A., Sloan T. M., Smartt S. J., Smith K. W., Young D. R., 2024, *RAS Techn. Instrum.*, 3, 362
- Wyatt S. D., Tohuvavohu A., Arcavi I., Lundquist M. J., Howell D. A., Sand D. J., 2020, *ApJ*, 894, 127
- Yuan W., Zhang C., Chen Y., Ling Z., 2022, in Bambi C., Sanganelo A., eds, *Handbook of X-ray and Gamma-ray Astrophysics*. Springer, Singapore, p. 86

APPENDIX A: ADDITIONAL FIGURES

A visualization of zeropoints and sensitivity of GOTO-[2-4] is shown analogously to Fig. 10 for GOTO-1 in Fig. A1.

Telescope-specific astrometric residual distributions, as additional insight into the distributions that produce the GOTO-level values shown in Fig. 6 are shown in Fig. A2.




Figure A1. Same as Fig. 10 but shown here for mounts GOTO-2, GOTO-3, and GOTO-4.




Figure A2. Astrometric residuals for each GOTO telescope in different regions of the image sensor. The regions are identical to the coloured heat map regions in Fig. 6. Distributions are shown as kernel-density estimates, with median values as black horizontal bars.

APPENDIX B: spalipy

SPALIPY (J. D. Lyman 2021) provides detection-based alignment of astronomical images, allowing for a flexible transformation of a ‘source’ image on to the pixel-coordinate system of a ‘template’ image. It was built initially from an implementation of image alignment by Malte Tewes. SPALIPY can perform its own detection routine to find common sources, or use externally supplied catalogues (e.g. output from SExtractor, E. Bertin & S. Arnouts 1996). The catalogues are searched for matching ‘quads’ following the algorithm of D. Lang et al. (2010). The best-matching quad in each image is used to derive an initial affine transformation that maps the ‘source’ detections into the coordinates of the ‘template’.

After applying the affine transform to all source detection coordinates, residual offsets to the template positions are evaluated. Because wide-field instruments, including GOTO, frequently exhibit spatially varying distortions, an affine transform is typically insufficient to achieve sub-pixel alignment across the field of view. To rectify this, SPALIPY can optionally fit two independent 2D spline surfaces (x and y alignment residuals) that capture non-uniform, spatially dependent deformations. These spline corrections are then applied to generate a final, refined mapping of the source image on to the template’s frame.

APPENDIX C: pympc

PYMP²⁷ (J. D. Lyman 2025) is a package to calculate bulk positions of major and minor Solar-system bodies for requested epochs and spatial coordinates. It interfaces with data products of the Minor Planet Center (MPC) and uses functionality of the PYEPHEM package²⁸ to achieve this. PYMP is used in the GOTO data pipeline to determine if a known minor or major body is in the vicinity of a new detection in a difference frame, using the astrometric position of that detection and time of observation.

The data from MPC comprises the `mpcorb`, `nea`, and `comets` catalogues, which, provide orbital elements of minor bodies (asteroids, near-earth asteroids, and comets, respectively). These are parsed into the file format of XEPHEM,²⁹ with internal handling then done by PYEPHEM to perform propagation of orbits. The underlying C-library (LIBASTRO) used to propagate orbits provides only geocentric astrometric coordinates, and so PYMP then performs its own topocentric correction based on the position of the observer provided, typically via a MPC Observatory code. The propagation of the orbits is from the epoch defined in the MPC data to the requested epoch by the user. MPC epochs are updated typically every ~ 200 d, and so for checking of the current epoch, there can be a significant propagation required. The offline propagation provided by PYEPHEM is not a full solution as performed by the MPC, and, owing to lack of treatment for planetary perturbations and other effects, the accuracy of a minor body's position is sensitively linked to the length of propagation required. Typically this is $\lesssim 1$ arcsec for epoch differences of several months or less, however it can sharply increase to $\sim$ arcminute levels when far from an MPC epoch update.

Additionally, positions of the planets and their major moons, for which pre-defined objects exist in PYEPHEM, are also checked during a position search. Since this list of moons is incomplete, PYMP allows one to determine if a detection resides (in projection) within the Hill sphere of a planet at that epoch, in which case one should be cautious about the transient nature of the detection.

Typically searches (with multiprocessing) takes of order a few seconds – dominated almost entirely by the minor body search. In practice, one can readily cache their positions for a given epoch to make light work of searching many detections in a given set of images.

¹Department of Physics, University of Warwick, Coventry CV4 7AL, UK

²School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

³Astrophysics Research Cluster, School of Mathematical and Physical Sciences, University of Sheffield, Sheffield S3 7RH, UK

⁴Centre for Electronic Imaging, School of Physical Sciences, The Open University, Walton Hall, Milton Keynes MK7 6AA, UK

⁵Research Software Engineering, University of Sheffield, Sheffield S1 4DP, UK

⁶School of Physics and Astronomy, Monash University, Clayton VIC 3800, Australia

⁷Institute for Globally Distributed Open Research and Education (IGDORE)

⁸Armagh Observatory and Planetarium, College Hill, Armagh BT61 9DG, Northern Ireland

⁹Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain

¹⁰School of Physics and Astronomy, University of Leicester, University Road, Leicester LE1 7RH, UK

¹¹National Astronomical Research Institute of Thailand (Public Organization), Chiang Mai 50180, Thailand

¹²Department of Physics and Astronomy, University of Turku, Vesilinnantie 5, Turku FI-20014, Finland

¹³Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

¹⁴Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK

¹⁵Departamento de Astrofísica, Univ. de La Laguna, E-38206 La Laguna, Tenerife, Spain

¹⁶Department of Astrophysics, American Museum of Natural History, Central Park West and 79th Street, New York, NY 10024-5192, USA

¹⁷Institute of Astronomy and Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

¹⁸School of Physics, University College Cork, Cork T12 K8AF, Ireland

¹⁹Radboud University, Postbus 9010, 6500 GL, Nijmegen, Netherlands

²⁰School of Sciences, European University Cyprus, Diogenes Street, Engomi, 1516 Nicosia, Cyprus

²¹School of Physics, Trinity College Dublin, College Green, Dublin 2, Ireland

²²Astronomy and Astrophysics Section, Dublin Institute for Advanced Studies, DIAS Dunsink Observatory, Dublin D15 XR2R, Ireland

²³Centre for Astrophysics Research, University of Hertfordshire, Hatfield AL10 9AB, UK

²⁷<https://github.com/lyalpha/pympc>

²⁸<https://rhodesmill.org/pyephem/>

²⁹<https://xephem.github.io/XEphem/Site/>

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.