Robotic Scheduling of Time-Critical Observations

Sebastian Buntin

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

April 2026

Publications

In the course of completeing the work presented in this thesis, the content of Chapter 3 has been submitted and accepted for publication in a refereed journal:

Buntin, S., Copperwheat, C. & Jermak, H., 2025, ‘Nighttime Cloud Detection, Tracking and Prediction with All-Sky Cameras’, *Royal Astronomical Society Techniques and Instruments*, Volume 4/2025

Sebastian Buntin May 2026

Declaration of Authorship

I, Sebastian Buntin, declare that this thesis titled, ‘Robotic Scheduling of Time Critical Observations’ and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

“The fact that you know something about the sky doesn’t change its magic.”

Prof. Dr. Harald Lesch

Liverpool John Moores University

Abstract

Faculty of Engineering and Technology
Astrophysics Research Institute

Doctor of Philosophy

by Sebastian Buntin

Astronomy has entered the era of time-domain surveys, in which fast and rare astrophysical transients are discovered on a nightly basis. The most informative phases of such events, including fast blue optical transients, kilonovae, the earliest stages of supernova explosions, and gamma-ray burst afterglows, unfold on timescales of minutes to hours. Capturing these fleeting signals requires robotic telescopes that can not only react rapidly but also make intelligent, fully autonomous scheduling decisions. Static queue schedulers and human-in-the-loop sequence preparation are fundamentally too slow for this purpose, creating a bottleneck that prevents telescopes from realising their full scientific potential.

This thesis develops and evaluates an integrated framework for atmospheric-aware robotic scheduling, designed and tested using the Liverpool Telescope (LT) as a platform and with direct application to the forthcoming New Robotic Telescope (NRT). Four key components are addressed. First, a real-time cloud detection and prediction system based on all-sky imaging provides actionable forecasts on ten- to twenty-minute timescales. Second, an automated extinction measurement pipeline using auxiliary Skycam imagers enables continuous, differential monitoring of atmospheric transparency. Third, empirical and machine-learning models of sky brightness are constructed, incorporating lunar and geometric parameters to deliver predictive estimates of background light levels. Finally, these data products are integrated into a signal-to-noise ratio (SNR) based exposure-time system, replacing fixed exposure assumptions with dynamic, condition-dependent calculations.

Together, these developments demonstrate that robotic telescopes can achieve environmental awareness and adaptive response, optimising efficiency and ensuring data quality even under variable atmospheric conditions. The results lay the foundation for a new generation of schedulers capable of autonomously deciding not only which target to observe, but also how to observe it. By eliminating the human bottleneck and integrating real-time environmental intelligence, the framework presented here positions the NRT - and robotic observatories more broadly - to meet the demands of modern time-domain astronomy.

Acknowledgements

Writing this thesis has been a long journey, and I owe more than I can express to those who have supported me along the way. It is a pleasure to take this opportunity to acknowledge the people without whom this work would not have been possible.

First and foremost, I am deeply grateful to my supervisors, Dr. Chris Copperwheat and Dr. Helen Jermak, for their guidance, patience, and encouragement. Chris, thank you for constantly reminding me of the bigger picture and for steering me towards the science that really matters. Helen, thank you for your steady support, insightful feedback, and for never losing faith in the project even when I did. I could not have asked for better mentors.

I also want to thank my colleagues and collaborators at the Liverpool Telescope and the wider NRT team. Working with you has been both a privilege and an inspiration, and I look forward to seeing how the ideas explored here will find their place in the telescopes of tomorrow.

To my family, I owe the greatest debt of gratitude.

To my wife, Nicola, and our children, Amelia and Oliver: you have been at the centre of this journey from beginning to end. For four long years you have carried me through the late nights, the weekends spent writing, and the countless moments when my thoughts were still at the telescope long after the working day was done. Thank you for your love, your patience, and for reminding me again and again that there is a world beyond La Palma weather statistics and SQL queries.

This thesis was only possible because of your unwavering support. Nicola, you have been my anchor, sharing every frustration and every small victory, and never letting me lose sight of why this work mattered. Amelia and Oliver, your joy, curiosity, and laughter kept me grounded and gave me the strength to keep going when the weight of the project felt overwhelming.

Whatever words appear on the following pages, they belong as much to you as they do to me. This thesis is, in the truest sense, a family achievement, and I dedicate it with all my love and gratitude to you.

To my brother, thank you for stepping in to take on parts of my responsibilities at the company while I focused on finishing this work. Your help made it possible for me

to dedicate the time and energy needed, and I will always be grateful for that act of generosity.

Finally, to my mum and extended family: thank you for always believing in me and for encouraging a curiosity that has carried me into both science and education.

This thesis is dedicated to all of you. Whatever its academic merits, it is the product of your support, sacrifice, and love.

To my father,
who first showed me the stars,
who shared his love of physics and astronomy,
and whose spark lit the flame that carried me here.
In loving memory.

Contents

List of Figures

Listings

Abbreviations

API	**A**pplication **P**rogramming **I**nterface
CCD	**C**harge **C**oupled **D**evice
CMT	**C**arlsberg **M**eridian **T**elescope
CUDA	**C**ompute **U**nified **D**evice **A**rchitecture
GLSL	Open**GL** **S**hading **L**anguage
GPU	**G**raphics **P**rocessing **U**nit
GPGPU	**G**eneral-**P**urpose computing on **G**raphics **P**rocessing **U**nits
LT	**L**iverpool **T**elescope
MML	**M**aritime **M**ixing **L**ayer
NRT	**N**ew **R**obotic **T**elescope
ORM	**O**bservatorio del **R**oque de los **M**uchachos
RCS	**R**obotic **C**ontrol **S**ystem
REST	**RE**presentational **S**tate **T**ransfer
RMSE	**R**oot **M**ean **S**quare **E**rror
RTML	**R**emote **T**elescope **M**arkup **L**anguage
SDSS	**S**loan **D**igital **S**ky **S**urvey
SNR	**S**ignal-to-**N**oise **R**atio
VLT	**V**ery **L**arge **T**elescope
VRO	**V**era **R**ubin **O**bservatory

Chapter 1

Introduction

The universe is an imposing reality,
and an answer to all questions.

Maria Montessori

1.1 Background and Context

Modern astrophysics has entered an era defined by time-domain astronomy. Large-scale surveys such as the Zwicky Transient Facility (ZTF) and, in the near future, the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST), are delivering unprecedented volumes of transient alerts. These discoveries demand rapid and efficient spectroscopic and photometric follow-up by ground-based facilities in order to classify events, constrain their physics, and capture short-lived phenomena. The scientific value of such discoveries is often maximised only if observations can be made within minutes to hours, and under conditions that guarantee usable data quality.

Robotic telescopes occupy a unique role in this landscape. Unlike classical observatories that rely on human observers to judge conditions and adjust observing strategies, robotic telescopes must operate fully autonomously. They must therefore be capable of responding instantly to new targets, assessing atmospheric conditions in real time, and adapting schedules without direct human intervention. This autonomy allows them to achieve high scientific throughput and to respond rapidly to transient alerts, but it also exposes them to significant vulnerabilities. Chief among these are the dynamic and often unpredictable effects of the Earth's atmosphere: cloud coverage, variable atmospheric extinction, and fluctuations in sky brightness.

At the Observatorio del Roque de los Muchachos (ORM) on La Palma, where the Liverpool Telescope (LT) is located, and its successor, the New Robotic Telescope (NRT) will be located, the atmosphere presents a continuously changing environment for astronomical observations. While robotic telescopes are highly efficient in principle, their performance ultimately depends on the scheduler's ability to respond to these changes. The LT's existing scheduler, first implemented in 2004, was pioneering in its day but reflects the design philosophy of its time: pre-defined observation groups are ranked and the best ranking group is executed, with little or no adaptation to evolving conditions. After a group is finished, the ranking algorithm is repeated. What is needed for the coming decade is a new generation of dynamic schedulers that move beyond static ranking. Such systems must incorporate environmental awareness and be capable of making fully autonomous decisions about not only which target to observe but also how to observe it: including whether conditions are suitable, what exposure time is required to achieve a given signal-to-noise ratio, and how best to allocate limited observing resources across competing scientific programmes.

1.2 Motivation

The central motivation of this thesis is to address these limitations by developing methods that allow robotic telescopes to become more environment-aware. The overall vision is to support next-generation robotic scheduling systems that can integrate real-time environmental monitoring with predictive models and adaptive exposure strategies. By closing the loop between environmental sensing, predictive modelling, and the scheduler itself, robotic telescopes will be able to make decisions that maximise both efficiency and data quality.

This problem is timely and of high practical importance. The NRT, currently in design and planning, is expected to be the largest robotic telescope in the world. It will inherit both the opportunities and the challenges of robotic operation in a time-domain astronomy era. Among the most compelling opportunities are the discovery and characterisation of fast and rare astrophysical events: fast blue optical transients (FBOTs, e.g. Ho et al., 2020), kilonovae associated with neutron-star mergers (e.g. Smartt et al., 2017), the earliest phases of supernova explosions (e.g. Groh, 2014), and rapidly fading gamma-ray burst afterglows (e.g. Kann et al., 2010). In each of these cases, the most informative signatures, whether spectroscopic, polarimetric, or photometric, emerge on timescales of minutes to hours after the event is first detected, and may fade or change irreversibly within a single night. Capturing these fleeting signals requires not only rapid telescope response but also a scheduler that can autonomously evaluate feasibility, determine the required exposure strategy to achieve a target signal-to-noise ratio, and allocate resources in real time. Without such capabilities, robotic telescopes risk missing the very phenomena they are best placed to study, or producing data of insufficient quality to address the underlying science questions. Capturing these fleeting signals requires not only rapid telescope response but also a scheduler that can autonomously evaluate feasibility, determine the required exposure time to achieve a target signal-to-noise ratio, and allocate resources in real time. In this context, there is no time for humans to configure or adjust observation groups in response to a new alert: the delays inherent in human decision-making and manual intervention constitute a bottleneck that is fundamentally incompatible with the demands of modern time-domain astronomy. Ensuring that the scheduler itself can act intelligently in response to both scientific priorities and evolving environmental conditions is therefore essential if the facility is to meet its goals. The research presented here, developed and tested using the Liverpool Telescope as a platform, provides a foundation for this capability.

1.3 Research Focus and Scope

This thesis explores four interconnected lines of research, each addressing a critical gap in the atmospheric awareness of robotic telescopes. The first is the detection, tracking, and prediction of clouds. Using all-sky cameras, inexpensive and robust monitoring of cloud coverage can be achieved. In this work, algorithms are developed to detect clouds in real time, track their evolution, and predict their short-term trajectories, providing actionable forecasts on ten- to twenty-minute timescales that can help avoid wasted exposures.

The second line of research concerns atmospheric extinction. Accurate photometry requires continuous knowledge of atmospheric transparency, yet traditional extinction measurements rely on standard-star observations, which are sparse and unsuited to fully robotic systems. Here, new methods are developed to extract extinction coefficients automatically and continuously from auxiliary wide- and narrow-field imagers (Skycams), enabling near-real-time calibration and distinguishing photometric from non-photometric conditions.

The third component investigates sky brightness modelling. Background light levels are a crucial factor for exposure time calculation, particularly for faint targets. This thesis develops both empirical and machine-learning-based models of sky brightness at the ORM, correlating observations with lunar altitude, angular separation, airmass, and other environmental parameters. The resulting models not only provide real-time corrections but also support long-term exposure forecasts and planning tools.

Finally, the fourth strand of this research addresses signal-to-noise ratio based exposure optimisation. Traditionally, astronomers specify fixed exposure times, relying on assumptions about average conditions. This approach is inherently inefficient in a variable environment. In this thesis, an adaptive exposure model is implemented that calculates exposure times dynamically according to the desired signal-to-noise ratio, the target's altitude and the prevailing atmospheric conditions. By integrating real-time extinction and sky brightness data, the system ensures that exposure times are neither unnecessarily long nor insufficient to reach scientific goals, thereby improving both data quality and telescope efficiency.

Taken together, these four components form the building blocks of a comprehensive atmospheric-aware scheduling framework. Each is investigated as a standalone method, but the overarching goal is to demonstrate how they can be integrated into a unified system that can support the operation of both the Liverpool Telescope and the New Robotic Telescope.

1.4 Thesis Structure

The chapters of this thesis follow a logical progression, beginning with the development of real-time monitoring methods and culminating in their integration into a scheduling framework. The first scientific chapter introduces the cloud detection and prediction system, based on all-sky imaging. The following chapter turns to atmospheric extinction, presenting the methodology for extracting real-time extinction coefficients from Skycams T and Z (co-pointing cameras installed at the Liverpool telescope) and validating them against catalogue photometry and IO:O standard-star observations. The next chapter focuses on modelling sky brightness, where empirical corrections are developed and tested against independent datasets. Building upon these components, the final technical chapter integrates the results into a signal-to-noise ratio based exposure time system, tested under real-world observing conditions on the Liverpool Telescope, including case studies of standard stars, a science proposal, and transient follow-up with the SPRAT spectrograph. The concluding chapter summarises the findings, discusses limitations, and outlines how the methods developed here can be integrated into the scheduling system of the New Robotic Telescope.

1.5 Aims and Significance

The aim of this thesis is to demonstrate the feasibility of integrating atmospheric awareness directly into robotic scheduling. More specifically, it seeks to determine whether real-time cloud detection can reduce wasted observing time, whether automated extinction measurements can provide stable and continuous calibration data, whether predictive models of sky brightness can enhance exposure planning, and whether adaptive signal-to-noise based exposure calculations can optimise telescope usage. The significance of this work lies in showing that these elements, when combined, form the basis of a new operational paradigm for robotic observatories.

By developing and testing these methods on the Liverpool Telescope, this research provides a prototype for how the New Robotic Telescope can operate in the 2030s. The outcomes are not limited to these two facilities: any robotic observatory working in the era of large transient surveys faces the same challenges, and the approaches presented here can serve as a model for enhancing efficiency and data quality in a dynamically changing atmospheric environment.

Chapter 2

Background

> *Any sufficiently advanced technology*
> *is indistinguishable from magic.*
>
> Arthur C. Clarke

2.1 Introduction

The landscape of observational astronomy has undergone a transformation with the advent of robotic telescopes. These automated systems have redefined how astronomical data is collected, enabling rapid-response observations and significantly enhancing the efficiency of astronomical research. This shift from traditional, human-operated telescopes to autonomous systems has been particularly impactful in the study of transient astronomical phenomena - events that are often too short or unpredictable for conventional observatories to capture in time. Robotic telescopes have played an important role in observing gamma-ray bursts, gravitational wave counterparts, and supernovae. Events that demand immediate attention to unravel their underlying physics.

Within this evolving context, the Liverpool Telescope (LT, Steele et al., 2004) stands as a distinguished pioneer, demonstrating the power of fully autonomous operations and pushing the boundaries of robotic observational capabilities. The upcoming New Robotic Telescope (NRT, Copperwheat et al., 2015) promises to build upon the LT's legacy, ushering in a new era of advanced robotic astronomical research.

The Liverpool Telescope, located at the Observatorio del Roque de los Muchachos on La Palma in the Canary Islands, Spain, at an altitude of 2363 metres, is a prime example of a fully robotic astronomical observatory. This 2-meter aperture telescope utilises a

Ritchey-Chrétien optical design, known for producing high-quality images over a wide field of view. First light occurred in 2003, with routine science operations commencing in 2004. The telescope employs an alt-azimuth mount and can host up to five scientific instruments simultaneously on its Acquisition and Guidance box (one straight-through port, currently occupied by the IO:O imager, Smith et al., 2017). This configuration allows for rapid switching between observing modes, enhancing responsiveness to a wide range of scientific requirements.

A distinctive feature of the LT is its clamshell enclosure, which opens in two halves from the top down, providing an unobstructed view of the entire sky and allowing for rapid slewing to new targets. Weighing approximately 24 metric tonnes and measuring 8.5 metres in height and 6.5 metres in width, the LT is designed for autonomous operation even under challenging environmental conditions, including wind speeds up to 60 km/h. The entire structure rests on a concrete pier levelled to within 20 microns, ensuring stability and minimising vibrations.

The design and operational features of the Liverpool Telescope prioritise robotic functionality. The combination of rapid slewing, multi-instrument support, and autonomous decision-making marks a significant departure from traditional telescope operations. The LT was purpose-built for autonomous, fast-response observing, enabling effective monitoring of variable objects and prompt reaction to unpredictable phenomena.

Unlike traditional telescopes, which typically require on-site human operators and adhere to static, pre-determined schedules, the LT is fully autonomous. Observations are scheduled by a software scheduler that considers target visibility, basic weather conditions, and scientific priorities. The system can automatically interrupt its regular programme to observe high-priority transient events (Targets of Opportunity, ToO), such as gamma-ray bursts, when triggered by external alerts. Observation requests can be submitted via a Java client or through RTML (Remote Telescope Markup Language, Pennypacker et al., 2002), an XML-based standard for defining target details, exposure requirements, and constraints in a machine-readable format.

This operational flexibility allows the LT to maximise observing time and adapt in real time to evolving scientific goals. Its Robotic Control System (RCS, Fraser et al., 2004) dynamically reprioritises observations based on alerts, weather, or updated science objectives. While robotic telescopes reduce the need for on-site personnel and enable cost-effective operations, they are not without limitations: maintenance, upgrades, and critical interventions still require human oversight. Additionally, while automation accelerates response time, it may lack the serendipity of human observers reacting intuitively to unexpected phenomena, though machine learning and alert systems increasingly mitigate this.

To further increase scientific efficiency, robotic telescopes must continuously monitor and adapt to atmospheric conditions. Parameters such as cloud cover, sky brightness, and atmospheric extinction directly affect image quality and the feasibility of executing queued observations. Understanding and managing these parameters is especially important when observations must meet specific photometric requirements, such as achieving a defined signal-to-noise ratio (SNR) within limited time constraints.

This chapter provides an overview of the technologies and strategies used to monitor these atmospheric conditions and their integration into robotic scheduling. It reviews current and historical systems for cloud detection, extinction monitoring, and sky brightness modelling, and introduces the concept of SNR-driven exposure planning as a pathway toward more intelligent, efficiency-maximising robotic observatories. These topics form the conceptual foundation for the contributions developed in the following chapters.

2.2 Time-Domain Science with Robotic Telescopes

Robotic telescopes can deliver science by ‘converting’ discovery alerts and monitoring campaigns into timely, usable data. For the Liverpool Telescope (LT) and the New Robotic Telescope (NRT) the dominant drivers are: fast follow-up of rapidly evolving transients (GRB afterglows, Fast Blue Optical Transients (FBOT), kilonovae), high-throughput spectroscopic classification of survey alerts, and long-baseline monitoring programmes at nightly or near-nightly cadence (e.g. AGN reverberation mapping). Each driver imposes specific operational requirements which a static, elevation-weighted queue cannot meet alone. The methods developed in this thesis target some of these requirements.

2.2.1 The Transient Phase Space

To place the science drivers in context, it is useful to consider the transient phase space: a schematic representation of characteristic timescales against peak luminosities for different classes of astrophysical transients. This view illustrates both the historical focus of the Liverpool Telescope and the new regimes opened by the New Robotic Telescope. Figure 2.1 shows an adapted version of such a diagram (Copperwheat et al., 2015). The horizontal axis represents the typical duration of the event, expressed as a logarithmic timescale in days, while the vertical axis gives the peak absolute magnitude, with brighter sources appearing higher. Each class therefore occupies a characteristic region of the diagram.

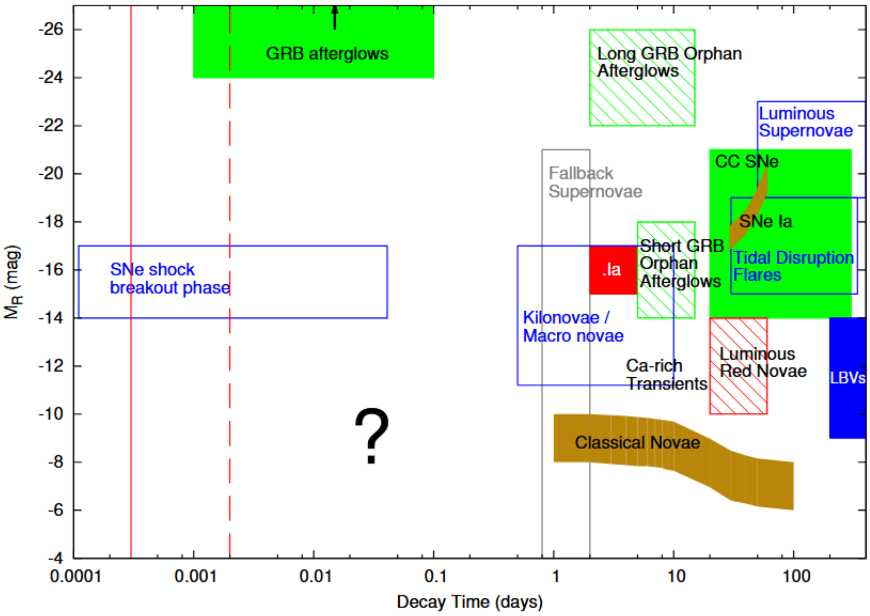

FIGURE 2.1: Illustrative phase space of astrophysical transients: characteristic timescale (log days) versus peak absolute magnitude (axis inverted; brighter upward). Classical novae and supernovae occupy the week-scale regime; FBOTs, kilonovae, and GRB afterglows evolve on day to hour timescales. The fast, high-throughput spectroscopic regime motivates NRT's design and the methods in Chapters 3–6. Adapted from Copperwheat et al. (2015).

Classical novae and supernovae, which motivated much of the early LT follow-up, evolve on timescales of weeks. In this regime, scheduling flexibility is less critical, since observations can be deferred without immediate loss of scientific value. By contrast, modern "fast" transients such as FBOTs (e.g. Perley et al., 2021), kilonovae (e.g. Tanvir et al., 2013), and gamma-ray burst afterglows (e.g. Dai et al., 1998) evolve an order of magnitude more quickly, in hours to days, and often fade below detectability on the timescale of a few nights. Spectroscopy in these early phases captures unique diagnostics, colour evolution in kilonovae, jet geometry and redshift for GRBs, or unusual spectral signatures in FBOTs. Features that cannot be recovered later. For these classes, latency rather than aperture is the limiting factor, and rapid, automated response becomes essential.

The diagram also illustrates the breadth of variability phenomena that fall between these extremes. Superluminous supernovae and luminous red novae evolve over many months, while active galactic nuclei vary stochastically on timescales of weeks to years. Such long-lived sources motivate stable monitoring programmes at regular cadence. Together, these regimes show that modern time-domain astronomy spans more than four orders of magnitude in both timescale and luminosity. No single observational strategy can address all of them; instead, facilities must be designed around the regions of phase space that align with their scientific priorities.

The NRT explicitly targets the fast, faint, and under-sampled part of this diagram. Its emphasis on sub-30 s response and high-throughput, low-resolution spectroscopy is intended to secure early classifications of large numbers of rapidly evolving transients, while still supporting longer-term monitoring through efficient scheduling. In this sense, the transient phase space not only contextualises past and current science, but also demonstrates why the methods developed in this thesis — cloud prediction, extinction monitoring, sky-brightness modelling, and SNR-based exposure calculation — are essential to exploit the NRT's capabilities in the fast-transient regime.

2.2.2 Fast transients (seconds to hours)

GRB afterglows fade on minute timescales (Groh, 2014), and the early spectral evolution of FBOTs and kilonovae unfolds over mere hours. These events define the fast end of the transient landscape and pose a clear challenge for robotic follow-up: the scientific leverage is highest in the first moments, yet traditional scheduling often introduces delays that are comparable to or longer than the timescales of interest. The NRT is therefore designed for sub-30 s response from trigger to exposure, with a focus on spectroscopic classification. Achieving this demands a scheduler that operates with minimal latency, computes exposures based on real-time conditions, and functions even when acquisition images are shallow or unavailable. This motivates the feasibility study in Chapter 4, which demonstrated that real-time extinction monitoring is possible using the co-pointing Skycams, provided suitable modifications are applied. Alongside this, the on-the-fly SNR-based exposure computation developed in Chapter 6 offers a practical system to integrating such environmental factors directly into scheduling decisions.

Among the fast transients, gamma-ray burst (GRB) afterglows remain a cornerstone: relativistic jets from compact-object explosions produce rapidly fading optical and near-infrared emission, often disappearing within minutes. Capturing early-time spectra, and, when available, fast polarimetry, reveals jet geometry, magnetic field structure, and redshift. In such cases, low-latency acquisition is as important as aperture. The NRT's response architecture continues the LT's legacy of rapid GRB follow-up, but with improved sensitivity and speed.

Kilonovae, arising from neutron-star mergers, exhibit their most informative colour and spectral evolution during the first 12 hours. Observations in this window can distinguish blue and red components, improving classification and association within the broad sky localisations of gravitational-wave events. For this reason, the NRT emphasises rapid-response spectroscopy.

Fast Blue Optical Transients, along with other recently uncovered faint-and-fast classes such as the rapidly evolving transients identified in Pan-STARRS1 (Drout et al., 2014) and in the Dark Energy Survey (Pursiainen et al., 2018), occupy a poorly explored regime of the transient phase space (see Figure 2.1). Their physical interpretation relies on capturing early spectra, ideally within hours of detection. However, manual setup or static queue scheduling often fails to meet this timescale. The NRT addresses this gap directly through its dedicated SPEC classification stream, designed to overcome the optical spectroscopy bottleneck identified by the transient community (Kulkarni, 2020).

Other phenomena, such as shock breakout and very early supernova phases, offer observational windows of less than an hour. Spectra taken during this time can reveal flash-ionisation features and circumstellar material, but these signatures fade rapidly. Any delay in execution forfeits key diagnostics, underscoring the need for scheduler-level observation setup and response.

Finally, the number of optical transients without detected γ-ray emissions, so-called orphaned afterglows, is expected to increase as wide-field surveys and X-ray missions expand their reach. In the Rubin-LSST era, alert streams will grow too large for human screening. Scientific return will depend not only on automation, but on how rapidly and intelligently the system can prioritise and respond.

Class	**Timescale**	**Peak M**	**Primary diagnostic**	**Latency need**
GRB afterglow	min–hrs	very bright	redshift, jet, polarimetry	$<$30 s to minutes
Kilonova	hrs–days	modest	early colour and spectra	$<$ 12 h (earliest best)
FBOT / faint&fast	hrs–days	bright	stectra for classification	minutes - hours
Shock breakout / early SN	hrs–day	bright	flash features (CSM)	hours

Table 2.1: Fast transient classes and scheduling implications.

2.2.3 Large-scale spectroscopic classification

Modern time-domain surveys such as ZTF, and soon LSST/VRO, generate hundreds of viable transient candidates per night, far more than can be followed up manually. Turning these alert streams into scientifically usable data at scale requires that classification spectra be both uniform and efficient: signal-to-noise ratios around 15 to 20 per resolution element in the continuum at low spectral resolution must be achieved reliably, despite ongoing changes in lunar background, airmass, and atmospheric extinction. This demands more than a static exposure strategy or per-target hand-tuning. Instead, a dynamic system is needed that responds in real time to the current state of the sky. This thesis addresses this requirement with an empirical sky-brightness model for the site (Chapter 5) and an adaptive exposure engine that computes SNR-matched integration times on a per-target basis (Chapter 6).

Within this context, the New Robotic Telescope formalises a high-throughput spectroscopic classification stream - *SPEC-Time*, as one of its key science modes. SPEC-Time is designed as a semi-autonomous programme that transforms alerts into uniform, low-resolution spectra. The operational concept assumes short exposures and aims to deliver up to 120 spectra per clear night at $R \approx 360$. The annual target is on the order of 10,000 classifications, which means that exposure times cannot be static or manually optimised; they must instead be determined at execution time using live sky data.

To enable this, the spectrograph is designed around a small-field integral-field unit (IFU), eliminating the ∼2-minute slit-placement loop that would otherwise dominate overhead. This supports the sub-30 s on-source latency goal, while maintaining high throughput across the 3750–7500 Å range. In addition, the IFU design is advantageous for classifying bright transients in crowded or high-background fields, since it allows local background subtraction and stabilises classification quality at low SNR.

Maintaining uniform classification quality across varying sky conditions is only possible if lunar brightness, airmass, and extinction are incorporated into the exposure decision. This thesis provides the necessary tools or feasibility analysis: a calibrated sky-brightness model (Chapter 5) and a real-time extinction signal from external, co-pointing Skycams (Chapter 4), decoupled from user science frames. Together, they ensure that the classification SNR remains predictable at run time, regardless of when a target is observed.

Throughput and overhead efficiency are equally critical. With exposures on the order of three minutes and short imaging setup sequences, the SPEC programme must minimise wasted time. The IFU acquisition model, paired with a scheduler that scores targets by expected science return per unit time, enables high nightly yield. Chapter 6 presents the SNR-driven exposure logic that replaces conservative padding, ensuring time is not lost to worst-case assumptions.

To fully automate this process, the SPEC programme must integrate directly with brokers such as Lasair (Smith, 2019), ALeRCE (Förster et al., 2021), or the Transient Name Server (Gal-Yam, 2021). Ingestion, exposure calculation, and execution must occur without human intervention, with spectra reduced within minutes for same-night follow-up. The exposure-time inversion logic developed in this thesis enables this pipeline: combining broker magnitudes with current sky diagnostics to generate an executable, condition-aware plan.

Finally, many scientifically valuable transients lie on bright host galaxies or in crowded fields. Achieving reliable classification in such environments depends not only on instrumental design, but also on accurate estimation of the local sky background and extinction. The IFU format helps mitigate host contamination, but exposure must still be matched to the real conditions. This again underscores the need for robust, real-time inputs to the SNR model—precisely the capability provided by the tools developed in this work.

To clarify how the methodological developments in this thesis map onto the operational requirements of large-scale transient classification at NRT, Table ?? summarises the key observational parameters of the proposed SPEC mode and links them directly to the enabling capabilities developed in Chapters 4–6. The table is not a performance validation, but an architectural mapping: it shows how each scientific or operational constraint (spectral resolution, cadence, classification depth, alert volume, and annual yield) depends on specific real-time environmental inputs and adaptive exposure control mechanisms introduced in this work.

Parameter / Constraint	**Typical value / goal**	**Capability from this thesis**
Spectral setup	$R \approx 360$, 3750–7500 Å, IFU 12×12 arcsec	Condition-agnostic SNR targeting (Chapter 6)
Exposure cadence	∼180 s per spectrum; ∼120/night	Dynamic exposure per target (Chapter 6)
Classification depth	SNR≈10–20 at low R	Real-time extinction + sky model (Chapters 4, 5)
Alert volume	Hundreds/night (ZTF→LSST)	Broker-to-scheduler path; auto inversion to SNR (Chapter 6)
Annual yield	∼10k classifications (to $r \approx 20.5 mag$)	Efficiency from eliminating worst-case padding (Chapter 6)

Table 2.2: SPEC classification at NRT: illustrative parameters and required capabilities.

In short, large-scale classification is not just “more targets”; it is a regime where *uniform data utility* and *overhead discipline* determine scientific output. The SPEC concept sets the science bar (SNR at low R, high nightly yield); the methods in this thesis - real-time extinction, an empirical sky model, and a dynamic SNR engine - are what make that bar reliably achievable under real sky variability at the cadences and volumes envisioned for NRT.

2.2.4 Monitoring programmes (cadence, short visits)

Active galactic nuclei (AGN) and tidal disruption event (TDE) programmes require a regular cadence over many weeks or months (up to several years) with relatively short visits (photometry, one to two images per filter). Variable extinction and Moon background can produce uneven SNR and hence uneven data utility if left unaccounted for. This motivates predictive sky models and extinction integration so visits are placed when efficiency is highest, or exposure is adapted to maintain the requested SNR (Chapters 4, 5, 6).

AGN reverberation mapping (e.g. Cackett et al., 2018) typically requires nightly or every-other-night cadence, where stable SNR is preferred over fixed exposure time to reduce flux-calibration scatter in light curves and improve line-to-continuum lag measurements. TDE follow-up spans weeks to months of spectro-photometric monitoring; early spectra evolve on day-to-week timescales, while photometry benefits from consistent SNR to track colour and temperature evolution. Other variable sources, including cataclysmic variables, compact binaries, and nuclear transients, likewise rely on short visits over long baselines, often constrained to windows near meridian passage and at minimal airmass.

Across a lunation the lunar background varies by orders of magnitude, and extinction can change on hour-timescales (e.g. caused by Calima). This thesis provides a site-specific sky-brightness model (Chapter 5) and a real-time, user-agnostic extinction signal (Chapter 4) to keep visit SNR close to the requested target. Many programmes allow a tolerance window, within this window the scheduler should either keep the exposure fixed and choose the most efficient time, or keep the time fixed and adapt the exposure to hit the SNR. A dispatcher selects whichever option minimises cost while meeting the SNR requirement. For very short integrations, overheads dominate; SNR targeting still prevents severe under- or over-exposure on bright or dark nights, but absolute efficiency gains are bounded by the overhead. This policy is applied in the scheduling strategy below.

In the case of AGN reverberation mapping, which typically requires nightly or every-other-night cadence with short photometric visits of one or two images per filter, the optimal strategy is to schedule the observation within the cadence window at the time of lowest predicted exposure based on sky brightness (and extinction). If the cadence window is tight (e.g. for low-altitude targets), the visit can still be executed at the planned time but with the exposure adapted to maintain the requested SNR using the live zero point and sky model.

For TDE monitoring (e.g. Bricman et al., 2020), which usually follows a cadence of two to three days and combines short spectra with photometric observations, the spectroscopic exposures are best calculated with the SNR engine using real-time extinction, while the photometry is placed away from peaks in lunar background or airmass whenever possible. If the monitoring window is restricted, the visit can still proceed by adjusting the exposure length to recover the target SNR.

For other repeating transients, where nightly short visits of only a few minutes are common, the limiting factor is often overhead rather than integration time. In this case, scheduling should prioritise the most efficient time within the cadence window, while for slightly longer exposures the time may be fixed and the exposure adapted to preserve the target SNR under variable conditions. Monitoring programmes therefore benefit directly from the same real-time environmental inputs and adaptive exposure logic that support fast transients and classification, ensuring stable data quality across extended campaigns.

Together, these science cases define the performance envelope that the NRT is built to meet: fast response to rapidly evolving transients, high-throughput spectroscopic classification of large alert streams, and stable monitoring of variable sources at nightly or near-nightly cadence. This thesis investigates and develops the enabling technologies required to make such a regime viable in practice. These include short-horizon cloud prediction (Chapter 3), real-time extinction monitoring independent of science frames (Chapter 4), an empirical sky-brightness model (Chapter 5), and a dynamic SNR engine that computes exposures at execution time (Chapter 6). Because acquisition images may be shallow, narrow, or omitted entirely, extinction must be determined independently of user frames and delivered continuously from auxiliary imagers. Exposure times are then calculated dynamically from these inputs so that the scheduler can execute with uniform SNR, efficient throughput, and principled trade-offs between immediate and deferred execution under variable conditions.

2.3 The Atmosphere over La Palma

The Liverpool Telescope is located on the island of La Palma, the westernmost island of the Canary Islands. The telescope is part of the Observatorio del Roque de los Muchachos (ORM) which is is close to the highest elevation of the island, located on the rim of a large volcanic caldera (Caldera de Taburiente) at an altitude of 2369m above sea level. The geographic coordinates of the ORM are $28°45'35''N$, $17°53'24''W$. According to night logs of the Liverpool Telescope, over 70% of the nights at the ORM provide conditions for astronomic observations. This number can reach up to 95% in summer nights. Up to 80% of those nights provide photometric conditions. The atmosphere and the meteorology of the site has been extensively monitored and analysed over the past 50 years (Kiepenheuer, 1972, Murdin, 1985, Pedani, 2004, Lombardi et al., 2006, Lombardi et al., 2007, Lombardi et al., 2008, Varela et al., 2008, Muñoz-Tuñón et al., 2015, Gaug et al., 2024). Assuming an average of ∼4300 hours of astronomical night per year, based on the previously named publications, at the ORM, a usable-night fraction of 70% corresponds to approximately 3000 hours of potential observing time annually. If up to 80% of those nights are photometric, this implies roughly 2400 hours per year under fully transparent conditions. These values place the ORM among the leading mid-latitude optical observatory sites, with long-term usability broadly comparable to other premier facilities in Chile and Hawaii, although exact fractions depend on the adopted definition of "usable" and "photometric" conditions. The ORM holds a position as a premier astronomical site globally. This reputation is largely attributed to the exceptional observing conditions caused by a combination of geographical and meteorological factors. The atmosphere around La Palma is characterised by a maritime mixing layer (MML), extending from sea-level to about 1500m. This layer contains mostly moist air and a stratocumulus cloud layer at the top. The MML is well-mixed due to trade winds which usually blow from the North-East direction (the Azores), solar heating and the island's geography. Above the MML is a temperature inversion layer (called "sea of clouds"), confining the clouds underneath. The temperature inversion layer is very stable and quasi-permanent throughout the year. It is thinner and higher (residing between 1400m and 1850m a.s.l. with a thickness of 350m) during the winter and thicker but lower in summer months (between 750m and 1400m above sea level (a.s.l.) with a thickness of about 550m; Font Tullot, 1956, Torres et al., 2001, Carrillo et al., 2016) and is mainly caused by the trade winds and the descending branch of the Hadley cell (Longo et al., 2019). Above the temperature inversion layer, the atmosphere is usually dry and behaves very stable, resulting in the high number of nights with photometric conditions and good seeing values.

2.3.1 Clouds

Although the ORM provides an exceptional high number of nights with photometric conditions, interruptions of observations due to bad weather are not uncommon. As shown in Figure 2.2, wind blowing from southern or western directions cause relatively more rain (and thus clouds) than wind coming from the eastern directions. The data is based on the accumulated measurements at the Liverpool Telescope from 2007 to 2025.

The percentages shown in Figure 2.2 are grouped by prevailing wind direction sector (E, NE, N, NW, W, SW, S, SE) as recorded in the LT weather database. An "observation" in this context refers to a completed scheduled exposure sequence (including associated overheads such as readout and slewing), rather than a fixed time bin. Observations at the LT typically lasts between 1 to 10 minutes but this metric is highly instrument- and

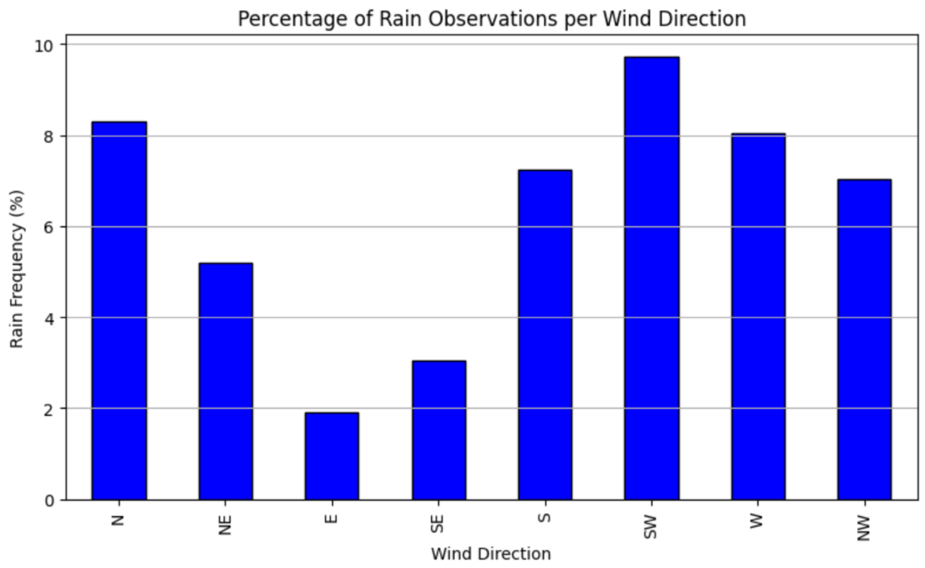

(a) Rain Frequency per Wind Direction

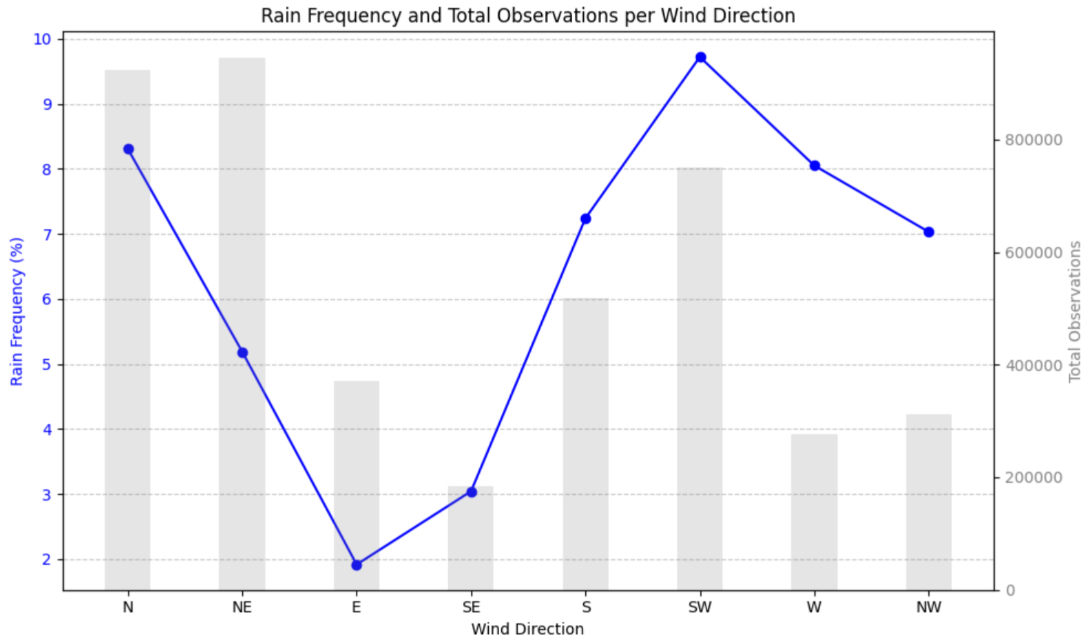

(b) Total Observations per Wind Direction

Figure 2.2: Relationship between wind direction and rain occurrence

programme dependent. It is also noted that East-wind usually carries a lot of Sahara dust (the phenomena is called "Calima"). A similar pattern is caused by "high humidity" events (see Figure 2.3). These are humidity levels above 70% where it is unsafe to operate the telescope as condensation will affect the optical and electronical parts of the system.

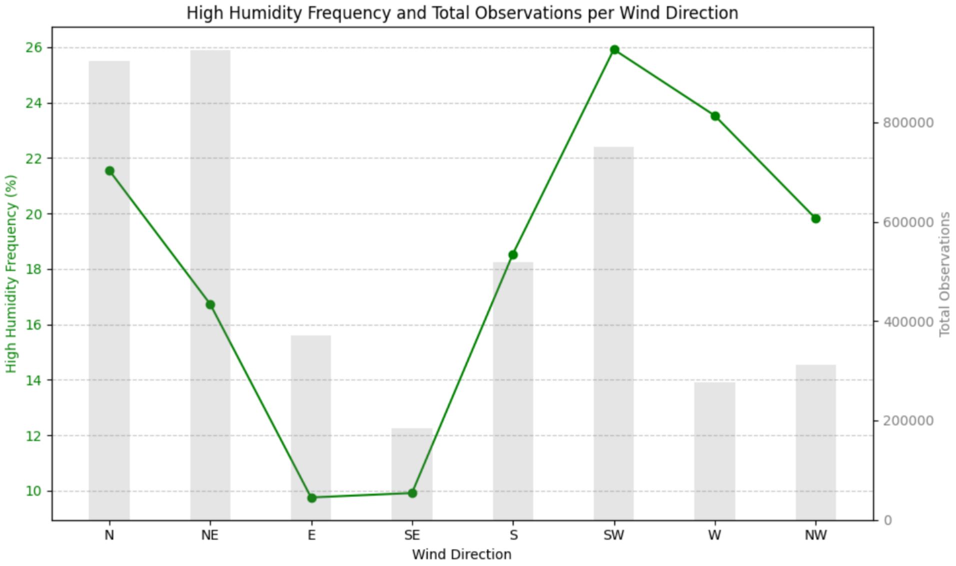

FIGURE 2.3: Humidity events with Humidity > 70% at the LT and Wind Direction

Both plots show that wind blowing from eastern directions causes dry conditions - in accordance with the literature mentioned before. The relatively low total number of events from eastern and south-eastern directions can be explained with the local geography. The Mercator Telescope and the Isaac Newton Telescope are located in eastern directions of the LT site. Also, the incline of the caldera is to the south-east of the LT site, both causing a "wind shadow". The relevant cloud levels for the system described in this thesis are clouds above 3000m a.s.l, so called high-level clouds. Clouds below 3000m usually cause high humidity at the ORM and thus no observations will be made in the first place.

High-level clouds can occur throughout the year. According to recent studies (e.g. Fromm et al., 2016), high-level Calima can trigger cloud formations at the Canary Islands. Calima is a wind blowing from eastern directions transporting huge amounts of Saharan dust to the Canary Islands. It is usually found in the lower levels but can reach heights similar to the ORM and beyond, causing astronomic observations to stop to protect the equipment or - if in even higher levels - increasing the atmospheric extinction. As Calima events occur more frequently during the summer months (Barreto et al., 2022), the fraction of nights affected by high-level dust layers and associated thin cloud structures is correspondingly higher than during winter times.

Operational statistics from the Liverpool Telescope night logs indicate that approximately 70 to 80% of nights at the ORM provide conditions suitable for astronomical observations, with seasonal variations reaching up to ∼95% during the summer months. When expressed in operational terms, this corresponds to the successful execution of the large majority of scheduled observations, with interruptions primarily caused by high humidity and related ice covering in winter, strong winds, or Calima-related dust events.

2.4 Cloud Detection in Astronomical Observatories

2.4.1 All-sky Camera Systems

Cloud detection plays a central role in robotic telescope operations, particularly for safeguarding instruments and optimising science return. One of the most widely adopted approaches to local sky monitoring is the use of all-sky cameras: imaging systems with wide-angle or fisheye lenses that capture the entire visible sky dome at regular intervals.

Several major observatories have implemented such systems. The SkyProbe system (Cuillandre et al., 2002) at the Canada-France-Hawaii Telescope (CFHT) was one of the earliest continuous optical all-sky monitors used for cloud detection and extinction measurements. The CONCAM project (Continuous Camera, Pereira, 2003) pioneered all-sky monitoring at multiple observatory sites simultaneously. These systems provided early examples of real-time monitoring, though they required manual interpretation to identify clouds.

In contrast, many smaller observatories rely on simpler cloud sensors based on infrared sky temperature measurements rather than full-sky imaging. These instruments measure the difference between ambient temperature and the effective radiative temperature of the sky; the presence of clouds typically raises the measured sky temperature due to increased infrared emission from water droplets. While such sensors provide a robust and inexpensive indicator of overall cloudiness, they do not offer good spatial information about the sky and therefore cannot identify clear regions within partially cloudy conditions.

More recently, instruments such as the Night Sky Live project (Pérez-Ramírez et al., 2004) and the MASCARA Project for finding exoplanets using all-sky cameras (Talens et al., 2017), have advanced the concept by incorporating real-time data pipelines and integrated alert systems.

2.4.2 Existing Methods for Night-Time Cloud Detection

The problem of automatic cloud detection from ground-based observatories is an active area of research, with diverse approaches proposed in recent years, particularly for night-time conditions relevant to astronomical operations. Most existing methods fall into three broad categories: star-counting techniques, image segmentation approaches, and machine-learning-based classifiers.

One widely used approach is to identify clouds by comparing the number of detected stars against a reference star catalogue. Adam et al. (2017), for example, detect clouds in all-sky images by identifying areas with a significant drop in the number of visible stars relative to expectations. This method performs well under dark-sky conditions and has the advantage of leveraging well-understood catalogues. However, it struggles during gibbous or full moon phases when overexposure and lens flare effects obscure large portions of the star field. Furthermore, this technique is sensitive to the limiting magnitude of the camera system and can misclassify regions where only bright stars remain detectable.

Alternative methods have focused on pixel-level or region-level segmentation of the sky image. Dev et al. (2017) propose a superpixel-based approach, in which the image is first partitioned into spatially coherent clusters of neighbouring pixels with similar colour and intensity properties. These “superpixels” are then classified as cloud or clear sky based on local statistical features. The method assumes that cloud structures exhibit distinct local contrast or texture relative to the background sky. However, under non-uniform illumination — particularly in the presence of the Moon or bright stars — strong radial intensity gradients are introduced across the image. These smooth gradients can artificially increase local contrast within superpixels, causing illumination variations to be misinterpreted as cloud boundaries. As a result, purely local segmentation approaches may incorrectly fragment clear sky regions or falsely identify bright gradient transitions as cloud structures.

More recently, machine learning techniques have been explored for cloud detection from all-sky images. Mommert (2020) evaluated two approaches: a convolutional ResNet model, and a feature-based gradient-boosted decision tree model (LightGBM). Their preferred solution, LightGBM achieves 95% accuracy on binary cloud detection performed on fixed spatial subregions of the image. Here, “subregions” denote predefined spatial tiles obtained by partitioning the all-sky image into a circular grid. Each tile, typically spanning several hundreds of pixels, is classified independently as cloud or clear sky. Their system uses a training set of approximately 1,000 manually labelled images.

However, the labelling process proved to be both labour-intensive and inconsistent, particularly under full-moon conditions, where overexposed sections with reduced contrast and faint sky features such as the Milky Way are easily mistaken for clouds. Mommert found that human annotators often struggled to interpret ambiguous regions reliably, and emphasised the need for explicit training of "labelers" to ensure high-quality annotations.

These observations are consistent with the results of a student annotation experiment conducted during the author's MSc thesis, which laid the groundwork for the cloud detection approach developed in this PhD project. Even with a brief training session, participants frequently misclassified low-contrast regions and bright astronomical features. This highlights the practical challenges of generating consistent labels, especially in scenarios involving rapid deployment or where annotators have limited domain expertise.

Additionally, the classifier in Mommert (2020) operates on a fixed spatial grid of 33 predefined image sectors aligned with the camera's orientation. For each sector, a set of hand-engineered statistical features — including colour ratios, mean intensity, and local contrast measures — is computed and provided as input to a LightGBM classifier, which outputs a binary label (cloud or clear sky) for that sector. The method therefore measures sector-level photometric and colour statistics rather than performing pixel-wise segmentation or modelling physical cloud properties. Because the sector geometry is hard-coded, any change in the imaging setup, such as camera rotation, replacement, or lens adjustment, alters the mapping between sky coordinates and sector indices, requiring re-annotation and retraining. Although the computational demands of LightGBM are relatively modest, these structural constraints limit portability and long-term maintainability in autonomous observatory environments.

Ye et al. (2022) introduced a self-training approach for daytime cloud detection using superpixel segmentation and semi-supervised learning. Although this reduces the need for full manual annotation, the algorithm relies heavily on colour features in visible light and thus is not applicable under night-time conditions with monochrome or low-contrast imaging.

Usually, existing methods either require strict lighting conditions (e.g., darkness), substantial calibration or training effort, or fail under the varying illumination introduced by moonlight. Furthermore, some are not designed for real-time execution at high cadence, as needed for robotic observatories. These limitations motivate the development of the method presented in this thesis, which leverages difference imaging and classical image processing to achieve robust, calibration-light, and computationally efficient cloud detection even under bright moonlit conditions.

In the context of this work, an all-sky camera system is used to derive both cloud masks and cloud motion vectors, enabling short-term forecasting of sky conditions. Unlike many traditional systems that rely on external calibration or auxillary infrared data, the system developed here operates in a purely optical domain with minimal calibration requirements, and forms a foundation for automated weather-aware scheduling. The technical implementation is detailed in Chapter 3.

2.4.3 Motion-based Prediction

Cloud tracking is also an active area of research with diverse applications across weather forecasting, the energy sector (e.g., photovoltaic energy prediction in smart grids), climate science, and astronomy. In the context of astronomical observatories, real-time cloud motion prediction is particularly valuable for short-term scheduling decisions and condition-aware observation planning.

One common approach to cloud tracking is based on optical flow, a computer vision technique that estimates the apparent two-dimensional velocity field of brightness patterns between consecutive image frames under the assumption of local intensity conservation. In its dense formulation, a motion vector is computed for nearly every pixel in the image, producing a spatially continuous flow field. This technique has been applied in both daytime and nighttime imaging. For instance, Zhang et al. (2019) use dense optical flow on high-cadence all-sky image sequences. However, this method requires a frame rate of at least 20 all-sky images per minute to yield physically meaningful results. However, this method requires a frame rate of at least 10 to 20 all-sky images per minute to yield physically meaningful results. This high cadence is necessary because optical flow assumes that image structures change only minimally between frames. In practice, cloud morphology evolves continuously; if the time interval between frames is too large, structural changes in the clouds violate the brightness-conservation assumption and lead to artefacts or physically implausible motion vectors.

Alternative methods use multiple cameras distributed across a large baseline to reconstruct cloud structures in three dimensions. Peng et al. (2015), for example, describe a system that combines several wide-angle cameras to produce a volumetric cloud model, including estimates of cloud base height. This approach enables highly accurate spatial tracking but is not applicable at high-altitude observatory sites such as the ORM, where physical constraints make it impractical to deploy and maintain a wide-area camera network.

Several studies have attempted to categorise and evaluate the wide range of cloud tracking techniques available. Zaher et al. (2017) conducted a comparative analysis of classical motion estimation methods, including block matching, gradient-based optical flow (e.g., Horn-Schunck), phase correlation, and statistical correlation techniques. Their study found that no single method universally outperforms the others; instead, the effectiveness of a given algorithm depends heavily on the temporal and spatial resolution of the input data, the variability of cloud structures, and the intended application. In general, methods like phase correlation and optical flow perform well under continuous illumination and smooth motion but degrade rapidly when applied to sparse (i.e. low-cadence), intermittent, or structurally evolving inputs - such as those common in nighttime all-sky imaging.

Building on this, Arrais et al. (2022) provide a systematic literature review focused specifically on short-term cloud motion forecasting using ground-based sensors. They classify the available methods into three broad categories: model-based approaches (e.g., based on cloud advection models), learning-based approaches (e.g., using neural networks or support vector machines), and hybrid methods that combine physical modeling with data-driven techniques. Their review highlights a common trade-off between accuracy and generalisability: learning-based methods can perform well in controlled or data-rich environments, but they often require extensive retraining when camera configurations or environmental conditions change. Furthermore, many of the reviewed systems are designed for daytime operation and rely on colour or thermal features not available in monochrome, low-light, or moonlit night-time imagery.

Taken together, these studies highlight the challenges of applying traditional motion estimation or forecasting methods to nighttime all-sky cloud tracking. Real-time deployment at observatory sites requires techniques that are robust to changing illumination, tolerant of sparse temporal sampling due to long exposure times, and capable of operating with minimal calibration.

To address these constraints, the approach developed in this thesis adopts a geometry-based tracking method applied to deprojected cloud features extracted from low-cadence all-sky images. Motion is estimated from the trajectories of detected cloud structures and smoothed using a Kalman filter to produce short-term forecasts. The details of this implementation are described in section 3.4.

2.4.4 Infrared and Satellite Data Comparison

In addition to localised optical imaging, many observatories utilise infrared (IR) cloud sensors (Maghrabi et al., 2009) and satellite-derived cloud products (McNally et al., 2003) as complementary sources for assessing sky conditions. Ground-based IR cloud sensors operate by measuring the apparent radiative temperature of the sky: under clear conditions the sky appears extremely cold in the infrared, whereas the presence of clouds increases the measured temperature due to thermal emission from cloud layers. Satellite-based systems apply similar radiometric principles but observe the atmosphere from above, where high cloud tops appear colder than the underlying surface and lower atmosphere.

Geostationary satellites such as EUMETSAT's Meteosat and NOAA's GOES series routinely produce full-disk thermal IR imagery with cadences of 5–15 minutes and spatial resolutions of approximately 1–5 km per pixel. These data are widely used for weather forecasting and have also been incorporated into observatory monitoring systems. For example, satellite-derived cloud products are used as part of the operational weather monitoring at facilities such as the Liverpool Telescope and the Las Cumbres Observatory network, where they complement local sensors in determining whether observing conditions are safe or scientifically usable.

However, satellite IR data come with notable limitations. The spatial resolution is generally too coarse to resolve local cloud structures above a specific telescope, particularly in mountainous terrain where cloud layers may form or dissipate rapidly. Furthermore, low-altitude warm clouds or fog may remain undetected due to their weaker temperature contrast with the ground or ocean below. Additionally, satellite images provide a top-down view that may not align well with a telescope's actual field of view.

In practice, robotic observatories like the Liverpool Telescope have historically relied on satellite overlays for cloud detection and weather safety. However, during periods without a working satellite feed, alternative ground-based solutions - such as the all-sky system described in Chapter 3 - become essential. These systems offer higher spatial and temporal resolution, and can detect transient or partial cloud cover that would go unnoticed in coarser IR data.

2.5 Atmospheric Extinction Monitoring

2.5.1 Overview

Atmospheric extinction, the reduction in the apparent brightness of celestial objects as their light traverses the Earth's atmosphere, stands as a fundamental consideration for ground-based astronomical observations. This phenomenon arises from the interplay of absorption and scattering processes in the atmosphere that change the flux of photons reaching the observer. Accurate quantification and correction for atmospheric extinction are important in achieving good photometric measurements, enabling meaningful comparisons of astronomical data acquired at different times and from various locations. The degree to which the atmosphere dims starlight directly impacts the precision of measurements aimed at determining the properties of celestial objects, such as their luminosity and distance.

The Earth's atmosphere, a complex mixture of gases, aerosols, and water vapour, interacts with incoming electromagnetic radiation from space through absorption and scattering. Absorption involves the conversion of light energy into other forms of energy, often heat, by atmospheric constituents such as ozone in the ultraviolet spectrum and water vapour in the infrared. Scattering, on the other hand, redirects photons from their original path. Rayleigh scattering, caused by air molecules, exhibits a strong wavelength dependence, preferentially scattering shorter (blue) wavelengths, which explains the blue colour of the daytime sky. Aerosols, including dust, pollutants, and other particulate matter, can both absorb and scatter light across a broader range of wavelengths, and their concentration in the atmosphere is subject to significant temporal and spatial variations (Mie-Scattering, see Tüg et al., 1977). The interplay of these processes dictates the overall transparency of the atmosphere at any given time.

The Airmass (X) provides a convenient measure of the relative atmospheric path length encountered by light from a celestial source. It is defined relative to the optical path at the zenith, where $X = 1$. For a celestial object at zenith angle z, the airmass can be approximated by the secant of the zenith angle, $X \approx \sec(z)$. This approximation holds reasonably well for zenith angles up to approximately 60°. At larger zenith angles, the curvature of the Earth and the vertical structure of the atmosphere require more accurate formulations to describe the effective optical path. Examples of such formulations are given by Young et al. (1967), Kasten et al. (1989), and Pickering (2002).

In observational astronomy, airmass is primarily used as a convenient parameterisation of atmospheric extinction, allowing the attenuation of starlight to be expressed as a function of viewing geometry.

Atmospheric extinction exhibits a strong dependence on the wavelength of light. Shorter wavelengths generally experience greater extinction due to the nature of Rayleigh scattering and absorption by various atmospheric molecules. Consequently, the atmospheric extinction coefficient (k), which quantifies the amount of extinction per unit airmass, is a function of wavelength, $k(\lambda)$. When employing a monochrome camera without filters, the measured instrumental magnitude represents an integrated response across the camera's sensitivity range. The effective wavelength of this integrated response is influenced by both the spectral sensitivity of the camera's sensor and the spectral energy distribution of the observed star. This implies that the effective wavelength might vary slightly depending on the colour of the star being observed.

The basic photometric equation establishes a linear relationship between a star's catalogue magnitude ($m_{catalogue}$), its instrumental magnitude ($m_{instrumental}$), the atmospheric extinction coefficient (k), the airmass (X), and the photometric zero point (ZP):

$$m_{catalogue} = m_{instrumental} + k * X + ZP \tag{2.1}$$

Equation 2.1 serves as the foundation for determining the atmospheric extinction coefficient from photometric observations of standard stars at varying airmasses. The photometric zero point (ZP) is an important parameter that accounts for the overall sensitivity of the observer's specific telescope and camera system, effectively calibrating the instrumental magnitude scale to the standard magnitude system for a given observation night (Chromey, 2010).

A variety of techniques have been developed to monitor extinction, ranging from classical standard-star observations to modern real-time sensors and data-driven calibration methods, each with its own advantages and limitations.

2.5.2 Classical Photometric Extinction

The traditional approach to determining atmospheric extinction is through photometry of standard stars at different airmass. By observing one or more standard stars (with known intrinsic magnitudes) over a range of altitudes, a fit of the observed magnitude as a linear function of air mass is used to derive the first-order extinction coefficient (Landolt, 1983 and Landolt, 1992). This so-called Bouguer or Beer's law method yields a nightly zero-point and extinction value for each filter. Classical extinction determination assumes that atmospheric transmission is uniform across the sky and stable over the night, so that the magnitude-airmass relation holds true. While robust on clear photometric nights, this method can be less reliable if conditions are variable or if spatial variations (e.g. patchy clouds or aerosols) are present. However, dedicating observing time to multiple standard stars incurs a cost to science time. Nevertheless, the standard-star technique has been the foundation of photometric calibration for decades and is implemented at most observatories (see Stubbs et al., 2006 for a review of photometric calibration methods).

2.5.3 Real-time Extinction Measurements

Modern observatories increasingly rely on dedicated monitors to measure atmospheric conditions in real time, including transparency. For example, combined MASS-DIMM units (Multi-Aperture Scintillation Sensor + Differential Image Motion Monitor) continuously observe bright stars to measure atmospheric seeing and can also track the flux attenuation, providing an estimate of the extinction throughout the night (Kornilov et al., 2010). The RoboDIMM system (Augusteijn, 2001) similarly offers automated seeing measurements at frequent intervals and can flag transparency dips when clouds or haze reduce a star's intensity. Another example is the calibration telescope for the Vera C. Rubin Observatory (VRO/LSST, Ivezić et al., 2019), a 1.2 m telescope (AuxTel, Ingraham et al., 2020) with a spectrophotometric instrument (LATISS) dedicated to measuring the atmospheric transmission on the same site in parallel with the main survey. These real-time data streams allow observers or automated pipelines to apply extinction corrections on the fly and to identify non-photometric periods immediately. By comparing simultaneous extinction readings from multiple instruments or sites, cross-calibration techniques can be used to maintain a consistent photometric scale across several telescopes.

2.5.4 Catalogue-Based Matching

An alternative strategy for extinction determination is to use external all-sky photometric catalogues as a reference for zero-point calibration. Large star catalogues such as Tycho-2 (Høg et al., 2000), APASS (Henden et al., 2018a), and more recently the Gaia catalogue (Gaia Collaboration et al., 2016), provide accurately calibrated magnitudes for millions of stars across the sky. By matching stars observed in a science exposure to their catalogued magnitudes, it is possible to infer the photometric zero-point offset and thus deduce the total extinction along that line of sight at the time of observation. In practice, if a field contains a sufficient number of catalogue stars of known brightness and colour, a simple fitting of the difference between instrumental magnitudes and catalogue magnitudes (as a function of air mass or colour) yields both the atmospheric extinction and the instrumental zero-point (Sterken et al., 1992). This catalogue-matching approach is widely used in automated survey pipelines to correct photometry on a frame-by-frame or nightly basis. It has the advantage of requiring no dedicated standard-star observations, since the calibration references are present in the science images themselves. However, its accuracy depends on the quality of the reference catalogue and the assumption that any systematic differences (such as colour term mismatches between the instrumental filter and the catalogue's passband) are either small or can be corrected separately. Still, with modern all-sky surveys providing dense grids of calibrator stars (for example, Gaia's

G-band photometry is internally calibrated to the millimagnitude level, see Gaia Collaboration et al., 2021), catalogue-based zero-point derivation has become a powerful tool for real-time extinction and zero-point monitoring in observational astronomy.

2.5.5 Differential Extinction Estimation

Beyond absolute calibration against standard catalogs or stars, one can also estimate extinction differentially by analyzing the observed magnitude residuals of many stars as a function of airmass. In this approach, the difference between an observed star's instrumental magnitude and its expected magnitude (from a catalog or from an assumed baseline) is tracked, and the trend of these residuals versus air mass is used to solve for the extinction coefficient. Essentially, if multiple stars are observed in the same frame or at close times but at slightly different air masses, their relative dimming can reveal the incremental effect of atmospheric path length. Plotting magnitude residuals against air mass and fitting a line yields an extinction slope without strictly requiring a prior absolute calibration for each star. This differential method can be applied on a single wide-field image that contains stars over a range of elevations, or across a set of images taken as a field traverses different air masses. The advantage of using wide-field cameras is that they capture numerous stars simultaneously, providing a large sample to average out measurement noise and temporal fluctuations. Wide-field observations may even cover enough sky area to include stars at notably different zenith angles in one shot, enhancing the leverage for an extinction fit. Meanwhile, employing a “near-field” reference - i.e. calibration stars observed in nearly the same direction and at the same time as the science target - ensures that local atmospheric effects (clouds or gradients in extinction across the sky) are nearly identical for both target and reference. By combining the breadth of wide-field measurements with the locality of near-field comparisons, this differential extinction estimation can yield robust, real-time extinction values tailored to the specific observation. This leverages the internal consistency of star brightness within each exposure (or set of exposures) to continuously solve for atmospheric transparency variations. This approach is conceptually similar to the self-calibration used in large surveys, which fit for zero-point and extinction across many overlapping fields (Padmanabhan et al., 2008), but here it can be implemented on a more immediate timescale and even on individual frames. The result is an ability to correct for atmospheric extinction dynamically, improving photometric accuracy without exclusively relying on traditional standard star observations.

2.5.6 Atmospheric Extinction at the ORM

The atmospheric extinction at the ORM has been analysed for several decades (for a general overview see Munoz-Tunón et al., 2007). Especially the Carlsberg Meridian Telescope (CMT), which was operated between 1984 and 2013, provided a database of 20 years of photometric measurements (Garcia-Gil et al., 2010) in the V band and r' band which is by today still the most comprehensive analysis of the atmospheric extinction over the ORM.

The analysis by Garcia-Gil et al. (2010) reveals a stable extinction coefficient over the 20-year period, with a median $k_V = 0.130\,\mathrm{mag\,airmass^{-1}}$ and no statistically significant long-term trend. Seasonal variations are clearly identifiable, with increased mean extinction values during the summer months (June-September) due to a higher frequency of Saharan dust intrusions. While the median extinction remains nearly constant across seasons, the mean v-Band extinction rises from $0.144\,\mathrm{mag\,airmass^{-1}}$ in the rest of the year to $0.183\,\mathrm{mag\,airmass^{-1}}$ in summer. The fraction of nights affected by dust or cirrus was estimated through log-normal fits to the extinction distributions, yielding 29% in summer and 13% in the remainder of the year. The CMT data also captured significant extinction increases following the eruptions of El Chichón (1982) and Mount Pinatubo (1991), demonstrating the influence of global-scale atmospheric events.

2.6 Sky Brightness Monitoring and Modelling

The brightness of the night sky is a fundamental limitation in ground-based astronomical observations. Even in the absence of moonlight and artificial light pollution, the night sky exhibits a persistent glow arising from a combination of natural atmospheric and astrophysical sources. This glow, often referred to as the night sky background, forms the baseline flux against which astronomical signals must be detected, and as such, it plays an important role in determining the achievable signal-to-noise ratio (SNR) in photometric and spectroscopic measurements.

Sky brightness is commonly expressed in units of astronomical magnitudes per square arcsecond ($\mathrm{mag\,arcsec^{-2}}$), a logarithmic surface brightness scale that facilitates comparison across different observing conditions and instruments. Under pristine dark-sky conditions, the zenith brightness in the V band typically lies in the range of 21.5-22.0 $\mathrm{mag\,arcsec^{-2}}$. Brighter sky conditions, whether due to atmospheric phenomena or anthropogenic light, degrade the contrast between astronomical sources and the background, requiring longer exposures to achieve a given SNR, or resulting in reduced photometric precision for fixed exposure times.

This section reviews the principal contributors to sky brightness, summarises common measurement techniques, and highlights key models used to predict sky background.

2.6.1 Natural and Artificial Contributions

Natural night sky brightness arises from several diffuse sources. The dominant contributor on a moonless night is airglow, originating from excited atoms and molecules. Airglow can account for up to 60% of zenith sky brightness, increasing toward the horizon due to increased atmospheric path length. Zodiacal light, caused by sunlight scattered by interplanetary dust, can account for up to 50%, especially near the ecliptic. Other contributors include diffuse starlight and galactic dust, as well as a negligible extragalactic background.

The Moon is the brightest natural source at night. A full Moon increases the sky brightness by up to 4 magnitudes in the optical bands; even a quarter Moon causes significant enhancement. The moon's effect is wavelength dependent, with a greater impact in bluer bands.

Twilight is caused by scattering of sunlight in the upper atmosphere (at dusk). Astronomical twilight ends when the Sun is 18° below the horizon. Before this, sky brightness decreases rapidly across civil ($<6°$) and nautical ($<12°$) twilight.

Sky brightness varies with time due to the solar cycle and seasonal effects. Airglow is stronger during solar maximum. For example, Mauna Kea shows V-band brightness ranging from 21.9 to 21.3 mag/arcsec^2 between solar minimum and maximum (Krisciunas, 1997).

Artificial light pollution elevates sky brightness due to scattered upward light from human settlements. The effect can extend tens to hundreds of kilometres. Pristine sites like Cerro Paranal or Mauna Kea have minimal artificial contributions due to remote locations and protective lighting laws. La Palma benefits from similar protections. Sky brightness in such sites reaches B$\approx$22.7, V$\approx$21.9 mag/arcsec^2 (Garstang, 1989).

2.6.2 Measurement Approaches

Quantifying night sky brightness is essential for assessing observational conditions and characterizing the impact of natural and artificial light sources. Several measurement techniques have been developed, each with specific advantages and limitations, depending on the required precision, spatial coverage, and practical constraints. A recent overview is given in Hänel et al., 2018.

Photoelectric photometers represent one of the earliest methods used in astronomy to measure sky brightness. These instruments use narrowband filters and photomultiplier tubes or photodiodes to obtain calibrated flux measurements in specific passbands, typically in the Johnson-Cousins system. Because they measure a small field of view, careful pointing to star-free sky regions is required. Long-term monitoring campaigns using photoelectric photometers have provided valuable historical records, but the method is labor-intensive and sensitive to calibration drifts (Sterken et al., 1992).

CCD imaging has become the standard approach in modern observatories (e.g. Aceituno et al., 2011). Sky background levels can be extracted from science images by measuring the mean or median pixel values in star-free regions or by masking stars algorithmically. This method allows spatially resolved measurements over large fields of view and enables correlation with other observational parameters (e.g. airmass, Moon phase). However, accurate background estimation requires proper flat-fielding and awareness of instrumental artifacts such as scattered light or residual bias structure.

All-sky cameras and DSLR-based systems (e.g. Smith et al., 2004 or Rabaza et al., 2010) with fisheye lenses enable wide-angle or full-sky monitoring. These instruments provide contextual information about the spatial distribution of sky brightness and cloud coverage. While relatively low cost and easy to deploy, such systems face challenges in calibration, particularly in correcting for vignetting and lens distortion. Nevertheless, they are used at observatories for qualitative monitoring and, when carefully calibrated, can yield quantitative sky brightness maps.

Sky Quality Meters (SQMs) (Sánchez de Miguel et al., 2017) are compact, commercially available devices based on photodiodes that report sky brightness in units of mag/arcsec^2. They offer a cost-effective solution for continuous monitoring and are popular among both professional astronomers and the amateur community. Although their spectral response is relatively broad and not strictly matched to standard filters, they are useful for tracking long-term trends, light pollution levels, and nightly variations when deployed in fixed, dark-sky locations.

Regardless of the instrument, accurate sky brightness measurements require careful calibration. This typically involves comparison with known standard stars or previously established zero points. Reported brightness values depend strongly on the filter system used; most professional data are presented in Johnson-Cousins UBVRI passbands or in a broadband visual-equivalent band. Variations between instruments and passbands must be considered when comparing data from different sources.

2.6.3 Empirical and Analytical Models

Predicting night sky brightness under varying conditions is important for observational planning and scheduling. Over the years, both empirical and physically motivated models have been developed to estimate sky brightness as a function of key parameters, such as lunar phase and atmospheric conditions.

One of the most widely adopted empirical models is the moonlight brightness model developed by Krisciunas et al. (1991). This model estimates the increase in sky brightness due to scattered moonlight based on a few easily accessible variables: the lunar phase, the altitude of the Moon, the altitude of the target, and the angular separation between the Moon and the target. It provides practical predictions that are sufficiently accurate for most scheduling and exposure time estimation purposes.

More advanced models aim to simulate sky brightness with greater physical fidelity by incorporating the underlying physics of atmospheric scattering and lunar reflection. Notably, the *ESO Sky Model* (Noll et al., 2013) provides a comprehensive framework that includes detailed radiative transfer calculations for both Rayleigh and aerosol scattering. This model accounts for the Moon's position, phase, and altitude, as well as the observer's location and atmospheric conditions. It has been successfully applied to estimate sky brightness at the Very Large Telescope (VLT) in Chile.

Other natural contributors to sky brightness are also modelled using empirical or semi-empirical approaches. Zodiacal light, which arises from sunlight scattered by interplanetary dust, is typically estimated using maps derived from satellite observations and photometric surveys, with brightness depending on solar elongation and the ecliptic latitude of the observed field (Leinert et al., 1998). In contrast, airglow, a complex phenomenon resulting from chemical reactions in the upper atmosphere, remains difficult to predict accurately. Nevertheless, its intensity is known to correlate with solar activity, increasing during solar maximum.

Artificial sky brightness due to light pollution has also been modelled extensively. One of the foundational works in this domain is by Garstang (1989), who combined atmospheric scattering models with empirical data on urban light emissions to estimate sky brightness gradients around populated areas. Subsequent models, such as those presented by Duriscoe et al. (2018), have refined these predictions using satellite-derived measurements of upward light emission, particularly from instruments such as the Visible Infrared Imaging Radiometer Suite (VIIRS, Cao et al., 2014) aboard the Suomi NPP satellite. These models enable global-scale mapping of light pollution and have become important tools for site characterisation and dark sky preservation efforts.

2.6.4 Integration into Robotic Scheduling

Accurate knowledge of sky brightness plays an important role in the scheduling of observations on robotic telescopes. Unlike classical observing, where a human observer may assess sky conditions in real time, robotic systems must rely on predictive models and real-time environmental data to make autonomous decisions. As such, sky brightness can be incorporated into multiple layers of the scheduling process to ensure optimal scientific return.

At the proposal stage, astronomers typically specify constraints related to acceptable sky brightness levels for their observations. These may take the form of categorical conditions such as "dark," "grey," or "bright" time, or they may define quantitative thresholds in units of mag/arcsec^2. Such constraints directly influence time allocation decisions, determining when a given target is eligible to be observed within the operational schedule.

During night-time operations, dynamic schedulers evaluate candidate observations by predicting sky brightness for each target at the current time and sky position. This real-time evaluation is important for robotic surveys such as the Zwicky Transient Facility (ZTF, Bellm et al., 2019) and the Vera C. Rubin Observatory (LSST, Ivezić et al., 2019), where rapid cadence and automated decision-making require fast yet accurate estimation of observing conditions. By computing expected sky brightness using empirical or analytical models, schedulers can prioritise targets that match the current sky conditions and defer those that do not.

Sky brightness also enters directly into exposure time calculators, particularly through its influence on the signal-to-noise ratio (SNR). As sky background increases, the SNR for a given target decreases, necessitating longer exposures to achieve the same photometric precision. Exposure calculators therefore incorporate sky brightness predictions or measurements to adjust exposure times dynamically, or to reject targets outright if the conditions are deemed unsuitable for the desired data quality.

Environmental monitoring systems provide real-time measurements of sky brightness, often using instruments such as Sky Quality Meters (SQMs) or calibrated all-sky cameras. These devices supply live data that can be fed into the scheduling system to update or override model-based predictions. By continually re-evaluating sky conditions, robotic systems can adapt to transient changes, such as the onset of moonlight or artificial light interference, and adjust the observing plan accordingly.

Through this multi-level integration of sky brightness, ranging from proposal constraints to real-time scheduling decisions, robotic observatories are able to maximise efficiency, protect sensitive programs, and allocate bright conditions to more tolerant observations.

This strategic coordination ensures that the telescope is utilised to its full potential under varying sky conditions.

2.6.5 Sky Brightness on La Palma

In recent decades, light pollution is a significant and growing contributor to sky brightness, especially in the proximity of urban areas. The upward scattering of artificial light by aerosols and molecules in the lower atmosphere creates a diffuse glow (referred to as skyglow) that can elevate the background by several magnitudes, particularly at low elevations and near the horizon. The spectral characteristics of this light pollution depend on the types of lighting in use, with modern white-light-emitting diodes (LEDs) posing new challenges due to their broad emission and strong blue component, which scatters efficiently in the atmosphere.

To prevent increased light pollution, the Spanish government introduced the "Ley del Cielo" ("The Sky Law") for the Island of La Palma (*Ley 31/1988, Ley del Cielo* 1988). The Sky Law represents one of the earliest and most comprehensive national efforts to legally protect astronomical infrastructure from environmental degradation. Administered in collaboration with the Instituto de Astrofísica de Canarias (IAC), the law applies specifically to the islands of La Palma and is enforced across multiple sectors that affect observational quality. It regulates four primary domains: light pollution, radio frequency interference, aerial traffic, and atmospheric pollution. The most visible and impactful component of the Sky Law concerns the control of artificial lighting, which poses a significant threat to the quality of optical and near-infrared observations through the enhancement of night sky brightness.

The enforcement of the Sky Law has demonstrable benefits. A spectroscopic survey conducted at the ORM before and after midnight (Pedani, 2004) revealed an approximately 50% reduction in mercury emission lines following the mandated light curfew, and a long-term decrease in total artificial sky brightness of < 0.1 at zenith in V and R bands.

2.7 Scheduling for Robotic Telescopes

2.7.1 Robotic Scheduling in Time-Domain Astronomy

2.7.1.1 From Nightly Scripts to Reactive Systems (1990–2010)

The earliest robotic telescope scheduling systems emerged in the 1990s, when automation was first applied to small-aperture observatories to improve efficiency and reduce the need for on-site personnel. These systems relied heavily on static scripts and predefined observing plans, typically generated during the day and executed without further decision-making.

Pioneering efforts included the RoboScope system (Honeycutt et al., 1992), which autonomously executed photometric observations of variable stars using a nightly target list and weather-based abort conditions. Similarly, the Global Oscillations Network Group (GONG, Harvey et al., 1996) and CONCAM (the Continuous Camera, Pereira, 2003) systems demonstrated how all-sky imaging and solar observations could be run in a fully unattended fashion. However, none of these systems featured real-time target selection based on scientific urgency or environmental feedback.

A major shift occurred with the development of dynamic queue schedulers: systems that could evaluate observation requests in real time, using current conditions, target visibility, and scientific priority. One of the earliest and most influential systems in this class was RAPTOR (Vestrand et al., 2002), which integrated transient detection and follow-up by automatically reacting to its own real-time alert stream.

Also, the MASTER network began implementing rapid response strategies for gamma-ray bursts and other transients (Lipunov et al., 2012), using custom schedulers that selected targets based on pre-assigned priority scores and sky location. These systems often used proprietary rules and hard-coded logic but laid the groundwork for later general-purpose robotic observatories.

The RoboNet-1.0 (Mottram et al., 2008) initiative marked another important step. Built around the Liverpool Telescope and the two Faulkes Telescopes, it offered a shared scheduling platform for science users across the UK and beyond. Although initially driven by human-curated scripts, it eventually incorporated a dynamic dispatch algorithm based on feasibility filtering and priority scoring. This structure - first filter out infeasible targets, then score and select the best one - became a defining feature of many robotic schedulers in the following decade.

2.7.1.2 Mature Queue Scheduling and Science Success (2010–2020)

By the early 2010s, dynamic queue-based scheduling had become a standard feature of many robotic telescope systems. The key architecture, a feasibility filter followed by a multi-factor scoring function, was adopted by diverse observatories ranging from small-aperture survey instruments to large robotic facilities.

The Liverpool Telescope's scheduler, described in Fraser et al. (2004), became one of the most widely referenced implementations. It provided real-time decision-making based on elevation, scientific priority, seeing constraints, and target cadence. Other facilities, including LCOGT (Brown et al., 2013) and the Faulkes Telescopes, built similar systems that allowed geographically distributed users to submit proposals and receive data with minimal human intervention. These systems were characterised by modularity, robustness, and operational simplicity.

Scientific output during this period highlighted the power of such systems. Robotic telescopes made critical contributions to the study of gamma-ray bursts (e.g., via Target-of-opportunity overrides), novae and recurrent novae (e.g., Darnley et al. 2014), and early-phase supernovae. Observations could be executed within minutes to hours of a trigger, and long-term campaigns (including AGN monitoring and microlensing) were reliably scheduled over months.

Despite differences in hardware and institutional context, most schedulers converged on a shared philosophy: score available targets using a weighted sum of heuristic features, then select the highest scorer. Flexibility was achieved by adjusting score weights or adding constraints (e.g., weather, moon avoidance). These systems worked well for observatories with moderate target volumes and relatively homogeneous science goals, as was typical before the rise of high-cadence transient surveys.

2.7.1.3 The New Frontier: Transient Floods and AI-Guided Scheduling

The past decade has brought a dramatic shift in the demands placed on robotic scheduling systems. High-cadence, wide-field surveys such as Pan-STARRS (Kaiser et al., 2002), ATLAS (Tonry et al., 2018), the Zwicky Transient Facility (ZTF; Bellm et al., 2019), and the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST; Ivezić et al. 2019) have transformed transient astronomy into a data-rich regime. These facilities together produce millions of alerts per night, many of which are only visible for brief intervals or evolve rapidly over hours or days.

To cope with this deluge, dedicated alert brokers have been developed to process, enrich, and classify transient alerts in real time. Notable among these are:

- **Lasair** (Smith, 2019), a broker optimised for UK-based follow-up, built and tested around the ZTF and ready for LSST streams.
- **ANTARES** (Matheson et al., 2021), developed in the US with strong integration into LSST operations planning.
- **ALeRCE** (Förster et al., 2021), a Chilean-led broker with particular focus on machine-learning-based real-time classification.

All three brokers implement automated classification pipelines using supervised machine learning trained on archival and simulated light curves. The typical goal is to assign transient candidates to probabilistic classes such as Type Ia or core-collapse supernovae, cataclysmic variables, AGN flares, or stellar variables. These brokers use decision trees, random forests, or neural networks depending on the use case and latency constraints. Features used include light curve shape descriptors, contextual catalogue matches, colour evolution, and variability metrics.

The classifications are not deterministic but probabilistic. For example, ALeRCE employs a two-stage classifier: a *Real-Bogus* filter to discard artefacts, followed by a *Hierarchical Light Curve Classifier* that assigns class probabilities from a taxonomy of known transient types. Lasair uses a similar two-step pipeline that includes a Gaia match and star-galaxy separation to inform its classification. ANTARES includes features such as the Mahalanobis distance of a light curve from known class templates.

These automated pipelines allow brokers to triage alerts into scientifically valuable categories. However, schedulers must still make complex decisions:

- **Uncertainty-aware execution**: Classifiers return probabilities, schedulers must decide how to act under ambiguity.
- **Science-prioritised scoring**: Should a high-probability Type Ia supernova be prioritised over a 50% chance of a rare transient?
- **Timing and evolution**: Some targets must be observed immediately (e.g. early spectra), others may benefit from delayed follow-up.
- **Exposure optimisation**: Based on transient brightness, required SNR, and real-time observing conditions.

Several telescope networks have begun adapting to this new landscape. The ZTF Bright Transient Survey (BTS) implemented a broker-to-telescope pipeline to prioritise Type Ia supernovae (Fremling et al., 2020a), while the GROWTH Marshal (Kasliwal et al., 2019) supports distributed manual and semi-automated follow-up planning. Experimental frameworks such as ALTSched (Rothchild et al., 2019) proposed goal-driven schedulers that integrate target utility functions.

These developments signal a fundamental shift: robotic scheduling is no longer about optimising a queue, but about making autonomous scientific decisions in real time. The scheduler becomes a reasoning engine that must integrate astrophysical classification, urgency, visibility, weather, and instrumentation constraints, often within seconds. Meeting these challenges requires new algorithmic approaches, including predictive scheduling and dynamic exposure planning.

It is in this context that the limitations of legacy systems such as the Liverpool Telescope's static scoring architecture, become increasingly apparent. The following sections explore these limitations and present possible pathways forward.

2.7.2 The Liverpool Telescope Scheduler

The Liverpool Telescope (LT), located at the Observatorio del Roque de los Muchachos (ORM) on La Palma, is one of the world's largest fully robotic telescopes. Since its first light in 2004, it has played a key role in time-domain astrophysics, particularly in the follow-up of gamma-ray bursts, novae, and supernovae. At the core of its autonomous operations lies a dynamic queue-based scheduler, which selects observations in real time based on a scoring algorithm and a set of feasibility checks (Fraser et al., 2004).

The LT's scheduler evaluates and executes observing groups (user-defined sequences of exposures) using a two-stage decision process. First, a `feasibility model` filters out groups that cannot be executed under current constraints (e.g. unavailable instrument, timing conflict, conditions, etc.). Then, a `scoring algorithm` assigns a numeric value to each candidate group, selecting the one with the highest score for execution.

Candidate groups are first placed into one of three categories:

- **Fixed Group Candidates**: groups with strict timing constraints,
- **Primary Candidates**: standard groups without priority Z,
- **Background Candidates**: low-priority fallback groups (priority Z).

The scheduler selects the best available list in that order of preference (Fixed > Primary > Background), and computes a score for each group in the list.

$$\text{score} = 0.5 \times S_E + 0.5 \times S_P + 0.15 \times S_{SM} + 0.05 \times \text{N} \tag{2.2}$$

Elevation Score S_E: This component favours targets currently near their maximum altitude. For each target within a group, the ratio of its current altitude to its predicted nightly maximum is computed. These ratios are then averaged across all targets in the group. This averaging reflects the fact that a group typically represents multiple observations of the same object (e.g. multi-filter or multi-instrument exposures) or of spatially co-located targets, which therefore share nearly identical visibility constraints. Using the mean ensures that the group is evaluated as a coherent observing unit rather than being dominated by a single component.

Because the score is normalised to each target's own nightly maximum altitude, targets with intrinsically low culminations are not penalised at transit. However, the metric still favours observations close to culmination, which may systematically deprioritise objects that are already descending in altitude, particularly when competing with rising targets that have not yet reached peak elevation.

The score naturally increases for rising targets and decreases for setting targets. As a result, targets in the eastern sky often receive a higher elevation score than western targets at comparable absolute altitudes. In time-critical campaigns, this temporal asymmetry may cause setting targets—despite being scientifically urgent—to be deferred in favour of objects that will remain visible for longer.

Priority Score S_P: Each group inherits the scientific priority (A, B, C, or Z) from its proposal and may be flagged as "urgent". The score is increased by:

- +4 for priority A, +2 for B, 0 for C, -100 for Z,
- +2 if marked urgent,
- +1 if the group is a monitor or minimum-interval observation.

Here, a monitor observation refers to a repeated observation of the same target as part of a time-series or long-term monitoring programme. A minimum-interval observation is one for which a specified minimum time separation between consecutive executions is enforced by the scheduler, ensuring controlled cadence and temporal spacing.

This value is then adjusted by the proposal's "priority offset" (-7 to +7) and scaled.

Seeing Match Score S_{SM}: Groups may specify a seeing constraint (e.g. max 1.2"). The scheduler compares this to the current seeing estimate and computes a match score:

$$\text{smscore} = \frac{1}{1 + (\text{maxSeeing} - \text{predictedSeeing})} \tag{2.3}$$

Here, `predictedSeeing` denotes the current estimate of atmospheric image quality provided to the scheduler. At the Liverpool Telescope, this value is derived automatically from recent science exposures through the Quicklook processing system, which extracts image-quality metrics from stellar profiles. The scheduler uses this rolling estimate as a real-time proxy for prevailing seeing conditions. Default seeing is assumed to be 1.5 arcsec if no constraint is set.

It should be noted that this formulation does not behave as a simple reward function for better seeing. Instead, it reflects the historical design of the LT scheduling system, where individual scoring components act as modifiers within a larger composite priority

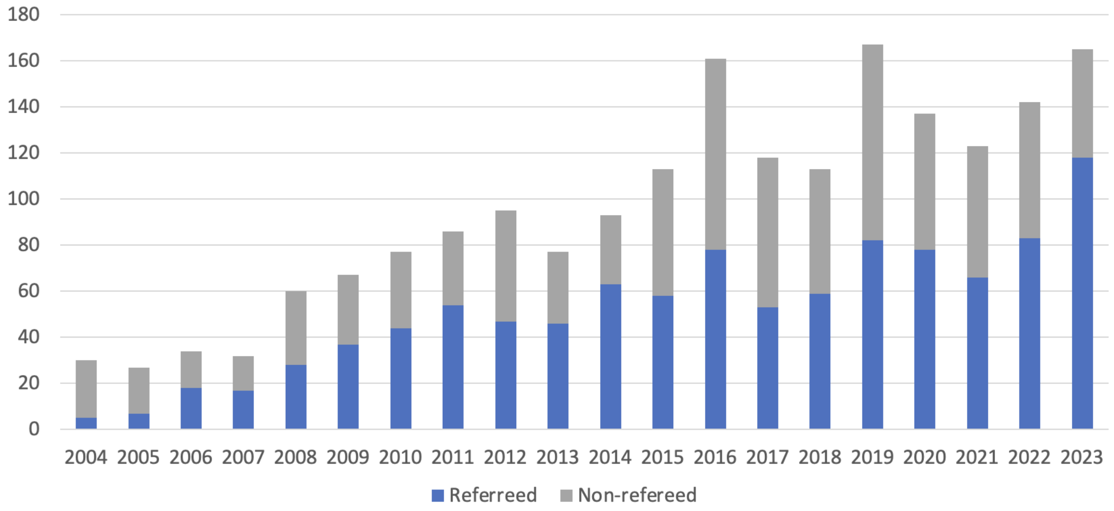

FIGURE 2.4: Number of publications related to the Liverpool Telescope in the years 2004 to 2023 (see `https://telescope.livjm.ac.uk/Pubs/`)

calculation. In practice, the scheduler combines multiple terms (priority class, elevation score, urgency, and constraint matches), and the final ranking is determined by the combined score rather than by the behaviour of any single component in isolation.

Random Noise N: A small random term ensures non-deterministic tie-breaking and avoids starvation of similar-scoring groups.

This scoring system is both light- weight and effective. It has served the LT well for nearly two decades, enabling significant science output with minimal human oversight (see Figure 2.4). However, as the demands of time-domain astronomy have evolved, this architecture shows signs of limitations, which will be explored in the following section.

2.7.3 Limitations of the Current Approach

The Liverpool Telescope's current scheduler, while effective for its original science goals, faces several limitations in the context of modern time-domain astronomy. These limitations are not due to technical failures, but rather reflect a change in the scientific landscape: new types of targets, faster timescales, and higher volumes of alerts demand a level of responsiveness and flexibility that the current system was not designed to provide.

Delayed response to fast transients: The scheduler is designed to favour targets near their meridian crossing, using elevation as a major component of the scoring function. While this maximises data quality by favouring observations near transit, it can delay follow-up of newly detected transients that are still at low elevation. In rapidly evolving events, waiting for improved altitude may result in missed early-time science. For fast-evolving transients such as Fast Blue Optical Transients (FBOTs, e.g., AT2018cow; Perley et al., 2019), gamma-ray bursts (GRB, i.e. 221009a “The BOAT”, Laskar et al., 2023), kilonovae, or early-phase novae (e.g., M31N 2012-10a; Darnley et al., 2014), this delay can result in missed scientific opportunities. Although the LT provides a dedicated override mode (Target of Opportunity) for gamma-ray bursts and other extreme events, most targets must still pass through the queue-based system, which may not prioritise them optimally.

Inadequate support for monitoring campaigns: Many long-term programs such as reverberation mapping of Active Galactic Nuclei (AGN), require regular, evenly spaced observations across weeks or months, sometimes regardless of weather or lunar phase. While the LT scheduler can repeat observations with specified intervals, the Phase II system (the interface through which accepted proposals are translated into detailed observing requests including targets, exposure times, constraints, and monitoring intervals) offers limited control over tolerances, priorities, or fallback behaviour under suboptimal conditions. As a result, programs like those described in Goicoechea et al. (2024) and McHardy et al. (2023) do require regular manual adjustment of exposure times and constraints or acceptance of an uneven cadence.

Bias against low-altitude and southern targets: As discussed in the previous section, the elevation scoring metric strongly favours high-altitude targets. This systematically devalues targets that transit low on the sky, typically southern objects from La Palma, or those visible only near the horizon. Scientifically important targets may be repeatedly outscored by better-placed alternatives, and eventually fall off the schedule altogether. While technically feasible to observe, such targets require active intervention (changing priority or priority offsets by the LT support staff) or relaxed constraints to be executed.

Human-in-the-loop classification bottleneck: The era of wide-field surveys has brought a deluge of transient alerts: ZTF, ATLAS, and now LSST at the Vera C. Rubin Observatory, generate thousands of candidate events per night. With LSST expected to produce up to 10,000 alerts per second, the volume of transients requiring follow-up will definitely exceed the capacity of traditional human-driven scheduling systems. However, only a small fraction of these alerts currently receive spectroscopic classification (Kulkarni, 2020), leaving a significant gap in our ability to characterise the transient sky.

The Liverpool Telescope’s Phase II interface requires manual preparation of observations, and the existing scheduler is not optimised for high-throughput classification work. This manual overhead creates a bottleneck in processing large alert streams. In contrast, the next-generation New Robotic Telescope (NRT, see **??**) is being designed with large-scale transient classification as a core goal. To achieve this, the follow-up loop must be closed:

from alert ingestion, to exposure time calculation, to immediate execution, all without human intervention.

This new operational paradigm requires a scheduler that not only reacts quickly, but also intelligently. For example, when an alert is received from brokers such as Lasair (Smith, 2019) or the Transient Name Server (TNS), it often includes a machine-learned pre-classification (e.g. Broccia, 2021; Williams et al., 2024) and an estimated magnitude in one or more bands (typically r or g). A modern scheduler must be capable of taking this input, computing the necessary exposure time to achieve a user-defined signal-to-noise ratio, and selecting the optimal execution window along the target's sky path. In some cases such as fast-evolving or rare transients, the observation must be triggered immediately, regardless of suboptimal conditions. In others, a more flexible window allows the scheduler to delay the observation for better sky placement and shorter exposure time.

Scientific motivation for a new paradigm: The limitations described above point toward a need for more responsive, predictive, and context-aware scheduling strategies. Rather than relying on static constraints and elevation-based scoring, future systems must integrate real-time atmospheric conditions, transient urgency, and science-driven priorities. This is particularly important for observations where speed matters more than quality, or where constraints are flexible but scientifically significant.

This kind of adaptive decision-making, where scientific urgency and exposure strategy are computed dynamically, forms the basis of the signal-to-noise ratio based scheduling system described later in this thesis.

2.8 Automated Transient Classification

In the modern era of time-domain astronomy, the rate at which new astrophysical transients are discovered has far surpassed the capacity to classify them spectroscopically. The combination of high-cadence optical surveys and multi-messenger alerts has transformed the field from one in which discoveries were rare and highly sought after, to one in which the nightly influx of candidates threatens to overwhelm available follow-up resources (Kulkarni, 2020). Facilities such as the Zwicky Transient Facility (ZTF) routinely generate between 10^5 and 10^6 alerts per clear night (Bellm et al., 2019), while the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST) will produce an anticipated 10^7 alerts per night (Ivezić et al., 2019). This represents a step change not only in data volume but also in operational philosophy: traditional workflows, in which human astronomers filter candidates and schedule observations, are no longer viable at these scales.

The so-called *classification bottleneck* has become a central challenge in this new landscape. While photometric filtering can eliminate obvious false positives and recover likely variables or moving objects, definitive classification of extragalactic transients almost always requires spectroscopy. In 2019, the ratio of reported optical transients to those receiving a spectroscopic classification was approximately ten to one (Kulkarni, 2020). With LSST-scale discovery rates, even maintaining this ratio would demand an unsustainable increase in spectroscopic throughput, and the inevitable consequence is that many rare or unusual phenomena will remain unidentified.

Efforts to automate classification have their origins in the late 2000s, when the Palomar Transient Factory (PTF; Law et al. 2009) integrated real-time candidate vetting, robotic scheduling, and immediate photometric and spectroscopic follow-up. PTF demonstrated that an end-to-end automated system could operate productively, with the Palomar 48-inch telescope providing discovery imaging and the robotic 60-inch telescope delivering rapid follow-up. Its successor, the intermediate Palomar Transient Factory (iPTF), improved the classification rate through more sophisticated filtering and scheduling, and ZTF extended this model to an order-of-magnitude higher alert rate by deploying a 47 deg^2 camera and a modernised alert-distribution infrastructure. Within ZTF, the Bright Transient Survey (BTS) exemplified the efficiency gains possible with a dedicated spectroscopic facility: using the low-resolution, high-throughput Spectral Energy Distribution Machine (SEDM) on the Palomar 60-inch, the BTS classified nearly every bright transient in its footprint, and in 2019 alone the SEDM was responsible for roughly 40% of all new supernova classifications worldwide (Fremling et al., 2020a).

As alert rates increased, it became evident that intermediate software layers would be essential to bridge the gap between survey discovery and telescope follow-up. Event brokers such as Lasair (Smith, 2019), ANTARES (Matheson et al., 2021), and ALeRCE (Förster et al., 2021) emerged to ingest the full survey alert stream, enrich each alert with contextual information, and apply machine-learned classification algorithms. These brokers typically cross-match transients with archival catalogues to identify likely hosts, separate Galactic from extragalactic sources, and flag known variables. Machine-learning components then analyse available photometry and colour information to assign probabilistic classifications, identifying, for example, a transient as a likely Type Ia supernova, cataclysmic variable, or tidal disruption event. The probabilistic nature of these classifications reflects the fact that alerts may contain only a few data points initially; brokers update their classifications dynamically as additional observations accumulate. In this way, candidate lists are continually reprioritised, with the most scientifically valuable events, however that value is defined, surfaced for rapid follow-up.

Automated classification is most effective when coupled directly to robotic telescopes. The Liverpool Telescope (LT) has demonstrated this through its low-latency queue injection for gravitational-wave and gamma-ray burst triggers (Copperwheat et al., 2016), in which alerts meeting predefined criteria can bypass manual proposal processes and enter the observing queue within seconds. Similar concepts are implemented in other domains, such as the H.E.S.S. gamma-ray observatory, whose transient follow-up system automatically filters incoming alerts according to science-case criteria, assesses visibility and conditions, and schedules observations accordingly (Hoischen et al., 2022). These systems are designed to operate as part of a closed loop: discovery leads to automatic classification and prioritisation, which in turn drives immediate observation requests; the resulting data are reduced in real time and fed back into the classification stage, potentially triggering further observations or community alerts.

Science-case-based scheduling forms a natural extension of this paradigm. A possible kilonova candidate from a gravitational-wave localisation might trigger a pre-scripted sequence of wide-field imaging followed by rapid spectroscopy if the object meets colour or brightness criteria. A bright, rising supernova could be queued for classification within 24 hours, prioritised according to host galaxy properties or redshift. A gamma-ray burst afterglow might prompt initial short exposures to track its rapid fading, followed immediately by a spectrum if it remains above a given magnitude threshold. By encoding such logic into the scheduler, robotic telescopes can adapt their behaviour to the transient type, ensuring that limited follow-up resources are used to maximum scientific effect.

The upcoming New Robotic Telescope (NRT) has been conceived with this fully integrated approach in mind. Its design calls for a 4-metre aperture with sub-30-second reaction times, optimised for both the rapid characterisation of fast-fading transients and the high-throughput classification of survey discoveries (Copperwheat et al., 2015). Operating in tandem with the existing LT, equipped with a wide-field imager to capture candidate fields, the NRT's main spectrograph would execute large numbers of classification observations per night, driven directly by brokered alert streams. Similar developments are underway or already deployed elsewhere: the GOTO (Dyer et al., 2018) and BlackGEM (Groot et al., 2024) arrays for gravitational-wave counterpart searches, the SOXS spectrograph (Schipani et al., 2018) for rapid classification on the ESO NTT 3.6-metre telescope, and network-level coordination efforts such as the Astronomical Event Observatory Network (AEON) (Street et al., 2020), which enables automated distribution of follow-up requests to the most appropriate telescope based on location, conditions, and capabilities.

In combination, these survey, broker, and robotic follow-up systems are converging on an operational model in which the latency from transient discovery to scientific exploitation is measured in minutes. By removing the human from the immediate decision-making loop, automated transient classification ensures that rare, short-lived, and scientifically valuable events are recognised and acted upon before they fade, addressing one of the most pressing challenges of time-domain astronomy in the 2020s.

2.9 Signal-to-Noise Ratio Based Scheduling

2.9.1 Overview

In ground-based optical and near-infrared (NIR) astronomy, the signal-to-noise ratio (SNR) is a fundamental metric that governs the quality and interpretability of observational data. The SNR quantifies the clarity with which an astronomical signal can be distinguished from various sources of noise and is thus important in both photometric and spectroscopic applications (Howell, 2006).

The *signal* in astronomical imaging and spectroscopy is defined as the number of photoelectrons generated by photons from the target source. Its value depends on the source brightness, the telescope aperture, the detector quantum efficiency, and the exposure time as well as the observational conditions such as atmospheric extinction, clouds, seeing and sky background. The *noise* encompasses all uncertainty contributors, and the total noise is typically the quadrature sum of several independent components: photon shot noise from both the source and sky background, dark current noise, and readout noise (Howell, 2006).

Photon shot noise arises due to the Poisson nature of photon arrival rates. For a detected flux of N electrons, the associated noise is $\sqrt{N}$. Sky background, a dominant noise source in faint object detection, contributes proportionally to the number of pixels over which the source is spread, and is heavily affected by atmospheric conditions and filter choice (Newberry, 1991). In the NIR, sky brightness can be orders of magnitude higher than in the optical, necessitating short exposures and rigorous background subtraction (Howell, 2006).

Detector dark current and readout noise add fixed noise components. While modern CCDs exhibit low dark currents when cooled, NIR detectors may still contribute significantly depending on the operational temperature. Readout noise, typically 3–10 electrons RMS per pixel, becomes significant in low-light or high-resolution observations, where signal levels are low or distributed over many pixels (Howell, 2006).

Atmospheric seeing, which broadens the point-spread function, degrades SNR by increasing the aperture area required to capture a majority of the source light, thus incorporating more background noise. In contrast, adaptive optics (AO) can counteract this effect by concentrating light into fewer pixels, improving SNR under suitable conditions.

Mathematically, the SNR can be approximated using the CCD equation:

$$\mathrm{SNR} = \frac{S}{\sqrt{S + B + D + N_{\mathrm{read}}^2}}$$

where S is the source signal, and the noise terms correspond to sky background (B), dark current (D), and readout contributions(N_{read}). Depending on observational parameters, SNR may be dominated by any one of these terms. For instance, faint-source observations under dark skies tend to be background-limited, while narrowband imaging or short exposures may be readout-noise dominated (Newberry, 1991; Howell, 2006).

SNR is directly linked to the precision of photometric and spectroscopic measurements. In photometry, the magnitude uncertainty σ_{mag} scales inversely with SNR: $\sigma_{\mathrm{mag}} \approx 1.0857/\mathrm{SNR}$. In spectroscopy, sufficient SNR is required to resolve weak spectral features, measure redshifts, and derive stellar parameters accurately. Optimal extraction techniques, such as those introduced by Horne (1986), further enhance SNR by weighting pixel contributions based on expected signal profiles.

2.9.2 Traditional Scheduling Approaches

In traditional observing, exposure lengths are typically fixed or planned in advance based on time (e.g. a certain number of seconds or minutes per target). This time-based scheduling assumes nominal conditions and does not adjust for real-time changes in sky quality. Classic scheduling methods like queue scheduling or fixed nightly plans treat observing time as static, which makes them vulnerable to variations in conditions (Granzer et al., 2010).

If the sky gets brighter (moonrise, twilight) or seeing worsens, a fixed exposure may fall short of the desired signal-to-noise ratio (SNR); conversely, if conditions improve, the exposure might collect more SNR than needed, effectively wasting telescope time. Time-based exposure scheduling lacks flexibility under varying conditions, leading to suboptimal data quality or efficiency when weather, sky background, or seeing deviate from the plan.

A major limitation of fixed schedules is that unpredictable environmental changes can "break" an optimal schedule, requiring on-the-fly adjustments. For example, a carefully optimised sequence of observations can be derailed by unforeseen clouds or poor seeing, since the original exposure times assumed better conditions. Ground-based observatories have long recognised that a rigid observing plan is difficult to maintain in practice due to weather variability. Traditional approaches often build in conservative exposure margins or simply accept the risk of variable SNR outcomes.

Other domains of astronomy have started to see the benefits of moving beyond purely time-based scheduling. For instance, in radio astronomy with very long baseline interferometry (VLBI), recent tests with the VLBI Global Observing System (VGOS, Petrachenko et al., 2013) replaced fixed scan durations with an SNR-based scheduling strategy. By scheduling each observation until a predefined SNR threshold was reached rather than allocating a fixed scan duration, the system reduced both under- and over-integration. In this context, “observing efficiency” refers to the number of successfully completed scans meeting the required SNR per unit telescope time. Under variable conditions, fixed-duration scans either fail to reach sufficient sensitivity (in poor conditions) or overshoot the requirement (in good conditions), wasting valuable time. The SNR-based strategy therefore increased the fraction of telescope time contributing directly to scientifically usable data (Schartner et al., 2025). This example underscores the potential gains of SNR-driven scheduling: improved data quality and optimised telescope time under variable conditions.

Schartner et al., 2025 propose and evaluate an SNR-based scheduling approach for the VGOS (a network of radio telescopes dedicated to geodetic and astrometric measurements). Traditional scheduling strategies for VGOS prioritise geometrical configurations to optimise Earth orientation parameters and station positions, often neglecting real-time observational conditions such as system sensitivity and source flux. The authors introduce a novel framework that integrates a predictive model for SNR estimation directly into the scheduling algorithm, allowing the selection of scans with higher expected SNRs without significantly compromising the geodetic performance of the network.

Through simulations and empirical data from VGOS sessions, their study demonstrated that incorporating SNR predictions improves the robustness of scheduled scans, particularly under sub-optimal observing conditions. The approach reduces the number of failed observations and increases the average SNR across the session. The paper highlights the trade-offs between purely geometric optimisation (where scan selection is driven primarily by array geometry, source position, and visibility constraints) and quality-aware scheduling (in which predicted signal-to-noise ratio is explicitly incorporated into scan duration and ordering decisions), and suggests that hybrid models—blending geometry and SNR considerations—could offer a superior balance for real-time operations. This work illustrates the potential of SNR-aware scheduling and reinforces the general principle that integrating sensitivity models into robotic scheduling frameworks can enhance the scientific return of observational campaigns.

2.9.3 Adaptive Exposure Systems

Modern robotic observatories are increasingly incorporating adaptive exposure systems to overcome the limitations of static scheduling. These systems adjust exposure times dynamically in response to actual conditions or data, rather than adhering to a fixed plan. Notable examples include the STELLA robotic observatory (Granzer et al., 2010) and the Las Cumbres Observatory Global Telescope (LCOGT, Brown et al., 2013) network.

Granzer et al., 2010 report on three years of fully autonomous operations of the STELLA robotic observatory, which consists of two 1.2-meter telescopes located at the Izaña observatory in Tenerife. STELLA is a fully autonomous observatory that reacts to real-time conditions. In operation, it continuously evaluates weather and seeing and chooses targets dynamically based on a dispatch-style scheduling system, in which the next observation is selected in real time from the set of currently feasible targets by evaluating a weighted merit function, rather than constructing a fixed schedule for the entire night in advance. This approach uses weighted merit functions to select observations in real time, balancing short-term constraints (e.g., airmass, moon avoidance) with long-term priorities (e.g., phase coverage, revisit intervals).

Although signal-to-noise ratio (SNR) is briefly mentioned in the context of guider exposure tuning (aiming for SNR 5), it is not used as a scheduling criterion. Instead, STELLA focuses on robust dynamic scheduling logic capable of handling diverse scientific campaigns and changing environmental conditions. The paper provides operational insights into robotic observatories and illustrates the flexibility and effectiveness of real-time merit-based scheduling in a production setting.

The LCOGT network employs a different adaptive strategy. It operates a distributed set of telescopes and uses a global adaptive scheduler to optimise target assignment. If one site becomes unavailable, the scheduler can reassign the observation to another telescope. This adaptivity lies in site selection and timing, rather than active adjustment of exposure durations. Exposure times are generally set in advance by users, with the help of calculators and guidelines.

A useful distinction can be made between forecasting and reactive adaptation. Forecasting approaches use predicted conditions (e.g., weather forecasts, expected seeing) to inform scheduling. Reactive systems, on the other hand, respond to real-time measurements. STELLA is a prime example of reactive control.

Schlekat et al., 2025 describe a major software and systems overhaul of the Skynet Robotic Telescope Network, a global network of 20 autonomous, queue-based optical telescopes. The upgraded infrastructure introduces new observing modes, scheduling capabilities, and a public API to facilitate flexible, scriptable access for time-domain astronomy. One of the key future capabilities under development is the ability to submit observation requests based on a desired signal-to-noise ratio (SNR).

2.9.4 SNR-Driven Models

The foundation of SNR-based scheduling is a formal model of the signal-to-noise ratio, which accounts for signal and noise contributions in an exposure:

$$\mathrm{SNR} = \frac{S}{\sqrt{S + B + D + N_{\mathrm{read}}^2}} \tag{2.4}$$

where S is the source signal, B the background sky signal, D the dark current, and N_{read} the total readout noise.

These terms can be expanded in terms of exposure time t, number of pixels n_{pix}, and noise rates. For instance,

$$\begin{aligned} S &= s \cdot t \\ B &= n_{\mathrm{pix}} \cdot b_{\mathrm{sky}} \cdot t \\ D &= n_{\mathrm{pix}} \cdot d \cdot t \\ N_{\mathrm{read}}^2 &= n_{\mathrm{pix}} \cdot N \cdot \sigma_{\mathrm{read}}^2 \end{aligned}$$

Combining these, the SNR becomes:

$$\mathrm{SNR} = \frac{st}{\sqrt{st + n_{\mathrm{pix}}(b_{\mathrm{sky}}t + dt) + n_{\mathrm{pix}} N \sigma_{\mathrm{read}}^2}} \tag{2.5}$$

This formulation reveals the dependencies:

- **Target Magnitude:** Fainter objects require longer exposure times to reach the same SNR.
- **Sky Background:** Bright conditions (e.g., moonlight) reduce SNR due to increased B.

- **Extinction and Airmass:** Higher airmass reduces S due to atmospheric attenuation.
- **Seeing:** Poor seeing spreads signal over more pixels, increasing n_{pix} and thus noise.
- **Detector Noise:** Read noise and dark current matter especially in low-background or short-exposure regimes.

These dependencies allow schedulers to compute the required exposure time for a given SNR, or conversely, estimate the SNR achieved for a given exposure.

2.9.5 Application in Robotic Telescopes

Robotic observatories benefit greatly from dynamic SNR-based scheduling. The Automated Planet Finder (APF, Vogt et al., 2014), for instance, now uses adaptive exposure control to determine the appropriate duration in real time. It can terminate exposures early or extend them to meet a target SNR, leading to more uniform data quality.

The Skynet network is developing similar functionality, allowing users to specify an SNR target rather than exposure time (Schlekat et al., 2025). The system evaluates incoming images and determines whether additional exposures are needed to meet the requirement.

Currently such a system is developed for the New Robotic Telescope (NRT), where the scheduler and SNR engine work in tandem. Based on catalogued target brightness, measured sky brightness, current seeing, and airmass, the system calculates the required exposure time to achieve the requested SNR. If clouds or transparency variations occur, the exposure is adjusted or repeated until the goal is met.

This approach can standardise data quality, improve observing efficiency, and make the system resilient to variable conditions. Exposure times become dynamic outputs, not static inputs, making sure the telescope autonomously delivers optimal data in any observing situation.

2.10 Small Telescopes installed at the Liverpool Telescope (STILT)

The Small Telescopes Installed at the Liverpool Telescope (STILT, Mawson et al., 2013) comprise a network of dedicated imaging systems designed to provide continuous and detailed monitoring of sky conditions at the Observatorio del Roque de los Muchachos (ORM), specifically the Liverpool Telescope. This network consists of Skycam A, an all-sky camera offering 180° sky coverage; Skycams T and Z, co-pointing wide-field and narrow-field cameras, aligned with the telescope's observing direction; and Skycam O, an additional external all-sky camera positioned outside the telescope enclosure.

These instruments support the operational efficiency and scientific output of the Liverpool Telescope by providing real-time environmental images which can be used for atmospheric condition measurements as well as cloud monitoring and prediction capabilities. The STILT system is specifically investigated as part of this thesis to determine whether the data produced by these skycams can meaningfully influence and improve scheduler decisions. By aiding the scheduler through adaptive exposure timing based on atmospheric parameters, cloud coverage detection and tracking, and accurate calibration of observational data via extinction measurements, STILT can enable dynamic decision-making processes.

2.10.1 Skycam A

Skycam A is an all-sky imaging instrument providing comprehensive, horizon-to-horizon coverage of sky conditions at the Observatorio del Roque de los Muchachos (ORM). Initially equipped with a different camera system, Skycam A was upgraded in the summer of 2014 from a previous configuration to a Starlight Xpress Oculus camera. This camera features a high-quality 1.55 mm, F/2.0, 180° fish-eye lens, significantly improving its sensitivity and field of view.

Skycam A effectively monitors atmospheric transparency, detecting cloud coverage down to approximately 6th magnitude stars, making it particularly suitable for identifying reasonably thin cirrus clouds. Its sensitivity and broad coverage make it an essential tool for continuously assessing atmospheric conditions. Consequently, Skycam A's data will be used for informing telescope scheduler decisions, particularly for adaptive planning of observations to ensure optimal exploitation of clear sky conditions.

Number of Images	2,665,044
Size of DB Table	777MB
Size of Files on Disk	6.5TB
First Image	2009-02-27 20:24:13

Table 2.3: Skycam A Data

2.10.2 Skycam T

Skycam T is a medium-field imaging system mounted on the top-end of the Liverpool Telescope, designed to co-point with the main telescope during observations. The current configuration utilizes a Starlight Xpress Trius SX-35 camera equipped with a Zeiss Planar T 85mm f/1.4 ZF2 lens, which provides a pixel scale of 44 arcseconds and a field of view of approximately 24 × 16°. With a detection limit reaching down to approx. 12 mag (V-Band), Skycam T offers a good balance between resolution and sensitivity.

This camera has undergone multiple iterations over time, with different optical and detector configurations, always maintaining a field of view between 10° and 20°. The current system was installed in February 2019. Skycam T supports the cloud transparency measurement as well as the measurement of atmospheric extinction by providing real-time data. The system generates one ten-second exposure per minute when the telescope is operating.

Number of Images	2,416,499
Number of Images (current Hardware)	840,683
Size of DB Table	965MB
Size of Files on Disk	7.2TB
First Image	2009-03-04 20:19:13
First Image (current Hardware)	2019-02-06 22:20:04
Extracted Stars (since Feb '19)	$9.8 \cdot 10^9$
Matched Stars	$8.2 \cdot 10^9$

Table 2.4: Skycam T Data

Once images from Skycam T become available on the Liverpool Telescope archive, they are instantly downloaded. This approach was used to avoid interfering with the primary data pipeline while ensuring timely access. The files are saved in the FITS file format (Wells et al., 1981). The first processing step involves extracting the World Coordinate System (WCS, Greisen et al., 2002) information from the FITS header. If no WCS data is found in the file header, the image is deemed unusable — likely taken during telescope slewing, under bright conditions (i.e. just after sunset or just before sunrise), or when clouds obscured the field of view.

Each downloaded image is stored on disk, and an entry is created in the PostgreSQL database, containing metadata such as the timestamp, right ascension (RA), declination (DEC), and various image quality parameters. A flag is set to indicate whether the image is usable.

For images with a valid WCS solution, Source Extractor (SExtractor, Bertin et al., 1996) is used to detect and extract all visible stars. The extracted sources are then cross-matched against the APASS and Tycho-2 star catalogues to obtain photometric reference data. The detected sources and their corresponding catalogue matches are stored in the database, linked via foreign keys to maintain relational integrity. This systematic approach ensures that Skycam T's data can be utilised for atmospheric transparency measurements and extinction analysis.

2.10.3 Skycam Z

Skycam Z is a narrow-field imaging system mounted next to Skycam T on the Liverpool Telescope. Like Skycam T, it co-points with the main telescope, providing complementary high-resolution imaging. The current system, installed in September 2024, consists of a ZWO ASI174MM monochrome uncooled CMOS camera mounted on an Orion Optics AG8 telescope. This setup provides a field of view of approximately $0.84 \times 0.55^{\circ}$ with a plate scale of 1.62 arcseconds per pixel. Skycam Z was not operating between October 2019 and September 2024 due to a hardware failure which could not be fixed due to Covid travel restrictions.

With a detection limit reaching down to approximately 17th magnitude and completeness to around 16.5 (R-Band), Skycam Z is particularly useful for capturing high-resolution images of faint celestial objects, including galaxies and nebulae. The system is capable of producing scientifically valuable images and is an important tool for studying atmospheric conditions and extinction effects. Given its finer resolution, Skycam Z provides complementary data to Skycam T, enabling a more detailed characterisation of atmospheric extinction and local sky conditions.

Number of Images	1,121,734
Size of DB Table	950MB
Size on Disk	2.3TB
First Image	2013-01-01 18:22:14
First Image (current Hardware)	2024-09-10 23:26:10.13
Extracted Stars (since 2024)	24,571,297
Matched Stars	24,571,297

Table 2.5: Skycam Z Data

2.10.4 Skycam O

In addition to the astrophysical instruments Skycam-A, Skycam-T, and Skycam-Z, the Liverpool Telescope also utilizes Skycam O, a Raspberry Pi-based HD camera system installed outside the telescope enclosure. Designed and built by the author, Skycam-O delivers high-resolution images of the sky and surrounding environment for real-time weather and situational monitoring.

Unlike its counterparts, Skycam-O is not used for photometric measurements but instead provides contextual, ground-level observations of local conditions such as cloud cover, haze, and weather fronts. This external visual system serves as a valuable complement to the Skycam array, enhancing operational decision-making and enabling more informed human oversight during robotic operations.

A detailed technical description and performance analysis of SkyCam-O is presented in Chapter 3, Section 3.7.

2.11 Chapter Summary and Transition

This chapter has reviewed the operational and environmental context of robotic time-domain astronomy. The discussion covered the atmospheric conditions at the Observatorio del Roque de los Muchachos, existing techniques for cloud detection and extinction monitoring, models of sky brightness, and the scheduling architectures used by robotic telescopes such as the Liverpool Telescope. Together, these elements define the observational constraints under which autonomous telescopes must operate.

A recurring theme throughout this chapter is that many key parameters affecting observational efficiency (including cloud cover, atmospheric extinction, and sky background) vary on timescales comparable to individual observations. Traditional scheduling systems often treat these quantities as static constraints or rely on coarse external estimates, limiting their ability to respond optimally to rapidly changing conditions.

The remaining chapters of this thesis investigate how real-time environmental information can be incorporated directly into robotic telescope operations. In particular, the following chapters develop methods for deriving actionable atmospheric diagnostics from auxiliary instruments installed at the Liverpool Telescope and integrating these measurements into adaptive scheduling decisions.

The next chapter begins this methodological work by presenting an all-sky cloud detection system based on optical imaging. The chapter describes the design of the system, the image-processing techniques used to identify cloud structures, and the methods used to track cloud motion and generate short-term forecasts of sky conditions.

Chapter 3

Real-Time Cloud Detection and Prediction Using All-Sky Cameras

> *The sky and the sun are always there.*
> *It's the clouds that come and go.*
>
> Rachel Joyce

3.1 Motivation and Role in Robotic Scheduling

Ground-based robotic telescopes must operate autonomously under changing atmospheric conditions, without the benefit of real-time human intervention. While this autonomy enables efficient execution of queued observations, it also introduces vulnerability to environmental variability - particularly cloud cover. At sites such as the Observatorio del Roque de los Muchachos (ORM), cloud coverage conditions can vary on short timescales, and even short-lived clouds can render science exposures unusable. In such cases, the telescope will not only lose valuable observation time but may also miss the opportunity to redirect efforts to clearer regions of the sky. Often, when one part of the sky is obscured, other regions remain usable. A scheduler equipped with real-time awareness of cloud positions could dynamically prioritise observations in those areas - potentially executing lower-ranked observations that would otherwise remain unscheduled. Similarly, in the presence of passing clouds, it may be preferable to perform many short observations, each with a reasonable chance of completion, rather than initiating a longer exposure likely to be interrupted. These adaptive strategies require knowledge not only of the current cloud distribution but also of its near-term evolution.

Conventional robotic schedulers, such as the dispatch system used by the Liverpool Telescope (LT), typically operate reactively: they select the next highest-priority observation once the current one concludes, without anticipating future conditions. This strategy can be suboptimal when atmospheric conditions degrade unexpectedly or when targets are only temporarily accessible. To address these limitations, the New Robotic Telescope (NRT) is being developed with a forward-planning, look-ahead scheduler capable of evaluating multiple candidate observations several steps in advance. However, such a scheduler requires reliable short-term forecasts of sky conditions in order to make effective decisions.

Large-scale weather prediction models are unsuitable for this task due to their coarse spatial and temporal resolution. Satellite imagery, while useful for large cloud systems, lacks the cadence and local sensitivity required for high-altitude observatory conditions. Instead, local all-sky cameras provide a promising alternative: they continuously monitor the visible sky dome from the observatory's location, offering high temporal resolution and local relevance.

The goal of this chapter is to present a system that uses all-sky fisheye imagery to detect, track, and predict cloud cover in real time. The developed system aims to deliver binary cloud maps and short-term (5–30 minute) forecasts with sufficient accuracy to support adaptive, condition-aware scheduling. This includes prioritising observations in temporarily clear regions, delaying or aborting exposures at risk of interruption, and reallocating resources when cloud movement can be anticipated.

The cloud forecasting capability developed in this chapter is not a scheduler in itself, but rather an important input to any future scheduling system capable of adaptive behaviour. It acts as the "eyes" of such a system, providing real-time and short-term predictive awareness of cloud conditions with spatial and temporal resolution far beyond that of conventional weather forecasts. The implementation presented here processes all-sky image sequences from the Skycam A instrument at the Liverpool Telescope to detect, track, and forecast cloud cover. Although developed and tested at a single site, the approach is broadly applicable to other observatories, provided a continuous time series of images is available and a World Coordinate System (WCS) solution can be derived using tools such as Astrometry.net (Lang et al., 2010).

The remainder of this chapter details the system in four stages: Section 3.2 provides the general system overview, section 3.3 describes the image processing pipeline used to generate binary cloud masks; Section 3.4 presents the tracking and motion prediction framework based on geometric deprojection and Kalman filtering; and Section 3.5.2 evaluates the system's detection performance, prediction accuracy, and real-time processing capabilities. These components together provide the situational awareness needed to enable informed and dynamic scheduling decisions at robotic telescope facilities.

3.2 System Overview

3.2.1 Skycam Hardware

The cloud detection and tracking system presented in this chapter is based on images acquired by Skycam A, the north-aligned all-sky camera mounted in the enclosure of the Liverpool Telescope. This instrument provides continuous imaging of the whole night sky, primarily for monitoring weather and observing conditions. A detailed description of the Skycam hardware, including field of view and image acquisition cadence, is provided in Section 2.10.1.

Skycam A is equipped with a fisheye lens that captures nearly the entire sky dome in a single exposure. The system records unfiltered, grayscale images with a cadence of one frame per minute and stores them in a rolling archive for processing. These images form the primary input to the real-time cloud detection and prediction pipeline described in the following sections.

3.2.2 Cloud Detection Pipeline

The core of the cloud detection system is a real-time image processing pipeline that operates on a continuous stream of all-sky images acquired by Skycam A. Due to the camera's static mounting and lack of a mechanical shutter, traditional calibration techniques such as dark-frame subtraction and flat-fielding are impractical. Instead, the pipeline operates directly on raw FITS images, relying on robust relative techniques to extract cloud structures.

Each new FITS image is downloaded from the Skycam archive to a dedicated analysis server and converted into a floating-point array. A difference image is then computed by subtracting a reference frame, typically the immediately preceding image in the time series. This suppresses the static star field and instrumental signatures, isolating temporal changes attributable to cloud motion.

To reduce the influence of bright outliers such as the Moon or optical artefacts near the camera housing, pixel intensities in the difference image are clipped at a configurable maximum value (referred to as the `MAX` parameter). This threshold is tuned to suppress saturation while preserving contrast and depends on the specific camera hardware. The clipped image is then normalised and thresholded using Otsu's method (Otsu, 1979), which partitions the histogram into two classes by maximising between-class variance.

To mitigate instability in threshold selection, particularly when bright sources like the Moon move across the field, the computed threshold is compared against recent historical values. If it deviates significantly, a clipped or smoothed fallback value is used to avoid erratic segmentation behaviour.

The resulting binary image is stored as a cloud mask that encodes the presence or absence of cloud cover on a per-pixel basis. Each mask is stored together with its original FITS WCS solution, either on the filesystem or in a database. These masks, along with their timestamps and projection metadata, serve as input to the cloud tracking and forecasting system described in the next sections. They are also made available to downstream applications via a lightweight Web API.

3.2.3 Imaging Characteristics and Limitations

Skycam A captures nearly the entire sky in a single frame using a fixed-position fisheye lens. While this provides full-sky coverage at a cadence of one image per minute, it comes with several optical and photometric limitations. Due to the lack of a mechanical shutter, dark and bias frames cannot be acquired, and the wide-angle lens geometry makes flat-fielding impractical. As a result, the images exhibit residual background structure and spatially varying sensitivity, particularly near the periphery.

The effective photometric range of the system is limited. On clear, moonless nights, stars down to approximately 6th magnitude can be reliably detected. A cross-match with the Tycho2 catalogue reveals that most matches with fainter stars represent noise peaks or spurious detections. While SExtractor applies significance thresholds during source detection, near the limiting magnitude statistical fluctuations in the background, optical artefacts, and residual gradients can produce low-significance peaks that pass the detection criteria. Catalogue cross-matching therefore provides an additional astrophysical validation layer, confirming that many faint detections do not correspond to genuine stellar sources.

To estimate the instrumental zeropoint, a subset of matched stars with $V_T < 4.5\,\mathrm{mag}$ and a instrumental magnitude smaller than $-8\,\mathrm{mag}$ were used to fit a linear model. The resulting fit, shown in Figure 3.1, yields a slope of -1.83 and a zeropoint of -16.76. While the slope has no real physical meaning due to uncorrected airmass variation and lack of photometric calibration, the intercept provides a rough estimate of the system's zeropoint under ideal conditions. The relatively small number of matched calibrators

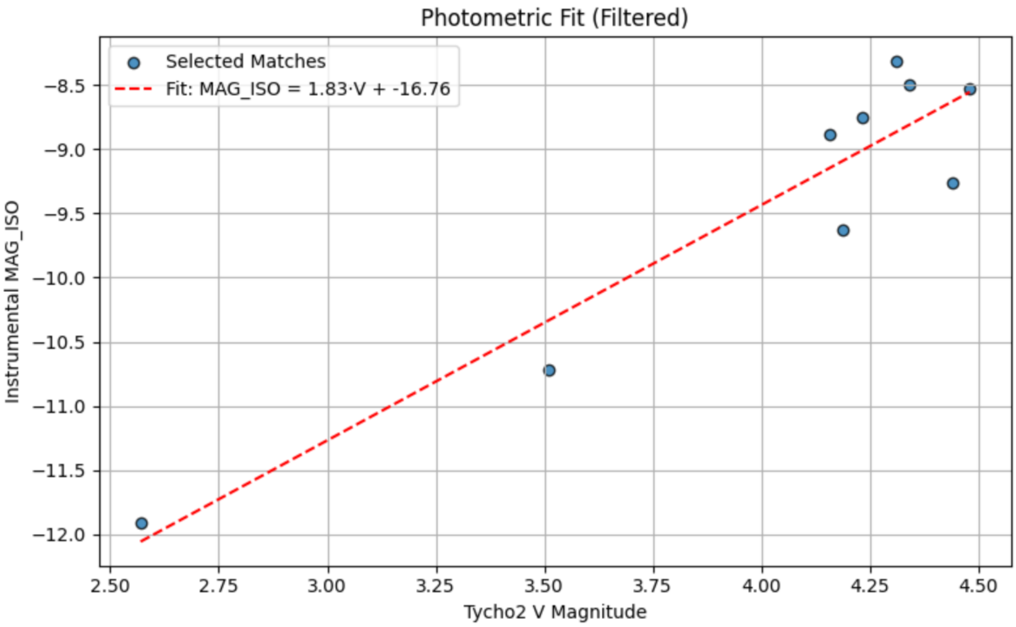

FIGURE 3.1: Instrumental magnitude `MAG_ISO` versus Tycho2 V_T magnitude for selected stars with $V_T < 4.5$ mag. A linear fit (dashed line) yields an estimated zeropoint of -16.76.

and the absence of flat-field correction contribute to the observed scatter. These limitations restrict Skycam A's use for precise photometry but are acceptable for the system's primary role in cloud detection and brightness-based masking.

3.3 Image-Based Cloud Detection

3.3.1 Preprocessing and Difference Imaging

The first stage of the cloud detection pipeline isolates temporal variations in the sky by computing the difference (see Figure 3.2) between consecutive all-sky images (see Figure 3.3). This technique leverages the stability of the celestial background relative to the rapid motion of clouds. Stars remain effectively fixed in the image due to the camera's fixed mounting, while cloud features appear as broad, evolving structures that vary between frames. By subtracting consecutive images, static elements such as stars and reflections are suppressed, highlighting transient features indicative of cloud presence.

This entire operation is executed on the graphics processing unit (GPU) using a single-pass GLSL shader. In this context, "single-pass" denotes that the pixel-wise computation is carried out in one rendering step, with each pixel processed independently and in parallel, allowing the full image difference to be computed in real time. The shader takes as input two floating-point textures, each representing a one-channel `float32` version of a FITS image - namely, the current image and the one acquired one minute earlier. The shader applies the FITS scaling parameters `BSCALE` and `BZERO` to both images using uniform variables `bscale1`, `bzero1`, `bscale2`, and `bzero2`. The shader computes the absolute difference between corresponding pixel values in the current and previous image, effectively measuring the magnitude of local variation while discarding the direction of change. The GLSL (OpenGL Shading Language) shader-code is presented in listing 3.1.

To reduce the influence of saturated regions caused by bright sources such as the Moon or optical artefacts, the difference image is clipped at a configurable threshold passed as a uniform (`clip_value`). This removes extreme outliers while preserving meaningful variations across the sky. The result is a detection-ready difference image in which clouds appear as coherent regions of nonzero intensity, suitable for thresholding and segmentation.

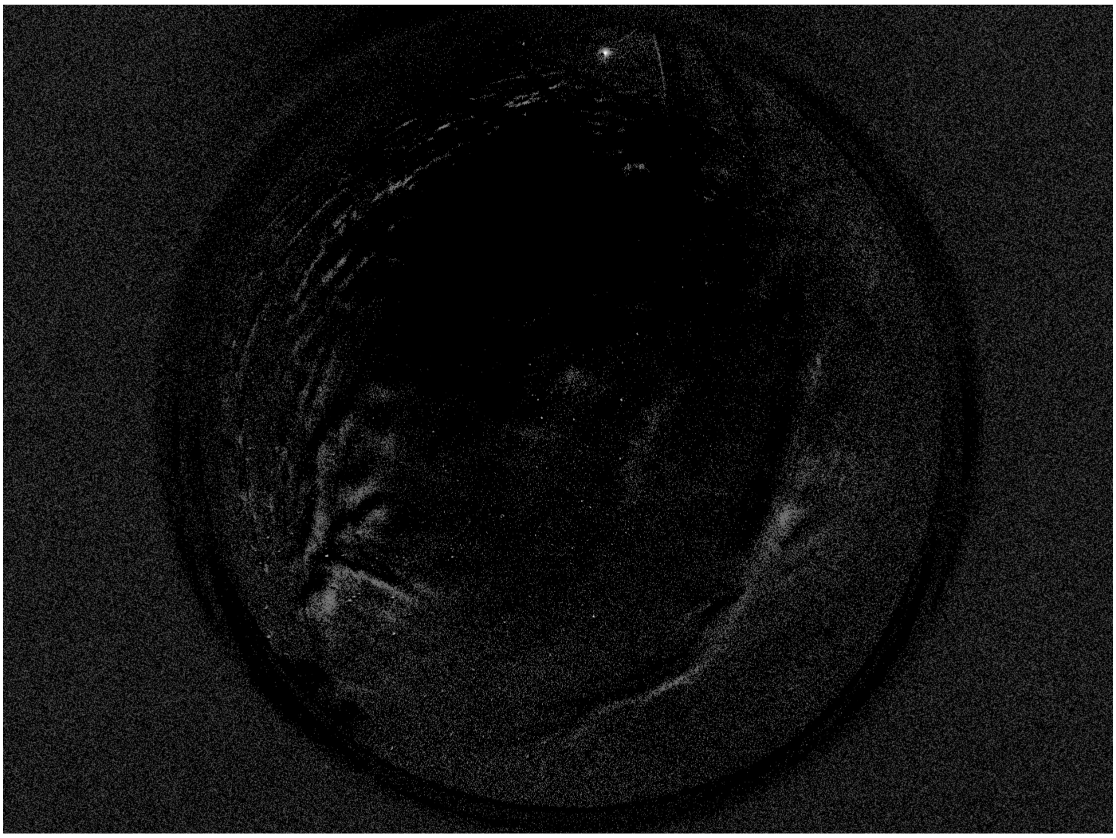

FIGURE 3.2: Difference image obtained by subtracting the image shown in Figure 3.3 from an image taken one minute earlier. Most of the structure visible in this image arises from the motion of cloud features between the two exposures. In the limiting case of completely stationary and non-evolving clouds the difference image would approach zero. In practice, however, some motion is always present due to atmospheric winds. If cloud motion is very slow, the effective time baseline between frames can be increased by skipping intermediate images. Apparent star trails caused by longer time gaps are removed during the morphological and heuristic filtering steps described later in this section.

```
uniform sampler2D image_now;
```

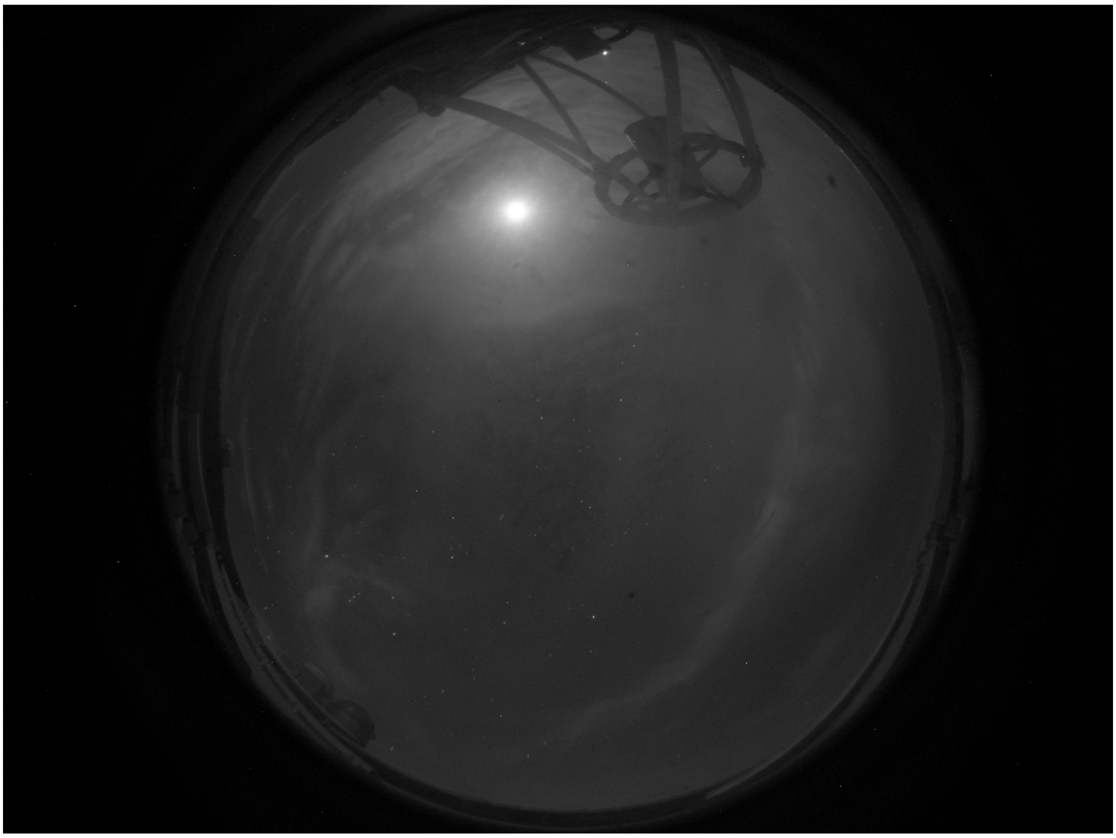

Figure 3.3: All-sky image taken using the Skycam A Camera at the Liverpool Telescope site at 22:45 UT on 2020 Jan 01. Some stars are visible as well as bands of clouds, with the cloud illuminated by the Moon at the top of the image

```
uniform sampler2D image_prev;
uniform float bscale1, bzero1;
uniform float bscale2, bzero2;
uniform float clip_value;

in vec2 texCoord;
out float out_diff;

void main() {
    float val1 = texture(image_now, texCoord).r * bscale1 + bzero1;
    float val2 = texture(image_prev, texCoord).r * bscale2 + bzero2;
    float diff = abs(val1 - val2);

    // Clip to configured threshold
    diff = clamp(diff, -clip_value, clip_value);

}
```

Listing 3.1: GLSL fragment shader used for real-time difference imaging and clipping

3.3.2 Otsu Thresholding

Following the construction of a normalised difference image, the next stage involves segmentation of the sky into regions of significant temporal change (potential clouds) and relatively static areas (clear sky). This is achieved through Otsu's method (Otsu, 1979), a non-parametric global thresholding algorithm. The method first computes the histogram of pixel intensities in the difference image and assumes that it contains two underlying classes: low-variation pixels (clear sky) and high-variation pixels (cloud motion). For every possible threshold value, the histogram is divided into these two classes, and the between-class variance is computed. The optimal threshold is defined as the value that maximises this between-class variance, thereby producing the strongest statistical separation between static and dynamic regions.

The algorithm assumes a bimodal distribution of pixel intensities within the difference image: one peak associated with low variation (clear regions), and another corresponding to significant changes (clouds, artefacts, or bright scattering sources). Let the normalised intensity range be divided into L histogram bins indexed by $i = 1, \ldots, L$, and let p_i denote the probability of a pixel falling into bin i. A threshold value k partitions the histogram into two classes:

- $C_0 = \{1, \ldots, k\}$: low-change regions,
- $C_1 = \{k + 1, \ldots, L\}$: high-change regions.

The cumulative class probabilities and class means are given by:

$$\omega_0(k) = \sum_{i=1}^{k} p_i, \qquad \omega_1(k) = \sum_{i=k+1}^{L} p_i = 1 - \omega_0(k),$$

$$\mu_0(k) = \frac{1}{\omega_0(k)} \sum_{i=1}^{k} i \cdot p_i, \qquad \mu_1(k) = \frac{1}{\omega_1(k)} \sum_{i=k+1}^{L} i \cdot p_i.$$

The between-class variance is then defined as:

$$\sigma_B^2(k) = \omega_0(k)\omega_1(k) \left[\mu_0(k) - \mu_1(k)\right]^2 . \tag{3.1}$$

The optimal threshold k^* maximises $\sigma_B^2(k)$ and is used to classify the difference image into cloud and non-cloud regions. All pixels with values above this threshold are marked as cloud-pixels.

To increase the robustness of the segmentation under variable sky conditions - such as rapid illumination changes due to the telescope moving in the line-of-sight between the Skycam A and the Moon or scattered light - the selected threshold is compared against a temporal history of recent threshold values. If a strong deviation is detected, a clipped or damped threshold may be substituted to reduce the impact of transient artefacts and to suppress segmentation instability.

This step is implemented on the CPU, as histogram analysis and variance optimisation are not efficiently parallelised for single-frame operations on the GPU. There are OTSU implementations on the GPU (e.g. Zhu et al. (2020)) but in the case presented here, there would be no speed benefit compared to a CPU based version due to multiple data copy operations and costly shader program switches. For colour images and images with much higher resolution than the ones provided by Skycam A, a GPU implementation might become feasible.

The resulting binary mask is then passed to the morphological filtering stage.

3.3.3 Morphological Processing and Component Filtering

The binary mask resulting from thresholding often contains noise artefacts, isolated pixels, bright (and thus large) stars, and incomplete cloud structures due to residual detector effects or faint stellar sources.To improve the spatial coherence and reliability of the detected regions, morphological image operations are applied. These operations, which modify binary images using structured neighbourhood transformations such as erosion and dilation, are designed to suppress small-scale features unlikely to correspond to cloud structures while enhancing connected regions consistent with cloud morphology.

Two classical operations are employed: erosion and dilation. Erosion removes isolated white pixels by requiring that all pixels within a defined neighbourhood satisfy the mask condition. Dilation, by contrast, expands the boundaries of detected regions. When applied in sequence (dilation followed by erosion, see equation 3.2), this process is known as closing and serves to fill small holes and connect nearby regions. Conversely, an erosion followed by dilation (opening, 3.3) can be used to remove small objects entirely. In this work, an opening operation is used to eliminating small-scale noise.

Let B be a structuring element (typically a square or disk of fixed radius), and let I denote the binary input image. The morphological dilation $\delta_B(I)$ and erosion $\varepsilon_B(I)$ are defined as:

$$\delta_B(I)(x, y) = \max_{(u,v)\in B} I(x - u, y - v),$$
$$\varepsilon_B(I)(x, y) = \min_{(u,v)\in B} I(x - u, y - v).$$

The closing operation is then expressed as:

$$I_{\text{closed}} = \varepsilon_B(\delta_B(I)). \tag{3.2}$$

And the respective opening operation as:

$$I_{\text{opened}} = \delta_B(\varepsilon_B(I)). \tag{3.3}$$

These operations are implemented on the GPU using fragment shaders for real-time performance. A fragment shader is a program executed once per pixel during the rendering process, allowing the morphological transformation to be applied independently and in parallel across the entire image.

Two discrete square kernels are used as structuring elements in the GPU implementation: a 3×3 (see Figure 3.4) and a 5×5 binary matrix. These define the local neighbourhood over which erosion and dilation are computed. Their structure is shown below:

$$\begin{array}{cc} 3 \times 3 \textbf{ Kernel} & 5 \times 5 \textbf{ Kernel} \\ \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{array}$$

Both kernels operate in pixel space and are isotropic in nature. The 3×3 kernel is preferred when fine detail preservation is critical, whereas the 5×5 kernel offers more aggressive smoothing and noise rejection. The choice of kernel size is configurable and depends on the expected spatial scale of cloud structures in the input image.

```
uniform sampler2D input_mask;
uniform ivec2 kernel_size;  // e.g., ivec2(1,1) for 3x3, ivec2(2,2) for 5x5
uniform int mode; // 0 = erosion, 1 = dilation
uniform vec2 texel_size; // 1.0 / image resolution

```

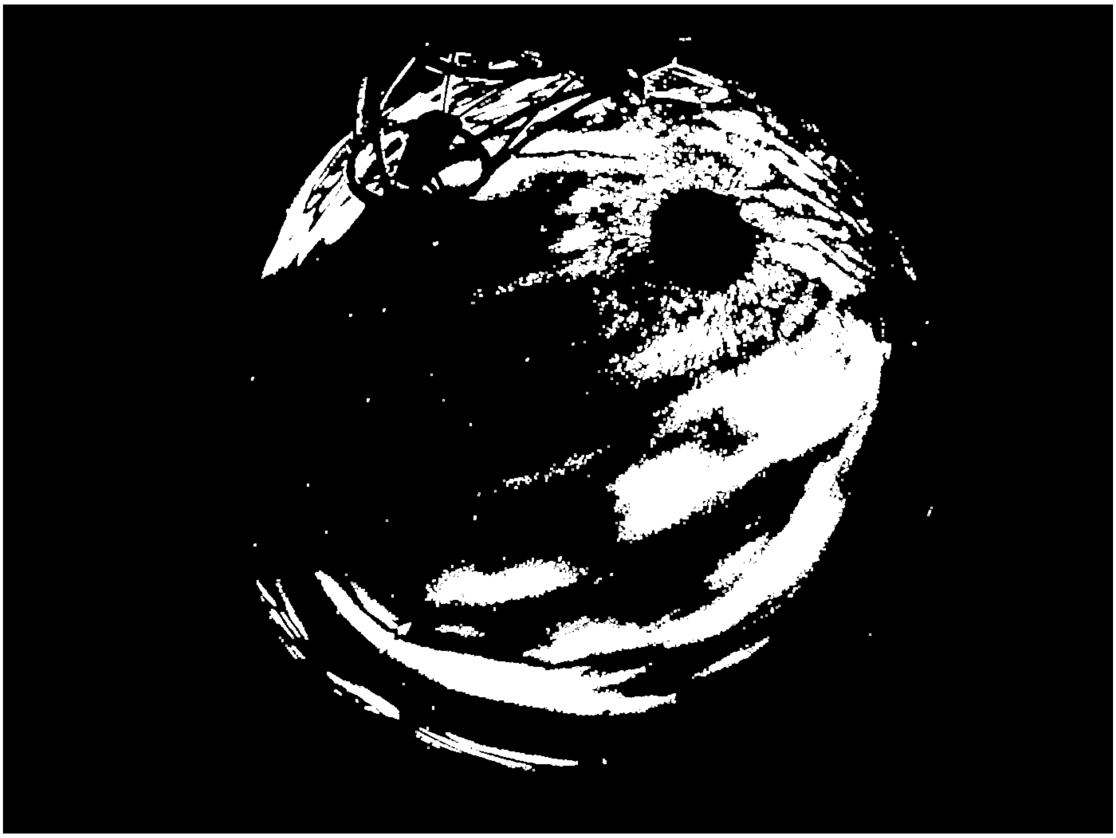

FIGURE 3.4: Binary image obtained after the application of the morphology methods erosion and dilation with a 3 × 3 filter kernel. The number of star pixels is reduced whilst the cloud structure remains intact. Larger filter kernels (e.g. 5 × 5 or 7 × 7 will remove all stars but also small cloud structures like the ones present in the middle of the image

```
in vec2 texCoord;
out float result;

void main() {
    float extreme = (mode == 0) ? 1.0 : 0.0;

    for (int dx = -kernel_size.x; dx <= kernel_size.x; ++dx) {
        for (int dy = -kernel_size.y; dy <= kernel_size.y; ++dy) {
            vec2 offset = vec2(dx, dy) * texel_size;
            float value = texture(input_mask, texCoord + offset).r;

            if (mode == 0) {
                extreme = min(extreme, value); // erosion
            } else {
                extreme = max(extreme, value); // dilation
            }
        }
    }

    result = extreme;
```

```
}
```

Listing 3.2: GLSL fragment shader for morphological erosion or dilation using a square kernel.

3.3.4 Binary Cloud Mask Output

After morphological processing, the resulting binary mask typically contains significantly fewer star-like artefacts. However, small residual detections originating from bright stars or image noise may still remain. While the use of larger structuring elements (e.g., 5×5 or 7×7 pixel kernels) can further suppress these detections, such an approach risks degrading the fidelity of the actual cloud structures and increases computational complexity.

To maintain processing speed and retain compatibility with GPU-based execution, a lightweight heuristic filtering step is applied. Here, “heuristic” denotes that the filtering rule is based on simple, empirically chosen logical criteria rather than on a formally optimal segmentation algorithm. This method examines the local connectivity and isolation of each foreground region using only simple logic operations over a small window.

For each foreground pixel (pixels with value $= 1$), a 21×21 pixel neighbourhood centred on that pixel is extracted. The purpose of this window is to determine whether the detected structure forms a small, isolated component or is part of a larger, extended feature. Rather than performing a full connected-component analysis, the algorithm inspects continuity in the horizontal and vertical directions. Rows and columns containing foreground pixels are identified, and gaps of entirely background pixels between such rows or columns are interpreted as evidence for multiple separate components within the window. In that case, the structure is left unchanged.

If the window contains only a single contiguous component, and that component does not touch the window boundary (indicating that it is fully contained within the local region), its total area is measured. If this area falls below a configurable threshold (typically 20 pixels), the component is removed from the mask. This effectively suppresses small, isolated stellar detections while preserving extended cloud structures.

While this method does not guarantee the removal of all stellar detections, particularly those with non-convex or elongated shapes, it has proven highly effective in practice. The vast majority of remaining star-like features are eliminated with minimal impact on cloud morphology. In contrast, recursive flood-filling or global connected-component labelling algorithms identify all contiguous pixel regions in the entire image by iteratively propagating connectivity through neighbouring pixels. While more rigorous, such methods require additional memory and control flow operations that are less suitable for efficient GPU-based implementation. (see Figure 3.5).

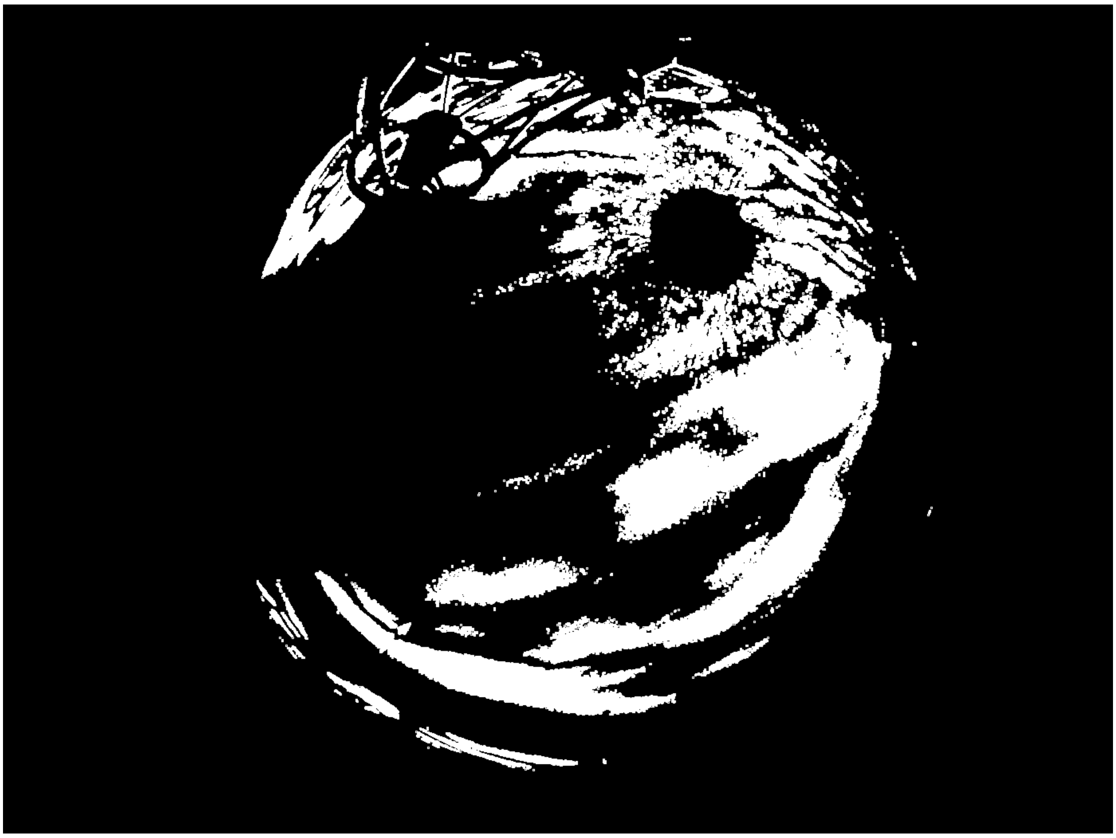

FIGURE 3.5: Binary image obtained after an additional star removal step, a simple heuristic is applied to a 21 × 21 pixel window moved over the whole image, again leveraging fast GPU hardware, to test for small star-shaped pixel areas

3.3.5 Storage and Query

Each binary cloud mask generated by the detection pipeline is stored as a `float32` FITS image, with pixel values of either 0.0 (clear) or 1.0 (cloud-covered). The original World Coordinate System (WCS) solution derived from the corresponding all-sky image is retained and embedded in the FITS header, allowing pixel-based cloud data to be projected onto the celestial sphere for subsequent analysis or visualisation.

Cloud masks and their associated difference images are saved to disk following a timestamp-based directory hierarchy. For quick inspection and web-based visualisation, corresponding PNG versions are generated in parallel.

Metadata for each cloud image is stored in a structured PostgreSQL database. The main table `AllskyImage` records the original FITS and JPEG filenames, storage paths, and acquisition timestamps. Each cloud mask is linked to its source image via the `Cloudpattern`

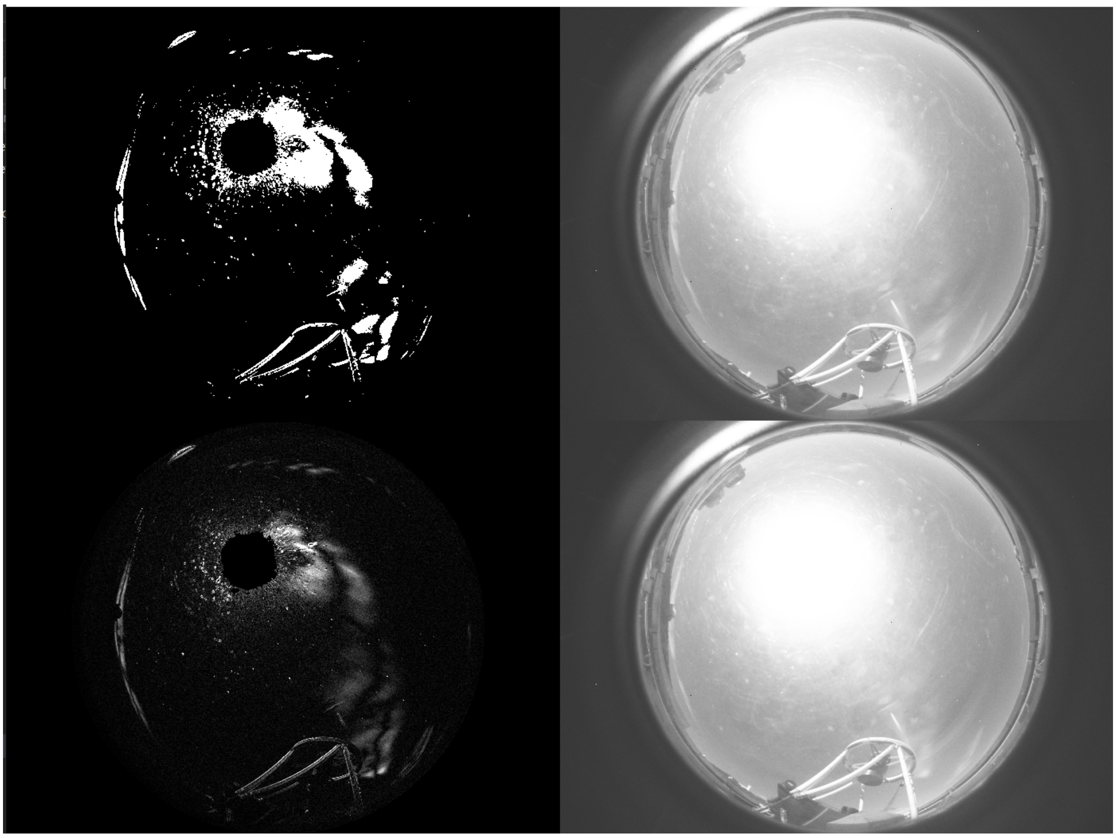

Figure 3.6: Principle work of the algorithm: The right column shows the input images at t and $t+1min$, the resulting difference image is shown on the bottom left and finally, the binarised cloud map is shown in the top left

table, which stores the Otsu threshold value used for segmentation (`OTSUThreshold`), the cloud mask filename and path, and a reference to the parent `AllskyImageId`.

In addition to raw filenames and thresholds, a derived parameter called *total cloudiness* is computed and stored. This value represents the ratio of cloud pixels (value 1.0) to total pixels within a predefined area of interest. In the case of the Liverpool Telescope installation, this area is defined as a circular region centred on the image with a radius corresponding to an altitude of 20°. This measure serves as a compact summary of overall sky coverage and is used to support scheduling decisions and cloud forecast scoring.

Cloud masks and associated metadata can be accessed programmatically via a RESTful Web API, allowing external systems (e.g. schedulers or visual dashboards) to retrieve cloud maps by timestamp or image ID. The API is exposed under the route `/api/pattern` and supports filtering based on Right Ascension and Declination, observation time, and coverage radius.

This architecture supports both archival access and real-time integration, providing a robust interface for downstream applications requiring cloud-state awareness. The whole principle of the presented algorithm can be seen in Figure 3.6.

3.4 Cloud Tracking and Forecasting

The cloud masks produced in the detection stage represent individual, time-stamped snapshots of sky coverage. In order to support short-term planning and adaptive scheduling, these masks must be transformed into a temporal sequence that captures cloud motion across the sky. This section describes the tracking system developed for this purpose.

The primary goal is to identify individual cloud structures and track their motion over time in order to predict their future positions. While static cloud detection provides a binary indication of sky visibility at a given moment, cloud tracking enables forward-looking scheduling decisions by estimating whether a given sky region will remain clear during the execution of an upcoming observation. Most observing groups at the Liverpool Telescope require approximately 10 to 15 minutes of execution time (including slewing, setup, and readout). As such, the system was designed to produce forecasts with a baseline horizon of 15 minutes. This enables the scheduler to make informed decisions about whether a group should be observed immediately or deferred in favour of a target that will remain cloud-free for longer.

The tracking system consists of three key stages: (1) projection of detected cloud regions from fisheye image space into a deprojected (“defisheyed”) space, (2) identification and temporal linking of coherent cloud structures across successive frames, and (3) motion estimation and short-term forecasting using a Kalman filter (Kalman, 1960). The output is a time series of labelled cloud features with predicted positions.

Meteorological observations support the underlying assumption of relatively stable cloud motion over short timescales. Wind direction and speed in the relevant atmospheric layers - particularly within the boundary layer, where clouds typically form - are generally governed by large-scale pressure gradients and topographic influences. These factors evolve slowly under stable conditions and do not exhibit abrupt directional changes unless disturbed by passing fronts or convective activity.

Empirical analysis of all-sky image sequences from the Skycam A system at the Liverpool Telescope, as well as independent recordings from an all-sky camera in Emleben (Thuringia, Germany), showed that the shortest interval during which the prevailing cloud direction changed by more than 60° was between two and three hours. This magnitude of directional change corresponds to an average rotation rate of only 0.3° per minute, which is consistent with the slow and steady drift of stratiform cloud types such as stratocumulus or altostratus under operationally safe weather conditions.

The assumption of constant cloud velocity vectors over short forecast intervals is further supported by work in the field of solar energy forecasting. Bosch et al. (2013) showed that cloud tracking models based on constant motion assumptions remain valid for prediction horizons of at least 15 minutes. Based on these considerations, it is deemed reasonable to adopt a similar assumption for astronomical cloud forecasting over timescales of 15 to 45 minutes.

To support real-time operation, the system includes a GPU-accelerated deprojection stage that transforms all-sky images into an undistorted representation of the sky. The transformation is performed in a fragment shader using empirically calibrated parameters, enabling each labelled pixel to be expressed in an unprojected coordinate space. The remainder of this section describes the projection, tracking, and prediction stages in detail.

3.4.1 Fisheye De-Projection

The deprojection step is implemented as a GPU fragment shader for real-time performance. The shader removes radial distortion by mapping each pixel of the fisheye image onto a corrected plane using an inverse projection model based on the arctangent function. The relevant GLSL implementation is shown in Listing 3.3.

All parameters, such as the optical centre, scale factor, and pixel scaling, were determined empirically through visual calibration. This involved aligning bright sources (e.g., the Moon, known stars) and adjusting the parameters to minimise residual geometric error.

A more rigorous calibration could, in principle, be achieved through a geometric approach using a known fiducial pattern such as a chessboard. In this method, the camera would be pointed at a planar calibration target taking images from multiple angles and distances. The intrinsic lens parameters (including distortion coefficients) would then be recovered using standard photogrammetric techniques (e.g., OpenCV's calibration routines[1]). However, this procedure requires daytime exposures under controlled conditions.

Such an approach was deemed infeasible for the Skycam A system, which is permanently mounted at the Liverpool Telescope in a weatherproof housing. Performing chessboard calibration in situ would have required deployment of a calibrated pattern at known positions and distances relative to the camera, likely involving scaffolding or elevated platforms. Additionally, remote reconfiguration of the acquisition system would have

---

[1] `https://docs.opencv.org/4.x/dc/dbb/tutorial_py_calibration.html`

been necessary to enable short daytime exposures, followed by restoration of the original nighttime settings. These operational and logistical constraints rendered physical calibration impractical in this context.

Nevertheless, the empirical method applied here yielded sufficiently accurate deprojection results for the intended application. Visual alignment of celestial sources indicates that geometric errors in the corrected coordinate map remain below roughly 0.5˘1° across the central field of view. Given that cloud structures typically extend over several degrees, this level of geometric uncertainty does not significantly affect the subsequent cloud detection or motion estimation.

The resulting coordinate transformation enables accurate reprojection of cloud-labelled pixels into horizontal coordinates, which is a prerequisite for consistent motion tracking across frames.

As the all-sky image covers a 180° field of view, the physical area represented by each pixel varies dramatically across the field. Pixels near the zenith correspond to relatively small regions of the atmosphere, whereas pixels near the horizon project to much larger areas on the ground or lower sky, often exceeding several hundred square metres per pixel. This distortion arises from the projection geometry: while angular resolution in square degrees remains approximately uniform across the image, the corresponding physical footprint increases sharply toward lower altitudes.

The defisheye transformation therefore introduces non-uniform scaling in the unprojected image, leading to over- or under-representation of cloud features depending on their location. To mitigate these effects and reduce edge-related artefacts, cloud identification and tracking are restricted to the region above 20° altitude. Pixels below this threshold are excluded from further processing.

```
#version 330 core

in vec2 fUv; // Normalised fragment coordinates (0.0..1.0)
out vec4 FragColor;

uniform sampler2D tex0;       // Input fisheye texture
uniform float pixelscale;     // Pixels per angular unit
uniform float scale;          // Global scaling factor (empirically calibrated)

const vec2 centre = vec2(696.0, 520.0); // Optical center in pixels
const vec2 resolution = vec2(1392.0, 1040.0);
const float PI = 3.1415926538;

void main() {
    // Convert UV to pixel coordinates
    vec2 imgPos = fUv * resolution;
    vec2 delta = imgPos - centre;

    // Radial distance from image centre
    float r = length(delta);
    float rt = r / pixelscale;

    // Inverse fisheye projection (arctangent model)
    float theta = atan(rt) / rt;

    // Apply radial correction and rescale
    vec2 corrected = centre + delta * theta * scale;

    // Back to UV coordinates
    vec2 sampleUv = corrected / resolution;

    // Sample corrected pixel
    FragColor = texture(tex0, sampleUv);
}
```

Listing 3.3: GLSL fragment shader for fisheye deprojection to remove radial distortion.

3.4.2 Labeling and Feature Identification

Once the binary cloud masks have been deprojected, each individual cloud feature must be isolated and uniquely identified. This is accomplished using a connected-component labelling (CCL) algorithm, which assigns a distinct label to each contiguous region of foreground pixels. In this work, the so-called *Spaghetti Labelling* algorithm, based on directed acyclic graphs and introduced by Bolelli et al. (2020), is employed. Although GPU-based variants of this method have been proposed (Bolelli et al., 2021), the CPU implementation was found to be sufficiently fast given the moderate image sizes, the image cadence and short processing times.

The CCL algorithm scans the binary mask and assigns labels to connected regions that represent individual cloud patches. For each labelled region, its centroid, bounding box, and pixel area are computed. These features are stored for each cloud map in the sequence and form the basis for motion tracking and prediction.

To eliminate noise and reduce computational load in subsequent stages, a minimum area filter is applied. Cloud components with fewer than 40 connected pixels are discarded, under the assumption that such regions are either residual stars or fragmented artefacts of the thresholding process. Future work may explore adaptive merging strategies to combine nearby small components with larger structures.

In practice, the segmentation process occasionally produces spurious detections caused by non-atmospheric features. Two recurrent issues were identified: (1) contamination from the telescope structure, especially during slews, and (2) overexposure or scattered light from the full Moon. Both effects can result in large, rapidly appearing "cloud-like" structures that are not consistent with the motion or morphology of real clouds. Figure 3.7 illustrates an example where the telescope structure was erroneously labelled as a cloud due to bright moonlight and rapid motion.

To mitigate these effects, detections in regions known to be affected by the telescope structure or the Moon are excluded from tracking. The Moon's position and phase are computed using standard ephemerides (Meeus, 1998), and used to define a dynamic exclusion zone in each frame. The region occupied by the telescope is estimated using WCS information derived from images recorded by Skycam T, a piggy-back instrument rigidly attached to the Liverpool Telescope structure (Mawson et al., 2013). Any cloud centroid falling within the Moon or telescope exclusion zones is flagged and excluded from further analysis.

Figure 3.8 shows an example cloud mask with labelled patches and their bounding boxes overlaid, illustrating the output of the component labelling process and subsequent filtering.

3.4.3 Temporal Matching of Cloud Patches

Before applying motion estimation, cloud patches must be associated across successive frames to establish temporal continuity. This association step matches labelled cloud

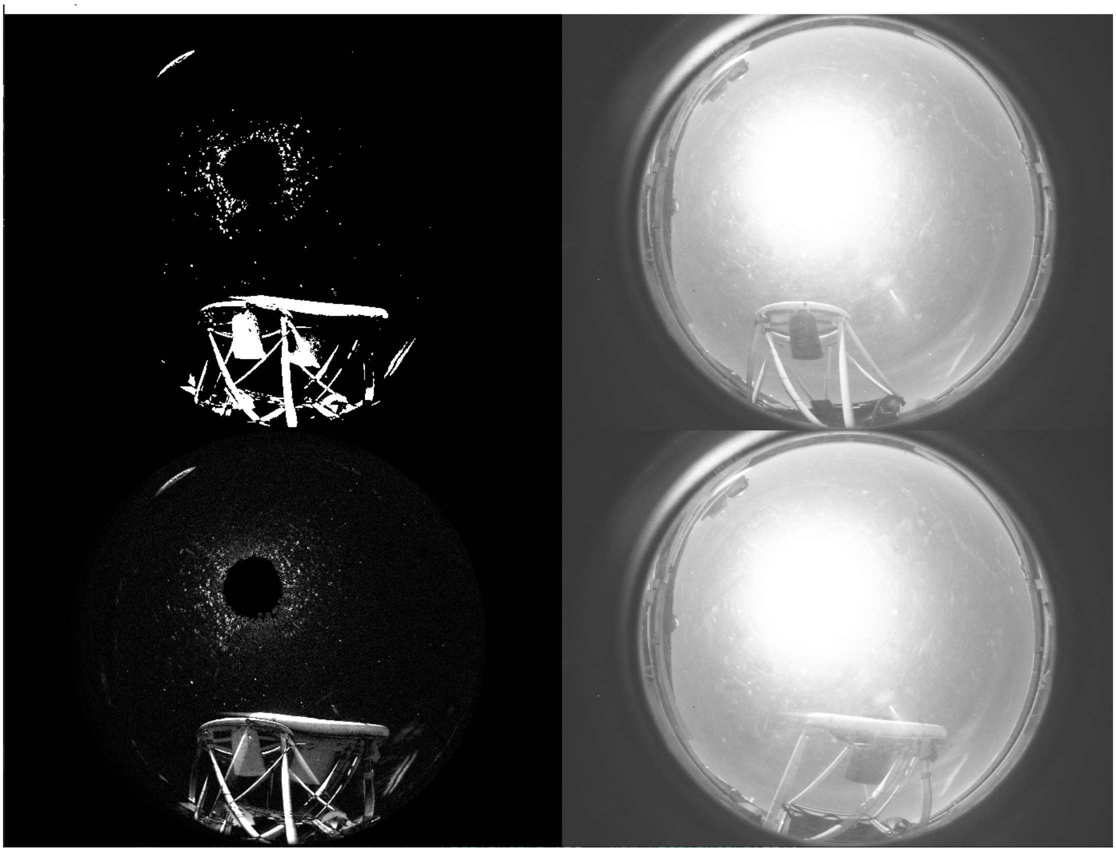

Figure 3.7: The fast slewing of the telescope and the full moon on the telescope structure caused the algorithm to detect the structure as a quite large cloud. The resulting cloud map can not be used for large parts of the sky. To detect this issue, large, sudden appearances of "cloud-like structures" will be discarded

components from the current frame to those detected in the previous one, enabling the construction of individual cloud trajectories over time.

For each patch in frame t, its centroid coordinates are compared to those of patches in frame $t-1$ using Euclidean distance. A nearest-neighbour approach is used, with a configurable maximum displacement threshold to avoid implausible associations. This threshold is based on expected wind speeds and frame cadence, and typically limits matching to displacements of less than 150 pixels per minute.

If multiple candidates fall within this threshold, the match with the smallest spatial distance is selected. This nearest-centroid assumption is justified by the relatively short temporal separation between frames (typically of order one minute), during which cloud structures are expected to undergo only limited displacement relative to their spatial extent. In cases where no suitable match is found (i.e. due to occlusion, segmentation error, or new cloud formation), a new track is initialized. Conversely, unmatched patches from the previous frame are assumed to have either dissipated or become temporarily

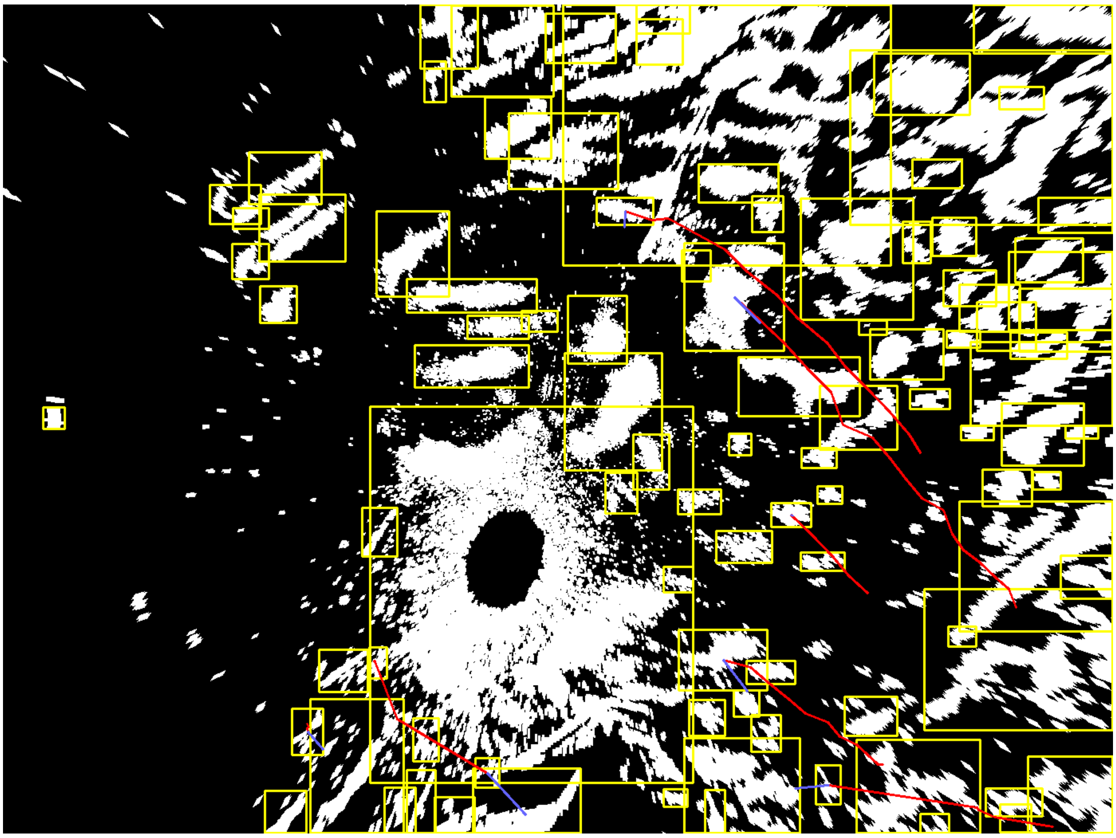

Figure 3.8: Detected cloud-like structures and their motion tracks in a deprojected (fish-eye corrected) cloud map. Yellow rectangles indicate the bounding boxes of connected cloud components detected in the current frame. Red curves represent the estimated motion tracks of these components, obtained by linking centroid positions across consecutive previous frames. Individual historical centroid positions are not marked explicitly; rather, the red line traces their temporal sequence. The large structure surrounding the Moon, caused by dust scattering and over-exposure, is also detected as a cloud-like region. However, its derived motion vector differs significantly from the dominant cloud flow and is therefore rejected during the filtering stage.

undetected; their trajectories are extrapolated by the Kalman filter for a limited number of frames before being removed.

This temporal matching ensures that each Kalman filter instance corresponds to a unique physical cloud feature with a consistent identity across time, allowing for stable velocity estimation and accurate short-term forecasting.

3.4.4 Track Smoothing

Before feeding cloud tracks into the Kalman filter, an optional geometric smoothing step can be applied to improve trajectory stability. This step leverages the assumption that clouds typically move in a consistent direction over short timescales, and erratic motion between frames is likely due to segmentation noise or partial occlusion.

Each track, composed of a sequence of centroid positions $(\vec{p}_1, \vec{p}_2, \ldots, \vec{p}_n)$, is evaluated for directional consistency by examining the angular deviation between successive motion vectors:

$$\vec{v}_i = \vec{p}_{i+1} - \vec{p}_i.$$

The angle θ_i between vectors $\vec{v}_i$ and $\vec{v}_{i+1}$ is calculated using the inner product:

$$\theta_i = \arccos\left(\frac{\vec{v}_i \cdot \vec{v}_{i+1}}{|\vec{v}_i||\vec{v}_{i+1}|}\right).$$

Tracks that contain angles exceeding a specified threshold (typically 10°) are flagged as noisy and excluded from prediction. This helps suppress false tracks caused by transient artefacts or ambiguous segmentation.

For the remaining tracks, a smoothing kernel is applied to reduce high-frequency jitter. A simple Gaussian filter is used to interpolate centroid positions over time:

$$\vec{p}_{i_{\text{smoothed}}} = \sum_{j=-k}^{k} w_j \cdot \vec{p}_{i+j},$$

where w_j are the Gaussian weights and k defines the window size (typically $k = 2$). This produces a refined trajectory (see Figure 3.9) that better reflects the true physical motion of the cloud and improves the robustness of subsequent velocity estimation.

This smoothing is particularly useful in marginal conditions or for partially obscured clouds near the horizon, where segmentation is less reliable. While not strictly required, it offers a practical trade-off between accuracy and complexity and improves the robustness of subsequent velocity estimation.

3.4.5 Kalman-Based Motion Tracking

To estimate the future positions of individual cloud patches and enable short-term forecasting, each labelled component is tracked across successive frames using a discrete Kalman filter. The Kalman filter provides a computationally efficient and statistically optimal method (under Gaussian noise assumptions) to estimate the evolving state of a dynamic system from noisy measurements.

Each cloud patch is modelled as a point object moving with approximately constant velocity in the deprojected coordinate system (altitude h, azimuth A). The state vector for a given patch at time t is defined as:

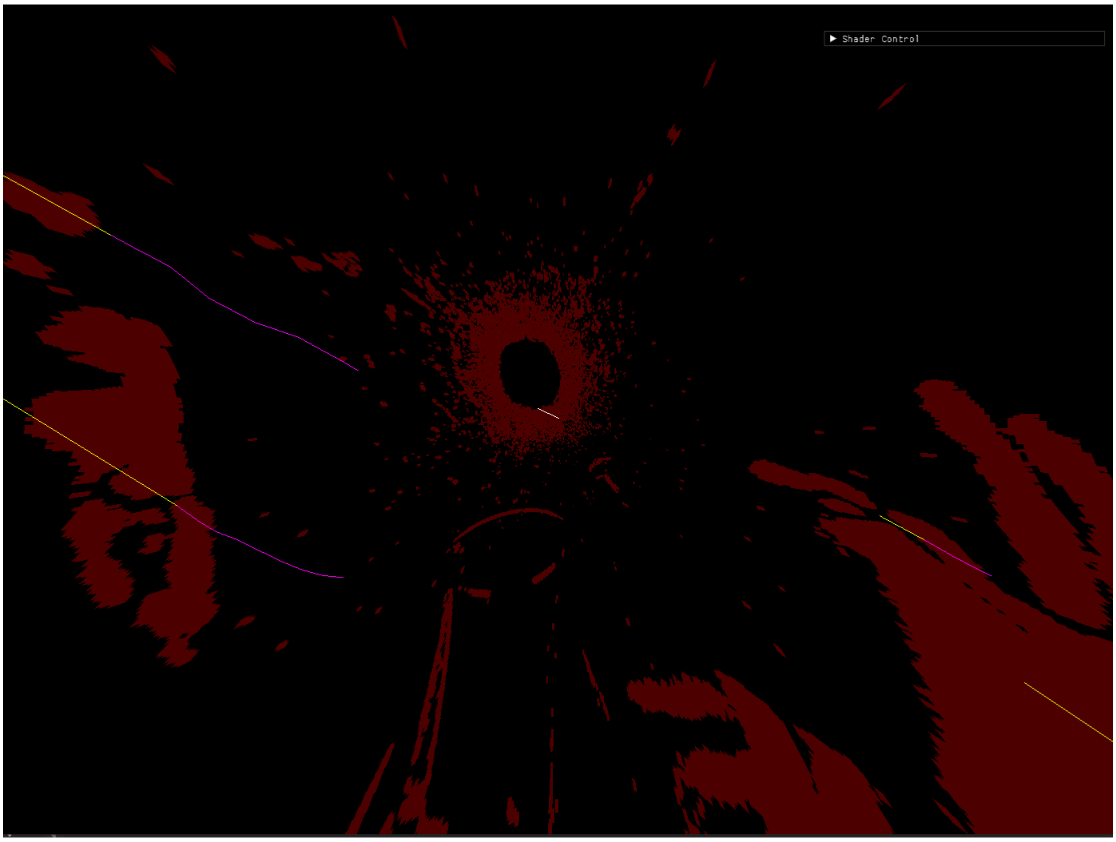

FIGURE 3.9: Image showing resulting cloud paths (magenta) with their predictions over 15 minutes (yellow). The obtained cloud paths are filtered to get smooth paths and then analysed to filter paths belonging to other structures (e.g. Moon, telescope structure) The remaining paths are shown here and used for the movement prediction

$$\mathbf{x}_t = \begin{bmatrix} A_t \\ h_t \\ v_{A,t} \\ v_{h,t} \end{bmatrix},$$

where (A_t, h_t) are the measured coordinates of the patch centroid in degrees, and $(v_{A,t}, v_{h,t})$ are the respective velocity components in pixels/step. Observations are derived directly from the labelled binary cloud masks as described in the previous subsection.

The filter operates in two stages: prediction and correction. In the prediction step, the state is propagated forward using a constant-velocity model:

$$\mathbf{x}_{t+1}^{-} = \mathbf{F} \cdot \mathbf{x}_t + \mathbf{w}_t, \quad \mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

where Δt is the time between frames and $\mathbf{w}_t$ is zero-mean process noise with covariance $\mathbf{Q}$. The matrix $\mathbf{F}$ encodes the physical assumption of linear motion with constant velocity.

When a new observation $\mathbf{z}_t = [A_t^{\text{obs}}, h_t^{\text{obs}}]^T$ is available from the current frame, the predicted state $\mathbf{x}_{t+1}^-$ is updated:

$$\mathbf{x}_{t+1} = \mathbf{x}_{t+1}^- + \mathbf{K}_t \cdot (\mathbf{z}_t - \mathbf{H} \cdot \mathbf{x}_{t+1}^-),$$

where $\mathbf{H}$ is the observation model:

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix},$$

and $\mathbf{K}_t$ is the Kalman gain, computed from the predicted error covariance and the observation noise covariance $\mathbf{R}$. This step adjusts the prediction based on the observed position of the patch centroid.

If no matching patch is found in a subsequent frame (e.g., due to partial occlusion or segmentation noise), the Kalman filter continues to propagate the cloud's position forward based on its internal velocity estimate, up to a configurable number of missed frames. This allows for temporally robust tracking even in the presence of incomplete observations.

The filter is initialized when a new cloud patch appears and matched to previously tracked patches using nearest-neighbour distance, constrained by motion consistency. Once initialized, the filter provides both a smoothed estimate of the cloud position and a forecast of its future location. These forecasts form the basis for the spatial prediction masks described in the following subsection.

3.4.6 Future Cloud Position Estimation

Once individual cloud patches have been matched and tracked across frames using the Kalman filter, their estimated state vectors provide not only smoothed positions but also velocity components. These can be used to predict the future location of each patch at arbitrary time horizons, typically 10 to 30 minutes ahead.

The predicted position $\mathbf{x}_{t+\Delta t}$ of each cloud patch is computed by propagating the estimated state forward using the motion model:

$$\mathbf{x}_{t+\Delta t} = \mathbf{F}^n \cdot \mathbf{x}_t, \quad \text{where } n = \Delta t/\delta t,$$

δt being the time between observations, and $\mathbf{F}$ the constant-velocity transition matrix defined previously. Since the Kalman filter maintains an internal uncertainty estimate (error covariance), each prediction can also be accompanied by a confidence region.

To visualise the predicted sky conditions, each forecasted centroid is converted back into fisheye image coordinates via the inverse of the defisheye projection described earlier. This allows the construction of a future binary cloud map consistent with the original image format and spatial layout. Predicted masks are rendered for forecast horizons ranging from 1 to 30 minutes in 1-minute increments.

During this reprojection, predicted patches are drawn into the forecast map using the last known size and shape of the corresponding cloud patch, optionally dilated to account for positional uncertainty. These future masks enable downstream applications - such as robotic scheduling systems - to query not only the current but also the expected cloud cover at specific coordinates and times.

While the system does not attempt to model cloud growth, dissipation, or complex deformation, the linear motion assumption has proven sufficient over short time intervals under stable conditions.

To account for varying prediction quality, a basic confidence filter can be applied to discard low-certainty forecasts. This includes rejecting predictions based on short or inconsistent tracks, as well as extrapolations beyond a configurable horizon (e.g. 20 minutes).

3.4.7 Prediction Confidence Filtering and Uncertainty Handling

While the presented prediction algorithm achieves reliable forecasts up to 15–20 minutes into the future, not all predicted cloud maps are equally trustworthy. To account for varying prediction quality, a basic confidence filtering mechanism is applied. This filter discards (or flags) predictions originating from insufficiently long or inconsistent cloud tracks, particularly those that span fewer than 4–5 frames or exhibit angular deviations above the tracking smoothness threshold (typically 10°). Additionally, tracks that terminate near the image edge are deprioritized, as their extrapolated motion may rely on distorted or incomplete data due to the fisheye projection's increased distortion near the horizon.

A configurable prediction horizon (e.g. 20 minutes) is enforced to prevent extrapolations beyond the validated temporal range of the Kalman filter's tracking model. Beyond this horizon, the prediction accuracy deteriorates significantly due to compounding errors and the inability to sense incoming clouds from off-frame regions.

For future versions, a probabilistic extension is envisioned: Instead of treating predicted cloud positions deterministically, the full error covariance matrix $\mathbf{P}_k$ of the Kalman state vector can be propagated forward in time. This would enable the construction of spatial uncertainty envelopes around each predicted cloud patch. By thresholding this covariance-derived uncertainty, one could generate a “confidence mask” for each pixel in the predicted cloud map.

Such an approach would allow the scheduler to distinguish between high-certainty and low-certainty predictions, facilitating probabilistic cloud cover assessments under marginal conditions. This aligns with recent developments in uncertainty-aware forecasting in both meteorology and computer vision domains.

3.5 Evaluation and Performance

The effectiveness of the proposed system is assessed by evaluating both its cloud detection capabilities and its short-term forecasting accuracy using real-world data from the Skycam A and Skycam T instruments at the Liverpool Telescope, located at the Observatorio del Roque de los Muchachos (ORM) on La Palma. The evaluation focuses on the reliability of cloud identification in individual frames and the quality of predicted cloud motion over forecast horizons of up to 30 minutes.

3.5.1 Evaluation Dataset and Operational Context

The performance of the cloud detection system was assessed using historical datasets acquired by Skycam A at the Liverpool Telescope. Two representative samples were selected for evaluation. The first covers the period from February 20th, 2019 to December 31st, 2022 and comprises a total of 619,421 all-sky images. The second spans from January 1st, 2024 to January 30th, 2024, and includes 16,937 images. Together, these samples reflect a broad range of atmospheric conditions, including photometric nights, nights with intermittent cloud cover, and nights dominated by full overcast.

The second sample is particularly noteworthy: during this period, the Liverpool Telescope operated without an external satellite-based cloud sensor due to a change in data provider. As a result, observations were conducted under conditions that would typically have triggered an automatic shutdown. This operational gap provided a valuable test bed for the autonomous all-sky cloud detection system, allowing for a real-world assessment of its robustness and accuracy under uncontrolled and often marginal conditions.

The Liverpool Telescope normally relies on an automated satellite-based cloud monitoring service as part of its weather safety system. This service periodically retrieves near-real-time cloud mask products derived from satellite imagery covering the Canary Islands region. The system queries the latest satellite data via automated scripts and evaluates the cloud probability values for the pixels corresponding to the observatory location and its immediate surroundings. If the retrieved cloud map indicates significant cloud coverage above predefined thresholds, the telescope enclosure is automatically closed to protect the instrument and prevent observations under unsuitable conditions.

The evaluation dataset was therefore intentionally chosen to stress-test the system and validate its capability to detect clouds under both nominal and adverse observing scenarios. This approach ensures that the reported performance metrics are representative of practical use cases in robotic telescope operations.

3.5.2 Cloud Detection Accuracy

The cloud detection algorithm was tested against the algorithm presented in Adam et al., 2017 (see section 2.4.2) who generously provided their code on Github[2], as well as a method using the Skycam T at the Liverpool Telescope.

The comparison of Adam et al., 2017 against the algorithm presented in this thesis, both using the Skycam A, could only be performed in dark nights when there was no visible moon. During times with visible moon, especially with a moon disk of 30% or more, their algorithm was unable to perform photometry in large areas of the sky. With the full moon up, the method did not provide any useful results at all. This is caused by hardware limitations of the all-sky camera equipment at the Liverpool Telescope (Starlight Xpress Oculus). However, the problem of over-exposed areas around the visible moon is a common problem on most typical all-sky cameras

In dark nights both algorithms matched against each other by around 92% of detected cloud area, depending on the density of catalogue stars in the area of the sky. A limitation of the algorithm presented in Adam et al., 2017 is the fact that it can only outline clouds based on visible and non-visible stars in a provided catalogue whilst the method presented in this paper is able to accurately detect the shape of the cloud which is the main cause for the average match of ‘only’ 92%.

In bright nights (moon disk $> 30\%$), whilst their algorithm failed to provide a cloud map (using the images recorded with the equipment at the Liverpool Telescope), the method presented in this thesis was able to accurately detect clouds around the actual

---

[2] `https://github.com/tudo-astroparticlephysics/starry_night`

moon disk. In Figure 3.6, the right side represents two all-sky images taken one minute apart. The bottom left side is the difference image and the top left picture contains the extracted cloud pattern. The cloud pattern contains various small white patches falsely identified as clouds. These are the result of moonlight scattering on dust accumulated on the camera protection over time. With an exposure time of 30s, the moon disk itself saturates the CCD completely and the ADU values in the pixel regions around the moon disk are also close to saturation levels resulting in visible noise. This noise is falsely labelled as a cloud but it has little to no influence on the observations itself as observations this close to the moon are very rare.

The next test was conducted with the help of our Skycam T. Skycam T takes a 10 second exposure every minute, at the same time as Skycam A takes its image) and stores it in a publicly available database on the website of the Liverpool Telescope. All Skycam T images are plate-solved using the Astromentry.net software suite (Lang et al., 2010) and a world coordinate system (WCS) is added to the FITS file.

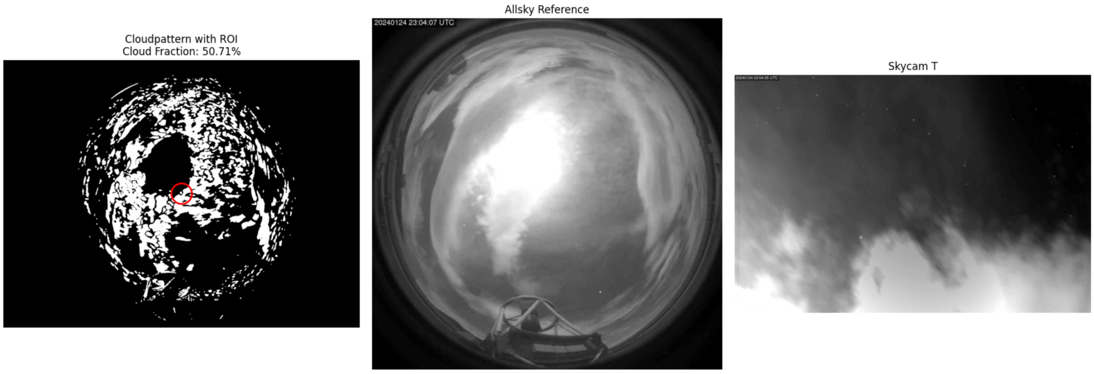

Figure 3.10: Example of a cloud detection result (typical case). **Left:** Binary cloud mask (white = cloud, black = clear sky) derived from the all-sky image using morphological and intensity-based filtering. The red circle marks the region of interest (ROI) centered on a scheduled observation target, with a radius corresponding to 20 pixels in the deprojected image. The detected cloud fraction within the ROI is 50.71%. **Center:** Original all-sky reference image taken by the main all-sky camera at 2024-01-24 23:04:07 UTC, showing significant cloud cover and partial moon illumination. **Right:** Simultaneous Skycam T image showing the corresponding sky segment near the zenith, where patchy clouds obscure several stars. The visual agreement with the all-sky detection validates the segmentation result in this case. The ROI in the left image is the area that corresponds to the Skycam T image (which is basically a zoomed in version of the ROI)

To validate the cloud detection system, a comparison was performed using images from Skycam T, a medium-field monitoring camera mounted on the Liverpool Telescope. Skycam T acquires a 10 s exposure every minute, synchronised with the acquisition of images from the all-sky camera (Skycam A), and stores the data in a publicly accessible archive on the Liverpool Telescope website. All Skycam T images are plate-solved using the

Astrometry.net software suite (Lang et al., 2010), which provides a World Coordinate System (WCS) solution for each image.

The validation procedure compares two independent indicators of cloud presence: (i) the cloud classification produced by the all-sky detection system (Skycam A), and (ii) the reduction in the number or brightness of stars detected in the corresponding Skycam T image.

First, a set of nights containing visible cloud structures was randomly selected from the Skycam T archive. For each Skycam T image, the sky coordinates of the image centre were obtained from the WCS solution. This position was then projected into the corresponding all-sky cloud map derived from Skycam A. A circular region with a radius of 20 pixels centred on this position was extracted from the cloud map. Within this region, a "cloudiness" parameter was computed, defined as the fraction of pixels classified as cloud by the detection algorithm (see Figure 3.10).

Independently, the Skycam T image was analysed photometrically using SExtractor (Bertin et al., 1996). The APASS catalogue (Henden et al., 2018a) was queried for the same sky region, retrieving all stars brighter than 10 mag. The catalogue query provides the expected number of detectable stars in the Skycam T field of view. This value was then compared to the number of stars actually detected in the image by SExtractor.

An image was labelled as "cloud covered" if the number of detected stars was below 75% of the expected catalogue count. The relatively high threshold accounts for the optical and photometric limitations of Skycam T. Because of its wide field of view and the use of aperture photometry, SExtractor may fail to detect some stars in crowded regions without employing more advanced techniques such as PSF photometry.

To account for thin clouds that dim stars without fully obscuring them, an additional check was performed for cases where more than 75% of the expected stars were detected. In these situations, the measured instrumental magnitudes were compared with catalogue magnitudes derived from calibration on a clear night. If the mean magnitude offset exceeded 3 mag, the image was also classified as cloud affected.

Finally, the classification obtained from the Skycam T analysis was compared with the cloudiness parameter derived from the all-sky cloud map. A region in the all-sky map was considered cloud-covered if more than 75% of the pixels within the corresponding region were labelled as cloud.

Cloud detection performance is evaluated based on a set of verified cases sampled across a variety of conditions, including moonless nights, moderate moon illumination, and challenging bright moonlight scenarios. For the purposes of analysis, the test cases are grouped into three categories:

- **Dark background:** Moon below horizon or new moon phase.
- **Bright, Moon far:** Moon above horizon, phase > 10%, and angular separation from the region of interest (ROI) greater than 25°.
- **Bright, Moon close:** Moon above horizon, phase > 10%, and angular separation from ROI less than 25°.

For the cloud detection evaluation, 47 nights from 2019 to 2023 were selected to represent a wide range of observing conditions, including bright nights with and without the Moon, as well as dark nights. These nights were identified via entries in the night log[3] and visual inspection of the nightly all-sky video sequences[4]. From a total of 32,307 all-sky images across these nights, the first and last 60 minutes of each night were excluded to remove twilight effects. From the remaining images, one 100-image set and two 1000-image sets were randomly selected:

- **Dark background**: 1000 images from dark nights (Moon altitude $< -5°$ and Moon disk $< 10\%$).
- **Bright, Moon far**: 1000 images from bright nights with Moon–target angular separation greater than 25°.
- **Bright, Moon close**: 100 images from bright nights (Moon altitude $> -5°$, Moon disk $> 10\%$) where the angular separation between the target and the Moon was less than 25°. This configuration is relatively rare.

Testing the cloud detection system using the Skycam T method, the following results were obtained:

For bright nights, a false negative rate (the sky area in the cloud map is labelled "clear" whilst is is actually cloud covered) of less than 1% of the analysed images and a false positive rate also of 1% outside a radius of 25° of the moon disk was found. Inside this area, the false positive rate can climb up to $10 - 15\%$ (especially with a visible moon disk of more than 50%) due to moonlight scattering on dirt and dust on the sensor cover as well as lens-flare effects. As astronomical observations are rarely performed so close to the moon, this relatively high false positive rate is thus of no concern. The rate could

	Bright > 25°		**Bright < 25°**		**Dark Nights**	
Actual / Predicted	C	O	C	O	C	O
C	0.46	0.01	0.41	0.13	0.41	0.01
O	0.01	0.52	0.00	0.46	0.14	0.44
Accuracy	0.98		0.87		0.85	
Precision	0.97		0.76		0.97	

Table 3.1: Confusion matrices for cloud detection under three illumination regimes: bright nights more than 25° from the Moon, bright nights within 25°, and dark nights. Rows correspond to the true class (C = cloud, O = open sky), columns to the predicted class. Values represent normalised fractions of classified sky regions.

be improved by frequently cleaning the camera's protective elements and using a sensor with a higher dynamic range.

A confusion matrix (see Table 3.1) was compiled to quantify the rates of true and false cloud detection under these conditions. High accuracy was observed under dark and moderately bright skies. Performance declines modestly in close proximity to the Moon due to residual artefacts and overexposure, although such conditions remain rare and have been addressed in earlier sections through filtering and predictive handling.

3.5.3 Tracking and Forecasting Accuracy

Short-term forecast accuracy is evaluated using 100 sequences from Skycam A, covering a range of observing conditions. For each sequence, the predicted cloud mask at various time steps is compared to the actual observed cloud mask using a spatial overlap metric in the defined ROI. Six representative examples are shown in Figure 3.11, with agreement curves and residuals plotted against time.

---

[3] `https://telescope.ljmu.ac.uk/Reports/`

[4] `https://telescope.livjm.ac.uk/data/webfiles/Skycam/browse2025.html`

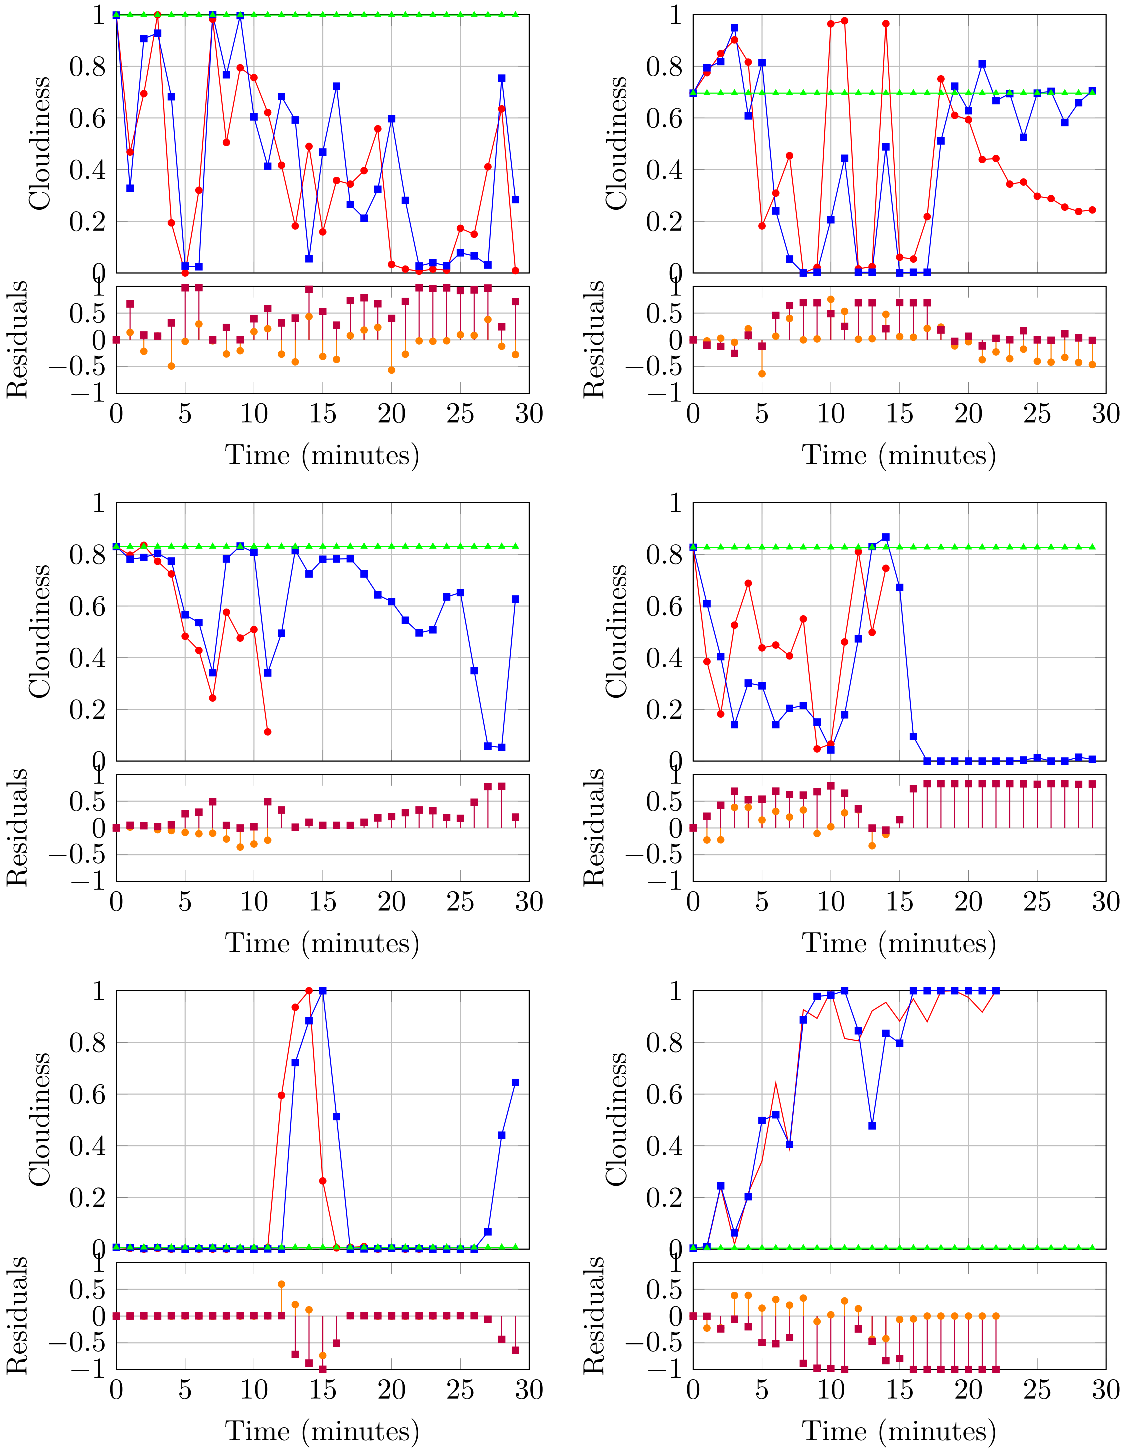
Cloudiness
0
0.2
0.4
0.6
0.8
1
Residuals
−1
−0.5
0
0.5
1
0
5
10
15
20
25
30
Time (minutes)

FIGURE 3.11

Figure 3.11: The plots show the accuracy for 6 different predictions: Plot 1 (top left) starting at 03:27 on 2020, April 9th, Plot 2 (top right) starting at 02:22 on 2020, April 8th, Plot 3 (middle left) starting at 20:15 on 2023, November 25th, Plot 4 (middle right) starting at 22:50 on 2021, April 26th, Plot 5 (bottom left) starting at 23:51 on 2020, December 31st, Plot 6 (bottom right) starting at 06:16 on 2020, October 18th, . The graphs in the upper panels show the predicted cloudiness (in percent) over the next x minutes. Here, *cloudiness* refers to the fraction of cloud-classified pixels within the observable sky region, defined as all-sky pixels from zenith down to an altitude of 20°, excluding the horizon region where distortions dominate. The cloudiness is a measure for the amount of cloud-pixels in the area of a scheduled observation (a circular region with a radius of 20 pixels around the target) over the total amount of pixels in this area. The green graph (baseline) represents a persistence forecast, in which the cloud map from the most recent all-sky image is assumed to remain unchanged over the prediction horizon. Consequently, the baseline curve is constant in time for each plot, reflecting the cloudiness measured in the last observed frame without incorporating any motion model. The red graph (prediction) shows the cloudiness obtained using the prediction method presented in this thesis. The blue graph (ground truth) shows the actual cloudiness measured at that point in time. This is put in as a reference, as this data would not be available in a real-time scenario. The residuals show the difference between the baseline and the prediction (orange) as well as the baseline and the ground truth (red), again, put in as a reference. The predictions in plots 3, 4, and 5 do not have the full prediction dataset, because the observations were made in the area of the sky where the wind was coming from, thus there is not enough cloud data for a longer prediction. Plot 5 shows passing clouds. Plot 6 ends after 22 minutes as the weather control system shut the telescope enclosure because of incoming severe conditions.

To generalise the tracking quality, the prediction agreement is averaged across all sequences and plotted as a function of forecast time (Figure 3.12). Agreement is defined as the fraction of matching cloud-labeled pixels in the ROI. Results show that agreement remains above 80% for up to 10–15 minutes and slowly declines thereafter, justifying the system's design horizon of 15–20 minutes.

Performance was found to be best for stable stratiform cloud patterns and degraded for fragmented or fast-evolving clouds, particularly near the horizon or under partial occlusion. These limitations are partially mitigated through prediction confidence filtering (see Section 3.4.7).

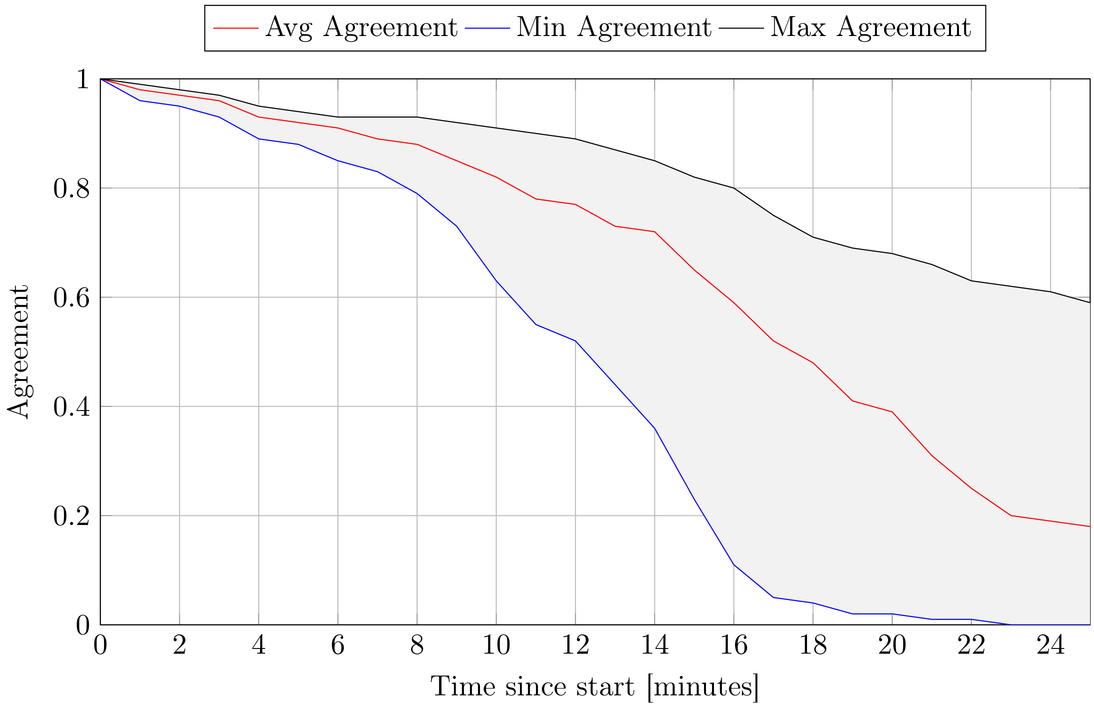

Figure 3.12: Agreement between predicted and actual cloud maps over time, averaged across 100 sequences. The red line shows the average agreement, while the shaded area spans the minimum and maximum values observed at each time step. Agreement is defined as the fraction of correctly matched cloud pixels in the region of interest.

3.5.4 Computational Efficiency and Real-Time Suitability

In addition to detection and prediction accuracy, the system's suitability for real-time operation was assessed. Although many algorithmic stages are implemented using GPU shaders to exploit parallelism, the overall processing pipeline remains lightweight and compatible with CPU-only deployments.

The core image processing chain, consisting of FITS normalization, difference imaging, morphological operations, and cloud masking, was designed to operate on all-sky images at a cadence of one frame per minute and a resolution of 1392 × 1040 pixels. Binary morphological operations, including erosion and dilation, are executed using OpenGL-based pixel shaders with a kernel size of 3 × 3 pixels. These operations were found to be effective in removing small-scale noise while preserving the structural integrity of cloud patterns. Larger kernels can be applied for more aggressive filtering, but this may attenuate finer cloud details and reduce detection granularity.

The motion prediction module, based on per-track Kalman filters, generates forecasts by shifting cloud pixels according to a smoothed average motion vector derived from recent cloud tracks. This forward-propagation is repeated in 1-minute increments for the full prediction horizon (typically up to 20 minutes). All shifting operations are carried out using GPU fragment shaders to ensure minimal computational overhead. Despite this use of GPU acceleration, benchmarks indicate that all processing stages can be executed comfortably in real time on modern multi-core CPUs, without requiring specialized hardware.

This efficient runtime behaviour ensures that the system can be deployed on modest observatory infrastructure, including systems without discrete GPUs. On the current implementation, processing a single all-sky frame (including image differencing, morphological filtering, segmentation, and object tracking) requires approximately $\sim$1 s on a standard GPU-based system. Given the one-minute acquisition cadence of the all-sky camera, this corresponds to less than $\sim$5% of the available processing interval, leaving ample margin for data transfer, logging, and prediction updates while maintaining real-time operation.

The computational cost of the algorithm scales approximately linearly with the number of pixels in the input image, as most operations consist of local filtering and segmentation steps. Consequently, doubling the image resolution approximately doubles the processing time. Similarly, increasing the frame cadence would primarily increase the total processing load proportionally to the number of frames processed per unit time, while the per-frame runtime remains unchanged. This behaviour allows the system to accommodate higher-resolution sensors or moderately increased frame rates without requiring specialised hardware.

3.5.5 Prediction Accuracy and Time Horizon

To assess the temporal reliability of the tracking and prediction pipeline, a large-scale benchmark was conducted using 100 cloud track sequences from Skycam A. For each sequence, the predicted cloud maps were compared against ground-truth maps at 1-minute intervals, up to a forecast horizon of 25 minutes.

Prediction accuracy was quantified using pixel-wise agreement in the binary cloud masks within the region of interest (altitude $> 20°$). As shown in Figure 3.12, the average agreement remains above 80% for the first 5 minutes and gradually declines to around 60% by minute 20, with most of the degradation attributable to untracked or newly formed clouds entering the frame.

For robotic scheduling purposes, the relevant quantity is not pixel-perfect structural agreement, but whether the predicted cloud coverage within a target's field of view crosses a decision threshold (e.g., clear pixels versus obstructed pixels). An average pixel-wise agreement above 80% over the first 5 minutes implies that the vast majority of target-area classifications remain unchanged relative to ground truth over typical exposure cadences (2 to 5 minutes). Even at 60% agreement after 20 minutes, most discrepancies arise from cloud boundaries entering or leaving the frame, rather than from large-scale misclassification of clear sky regions.

Typical operations consist of 1 to 5 exposures per group, with individual exposure times usually below 5 minutes, resulting in an average group execution time of approximately 10 to 15 minutes. Only a small fraction of programmes require uninterrupted executions approaching longer times. The prediction horizon therefore comfortably covers the decision-relevant timescale for the majority of scheduling scenarios, with agreement remaining above 70 to 80% during the interval most critical for execution.

3.5.6 Known Limitations

Despite its overall robustness, the cloud monitoring system has several limitations:

- **Moonlight contamination**: During full Moon phases, lens flare and sensor blooming can lead to spurious cloud detections. These are mitigated via region masking and component filtering but are not always fully suppressible.
- **Telescope contamination**: Bright reflections from the telescope structure (especially during slews) can be misclassified as clouds. These can not be masked with the current system as the camera system does not have real-time access to the telescope RCS and thus is unable to determine the live pointing coordinates.
- **Incoming clouds at horizon**: Due to optical distortion and undersampling near the horizon, clouds entering the field of view from below 20° altitude are hard to detect early. This imposes a natural limit on look-ahead capability.
- **Thin slow-moving clouds**: In very thin cirrus conditions, especially during dark time, low-contrast cloud features may not register significantly in the difference image and thus escape detection. This problem can be partially solved with better camera hardware.

These issues are inherent to the limitations of ground-based, wide-angle imaging systems without active calibration. However, their impact is mitigated by filtering, regional masking, and short forecast horizons aligned with operational decision timescales.

3.6 Integration with Scheduler and Use Cases

The cloud detection and forecasting system currently operates as a stand-alone monitoring tool that produces cloud maps, quantitative “cloudiness” metrics for user-defined regions, and plots of predicted future sky conditions. At present, these outputs are generated in real time and can be accessed via a lightweight web interface. While the Liverpool Telescope scheduler does not yet consume this information directly, the system provides valuable situational awareness for operators and forms the basis for future integration.

In the short term, the cloud monitor is intended to complement the existing satellite-based decision system that governs the operational state of the telescope enclosure. The satellite feed primarily serves a safety function, providing large-scale weather awareness relevant to precipitation and severe atmospheric conditions. However, it does not resolve local, transient cloud structures at the spatial and temporal scales relevant for observation planning. The all-sky cloud monitor therefore does not replace safety-critical weather systems, but augments them by providing high-resolution, site-specific awareness of passing cloud structures. This additional layer of information can support operational decisions such as delaying enclosure opening during transient cloud cover or suspending observations when localised cloud bands approach, without altering the primary safety safeguards.

In the longer term, the predictive cloud maps generated here are designed to be compatible with adaptive schedulers, such as the system under development for the New Robotic Telescope (NRT). In that context, short-horizon forecasts (5 to 10 minutes) can be queried programmatically as part of the scheduling loop, allowing the system to anticipate imminent cloud motion and adjust observation selection accordingly. Given the demonstrated forecast accuracy, such integration would not eliminate uncertainty, but could reduce avoidable failed observations and improve short-term decision robustness compared to purely reactive strategies. Rather than replacing reactive scheduling, the predictive layer is intended to provide incremental gains in efficiency under variable sky conditions.

In summary, while the present implementation primarily delivers real-time monitoring outputs, its architecture is designed to support integration into operational awareness systems and, in the longer term, into adaptive robotic schedulers operating on short prediction horizons.

Figure 3.13: Skycam O installed on the weather mast at the Liverpool Telescope. The camera is housed in a weatherproof enclosure and connected to the observatory network for real-time data transmission.

3.7 Skycam O

To complement the upward-facing SkyCam instruments (A, T, and Z), a new camera system was developed and deployed by the author to provide real-time visual context of the surrounding sky and local environment. This system, named Skycam O (for "Outside", see Figure 3.13), is designed for robustness, energy independence, and continuous operation in the harsh outdoor conditions of the Observatorio del Roque de los Muchachos.

3.7.1 Hardware Design and Installation

At its core, Skycam O is built around a Raspberry Pi 4 Model B, chosen for its low power consumption, strong community support, and GPIO flexibility. To further increase the system's resilience against power disruptions and voltage fluctuations, a StromPi 3 HAT

(**H**ardware **A**ttached on **T**op expansion board) is integrated. A HAT is an add-on circuit board designed to plug directly onto the Raspberry Pi's GPIO header, extending its functionality. In this case, the StromPi 3 enables seamless switching between multiple power sources and includes a real-time clock for accurate timekeeping. It also supports an optional $LiFePO_4$ battery backup, allowing the system to survive short-term outages without data loss or corruption.

The imaging component is a Raspberry Pi High-Quality Camera Module equipped with a fisheye lens providing an approximately 180° diagonal field of view, corresponding to near-hemispherical sky coverage. The lens specification refers to the nominal optical field of view of the fisheye system; the effective usable sky region is slightly smaller due to mechanical constraints of the housing and dome. The camera is mounted within a custom-built enclosure consisting of a standard white UV-resistant DN100 pipe (a 100 mm drainage pipe), repurposed as a weatherproof cylindrical casing. Internally, a 3D-printed mounting structure (see Figure 3.14) securely supports the Raspberry Pi and camera, with a circular base that fits precisely into the pipe to prevent movement. The upper end of the housing is sealed with a clear acrylic dome, providing an unobstructed field of view while protecting the optics from environmental exposure.

Skycam O is physically installed on the outdoor weather mast near the Liverpool Telescope. It is mounted on the original base plate of a decommissioned pyranometer, ensuring structural stability and optimal exposure to the sky. The system is connected via a dedicated LAN cable and powered through a regulated 24V power supply routed through the mast infrastructure.

3.7.2 Software and Operation

Skycam O runs Raspberry Pi OS, a Debian Linux-based distribution optimized for headless operation. The camera control is managed via the open-source `allsky` software suite[5], which supports timelapse imaging, automatic exposure control, and dynamic overlays.

To integrate the data into the observatory workflow, a set of custom Python and shell scripts has been developed. These scripts:

- Convert JPEG frames into `FITS` format for archival and scientific compatibility.
- Copy both JPEG and FITS files to the Liverpool Telescope backend servers via scheduled secure transfer.

---

[5] `https://github.com/AllskyTeam/allsky`

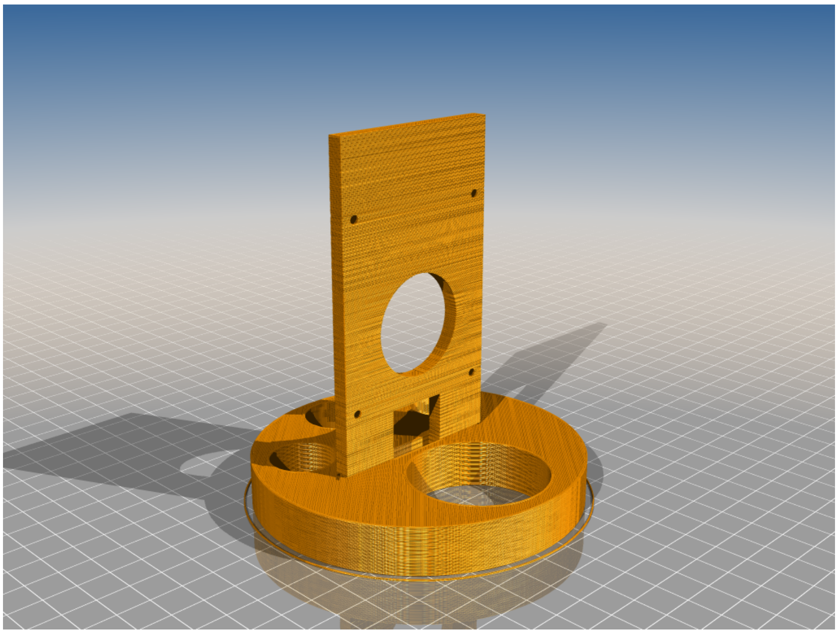

FIGURE 3.14: 3D-printed mounting structure for the Raspberry Pi and camera inside the Skycam O enclosure. The design ensures secure placement and optimal alignment of the camera lens with the acrylic dome.

- Perform regular housekeeping, such as rotating log files, cleaning up temporary data, and monitoring system health metrics.

The entire system is designed to operate with minimal human intervention. Reboots, service restarts, and updates can be issued remotely via SSH, and uptime is continuously monitored.

3.7.3 Role in the Observing Workflow

Skycam O provides an external view of the sky and horizon (see Figure 3.15), capturing clouds, haze, Calima, fog banks, and approaching weather fronts. At present, these images primarily support situational awareness for operators under uncertain weather conditions. However, the system architecture allows for automated extraction of cloud and visibility metrics, enabling future integration into decision-support or enclosure-control logic without requiring continuous human supervision.

Though not used for photometry, the system adds valuable qualitative data, and unlike many commercial all-sky systems, it is fully open, maintainable, and reproducible.

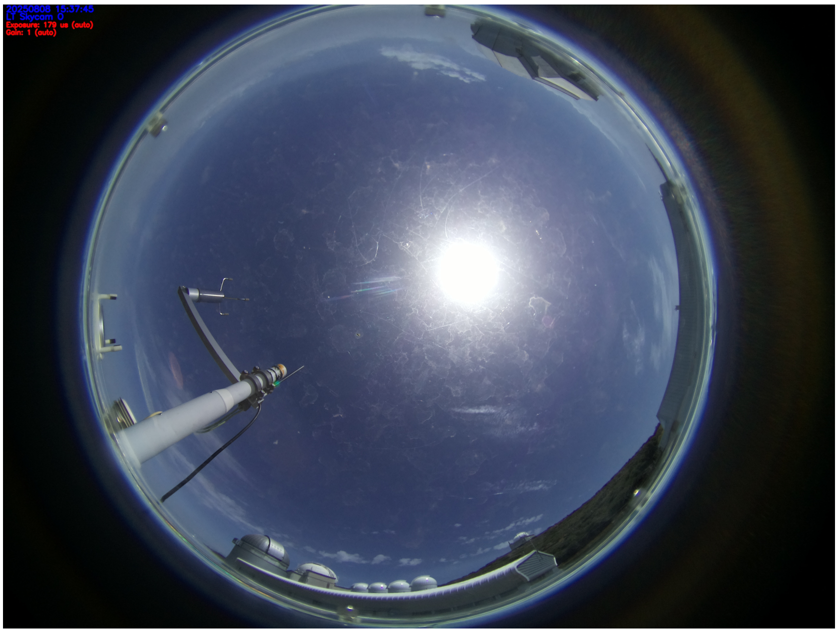

Figure 3.15: Example all-sky image captured by Skycam O, showing the sky and horizon at the Liverpool Telescope site. The image captures clouds, haze, and atmospheric conditions in real time.

3.8 Discussion

The results presented in this chapter demonstrate that real-time cloud monitoring with all-sky cameras can provide reliable, actionable input for robotic telescope scheduling. By combining difference imaging, adaptive thresholding, and morphological filtering, the system successfully reduces stellar contamination and produces robust binary cloud masks across a wide range of observing conditions. The GPU-based implementation ensures that these steps can be executed at the native one-minute cadence of Skycam A, even on modest hardware, which is essential for operational integration.

The tracking and prediction framework, based on deprojection, connected-component labeling, and Kalman filtering, extends this capability beyond static detection. Forecast horizons of 10 to 20 minutes show agreement rates exceeding 70-80% under stable conditions, validating the constant-velocity assumption for short-term predictions. These results are particularly relevant given the timescales of typical Liverpool Telescope observations, where setup and execution of groups often span similar intervals. The evaluation also highlights the system's ability to outperform simple persistence baselines, offering genuine predictive value rather than merely describing current sky conditions.

At the same time, several limitations remain. Prediction quality degrades for fragmented or fast-evolving clouds, especially near the horizon where geometric distortion and undersampling reduce sensitivity. Moonlight contamination and telescope reflections occasionally produce spurious detections, though these can be partially mitigated through exclusion zones and heuristic filters. The system is also inherently constrained by its reliance on a single upward-facing fisheye camera: incoming clouds remain undetectable until they enter the field of view, limiting the maximum achievable forecast horizon.

Despite these limitations, the system demonstrates that local, high-cadence imaging provides a practical and effective approach to real-time sky condition awareness. The integration of cloud masks and short-term forecasts into robotic scheduling workflows has the potential to reduce wasted exposures, prioritize targets in temporarily clear regions, and increase overall efficiency. In this sense, the cloud detection pipeline can be regarded as a foundational capability for the next generation of adaptive robotic observatories, complementing extinction monitoring and sky brightness modeling to form a complete situational awareness framework.

3.9 Future Cloud Sensor Design

While the current cloud monitoring system based on the Skycam A infrastructure has demonstrated robust performance, it is inherently limited by legacy hardware, fixed exposure configurations, and the lack of onboard computational capabilities. To enhance long-term maintainability and allow for future algorithmic expansion, a new cloud sensing system is proposed.

The core of the redesigned system is a high-quality CMOS or CCD all-sky imaging camera equipped with a fisheye lens, offering high sensitivity and dynamic range under a wide range of illumination conditions (from new Moon to full Moon). Modern sensor modules, such as the ZWO ASI series (e.g. ASI715MM) or industrial-grade USB3.0 cameras with global shutter capabilities, provide significantly improved signal-to-noise ratio and frame rate performance compared to the current hardware.

Image acquisition, real-time processing, and data transmission are managed by a low-power embedded compute platform, such as a Raspberry Pi 5 or an industrial fanless mini-PC. This unit performs local preprocessing, including bias and flatfield correction, difference image generation, and binary cloud mask extraction. The existing detection and tracking pipeline described in this chapter can be directly ported to this embedded platform, utilising GPU acceleration (e.g. via OpenGL ES or Vulkan) for real-time defishing and cloud displacement estimation.

Each processed frame is timestamped and annotated with metadata including thresholding parameters, cloud fraction, and system diagnostics. The resulting cloud mask and associated metadata are transmitted over a local network or to a central backend server. This server hosts a relational database and object storage service for long-term archiving, as well as a web-facing API to enable programmatic access by other observatory systems and external users.

The modular architecture ensures that additional sensors—such as a pyrometer, sky temperature sensor, or lightning detector; can be integrated via the same edge compute unit. Moreover, redundant deployment at multiple telescope sites would enable cross-validation and facilitate distributed sky condition monitoring across a larger observing region.

This upgraded platform would not only improve image quality and processing flexibility but also reduce maintenance overhead and eliminate the reliance on monolithic legacy systems. It opens the door for future features such as uncertainty quantification, machine-learning-based cloud classification, and direct coupling with robotic telescope schedulers via forecast-aware observation planning APIs.

Chapter 4

Real-time Atmospheric Extinction Measurements

For my part I know nothing with any certainty, but the sight of the stars makes me dream.

Vincent van Gogh

4.1 Introduction

While the previous chapter focused on detecting, tracking, and predicting cloud cover - enabling the identification of unsuitable observing conditions in parts of the sky—this chapter addresses the continuous attenuation of starlight under clear skies: atmospheric extinction. Together, cloud detection and extinction monitoring provide a comprehensive assessment of sky transparency, forming the foundation for adaptive scheduling in robotic telescope operations.

Atmospheric extinction is caused by the absorption and scattering of incoming starlight as it passes through the Earth's atmosphere. The amount of attenuation increases with airmass—the optical path length through the atmosphere—and varies with wavelength, atmospheric composition, and local conditions. Quantifying this effect is essential for photometric calibration and for ensuring the scientific reliability of ground-based observations.

At the Observatorio del Roque de los Muchachos (ORM), precision atmospheric monitoring has long played a crucial role in supporting astronomical observations. The

Carlsberg Meridian Telescope (CMT), operated from 1984 to 2013, was instrumental in characterising long-term atmospheric extinction in the Johnson V band. Its nightly extinction measurements, derived from standard stars crossing the meridian, form the most comprehensive dataset of its kind for the site (Garcia-Gil et al., 2010). Although the CMT has been decommissioned, its legacy endures. Notably, the New Robotic Telescope (NRT) is set to be built at the very site where the CMT once stood—symbolically linking the future of robotic astronomy to the past of precision atmospheric monitoring.

For the LT, photometric science frames are common, whereas the NRT will prioritise spectroscopy and rapid response. In such regimes, acquisition images are often shallow, taken through narrow slits or non-photometric configurations, or omitted entirely for fast triggers. As a result, attempting to derive transparency or zero-point (ZP) information from user science frames becomes sporadic and instrument-dependent. Tying scheduler state (e.g. transparency, ZP) to these frames introduces latency—because the system must wait for suitable exposures—and brittleness—because the inferred calibration is sensitive to filter changes, grism/slit choices, exposure strategy, and target selection.

External but co-pointing wide-field imagers, specifically such like Skycam T and Skycam Z, provide a clean decoupling between science acquisition and transparency sensing. Their continuous, independent data stream yields a high-cadence record of atmospheric transmission throughout the night that is independent from the active instrument or observing mode. Because the Skycams co-point with the main telescope, their measurements follow the same line of sight and airmass as the current target, offering operationally relevant, near-real-time estimates of the extinction coefficient k without consuming additional telescope time or altering the observing sequence.

A stable extinction feed from the Skycams enables immediate SNR-based exposure adaptation and objective quality rating (photometric versus non-photometric) within the scheduling loop, complementing cloud detection and sky-brightness forecasting (Chapers 3, 5, 6). Additionally, this remains available even for short visits and rapid spectroscopic follow-up where no suitable photometric frames exist, thereby improving both responsiveness and uniformity of data quality under realistic robotic operations.

The objective of this part of the project is to evaluate whether wide-field auxiliary imaging systems (specifically systems like the existing Skycam T and Skycam Z) can be repurposed to deliver reliable, real-time atmospheric extinction coefficients. These instruments, originally designed for co-aligned sky imaging with the Liverpool Telescope (LT), continuously acquire unfiltered images throughout the night. Although their spectral response and optical characteristics currently differ from standard photometric systems as the systems operate unfiltered, they provide uninterrupted data streams that can, in principle, be mined for transparency information.

This chapter therefore continues, in spirit, the mission of the CMT by investigating whether Skycam data can be used to derive physically meaningful extinction measurements under realistic operating conditions. The key idea is to compare instrumental magnitudes of detected stars (derived via SExtractor photometry) to catalogue magnitudes from Tycho2 and APASS, and to model the dependence of the residuals on airmass. If a coherent extinction trend emerges, it would demonstrate the feasibility of integrating real-time extinction monitoring into the robotic scheduling framework.

To ensure robustness and explore the reliability of results under varying sky conditions, three complementary modelling approaches are employed: (1) a full linear least-squares regression over all matched stars, (2) a RANSAC-based robust regression to exclude outliers and reduce contamination from clouds or misidentifications, and (3) a median-based method that fits extinction using per-image median residuals binned by airmass. These techniques provide a multi-perspective assessment of each night's transparency and allow internal cross-validation of the extinction slopes derived.

While the unfiltered nature and wide field of the Skycams introduce complexity and limitations compared to traditional filtered photometry, the ability to obtain continuous, automated measurements with minimal additional overhead offers a compelling advantage. This chapter documents the methods, results, and limitations encountered in the extinction analysis and explores how such systems may contribute to the operational intelligence of future robotic observatories.

4.2 Instruments and Data Sources

The primary data sources used in this study are Skycam T and Skycam Z, two auxiliary imaging systems permanently mounted on the top-end of the Liverpool Telescope (LT). Both cameras operate continuously during observations and acquire images at approximately one-minute cadence. A summary of the main instrument characteristics is provided in Table 4.1.

Skycam T is a medium-field Starlight Xpress Trius SX-35 CCD camera fitted with a Zeiss Planar T 85mm f/1.4 ZF2 lens. This setup delivers a field of view of approximately 24×16 degrees and a pixel scale of 44 arcseconds per pixel. The system achieves a limiting magnitude around V $\approx$ 12mag and is particularly well-suited for identifying thin cloud cover due to its broad coverage and sensitivity to changes in transparency. Historical iterations of Skycam T used various optical and detector combinations, maintaining field sizes generally between 10 to 20 degrees[1].

	Skycam T	**Skycam Z**
Camera model	Starlight Xpress Trius SX-35	ZWO ASI 174MM
Lens	Zeiss Planar T 85mm f/1.4	Orion Optics AG8 telescope
Field of view	$24 \times 16^\circ$	$\sim 1^\circ$
Pixel scale	$\sim 44''$/pixel	$\sim 1.62''$/pixel
Exposure time	10 s	10s
Cadence	1 min	1 min
Magnitude range	$\sim$ 4–12 mag	$\sim$ 10–16 mag
Primary use	wide-field monitoring	near-field monitoring

Table 4.1: Summary of the main characteristics of the Skycam instruments used in this work.

Skycam Z, re-commissioned in September 2024, is a narrow-field instrument based on a ZWO ASI174MM uncooled monochrome CMOS camera coupled with an Orion Optics AG8 telescope. It provides a 0.84×0.55 degree field of view and a plate scale of 1.62 arcseconds per pixel. Skycam Z is capable of detecting stars down to $V \approx 16.5 - 17$mag, offering high spatial resolution and depth, with sufficient photometric precision to reveal galaxies and extended sources under good conditions[1].

Both systems acquire one ten-second exposure every minute during LT operations, producing an immense dataset. At the time of analysis, Skycam T had collected over 2.4 million images, while Skycam Z had accumulated more than 1.1 million exposures. Skycam Z was commissioned later than Skycam T and was not mounted on the telescope between 2019 and 2024 due to a camera failure and Covid-19 travel restrictions.

Raw data are stored in both FITS and JPEG formats, enabling both scientific processing and quick-look diagnostics. All image metadata and source extraction results are ingested into a custom PostgreSQL database developed for this project. Standard calibration frames such as bias and darks are routinely applied, and twilight flats are occasionally available for flat-field correction. The resulting dataset forms the foundation for real-time photometric monitoring and extinction analysis.

4.3 Data Acquisition and Processing

Before photometric analysis can be performed, raw Skycam images must be corrected for instrumental effects to ensure accurate and reliable measurements. The preprocessing pipeline applies a series of calibration steps designed to mitigate fixed-pattern noise, pixel sensitivity variations, and geometric distortions. This process prepares the data for subsequent astrometric calibration, source detection, and photometric extraction.

---

[1] https://telescope.livjm.ac.uk/TelInst/Inst/SkyCam/

The key preprocessing steps include dark and bias correction, flat-field correction, astrometric calibration, source detection and catalogue cross-matching.

As neither Skycam T nor Skycam Z is equipped with a mechanical shutter, traditional dark frame acquisition is not possible. Instead, an alternative dark subtraction method based on dome images was implemented by the LT Project Scientist Dr. Robert Smith. While this approach mitigates fixed-pattern noise, it introduces a systematic offset in sky background measurements due to over-subtraction. A detailed discussion of this procedure and its photometric impact is provided in Chapter 5.

Flat-field correction is performed using twilight sky flats obtained during the telescope's automated flat-fielding routine at dusk and dawn. Because the Skycams operate with a fixed exposure time, typically only a few images each night achieve the desired illumination having high enough counts without reaching saturation levels.

To assemble reliable flat-field frames, the evening and morning images are collected over several months. Then, frames are manually reviewed to select those with suitable brightness and minimal gradients and then, these selected frames are stacked to produce a master twilight flat (see Figures 4.1 and 4.2).

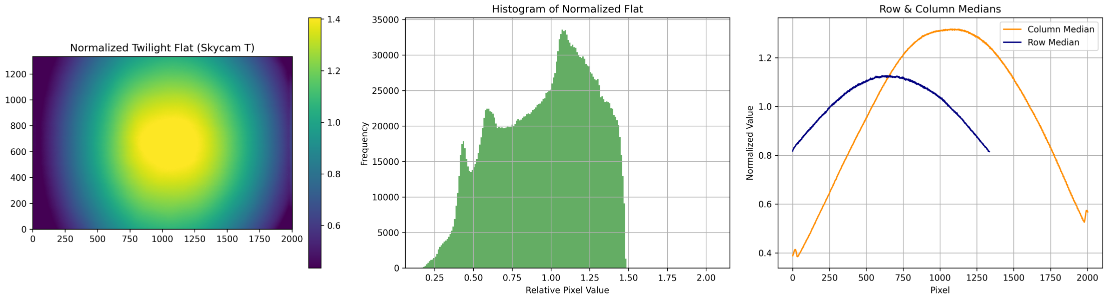

FIGURE 4.1: Analysis of the normalised master twilight flat for Skycam T. **Left:** The flat field exhibits smooth radial vignetting centred near the optical axis, without evidence of dust shadows or fringing. **Centre:** The histogram of normalised pixel values shows a broad, symmetric distribution consistent with gradual optical fall-off across the field. **Right:** Median row and column profiles confirm the radial symmetry and reveal a steady decrease in response toward the edges, with peripheral pixels dropping to approximately 40% of the peak—typical for fast, wide-angle optical systems.

As noted by the observatory's project scientist Robert J. Smith, this approach works reasonably well but is susceptible to contamination (e.g., moonlight or cloud features). The master flat is updated periodically when data quality monitoring suggests significant changes or degradation.

It is noted that the resulting dark frame also contains the bias signal (since no separate bias frames are acquired), and thus represents a combined dark+bias calibration frame. This combined frame is applied during preprocessing to correct each science image before subsequent extinction and sky brightness measurements.

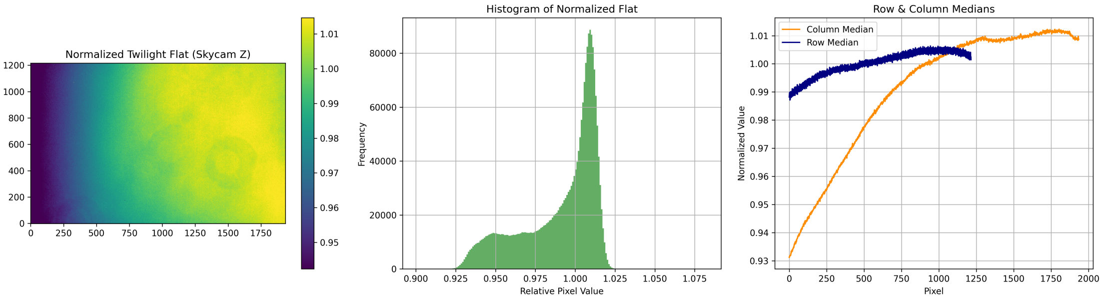

Figure 4.2: Analysis of the normalised twilight flat field for Skycam Z. **Left:** The flat image shows smooth large-scale vignetting, increasing in sensitivity from the lower left to upper right, with several faint circular features likely caused by internal reflections or 'dust bunnys'. **Centre:** The histogram of pixel values is peaked just above unity, with an asymmetric tail toward lower values and minor secondary features, consistent with spatially varying throughput. **Right:** Median row and column profiles confirm the smooth radial structure, with column medians rising from ≈ 0.93 to over 1.01 across the field.

In order to assess the quality and stability of the dark+bias calibration frame, a statistical and visual analysis was performed. The combined dark+bias image shows a uniform background level with pixel values centred around 1970 ADU. A histogram of pixel intensities (see Figure 4.3 Image 2) reveals a sharp peak, indicative of a stable and well-behaved bias level, with only minor outliers. A visual inspection of the image (Figure 4.3 image 1) shows minimal spatial structure, as expected, except for a single column (column 1399) that appears significantly elevated across all rows. This behaviour is consistent with a defective readout column and is evident both in the 2D image and in the row/column median profile plots.

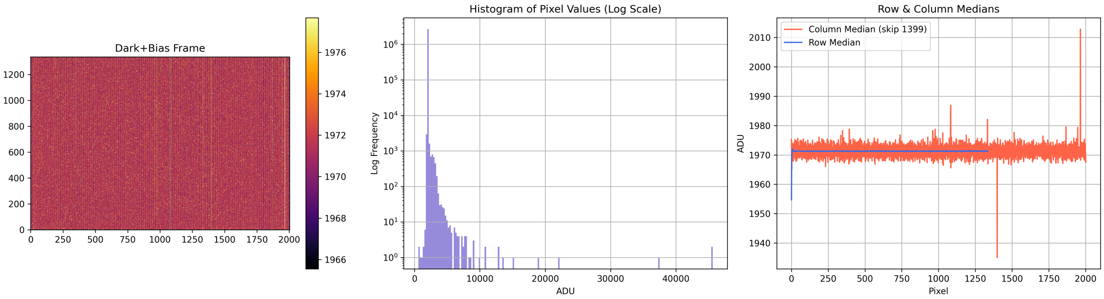

Figure 4.3: Diagnostic analysis of the Skycam T dark+bias calibration frame from the night of 2025-04-14. **Left:** The dark+bias image, shown with a `magma` colormap, displays a spatially uniform background with low-amplitude fixed-pattern noise and one clearly visible defective column (1399) reading significantly higher than its surroundings. **Centre:** The histogram of pixel values, plotted on a logarithmic scale, is sharply peaked around 1970 ADU, confirming a stable and narrow electronic bias distribution, with a faint high-ADU tail due to hot pixels or cosmic ray events. **Right:** Median column and row profiles show a flat and uniform bias structure across the sensor, except for minor readout-related deviations and edge effects. Column 1399 was explicitly excluded from the column median profile to remove its dominant influence.

To prevent this artefact from skewing downstream statistical analyses, column 1399 was

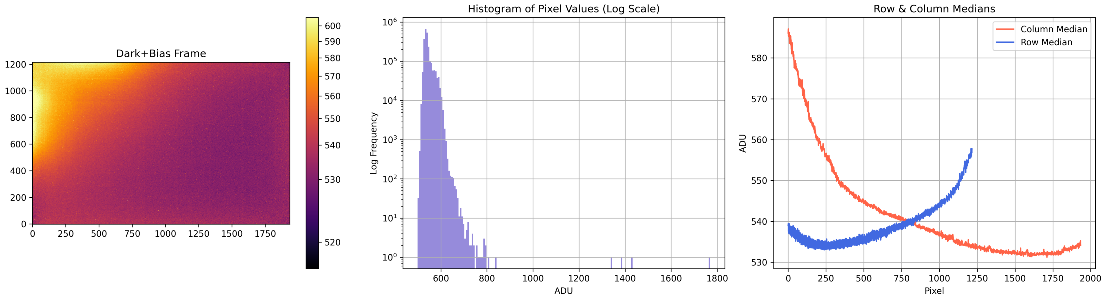

FIGURE 4.4: Diagnostic analysis of the Skycam Z dark+bias calibration frame. **Left:** The image reveals a clear spatial gradient, with elevated background levels in the upper-left corner exceeding 600 ADU. This feature is consistent with detector glow, a known phenomenon in uncooled CMOS sensors where internal electronics emit thermal photons. The median background across the frame is 536 ADU. Given this structure and intensity, the dark subtraction process may inadvertently remove genuine sky background flux, particularly in images with low sky signal, thereby contributing to the systematic underestimation of sky brightness in Skycam Z data. **Centre:** The histogram of pixel values, shown on a logarithmic scale, peaks sharply around 540 ADU, with a faint high-value tail attributable to hot pixels or sporadic noise events. **Right:** The dark+bias image reveals a smooth but asymmetric background gradient, with elevated values in the upper-left corner consistent with detector glow. The overall structure suggests that a significant component of the dark signal is spatially structured, likely due to thermal effects in the uncooled CMOS detector.

excluded from all column profile plots and can be optionally masked during preprocessing. Apart from this single artefact, the dark+bias frame demonstrates low fixed-pattern noise and provides a sufficiently stable calibration reference for Skycam preprocessing.

To characterise the twilight flat field of Skycam T, a master flat was created by median-stacking selected twilight exposures with uniform illumination. The resulting frame was normalised by its median and visually inspected for large-scale gradients and artefacts. As shown in Figure 4.1, the flat field exhibits smooth radial vignetting typical of wide-field camera lenses, with no significant dust shadows or discontinuities. The histogram of normalised values peaks near unity and extends down to ≈ 0.4, indicating up to 60% loss of sensitivity at the field edges. Median row and column profiles show smooth variation, confirming optical vignetting as the dominant effect and validating the flat's suitability for pixel sensitivity correction.

Quantitatively, the normalised flat has a mean of 0.96, a standard deviation of 0.30, and a flatness ratio (max/min) of 13.17. These values reflect both the optical vignetting and the lack of additional pixel-level noise structures.

4.4 Photometric Calibration

The goal of photometric calibration in this study is not to establish an absolute magnitude scale, but to determine relative zeropoints for each image and to track systematic shifts in brightness that reflect atmospheric transparency and system degradation (e.g. due to the accumulation of dust on the lenses). The instrumental magnitudes derived from the Skycam data serve as a basis for identifying extinction trends and variability across different observing nights.

Instrumental magnitudes were extracted using SExtractor (Bertin et al., 1996) and represent the ISO photometry ("MagIso"), which approximates the total flux for each detected object using an isophotal aperture. Because the Skycams operate without filters, these magnitudes reflect the system's broad unfiltered spectral response. No color transformations or attempts to convert to standard photometric systems were applied. Instead, instrumental magnitudes were compared directly to catalogue V-band magnitudes from the Tycho 2 cyatalogue (Høg et al., 2000) and the APASS catalogue (Henden et al., 2018b).

The use of unfiltered detectors in both Skycam T and Skycam Z necessitates a careful consideration of their effective spectral response when interpreting instrumental magnitudes. As illustrated in Figure 4.5, the total throughput of each system arises from the product of optical transmission and sensor quantum efficiency. Skycam Z, in particular, shows a response that closely overlaps with the Johnson V band, peaking near 540 nm (Figure 4.6). However, its sensitivity extends well beyond the V band into the red, resulting in a much broader overall response curve. Skycam T exhibits an even broader response, with enhanced sensitivity at longer wavelengths due to its optics and sensor characteristics, although no direct comparison with standard filter systems is available for it.

This mismatch between instrumental response and standard photometric passbands introduces a colour-dependent offset in the instrumental magnitudes when compared to catalogue V magnitudes. In particular, redder stars are expected to appear artificially brighter in the Skycam systems due to their broader throughput, especially beyond 600 nm. Nevertheless, extinction determination is based on differential trends (the change in photometric zeropoint with airmass) rather than the absolute calibration.

The homogeneity of the Tycho2 and APASS catalogues in providing calibrated V-band magnitudes enables reliable relative calibration across a broad range of stellar types. In this context, Skycam Z proves particularly valuable: its throughput overlaps the V band more closely than Skycam T, and its sharper system response and higher signal-to-noise at fainter magnitudes reduce photometric scatter.

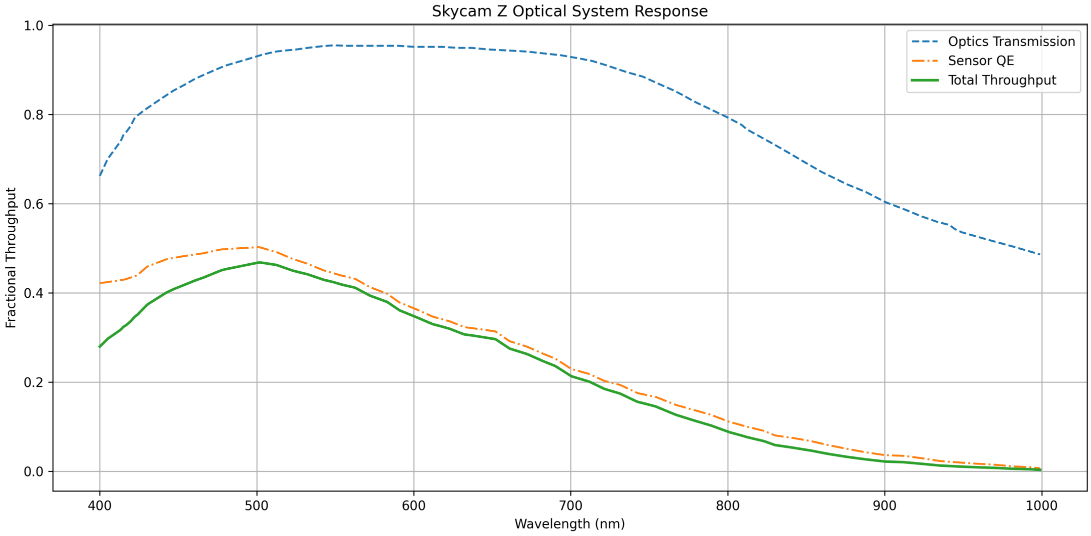

(a) Skycam T transmission curve

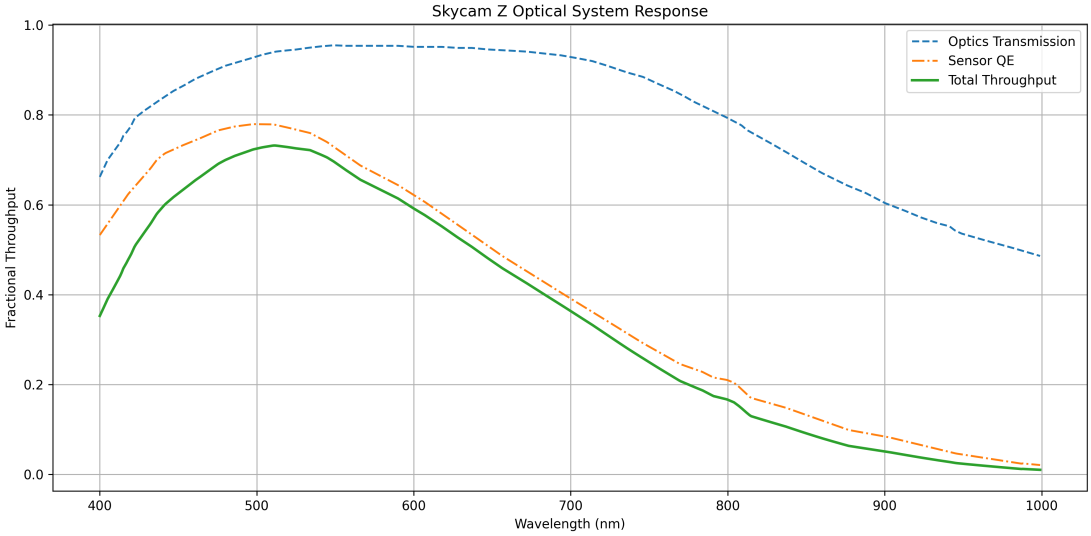

(b) Skycam Z transmission curve

Figure 4.5: Transmission curves of Skycam T and Skycam Z.

Catalogue matching was performed for each image using positional cross-referencing. For Skycam T, the Tycho2 catalogue was used for brighter stars and the APASS catalogue for fainter stars, while Skycam Z relied exclusively on APASS due to its deeper limiting magnitude. Matches were filtered by angular separation and assigned a “GoodFit” quality metric. The GoodFit value corresponds to the angular distance between the detected source and the matched catalogue star. If multiple detected sources fall within the matching radius of the same catalogue star, only the closest source (and, where applicable, the one consistent with the expected magnitude range) is retained as the valid match. All other candidate matches are flagged with `GoodFit` $= -1$ and excluded from

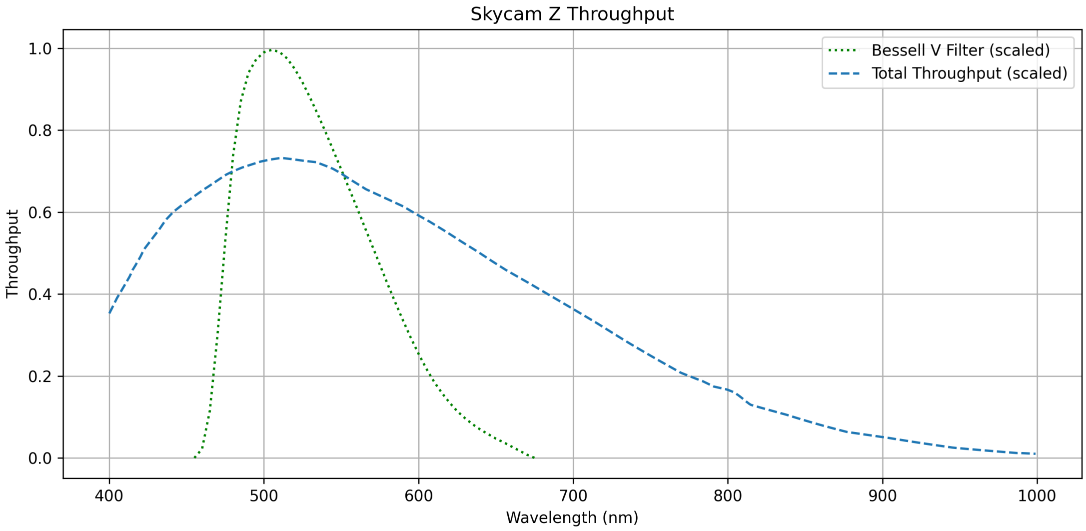

Figure 4.6: Comparison between the total system throughput of Skycam Z (dashed blue line) and the standard Bessell V filter response curve (dotted green line), expressed as dimensionless throughput (0–1). Both curves are normalised such that their maximum transmission value is equal to unity, allowing the relative shape and wavelength coverage of the responses to be compared independently of their absolute efficiency. The Skycam Z response reflects the combined effect of optics, sensor quantum efficiency, and atmospheric transmission, and shows a much broader wavelength sensitivity compared to the Bessell V band. This broader spectral response limits direct transformation to standard photometric systems but enables higher throughput for differential photometry. All data were obtained from values presented in the hardware documentation or from manufacturers' specifications (e.g. CCD, lens, OTA).

further calibration. Only entries with `GoodFit` > -1 were therefore used in the subsequent analysis. This step also removes duplicate matches and prevents faint spurious detections from being associated with bright catalogue stars.

For each matched star, the difference between the instrumental magnitude and the catalogue magnitude was computed:

This residual serves as a proxy for extinction once instrumental systematics are removed. Initial analysis revealed that these residuals can vary significantly within a night and across the field of view, particularly in Skycam T data. This increased variability is primarily attributable to Skycam T's much wider field of view, which introduces stronger vignetting effects, larger airmass gradients across the image, and increased susceptibility to scattered moonlight and sky-brightness gradients. In addition, its broader spectral response compared to the standard V band enhances colour-dependent offsets, further increasing photometric scatter.

Zeropoints were determined using three complementary strategies:

- A full linear regression of residuals against airmass.

- A robust RANSAC fit to remove outliers and spurious matches.
- A median-based approach that computes per-image residuals, bins them by airmass, and fits a line to the image-level medians.

These methods allowed for cross-validation and provided insight into the consistency and robustness of the derived zeropoints. The scatter of residuals and the variability of the fitted zeropoints from one model to another served as an indicator of photometric stability.

For quality control and diagnostic purposes, residual histograms, MagIso vs. Vmag scatter plots, and per-image zeropoint distributions were also examined. These visualisations helped identify problematic nights, saturation effects, and cases of poor catalogue matching.

4.5 Sky Brightness Measurement

Although primarily designed for extinction monitoring, the Skycam system inherently captures data suitable for estimating the night sky surface brightness. Accurate measurements of sky brightness, typically expressed in magnitudes per square arcsecond (mag/arcsec^2), are essential for evaluating the quality of astronomical observations, particularly in the context of faint-object imaging, background-limited exposures, and robotic scheduler-level signal-to-noise ratio (SNR) calculations.

In this section, the methodology used to derive sky brightness values from Skycam T and Skycam Z data is presented, and a systematic discrepancy observed in the Skycam Z results is investigated.

4.5.1 Methodology

Sky brightness was estimated using the background flux values reported by SExtractor, which computes a clipped mean in a mesh grid across the image to suppress contamination from stars and artefacts. The median of these local background values was adopted as the global background level per frame. This per-image background estimate, measured in analog-to-digital units (ADU), was then converted into astronomical units using the photometric zeropoint (ZP) derived individually for each image via catalogue matching.

Instrumental magnitudes were computed from matched stars by cross-referencing detected sources with the APASS or Tycho2 V-band catalogues. For each image, a linear regression between instrumental and catalogue magnitudes was performed, and the intercept was adopted as the image's zeropoint. For Skycam Z, this zeropoint was found to be remarkably stable, with a mean around +25.2 mag on dark, photometric nights, confirming the relative photometric stability of the system.

To express sky brightness in units of mag/arcsec^2, the following formula (4.1) was applied:

$$m_{\text{sky}} = \text{ZP} - 2.5 \log_{10} \left(\frac{B}{t_{\text{exp}} \cdot s^2} \right) \tag{4.1}$$

where B is the SExtractor background (in ADU), t_{exp} is the exposure time (10 seconds), and s is the plate scale in arcseconds per pixel (44.0 arcsec per pixel for Skycam T and 1.62 arcsec per pixel for Skycam Z). This equation effectively normalises the flux per square arcsecond and places it on an astronomical magnitude scale.

Using this method, time series of sky brightness were constructed with one data point per image (typically every minute). These values were binned by night, and for each night, both the median and 10th percentile sky brightness were computed. The 10th percentile serves as a proxy for the darkest portions of the night, free from transient contamination by the Moon, clouds, or telescope pointings at bright fields. This dual-statistic approach enables robust tracking of natural brightness cycles, seasonal variability, and the impact of lunar illumination.

4.5.2 Sky Background Levels

While working with Skycam Z data, it was noted that measured sky background levels during dark, photometric nights often fell below the expected sky brightness at the Observatorio del Roque de los Muchachos (ORM). Under moonless conditions, typical 'dark' V-band sky brightness at ORM is reported to be approximately 21.9 mag arcsec^{-2} (**benn1998**). When converted to instrumental background levels, the Skycam Z measurements implied systematically darker values than physically plausible for the site. This discrepancy raised questions about the accuracy of the dark subtraction pipeline and its impact on background-limited photometry.

A detailed cross-calibration with IO:O data, as well as an investigation into the Skycam dark frame methodology, revealed a systematic 1.8 mag offset attributable to over-subtraction. This analysis, along with the resulting correction and its implications for background modelling, is discussed in Chapter 5.

4.6 Atmospheric Extinction Calculation

4.6.1 Extinction Estimation Workflow

Atmospheric extinction was estimated from the difference between instrumental and catalogue magnitudes of stars detected in Skycam images. The images were processed using the source extraction and catalogue matching pipeline described in Section **??**, which performs astrometric calibration, source detection using SExtractor (Bertin et al., 1996), and positional cross-matching to reference catalogues.

For Skycam T, stars were matched against both the Tycho2 catalogue (Høg et al., 2000) and the APASS catalogue (Henden et al., 2018b), while Skycam Z relied exclusively on APASS due to its deeper sensitivity. Matches were filtered using a conservative angular separation threshold ($< 3\,\text{arcsec}$) and the GoodFit quality flag described earlier; only entries with `GoodFit` > -1 were retained.

For each matched star, the photometric residual

$$r = m_{\text{instr}} - m_{\text{catalogue}} \tag{4.2}$$

was computed, where m_{instr} denotes the instrumental magnitude derived from the unfiltered image (typically the SExtractor `MAG_ISO`) and $m_{\text{catalogue}}$ is the V-band magnitude from Tycho2 or APASS.

The residuals were then analysed as a function of airmass, calculated from the celestial coordinates of each star, the observation timestamp, and the observatory location. The coordinate transformation to horizontal coordinates was performed using the `AltAz` frame implemented in `astropy`.

4.6.2 Initial Photometric Consistency Check

To begin the extinction analysis, a representative photometric night - 30th September 2024 - was selected for detailed inspection. This night was characterized by excellent seeing, no visible cloud cover, and stable atmospheric conditions throughout. These ideal circumstances provided a well-controlled scenario to assess the photometric behaviour of both Skycam T and Skycam Z under clear sky conditions.

Skycam T: Tycho2 Comparison For Skycam T, all images from the selected night (e.g. see Figure 4.7) were retrieved together with their corresponding source detections obtained from SExtractor. These detections were cross-matched with stars from the

Tycho2 catalogue using positional association. For each matched star, the instrumental magnitude (MagIso) measured in the Skycam image was compared with the catalogue V_T magnitude from Tycho2.

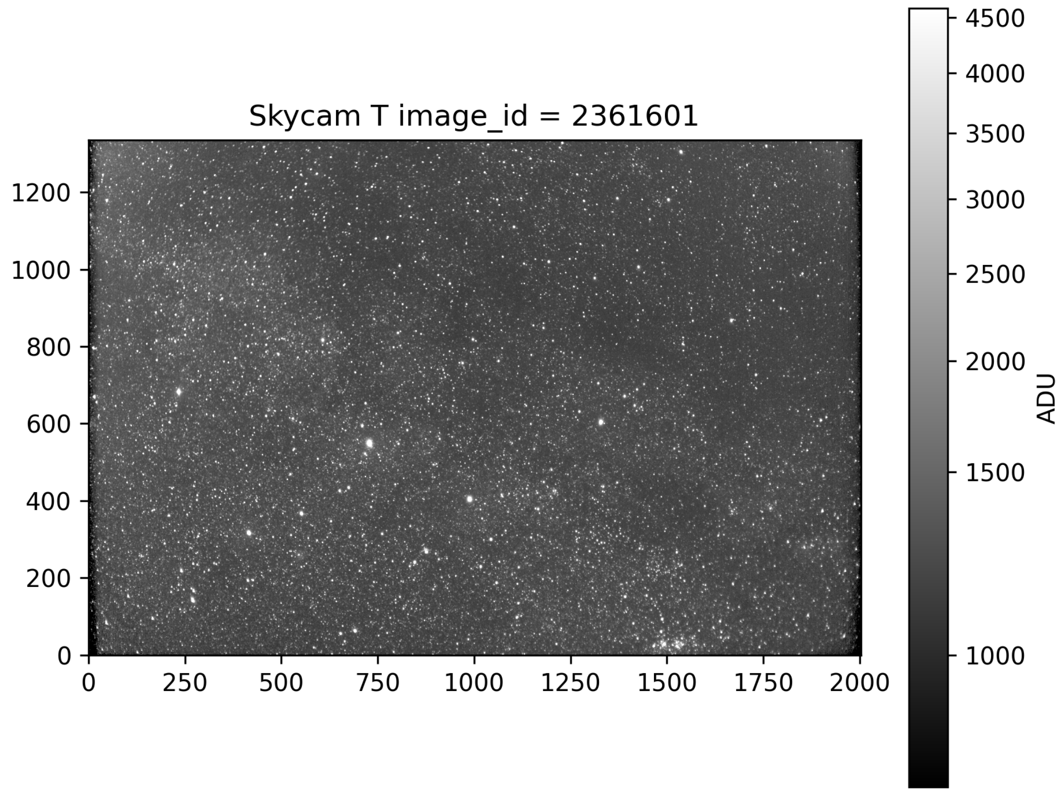

FIGURE 4.7: Skycam T image from 2024-09-30 22:23 used for the photometric consistency check. The frame shows a wide-field star field with no cloud contamination and good image quality, representative of stable observing conditions. This image yielded the highest number of Tycho2 catalogue matches on the night and served as the reference frame for validating the linearity of instrumental magnitudes.

The resulting distribution (see Figure 4.8) reveals a nearly linear trend in the reliable magnitude regime, with substantial deviation and scatter beyond MagIso > -11 and $V_T > 9$. These thresholds were therefore adopted throughout the analysis to filter unreliable photometric matches. This preliminary check confirms that Skycam T, despite its broad unfiltered response, captures a usable photometric signal — provided appropriate quality cuts are applied.

Skycam Z: APASS Comparison A similar procedure was applied to Skycam Z using images (e.g. Figure 4.9) from the same night. Image metadata were sourced from the `ZcamImage` table, and sources were loaded from the `StarsZSX` table. Cross-matches to the APASS catalogue were obtained from the `ZPSkycamZ` table in the database. The instrumental magnitude was again plotted against the catalogue V magnitude (see Figure 4.10).

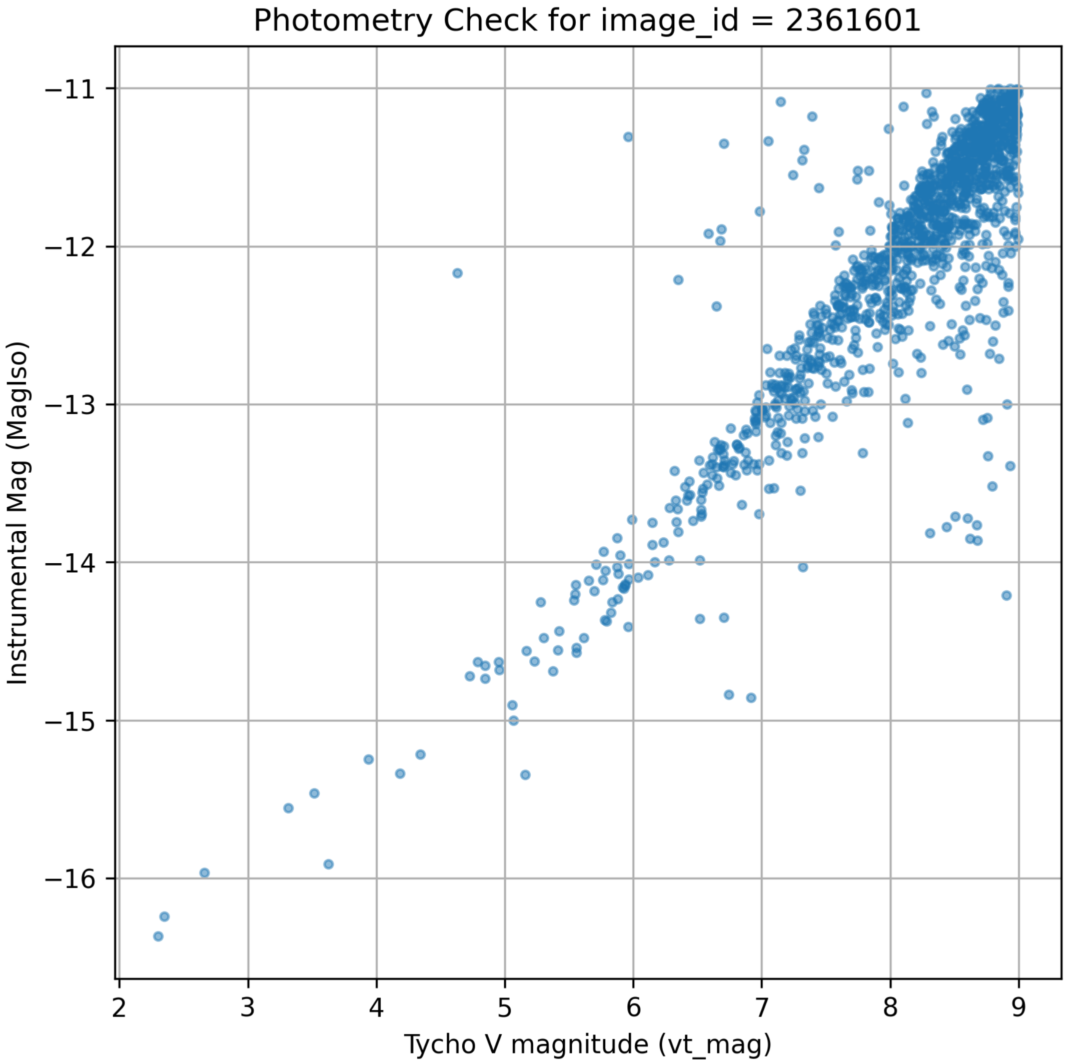

Figure 4.8: Photometric consistency check for the Skycam T image taken on 2024-09-30 22:23. Instrumental magnitudes (MagIso) are plotted against Tycho2 V_T magnitudes. A linear trend is visible in the reliable regime (MagIso > -11, $V_T < 9$), while the scatter increases significantly beyond this range.

Skycam Z displayed significantly better photometric performance. A tight linear relation between instrumental and catalogue magnitudes was observed across a wider dynamic range. Unlike Skycam T, no artificial truncation or scatter floor was evident in the bright or faint ends of the distribution. The system reliably tracked stars from $V \sim 10$ to $V \sim 16$, with mild scatter beyond that range.

These preliminary checks establish that Skycam Z provides superior photometric quality, likely due to its narrower field of view, improved optics, and higher-resolution sensor. Consequently, the extinction analysis described in the following sections emphasises Skycam Z data, with Skycam T results used primarily for comparison and qualitative insights.

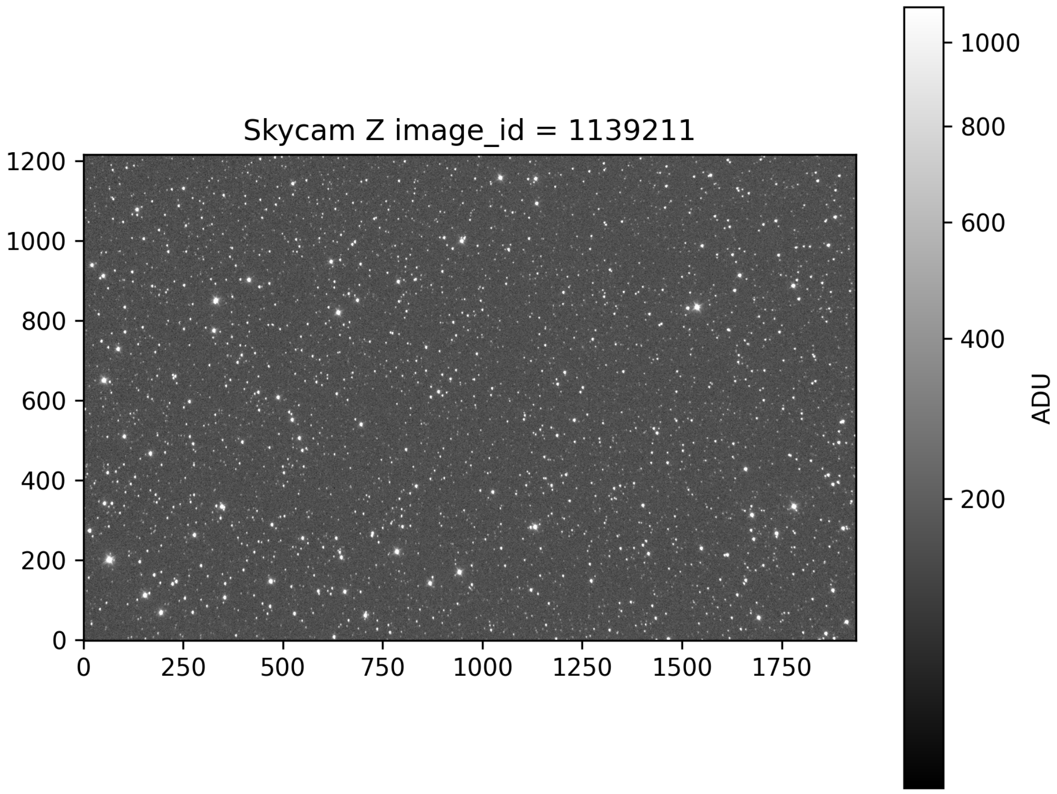

Figure 4.9: Skycam Z image from 2024-10-01 05:54, which yielded the highest number of APASS catalogue matches. The field is free of cloud and shows excellent star quality, representative of optimal observing conditions.

4.6.3 Full-Night Photometry: Instrumental vs. Catalogue Magnitude

Following the initial single-frame consistency checks, the full-night dataset for each instrument was analysed. All images taken on 2024-09-30 (Skycam T) and 2024-10-01 (Skycam Z) were processed, and the resulting matched stars were aggregated. For each matched star, the instrumental magnitude (MagIso) was plotted against the corresponding catalogue V magnitude: Tycho2 V_T for Skycam T (see Figure 4.11), and APASS V for Skycam Z (see Figure 4.12).

To evaluate both the match quality and the density distribution of photometrically useful sources, two types of visualisations were employed. First, scatter plots coloured by match distance were used to investigate spatial matching errors. Second, hexbin density plots were produced to assess the overall structure of the data and the distribution of matches across the photometric dynamic range.

Both instruments exhibit a high concentration of stars along a well-defined linear relation, confirming that the majority of matched sources are photometrically consistent. In both

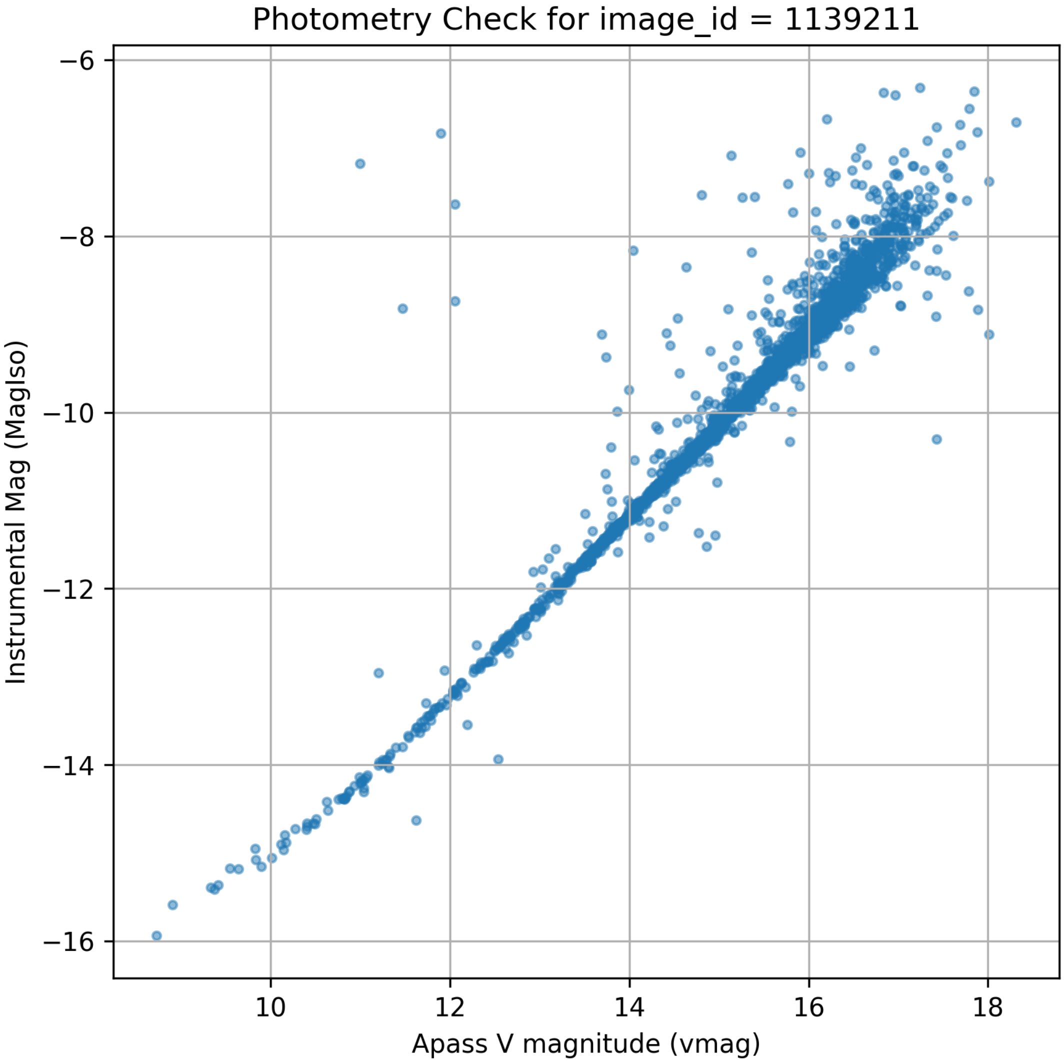

Figure 4.10: Photometric consistency check for Skycam Z on 2024-10-01 05:54. Instrumental magnitudes are plotted against APASS V magnitudes. A mostly linear relation is visible over the range $10 < V < 16$, with some scatter beyond $V > 15$, indicating the effective photometric range of the system.

cases, however, a deviation from strict linearity is observed at the bright end (i.e., lower-left region of the plots). This bending of the trend may be due to nonlinear detector response, blooming, or local pixel saturation effects, which are common when observing bright stars with fixed exposure times. Additionally, catalogue magnitudes (particularly from APASS) may be less reliable for bright sources due to issues with saturation and blending.

These density plots confirm that the magnitude cuts applied during filtering (e.g., MagIso < -11 for Skycam T) effectively isolate the usable photometric range and exclude the most distorted regions of parameter space. Skycam Z's sharper trend and wider dynamic range once again reinforce its suitability for extinction fitting and night-to-night

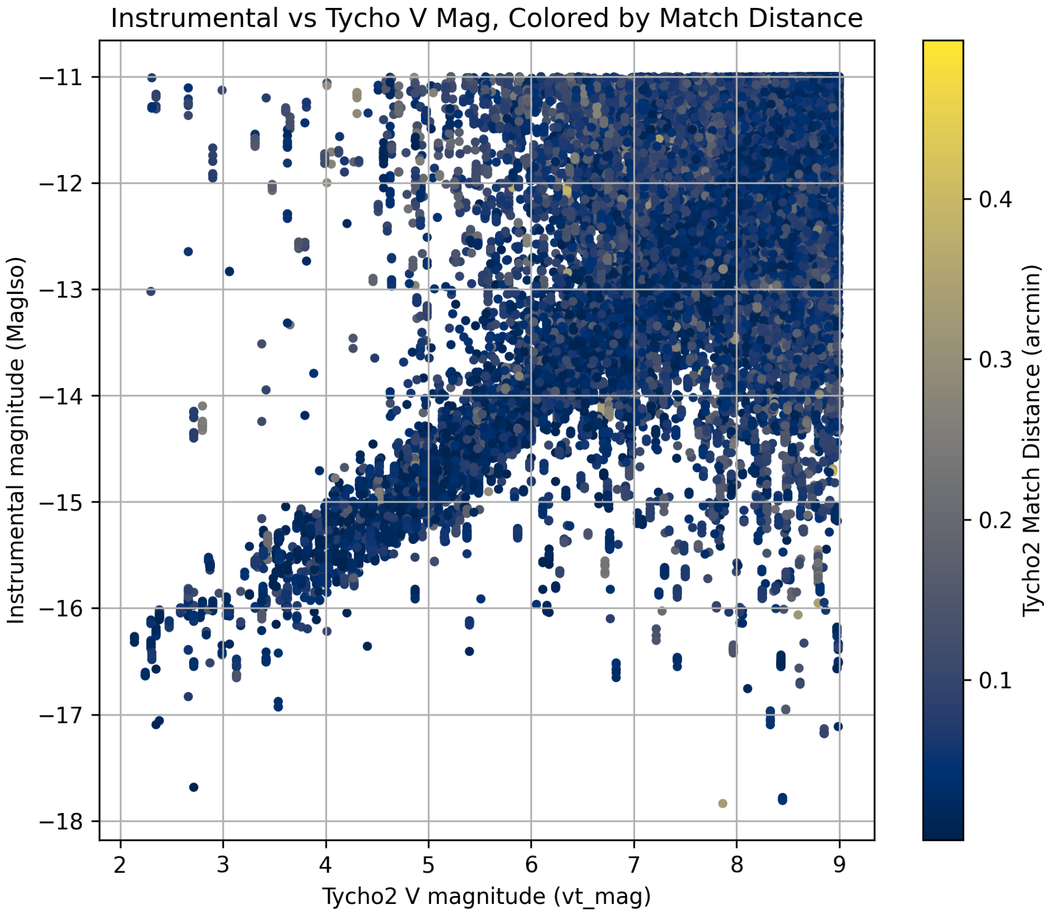

FIGURE 4.11: Skycam T: Instrumental magnitude (MagIso) vs. Tycho2 V_T magnitude for all matched stars on the night of 2024-09-30. Point colour encodes match distance (in arcminutes). No clear systematic trend between photometric residuals and match distance is visible, indicating that positional cross-matching errors are not the dominant contributor to the observed scatter. Increased dispersion beyond $V_T > 6$ is therefore primarily attributable to photometric noise, colour-dependent systematics, and the broad instrumental response rather than astrometric mismatch. The apparent truncation of the distribution at MagIso ≈ -11 reflects the selection limits of the quality filtering criteria rather than a physical cutoff in the detector response.

photometric monitoring.

4.6.4 Image-level Quality Assessment via Residual Scatter

To assess the photometric reliability of individual Skycam images, the standard deviation of the photometric differences (instrumental minus catalogue magnitude) was computed for each image in the night. These per-image scatter values serve as a proxy for photometric quality and were analysed statistically to identify and exclude anomalously noisy frames.

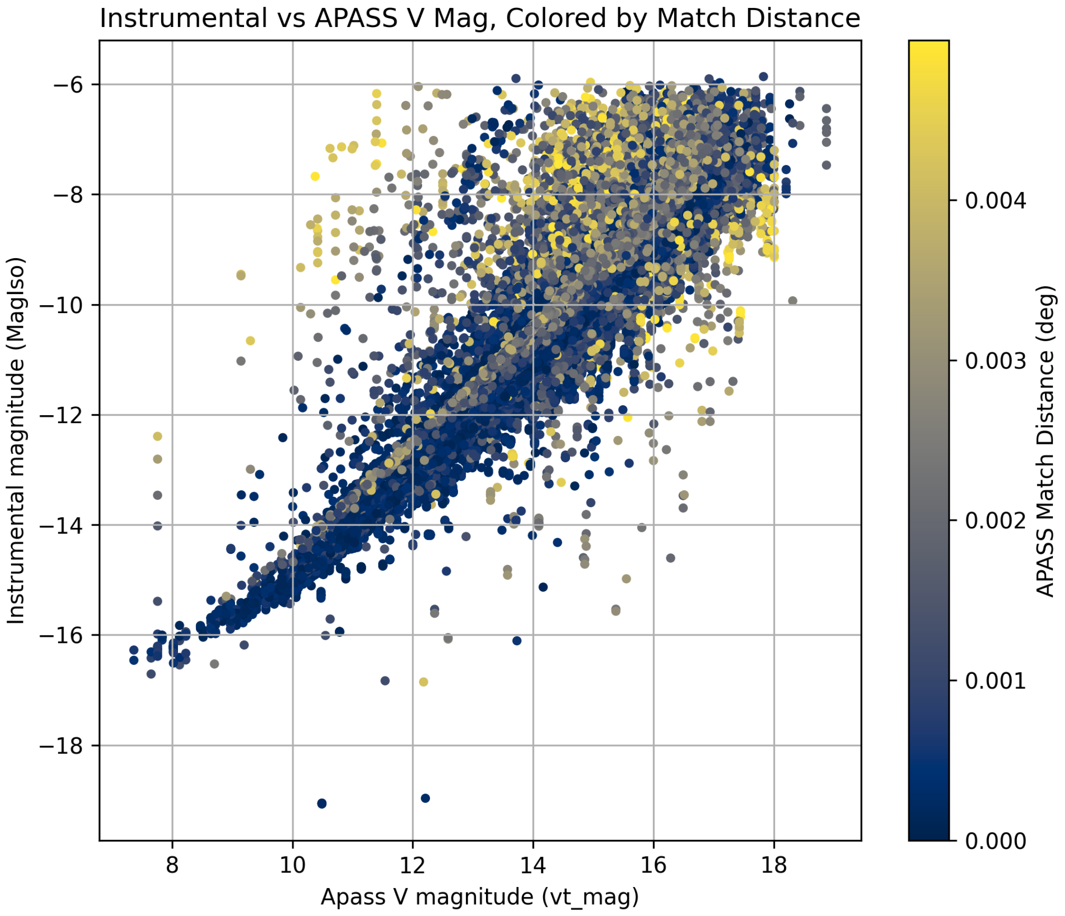

Figure 4.12: Skycam Z: Instrumental magnitude versus APASS V magnitude for all matched stars on 2024-10-01, coloured by match distance (in degrees). A clear linear relation is observed across the calibrated magnitude range. In contrast to the Tycho2 comparison, increased scatter and outliers correlate with larger match distances, indicating that positional cross-matching uncertainty contributes significantly to photometric residuals in the fainter APASS regime.

	Skycam T	**Skycam Z**
Usable magnitude range	$3 < V_T < 9$	$10 < V < 16$
Images passing quality cuts	312	311
Linear correlation (R^2)	0.36	0.72
Match radius threshold	< 0.5 arcmin	< 0.02 arcmin
Pixel scale	$\sim 44''$/pixel	$\sim 1.62''$/pixel
Field of view	$24^\circ \times 16^\circ$	$0.84^\circ \times 0.55^\circ$

Table 4.2: Comparison of Skycam T and Z photometric performance for a single photometric night

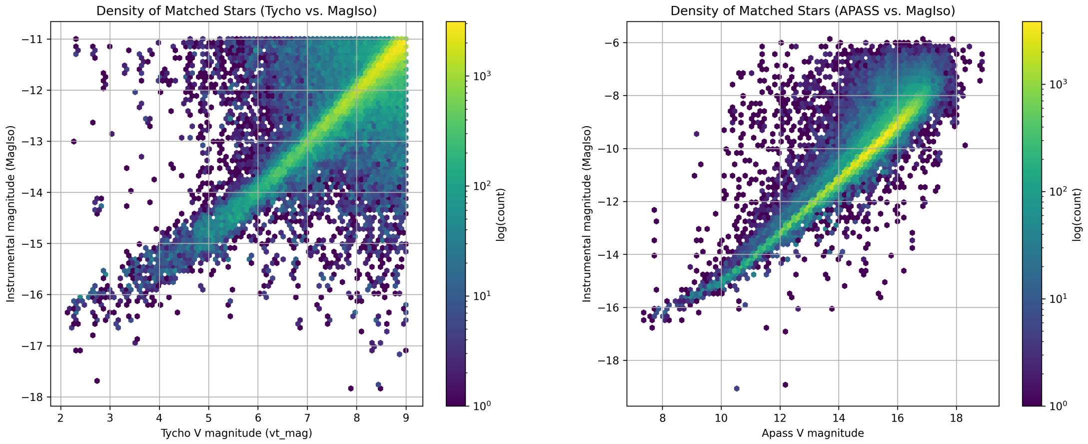

(a) Skycam T

(b) Skycam Z

FIGURE 4.13: Density plots (hexbin) of matched stars for a full photometric night. The logarithmic colour scale indicates the number of matched stars per bin. Skycam Z exhibits a tighter, more stable linear trend over a broader dynamic range. Curvature at the bright end is visible in both cameras.

Figure 4.14 shows the distribution of residual scatter values for both Skycam T and Skycam Z, along with Gaussian fits. For Skycam T, the distribution is broader and more asymmetric, with a long tail toward high-scatter values. The fitted parameters were $\mu = 0.56$ mag and $\sigma = 0.25$ mag, with many frames exceeding the typical noise floor. In contrast, Skycam Z displayed a more compact and symmetric distribution centered at $\mu = 0.35$ mag with $\sigma = 0.20$ mag.

While neither distribution is perfectly Gaussian, the comparative shapes are instructive. Skycam Z's tighter peak and lower mean scatter reflect improved frame-to-frame consistency, attributable to its narrower field, higher resolution, and better optical performance. Based on these distributions, a quality cut at $\mu + \sigma$ was adopted to exclude unreliable images prior to extinction fitting.

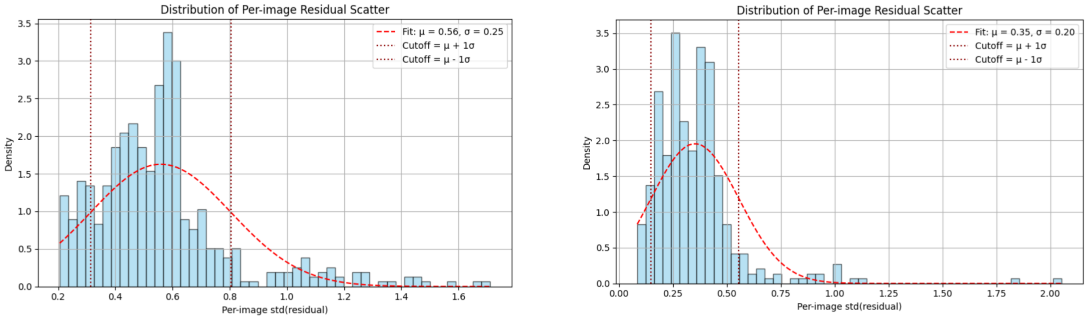

(a) Skycam T: $\mu = 0.56$, $\sigma = 0.25$

(b) Skycam Z: $\mu = 0.35$, $\sigma = 0.20$

FIGURE 4.14: Histogram of per-image photometric residual scatter. Dashed red curves show Gaussian fits. Skycam Z displays a narrower and more symmetric distribution, indicating greater frame-level photometric stability.

4.6.5 Photometric Zeropoint Distributions

To assess the overall photometric zeropoint offset and residual noise characteristics of each instrument, histograms of all individual star-by-star residuals were produced for a single photometric night (2024-09-30). Each residual was calculated as the difference between the instrumental magnitude (MagIso) and the catalogue V magnitude. This approach provides a global view of the magnitude zeropoint consistency and the intrinsic scatter of the photometric measurements for the night.

Figure 4.15 shows the residual distributions for Skycam T and Skycam Z. Both distributions exhibit approximate Gaussian shapes, with the fitted mean corresponding to the system's zeropoint (ZP) and the width (σ) indicating overall scatter due to measurement noise, optical distortions, or imperfect catalogue matching.

Skycam T shows a broader residual distribution centred at $\mu = -20.10$ mag with $\sigma = 0.54$ mag, while Skycam Z exhibits a tighter core distribution centred at $\mu = -25.06$ mag with $\sigma = 0.45$ mag. The Gaussian curves overplotted serve as simple parametric summaries of the central tendency and dispersion; however, the residual distributions are not perfectly Gaussian. In particular, the Skycam Z histogram exhibits noticeable asymmetry and extended wings, likely arising from colour-dependent offsets, residual matching errors, and spatial systematics across the field of view. The fitted μ therefore represents an effective zeropoint estimate for the dominant population, while σ approximates the core photometric scatter rather than a strict statistical standard deviation of a normal distribution.

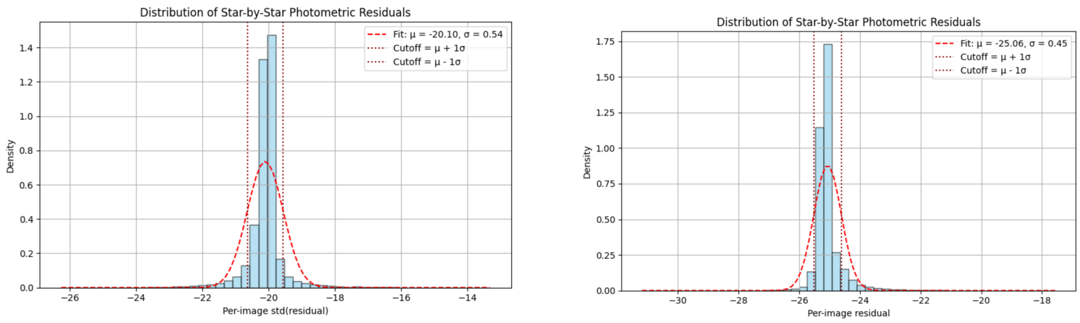

(a) Skycam T: $\mu = -20.10$, $\sigma = 0.54$

(b) Skycam Z: $\mu = -25.06$, $\sigma = 0.45$

Figure 4.15: Histogram of all photometric residuals (MagIso – catalogue V) for one photometric night. Overplotted Gaussian curves provide a descriptive estimate of the effective zeropoint (μ) and core scatter (σ), although the residual distributions show mild non-Gaussian structure and asymmetry.

4.6.6 Regression Approaches and Model Stability

To estimate the extinction coefficient k (in magnitudes per airmass) and the zeropoint ZP, three complementary regression strategies were employed, each method was designed to balance robustness against noise and computational efficiency, while leveraging different statistical assumptions about the nature of the data.

The first approach applied a standard Ordinary Least Squares (OLS) regression to all matched stars in a given night. For each star, the difference between its instrumental magnitude (MagIso) and catalogue magnitude (from APASS or Tycho2) was plotted against the computed airmass. The resulting cloud of points was fit with a linear model, yielding a slope corresponding to the extinction coefficient (k) and an intercept representing the system zeropoint (ZP). This approach has the advantage of simplicity and completeness, it considers the entire dataset without filtering. But it is inherently sensitive to outliers. These include mismatched stars, partially clouded frames, or observations near the horizon where refraction and extinction can vary.

To mitigate the influence of such outliers, the second method employed RANSAC (Random Sample Consensus) regression. This robust fitting technique iteratively selects random subsets of the data to propose candidate models and then evaluates which points (inliers) agree with that model within a specified residual threshold. Stars classified as outliers (e.g. due to poor match quality, extreme brightness deviations, or atmospheric effects) are excluded from the final regression. This method proved particularly effective for nights affected by thin clouds or for datasets where a small number of poor matches would otherwise dominate the fit. RANSAC was implemented using the scikit-learn library with a residual threshold typically set to 0.75 mag, and a minimum inlier fraction of 50% to ensure model stability.

The third strategy introduced an additional layer of data aggregation. Instead of fitting individual star residuals, the pipeline computed the median residual per image and associated each with the median airmass of stars detected in that frame. This image-level view significantly reduces the impact of outlier stars within a given frame, producing a smoothed sequence of per-image median residuals over the course of the night. A linear regression was then applied to these median values. This method effectively captures the overall extinction trend for photometric sequences and is more resilient to image-to-image noise caused by slewing, catalog cross-match failures, or transient clouds. It also proved useful for visualizing temporal trends in zeropoint stability.

Together, these three regression techniques allowed for cross-validation of extinction estimates. Agreement between the methods indicated a stable and photometric night, while discrepancies helped flag issues such as localized cloud cover or inconsistent photometry. Moreover, the methods revealed that while Skycam Z produced internally consistent fits across all three models, Skycam T often resulted in significant variation, highlighting the latter's limited suitability for quantitative extinction analysis without additional filtering.

A post-fit outlier rejection using the Median Absolute Deviation (MAD) of the residuals was also applied to retain high-airmass data points that follow the extinction trend while excluding anomalous measurements.

4.6.7 Instrumental Comparison: Skycam T vs. Skycam Z

Skycam T, while offering a longer time baseline, proved too noisy for consistent absolute extinction estimation. Residuals showed significant scatter, particularly at low airmass. Although relative transparency trends (e.g., whether a night was better or worse than the previous one) could be inferred, robust per-night extinction coefficients were unreliable in many cases.

In contrast, Skycam Z data (although limited to observations since September 2024) exhibited significantly better performance. Extinction fits were more stable, and the residuals displayed a consistent linear correlation with airmass under photometric conditions. A statistical comparison with historical extinction measurements from the Carlsberg Meridian Telescope (CMT) showed encouraging alignment, though the shorter data span currently limits definitive conclusions.

As a result, Skycam Z is identified as the more promising platform for future real-time extinction monitoring, while Skycam T is best suited for qualitative assessments and relative transparency tracking.

4.7 Benchmarking Skycam Z Extinction Against Calibrated Standards

4.7.1 Methodology: Cross-Matching with IO:O Standard Star Observations

To benchmark the atmospheric extinction coefficients derived from Skycam Z, a cross-comparison was performed against calibrated observations taken with the Liverpool Telescope's IO:O instrument using the Bessell V filter. The comparison focused on two standard stars observed regularly in the first half of 2025: `PG1047` and `TYC 568-1416-1` (internal name `114_654`). These stars are part of the calibration programme.

The selected time span covered a six-month interval from January to July 2025. Only observations taken under photometric conditions were retained, identified by inspection of the night logs[1]. Each standard star observation by IO:O was matched to a Skycam Z exposure taken within ±30 seconds, and the corresponding Skycam Z image was identified.

The zeropoint for each observation (IO:O and Skycam Z) was computed using:

$$\mathrm{ZP} = m_{\mathrm{catalogue}} - m_{\mathrm{instrumental}}, \tag{4.3}$$

where $m_{\mathrm{catalogue}}$ is the known V-band magnitude of the standard star, and $m_{\mathrm{instrumental}}$ is the measured instrumental magnitude.

The FITS files were analysed using Source Extractor (SExtractor, Bertin et al., 1996) to find and extract the instrument magnitudes of the standard stars.

To derive extinction coefficients, a linear regression was performed for each image and each instrument separately, fitting:

$$\mathrm{ZP} = \mathrm{ZP}_0 - k \cdot X,$$

where X is the airmass of the observation (derrived from the `Alt` FITS header of the IO:O observation) and k is the extinction coefficient.

A short statistical summary of the resulting extinction fits (combining both stars and all photometric nights) is given below:

- **Skycam Z**: Mean $k = 0.285$, standard deviation $\sigma_k = 0.014$, R^2 range: 0.4–0.5
- **IO:O**: Mean $k = 0.156$, standard deviation $\sigma_k = 0.012$, R^2 range: 0.3–0.4

The quoted R^2 range refers to the coefficient of determination obtained from individual nightly linear fits of residual versus airmass.

Non-parametric correlation tests confirmed the robustness of the extinction trend. For instance, the extinction series for `114_654` from Skycam Z yielded a Spearman rank correlation of $\rho = -0.654$ with $p = 0.0018$, confirming a statistically significant monotonic decline of zeropoint with increasing airmass.

---

[1] Available at `https://telescope.ljmu.ac.uk/Reports/`

4.7.2 Results: System Comparison

The extinction coefficients derived are summarised in Figures 4.16 and 4.17, showing the individual fits for each target star. The results illustrate that Skycam Z generally produces consistent extinction trends despite the lack of a physical filter in Skycam Z.

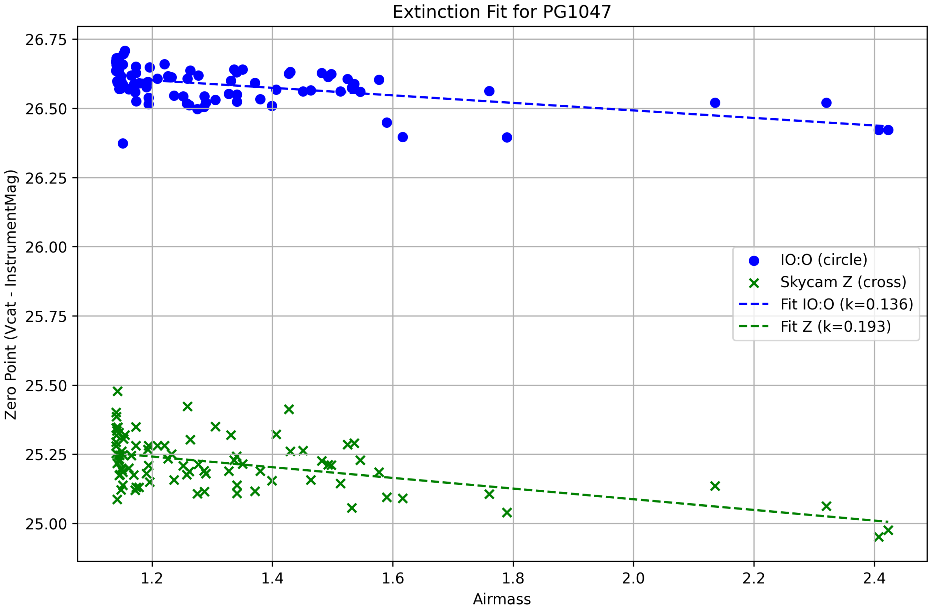

FIGURE 4.16: Extinction fit for the standard star `PG1047`, comparing IO:O observations (blue circles) with Skycam Z measurements (green crosses). Dashed lines indicate fitted linear regressions, yielding extinction coefficients of $k = 0.136$ (IO:O) and $k = 0.193$ (Skycam Z).

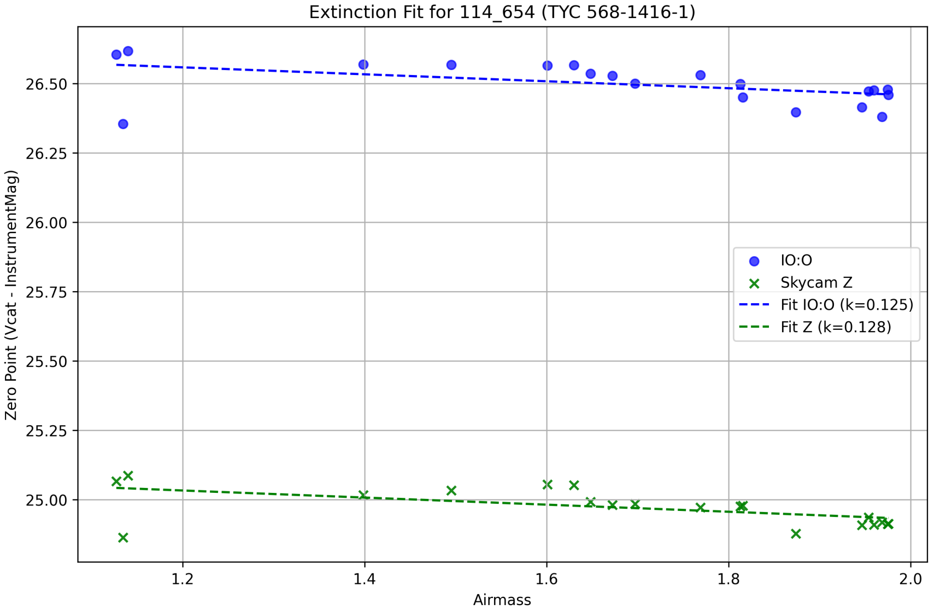

FIGURE 4.17: Extinction fit for the standard star `114_654` (TYC 568-1416-1). IO:O and Skycam Z yield consistent extinction coefficients of $k = 0.125$ and $k = 0.128$, respectively, with well-behaved residuals and linear trends.

Inst.	k	R^2	Spearman ρ	RMSE	MAE	Std. Dev.
IO:O	0.156	0.377	-0.614 ($p = 1.6 \times 10^{-12}$)	0.0607	0.0458	0.0610
Skycam Z	0.285	0.427	-0.544 ($p = 1.2 \times 10^{-9}$)	0.0998	0.0797	0.1002

TABLE 4.3: Extinction fitting and residual statistics for the combined sample of two standard stars (PG1047 and TYC 568-1416-1), observed between January and July 2025. k in mag/airmass, RMSE, MAE and Std.Dev. in mag

A combined extinction analysis was performed using observations of two standard stars, PG1047 and TYC 568-1416-1, taken between January and July 2025. The results are summarised in Table 4.3. Both IO:O and Skycam Z exhibit a clear correlation between zeropoint and airmass, with extinction coefficients of $k = 0.156$ mag/airmass for IO:O and $k = 0.285$ mag/airmass for Skycam Z. Despite Skycam Z's wider field and coarser photometric resolution, the measured extinction values are consistent with expectations. The Spearman rank correlation coefficients confirm statistically significant monotonic relationships in both cases ($\rho = -0.614$ for IO:O and $\rho = -0.544$ for Skycam Z, with $p \ll 0.01$), supporting the reliability of the fitted trends. Residual analysis shows that IO:O yields a lower scatter (RMSE $= 0.0607$ mag) compared to Skycam Z (RMSE $=$ 0.0998 mag), as expected from the higher precision of the primary instrument. Nevertheless, the performance of Skycam Z remains robust, highlighting its utility as an extinction monitor under typical observing conditions.

These results demonstrate that Skycam Z is capable of tracking atmospheric extinction even without a dedicated photometric filter. The agreement with IO:O is particularly notable given that no colour correction or filter compensation was applied to the Skycam Z data. This benchmark establishes confidence in the use of Skycam Z for extinction monitoring under photometric conditions.

4.8 Recommendations for Future Extinction Monitoring with Skycams

While the current Skycam Z system already provides extinction estimates under clear conditions, its performance is inherently limited by the lack of a well-defined spectral bandpass. The following subsections explore potential design upgrades aimed at improving the scientific utility and calibration accuracy of the system, ranging from a simple fixed-filter addition to a more advanced multi-band approach.

4.8.1 Fixed-Filter Upgrade

The simplest and most robust improvement would be the insertion of a fixed photometric filter, such as a Bessell V, into the optical path of Skycam Z. This would narrow the system's effective bandpass, reduce the colour dependence of instrumental magnitudes, and bring the spectral response into alignment with standard V-band photometry.

The main advantages of a fixed-filter upgrade are its simplicity and immediate photometric benefit. By aligning the spectral response of Skycam Z more closely with standard systems like the APASS V band, the comparability of instrumental magnitudes with catalogue values is significantly improved. This change reduces colour-dependent zeropoint offsets and allows extinction coefficients to be determined with lower residual scatter. Importantly, the upgrade introduces no moving parts, making it mechanically robust and well suited to the dynamic environment at the top-end of the telescope.

However, there are also important limitations to consider. The introduction of a filter inevitably reduces overall system throughput, which is particularly relevant given the fixed 10-second exposure time imposed by the Skycam software. A new (longer) exposure time would have to be determined. This will bring Skycam Z out of synchronisation with Skycam T but this should not matter to any currently running scientific projects with those instrument but should be considered if a future experiment will make use of both images. Depending on the filter transmission curve and its overlap with the sensor's quantum efficiency, the limiting magnitude may degrade by approximately 0.5 to 1.0mag. In addition, careful mechanical integration is required to avoid vignetting or alignment issues, especially since the optical tube offers only a narrow 2" clear aperture and already contains a collimator lens assembly.

Despite the expected loss in sensitivity, the addition of a single fixed filter is well-justified for extinction monitoring, where relative trends and stability are more critical than deep imaging performance.

4.8.2 Advanced Multi-Filter Concept: 2×2 Lenslet Array

A more ambitious upgrade path involves the use of a 2×2 lenslet array, with each quadrant of the image passing through a different broadband filter (e.g., B, V, R, I). This would enable simultaneous multi-band imaging without moving parts, allowing for colour-term calibration and independent extinction estimation in each band.

The proposed concept involves using a compact 2×2 lenslet array placed in the collimated beam path to replicate the same astronomical field onto four distinct sub-regions of a high-resolution CMOS detector. Each optical channel hereby consisting of a miniature lens and a fixed broadband filter (e.g., Bessell B, V, R, or I), projects a filtered version of the full image onto a different quadrant of the sensor. In this configuration, the same set of stars appears four times on the detector, each instance corresponding to a different photometric band.

This design offers several advantages. It enables true simultaneous multi-band photometry without any moving parts or sequential exposures, allowing extinction trends to be tracked independently in multiple filters under identical atmospheric conditions. Furthermore, the availability of colour information supports internal consistency checks and allows separation of extinction components (e.g., Rayleigh vs. aerosol scattering).

There are, however, several implementation challenges. Achieving uniform and co-registered imaging across all four channels requires highly precise alignment of the lenslets and filters, as well as careful control of optical path length and focus. Differences in PSF shape or optical distortion across quadrants may introduce systematic offsets in photometry. Additionally, since the full field is replicated multiple times, the effective image scale and spatial coverage are reduced, demanding a high-resolution sensor (e.g., ZWO ASI294MM or ASI2600MM) to retain adequate sampling and star count per sub-image. Furthermore, because the incoming photon flux is divided among the four optical channels, the per-band sensitivity is reduced compared to a single broadband detector, leading to lower signal-to-noise ratios in each individual band. The data reduction pipeline must segment the quadrants and associate each detected source across bands, which is manageable using modern image processing libraries but adds complexity to the workflow.

Mechanically, this upgrade is feasible within the 2" optical tube of the Orion AG8 system, particularly given the lack of thermal or contamination issues and the availability of custom 3D-printed mounts. Careful optical prototyping and calibration would be required to evaluate PSF stability across subfields.

4.8.3 Alternative Approach: Bayer-Matrix Photometry with a Colour CMOS Sensor

A practical intermediate solution between fully filtered photometry and unfiltered imaging is the use of a one-shot colour (OSC) CMOS camera with a Bayer matrix, such as the ZWO ASI2600MC. This camera uses an RGGB Bayer filter pattern integrated into the sensor, which separates the incoming light into red, green, and blue channels at the pixel level. After demosaicing, these measurements can be combined into three broad colour bands (RGB).

The quantum efficiency (QE) curve of the sensor, shown in Figure 4.18, reveals distinct spectral responses for the red, green, and blue channels, allowing them to act as coarse analogues to standard photometric bands. While these responses do not perfectly align with Bessell BVR filters, they provide sufficient wavelength separation to enable multi-band extinction tracking and approximate colour correction.

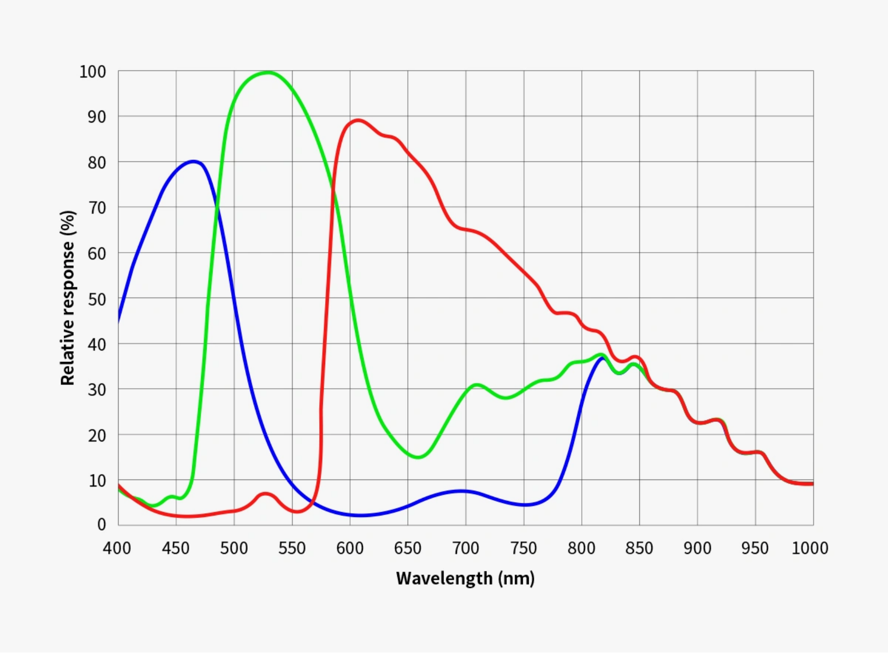

FIGURE 4.18: Quantum efficiency curves for the red, green, and blue channels of the Bayer matrix on the Sony IMX571 sensor (used in the ASI2600MC Pro). While the transmission curves are broader and differ from standard filters, they provide coarse spectral separation suitable for quasi-3-band photometry. Source: Suzhou ZWO Co. Ltd.

This approach offers several practical advantages. It enables simultaneous multi-band measurements in a single exposure, with no moving parts or filter wheels required. Since the colour separation is inherent to the sensor itself, the system remains mechanically simple and robust. Moreover, one-shot colour cameras such as the ASI2600MC are relatively cost-effective compared to monochrome cameras paired with filter wheels or custom optics. The camera can also be integrated into the existing Skycam Z optical configuration without major mechanical modifications.

However, the method also has notable limitations. The spectral response curves of the red, green, and blue Bayer channels are broad (especially the red channel) and overlapping, which reduces the precision of colour discrimination compared to standard photometric filters. In addition, the Bayer matrix reduces spatial resolution, since only half the pixels sample green and only a quarter each sample red and blue. Debayering and per-channel extraction (e.g. Ramanath et al., 2002) are required during preprocessing, adding some complexity to the pipeline. Finally, the effective bandpasses do not correspond directly to any standard photometric system and therefore require empirical calibration to establish colour transformations or extinction coefficients.

Despite these limitations, Bayer-matrix photometry represents an option for systems where mechanical simplicity and cost effectiveness are important. With appropriate calibration, such systems can deliver useful extinction and colour-trend information at low overhead, offering a promising upgrade path for instruments like Skycam Z.

4.8.4 Operational Considerations

All three proposed upgrades (the fixed photometric filter, the 2×2 lenslet array, and the Bayer-matrix colour CMOS camera) share a key operational advantage over traditional filter wheels: complete immunity to mechanical failure under fast-moving telescope conditions. Filter wheels pose a significant risk in the current Skycam Z configuration, which is mounted at the top-end of the Liverpool Telescope and is subjected to frequent slewing, vibrations, and varying gravitational orientations. Any mechanical jamming, misalignment, or cable wearing in such an environment could disable the system for extended periods and undermine its value as a passive monitoring instrument.

Consequently, any viable enhancement to Skycam Z must prioritise the following:

- **Mechanical simplicity**: avoiding motors, gears, or actuators that could fail or shift during telescope movement.
- **Robustness under dynamic load**: fixed components must retain optical alignment under acceleration and vibration.
- **Software compatibility**: modifications should maintain or minimally impact the existing data pipeline, which assumes a fixed field-of-view, stable orientation, and simple exposure cadence.

Among the upgrade paths, a fixed V-band filter remains the most straightforward and robust option. It requires minimal mechanical modification and would significantly improve the scientific calibration of extinction measurements. The lenslet array, while significantly more complex, offers multi-band capability and could serve as a prototype for next-generation auxiliary instruments. Finally, the use of a Bayer-matrix colour CMOS sensor presents an interesting intermediate solution. Mechanically simple, capable of quasi-3-band photometry, and compatible with existing data acquisition systems. With appropriate calibration, all three options seem compatible with the long-term goal of integrating real-time extinction monitoring into the robotic scheduling infrastructure of the Liverpool Telescope and the upcoming New Robotic Telescope.

At present, no final deployment strategy has been formally adopted. One possible pathway would involve retrofitting the existing Skycam Z unit at the Liverpool Telescope to provide calibrated real-time extinction measurements that could inform both LT operations and serve as a development testbed for NRT integration. Alternatively, a similar but purpose-built auxiliary instrument could be installed directly on the New Robotic Telescope. The hardware concepts discussed here are intentionally modular and mechanically simple, allowing adaptation to either scenario. The key objective is not tied to a specific telescope platform, but to establish a reliable, low-maintenance extinction monitoring capability compatible with modern robotic scheduling systems.

Chapter 5

Sky Brightness Modelling for La Palma

> We are trying to measure the light of things that are not there.
>
> Anonymous astronomer (unverified)

5.1 Introduction

The brightness of the night sky is a limiting factor in ground-based astronomical observations. Even under clear, moonless skies at premier observatory sites such as the Observatorio del Roque de los Muchachos (ORM), the sky exhibits a non-negligible background glow arising from atmospheric and astrophysical processes. This sky background contributes significantly to the noise budget of an observation, particularly for faint targets and in cases where exposure time is limited, such as during time-critical or fast-cadence observations.

Accurate estimation of the sky background is essential for computing reliable signal-to-noise ratios, particularly when observing those faint targets or during bright lunar phases. The sky background enters the SNR calculation as a dominant noise source for fainter targets and is highly sensitive to lunar illumination, Moon altitude, and the angular separation between the target and the Moon.

For robotic telescopes tasked with executing diverse observing programmes under dynamic environmental conditions, accurate knowledge of sky brightness is essential. While static sky models which are often parameterised by moon phase, target airmass, and angular separation from the Moon, provide rough estimates, they frequently fail to capture site-specific effects and short-term variability. Empirical corrections based on local data are therefore necessary to improve reliability.

Initial attempts to build the model based on Skycam zeropoints proved unsatisfactory (see Chapter 4), as the per-image scatter and calibration uncertainties (especially the per-image ZP) limited the accuracy of any derived trends. The modelling process was therefore changed: IO:O standard star observations served as the photometric anchor, and Skycam measurements were interpreted in relation to it. This approach avoids the pitfalls of relying on noisy zeropoints and enables meaningful cross-instrument comparisons.

5.2 Data Sources

5.2.1 IO:O Standard Star Observations

IO:O (Smith et al., 2017) is a calibrated optical imager on the Liverpool Telescope, equipped with standard photometric filters. For this study, a dataset of 4222 images of standard star observations taken in the Bessell V filter (to ensure maximum compatibility with Skycam observations) of photometric standard stars between March 2019 and July 2025 was used. Each exposure includes calibrated metadata and is processed using SExtractor to extract the background levels, expressed in units of magnitude per square arcsecond. The following standard stars were used in the analysis: `GSC 00048-00918`, `TYC 568-1416-1`, `BD+18 3407` and `GSC 04914-00008` from the catalogue presented in Landolt, 1992.

These IO:O measurements are treated as the reference for all subsequent comparisons and model calibration.

5.2.2 Skycam T and Z

Skycam T and Z are auxiliary cameras mounted on the LT. An in-depth description is given in Chapter 2.10. Neither camera is filtered, and their effective bandpasses are determined by optics and detector response.

Due to the lack of a mechanical shutter, dark current subtraction is handled via synthetic methods, resulting in systematic biases in the measured background level. In particular, a consistent offset of ∼1.8 mag has been observed between Skycam Z measurements and IO:O-calibrated sky brightness values (see section 5.3.1).

Despite these limitations, Skycam T and Z offer a uniquely rich dataset: images at 1-minute cadence spanning years, covering a wide range of lunar phases, target altitudes, and atmospheric conditions.

Rather than serving as photometric references or validation tools, the Skycam instruments are considered as instruments for monitoring relative changes in sky brightness over time. Their high cadence make them well suited for detecting short- and long-term temporal trends, such as the progression of lunar illumination, gradual shifts in seasonal sky brightness, or the influence of atmospheric conditions. While their unfiltered nature and calibration limitations preclude direct photometric comparison with standard systems, they nonetheless provide a valuable statistical overview of night sky conditions at ORM. Furthermore, when cross-calibrated with a well-characterised instrument such as IO:O, the Skycams can also be used to track atmospheric extinction trends. As demonstrated in Section 4.7, Skycam Z in particular shows a stable correlation with calibrated extinction measurements, indicating that unfiltered wide-field imagers can provide useful transparency monitoring even without dedicated photometric filters, provided that appropriate reference calibration is applied.

5.3 Cross-Calibration with the Skycams

5.3.1 Discovery of a Systematic Offset in Skycams T and Z

Application of the background estimation method to Skycam Z data revealed that calculated sky brightness values during dark, photometric nights frequently fell below physically plausible thresholds and often reached 23.7 mag arcsec^{-2}. This conflicts with established literature values for ORM, where typical dark-sky surface brightness in the V band lies around 21.9 mag arcsec^{-2} (Benn et al., 1999).

To investigate this discrepancy, independent observations from the IO:O instrument in the Bessell V filter were utilised for cross-validation. These observations provide a photometrically calibrated and spectrally comparable dataset. A representative example was drawn from an observation of the standard star HD 260655[1] (another frequently used Landolt star) on 29 January 2025 under good observing conditions (airmass 1.04, no moon) with the Bessel V filter. Using the published zeropoint of 24.95 mag (Smith et al.,

---

[1] `https://simbad.cds.unistra.fr/simbad/sim-basic?Ident=HD+260655`

2017), a sky brightness of 21.7 mag/arcsec^2(V − Band) was derived from the FITS file, consistent with expectations.

A systematic cross-match was then performed between this IO:O Bessell V standard star observation and the temporally closest Skycam Z image, revealing a 1.8mag offset (see Figure 5.1). To further investigate this offset, a large sample of IO:O Bessell V standard star observations from October 2024 to July 2025 was compared against Skycam Z images taken within ±30 seconds of each other. The resulting dataset contained 320 matched pairs.

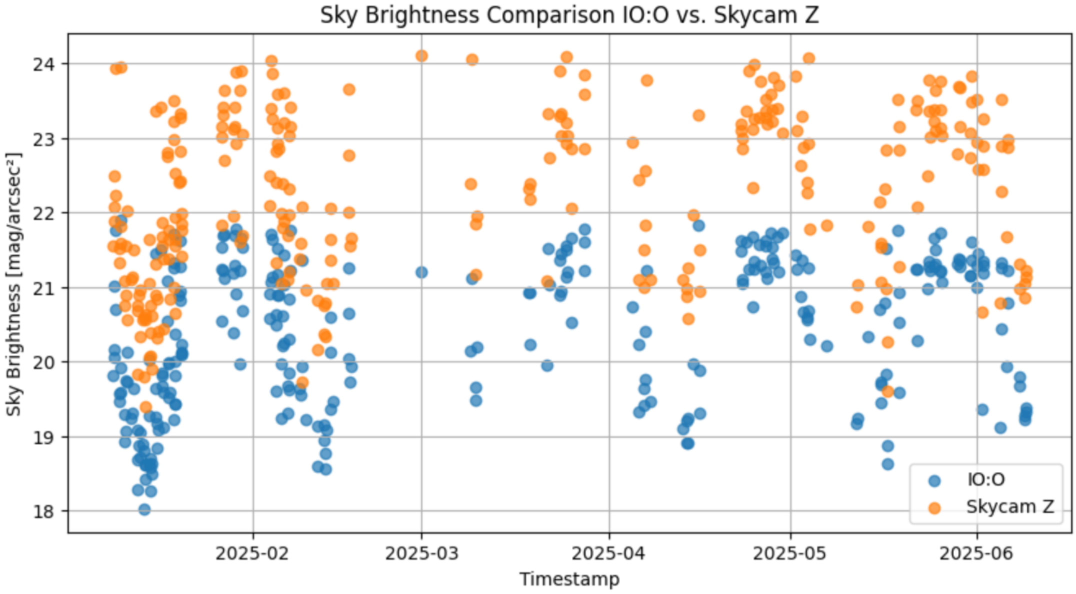

FIGURE 5.1: Comparison of sky brightness measurements from IO:O (blue) and Skycam Z (orange) between January and June 2025. Each point represents a calibrated estimate of sky surface brightness in magnitudes per square arcsecond (mag/arcsec^2), with IO:O using Bessell V filter photometry and Skycam Z measurements derived from matched exposures. A consistent offset is visible, with Skycam Z systematically reporting darker skies by approximately 1.8 mag due to over-subtraction in its dark correction pipeline.

An automated analyis of these 320 IO:O and Skycam Z images was performed to compute the difference in sky brightness between the two instruments. Both images were processed with the SExtractor tool. The median `Background` value from the SExtractor results were used in both cases and compared. The differences in sky brightness were then computed for each matched pair, resulting in a distribution of offsets.

The resulting offset distribution exhibited a Gaussian-like peak centred at +1.8 mag (see Figure 5.2), indicating that Skycam Z consistently underestimates the true sky brightness by a factor of approximately 5.25 in flux.

To further characterise this offset, several descriptive statistics were computed. The median offset was found to be $\langle \Delta m \rangle = 1.80$ mag, closely matching the peak of the distribution. The mean offset was slightly higher at 1.84 mag, with a standard deviation of $\sigma = 0.35$ mag. The interquartile range (IQR) was 0.41 mag, and the mean absolute deviation (MAD), a robust measure less sensitive to outliers, was 0.27 mag. Assuming normally distributed residuals, the 95% confidence interval for the mean offset was estimated as $1.80 \leq \mu \leq 1.88$ mag. These results confirm that the offset between Skycam Z and calibrated IO:O backgrounds is stable and systematic, supporting the use of a global correction factor of -1.8 mag to align Skycam Z background measurements with physically plausible values.

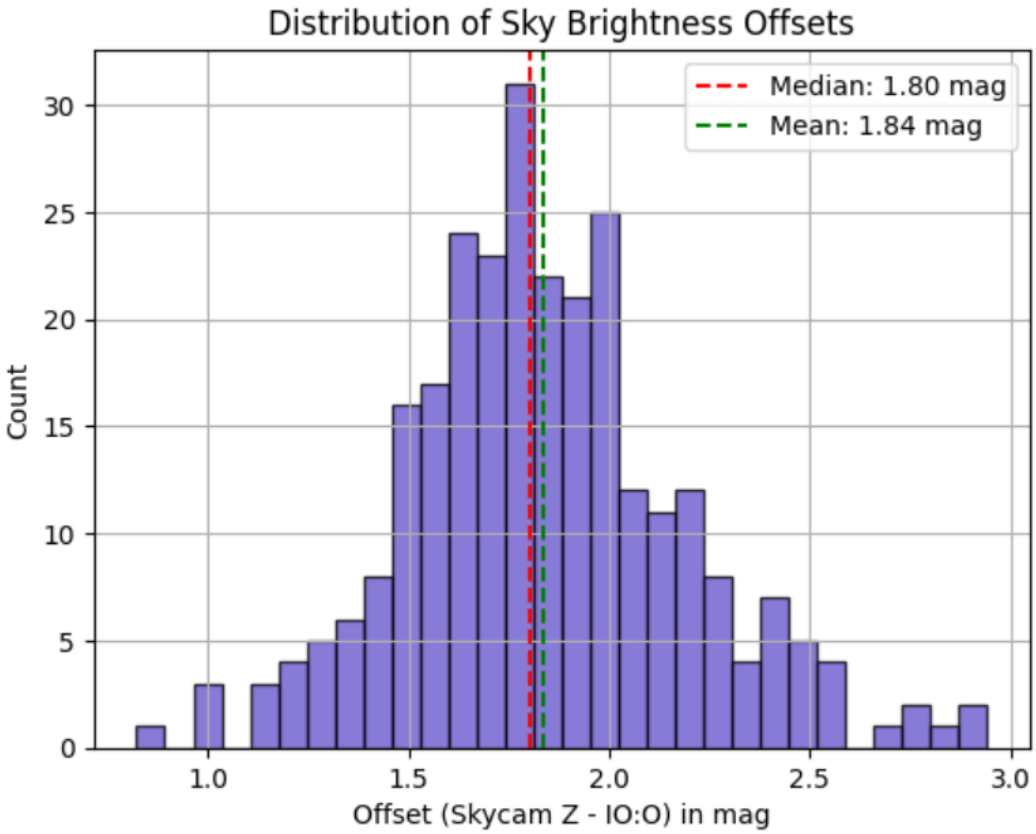

FIGURE 5.2: Distribution of sky brightness offsets between Skycam Z and temporally matched IO:O Bessell V standard star observations. The peak at +1.8 mag indicates a consistent underestimation of the sky background in Skycam Z images, attributed to systematic over-subtraction during the dark correction stage.

This systematic discrepancy prompted an investigation into the dark subtraction process employed in the Skycam Z pipeline. As the camera lacks a mechanical shutter, dark frames are generated by imaging the closed enclosure around midnight on nights where the LT is not operating (e.g. due to bad weather or technical circumstances) and applying heavy median filtering to infer a synthetic dark frame (see Section 4.3). Analysis of the darkbias frames (see Figure 4.4, Image 1) revealed a median value of approximately

540 ADU, while science image backgrounds routinely lie between 110 - 120 ADU. This subtraction therefore removes a substantial portion of the true background flux. The resulting over-subtraction is sufficient to account for the observed 1.8 mag underestimation in the sky brightness measurements.

To validate the persistence of this offset, the full archive of IO:O Bessell V standard star observations from January to June 2025 was cross-referenced with corresponding Skycam Z images. The observed offset remained consistently centred at +1.8 mag throughout the dataset (see Figure 5.2), confirming the systematic nature of the discrepancy and ruling out variability due to changing conditions or photometric uncertainty.

To evaluate whether a similar offset was present in Skycam T data, the same procedure was applied using the same IO:O Bessell V standard star observations as reference (see Figure 5.3). Despite Skycam T's wider field of view and different optical characteristics, a consistent underestimation of the sky brightness was also observed. The resulting distribution of magnitude offsets between Skycam T and IO:O backgrounds exhibited a peak near +1.6 mag (see Figure 5.4). Specifically, the median offset was 1.59 mag, with a mean of 1.57 mag and a standard deviation of 0.40 mag. The interquartile range (IQR) was 0.55 mag, and the mean absolute deviation (MAD) was 0.32 mag. The 95% confidence interval for the mean offset was estimated as $1.53 \leq \mu \leq 1.62$ mag. These values confirm that both Skycam T and Skycam Z exhibit a similar systematic offset relative to calibrated IO:O backgrounds, reinforcing the conclusion that the dark correction procedure - which is the same for both cameras - introduces a consistent flux underestimation across both instruments.

5.3.2 Correction and Implications

Given the systematic nature of the offset and the clear instrumental origin of the bias, a uniform correction of +1.8 mag was applied to all Skycam Z background measurements and a uniform correction of +1.59 mag was applied to all Skycam T background measurements. These values correspond to the median offsets derived from the cross-instrument comparison. Although individual measurements exhibit measurable dispersion about these central values, the dominant component of the discrepancy is a stable zeropoint shift. The correction therefore aligns the surface brightness estimates with independent IO:O calibrations and published site values at a first-order level. While this approach does not replace a proper dark-frame calibration, it provides an empirically motivated bias correction based on cross-instrument analysis.

It is noted that the Skycam instruments are unfiltered, with an effective bandpass determined by the convolution of sensor quantum efficiency, optical response, atmospheric

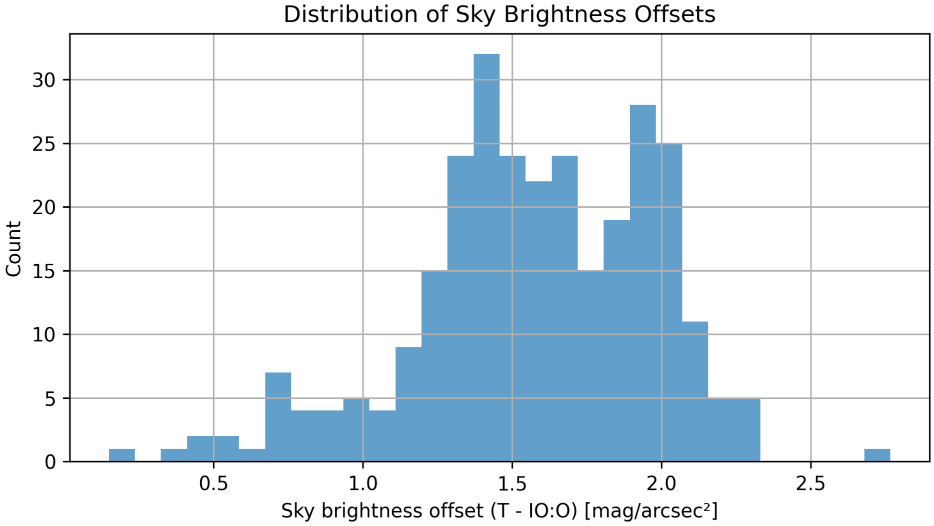

FIGURE 5.3: Comparison of sky brightness measurements from IO:O (blue) and Skycam T (orange) between January and June 2025. Each point represents a calibrated estimate of sky surface brightness in magnitudes per square arcsecond (mag/arcsec^2), with IO:O using Bessell V filter photometry and Skycam Z measurements derived from matched exposures. A consistent offset is visible, with Skycam T systematically reporting darker skies by approximately 1.8 mag due to over-subtraction in its dark correction pipeline.

transmission, and the spectral energy distribution of the sky. Although the measured brightness does not map directly to standard photometric systems (e.g., Johnson V), the internally consistent zeropoint calibration and comparisons with filtered IO:O images allow for a reliable brightness scale suitable for trend analysis and robotic scheduling applications.

5.4 Empirical Sky Brightness Model

5.4.1 Existing LT Sky Brightness Model

The Liverpool Telescope (LT) scheduling system currently employs an empirical sky brightness model developed by Ivan Baldry[2], based on several years of RATCam observations. This model estimates sky brightness in the Sloan r' band as a function of key astronomical parameters: solar elevation, lunar phase, lunar altitude, angular separation between the Moon and the target, and the target's own altitude (see table 5.1). It captures the complex interplay of twilight and moonlight conditions, acknowledging, for example, that a gibbous Moon near the horizon can produce similar sky brightness as a high-altitude crescent Moon.

---

[2] `https://telescope.ljmu.ac.uk/PropInst/Phase2/?sf=SkyBrightness`

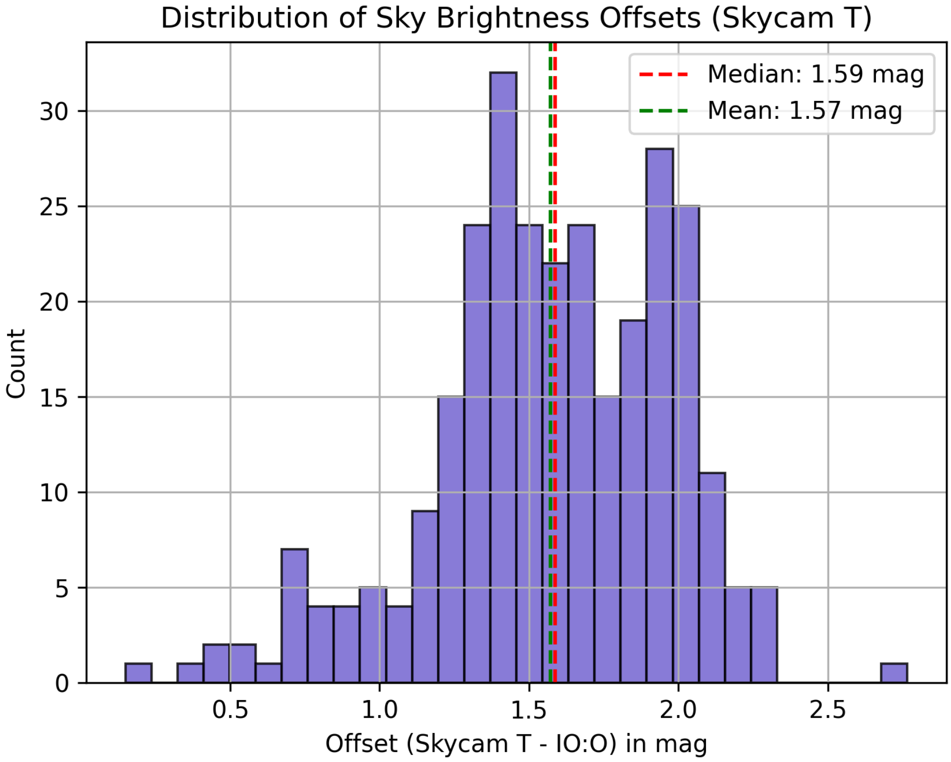

Figure 5.4: Distribution of sky brightness offsets between Skycam T and temporally matched IO:O Bessell V standard star observations. The peak at +1.57 mag indicates a consistent underestimation of the sky background in Skycam T images, attributed to systematic over-subtraction during the dark correction stage.

Sky brightness is treated as a constraint within the LT scheduler, expressed relative to the best achievable dark-sky conditions. Specifically, users set an upper limit on how many magnitudes brighter than dark sky the observation can tolerate, e.g. a "2 mag" constraint means the scheduler will only consider times when the sky is at most 2 magnitudes brighter than ideal conditions.

Traditional observing categories such as "civil", "nautical", and "astronomical" twilight are approximately mapped to numerical thresholds in measured sky brightness. The following correspondence is used:

The brightness constraint is always applied in the r' band, regardless of the instrument or filter actually used. Empirically, the u' and g' bands are observed to be about 1 magnitude more sensitive to twilight and moonlight, while the z' band is about 1 magnitude less affected. The criterion remains relative to the optimal dark sky in the respective band.

Observing groups that do not specify a sky brightness constraint are scheduled without restriction as soon as the Sun drops below the horizon. While technically valid, this

Twilight Type	**Sun Alt. (deg)**	**Moon Cond,**	**Measured r'**	**Scheduler Class**
Civil	< 4	N/A	< 10 mag	10 mag
Nautical	8–12	N/A	< 6 mag	6 mag
Astronomical	12–16	Bright / close	< 4 mag	4 mag
Astronomical	12–16	None / far	< 2 mag	2 mag
Dark	> 16	Gibbous, near	< 2 mag	2 mag
Dark	> 16	Gibbous, far	< 1.5 mag	1.5 mag
Dark	> 16	Crescent, near	< 0.75 mag	0.75 mag
Dark	> 16	None or very far	< 0.3 mag	dark

Table 5.1: Empirical classification of sky brightness conditions at the Liverpool Telescope (LT). The scheme combines Sun altitude, lunar phase and separation, and measured r'-band background levels. Here r' denotes the increase in sky brightness in the Sloan r' band relative to an ideal dark sky, expressed as a magnitude difference. The resulting categories are mapped to scheduler classes, which define the limiting background brightness under which an observation can be executed.

setting is typically discouraged unless explicitly intended.

Availability of each sky brightness class depends on lunar phase, moon-target separation, and target declination. For instance, high-declination fields are less affected by the Moon's path than equatorial ones. On average, conditions equivalent to or darker than the "2 mag" class are available approximately 65% of the time throughout the year, though availability may drop sharply near full Moon. This temporal variation is visualized in scheduling tools through cumulative and differential sky brightness availability plots across lunations.

To assist Phase II users, the LT Phase2UI provides illustrative pictograms that map common observing scenarios (e.g. full Moon, twilight, crescent Moon near target) to brightness classes, allowing observers to make informed decisions about constraint selection.

This empirical model provides a practical and well-tested baseline for sky brightness classification and constraint handling in robotic scheduling. However, it remains a heuristic model with some simplifying assumptions. In the following section, the Krisciunas & Schaefer (1991) model is described, which offers a more physically motivated approach than that used as the foundation for the improvements presented in this chapter.

5.4.2 Baseline Model: Krisciunas & Schaefer (1991)

A widely used semi-empirical model for estimating night sky brightness under moonlit conditions was developed by Krisciunas & Schaefer (Krisciunas et al., 1991). The model provides a formula to estimate the additional brightness contribution of scattered moonlight, expressed in magnitudes per square arcsecond, as a function of geometric and atmospheric parameters.

The model accounts for the following key variables:

- the lunar phase (as illuminated fraction or phase angle),
- the Moon's altitude above the horizon,
- the angular separation between the Moon and the target field,
- the target's altitude (airmass), and
- the atmospheric extinction coefficient k.

The formulation is based on physical principles of Rayleigh and Mie scattering, with empirically tuned coefficients to match observations made at Mauna Kea Observatory using a 15 cm Newtonian telescope and a photometer with UBV filters. The model outputs an additive term ΔV representing the increased brightness of the night sky relative to dark-sky conditions, typically assuming a nominal base value of $21.9\,\mathrm{mag\,arcsec^{-2}}$ in the V band.

In this work, the Krisciunas & Schaefer model is adopted as a reference for describing sky brightness at the Observatorio del Roque de los Muchachos (ORM). While the general behaviour of the model is consistent with trends observed in Skycam T and Z data - such as the dependence on Moon phase and angular separation - systematic deviations were identified when compared to actual measurements. These include a consistent brightness underestimation during partial moon phases and at low lunar altitudes. As a result, a regression-based empirical correction was developed to improve the agreement between modelled and observed brightness values, as described in the following sections.

5.4.3 Empirical Correction to the Baseline Model

While the Krisciunas & Schaefer model provides a physically motivated prediction of sky brightness under moonlit conditions, systematic discrepancies were observed when comparing its output to calibrated IO:O observations at the ORM. In particular, the model tended to over- or underestimate sky brightness depending on lunar phase (Moon disk size), Moon altitude, target altitude and angular separation of the Moon and the target.

To account for these residual trends, an empirical correction was derived by analysing the difference $\Delta m = M - C$, where M denotes the measured sky brightness from IO:O and

C the calculated value from the original Krisciunas et al., 1991 model, both expressed in units of magnitude per square arcsecond. This residual was regressed against four parameters: fractional lunar disk illumination d, Moon altitude h_m, target altitude h_t and the angular separation Θ between Moon and target.

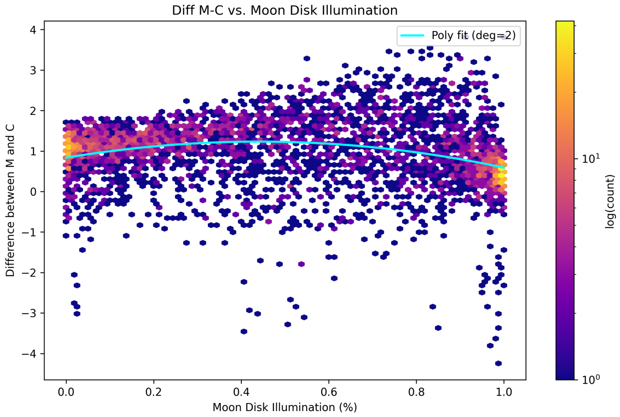

FIGURE 5.5: Residuals $\Delta m = M - C$ between the Krisciunas & Schaefer model and calibrated IO:O sky brightness measurements as a function of Moon disk illumination. Cyan line shows the fitted polynomial used in the empirical correction model.

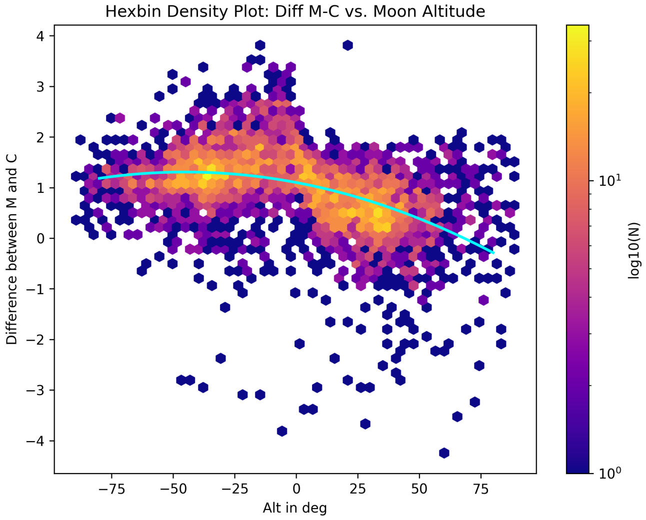

FIGURE 5.6: Residuals $\Delta m = M - C$ between the Krisciunas & Schaefer model and calibrated IO:O sky brightness measurements as a function of Moon altitude. Cyan line shows the fitted polynomial used in the empirical correction model.

Moon Disk Illumination The residuals exhibit a non-linear dependence on lunar disk illumination. As shown in Figure 5.5, the difference between modelled and measured brightness is smallest during crescent phases and increases toward intermediate phases before slightly decreasing again near full Moon. A second-degree polynomial provides a

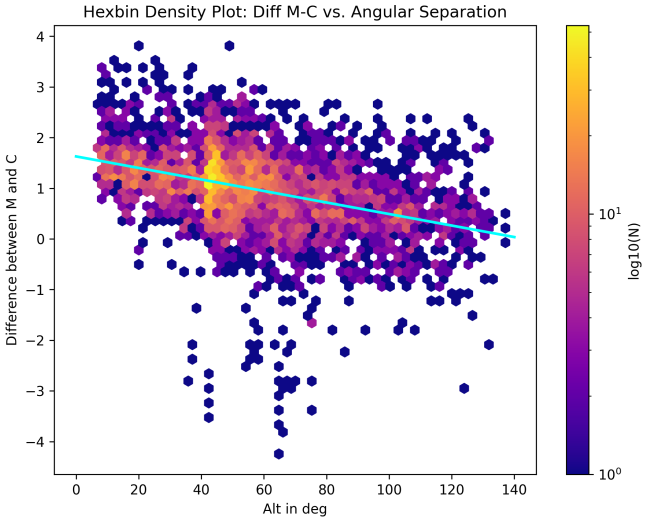

Figure 5.7: Residuals $\Delta m = M - C$ between the Krisciunas & Schaefer model and calibrated IO:O sky brightness measurements as a function of Moon–target angular separation. Cyan line shows the fitted polynomial used in the empirical correction model.

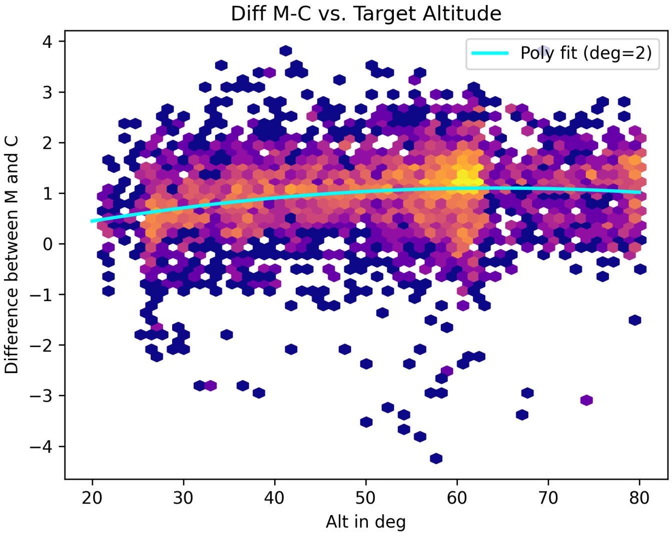

Figure 5.8: Residuals $\Delta m = M - C$ between the Krisciunas & Schaefer model and calibrated IO:O sky brightness measurements as a function of target altitude. Cyan line shows the fitted polynomial used in the empirical correction model.

good fit to this trend, capturing the asymmetric response of the sky brightness to partial lunar illumination.

Moon Altitude A similar second-degree trend is observed with respect to Moon altitude (Figure 5.6). Residuals are largest when the Moon is near the horizon or below it and decrease when the Moon is high in the sky.

Angular Separation The residuals show a clear negative correlation with angular separation between the Moon and the target (Figure 5.7). Smaller separations result in larger residuals, indicating that the original model underestimates the effect of scattered moonlight at small Moon - target angles. A linear fit is sufficient to capture this trend and was incorporated directly into the correction model.

Target Altitude Figure 5.8 shows the residuals as a function of target altitude. A weak second-order trend is present, but the overall dependence is modest compared to the other parameters. Nevertheless, a quadratic term was retained in the model to account for small altitude-dependent effects not captured by the baseline model.

5.4.4 Final Correction Model

Second-degree polynomials were fitted to the residuals as functions of d, h_m, and h_t, while a linear term was used for Θ. The resulting correction function takes the form:

$$\Delta m = a_1 d^2 + a_2 d + a_3 h_m^2 + a_4 h_m + a_5 h_t^2 + a_6 h_t + a_7 \Theta + a_0 \tag{5.1}$$

All angles are expressed in degrees, and the fitted coefficients a_i were obtained through ordinary least squares regression using the full dataset of IO:O standard star observations from March 2019 to July 2025. The final sky brightness prediction is then computed as:

$$S_{\text{corrected}} = S_{\text{K\&S}} + \Delta m \tag{5.2}$$

In principle, the site-specific discrepancies between the predictions of the Krisciunas & Schaefer model and the measured sky brightness at ORM could also be addressed by re-fitting the original model parameters for the local site conditions. However, the Krisciunas & Schaefer formulation contains several empirically derived coefficients embedded within exponential and geometric terms whose physical interpretation is not straightforward. A direct re-optimisation of these parameters would therefore risk obscuring the original physical structure of the model and complicating comparisons with previous studies. For this reason, the approach adopted here preserves the published formulation and introduces an empirical correction layer derived from ORM observations. This strategy retains the physical interpretability of the baseline model while accounting for systematic site-specific deviations.

5.4.5 Model Evaluation

5.4.5.1 Comparison with LT Heuristic Model

To benchmark the corrected sky brightness model against operational scheduling tools, a comparison was performed with the empirical model currently used by the Liverpool Telescope (LT) scheduler. This model, derived from RATCam data, classifies observing conditions into discrete brightness classes based on Moon phase and Moon–target angular separation. For this comparison, a simplified version of the LT model was implemented for astronomical night (Sun altitude $< -16°$) and converted into approximate sky brightness predictions in the Bessell V band.

The conversion assumes that the LT model operates in the Sloan r' band and that g'-like bands such as Bessell V exhibit approximately 1 magnitude greater sensitivity to lunar illumination. Each LT brightness class was therefore offset by $+1$ mag and subtracted from the canonical dark-sky level of 21.9 mag arcsec^{-2} to yield comparable sky brightness estimates:

$$S_V = 21.9 - (\Delta m_{\text{class}} + 1.0)$$

Figure 5.9 shows the measured IO:O sky brightness versus the predicted values from the LT model. Due to the class-based nature of the model, the predictions cluster at discrete levels and show substantial residual scatter when compared to calibrated observations.

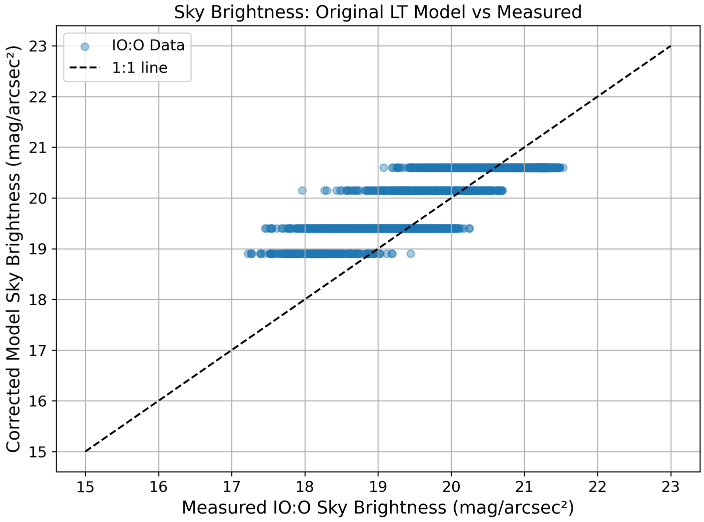

FIGURE 5.9: Measured sky brightness from IO:O vs. LT scheduler model predictions (converted to Bessell V). The discretised nature of the LT model introduces visible banding and limits predictive accuracy.

Quantitatively, the LT model yields a root mean square error (RMSE) of 1.20 mag, a mean absolute deviation (MAD) of 1.02 mag, and a coefficient of determination R^2 =

0.086. These metrics are comparable to or slightly worse than the Krisciunas & Schaefer model, and significantly inferior to the corrected empirical model developed in this work (see Table 5.2).

Table 5.2: Final comparison of sky brightness models against IO:O data.

Model	**RMSE (mag)**	**MAD (mag)**	R^2
LT Heuristic Model (approx. V)	1.20	1.02	0.086
Krisciunas & Schaefer (1991)	1.27	1.09	-0.02
Corrected Model (this work)	0.61	0.41	0.76

These results support the use of the corrected empirical model as a more accurate and physically interpretable alternative to the current heuristic classification approach. Its continuous output in calibrated units makes it particularly suitable for dynamic scheduling and SNR-based exposure calculations.

5.4.6 Comparison with the Krisciunas & Schaefer Model

The performance of the empirically corrected sky brightness model was evaluated by comparing its predictions against the calibrated measurements from the IO:O instrument. For reference, the original Krisciunas & Schaefer (1991) model was also evaluated on the same dataset.

Figures 5.10 and 5.11 show scatter plots of predicted versus measured sky brightness for both the original and the corrected models. The Krisciunas model systematically overestimates sky brightness during many lunar conditions and shows a large spread around the 1:1 line. In contrast, the corrected model shows significantly improved alignment, with reduced scatter and better adherence to the expected diagonal.

A histogram of residuals (Figure 5.12) further illustrates the improved performance of the corrected model. The residuals from the Krisciunas model exhibit a broad, asymmetric distribution centred around $+1$ mag, while those from the corrected model are sharply peaked around zero with a narrower spread.

Quantitatively, the improvement is summarised in Table 5.2, showing root mean square error (RMSE), mean absolute deviation (MAD), and the coefficient of determination (R^2) for both models. The corrected model reduces the RMSE by more than half and achieves an R^2 of 0.76, indicating that over 75% of the variance in sky brightness can be explained by the adapted model.

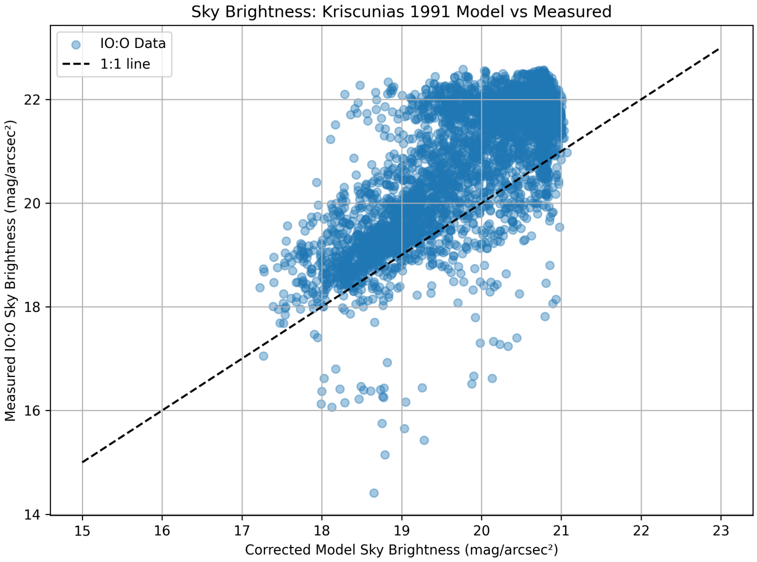

FIGURE 5.10: Comparison of model-predicted sky brightness with IO:O measurements. The corrected model shows improved correlation and reduced bias relative to the original Krisciunas & Schaefer formulation.

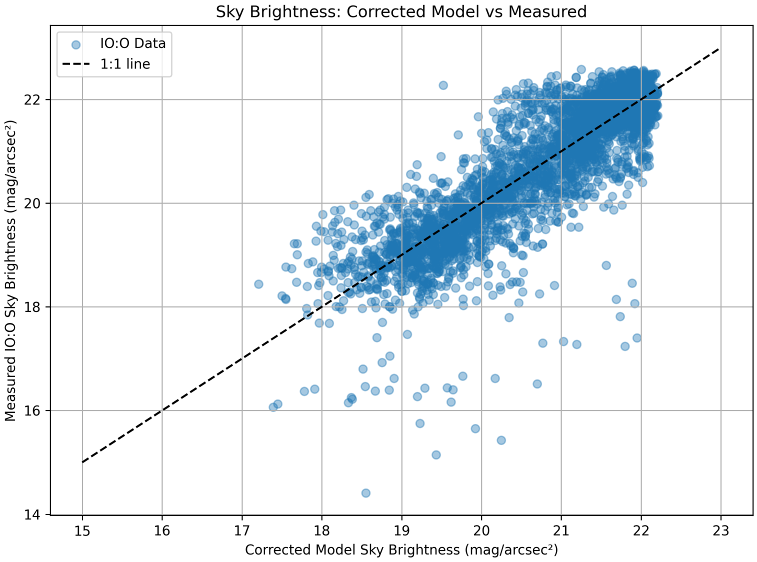

FIGURE 5.11: Comparison of Krisciunas & Schaefer model-predicted sky brightness with IO:O measurements.

5.4.7 Cross-Instrument Validation with Skycam Z and T

To assess the transferability and robustness of the sky brightness model beyond IO:O, a cross-instrument validation was performed using data from the Skycam Z and Skycam T imagers. Sky brightness estimates from both Skycam Z and Skycam T were compared to three reference models:

1. the measured sky brightness from IO:O standard star observations,

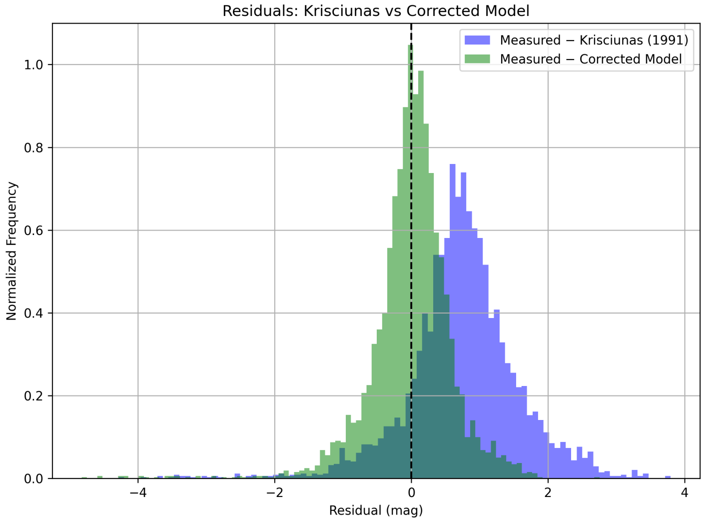

Figure 5.12: Histogram of residuals between measured sky brightness and model predictions. The corrected model (green) shows a narrower, symmetric distribution centred near zero, while the original model (blue) shows a broad, biased offset.

2. the Krisciunas & Schaefer (1991) model ,
3. and the corrected empirical model developed in this work .

The statistical comparison is summarised in Table 5.3, using root mean square error (RMSE), mean absolute deviation (MAD), and the coefficient of determination (R^2) as metrics.

Table 5.3: Statistical comparison of Skycam Z and T against IO:O and model predictions.

Instrument	**Reference**	**RMSE (mag)**	**MAD (mag)**	R^2
IO:O	Krisciunas	1.27	1.09	−0.02
IO:O	Corrected	0.61	0.41	0.76
Skycam Z	IO:O	0.54	0.37	0.78
Skycam Z	Krisciunas	1.02	0.82	0.23
Skycam Z	Corrected	0.77	0.55	0.55
Skycam T	IO:O	1.63	0.76	0.35
Skycam T	Krisciunas	1.94	1.15	0.07
Skycam T	Corrected	1.70	0.84	0.29

Skycam Z shows good agreement with IO:O observations, achieving an RMSE of 0.54 mag and an R^2 of 0.78 when compared directly to calibrated IO:O measurements. It also demonstrates improved performance when using the corrected model (RMSE 0.77 mag) over the original Krisciunas formulation (RMSE 1.02 mag), though residual trends remain due to differences in field of view, exposure settings, and calibration precision.

Skycam T performs noticeably worse, with an RMSE of 1.63 mag against IO:O and only modest improvements when using the corrected model. Its wider field of view and systematic calibration issues likely limit its utility for precise sky brightness estimation. Nonetheless, Skycam T still captures broad trends and can serve as a real-time monitor for large-scale brightness variations.

Figure 5.13 illustrates the improved agreement of Skycam Z with the corrected model and its relatively close match to IO:O-calibrated brightness levels.

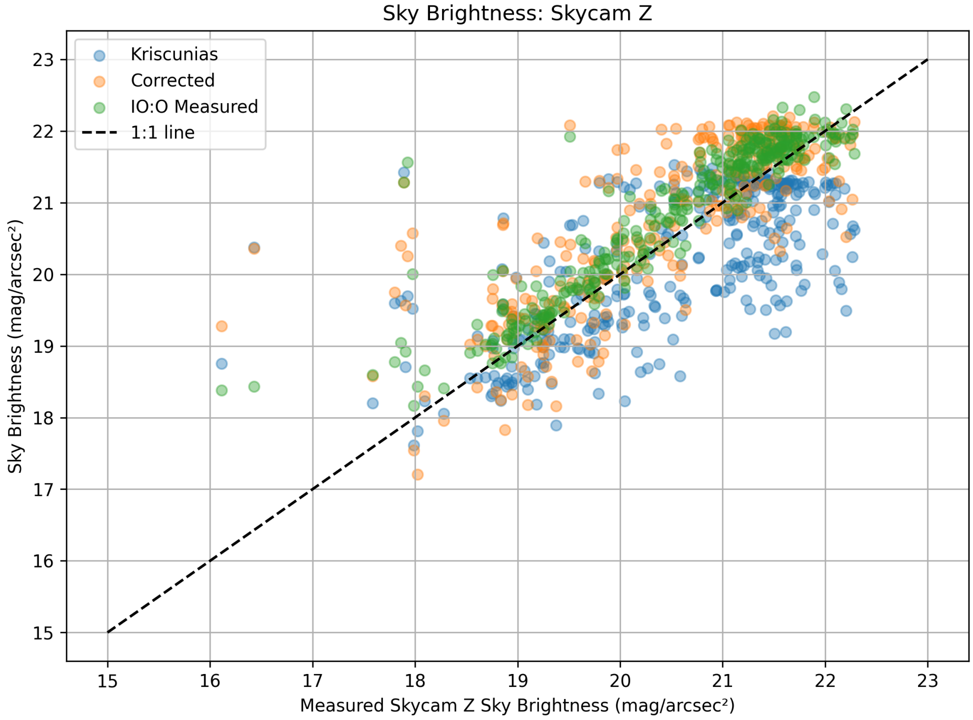

Figure 5.13: Skycam Z comparison against IO:O measured sky brightness and model predictions.

Figure 5.14 illustrates the agreement of Skycam T with the corrected model and its comparison to IO:O-calibrated brightness levels.

These findings support the use of Skycam Z as a secondary validation tool and a potential input source for sky brightness trend analysis. While it cannot fully replace calibrated photometric measurements, it provides an independent check on model behaviour and can be used to detect systematic offsets or seasonal variations in sky conditions.

5.5 Integration into Robotic Scheduling Systems

Sky brightness plays a critical role in signal-to-noise calculations for imaging and spectroscopic observations. Elevated background levels reduce contrast, necessitating longer

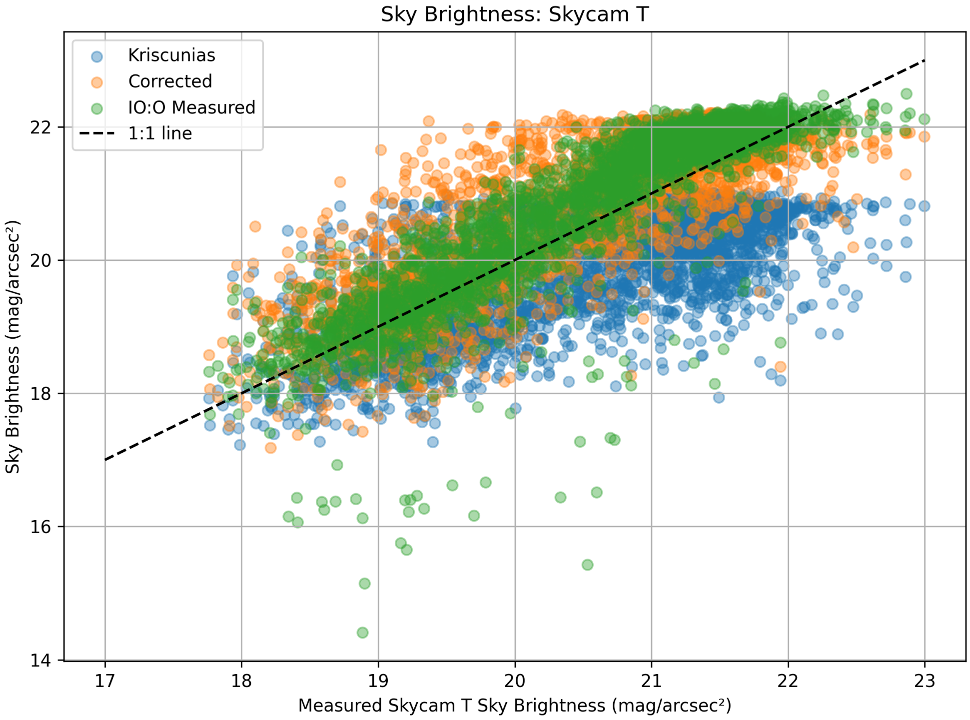

Figure 5.14: Skycam T comparison against IO:O measured sky brightness and model predictions.

exposures to achieve the same Signal-to-Noise Ratio, particularly for faint sources. In robotic observatories where schedule efficiency is vital, the ability to anticipate and adapt to changing sky brightness conditions provides a distinct operational advantage.

The corrected sky brightness model developed in this work enables both predictive and reactive scheduling strategies. For example, in offline mode, the model can be used to evaluate observing proposals during semester planning. Targets that require dark sky conditions can be flagged and corrected if scheduled too close to full moon or at low angular separation from the lunar position.

In real-time operation, the model can be queried dynamically to inform the scheduler of current or predicted background levels. Although sky brightness typically evolves gradually over timescales of tens of minutes (primarily driven by solar elevation and lunar geometry) incorporating this information allows exposure times to be adjusted continuously rather than relying on fixed worst-case assumptions. This enables more accurate SNR matching throughout the night, particularly during twilight transitions or periods of changing lunar altitude. Rather than triggering abrupt rescheduling, the primary benefit lies in improved efficiency and more consistent data quality through adaptive exposure scaling. This functionality integrates naturally with the SNR-based scheduling system described in Chapter 6 for both the Liverpool Telescope and the future New Robotic Telescope.

5.5.1 Worked Example

Considering a 20 magnitude object observed in the Bessell V band with a required SNR of 10. Under dark sky conditions (21.9 mag/arcsec^2), the exposure time needed is approximately 300 seconds. However, on a night with 80% moon illumination and a 40° Moon-target separation, the corrected model predicts a background brightness of 19.8 mag/arcsec^2. In this case, the same SNR would require an exposure time of 640 seconds, more than double. The scheduler, using this model, may choose to delay the observation to a darker window or reallocate the time to a brighter, less background-sensitive target.

Such decisions improve scientific return, reduce wasted observing time, and allow for more robust automated scheduling. Integration into existing LT scheduling logic would only require minor extensions to the exposure time estimation and target scoring modules, as demonstrated in Section 7.7.1.

5.5.2 Replacing the Heuristic Model of the LT with a Continuous Estimator

While the existing LT sky brightness model has proven useful for broad classification of observing conditions, it operates on discrete brightness bins and applies fixed thresholds that do not fully capture the geometric and photometric complexity of night sky brightness. As robotic scheduling moves toward continuous quality-aware metrics such as signal-to-noise ratio (SNR), a more precise and dynamically evaluable model is needed.

To that end, an improved empirical sky brightness model was developed in this chapter and validated against calibrated IO:O observations in the Bessell-V filter. The model predicts sky brightness in mag/arcsec^2 based on real-time lunar geometry (see 5.4.3).

This formulation can be evaluated rapidly at scheduling time for any target and timestamp, requiring only lunar ephemerides and target coordinates. The resulting sky brightness prediction can be used in two distinct ways:

1. As input to the SNR model for dynamic exposure time calculation (Chapter 6),
2. As a continuous constraint against which to validate observer-specified brightness limits, replacing the previous threshold table.

This integration enables fine-grained scheduling decisions and improves the fidelity of predicted SNR values under varying lunar conditions.

Unlike the previous model, which operated entirely in the Sloan r' band, the corrected model is band-independent and can be fitted per filter. In practice, a version fitted for Bessell-V was used to support comparison with the Skycam sytems and to align with the SNR-based exposure time system.

This real-time estimator is architecturally compatible with both the Liverpool Telescope (LT) and the New Robotic Telescope (NRT) scheduling pipelines, and has been developed with the intention of serving as a continuous alternative to the existing heuristic sky-brightness model. While it has not yet formally replaced the operational LT model, its lightweight, geometry-driven formulation allows straightforward integration into both Java- and Python-based scheduling systems. Ongoing post-thesis development work has focused on enabling such integration within modernised scheduling infrastructure, ensuring that the estimator can be evaluated thousands of times per night without measurable performance overhead.

Chapter 6

Signal-to-Noise Ratio Based Scheduling for Robotic Telescopes

> *Prediction is very difficult, especially about the future.*
>
> Niels Bohr

6.1 Introduction

A central challenge in observational astronomy is the optimisation of limited telescope time while ensuring that observations meet scientific requirements. For the New Robotic Telescope (NRT), the primary science drivers are large-scale spectroscopic follow-up of transients discovered by facilities such as LSST, and long-term monitoring of variable sources requiring a regular (nightly) cadence. Both modes necessitate hundreds of observations per night with minimal human intervention. In this regime, efficiency is measured not only in raw observing time but also in the ability to deliver spectra and light curves of consistent scientific quality across a wide range of conditions.

A key metric for data utility is the Signal-to-Noise Ratio (SNR), which quantifies the clarity and usability of an observation. Current practice on the Liverpool Telescope (LT) and similar facilities relies on fixed exposure times specified by users, often derived from exposure time calculators and padded with conservative safety margins to compensate for uncertain conditions. These exposure times are frequently rounded to convenient values (e.g. 90 s instead of 84.3 s), introducing a systematic bias towards overexposure. This reduces scheduling flexibility and consumes telescope time without proportionate gains in scientific utility (Figure 6.1).

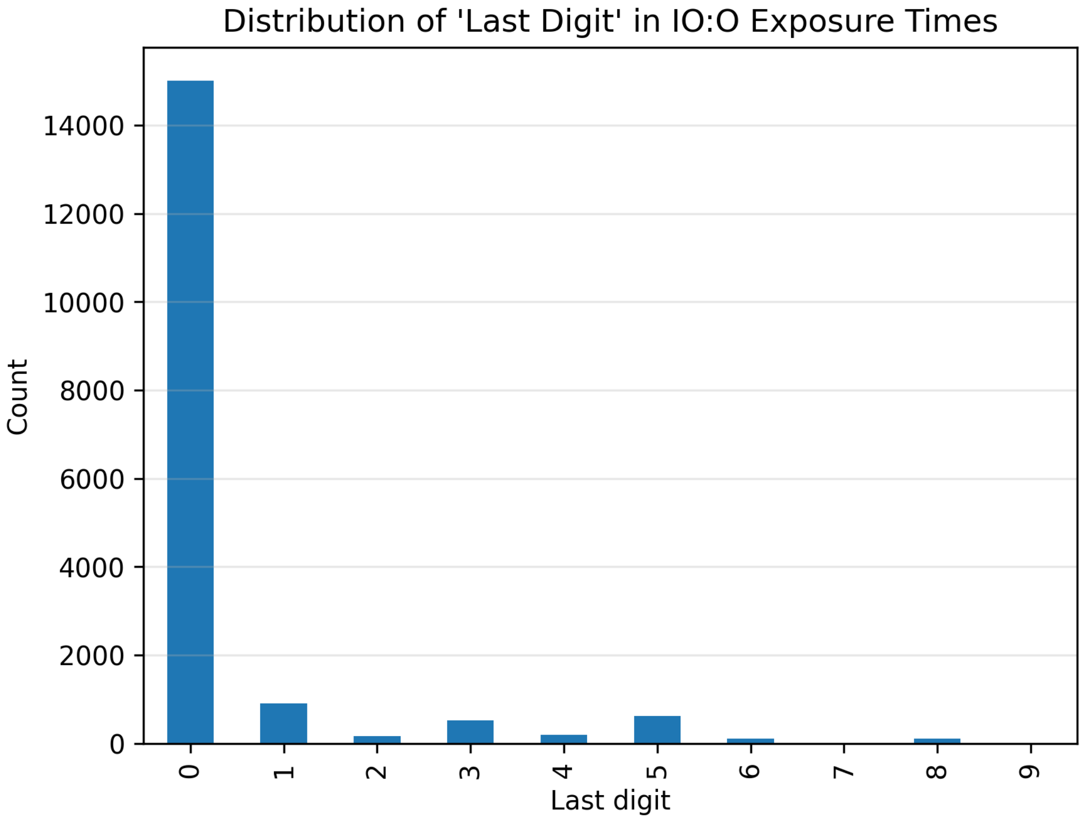

Figure 6.1: "Last digit" distribution of IO:O science exposures from January to July 2025. Exposures ending in “0” comprise 85.1% of observations and 98.1% of total time, reflecting a strong bias toward rounded exposure times.

This chapter introduces a dynamic scheduling approach based on target SNR rather than fixed exposure times. Users specify coordinates, approximate magnitude, and desired SNR; the scheduler then determines the exposure time in real time using measured conditions such as airmass, sky brightness, seeing, and atmospheric transparency. This removes the need for conservative buffers, ensures that observations are matched to prevailing conditions, and enables the ingestion of large target lists without bespoke human planning.

The approach integrates the atmospheric characterisation methods developed in previous chapters: cloud detection (Chapter 3), extinction (Chapter 4), and sky brightness modelling (Chapter 5); into a holistic, condition-aware scheduling framework. A proof-of-concept has been demonstrated on the Liverpool Telescope, and the system is designed for deployment within the NRT scheduling architecture.

6.2 SNR-Based Exposure Time Calculation

To dynamically determine the required exposure time for a given observation, a model was implemented to both calculate and invert the signal-to-noise ratio (SNR), using instrumental parameters specific to the IO:O imager on the Liverpool Telescope. The calculation incorporates real-time observing conditions including sky brightness, airmass, and seeing, as well as fixed instrumental constants such as read noise, dark current, and filter-dependent zero points.

6.2.1 Input Parameters

The exposure time calculation requires the following inputs:

- Target coordinates (RA, Dec)
- Apparent magnitude in the relevant filter
- Desired SNR
- Real-time sky conditions: airmass, sky brightness, seeing
- Instrumental parameters: zero point (ZP), read noise (R), dark current (D)

These inputs are combined to compute the photon flux from the target and the expected noise contributions from the background, detector, and readout processes.

6.2.2 SNR Equation

The SNR for a point source is computed as:

$$\mathrm{SNR} = \frac{S \cdot t}{\sqrt{S \cdot t + B_{\mathrm{sky}} \cdot t \cdot n_{\mathrm{pix}} + \left(\frac{R}{g}\right)^2 \cdot n_{\mathrm{pix}} + \frac{D \cdot t \cdot n_{\mathrm{pix}}}{g}}} \tag{6.1}$$

where:

- S is the source signal (in counts per second)
- t is the exposure time
- B_{sky} is the sky background (in counts per pixel per second)

- n_{pix} is the number of pixels encompassing the source
- R is the read noise (in electrons)
- D is the dark current (in electrons per pixel per second)
- g is the detector gain (in electrons per ADU)

For the IO:O camera, the dark current is negligible ($D \approx 0$) under normal operating conditions and is omitted from further calculations.

6.2.3 Exposure Time Inversion

To determine the exposure time required to achieve a given SNR, Equation **??** is algebraically inverted, yielding a quadratic form:

$$A = S^2, \quad B = -\text{SNR}^2 \cdot (S + B_{\text{sky}} \cdot n_{\text{pix}}), \quad C = -\text{SNR}^2 \cdot n_{\text{pix}} \cdot R^2 \tag{6.2}$$

Solving for t using the quadratic formula:

$$t = \frac{-B + \sqrt{B^2 - 4AC}}{2A} \tag{6.3}$$

This formulation allows the exposure time to be determined dynamically for each observation, based on both target properties and prevailing observing conditions. Values for S and B_{sky} are estimated from the apparent magnitude, instrumental zero point, and either the measured sky brightness or the output of the corrected model specified in Chapter 5 in magnitudes per square arcsecond. The area of the seeing disk is estimated using a circular aperture with a diameter equal to the full width at half maximum (FWHM) of the seeing profile.

6.3 Exposure Time Rounding Bias in Practice

An empirical analysis of exposure times on the Liverpool Telescope reveals that manual specification introduces a strong and previously unquantified rounding bias. Rather than entering exact values, users overwhelmingly select convenient durations such as 60, 90, or 120 seconds (see Figure 6.1). This behaviour is not surprising given the constraints of the Phase II interface, but it introduces a systematic bias towards longer exposures and therefore reduces overall scheduling efficiency.

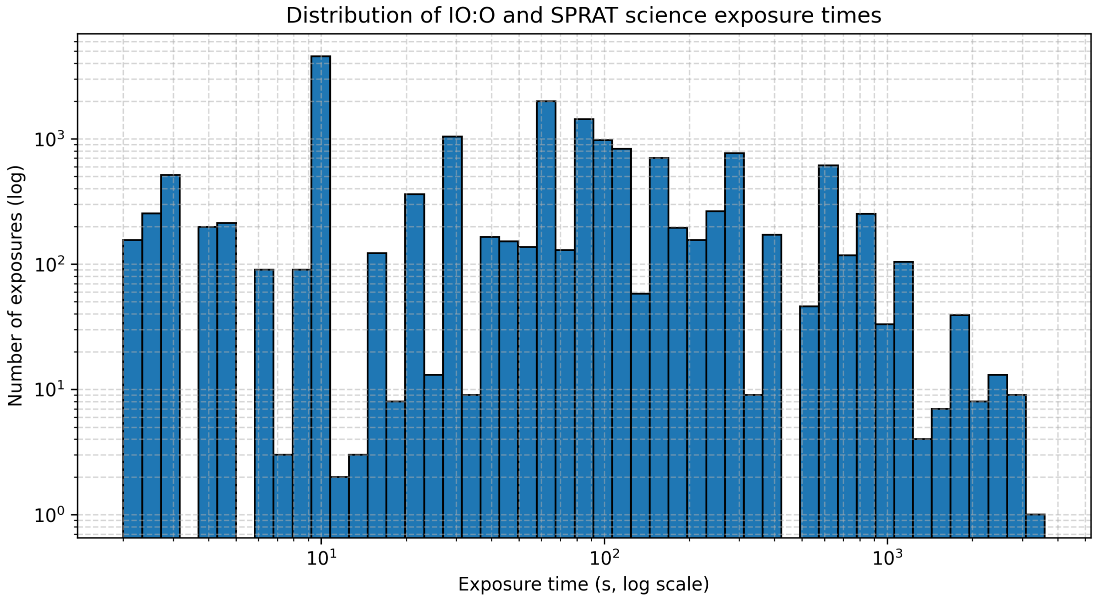

FIGURE 6.2: Histogram of all IO:O and SPRAT science observation times between January and July 2025.

To quantify the effect, all IO:O and SPRAT science observations between January and July 2025 were analysed. Exposure times were grouped by their last digit:

Category	**N (Exposures)**	**Total Time (s)**	**Fraction (%)**
Ends in 0	15,016	2,082,710	98.07
Ends in 5	620	33,430	1.57
Other	2,010	7,661	0.36

TABLE 6.1: Exposure time rounding statistics for IO:O and SPRAT (January–July 2025).

The results show that 85% of exposures ended in a zero and a further 3.5% ended in a five, meaning that 88.5% of individual exposures used rounded durations. However, because these rounded exposures correspond primarily to longer integrations, they account for more than 98% of the total exposure time executed. The remaining 11.4% of exposures used non-rounded durations, but these are typically short observations (e.g. standard stars) and together represent only 0.36% of the total integrated time. Assuming a conservative 5% excess per rounded observation (e.g. 90 s instead of the required 84.3 s), the cumulative overhead for the dataset amounts to ∼29.4 hours, equivalent to nearly four full nights of telescope time 'lost' in only six months.

The assumed 5% excess per rounded observation represents a conservative estimate of the typical rounding offset relative to the SNR-derived exposure time. A small validation sample, including proposals with explicitly stated SNR goals, suggests that the deviation between rounded and model-derived durations can in some cases exceed this value. However, a comprehensive reassessment of proposal-level exposure adequacy lies beyond the scope of this study. The 5% assumption should therefore be interpreted as a representative first-order estimate for quantifying rounding-induced overhead, rather than as an exact universal correction factor.

This finding highlights a structural limitation of manual Phase II planning: exposure time entry is coarse-grained, convenience-driven, and systematically inefficient. For current instruments this loss is already non-negligible; for the NRT, which must handle hundreds of transient spectra and large monitoring campaigns every night, such inefficiency would be unsustainable. A condition-aware, SNR-driven system eliminates the rounding bias entirely, allocating exposure time precisely to the level required by prevailing conditions and scientific goals. It also elliminates the "human bottleneck" in target specification, enabling the processing of real-time alerts without human intervention.

6.4 Validation of the SNR-Based Scheduling Model

The effectiveness of the SNR-based exposure time calculation model was assessed through a series of validation experiments conducted with the Liverpool Telescope (LT). Although the system was not yet integrated into the operational scheduler, targeted follow-up observations were triggered via Robotic Telescope Markup Language (RTML) requests submitted under a high-priority proposal (called `SchedEfficiency`, data available in the LT archive[1]). This approach enabled empirical testing of the model's predictions under real observing conditions while retaining control over target selection and experimental design.

Three complementary experiments were designed to probe different aspects of model performance. First, controlled tests with a well-characterised photometric standard star provided a baseline under stable conditions. Second, a simulated comparison against fixed exposure strategies defined in an active proposal evaluated how the model performs relative to conventional observing practice. Finally, a real-time demonstration linked transient alerts from the Lasair broker to dynamically calculated exposure times for spectroscopic follow-up with SPRAT, testing the system in a realistic, time-critical scenario.

Together, these experiments span the range from controlled calibration tests to fully automated transient response, providing a comprehensive assessment of the model's robustness. Beyond validating the method on the LT, these tests also illustrate the capability

---

[1] `https://telescope.ljmu.ac.uk/cgi-bin/lt_search`

required by the New Robotic Telescope (NRT): to ingest large volumes of alerts and monitoring targets, and to autonomously decide on exposure times without human intervention. The results presented here therefore serve both as a proof of concept and as a demonstration of feasibility for the fully autonomous scheduling architecture envisioned for the NRT.

6.4.1 Empirical Tests with Standard Star Re-observations

To test the system under controlled conditions, a dedicated experiment was carried out using the photometric standard star SA 94-702 (GSC 00048-00918). This target is part of the Landolt catalogue and provides a well-characterised, non-variable reference source. It was selected for repeated observations in two filters (SDSS-R and SDSS-I) to verify whether the SNR-based exposure time calculations reliably meet the assumed target of SNR = 100.

6.4.1.1 Methodology

The procedure for each test began with monitoring the publicly accessible observation log of the Liverpool Telescope during the night. As soon as an entry corresponding to a standard star was recorded (specifically SA 94-702 in this experiment) the system parsed the log to extract relevant metadata such as the timestamp, filter, and exposure time. Since the actual FITS image was not yet available at this stage, the apparent magnitude of the target in the corresponding filter was retrieved from the SIMBAD Web-API (as a proof-of-concept for getting live magnitude data from future surveys such as the VRO alert streams) and used as the input to compute the exposure time required to achieve an SNR of 100. This calculation followed the formalism outlined in Section 6.2, incorporating instrumental parameters and a nominal seeing value. The seeing was read from previous entries in the observation log and averaged over the latest ten entries. The sky brightness was estimated using the model proposed in Chapter 5. To ensure hardware safety and scheduling feasibility, a minimum exposure time of 1 s was enforced, even if the computed value was shorter.

Once the new exposure time was determined, a Robotic Telescope Markup Language (RTML) request in a high-priority proposal called `SchedEfficiency` was issued, specifying the target coordinates, filter, and desired exposure time. These re-observations were typically executed within 30 to 120 minutes of the original observation, ensuring that sky conditions such as airmass and transparency remained as consistent as possible.

A shorter response time would be desirable, but is limited by the constraints of the existing LT scheduling system, which does not natively support immediate re-observation triggers.

After both the original and re-observed FITS images were available after the data reduction, typical on the next day, they were processed offline for analysis. The standard star was located in each image using its known coordinates, and the source closest to this position was selected. Source extraction was performed using `SExtractor` (Bertin et al., 1996), which yielded photometric parameters including the isophotal flux, local background level, and source area. These parameters were then used to compute the achieved SNR using Equation **??**, enabling a direct comparison with the predicted values derived from the model.

6.4.1.2 Results and Discussion

A total of 14 re-observations were performed over the nights of January 6th and 7th, 2025. Exposure times for SA 94-702 ranged around 0.01 s to 0.03 s and thus were clipped to 1s. For each case, the actual SNR achieved was compared to:

1. The target value of SNR = 100 (based on the SIMBAD magnitude)
2. The predicted scaling from the original observation, assuming linear SNR growth with $\sqrt{t}$

While some deviations are present, especially in the low-exposure regime limited by the 1 s floor, the system successfully predicts appropriate exposure times within the expected tolerance for atmospheric and instrumental variability.

6.4.1.3 Conclusion

This experiment demonstrates the feasibility of the SNR-based exposure model in a realistic robotic telescope environment. Even without full integration into the scheduling queue, re-observation of standard stars with computed exposure times showed consistency with the theoretical predictions.

6.4.2 Evaluation Against Proposal Exposure Strategies

One of the motivations behind the signal-to-noise ratio (SNR) driven scheduling approach is to replace conservative, fixed exposure time estimates with condition-aware

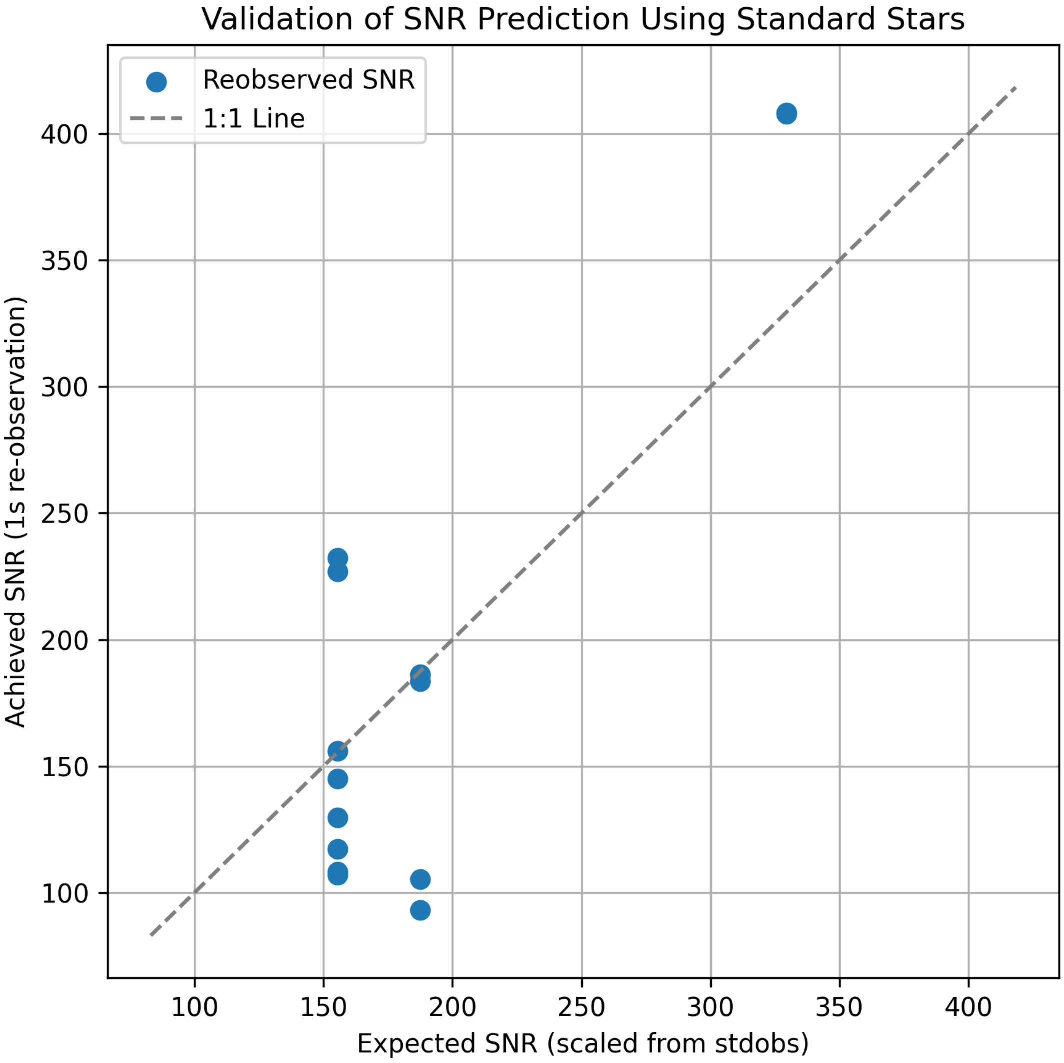

Figure 6.3: Achieved versus predicted SNR for re-observations of the standard star SA 94-702. The dashed line indicates the ideal 1:1 relationship

predictions that are optimised per observation. Many observing proposals continue to rely on static exposure values, derived under simplified assumptions (median seeing, dark sky, low airmass) and often padded with generous safety margins. This introduces a mismatch between requested time and actual sky conditions, a gap the SNR-based system is designed to close.

By dynamically adjusting exposure times, the system avoids the systematic overexposure that arises from conservative user estimates, freeing telescope time without compromising scientific quality. More importantly, it enables a shift in how proposals are formulated. Rather than requiring principal investigators to manually estimate detailed exposure times for each target, large-scale programmes can specify the desired data quality (e.g. SNR = 20 in a given band), while the scheduler determines the required integration time based on instrument sensitivity and prevailing conditions. This does not eliminate time accounting or allocation constraints: at the proposal stage, expected integration times are still estimated using representative conditions to ensure feasibility and fair resource distribution. During execution, the scheduler adjusts exposure durations dynamically within the allocated time budget, compensating for changing sky brightness or transparency. In this framework, SNR-based specification refines exposure control without undermining the principle that each accepted programme operates within a defined and finite share of telescope time. This paradigm is particularly important for the New Robotic Telescope (NRT) and its SPEC programme, which is designed to classify hundreds to thousands of transient alerts per night. In such a regime, manual specification of exposure times is neither practical nor efficient; the scheduler must autonomously compute and assign optimal exposures in real time. The simulations below demonstrate how this approach performs when applied to a concrete example from an LT proposal, illustrating both the inefficiencies of traditional fixed strategies and the advantages of condition-aware scheduling.

6.4.2.1 Methodology

To test the model against real observing practice, a simulation study was performed based on Proposal JL25B10 (PI: K. Ryan Hinds), which targets Type IIP supernovae in the radioactive tail phase. The proposal specifies fixed exposure blocks designed to reach SNR $\approx$ 15 under dark-time, $2''$ seeing conditions. Using the model from Section 6.2, the expected performance of these blocks was simulated under a range of conditions (dark vs. bright sky, varying airmass). Instrumental parameters for IO:O (gain $= 1.6\,\mathrm{e}^-/\mathrm{ADU}$, read noise $= 16\,\mathrm{e}^-$, pixel scale $= 0.3$ arcsec/px, zeropoints from Smith et al., 2017) were adopted.

6.4.2.2 Results and Discussion

The simulations show that under favourable conditions the proposal exposures are substantially conservative, often delivering significantly higher SNR than requested. In contrast, under bright-sky or high-airmass conditions, especially in the g-band, the fixed blocks can underperform, failing to reach SNR 15.

To further illustrate this, the model was run across a full night for a virtual $g = 20\,\mathrm{mag}$ target. Figure 6.4 shows the achieved SNR for a fixed 180 s exposure compared to the dynamically required exposure time, as a function of time, airmass, and sky brightness. The plots highlight both excessive overexposure near transit under dark conditions and underexposure at high airmass under bright conditions.

Filter	Mag	Sky (mag/arcsec^2)	Proposal exp. (s)	Proposal SNR	Req. exp. (s)
g	20	22.0	180	30.8	70
g	20	20.5	180	19.9	113
r	19	21.5	90	47.1	21
r	19	20.0	90	31.0	29

Table 6.2: Comparison between proposal-defined exposure blocks and SNR-based requirements for tail-phase SN magnitudes. Simulations assume $g = 20$ mag and $r = 19$ mag, an aperture radius of 6 px (1.8 arcsec), and sky brightness values representative of dark and bright conditions.

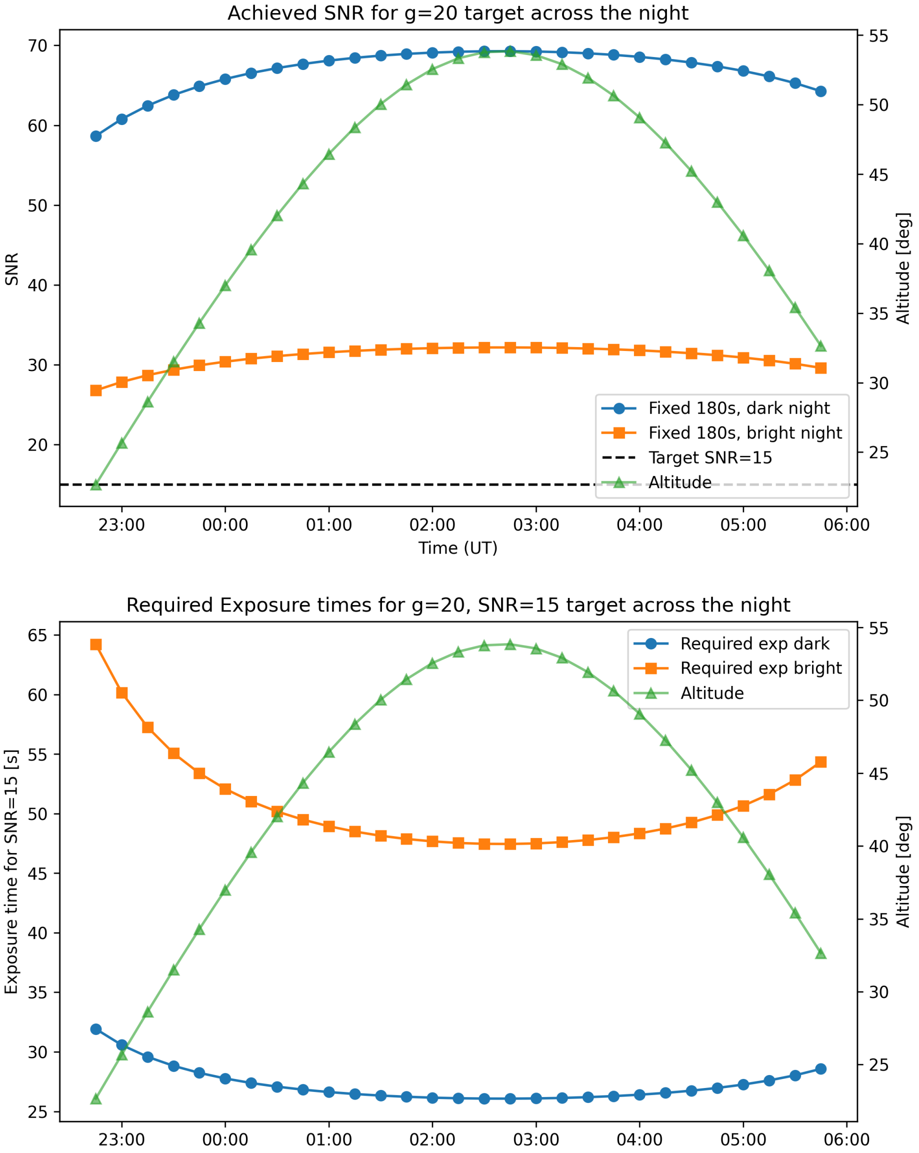

Figure 6.4: Simulated performance of a fixed 180 s exposure in the g-band for a $g = 20$ mag target across one night at ORM. Top: achieved SNR under dark and bright conditions; the dashed line marks the target SNR=15. Bottom: required exposure times for SNR=15 under the same conditions. Both plots with target altitude overplotted (green).

6.4.2.3 Conclusion

This experiment demonstrates that the SNR-based model adapts in both directions: reducing excessive overexposure under optimal conditions and extending exposure times when conditions degrade. Even with a limited empirical dataset, the consistency between simulations and re-observed data supports the model's predictive power. In practice, such adaptability would prevent the systematic inefficiencies inherent in fixed proposal blocks, enabling the scheduler to optimise resource use while maintaining the requested SNR.

6.4.3 Real-Time Transient Follow-Up with Pre-Classified Alerts

The third experiment evaluated the system in the most realistic regime: live transient alerts and rapid spectroscopic response, directly reflecting the requirements of the NRT SPEC programme. This experiment linked the Lasair transient alert stream to the exposure time model and robotic triggering system for the Liverpool Telescope's spectrograph SPRAT (Spectrograph for the Rapid Acquisition of Transients; Piascik et al., 2014).

6.4.3.1 Alert Stream and Classification

Lasair provides real-time alert packets for transient candidates detected by the Zwicky Transient Facility (ZTF) and will also provide similar services for the VRO/LSST, including preliminary classifications based on photometric behaviour. For this test, the alert listener ingested alerts using a user-defined filter, filtering the g band magnitude in the range $16 < g < 19$. Each alert was enriched via the ALeRCE[2] API, which returns a full set of classification probabilities from multiple machine-learning classifiers. Both the complete classification output and the top-3 ranked classes were stored in the database (the latter as JSON for rapid inspection; see Table 6.3).

Field	**Description**
`Name`	Alert object identifier (e.g. ZTF17aaawqgi)
`RA`, `Dec`	J2000 coordinates in degrees
`MagnitudeG`, `MagnitudeR`	Most recent g/r-band photometry
`Classification`	Lasair classification string
`AlerceClassification`	ALeRCE top-3 ranked classes (JSON)
`LatestDetection`	Timestamp of latest detection (r or g mag)
`FirstDetection`	Timestamp of first detection

TABLE 6.3: Core schema for transient alerts ingested into the follow-up system.

---

[2] `https://alerce.science`

6.4.3.2 Science Case Matching

A simple matching rule was applied to select only alerts corresponding to supernovae ("SN") in the ALeRCE classification output. If the highest probability from all classificators suggested a supernova and if the probability was above 0.75, the alert was considered a candidate for follow-up. This approach allowed the system to focus on high-priority transients while avoiding false positives from other classes such as asteroids. In the process, mis-classifications (e.g. Long Periodic Variables) occurred, underscoring the need for a more sophisticated decision logic of which new alerts to follow up but which lies beyond the scope of this thesis.

6.4.3.3 Exposure Time Calculation and Scheduling

For each selected alert, the most recent magnitude was passed into the SNR-based exposure time calculator described in Section 6.2, with the target instrument set to SPRAT and the desired signal-to-noise ratio fixed at SNR = 25 per spectral resolution element at 5500 Å. The calculation assumes a flat continuum spectral energy distribution across the SPRAT wavelength range. The modified sky brightness model developed in Chapter 5 was incorporated using the current Moon phase, altitude, and target airmass to refine the sky background estimate.

The resulting exposure time was then used to construct a RTML observation request for SPRAT, submitted to the LT robotic scheduler with high priority. This closed-loop from alert to observation was fully automated, with the only human interaction being the activation of the alert listener.

6.4.3.4 Operational Demonstration

The system successfully triggered test observations of multiple ZTF supernova candidates, including the example shown in Figure 6.5. In each case, the dynamically calculated exposure time matched well with post-observation SNR measurements, confirming that the model could be applied in rapid-response scenarios without conservative overexposure.

This experiment provides a proof-of-concept for autonomous large-scale transient spectroscopy at the NRT. It showed that the SNR-based scheduling approach can extend beyond planned observations to on-the-fly transient follow-up, directly linking machine-learned classifications in alert streams to optimised exposure times for rapid spectroscopy.

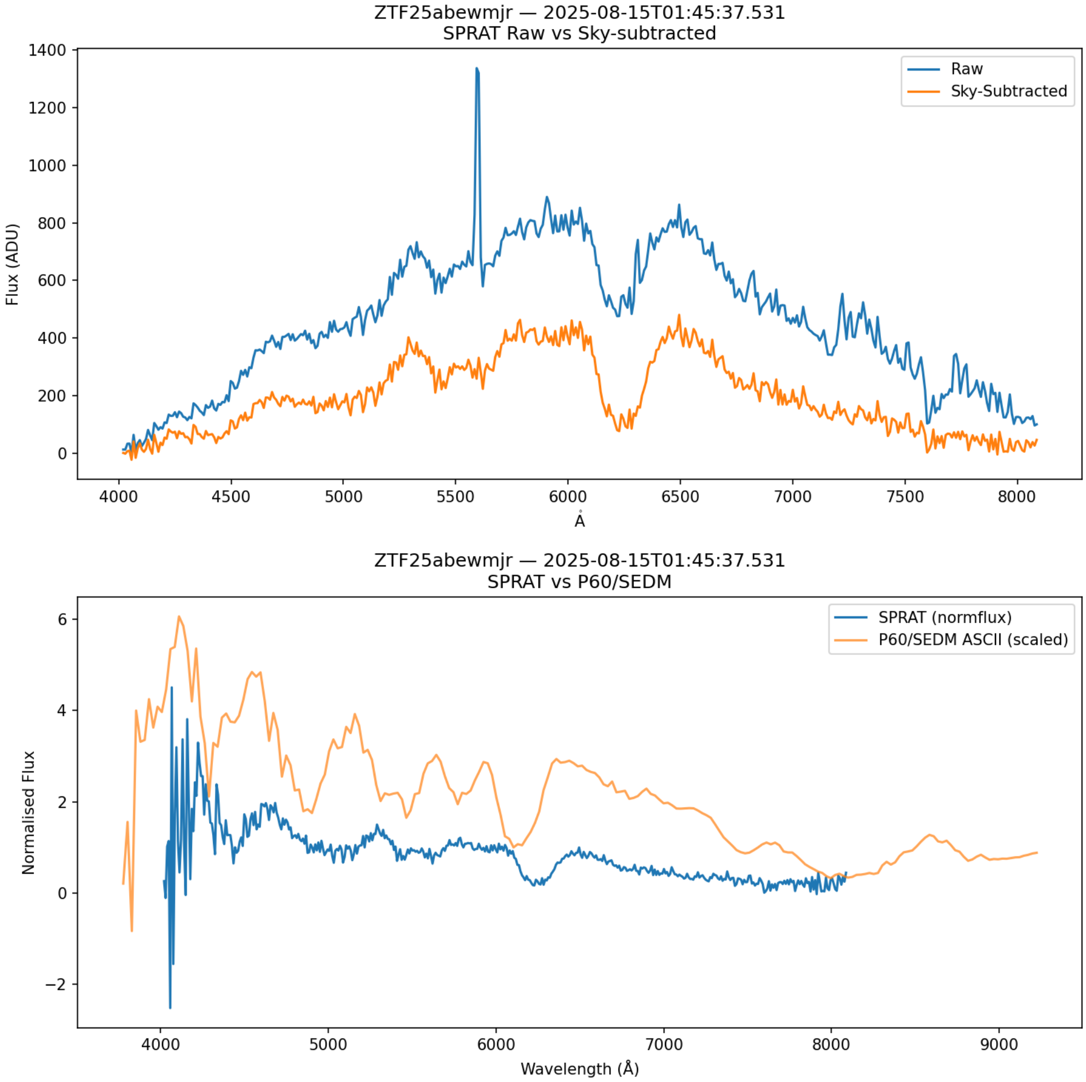

Figure 6.5: Comparison of spectra of ZTF25abewmjr obtained with two different instruments. The blue line shows the SPRAT spectrum from the Liverpool Telescope, plotted in normalised flux (extension 4 of the pipeline output). The orange line shows the publicly available P60/SEDM spectrum from the Transient Name Server (TNS), scaled for visual comparison. The overall agreement in broad spectral features and continuum shape demonstrates the feasibility of the automated pipeline from alert ingestion and magnitude query (via ALeRCE in this case) to autonomous scheduling of an exposure calculated for a specified SNR (20). No additional cross-instrument wavelength alignment or rest-frame correction was applied for this comparison; small-to-moderate discrepancies in feature centroids likely reflect differences in wavelength calibration, resolution, and reduction pipelines between the two instruments. The comparison is therefore qualitative rather than a precision spectroscopic validation.

6.4.4 Model Performance towards Autonomous Scheduling

The validation experiments collectively demonstrate that the signal-to-noise ratio based scheduling system performs robustly across a range of use cases, from controlled calibration to realistic transient response. Each test highlighted complementary strengths

Alert Name	Mag_{g}	Observed (UTC)	Exp. t (s)	Airmass	Altitude (deg)
ZTF25abewmjr	16.69	2025-08-15 01:45:37	207.0	1.007	83.56
ZTF19aatihef	17.63	2025-08-12 22:05:30	274.0	1.546	40.44
ZTF25abewmjr	16.69	2025-08-15 01:49:13	207.0	1.007	83.56
ZTF17aaarisn	17.61	2025-08-12 05:35:38	279.0	1.604	38.75

Table 6.4: Set of transient alerts ingested from Lasair, enriched via ALeRCE classification, and observed with SPRAT using dynamically calculated exposure times for SNR=20.

of the approach: standard star re-observations confirmed the internal consistency of the exposure time inversion; proposal-based simulations showed clear efficiency gains compared to static exposure blocks; and the transient follow-up pipeline illustrated seamless integration of real-time alert streams into condition-aware scheduling. Together, these results confirm both the feasibility of the method on the Liverpool Telescope and its readiness for large-scale deployment on the New Robotic Telescope.

Several practical considerations emerged during testing. A minimum exposure time of one second necessarily limited validation for bright targets, but this is an instrumental safeguard rather than a limitation of the model itself. Catalogue mismatches and intrinsic variability of some targets introduced scatter, yet the predicted and achieved SNRs remained consistent within expected tolerances. Most importantly, the system repeatedly demonstrated that it can adapt to varying airmass, transparency, and sky brightness in real time, assigning exposures that meet the requested SNR without conservative safety margins or convenience-driven rounding.

The broader significance of these results lies in their ability to address the central bottleneck of robotic astronomy in the time-domain era. Existing Phase II paradigms require human specification of exposure times and therefore cannot scale to the alert rates expected from surveys such as the LSST at the Vera C. Rubin Observatory. Moreover, they assume the user can anticipate observing conditions in advance, including the altitude of the target at the moment of execution. In practice, this is rarely the case: a target may be observed at higher airmass if other, higher-ranking proposals consume time around transit, or under brighter skies than assumed. Fixed exposure times therefore either underperform or require conservative padding. The SNR-based framework removes both problems by allowing users to request data quality directly, leaving the scheduler to determine the appropriate exposure in situ. By calculating integration times for the actual altitude, sky brightness, and extinction at the scheduled instant, the system eliminates inefficiencies caused by safety margins and convenience rounding, while simultaneously enabling ultra-fast transient classification at the scale demanded by the SPEC programme.

Fast transients such as the early stages of fast blue optical transients (FBOTs) or kilonovae do not wait for users to calculate exposure parameters, and a fully autonomous, condition-aware system is therefore essential if these fleeting events are to be captured at scientifically useful signal-to-noise.

The architecture underpinning this framework is summarised in Figure 6.6, which brings together the preceding chapters into a single workflow. Cloud detection (Chapter 3), extinction monitoring (Chapter 4), and sky brightness modelling (Chapter 5) feed directly into the SNR-based exposure calculation (this chapter), producing an integrated, condition-aware scheduler capable of fully autonomous operation. This simple flow encapsulates the trajectory of the thesis: the transformation of environmental characterisation into actionable, science-optimised scheduling decisions.

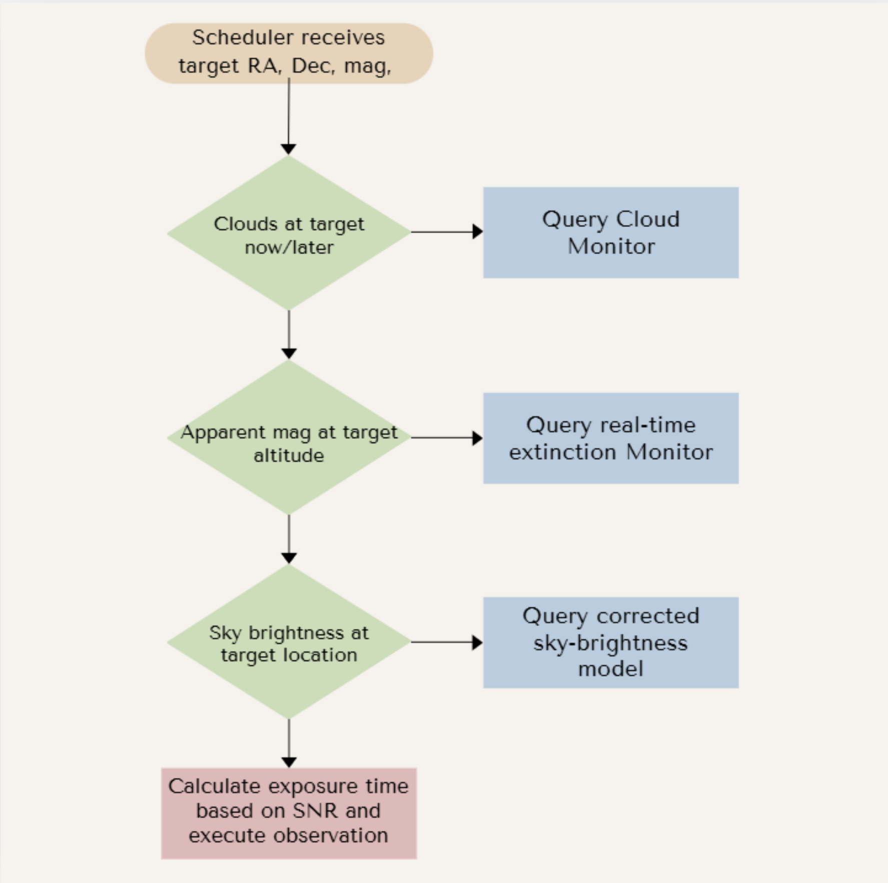

Figure 6.6: Workflow of the SNR-based scheduling framework. Environmental monitors developed in earlier chapters feed into the exposure time calculation, which dynamically assigns integration times to meet user-specified SNR requirements.

Chapter 7

Results and Discussions, Integrations

Science is not about certainty.
Science is about finding the most
reliable way of thinking at the present
level of knowledge.

Carlo Rovelli

7.1 Introduction

The preceding chapters presented the development of four core systems designed to enable condition-aware robotic scheduling: real-time cloud detection and forecasting, atmospheric extinction monitoring with Skycams, an empirical model of sky brightness, and a dynamic signal-to-noise ratio (SNR) based exposure time estimator. Each module was developed, tested, and validated independently. In this chapter, these results are evaluated together in the context of their scientific motivation and their implications for robotic telescope operations.

The central question addressed here is whether these systems, individually and in combination, advance the ability of robotic telescopes to respond effectively to the demands of modern time-domain astronomy. When the Liverpool Telescope was commissioned in 2004, the dominant science drivers were long-lived transients such as supernovae and classical novae, whose evolution could be followed over weeks. Under those conditions, simple altitude-based scheduling and fixed exposure times were sufficient. Today, however, the discovery space has shifted toward faster and fainter phenomena - from fast blue optical transients (FBOTs) to kilonovae and gamma-ray burst afterglows - which evolve on timescales of hours or even minutes. Responding to these events requires not only rapid telescope reaction, but also an observing strategy that adapts in real time to atmospheric transparency, sky brightness, and the onset of cloud cover.

This chapter therefore goes beyond reporting numerical performance. Each system is assessed in terms of what it enables scientifically: how well cloud forecasts align with the cadence of transient follow-up, how stable extinction monitoring is under operational conditions, how accurately the empirical brightness model captures lunar and atmospheric effects, and to what extent the SNR-based exposure estimator can replace manual observer judgement. The discussion highlights strengths and limitations, and considers how these results map onto the practical requirements of both the current Liverpool Telescope scheduler and the future New Robotic Telescope. In doing so, it aims to close the loop between the "science story", the motivation for condition-aware scheduling in the time-domain era, and the "engineering story" of how such a system can be implemented in practice.

7.2 Cloud Detection and Forecasting

Reliable knowledge of sky conditions on short timescales is a prerequisite for autonomous operation in time-domain astronomy. While traditional queue schedulers assume either stable conditions or rely on all-sky monitors to flag cloudy periods, they lack predictive power. For transient science, where a decision to observe or defer must be made within minutes, a forecasting capability is essential. The cloud detection and forecasting pipeline developed in this thesis provides such a capability by combining difference imaging with binary thresholding and extrapolating cloud motion using connected-component labelling and Kalman filters.

The system demonstrates robust performance under a wide range of observing conditions. For extended and slowly evolving cloud structures, forecasts of cloud arrival and departure times were accurate on horizons of up to 20 minutes (see Chapter 3). This level of accuracy is well matched to the operational cadence of the Liverpool Telescope, where the median observing group lasts less than half an hour. In this regime, the forecast enables informed "now/next" scheduling decisions, allowing the scheduler to defer or advance observations in anticipation of cloud passage rather than reacting only after conditions deteriorate.

Limitations become apparent when cloud structures evolve rapidly or when small, fast-moving clouds dominate. In such cases, forecast skill declines significantly beyond about

15 to 20 minutes, and the uncertainty in predicted cloud patterns increases. From an operational perspective, however, this is not a fundamental shortcoming: the scheduling problem at the ORM rarely requires cloud predictions on hour-long timescales, but rather accurate knowledge of conditions over the duration of the next group. In that respect, the achieved forecast horizon is both sufficient and scientifically meaningful.

Another strength of the system is its modest calibration requirements. Only an initial set-up is needed, and no site-specific retraining is required. This reduces operational overhead and increases the likelihood of successful long-term deployment. Taken together, these results show that cloud forecasting is feasible in real time and can provide actionable input to robotic schedulers, improving efficiency and data yield in time-critical programmes.

7.3 Extinction Monitoring

Atmospheric extinction directly affects the quality and interpretability of both photometric and spectroscopic data. Classical approaches rely on standard star observations to derive extinction coefficients for each night, a method that is accurate but inefficient and not well suited to fully robotic operation. The goal of this work was to test whether wide-field auxiliary cameras such as Skycam T and Z could provide continuous, automated extinction information in real time.

The results show that the system can indeed recover extinction behaviour, but with important caveats. Linear fits between instrumental magnitude and airmass yielded extinction coefficients in the expected range of $k_V \approx 0.13$–0.15, broadly consistent with values reported in the literature for the Observatorio del Roque de los Muchachos. However, the scatter in these fits was higher than for dedicated standard star observations, with typical R^2 values between 0.6 and 0.75. This noise is most likely a consequence of the unfiltered nature of the Skycams: without a well-defined passband, variations in stellar colour introduce systematic uncertainties that limit the precision of absolute extinction coefficients.

Despite this limitation, the system performs well in tracking differential extinction. Changes in transparency on sub-night timescales were clearly detected, allowing the identification of intervals of improving or deteriorating conditions. For robotic scheduling, this capability is often more important than measuring an exact extinction coefficient, since it enables adaptive decisions about whether to continue, suspend, or prioritise observations. In this respect, the Skycams provide a valuable and low-overhead transparency monitor.

Nevertheless, extracting robust absolute extinction values requires further development. A straightforward path would be to introduce filters to define a stable photometric system—either through a fixed Bessell V filter, a 2×2 lenslet array with different filters, or a filter wheel coupled to a more sensitive detector. Such modifications would significantly reduce colour-dependent scatter and allow the system to deliver extinction coefficients comparable in reliability to those from standard star observations.

In summary, the current Skycam system demonstrates that real-time extinction monitoring is feasible, but its primary strength lies in differential transparency tracking. With modest hardware upgrades, it could evolve into a robust source of absolute extinction coefficients for both photometric calibration and automated scheduling.

7.4 Sky Brightness Modeling

The natural sky background is one of the dominant noise sources in ground-based optical astronomy and therefore a critical input to exposure time estimation. Its level varies with lunar phase and altitude, with angular separation between the Moon and the target, and with atmospheric conditions. Classical models, such as that of Krisciunas and Schaefer, capture the broad trends but systematically underperform at a given site due to local atmospheric differences and simplifications in the analytic formulation. For robotic scheduling, where exposure times must be computed in real time and without human intervention, a more accurate and locally calibrated model is essential.

The empirical model developed in this thesis incorporates target altitude, lunar altitude, lunar phase, and the angular separation between target and Moon. Applied to LT background measurements, it reduced residual scatter by about 30% compared to the Krisciunas and Schaefer model. In particular, including the angular separation as a predictor was found to significantly improve the fit, capturing the strong directional dependence of scattered moonlight. The model reproduced both dark-time and bright-time trends, with residuals that were small enough to make it directly usable in an exposure time calculator.

From a scheduling perspective, this improvement has a tangible impact. In the legacy system, observers typically adopt conservative exposure times to guard against underestimation of sky brightness. These safety margins consume telescope time and reduce overall efficiency. By providing more accurate real-time predictions, the empirical model narrows the uncertainty envelope, reducing the need for such margins and allowing exposure times to be tuned more closely to actual conditions. This directly translates into higher observing efficiency, particularly in time-critical programmes where the number of spectra that can be obtained in a night is limited.

The model is not without limitations. Its accuracy does not account for rapid short-term fluctuations such as those caused by thin cirrus. Nevertheless, in combination with the cloud detection and extinction monitoring systems, it represents a substantial step toward a fully condition-aware scheduler. Future extensions could incorporate additional predictors such as atmospheric transparency indices or machine-learning refinements accounting for e.g. high-level Calima and the current sun-cycle, but the current model already offers a practical and deployable improvement over existing analytic prescriptions.

7.5 SNR-Based Exposure Estimation

At the heart of this thesis is the development of a dynamic exposure time estimator based on signal-to-noise ratio (SNR). Unlike traditional approaches, where observers manually specify exposure times in advance, often with generous safety margins, the SNR-based method adapts exposure length to the prevailing conditions in real time. This approach is central to robotic scheduling: it allows programmes to run independently without human oversight, ensures consistent data quality, and maximises efficiency by reducing wasted telescope time.

The estimator integrates inputs from the preceding modules. Atmospheric extinction coefficients, provide a measure of transparency; the empirical sky brightness model supplies background noise levels; and cloud forecasts inform short-term feasibility. Together with instrumental constants such as telescope aperture, throughput, and zeropoints, these inputs allow exposure times to be predicted for arbitrary targets and conditions. The calculation proceeds by inverting the standard SNR equation, solving for the exposure time required to reach a user-specified SNR threshold.

Validation was carried out on both historical observations and simulated test cases. For science targets in the 12 to 20 mag range, predicted exposure times closely matched those empirically required to achieve the requested SNR. Incorporating extinction and sky brightness parameters was especially beneficial for fainter targets, where transparency and background noise dominate the error budget. The system successfully adapted to changing conditions, with exposure times increasing under brighter sky and decreasing when conditions improved. This dynamic adjustment is a key advantage over static exposure times, which can lead to under- or over-exposure if conditions deviate from initial assumptions.

The scientific implications are significant. Analysis of LT archive data shows that nearly all manually chosen exposure times end in rounded multiples of 10 seconds, reflecting observer preference rather than numerical optimisation. Comparison with SNR-model-derived exposure times indicates that these rounded values frequently exceed the duration required to achieve the target signal-to-noise ratio under prevailing conditions, consistent with a conservative bias in manual estimation. The SNR-based tool reduces this quantisation inefficiency by tailoring exposures to real-time conditions, thereby decreasing the overhead associated with heuristic judgement. More importantly, it enables programmes such as the SPEC time programme on the New Robotic Telescope, where alerts with preliminary classifications can be ingested, matched to a science case, and executed without human intervention. In this context, the exposure calculator becomes the linchpin that translates a scientific requirement (e.g. $\mathrm{SNR} = 20$ for classification) into an executable command.

The SNR-based exposure estimator thus operationalises the core argument of this thesis: that condition-aware scheduling can move beyond static assumptions and observer heuristics to deliver autonomous, efficient, and scientifically competitive robotic astronomy. While edge cases remain, particularly under extreme transparency variations, the demonstrated performance shows that the approach is ready for deployment and provides a clear path forward for integration into next-generation scheduling systems.

7.6 Integration and Implications

Taken together, the four systems developed in this thesis represent the building blocks of a condition-aware scheduling framework. Each module addresses a different aspect of the observing environment: cloud forecasting constrains short-term feasibility, extinction monitoring provides real-time transparency, the empirical sky brightness model quantifies background noise, and the SNR-based estimator translates these inputs into actionable exposure times. Individually, each system demonstrates clear utility; in combination, they form a coherent toolkit for autonomous decision-making in robotic astronomy.

The integration of these modules is not merely an engineering exercise but a scientific necessity. When the Liverpool Telescope was commissioned in 2004, its queue scheduler was designed for long-lived transients such as supernovae, whose evolution could be tracked over weeks. Under those conditions, static altitude constraints and fixed exposure times were sufficient. In 2025, however, the discovery space has shifted decisively toward faster and fainter phenomena: fast blue optical transients, kilonovae, and gamma-ray burst afterglows evolve on timescales of hours or even minutes. For such sources, the bottleneck is no longer the ability to observe at all, but the ability to optimise observations in real time under variable conditions.

The results presented here demonstrate that this optimisation is achievable. Cloud forecasts, although limited to horizons of 15–30 minutes, align well with the median group length at the LT and are thus directly useful for "now/next" decisions. Extinction monitoring, while noisy in absolute terms due to the unfiltered Skycams, reliably detects differential transparency, which is sufficient for deciding whether to continue or suspend a spectroscopic sequence. The sky brightness model reduces prediction scatter by about 30%, narrowing the uncertainty envelope that observers would otherwise buffer with safety margins. The SNR-based exposure calculator integrates these inputs, replacing manual heuristics with condition-aware predictions. Together, these modules demonstrate a clear pathway to autonomous programme execution, where an alert with a preliminary classification can be matched to a science case and translated into an optimised observation without human intervention.

From an operational standpoint, this integration has two immediate implications. First, it offers a pathway to improving the scientific efficiency and effective observing yield of the Liverpool Telescope. For example, analysis of archival exposure rounding alone suggests that systematic inefficiencies amount to the equivalent of several nights of telescope time over a six-month period. Condition-aware exposure control and improved background prediction reduce the need for conservative safety margins, thereby increasing the fraction of time spent acquiring scientifically optimal data. While these gains do not alter the physical lifetime of the facility, they enhance its effective competitiveness by maximising science return within existing time allocations.

Looking forward, further refinements will strengthen this framework. Introducing filters to the Skycams would improve extinction accuracy; machine-learning methods may extend cloud forecast horizons; and integration with broker streams could allow the scheduler to directly incorporate classification probabilities into decision-making. These developments are future work, but the present results already show that condition-aware scheduling is both feasible and beneficial. In this way, the thesis delivers not only technical tools but also a demonstration of their scientific value: enabling robotic telescopes to respond efficiently and autonomously to the demands of modern time-domain astronomy.

7.7 Integration into Robotic Scheduling Systems

7.7.1 Adapting the System for the Liverpool Telescope

While the LT is not the long-term target, it provides a testbed for integration pathways that inform NRT design. This subsection does not introduce new experimental results; rather, it formalises how the proof-of-concept experiments described in previous chapters can be interpreted in terms of operational integration within the LT framework. The implementation of the SNR-based exposure time calculator as a standalone system demonstrated its potential to improve scheduling efficiency on the Liverpool Telescope. While the system was not (yet) integrated into the live scheduling queue, the architecture was designed to be compatible with the LT's existing infrastructure.

The system operates independently of the internal scheduler, using the publicly available observation logs to detect completed observations in real time. Upon detecting the observation of a target with a known magnitude such as a standard star or a science object with a SIMBAD or Lasair magnitude entry, the system calculates the exposure time required to reach a user-defined SNR based on prevailing sky conditions. Follow-up observations are then submitted via the LT's Robotic Telescope Markup Language (RTML) interface.

This external scheduling loop enabled a practical demonstration of real-time reobservation under dynamic conditions, without requiring internal modification to the queue system. The system successfully triggered dozens of follow-up observations during late 2024 and 2025, including standard stars and selected science targets. These reobservations were used to validate the underlying SNR model and test the accuracy of the exposure time predictions.

Although no changes to the internal LT scheduler were made during this project, several integration pathways are envisioned:

- **Internal exposure time override:** The scheduler could optionally allow an external SNR calculator to override the exposure time field in a queued observation request, based on current conditions.
- **SNR-aware proposal validation:** During proposal preparation, observers could specify the desired data quality (e.g. target SNR) rather than manually estimating exposure time. The scheduler would then convert this requirement into an expected integration time under representative conditions for feasibility assessment and time allocation. During execution, exposure durations would be adjusted dynamically

within the awarded time budget to achieve the requested data quality without exceeding the allocated observing time.

- **Live condition feed:** Integration with sky transparency and cloud monitors would enable the system to track real-time sky conditions more accurately, reducing prediction error.

The success of the reobservation experiments demonstrates the viability of these extensions. Their implementation would require careful coordination with the LT operations team to ensure consistency with existing queue policies and data quality standards. Nonetheless, the demonstrated compatibility with RTML interfaces and external scheduling logic makes future integration technically straightforward.

7.7.2 Integration into the Legacy LT Scheduler

While the proof-of-concept SNR-based exposure system has demonstrated its compatibility with the Liverpool Telescope's infrastructure through RTML-based observation injection, a more seamless integration into the core scheduling logic would enable greater efficiency and scientific return. This section outlines a practical pathway for modifying the legacy scheduler, based on the architectural principles described by Fraser et al. (2004).

The LT scheduler, originally designed in Java over two decades ago, employs a modular architecture in which scheduling decisions are made by assigning weights to competing observation groups. The weighting functions operate within a defined "schedule context" that includes environmental inputs, queue constraints, and telescope/instrument status. The decision logic is executed in real time to generate observation sequences dynamically throughout the night.

To integrate SNR-based exposure time prediction into this framework, the following modifications are proposed:

- **Augment the Schedule Context:** Extend the schedule context object to include cloud maps, real-time sky transparency, sky brightness and seeing estimates, either via existing monitor feeds, models or all-sky camera predictions. These variables are essential inputs to accurate SNR estimation.

- **Introduce a Dynamic Exposure Module:** Develop a new Java module within the scheduler that computes the required exposure time for a given SNR target based on current conditions, object brightness, and instrument parameters. This module should be accessible from the weighting function layer.

- **Extend Proposal Specification:** Allow observers to specify SNR-based goals in their Phase II submissions, using an extended schema. If a target specifies SNR rather than exposure time, the scheduler should query the dynamic exposure module at scheduling time to compute the appropriate exposure length.
- **Optional Weighting Bias:** Include an optional weighting bias in the scoring function that favours observations for which the predicted SNR is currently optimal. This would allow the scheduler to opportunistically prioritise observations when conditions are favourable.

None of these changes require a wholesale rewrite of the scheduler. The modularity described by Fraser et al. (2004) allows for targeted enhancements, particularly within the weighting function implementation layer. The exposure time override could be implemented analogously to existing flexible group scoring, with hooks into the real-time environmental context.

Moreover, the scheduler's compatibility with RTML, already exploited in the proof-of-concept system, ensures that fallback operation via external submission remains possible during transitional deployment phases. This dual-mode operation could ease testing and validation before committing to full internal integration.

> *"The key to robotic scheduling is the ability to make informed, adaptive decisions based on environmental conditions and observation requirements."* - adapted from Fraser et al. (2004)

In this spirit, integrating SNR-based scheduling logic directly into the Liverpool Telescope's core system represents a natural evolution of its design, and aligns with its longstanding commitment to efficiency, responsiveness, and scientific value.

7.7.2.1 SNR Logic for the Legacy Scheduler

The Liverpool Telescope scheduler, written in Java and maintained since the early 2000s, operates by dynamically selecting the next best observation group during each scheduling sweep. This selection is handled in the method `nextScheduledJob()` of the `BasicDespatchScheduler` class, which constructs a view of the current Phase II database, assesses feasibility, and ranks candidate observations using a scoring model.

To integrate SNR-based exposure estimation into this logic, a minimal and targeted modification is proposed. After retrieving the list of candidate groups via the `Phase2CompositeModel`, a new pass can be inserted to adjust exposure times for those observations that specify an SNR goal rather than a fixed exposure time.

This process involves:

1. Inspecting each observation group for an SNR goal (e.g. via a new field `snrTarget`).
2. Estimating the target's apparent magnitude (e.g. from the object catalogue, recent photometry or another new field `targetMag` for each available filter).
3. Retrieving current environmental conditions (transparency, seeing, sky brightness) from the scheduler's internal environment model or the external monitors.
4. Computing the required exposure time using the known instrumental parameters and the user-defined SNR.
5. Updating the group's exposure time field before the scoring phase.

This logic can be implemented in a utility class `DynamicExposureAdjuster`, which applies the SNR logic to all eligible groups prior to the scoring pass. Integration is limited in scope, requires no modification to the weighting function logic, and is consistent with the modular architecture described in the original LT scheduling framework (Fraser et al., 2004).

The change would allow observers to request photometry and spectroscpoy by SNR target rather than by trial-and-error exposure selection, and would support a broader shift toward condition-aware observing on robotic telescopes.

7.7.2.2 User-Level Integration via IO:O Exposure Panel

While backend integration of SNR-aware exposure calculation is critical for execution, user adoption depends on intuitive configuration at the proposal stage. One viable integration point is a IO:O or SPRAT-specific extension to the Liverpool Telescope Phase II configuration system, particularly the *Exposure* button in the sequencer interface.

To support SNR-based requests, this interface could offer an additional mode alongside fixed exposure time, allowing observers to specify:

- A target Signal-to-Noise Ratio (e.g. SNR = 100),
- The filter to be used (e.g. SDSS-g, Bessell-V),
- The target magnitude, either:

  - Entered manually by the user, or
  - Retrieved dynamically via a SIMBAD or Lasair query at scheduling time.

The manually specified magnitude approach ensures offline compatibility and reproducibility, while the dynamic lookup allows for real-time reactivity (e.g. for variables or transients). However, the latter option would require internet access within the scheduler subsystem, which is currently isolated from external networks for operational reasons.

An XML schema extension for the Phase II observation configuration might look as follows:

```
<exposure mode="snr">
    <snrTarget>100</snrTarget>
    <filter>sdss-g</filter>
    <magnitude source="manual">15.2</magnitude>
</exposure>
```

This information would be passed down into the Phase2CompositeModel and read by the scheduling engine. Prior to scoring, the `DynamicExposureAdjuster` module (Section 7.7.2.1) would interpret these fields, query current conditions, and calculate the required exposure time accordingly.

To maintain operational robustness, the scheduling system should support fallback logic in case dynamic lookup fails or network access is unavailable. This ensures consistent behaviour whether the system is operating in an “offline” environment or during transient follow-up campaigns where internet access is required.

This pathway enables a seamless user experience for requesting science observations based on data quality (SNR), rather than arbitrary or outdated estimates of exposure length, and marks a significant step toward fully adaptive robotic scheduling.

7.7.3 Integration into the New Robotic Telescope

Building on the successful proof-of-concept implementation on the Liverpool Telescope, the SNR-based exposure time system is intended to be a native component of the New Robotic Telescope (NRT). Unlike the LT, where integration requires retrofitting into an established architecture, the NRT offers a clean-slate opportunity to embed condition-aware scheduling logic from the outset.

The NRT is being developed with a modern, modular architecture based on containerised microservices orchestrated by Kubernetes (Bento et al., 2024; Copley et al., 2024). Each functional component of the observatory, including scheduling, environment monitoring, data processing, and observation control, will operate as an independent service, communicating via well-defined APIs. This architectural choice provides scalability, robustness, and flexibility in deployment, while also facilitating the integration of third-party modules such as the SNR-based exposure time estimator.

At the heart of the scheduling system will be a look-ahead scheduler, capable of forecasting and optimizing observation plans over a range of time horizons. This scheduler will natively incorporate the cloud detection and extinction monitoring system developed earlier in this thesis (Chapters 3 and 4). Real-time inputs from the all-sky cloud tracking camera and Skycam-based extinction measurements will be continuously put into the scheduling context, allowing the system to anticipate and respond to evolving sky conditions.

In this setting, observers will be able to submit science programmes by specifying target coordinates, estimated magnitudes, and required SNR thresholds. The scheduler will then:

1. Forecast observing conditions for all viable time windows, including expected extinction, transparency, sky brightness, and lunar geometry.
2. Estimate the exposure time needed to reach the requested SNR using the current instrumental parameters and predicted environmental conditions.
3. Score and prioritise observations based on schedulability, expected efficiency, urgency, and science priority.

The SNR-based exposure system will be deployed as a dedicated service within the Kubernetes cluster. It will be invoked during the scheduling sweep to compute optimal exposure durations per target and filter. As part of the Phase II interface, observers will have the option to specify either a fixed exposure time or an SNR target. The latter mode will enable fully dynamic, quality-driven observing.

Furthermore, the integration of real-time cloud and extinction data enables the system to make scientifically informed trade-offs. For instance, in the presence of mild extinction, the scheduler may choose to increase exposure times to preserve SNR, or postpone observations until transparency improves, depending on the scientific urgency and other constraints.

This architecture not only improves individual observation efficiency, but also enables strategic planning over entire nights, weeks, or semesters. It represents a paradigm shift in robotic telescope scheduling: from static queues and conservative estimates to a dynamic, data-driven system that aligns telescope operation with the real-time sky.

7.7.4 Integration of automated transient classification

Within the Liverpool Telescope and New Robotic Telescope environment, the principles of automated transient classification can be operationalised through a structured, priority-based approach to proposal management. One core element of this approach would be the maintenance of two complementary proposal streams, each designed to address a distinct stage in the transient classification process while making good use of available conditions.

The first stream, designated as Priority A and granted Target-of-Opportunity (ToO) capabilities, would be reserved for the initial spectroscopic classification of newly discovered transients for which no prior spectrum exists. To avoid choking the scheduling queue, this can be limited to e.g. five classifications per night. These targets, typically identified through real-time broker filtering and photometric classification, would be injected directly into the observing queue with minimal latency. Given their scientific importance and the rapid evolution of many classes of transient, particularly core-collapse supernovae, kilonovae, and gamma-ray burst afterglows, these observations would be executed under the currently available conditions, ideally photometric and with low sky brightness, to ensure the highest possible data quality and maximise the chance of successful classification on the first attempt.

The second stream would target transients for which a classification spectrum has already been obtained, but for which further spectroscopy remains scientifically valuable. This could include monitoring the spectral evolution of rare events, confirming marginal classifications, or capturing additional phases of an outburst. By assigning these observations a Priority C or Priority Z designation, they can be scheduled flexibly: not relegated to background “filler” tasks, but preferentially executed during bright time or under non-photometric conditions when high-priority work is not possible. This ensures that spectroscopic resources are still productively employed in sub-optimal conditions without displacing urgent classification targets.

Central to the efficiency of both streams is the integration of an automated exposure time calculator directly into the SPRAT spectrograph control logic. Relying solely on an external, pre-scheduled exposure calculation can be sub-optimal, as it necessarily assumes conditions at the time of scheduling rather than at execution. For Rank A ToO triggers, which are typically observed within minutes of group creation, this is not a significant drawback, since the calculated exposure closely reflects the actual conditions. However, for lower-priority Priority C/Z observations, which may execute hours or even a night later, the prevailing conditions can differ substantially from those assumed. Recalculating exposures too frequently risks generating large numbers of database updates and resubmissions, while relying on a stale calculation risks under- or over-exposure. A practical compromise is to compute exposure times for the worst expected conditions during the intended execution window and to constrain these values for up to two nights, avoiding repeated queue updates while ensuring that data remain scientifically usable under a wide range of conditions. An alternative approach would be to extend the Remote Telescope Markup Language (RTML) specification to allow group updates, enabling direct modification of exposure times for already-scheduled observations. As far as the author is aware, this feature is not currently supported, and its absence limits the degree to which exposure plans can be adapted without re-submitting entire groups.

The system would dynamically determine the required exposure time to achieve a user-defined signal-to-noise ratio, based on real-time estimates of transparency, seeing, airmass, and sky brightness. Under favourable conditions, the calculator would shorten exposures to avoid overshooting the SNR requirement, thereby increasing nightly throughput. In poor conditions, it would extend exposures appropriately, or in extreme cases defer the observation until conditions improve, thereby protecting the scientific viability of the data. The use of such a system is particularly well-suited to classification work, where a moderate but reliable SNR threshold is often the critical determinant of success.

When combined with automated alert ingestion, broker-based candidate selection, and science-case-driven scheduling, this two-tiered proposal structure provides a coherent framework for large-scale transient classification. The Priority A/ToO stream ensures rapid response and maximises the probability of capturing fast-evolving phenomena at the optimal moment, while the Priority C/Z stream extends scientific return by exploiting less-than-ideal conditions for secondary classification work. By embedding the SNR-based exposure model into both, the system adapts seamlessly to the prevailing environment, maximising efficiency across the full range of observing conditions. In the LSST era, when the classification bottleneck will be defined as much by operational agility as by raw telescope time, such a strategy offers a pragmatic and scalable path to closing the loop from discovery to classification.

7.8 Semester-Time Estimation and Validation Tool

To support the planning and validation of observing proposals, a complementary simulation tool was developed that estimates the total telescope time required to observe a target sample with a user-defined SNR. The tool combines the SNR-based exposure time calculations described in this chapter with statistical site conditions at the Observatorio del Roque de los Muchachos (ORM), enabling realistic assessments of observational feasibility.

The simulator operates on target lists defined by right ascension, declination, and apparent magnitude, and supports multiple instrument/filter configurations. For each target, the number of exposures and desired SNR per filter are specified. The system then calculates:

- The exposure time required per target/filter
- The visibility window and airmass profile for each target
- Expected sky brightness and seeing conditions (based on ORM statistics)
- Total required telescope time, including overheads

The tool currently supports photometric observations with the IO:O instrument (ugriz filters) and is being extended to accommodate spectroscopic use cases (e.g., SPRAT). A graphical visibility calendar and CSV export allow integration into proposal preparation workflows.

The simulator serves two primary functions:

1. **Proposal design:** Enabling observers to justify requested time based on realistic, condition-aware exposure time predictions.
2. **Schedulability validation:** Providing the observatory with a mechanism to assess whether the proposed programme fits within expected conditions and time allocations.

This validation layer bridges the gap between theoretical exposure time calculators and the practical realities of queue-scheduled, weather-dependent observing. It also supports long-term planning for robotic telescopes, ensuring that science goals are achievable within the constraints of atmospheric conditions and operational overheads.

It is important to note that semester-time estimates derived from statistical site conditions represent expectation values rather than guarantees. Actual observing efficiency will inevitably vary due to stochastic weather, transparency fluctuations, and operational interruptions. In cases where realised conditions deviate significantly from the statistical assumptions used at proposal stage, programmes may require either additional time to reach their requested SNR or acceptance of reduced data quality within the originally allocated time budget. The simulator does not resolve such allocation conflicts; rather, it provides a quantitative framework that makes these trade-offs explicit. Final decisions regarding time redistribution or completion thresholds remain a matter of observatory policy and scheduling governance.

It is further recognised that deviations from projected semester-time estimates are not driven solely by atmospheric variability. Technical downtime, scheduled maintenance, instrument availability, and unforeseen operational interruptions are intrinsic features of long-term robotic facilities. Such stochastic losses occur independently of atmospheric modelling and are typically accounted for in observatory-level time accounting and subscription strategies. The simulator therefore models the atmospheric and geometric components of observing feasibility, while broader operational contingencies remain governed by institutional scheduling policies and risk buffers applied at the programme-allocation level.

7.9 Performance and Limitations of the SNR-Based Systems

The implementation and validation of an SNR-driven exposure time system demonstrate a viable path toward more efficient and adaptive telescope scheduling. While the current system was tested as a standalone proof-of-concept on the Liverpool Telescope (LT), its core components are designed with future integration into fully autonomous facilities such as the New Robotic Telescope (NRT) in mind.

7.9.1 Advantages of SNR-Based Scheduling

The results presented in this chapter highlight several key benefits over traditional fixed-time scheduling:

- **Efficiency:** Observations are matched to the actual conditions, avoiding unnecessary overexposures while ensuring sufficient depth under suboptimal conditions.
- **Scientific quality:** Data products achieve a consistent and meaningful SNR, reducing the risk of failed science observations due to underexposure.

- **Automation:** The system removes the need for manual exposure time tuning and accommodates targets with uncertain visibility or scheduling flexibility.
- **Transparency-aware:** Sky brightness and airmass are directly accounted for, and future extensions will incorporate real-time extinction and cloud measurements.

These features make the system particularly well-suited to robotic or queue-scheduled telescopes where observation parameters must be optimised automatically without user intervention during execution.

7.9.2 Limitations and Uncertainties

Despite the promising results, several limitations remain:

- **Transparency variation:** Real-time extinction due to thin clouds or aerosols is not currently measured dynamically but assumed to be stable or photometric.
- **Catalogue uncertainty:** Apparent magnitudes obtained from external databases (e.g., Lasair, SIMBAD) may be outdated or imprecise, especially for variable or transient sources.
- **Instrument calibration:** The zero points used are fixed per filter and do not account for night-to-night variability, which may introduce minor errors in predicted signal.
- **Sky brightness model bias:** Although corrected empirically, the model may still underperform during unusual lunar geometries or during fast-changing twilight conditions.

These sources of uncertainty are addressed conservatively by the pipeline (e.g., through re-observation of standard stars), but further refinement is planned.

7.10 Scheduler Scoring Framework

Any robotic telescope scheduler must resolve a fundamental tension: many scientifically valuable jobs are simultaneously feasible, yet only one can be executed at a time. At the Liverpool Telescope (LT) this has traditionally been handled by a priority queue, where each job is assigned a static rank by its Phase II definition and then executed in order of feasibility. While this approach was pioneering in the early 2000s, it has become increasingly inadequate in the era of time-domain astronomy, where conditions evolve rapidly and new transient alerts arrive unpredictably.

The environmental context developed in this thesis enables a move beyond static queue scheduling. Previous chapters have demonstrated that:

- Clouds can be detected and forecast in real time using all-sky cameras.
- Extinction can be measured continuously with Skycam Z/T, allowing transparency losses to be quantified.
- Sky brightness can be modelled empirically as a function of lunar geometry, improving predictions for background-limited exposures.
- Required exposure times can be computed dynamically to achieve a specified signal-to-noise ratio (SNR).

Together, these modules provide the real-time inputs required for a scheduler that not only checks feasibility but also ranks competing jobs by scientific priority, observational opportunity, and achievable data quality.

To formalise this process, a **scoring function** is introduced. Each candidate job is evaluated against a set of criteria, mapped to normalised scores in the range $[0,1]$, and combined with tunable weights to produce an overall merit value. The job with the highest score is selected for execution. The framework is modular (terms can be added or removed as required), tunable (weights can be adapted to programme or facility priorities), and predictive (forward-looking terms such as setting urgency or forecasted extinction can be included).

7.10.1 Scoring Function

The scoring function assigns a numerical value S to every feasible job in the queue:

$$S = w_{\mathrm{SP}} \cdot SP + w_{\mathrm{V}} \cdot V + w_{\mathrm{CM}} \cdot CM + w_{\mathrm{SNReff}} \cdot SNReff + w_{\mathrm{U}} \cdot U + w_{\mathrm{C}} \cdot C + \epsilon, \quad (7.1)$$

where each component is normalised to the range $[0,1]$, w_i are tunable weights, and ϵ is a small random jitter to break ties. The core terms are described in detail below:

Science Priority (SP). This term encodes the human-defined ranking of scientific importance, while allowing urgent overrides. Proposal classes are mapped to scores via

$$SP_{\text{base}} = \begin{cases} 1.0 & \text{Priority A} \\ 0.7 & \text{Priority B} \\ 0.4 & \text{Priority C} \\ 0.1 & \text{Priority Z} \end{cases}$$

An urgency flag (e.g. transient follow-up) introduces a bonus of +0.2, capped at unity:

$$SP = \min(1, SP_{\text{base}} + 0.2).$$

Visibility (V). Visibility measures how favourable the current position of the target is on the sky. A naïve altitude-only score systematically disfavors western targets, since they are always "worse" than their symmetric eastern counterparts until it is too late. To counteract this, a two-part measure is introduced.

First, a normalized altitude factor:

$$V_{\text{alt}} = \frac{\sin(h) - \sin(h_{\text{min}})}{\sin(h_{\text{max}}) - \sin(h_{\text{min}})},$$

with h the target altitude and h_{min} and h_{max} denote the minimum and maximum altitudes reached by the specific target during the current night (i.e. the observing limit and nightly culmination altitude, respectively). The altitude term is therefore normalised per target, ensuring that objects are prioritised relative to their own attainable geometry rather than against an absolute zenith reference.

Second, a setting-urgency factor:

$$V_{\text{set}} = \exp\left(-\frac{t_{\text{set}} - t_{\text{now}}}{\tau}\right),$$

where $(t_{\text{set}} - t_{\text{now}})$ is the time remaining until the object sets, and τ is a tunable decay constant (typically $\tau = 1\,\text{h}$).

The final visibility score is then

$$V = \alpha \cdot V_{\text{alt}} + (1 - \alpha) \cdot V_{\text{set}},$$

with $\alpha \in [0, 1]$ controlling the balance between altitude optimisation and urgency to observe targets before they set.

Condition Matching (CM). This term quantifies how well the current observing conditions align with the constraints of the job. Examples include requirements on image quality (seeing $< 1.5''$), sky brightness (e.g. V-band darker than 20 mag arcsec^{-2}), or photometric transparency. The value of CM is defined on the interval $[0, 1]$: 1 if all requested conditions are satisfied, fractional if only partially met, and 0 if the job is infeasible. This ensures that programmes requiring excellent conditions are not displaced by less demanding observations when conditions are suboptimal.

SNR Efficiency (SNReff). The $SNReff$ term measures how efficiently a job can be executed under current conditions. It is defined as the ratio of achievable to requested signal-to-noise ratio,

$$SNReff = \min\left(1, \frac{\mathrm{SNR}_{\mathrm{achievable}}}{\mathrm{SNR}_{\mathrm{required}}}\right).$$

Here, $\mathrm{SNR}_{\mathrm{achievable}}$ denotes the signal-to-noise ratio predicted by the SNR model for the exposure time either currently specified in the job request, evaluated under prevailing atmospheric conditions or a “global” maximum specified for the instrument. The metric therefore quantifies how efficiently the requested observation can be executed at the present time without modifying its nominal integration time.

If the achievable SNR exceeds the requested value, the efficiency saturates at unity; if it falls short, $SNReff$ decreases proportionally. This factor favours jobs that can be completed efficiently in situ, given the prevailing seeing, extinction, and background levels.

Urgency (U). Urgency reflects the time sensitivity of an observation. For transient phenomena, U increases as the target fades or as a deadline approaches. For monitoring programmes, U rises when cadence requirements are at risk of being missed. A common formulation is an exponential or logistic function of time, which causes U to grow smoothly as the window of opportunity closes.

As an illustrative example, a fast blue optical transient (FBOT) evolves on timescales of only a few hours: spectra obtained within the first night may reveal unique high-velocity signatures that are no longer detectable the following day. Similarly, a kilonova typically peaks in brightness within $\sim$12 to 24 hours after a neutron star merger and then fades rapidly by several magnitudes over the subsequent days. In both cases, the scientific return is maximised only if the scheduler promotes the target immediately, a behaviour naturally captured by a rising U term. This ensures that fast transients rise to the top of the queue without the need for manual intervention.

Cadence Adherence (C). The C term encodes how well a candidate observation fulfils the cadence requirements of long-term monitoring programmes. If the elapsed time since the last observation closely matches the requested cadence, C approaches unity. If the programme has already been observed recently, or if the cadence requirement is otherwise satisfied, C is reduced. This prevents oversampling of well-covered programmes and promotes fair temporal distribution across the monitoring set.

Random Jitter (ϵ). Finally, a small stochastic perturbation ϵ is added to break ties between jobs with identical scores. This jitter is typically drawn from a uniform distribution with very small amplitude (e.g. $\epsilon \sim \mathcal{U}(-0.01, 0.01)$). Its role is purely technical: it prevents deterministic cycling or persistent bias towards specific targets when two or more jobs are otherwise indistinguishable.

7.10.2 Slew Cost as a Secondary Metric

In addition to scientific priority and visibility, the time required for the telescope to move from its current position to the next target represents a natural source of inefficiency. On an alt–azimuth mount such as the Liverpool Telescope, slewing involves both azimuthal rotation and altitude motion, with accelerations, decelerations, and mechanical limits making the total time somewhat more complex than a simple angular separation. Nevertheless, as a first approximation the slew cost can be estimated from the great-circle angular separation between successive targets, translated into an expected time penalty by the mount control system.

Two approaches can be considered for incorporating slew cost into the scheduler:

1. **Score-based penalty:** A normalised slew-efficiency term can be added directly to the scoring function,

   $$S = \ldots + w_{\text{slew}} \cdot \left(1 - \frac{t_{\text{slew}}}{t_{\max}}\right),$$

   where t_{slew} is the estimated slew time and $t_{\max}$ is a normalisation constant corresponding to the maximum expected slew. In this formulation, closer targets receive a slight score advantage, while distant targets are penalised.

2. **Tie-breaking rule:** The scheduler can first compute scores without considering slew, then among the top-N candidates select the job with the smallest expected slew time. In this case, slew acts as a secondary discriminator, only affecting decisions when other scientific and observational metrics are approximately equal.

Both approaches are mathematically valid but encode slightly different philosophies. The score-based penalty allows slewing efficiency to influence the global optimisation directly, while the tie-breaking rule treats slewing as a secondary concern that should never override science-driven priorities. For the LT and similar facilities, where typical slew times are on the order of tens of seconds and individual exposures are several minutes long, the latter strategy is generally sufficient.

It is therefore concluded that slewing should be considered as a secondary metric within the scheduling framework. Its inclusion acknowledges the potential for wasted time during frequent target changes, while ensuring that scientific drivers such as priority, visibility, and SNR efficiency remain dominant. A more sophisticated treatment—for example, a dynamic model that minimises total slew time across an entire night—would constitute a significant optimisation problem in its own right, and is best considered as a future research avenue.

7.10.2.1 Discussion

The scoring function design illustrates a compromise between hard constraints and flexible optimisation. Science Priority captures human strategic decisions, while Visibility ensures efficient exploitation of real-time observing windows. The introduction of the setting-urgency factor V_{set} directly addresses a long-standing weakness of altitude-only ranking, and aligns the scheduler's behaviour with intuitive practice ("take it now before it sets"). Both terms remain normalised and modular, which facilitates adjustment relative to other factors such as condition matching or SNR efficiency.

7.10.2.2 Future Work

Future refinements of the scoring function will focus on three directions:

1. **Dynamic Weighting:** Allowing w_i to adapt automatically based on queue composition (e.g. increasing w_{V} if many short-lived western targets are available).
2. **Learning-based Optimisation:** Training weightings on historical schedule outcomes, using metrics such as science return or completion fraction.
3. **Forecast Integration:** Coupling V_{set} to short-term forecasts of cloud cover or extinction, yielding a look-ahead capability where western targets are promoted only if clear sky is predicted.

Such extensions will enable the scheduler to evolve beyond fixed heuristics and towards a predictive, self-optimising system.

7.10.2.3 Weight Optimisation

The choice of weights w_i strongly influences scheduler behaviour. A data-driven calibration is proposed, making use of archival Liverpool Telescope observations. Each completed observation provides:

- Target properties (magnitude, filter, requested exposure).
- Instrument used (e.g. IO:O, SPRAT).
- Observing conditions at the time (airmass, extinction, sky brightness).
- Achieved signal-to-noise ratio or quality flag.

A loss function L can be defined that penalises mismatches between the score assigned by the scheduler and the actual scientific utility of the observation, for example:

$$L = \sum_j \left(\hat{S}_j - U_j \right)^2 ,$$

where $\hat{S}_j$ is the normalised score assigned to job j by the scheduler and U_j is the observed utility (e.g. binary usable/unusable, or a continuous function of achieved SNR / required SNR).

The optimal weight vector $\mathbf{w} = (w_{\text{SP}}, w_{\text{V}}, w_{\text{CM}}, \ldots)$ is then obtained by minimising L across the historical dataset. In practice, this optimisation can be implemented via gradient descent, Bayesian optimisation, or grid search, depending on the dimensionality and noise level of the training data. The resulting weights reflect both human-defined priorities and the empirical efficiency of the telescope under real observing conditions. Validation on a hold-out set of nights ensures that the scheduler generalises to unseen scenarios.

7.11 Summary and Outlook

This chapter has evaluated the performance of the four systems developed in this thesis and considered their implications for robotic scheduling. The results show that each module achieves its core objective: clouds can be detected and forecast on timescales that match the cadence of typical observing groups; extinction can be tracked continuously, albeit with noisy absolute coefficients, enabling real-time assessment of transparency; the empirical sky brightness model reduces prediction scatter compared to analytic prescriptions, narrowing uncertainty in background noise estimates; and the SNR-based exposure time estimator successfully integrates these inputs to deliver condition-aware predictions of data quality.

Taken individually, these results demonstrate proof of concept for each subsystem. Taken together, they illustrate a coherent framework for adaptive scheduling in the time-domain era. The transition from long-lived to fast-evolving transients has shifted the scheduling challenge from planning days or even weeks in advance to optimising decisions minute by minute. The systems developed here address that challenge directly: they enable observations to be matched dynamically to both scientific requirements and prevailing conditions, reducing inefficiency and human overhead. In doing so, they lay the foundation for autonomous execution of programmes such as SPEC on the New Robotic Telescope, where success depends on rapid, independent, and robust follow-up.

Looking forward, several extensions will strengthen this framework. The exposure time model can be expanded with short-term forecasts of sky transparency, integrating extinction values from Skycam Z and cloud predictions from all-sky systems to anticipate conditions over the next observing blocks. The system can also be adapted to estimate achievable SNR across an entire queue in advance, supporting global optimisation of nightly and semester-level efficiency. Improved coupling with photometric and spectroscopic pipelines will allow direct quality-assurance feedback loops, so that science products inform subsequent scheduling in real time. These developments point toward a more intelligent, responsive, and efficient telescope network.

The outlook is therefore twofold. In the near term, incremental integration of these modules into the Liverpool Telescope offers a route to immediate efficiency gains, extending the competitiveness of its existing infrastructure. In the longer term, the same framework provides a blueprint for the New Robotic Telescope, where condition-aware scheduling will be central to achieving its science goals. While refinements such as filters for extinction monitoring, longer cloud forecast horizons, and machine-learning approaches will further improve performance, the systems presented here already demonstrate both feasibility and benefit. The work thus delivers not only practical tools but also a scientific proof of principle: robotic telescopes can be made not only faster, but also smarter in the way they allocate their most precious resource, observing time.

Chapter 8

Conclusion

> *I may not have gone where I intended to go, but I think I have ended up where I needed to be.*
>
> Douglas Adams

8.1 Conclusion

This thesis has presented a framework for condition-aware scheduling of robotic telescopes, developed and tested using the Liverpool Telescope and its auxiliary instruments. The work addressed four central challenges: detecting and forecasting clouds, monitoring atmospheric extinction, sky brightness estimation, and dynamically optimising exposure times through an SNR-based approach. Together, these components close the loop between environmental monitoring, predictive modelling, and scheduling decisions.

The cloud detection and forecasting pipeline demonstrated that real-time cloud masks and short-term forecasts are achievable with modest calibration. Predictions of cloud arrival and departure times were accurate on horizons of 15 to 30 minutes, aligning closely with the median duration of observing groups on the LT. This shows that predictive cloud monitoring can be scientifically useful even if it does not extend to hour-long horizons: for transient follow-up, what matters most is whether the next block of data will be cloud-free.

The extinction monitoring experiments showed that Skycams can provide continuous, automated transparency estimates, though not without limitations. While extinction coefficients derived from Skycam Z were consistent with literature values ($k_V \approx 0.13$–0.15), the scatter was higher than for dedicated standard star observations. This noise arises from the unfiltered nature of the instruments, which introduces colour-dependent systematics. Despite this, the system performed reliably for differential extinction: detecting whether conditions were improving or deteriorating in real time. This capability is directly relevant to robotic scheduling, where the decision to continue or suspend an observation is often more important than measuring an exact coefficient.

The empirical sky brightness model provided a significant improvement over analytic prescriptions such as Krisciunas & Schaefer. By incorporating lunar altitude, phase, and angular separation, residual scatter was reduced by 30 to 40%. This improvement has a direct operational consequence: observers no longer need to add large safety margins to exposure times, which frees additional telescope time and increases the number of targets that can be observed. While the model cannot capture every short-term fluctuation, it provides a deployable baseline for real-time scheduling.

The culmination of this work was the development of an SNR-based exposure time estimator. By integrating extinction and sky brightness inputs into the standard SNR equation, exposure times can be calculated dynamically to meet user-specified quality requirements. Tests showed that the tool reproduces empirical exposure needs for targets between 12–19 mag, and in many cases achieves the intended SNR with shorter exposures than those chosen manually. The inefficiency of human-selected exposure times, often rounded to convenient multiples, was explicitly demonstrated in archival LT data. The SNR-based system removes this inefficiency and, more importantly, enables autonomous programme execution such as the SPEC programme on the NRT, where classification spectra must be obtained without human intervention.

In scientific terms, this thesis demonstrates that robotic telescopes can adapt to the demands of modern time-domain astronomy. The LT scheduler, designed in 2004 for long-lived supernovae and novae, is no longer sufficient in an era where instrumentation and surveys now allow, and indeed require, measurements of fast blue optical transients, kilonovae, and gamma-ray burst afterglows on their intrinsic hour-scale timescales. The systems developed here show how telescopes can become not only faster in their response, but also smarter in their decision-making, matching observations dynamically to prevailing conditions. This represents a shift from static, altitude-driven scheduling to a condition-aware, science-driven paradigm.

8.2 Future Work

Several avenues remain open for development and refinement:

- **Cloud forecasting:** extending the deterministic Kalman approach to probabilistic forecasts with confidence intervals, enabling schedulers to quantify risk when deciding between high-priority but condition-sensitive observations.
- **Extinction monitoring:** reducing noise by introducing filters (e.g. a fixed Bessell V filter, a 2×2 lenslet array, or a filter wheel), thereby allowing both absolute extinction coefficients and differential trends to be measured reliably.
- **Sky brightness modeling:** incorporating machine-learning techniques, seasonal priors, and real-time updates from Skycams to generate short-term forecasts of background levels, which can be integrated directly into the scheduler's cost function.
- **Scheduler integration:** embedding the modules into the New Robotic Telescope's look-ahead scheduler, influencing both real-time execution and semester-level planning through predictions of achievable SNR and expected yield.
- **Generalisability:** adapting the framework to other facilities with wide-field monitors, creating a globally distributed network of condition-aware telescopes.

Closing Remarks

The journey of this thesis has been both technical and scientific: from designing algorithms for cloud masks and extinction fits, to asking what those numbers mean in the context of real observing programmes. The central conclusion is that robotic telescopes can do more than react quickly: they can also make automated, informed choices about how best to use their time. By integrating environmental monitoring, predictive modelling, and dynamic exposure control, telescopes become autonomous partners in discovery rather than passive instruments awaiting human commands.

The systems presented here are not the final word. They are prototypes, built on auxiliary instruments and validated under real conditions, but with clear scope for refinement. Yet even in their current form, they demonstrate a new operational paradigm: one where the scientific yield of a night's observing is maximised not by human intuition or conservatism, but by data-driven adaptation in real time. As new generations of fast transient surveys continue to expand the discovery space, such condition-aware scheduling will be essential to ensure that telescopes can keep pace. In this sense, the work has not only ended up where it needed to be, but also points toward where the field is going next.

Bibliography

Aceituno, J., S. F Sánchez, F. J. Aceituno, D. Galadí-Enríquez, J. J. Negro, R. C. Soriguer, and G. Sanchez Gomez (Sept. 2011). “An All-Sky Transmission Monitor: ASTMON”. en. In: *Publications of the Astronomical Society of the Pacific* 123.907, pp. 1076–1086. ISSN: 00046280, 15383873. DOI: `10.1086/661918`. URL: `http://iopscience.iop.org/article/10.1086/661918` (visited on 03/09/2022).

Adam, Jan, Jens Buss, Kai Brügge, Max Nöthe, and Wolfgang Rhode (2017). “Cloud Detection and Prediction with All Sky Cameras”. In: *EPJ Web of Conferences* 144. Ed. by P. Trávníček, M. Prouza, M. Gaug, and B. Keilhauer, p. 01004. ISSN: 2100-014X. DOI: `10.1051/epjconf/201714401004`. (Visited on 05/09/2020).

Arrais, Juliana Marian, Bruno Juncklaus Martins, Thiago Zimmermann Loureiro Chaves, Allan Cerentini, Sylvio Luiz Mantelli Netto, and Aldo Von Wangenheim (2022). “Systematic Literature Review on Ground-Based Cloud Tracking Methods for Nowcasting and Short-term Forecasting”. In: DOI: `10.13140/RG.2.2.17730.25281`. (Visited on 01/12/2023).

Augusteijn, T. (Mar. 2001). “RoboDIMM”. In: *The Newsletter of the Isaac Newton Group of Telescopes* 4. ADS Bibcode: 2001INGN....4...27A, pp. 27–28. URL: `https://ui.adsabs.harvard.edu/abs/2001INGN....4...27A` (visited on 07/19/2025).

Barreto, África, Emilio Cuevas, Rosa D. García, Judit Carrillo, Joseph M. Prospero, Luka Ilić, Sara Basart, Alberto J. Berjón, Carlos L. Marrero, Yballa Hernández, Juan José Bustos, Slobodan Ničković, and Margarita Yela (Jan. 2022). “Long-Term Characterisation of the Vertical Structure of the Saharan Air Layer over the Canary Islands Using Lidar and Radiosonde Profiles: Implications for Radiative and Cloud Processes over the Subtropical Atlantic Ocean”. In: *Atmospheric Chemistry and Physics* 22.2, pp. 739–763. ISSN: 1680-7316. DOI: `10.5194/acp-22-739-2022`. (Visited on 05/14/2025).

Bellm, Eric C. et al. (Jan. 2019). “The Zwicky Transient Facility: System Overview, Performance, and First Results”. In: *Publications of the Astronomical Society of the Pacific* 131.995, p. 018002. ISSN: 0004-6280, 1538-3873. DOI: `10.1088/1538-3873/aaecbe`. (Visited on 01/26/2020).

Benn, Chris R. and Sara L. Ellison (Sept. 1999). *LA Palma Night-Sky Brightness.* DOI: `10.48550/arXiv.astro-ph/9909153`. arXiv: `astro-ph/9909153`. (Visited on 10/25/2024).

Bento, João, David Heffernan, César Quintana Rivero, Alberto Prieto Antúnez, Adam Garner, David Copley, J. J. Fernández-Valdivia, Javier León Gil, Josué Barrera Martín, and Miguel Torres (July 2024). "A Modern DevOps and Serverless Architecture for the New Robotic Telescope Software Infrastructure". In: *Software and Cyberinfrastructure for Astronomy VIII.* Vol. 13101. SPIE, pp. 562–568. DOI: `10.1117/12.3018820`. (Visited on 10/02/2025).

Bentz, Misty C., Kelly D. Denney, Catherine J. Grier, Aaron J. Barth, Bradley M. Peterson, Marianne Vestergaard, Vardha N. Bennert, Gabriela Canalizo, Gisella De Rosa, Alexei V. Filippenko, Elinor L. Gates, Jenny E. Greene, Weidong Li, Matthew A. Malkan, Richard W. Pogge, Daniel Stern, Tommaso Treu, and Jong-Hak Woo (Apr. 2013). "THE LOW-LUMINOSITY END OF THE RADIUS–LUMINOSITY RELATIONSHIP FOR ACTIVE GALACTIC NUCLEI". In: *The Astrophysical Journal* 767.2, p. 149. ISSN: 0004-637X. DOI: `10.1088/0004-637X/767/2/149`. (Visited on 09/10/2025).

Bertin, E. and S. Arnouts (June 1996). "SExtractor: Software for source extraction." In: *A&AS* 117, pp. 393–404. DOI: `10.1051/aas:1996164`.

Bolelli, Federico, Stefano Allegretti, Lorenzo Baraldi, and Costantino Grana (2020). "Spaghetti Labeling: Directed Acyclic Graphs for Block-Based Connected Components Labeling". In: *IEEE Transactions on Image Processing* 29. tex.ids= bolelli2020a conferenceName: IEEE Transactions on Image Processing, pp. 1999–2012. ISSN: 1941-0042. DOI: `10.1109/TIP.2019.2946979`.

Bolelli, Federico, Stefano Allegretti, and Costantino Grana (2021). "One DAG to Rule Them All". en. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1. ISSN: 0162-8828, 2160-9292, 1939-3539. DOI: `10.1109/TPAMI.2021.3055337`. URL: `https://ieeexplore.ieee.org/document/9339949/` (visited on 02/20/2024).

Bosch, J.L. and J. Kleissl (Sept. 2013). "Cloud Motion Vectors from a Network of Ground Sensors in a Solar Power Plant". In: *Solar Energy* 95, pp. 13–20. ISSN: 0038092X. DOI: `10.1016/j.solener.2013.05.027`. (Visited on 06/17/2025).

Bricman, Katja and Andreja Gomboc (Feb. 2020). "The Prospects of Observing Tidal Disruption Events with the Large Synoptic Survey Telescope". In: *The Astrophysical Journal* 890.1, p. 73. ISSN: 0004-637X, 1538-4357. DOI: `10.3847/1538-4357/ab6989`. (Visited on 09/28/2025).

Broccia, Gianmario (Apr. 2021). *Real-Time Detection and Classification of Astronomical Transient Events: The State-of-the-Art.* DOI: `10.48550/arXiv.2105.00089`. arXiv: `2105.00089 [astro-ph]`. (Visited on 09/28/2025).

Brown, T. M., N. Baliber, F. B. Bianco, M. Bowman, B. Burleson, P. Conway, M. Crellin, É. Depagne, J. De Vera, B. Dilday, D. Dragomir, M. Dubberley, J. D. Eastman, M. Elphick, M. Falarski, S. Foale, M. Ford, B. J. Fulton, J. Garza, E. L. Gomez, M. Graham, R. Greene, B. Haldeman, E. Hawkins, B. Haworth, R. Haynes, M. Hidas, A. E. Hjelstrom, D. A. Howell, J. Hygelund, T. A. Lister, R. Lobdill, J. Martinez, D. S. Mullins, M. Norbury, J. Parrent, R. Paulson, D. L. Petry, A. Pickles, V. Posner, W. E. Rosing, R. Ross, D. J. Sand, E. S. Saunders, J. Shobbrook, A. Shporer, R. A. Street, D. Thomas, Y. Tsapras, J. R. Tufts, S. Valenti, K. Vander Horst, Z. Walker, G. White, and M. Willis (Sept. 2013). "Las Cumbres Observatory Global Telescope Network". en. In: *Publications of the Astronomical Society of the Pacific* 125.931. Publisher: The Astronomical Society of the Pacific, p. 1031. ISSN: 1538-3873. DOI: `10.1086/673168`. URL: `https://dx.doi.org/10.1086/673168` (visited on 07/19/2025).

Buntin, Sebastian (Sept. 2021). "Robotic Scheduling of Time Critical Observations". en. MA thesis. Liverpool John Moores University.

Cackett, Edward M., Chia-Ying Chiang, Ian McHardy, Rick Edelson, Michael R. Goad, Keith Horne, and Kirk T. Korista (Apr. 2018). "Accretion Disk Reverberation with Hubble Space Telescope Observations of NGC 4593: Evidence for Diffuse Continuum Lags". In: *The Astrophysical Journal* 857.1, p. 53. ISSN: 0004-637X, 1538-4357. DOI: `10.3847/1538-4357/aab4f7`. (Visited on 09/28/2025).

Cao, Changyong, Frank J. De Luccia, Xiaoxiong Xiong, Robert Wolfe, and Fuzhong Weng (Feb. 2014). "Early On-Orbit Performance of the Visible Infrared Imaging Radiometer Suite Onboard the Suomi National Polar-Orbiting Partnership (S-NPP) Satellite". In: *IEEE Transactions on Geoscience and Remote Sensing* 52. Publisher: IEEE ADS Bibcode: 2014ITGRS..52.1142C, pp. 1142–1156. ISSN: 0196-2892. DOI: `10.1109/TGRS.2013.2247768`. URL: `https://ui.adsabs.harvard.edu/abs/2014ITGRS..52.1142C` (visited on 07/19/2025).

Carrasco-Davis, Rodrigo, Esteban Reyes, Camilo Valenzuela, Francisco Förster, Pablo A. Estévez, Giuliano Pignata, Franz E. Bauer, Ignacio Reyes, Paula Sánchez-Sáez, Guillermo Cabrera-Vives, Susana Eyheramendy, Márcio Catelan, Javier Arredondo, Ernesto Castillo-Navarrete, Diego Rodríguez-Mancini, Daniela Ruz-Mieres, Alberto Moya, Luis Sabatini-Gacitúa, Cristóbal Sepúlveda-Cobo, Ashish A. Mahabal, Javier Silva-Farfán, Ernesto Camacho-Iñiquez, and Lluís Galbany (June 2021). "Alert Classification for the ALeRCE Broker System: The Real-time Stamp Classifier". In: *arXiv:2008.03309 [astro-ph]*. arXiv: `2008.03309 [astro-ph]`. (Visited on 06/27/2021).

Carrillo, J., J. C. Guerra, E. Cuevas, and J. Barrancos (Feb. 2016). "Characterization of the Marine Boundary Layer and the Trade-Wind Inversion over the Sub-tropical North Atlantic". In: *Boundary-Layer Meteorology* 158.2, pp. 311–330. ISSN: 0006-8314, 1573-1472. DOI: `10.1007/s10546-015-0081-1`. (Visited on 05/14/2025).

Chromey, Frederick R. (2010). *To Measure the Sky: An Introduction to Observational Astronomy.* Cambridge ; New York: Cambridge University Press. ISBN: 978-0-521-76386-8.

Copley, David, Richard Ashley, Stuart D. Bates, Joao Bento, Sebastian Buntin, Chris M. Copperwheat, Adam Garner, Beth Garton, David Heffernan, Helen E. Jermak, David Law, Adrian McGrath, Chloe Miossec, Ali Ranjbar, Robert Smith, Iain A. Steele, Alejandra Goded, Carlos M. M. Gutiérrez, Alberto Prieto, Marta Puga, Cesar Quintana, Rafael Rebolo, Francisco Javier De Cos Juez, Javier Gracia Rodríguez, Cesar Rodríguez Pereira, Fernando Sánchez Lasheras, and Fernando Abarzuza (Sept. 2024). "New Robotic Telescope system progress towards critical design review". en. In: *Ground-based and Airborne Telescopes X.* Ed. by Heather K. Marshall, Jason Spyromilio, and Tomonori Usuda. Yokohama, Japan: SPIE, p. 217. DOI: `10.1117/12.3018999`. URL: `https://www.spiedigitallibrary.org/conference-proceedings-of-spie/13094/3018999/New-Robotic-Telescope-system-progress-towards-critical-design-review/10.1117/12.3018999.full` (visited on 07/19/2025).

Copperwheat, C. M., I. A. Steele, R. M. Barnsley, S. D. Bates, D. Bersier, M. F. Bode, D. Carter, N. R. Clay, C. A. Collins, M. J. Darnley, C. J. Davis, C. M. Gutierrez, D. J. Harman, P. A. James, J. Knapen, S. Kobayashi, J. M. Marchant, P. A. Mazzali, C. J. Mottram, C. G. Mundell, A. Newsam, A. Oscoz, E. Palle, A. Piascik, R. Rebolo, and R. J. Smith (Mar. 2015). "Liverpool Telescope 2: A New Robotic Facility for Rapid Transient Follow-Up". In: *Experimental Astronomy* 39.1, pp. 119–165. ISSN: 0922-6435, 1572-9508. DOI: `10.1007/s10686-015-9447-0`. arXiv: `1410.1731`. (Visited on 01/06/2022).

Copperwheat, C. M., I. A. Steele, A. S. Piascik, D. Bersier, M. F. Bode, C. A. Collins, M. J. Darnley, D. K. Galloway, A. Gomboc, S. Kobayashi, G. P. Lamb, A. J. Levan, P. A. Mazzali, C. G. Mundell, E. Pian, D. Pollacco, D. Steeghs, N. R. Tanvir, K. Ulaczyk, and K. Wiersema (Nov. 2016). "Liverpool Telescope Follow-up of Candidate Electromagnetic Counterparts during the First Run of Advanced LIGO". In: *Monthly Notices of the Royal Astronomical Society* 462.4, pp. 3528–3536. ISSN: 0035-8711. DOI: `10.1093/mnras/stw1849`. (Visited on 08/10/2025).

Cuillandre, Jean-Charles, Eugene A. Magnier, Sidik Isani, Daniel Sabin, Wiley Knight, Simon Kras, and Kamson Lai (Dec. 2002). "CFHT's SkyProbe: a real-time sky-transparency monitor". en. In: *SPIE Proceedings.* Ed. by Peter J. Quinn. Vol. 4844. ISSN: 0277-786X. Waikoloa, Hawai'i, United States: SPIE, p. 501. DOI: `10.1117/12.460613`. URL: `http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.460613` (visited on 07/18/2025).

Dai, Z. G. and T. Lu (July 1998). "Gamma-Ray Burst Afterglows: Effects of Radiative Corrections and Non-Uniformity of the Surrounding Medium". In: *Monthly Notices of* *the Royal Astronomical Society* 298.1, pp. 87–92. ISSN: 0035-8711. DOI: `10.1046/j.1365-8711.1998.01681.x`. (Visited on 09/28/2025).

Darnley, M. J., S. C. Williams, M. F. Bode, M. Henze, J.-U. Ness, A. W. Shafter, K. Hornoch, and V. Votruba (Mar. 2014). "A Remarkable Recurrent Nova in M 31: The Optical Observations". In: *Astronomy & Astrophysics* 563, p. L9. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361/201423411`. (Visited on 08/07/2025).

Dev, Soumyabrata, Florian M. Savoy, Yee Hui Lee, and Stefan Winkler (May 2017). "Nighttime Sky/Cloud Image Segmentation". In: *arXiv:1705.10583 [cs]*. arXiv: `1705.10583 [cs]`. (Visited on 05/09/2020).

Drout, M. R., R. Chornock, A. M. Soderberg, N. E. Sanders, R. McKinnon, A. Rest, R. J. Foley, D. Milisavljevic, R. Margutti, E. Berger, M. Calkins, W. Fong, S. Gezari, M. E. Huber, E. Kankare, R. P. Kirshner, C. Leibler, R. Lunnan, S. Mattila, G. H. Marion, G. Narayan, A. G. Riess, K. C. Roth, D. Scolnic, S. J. Smartt, J. L. Tonry, W. S. Burgett, K. C. Chambers, K. W. Hodapp, R. Jedicke, N. Kaiser, E. A. Magnier, N. Metcalfe, J. S. Morgan, P. A. Price, and C. Waters (Sept. 2014). "RAPIDLY EVOLVING AND LUMINOUS TRANSIENTS FROM PAN-STARRS1". In: *The Astrophysical Journal* 794.1, p. 23. ISSN: 0004-637X. DOI: `10.1088/0004-637X/794/1/23`. (Visited on 09/10/2025).

Duriscoe, Dan M., Sharolyn J. Anderson, Christian B. Luginbuhl, and Kimberly E. Baugh (July 2018). "A simplified model of all-sky artificial sky glow derived from VIIRS Day/Night band data". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 214, pp. 133–145. ISSN: 0022-4073. DOI: `10.1016/j.jqsrt.2018.04.028`. URL: `https://www.sciencedirect.com/science/article/pii/S0022407317308749` (visited on 07/19/2025).

Dyer, Martin J., Vik S. Dhillon, Stuart Littlefair, Danny Steeghs, Krzysztof Ulaczyk, Paul Chote, Duncan Galloway, and Evert Rol (July 2018). "A Telescope Control and Scheduling System for the Gravitational-wave Optical Transient Observer (GOTO)". In: *Observatory Operations: Strategies, Processes, and Systems VII*. Vol. 10704. SPIE, pp. 124–137. DOI: `10.1117/12.2311865`. (Visited on 08/10/2025).

Egret, Daniel, Marc Wenger, and Pascal Dubois (1991). "The SIMBAD astronomical database". In: *Databases & On-line Data in Astronomy*. Springer, pp. 79–88.

Font Tullot, Inocencio (1956). "El tiempo atmosférico en las Islas Canarias". In.

Förster, F., G. Cabrera-Vives, E. Castillo-Navarrete, P. A. Estévez, P. Sánchez-Sáez, J. Arredondo, F. E. Bauer, R. Carrasco-Davis, M. Catelan, F. Elorrieta, S. Eyheramendy, P. Huijse, G. Pignata, E. Reyes, I. Reyes, D. Rodríguez-Mancini, D. Ruz-Mieres, C. Valenzuela, I. Álvarez-Maldonado, N. Astorga, J. Borissova, A. Clocchiatti, D. De Cicco, C. Donoso-Oliva, L. Hernández-García, M. J. Graham, A. Jordán, R. Kurtev, A. Mahabal, J. C. Maureira, A. Muñoz-Arancibia, R. Molina-Ferreiro, A. Moya, W. Palma, M. Pérez-Carrasco, P. Protopapas, M. Romero, L. Sabatini-Gacitua,

A. Sánchez, J. San Martín, C. Sepúlveda-Cobo, E. Vera, and J. R. Vergara (May 2021). "The Automatic Learning for the Rapid Classification of Events (ALeRCE) Alert Broker". In: *The Astronomical Journal* 161, p. 242. ISSN: 0004-6256. DOI: `10.3847/1538-3881/abe9bc`. (Visited on 08/08/2025).

Fraser, Stephen and Iain A. Steele (Sept. 2004). "Robotic Telescope Scheduling: The Liverpool Telescope Experience". In: *SPIE Astronomical Telescopes + Instrumentation*. Kona, United States, p. 331. DOI: `10.1117/12.551380`. (Visited on 08/07/2024).

Fremling, C., A. A. Miller, Y. Sharma, A. Dugas, D. A. Perley, K. Taggart, J. Sollerman, A. Goobar, M. L. Graham, J. D. Neill, J. Nordin, M. Rigault, R. Walters, I. Andreoni, A. Bagdasaryan, J. Belicki, C. Cannella, E. C. Bellm, S. B. Cenko, K. De, R. Dekany, S. Frederick, V. Z. Golkhou, M. J. Graham, G. Helou, A. Y. Q. Ho, M. M. Kasliwal, T. Kupfer, R. R. Laher, A. Mahabal, F. J. Masci, R. Riddle, B. Rusholme, S. Schulze, D. L. Shupe, R. M. Smith, S. van Velzen, Lin Yan, Y. Yao, Z. Zhuang, and S. R. Kulkarni (May 2020a). "The Zwicky Transient Facility Bright Transient Survey. I. Spectroscopic Classification and the Redshift Completeness of Local Galaxy Catalogs". In: *The Astrophysical Journal* 895, p. 32. ISSN: 0004-637X. DOI: `10.3847/1538-4357/ab8943`. (Visited on 08/08/2025).

Fremling, C., A. A. Miller, Y. Sharma, A. Dugas, D. A. Perley, K. Taggart, J. Sollerman, A. Goobar, M. L. Graham, J. D. Neill, J. Nordin, M. Rigault, R. Walters, I. Andreoni, A. Bagdasaryan, J. Belicki, C. Cannella, E. C. Bellm, S. B. Cenko, K. De, R. Dekany, S. Frederick, V. Z. Golkhou, M. J. Graham, G. Helou, A. Y. Q. Ho, M. M. Kasliwal, T. Kupfer, R. R. Laher, A. Mahabal, F. J. Masci, R. Riddle, B. Rusholme, S. Schulze, D. L. Shupe, R. M. Smith, S. Van Velzen, Lin Yan, Y. Yao, Z. Zhuang, and S. R. Kulkarni (May 2020b). "The Zwicky Transient Facility Bright Transient Survey. I. Spectroscopic Classification and the Redshift Completeness of Local Galaxy Catalogs". In: *The Astrophysical Journal* 895.1, p. 32. ISSN: 0004-637X, 1538-4357. DOI: `10.3847/1538-4357/ab8943`. (Visited on 09/28/2025).

Fromm, Michael, George Kablick, and Peter Caffrey (Dec. 2016). "Dust-infused Baroclinic Cyclone Storm Clouds: The Evidence, Meteorology, and Some Implications". In: *Geophysical Research Letters* 43.24. ISSN: 0094-8276, 1944-8007. DOI: `10.1002/2016GL071801`. (Visited on 05/14/2025).

Gaia Collaboration et al. (Nov. 2016). "The Gaia mission". In: *Astronomy and Astrophysics* 595. Publisher: EDP ADS Bibcode: 2016A&A...595A...1G, A1. ISSN: 0004-6361. DOI: `10.1051/0004-6361/201629272`. URL: `https://ui.adsabs.harvard.edu/abs/2016A&A...595A...1G` (visited on 07/19/2025).

Gaia Collaboration et al. (May 2021). "Gaia Early Data Release 3. Summary of the Contents and Survey Properties". In: *Astronomy and Astrophysics* 649, A1. ISSN: 0004-6361. DOI: `10.1051/0004-6361/202039657`. (Visited on 09/28/2025).

Gal-Yam, A. (Jan. 2021). “The TNS Alert System”. In: *American Astronomical Society Meeting Abstracts #237*. Vol. 237, p. 423.05. (Visited on 09/11/2025).

Gal-Yam, Avishay, I. Arcavi, E. O. Ofek, S. Ben-Ami, S. B. Cenko, M. M. Kasliwal, Y. Cao, O. Yaron, D. Tal, J. M. Silverman, A. Horesh, A. De Cia, F. Taddia, J. Sollerman, D. Perley, P. M. Vreeswijk, S. R. Kulkarni, P. E. Nugent, A. V. Filippenko, and J. C. Wheeler (May 2014). “A Wolf–Rayet-like Progenitor of SN 2013cu from Spectral Observations of a Stellar Wind”. In: *Nature* 509.7501, pp. 471–474. ISSN: 1476-4687. DOI: `10.1038/nature13304`. (Visited on 09/10/2025).

Garcia-Gil, A., C. Munoz-Tunon, and A. M. Varela (Sept. 2010). “Atmosphere Extinction at the ORM on La Palma: A 20 Yr Statistical Database Gathered at the Carlsberg Meridian Telescope”. In: *Publications of the Astronomical Society of the Pacific* 122.895, pp. 1109–1121. ISSN: 0004-6280, 1538-3873. DOI: `10.1086/656329`. arXiv: `1009.4056 [astro-ph]`. (Visited on 05/24/2022).

Garstang, R. H. (Mar. 1989). “NIGHT SKY BRIGHTNESS AT OBSERVATORIES AND SITES”. en. In: *Publications of the Astronomical Society of the Pacific* 101.637. Publisher: The Astronomical Society of the Pacific, p. 306. ISSN: 1538-3873. DOI: `10.1086/132436`. URL: `https://dx.doi.org/10.1086/132436` (visited on 10/26/2024).

Gaug, Markus, Alessandro Longo, Stefano Bianchi, Lluís Font, Sofia Almirante, Harald Kornmayer, Michele Doro, Alexander Hahn, Oscar Blanch, Wolfango Plastino, and Daniela Dorner (Nov. 2024). “Detailed Analysis of Local Climate at the CTAO-North Site on La Palma from 20 Yr of MAGIC Weather Station Data”. In: *Monthly Notices of the Royal Astronomical Society* 534.3, pp. 2344–2377. ISSN: 0035-8711. DOI: `10.1093/mnras/stae2214`. (Visited on 05/14/2025).

Gehrels, N., E. Ramirez-Ruiz, and D. B. Fox (Sept. 2009). “Gamma-Ray Bursts in the Swift Era”. In: *Annual Review of Astronomy and Astrophysics* 47.Volume 47, 2009, pp. 567–617. ISSN: 0066-4146, 1545-4282. DOI: `10.1146/annurev.astro.46.060407.145147`. (Visited on 09/10/2025).

Goicoechea, L. J., V. N. Shalyapin, and A. Oscoz (May 2024). “Apparent Correlation between Extrinsic and Intrinsic Flux Variations in the First Gravitationally Lensed Quasar”. In: *Monthly Notices of the Royal Astronomical Society* 530, pp. 2273–2281. ISSN: 0035-8711. DOI: `10.1093/mnras/stae952`. (Visited on 08/07/2025).

Granzer, Thomas, Michael Weber, and Klaus G. Strassmeier (2010). “Three Years of Experience with the STELLA Robotic Observatory”. en. In: *Advances in Astronomy* 2010.1. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2010/980182, p. 980182. ISSN: 1687-7977. DOI: `10.1155/2010/980182`. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1155/2010/980182` (visited on 07/19/2025).

Greisen, Eric W. and Mark R. Calabretta (Dec. 2002). “Representations of world coordinates in FITS”. In: *Astronomy & Astrophysics* 395.3. arXiv: astro-ph/0207407, pp. 1061–1075. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361:20021326`. URL: `http://arxiv.org/abs/astro-ph/0207407` (visited on 02/14/2022).

Groh, Jose H. (Dec. 2014). "Early-Time Spectra of Supernovae and Their Precursor Winds - The Luminous Blue Variable/Yellow Hypergiant Progenitor of SN 2013cu". In: *Astronomy & Astrophysics* 572, p. L11. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361/201424852`. (Visited on 10/02/2025).

Groot, P. J., S. Bloemen, P. M. Vreeswijk, J. C. J. van Roestel, P. G. Jonker, G. Nelemans, M. Klein-Wolt, R. Lepoole, D. L. A. Pieterse, M. Rodenhuis, W. Boland, M. Haverkorn, C. Aerts, R. Bakker, H. Balster, M. Bekema, E. Dijkstra, P. Dolron, E. Elswijk, A. van Elteren, A. Engels, M. Fokker, M. de Haan, F. Hahn, R. ter Horst, D. Lesman, J. Kragt, J. Morren, H. Nillissen, W. Pessemier, G. Raskin, A. de Rijke, L. H. A. Scheers, M. Schuil, S. T. Timmer, L. Antunes Amaral, E. Arancibia-Rojas, I. Arcavi, N. Blagorodnova, S. Biswas, R. P. Breton, H. Dawson, P. Dayal, S. De Wet, C. Duffy, S. Faris, M. Fausnaugh, A. Gal-Yam, S. Geier, A. Horesh, C. Johnston, G. Katusiime, C. Kelley, A. Kosakowski, T. Kupfer, G. Leloudas, A. Levan, D. Modiano, O. Mogawana, J. Munday, J. Paice, F. Patat, I. Pelisoli, G. Ramsay, P. T. Ranaivomanana, R. Ruiz-Carmona, V. Schaffenroth, S. Scaringi, F. Stoppa, R. Street, H. Tranin, M. Uzundag, S. Valenti, M. Veresvarska, M. Vucković, H. C. I. Wichern, R. A. M. J. Wijers, R. A. D. Wijnands, and E. Zimmerman (Nov. 2024). "The BlackGEM Telescope Array. I. Overview". In: *Publications of the Astronomical Society of the Pacific* 136.11, p. 115003. ISSN: 1538-3873. DOI: `10.1088/1538-3873/ad8b6a`. (Visited on 08/10/2025).

Guerrero, M.A., R.J. García-López, R.L.M. Corradi, A. Jiménez, J.J. Fuensalida, J.M. Rodríguez-Espinosa, A. Alonso, M. Centurión, and F. Prada (Nov. 1998). "Extinction over the Canarian Observatories: The Limited Influence of Saharan Dust". In: *New Astronomy Reviews* 42.6-8, pp. 529–532. ISSN: 13876473. DOI: `10.1016/S1387-6473(98)00066-9`. (Visited on 01/07/2024).

Gutiérrez, C. M., D. Arnold, D. Copley, C. M. Copperwheat, E. Harvey, H. Jermak, J. Knapen, A. McGrath, A. Oria, R. Rebolo, I. A. Steele, and M. Torres (Jan. 2019). "The New 4-m Robotic Telescope". In: *Astronomische Nachrichten* 340.1-3, pp. 40–45. ISSN: 0004-6337, 1521-3994. DOI: `10.1002/asna.201913556`. (Visited on 08/25/2024).

Hänel, Andreas, Thomas Posch, Salvador J. Ribas, Martin Aubé, Dan Duriscoe, Andreas Jechow, Zoltán Kollath, Dorien E. Lolkema, Chadwick Moore, Norbert Schmidt, Henk Spoelstra, Günther Wuchterl, and Christopher C.M. Kyba (Jan. 2018). "Measuring night sky brightness: methods and challenges". en. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 205, pp. 278–290. ISSN: 00224073. DOI: `10.1016/j.jqsrt.2017.09.008`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0022407317304442` (visited on 12/29/2023).

Harvey, J W, F Hill, R P Hubbard, J R Kennedy, J W Leibacher, J A Pintar, P A Gilman, R W Noyes, A M Title, J Toomre, R K Ulrich, A Bhatnagar, J A Kennewell, W Marquette, J Patron, and E Yasukawa (1996). “The Global Oscillation Network Group (GONG) Project”. In.

Henden, Arne A., Stephen Levine, Dirk Terrell, Douglas L. Welch, Ulisse Munari, and Brian K. Kloppenborg (June 2018a). “APASS Data Release 10”. In: 232. Conference Name: American Astronomical Society Meeting Abstracts #232 ADS Bibcode: 2018AAS...23222306H, p. 223.06. URL: `https://ui.adsabs.harvard.edu/abs/2018AAS...23222306H` (visited on 02/22/2023).

Henden, Arne A., Stephen Levine, Dirk Terrell, Douglas L. Welch, Ulisse Munari, and Brian K. Kloppenborg (June 2018b). “APASS Data Release 10”. In: *American Astronomical Society Meeting Abstracts #232*. Vol. 232, p. 223.06. (Visited on 05/19/2025).

Ho, Anna Y. Q., Daniel A. Perley, S. R. Kulkarni, Dillon Z. J. Dong, Kishalay De, Poonam Chandra, Igor Andreoni, Eric C. Bellm, Kevin B. Burdge, Michael Coughlin, Richard Dekany, Michael Feeney, Dmitry D. Frederiks, Christoffer Fremling, V. Zach Golkhou, Matthew J. Graham, David Hale, George Helou, Assaf Horesh, Mansi M. Kasliwal, Russ R. Laher, Frank J. Masci, A. A. Miller, Michael Porter, Anna Ridnaia, Ben Rusholme, David L. Shupe, Maayane T. Soumagnac, and Dmitry S. Svinkin (May 2020). “The Koala: A Fast Blue Optical Transient with Luminous Radio Emission from a Starburst Dwarf Galaxy at z = 0.27”. In: *The Astrophysical Journal* 895, p. 49. ISSN: 0004-637X. DOI: `10.3847/1538-4357/ab8bcf`. (Visited on 09/25/2025).

Høg, E., C. Fabricius, V. V. Makarov, S. Urban, T. Corbin, G. Wycoff, U. Bastian, P. Schwekendiek, and A. Wicenec (Mar. 2000). “The Tycho-2 Catalogue of the 2.5 Million Brightest Stars”. In: *Astronomy and Astrophysics* 355, pp. L27–L30. ISSN: 0004-6361. (Visited on 05/19/2025).

Hoischen, C., M. Füßling, S. Ohm, A. Balzer, H. Ashkar, K. Bernlöhr, P. Hofverberg, T. L. Holch, T. Murach, H. Prokoph, F. Schüssler, S. J. Zhu, D. Berge, K. Egberts, and C. Stegmann (Oct. 2022). “The H.E.S.S. Transients Follow-up System”. In: *Astronomy & Astrophysics* 666, A119. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361/202243092`. (Visited on 08/10/2025).

Honeycutt, R. Kent and George W. Turner (Jan. 1992). “Architecture of the Software for the Indiana CCD Automated Telescope”. In: vol. 34. ADS Bibcode: 1992ASPC...34...77H, p. 77. URL: `https://ui.adsabs.harvard.edu/abs/1992ASPC...34...77H` (visited on 08/07/2025).

Horne, K. (June 1986). “An Optimal Extraction Algorithm for CCD Spectroscopy.” In: *Publications of the Astronomical Society of the Pacific* 98, pp. 609–617. ISSN: 0004-6280. DOI: `10.1086/131801`. (Visited on 05/16/2025).

Howell, Steve B (2006). *Handbook of CCD Astronomy, Second Edition*. 2nd ed. Cambridge Observing Handbooks for Research Astronomers. New York: Cambridge University Press. ISBN: 0-521-85215-3.

Ingraham, Patrick J., Andy W. Clements, Tiago Ribeiro, Michael A. Reuter, Merlin Fisher-Levine, Joshua Hoblitt, Robert H. Lupton, Sandrine Thomas, Christopher W. Stubbs, Kirk Arndt, Nick Callahan, CHarles F. Claver, Franco Colleoni, Luis Corral, Peter Doherty, Frossie Economou, Angel Fausti, Tim Jeness, Htut Khine, Simon Krughoff, Nicholas Mondrik, Felipe Menanteau, David J. Mills, William O'Mullane, Kevin A. Reil, Mario Rivera, Brian Stalder, Jacques Sebag, Ian Shipsey, Roberto Tighe, Adam J. Thornton, Andres Villalobos, and Oliver Wiecha (Dec. 2020). "Vera C. Rubin Observatory auxiliary telescope commissioning as a control system pathfinder". en. In: *Software and Cyberinfrastructure for Astronomy VI*. Ed. by Juan C. Guzman and Jorge Ibsen. Online Only, United States: SPIE, p. 31. DOI: `10.1117/12.2561112`. URL: `https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11452/2561112/Vera-C-Rubin-Observatory-auxiliary-telescope-commissioning-as-a-control/10.1117/12.2561112.full` (visited on 07/19/2025).

Ivezić, Željko et al. (Mar. 2019). "LSST: From Science Drivers to Reference Design and Anticipated Data Products". In: *The Astrophysical Journal* 873. Publisher: IOP ADS Bibcode: 2019ApJ...873..111I, p. 111. ISSN: 0004-637X. DOI: `10.3847/1538-4357/ab042c`. URL: `https://ui.adsabs.harvard.edu/abs/2019ApJ...873..111I` (visited on 07/19/2025).

Jia, Peng, Qiang Liu, and Yongyang Sun (Apr. 2020). "Detection and Classification of Astronomical Targets with Deep Neural Networks in Wide-field Small Aperture Telescopes". In: *The Astronomical Journal* 159.5, p. 212. ISSN: 1538-3881. DOI: `10.3847/1538-3881/ab800a`. (Visited on 06/27/2021).

Kaiser, Nicholas, Herve Aussel, Barry E. Burke, Hans Boesgaard, Ken Chambers, Mark Richard Chun, James N. Heasley, Klaus-Werner Hodapp, Bobby Hunt, Robert Jedicke, D. Jewitt, Rolf Kudritzki, Gerard Anthony Luppino, Michael Maberry, Eugene Magnier, David G. Monet, Peter M. Onaka, Andrew J. Pickles, Pui Hin H. Rhoads, Theodore Simon, Alexander Szalay, Istvan Szapudi, David J. Tholen, John L. Tonry, Mark Waterson, and John Wick (Dec. 2002). "Pan-STARRS: A Large Synoptic Survey Telescope Array". In: *Survey and Other Telescope Technologies and Discoveries*. Vol. 4836. SPIE, pp. 154–164. DOI: `10.1117/12.457365`. (Visited on 08/08/2025).

Kalman, R. E. (Mar. 1960). "A New Approach to Linear Filtering and Prediction Problems". en. In: *Journal of Basic Engineering* 82.1, pp. 35–45. ISSN: 0021-9223. DOI: `10.1115/1.3662552`. URL: `https://asmedigitalcollection.asme.org/fluidsengineering/article/82/1/35/397706/A-New-Approach-to-Linear-Filtering-and-Prediction` (visited on 03/02/2023).

Kann, D. A., S. Klose, B. Zhang, D. Malesani, E. Nakar, A. Pozanenko, A. C. Wilson, N. R. Butler, P. Jakobsson, S. Schulze, M. Andreev, L. A. Antonelli, I. F. Bikmaev, V. Biryukov, M. Böttcher, R. A. Burenin, J. M. Castro Cerón, A. J. Castro-Tirado, G. Chincarini, B. E. Cobb, S. Covino, P. D'Avanzo, V. D'Elia, M. Della Valle, A. de Ugarte Postigo, Yu. Efimov, P. Ferrero, D. Fugazza, J. P. U. Fynbo, M. Gålfalk, F. Grundahl, J. Gorosabel, S. Gupta, S. Guziy, B. Hafizov, J. Hjorth, K. Holhjem, M. Ibrahimov, M. Im, G. L. Israel, M. Jelínek, B. L. Jensen, R. Karimov, I. M. Khamitov, Ü. Kiziloğlu, E. Klunko, P. Kubánek, A. S. Kutyrev, P. Laursen, A. J. Levan, F. Mannucci, C. M. Martin, A. Mescheryakov, N. Mirabal, J. P. Norris, J.-E. Ovaldsen, D. Paraficz, E. Pavlenko, S. Piranomonte, A. Rossi, V. Rumyantsev, R. Salinas, A. Sergeev, D. Sharapov, J. Sollerman, B. Stecklum, L. Stella, G. Tagliaferri, N. R. Tanvir, J. Telting, V. Testa, A. C. Updike, A. Volnova, D. Watson, K. Wiersema, and D. Xu (Aug. 2010). "THE AFTERGLOWS OF SWIFT-ERA GAMMA-RAY BURSTS. I. COMPARING PRE-SWIFT AND SWIFT-ERA LONG/SOFT (TYPE II) GRB OPTICAL AFTERGLOWS*". In: *The Astrophysical Journal* 720.2, p. 1513. ISSN: 0004-637X. DOI: `10.1088/0004-637X/720/2/1513`. (Visited on 09/10/2025).

Kasliwal, M. M., C. Cannella, A. Bagdasaryan, T. Hung, U. Feindt, L. P. Singer, M. Coughlin, C. Fremling, R. Walters, D. Duev, R. Itoh, and R. M. Quimby (Feb. 2019). "The GROWTH Marshal: A Dynamic Science Portal for Time-domain Astronomy". In: *Publications of the Astronomical Society of the Pacific* 131.997, p. 038003. ISSN: 1538-3873. DOI: `10.1088/1538-3873/aafbc2`. (Visited on 08/08/2025).

Kasten, Fritz and Andrew T. Young (Nov. 1989). "Revised Optical Air Mass Tables and Approximation Formula". In: *Applied Optics* 28, pp. 4735–4738. ISSN: 0003-6935. DOI: `10.1364/AO.28.004735`. (Visited on 05/14/2025).

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu (2017). "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc. URL: `https://proceedings.neurips.cc/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html` (visited on 07/20/2025).

Kiepenheuer, K. O. (Jan. 1972). "Fortgang Des Projektes JOSO (Joint Organization for Solar Observation)". In: *Mitteilungen der Astronomischen Gesellschaft Hamburg* 31, p. 147. ISSN: 0374-1958. (Visited on 05/14/2025).

King, D L (Sept. 1985). "Atmospheric Extinction at the Roque de Los Muchachos Observatory, La Palma". In: p. 5.

Kornilov, V., N. Shatsky, O. Voziakova, B. Safonov, S. Potanin, and M. Kornilov (Oct. 2010). "First results of site testing program at Mt. Shatdzhatmaz in 2007 - 2009". In: *Monthly Notices of the Royal Astronomical Society* 408.2. arXiv:1006.3528 [astro-ph], pp. 1233–1248. ISSN: 0035-8711. DOI: `10.1111/j.1365-2966.2010.17203.x`. URL: `http://arxiv.org/abs/1006.3528` (visited on 07/19/2025).

Krisciunas, K. (Oct. 1997). “OPTICAL NIGHT-SKY BRIGHTNESS AT MAUNA KEA OVER THE COURSE OF A COMPLETE SUNSPOT CYCLE”. In: *Publications of the Astronomical Society of the Pacific* 109.740, p. 1181. ISSN: 1538-3873. DOI: `10.1086/133993`. (Visited on 03/12/2026).

Krisciunas, Kevin and Bradley E. Schaefer (Sept. 1991). “A MODEL OF THE BRIGHTNESS OF MOONLIGHT”. In: *Publications of the Astronomical Society of the Pacific* 103.667, p. 1033. ISSN: 1538-3873. DOI: `10.1086/132921`. (Visited on 10/25/2024).

Kulkarni, S. R. (Apr. 2020). *Towards An Integrated Optical Transient Utility.* DOI: `10.48550/arXiv.2004.03511`. arXiv: `2004.03511 [astro-ph]`. (Visited on 08/07/2025).

Laken, Benjamin A., Hannu Parviainen, Enric Pallé, and Tariq Shahbaz (Apr. 2014). “Saharan Mineral Dust Outbreaks Observed over the North Atlantic Island of La Palma in Summertime between 1984 and 2012”. In: *Quarterly Journal of the Royal Meteorological Society* 140.680, pp. 1058–1068. ISSN: 0035-9009, 1477-870X. DOI: `10.1002/qj.2170`. arXiv: `1306.3755 [physics]`. (Visited on 01/07/2024).

Landolt, A. U. (Mar. 1983). “UBVRI Photometric standard stars around the celestial equator.” In: *The Astronomical Journal* 88. Publisher: IOP ADS Bibcode: 1983AJ.....88..439L, pp. 439–460. ISSN: 0004-6256. DOI: `10.1086/113329`. URL: `https://ui.adsabs.harvard.edu/abs/1983AJ.....88..439L` (visited on 07/19/2025).

Landolt, Arlo U. (July 1992). “UBVRI Photometric Standard Stars in the Magnitude Range 11.5 < V < 16.0 Around the Celestial Equator”. In: *The Astronomical Journal* 104. Publisher: IOP ADS Bibcode: 1992AJ....104..340L, p. 340. ISSN: 0004-6256. DOI: `10.1086/116242`. URL: `https://ui.adsabs.harvard.edu/abs/1992AJ....104..340L` (visited on 07/19/2025).

Lang, Dustin, David W. Hogg, Keir Mierle, Michael Blanton, and Sam Roweis (May 2010). “Astrometry.Net: Blind Astrometric Calibration of Arbitrary Astronomical Images”. In: *The Astronomical Journal* 139, pp. 1782–1800. ISSN: 0004-6256. DOI: `10.1088/0004-6256/139/5/1782`. (Visited on 03/12/2024).

Laskar, Tanmoy, Kate D. Alexander, Raffaella Margutti, Tarraneh Eftekhari, Ryan Chornock, Edo Berger, Yvette Cendes, Anne Duerr, Daniel A. Perley, Maria Edvige Ravasio, Ryo Yamazaki, Eliot H. Ayache, Thomas Barclay, Rodolfo Barniol Duran, Shivani Bhandari, Daniel Brethauer, Collin T. Christy, Deanne L. Coppejans, Paul Duffell, Wen-fai Fong, Andreja Gomboc, Cristiano Guidorzi, Jamie A. Kennea, Shiho Kobayashi, Andrew Levan, Andrei P. Lobanov, Brian D. Metzger, Eduardo Ros, Genevieve Schroeder, and P. K. G. Williams (Mar. 2023). “The Radio to GeV Afterglow of GRB 221009A”. In: *The Astrophysical Journal Letters* 946.1, p. L23. ISSN: 2041-8205, 2041-8213. DOI: `10.3847/2041-8213/acbfad`. (Visited on 04/08/2023).

Law, Nicholas M., Shrinivas R. Kulkarni, Richard G. Dekany, Eran O. Ofek, Robert M. Quimby, Peter E. Nugent, Jason Surace, Carl C. Grillmair, Joshua S. Bloom, Mansi M. Kasliwal, Lars Bildsten, Tim Brown, S. Bradley Cenko, David Ciardi, Ernest Croner,

S. George Djorgovski, Julian van Eyken, Alexei V. Filippenko, Derek B. Fox, Avishay Gal-Yam, David Hale, Nouhad Hamam, George Helou, John Henning, D. Andrew Howell, Janet Jacobsen, Russ Laher, Sean Mattingly, Dan McKenna, Andrew Pickles, Dovi Poznanski, Gustavo Rahmer, Arne Rau, Wayne Rosing, Michael Shara, Roger Smith, Dan Starr, Mark Sullivan, Viswa Velur, Richard Walters, and Jeff Zolkower (Dec. 2009). “The Palomar Transient Factory: System Overview, Performance, and First Results”. In: *Publications of the Astronomical Society of the Pacific* 121, p. 1395. ISSN: 0004-6280. DOI: `10.1086/648598`. (Visited on 09/03/2021).

Leinert, Ch. and K. Mattila (Jan. 1998). “Natural Optical Sky Background”. In: vol. 139. ADS Bibcode: 1998ASPC..139...17L, p. 17. URL: `https://ui.adsabs.harvard.edu/abs/1998ASPC..139...17L` (visited on 07/19/2025).

Ley 31/1988, Ley del Cielo (1988). *Protección de la Calidad Astronómica de los Observatorios del Instituto de Astrofísica de Canarias*. Published in Boletín Oficial del Estado, No. 262, 1 November 1988. URL: `https://www.boe.es/buscar/doc.php?id=BOE-A-1988-24790`.

Lipunov, V., V. Kornilov, E. Gorbovskoy, A. Belinski, D. Kuvshinov, N. Tyurina, A. Krylov, N. Shatsky, P. Balanutsa, V. Chazov, A. Kuznetsov, D. Zimnuhov, A. Tlatov, A. Parkhomenko, D. Dormidontov, V. Krushinsky, I. Zalozhnyh, A. Popov, S. Yazev, N. Budnev, K. Ivanov, E. Konstantinov, O. Gress, O. Chvalaev, V. Yurkov, Y. Sergienko, and I. Kudelina (Jan. 2012). “MASTER Global Robotic Net”. In: *Astronomical Society of India Conference Series*. Vol. 7, p. 275. (Visited on 08/08/2025).

Lombardi, G., V. Zitelli, S. Ortolani, and M. Pedani (Aug. 2006). “El Roque de Los Muchachos Site Characteristics. I. Temperature Analysis”. In: *Publications of the Astronomical Society of the Pacific* 118, pp. 1198–1204. ISSN: 0004-6280. DOI: `10.1086/507344`. (Visited on 05/13/2025).

Lombardi, G., V. Zitelli, S. Ortolani, and M. Pedani (Mar. 2007). “El Roque de Los Muchachos Site Characteristics. II. Analysis of Wind, Relative Humidity, and Air Pressure”. In: *Publications of the Astronomical Society of the Pacific* 119, pp. 292–302. ISSN: 0004-6280. DOI: `10.1086/513079`. (Visited on 05/13/2025).

Lombardi, G., V. Zitelli, S. Ortolani, M. Pedani, and A. Ghedina (May 2008). “El Roque de Los Muchachos Site Characteristics. III. Analysis of Atmospheric Dust and Aerosol Extinction”. In: *Astronomy and Astrophysics* 483, pp. 651–659. ISSN: 0004-6361. DOI: `10.1051/0004-6361:20078372`. (Visited on 05/13/2025).

Longo, Alessandro, Stefano Bianchi, and Wolfango Plastino (June 2019). “Tvf-EMD Based Time Series Analysis of 7Be Sampled at the CTBTO-IMS Network”. In: *Physica A: Statistical Mechanics and its Applications* 523, pp. 908–914. ISSN: 0378-4371. DOI: `10.1016/j.physa.2019.04.111`. (Visited on 05/14/2025).

Maghrabi, A., R. Clay, N. Wild, and B. Dawson (Nov. 2009). “Design and development of a simple infrared monitor for cloud detection”. en. In: *Energy Conversion* *and Management* 50.11. Publisher: Elsevier BV, pp. 2732–2737. ISSN: 0196-8904. DOI: `10.1016/j.enconman.2009.06.029`. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0196890409002441` (visited on 07/18/2025).

Marquez, Ricardo and Carlos F.M. Coimbra (May 2013). “Intra-Hour DNI Forecasting Based on Cloud Tracking Image Analysis”. In: *Solar Energy* 91, pp. 327–336. ISSN: 0038092X. DOI: `10.1016/j.solener.2012.09.018`. (Visited on 09/06/2022).

Matheson, Thomas, Carl Stubens, Nicholas Wolf, Chien-Hsiu Lee, Gautham Narayan, Abhijit Saha, Adam Scott, Monika Soraisam, Adam S. Bolton, Benjamin Hauger, David R. Silva, John Kececioglu, Carlos Scheidegger, Richard Snodgrass, Patrick D. Aleo, Eric Evans-Jacquez, Navdeep Singh, Zhe Wang, Shuo Yang, and Zhenge Zhao (Feb. 2021). “The ANTARES Astronomical Time-domain Event Broker”. In: *The Astronomical Journal* 161.3, p. 107. ISSN: 1538-3881. DOI: `10.3847/1538-3881/abd703`. (Visited on 08/08/2025).

Mawson, N. R., I. A. Steele, and R. J. Smith (Aug. 2013). “STILT: System Design & Performance”. In: *Astronomische Nachrichten* 334.7, pp. 729–737. ISSN: 0004-6337, 1521-3994. DOI: `10.1002/asna.201311907`. arXiv: `1305.0573 [astro-ph]`. (Visited on 12/14/2023).

McHardy, I M, M Beard, E Breedt, J H Knapen, F M Vincentelli, M Veresvarska, V S Dhillon, T R Marsh, S P Littlefair, K Horne, R Glew, M R Goad, E Kammoun, and D Emmanoulopoulos (Mar. 2023). “First detection of the outer edge of an AGN accretion disc: very fast multiband optical variability of NGC 4395 with GTC/HiPERCAM and LT/IO:O”. In: *Monthly Notices of the Royal Astronomical Society* 519.3, pp. 3366–3382. ISSN: 0035-8711. DOI: `10.1093/mnras/stac3651`. URL: `https://doi.org/10.1093/mnras/stac3651` (visited on 08/08/2025).

McHardy, I. M., E. Koerding, C. Knigge, P. Uttley, and R. P. Fender (Dec. 2006). “Active Galactic Nuclei as Scaled-up Galactic Black Holes”. In: *Nature* 444.7120, pp. 730–732. ISSN: 1476-4687. DOI: `10.1038/nature05389`. (Visited on 09/10/2025).

McNally, AP and PD Watts (2003). “A cloud detection algorithm for high-spectral-resolution infrared sounders”. In: *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* 129.595, pp. 3411–3423.

Meeus, Jean (1998). *Astronomical algorithms.* 2nd ed. Richmond, Va: Willmann-Bell. ISBN: 978-0-943396-61-3.

Mommert, Michael (Mar. 2020). “Cloud Identification from All-sky Camera Data with Machine Learning”. In: *The Astronomical Journal* 159.4, p. 178. ISSN: 1538-3881. DOI: `10.3847/1538-3881/ab744f`. arXiv: `2003.11109`. (Visited on 05/29/2020).

Mottram, CJ and SN Fraser (2008). “Robonet-1.0”. In: *Astronomische Nachrichten: Astronomical Notes* 329.3, pp. 317–320.

Mundell, C. G., D. Kopač, D. M. Arnold, I. A. Steele, A. Gomboc, S. Kobayashi, R. M. Harrison, R. J. Smith, C. Guidorzi, F. J. Virgili, A. Melandri, and J. Japelj (Dec. 2013). “Highly Polarized Light from Stable Ordered Magnetic Fields in GRB 120308A”. In: *Nature* 504, pp. 119–121. ISSN: 0028-0836. DOI: `10.1038/nature12814`. (Visited on 10/02/2025).

Munoz-Tunón, Casiana, Antonia M Varela, and Jesús J Fuensalida (2007). “Recent results at the Canarian Observatories”. In: *Revista Mexicana de Astronomía y Astrofísica* 31, pp. 36–46.

Muñoz-Tuñón, C, A M Varela, and J A Castro-Almazán (Apr. 2015). “Astronomical Site Characterization at the Canarian Observatories”. In: *Journal of Physics: Conference Series* 595, p. 012042. ISSN: 1742-6588, 1742-6596. DOI: `10.1088/1742-6596/595/1/012042`. (Visited on 05/14/2025).

Murdin, Paul (Jan. 1985). “Nighttime Skies above the Canary Islands”. In: *Vistas in Astronomy* 28, pp. 449–465. ISSN: 00836656. DOI: `10.1016/0083-6656(85)90069-8`. (Visited on 01/07/2024).

Newberry, Michael V. (Jan. 1991). “SIGNAL-TO-NOISE CONSIDERATIONS FOR SKY-SUBTRACTED CCD DATA”. In: *Publications of the Astronomical Society of the Pacific* 103.659, p. 122. ISSN: 1538-3873. DOI: `10.1086/132801`. (Visited on 05/16/2025).

Noll, S., W. Kausch, M. Barden, A. M. Jones, C. Szyszka, and S. Kimeswenger (July 2013). *The Cerro Paranal Advanced Sky Model*. Tech. rep. Publication Title: VLT-MAN-ESO-19550-5339 ADS Bibcode: 2013asm..rept.....N. URL: `https://ui.adsabs.harvard.edu/abs/2013asm..rept.....N` (visited on 07/19/2025).

Otsu, Nobuyuki (Jan. 1979). “A Threshold Selection Method from Gray-Level Histograms”. en. In: *IEEE Transactions on Systems, Man, and Cybernetics* 9.1, pp. 62–66. DOI: `10.1109/TSMC.1979.4310076`.

Padmanabhan, Nikhil, David J. Schlegel, Douglas P. Finkbeiner, J. C. Barentine, Michael R. Blanton, Howard J. Brewington, James E. Gunn, Michael Harvanek, David W. Hogg, Željko Ivezić, David Johnston, Stephen M. Kent, S. J. Kleinman, Gillian R. Knapp, Jurek Krzesinski, Dan Long, Eric H. Neilsen Jr., Atsuko Nitta, Craig Loomis, Robert H. Lupton, Sam Roweis, Stephanie A. Snedden, Michael A. Strauss, and Douglas L. Tucker (Feb. 2008). “An Improved Photometric Calibration of the Sloan Digital Sky Survey Imaging Data”. en. In: *The Astrophysical Journal* 674.2, p. 1217. ISSN: 0004-637X. DOI: `10.1086/524677`. URL: `https://dx.doi.org/10.1086/524677` (visited on 07/19/2025).

Pedani, Marco (Oct. 2004). “Light Pollution at the Roque de Los Muchachos Observatory”. In: *New Astronomy* 9.8, pp. 641–650. ISSN: 1384-1076. DOI: `10.1016/j.newast.2004.07.001`. (Visited on 05/13/2025).

Peng, Zhenzhou, Dantong Yu, Dong Huang, John Heiser, Shinjae Yoo, and Paul Kalb (Aug. 2015). "3D Cloud Detection and Tracking System for Solar Forecast Using Multiple Sky Imagers". In: *Solar Energy* 118, pp. 496–519. ISSN: 0038092X. DOI: `10.1016/j.solener.2015.05.037`. (Visited on 09/06/2022).

Pennypacker, C., M. Boer, R. Denny, F. V. Hessman, J. Aymon, N. Duric, S. Gordon, D. Barnaby, G. Spear, and V. Hoette (Nov. 2002). "RTML – a Standard for Use of Remote Telescopes: Enabling Ubiquitous Use of Remote Telescopes". In: *Astronomy & Astrophysics* 395.2, pp. 727–731. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361:20021318`. (Visited on 12/07/2024).

Pereira, Wellesley Ernesto (2003). *The CONCAM global sky monitoring network: Its evolution and the creation of a performance model.* Michigan Technological University.

Pérez-Ramírez, D., R. J. Nemiroff, and J. B. Rafert (Oct. 2004). "nightskylive.net: The Night Sky Live project". en. In: *Astronomische Nachrichten* 325.6-8. Publisher: Wiley, pp. 568–570. ISSN: 0004-6337, 1521-3994. DOI: `10.1002/asna.200410292`. URL: `https://onlinelibrary.wiley.com/doi/10.1002/asna.200410292` (visited on 07/18/2025).

Perley, Daniel A, Anna Y Q Ho, Yuhan Yao, Christoffer Fremling, Joseph P Anderson, Steve Schulze, Harsh Kumar, G C Anupama, Sudhanshu Barway, Eric C Bellm, Varun Bhalerao, Ting-Wan Chen, Dmitry A Duev, Lluís Galbany, Matthew J Graham, Mariusz Gromadzki, Claudia P Gutiérrez, Nada Ihanec, Cosimo Inserra, Mansi M Kasliwal, Erik C Kool, S R Kulkarni, Russ R Laher, Frank J Masci, James D Neill, Matt Nicholl, Miika Pursiainen, Joannes van Roestel, Yashvi Sharma, Jesper Sollerman, Richard Walters, and Philip Wiseman (Dec. 2021). "Real-Time Discovery of AT2020xnd: A Fast, Luminous Ultraviolet Transient with Minimal Radioactive Ejecta". In: *Monthly Notices of the Royal Astronomical Society* 508.4, pp. 5138–5147. ISSN: 0035-8711. DOI: `10.1093/mnras/stab2785`. (Visited on 09/25/2025).

Perley, Daniel A., Paolo A. Mazzali, Lin Yan, S. Bradley Cenko, Suvi Gezari, Kirsty Taggart, Nadia Blagorodnova, Christoffer Fremling, Brenna Mockler, Avinash Singh, Nozomu Tominaga, Masaomi Tanaka, Alan M. Watson, Tomás Ahumada, G. C. Anupama, Chris Ashall, Rosa L. Becerra, David Bersier, Varun Bhalerao, Joshua S. Bloom, Nathaniel R. Butler, Chris Copperwheat, Michael W. Coughlin, Kishalay De, Andrew J. Drake, Dmitry A. Duev, Sara Frederick, J. Jesús González, Ariel Goobar, Marianne Heida, Anna Y. Q. Ho, John Horst, Tiara Hung, Ryosuke Itoh, Jacob E. Jencson, Mansi M. Kasliwal, Nobuyuki Kawai, Shrinivas R. Kulkarni, Brajesh Kumar, Harsh Kumar, Alexander S. Kutyrev, Tanazza Khanam, William H. Lee, Keiichi Maeda, Ashish Mahabal, Katsuhiro L. Murata, James D. Neill, Chow-Choong Ngeow, Bryan Penprase, Elena Pian, Robert Quimby, Enrico Ramirez-Ruiz, Michael Richer, Carlos G. Román-Zúñiga, Shubham Srivastava, Quentin Socia, Jesper Sollerman, Yutaro Tachibana, Francesco Taddia, Samaporn Tinyanont, Eleonora Troja, Charlotte Ward, and Jerrick Wee (Mar. 2019). “The Fast, Luminous Ultraviolet Transient AT2018cow: Extreme Supernova, or Disruption of a Star by an Intermediate-Mass Black Hole?” In: *Monthly Notices of the Royal Astronomical Society* 484.1, pp. 1031–1049. ISSN: 0035-8711, 1365-2966. DOI: `10.1093/mnras/sty3420`. arXiv: `1808.00969 [astro-ph]`. (Visited on 08/07/2025).

Peterson, Bradley M. (Mar. 1993). “REVERBERATION MAPPING OF ACTIVE GALACTIC NUCLEI”. In: *Publications of the Astronomical Society of the Pacific* 105.685, p. 247. ISSN: 1538-3873. DOI: `10.1086/133140`. (Visited on 09/10/2025).

Petrachenko, William, Dirk Behrend, Hayo Hase, Chopo Ma, Arthur Niell, Harald Schuh, and Alan Whitney (Apr. 2013). “The VLBI2010 Global Observing System (VGOS)”. In: ADS Bibcode: 2013EGUGA..1512867P, EGU2013–12867. URL: `https://ui.adsabs.harvard.edu/abs/2013EGUGA..1512867P` (visited on 07/19/2025).

Pian, E., P. D’Avanzo, S. Benetti, M. Branchesi, E. Brocato, S. Campana, E. Cappellaro, S. Covino, V. D’Elia, J. P. U. Fynbo, F. Getman, G. Ghirlanda, G. Ghisellini, A. Grado, G. Greco, J. Hjorth, C. Kouveliotou, A. Levan, L. Limatola, D. Malesani, P. A. Mazzali, A. Melandri, P. Møller, L. Nicastro, E. Palazzi, S. Piranomonte, A. Rossi, O. S. Salafia, J. Selsing, G. Stratta, M. Tanaka, N. R. Tanvir, L. Tomasella, D. Watson, S. Yang, L. Amati, L. A. Antonelli, S. Ascenzi, M. G. Bernardini, M. Boër, F. Bufano, A. Bulgarelli, M. Capaccioli, P. Casella, A. J. Castro-Tirado, E. Chassande-Mottin, R. Ciolfi, C. M. Copperwheat, M. Dadina, G. De Cesare, A. Di Paola, Y. Z. Fan, B. Gendre, G. Giuffrida, A. Giunta, L. K. Hunt, G. L. Israel, Z.-P. Jin, M. M. Kasliwal, S. Klose, M. Lisi, F. Longo, E. Maiorano, M. Mapelli, N. Masetti, L. Nava, B. Patricelli, D. Perley, A. Pescalli, T. Piran, A. Possenti, L. Pulone, M. Razzano, R. Salvaterra, P. Schipani, M. Spera, A. Stamerra, L. Stella, G. Tagliaferri, V. Testa, E. Troja, M. Turatto, S. D. Vergani, and D. Vergani (Nov. 2017). “Spectroscopic Identification of R-Process Nucleosynthesis in a Double Neutron-Star Merger”. In: *Nature* 551.7678, pp. 67–70. ISSN: 1476-4687. DOI: `10.1038/nature24298`. (Visited on 09/10/2025).

Piascik, A. S., Iain A. Steele, Stuart D. Bates, Christopher J. Mottram, R. J. Smith, R. M. Barnsley, and B. Bolton (July 2014). “SPRAT: Spectrograph for the Rapid Acquisition of Transients”. In: 9147. Conference Name: Ground-based and Airborne Instrumentation for Astronomy V ADS Bibcode: 2014SPIE.9147E..8HP, 91478H. DOI: `10.1117/12.2055117`. URL: `https://ui.adsabs.harvard.edu/abs/2014SPIE.9147E..8HP` (visited on 12/29/2023).

Pickering, Keith A. (Sept. 2002). “The Southern Limits of the Ancient Star Catalog and the Commentary of Hipparchos”. In: *DIO* 12, pp. 3–27. (Visited on 05/14/2025).

Pursiainen, M, M Childress, M Smith, S Prajs, M Sullivan, T M Davis, R J Foley, J Asorey, J Calcino, D Carollo, C Curtin, C B D’Andrea, K Glazebrook, C Gutierrez, S R Hinton, J K Hoormann, C Inserra, R Kessler, A King, K Kuehn, G F Lewis, C Lidman, E Macaulay, A Möller, R C Nichol, M Sako, N E Sommer, E Swann, B E

Tucker, S A Uddin, P Wiseman, B Zhang, T M C Abbott, F B Abdalla, S Allam, J Annis, S Avila, D Brooks, E Buckley-Geer, D L Burke, A Carnero Rosell, M Carrasco Kind, J Carretero, F J Castander, C E Cunha, C Davis, J De Vicente, H T Diehl, P Doel, T F Eifler, B Flaugher, P Fosalba, J Frieman, J García-Bellido, D Gruen, R A Gruendl, G Gutierrez, W G Hartley, D L Hollowood, K Honscheid, D J James, T Jeltema, N Kuropatkin, T S Li, M Lima, M A G Maia, P Martini, F Menanteau, R L C Ogando, A A Plazas, A Roodman, E Sanchez, V Scarpine, R Schindler, R C Smith, M Soares-Santos, F Sobreira, E Suchyta, M E C Swanson, G Tarle, D L Tucker, A R Walker, and (DES Collaboration) (Nov. 2018). "Rapidly Evolving Transients in the Dark Energy Survey". In: *Monthly Notices of the Royal Astronomical Society* 481.1, pp. 894–917. ISSN: 0035-8711. DOI: `10.1093/mnras/sty2309`. (Visited on 09/10/2025).

Pursiainen, M, T L Killestein, H Kuncarayakti, P Charalampopoulos, B Warwick, J Lyman, R Kotak, G Leloudas, D Coppejans, T Kravtsov, K Maeda, T Nagao, K Taguchi, K Ackley, V S Dhillon, D K Galloway, A Kumar, D O'Neill, G Ramsay, and D Steeghs (Mar. 2025). "Optical Evolution of AT 2024wpp: The High-Velocity Outflows in Cow-like Transients Are Consistent with High Spherical Symmetry". In: *Monthly Notices of the Royal Astronomical Society* 537.4, pp. 3298–3309. ISSN: 0035-8711. DOI: `10.1093/mnras/staf232`. (Visited on 09/10/2025).

Rabaza, O, David Galadí-Enríquez, A Espín Estrella, and F Aznar Dols (2010). "All-sky brightness monitoring of light pollution with astronomical methods". In: *Journal of environmental management* 91.6, pp. 1278–1287.

Ramanath, Rajeev, Wesley E. Snyder, Griff L. Bilbro, and William A. Sander (2002). "Demosaicking methods for Bayer color arrays". en. In: *Journal of Electronic Imaging* 11.3. Publisher: SPIE-Intl Soc Optical Eng, p. 306. ISSN: 1017-9909. DOI: `10.1117/1.1484495`. URL: `http://electronicimaging.spiedigitallibrary.org/article.aspx?doi=10.1117/1.1484495` (visited on 07/18/2025).

Rau, Arne, Shrinivas R. Kulkarni, Nicholas M. Law, Joshua S. Bloom, David Ciardi, George S. Djorgovski, Derek B. Fox, Avishay Gal-Yam, Carl C. Grillmair, Mansi M. Kasliwal, Peter E. Nugent, Eran O. Ofek, Robert M. Quimby, William T. Reach, Michael Shara, Lars Bildsten, S. Bradley Cenko, Andrew J. Drake, Alexei V. Filippenko, David J. Helfand, George Helou, D. Andrew Howell, Dovi Poznanski, and Mark Sullivan (Oct. 2009). "Exploring the Optical Transient Sky with the Palomar Transient Factory". In: *Publications of the Astronomical Society of the Pacific* 121.886, p. 1334. ISSN: 1538-3873. DOI: `10.1086/605911`. (Visited on 09/10/2025).

Rothchild, Daniel, Christopher Stubbs, and Peter Yoachim (Nov. 2019). "ALTSched: Improved Scheduling for Time-domain Science with LSST". In: *Publications of the Astronomical Society of the Pacific* 131.1005, p. 115002. ISSN: 0004-6280, 1538-3873. DOI: `10.1088/1538-3873/ab3300`. (Visited on 08/08/2025).

Sánchez de Miguel, A., M. Aubé, J. Zamorano, M. Kocifaj, J. Roby, and C. Tapia (June 2017). "Sky Quality Meter measurements in a colour-changing world". In: *Monthly Notices of the Royal Astronomical Society* 467.3, pp. 2966–2979. ISSN: 0035-8711. DOI: `10.1093/mnras/stx145`. URL: `https://doi.org/10.1093/mnras/stx145` (visited on 07/19/2025).

Sauvola, J. and M. Pietikäinen (Feb. 2000). "Adaptive document image binarization". en. In: *Pattern Recognition* 33.2, pp. 225–236. ISSN: 00313203. DOI: `10.1016/S0031-3203(99)00055-2`.

Schartner, Matthias, Bill Petrachenko, Mike Titus, Hana Krásná, John Barrett, Dan Hoak, Dhiman Mondal, Ming Hui Xu, and Benedikt Soja (May 2025). "Optimizing VGOS observations using an SNR-based scheduling approach". en. In: *Earth, Planets and Space* 77.1, p. 61. ISSN: 1880-5981. DOI: `10.1186/s40623-025-02158-0`. URL: `https://doi.org/10.1186/s40623-025-02158-0` (visited on 07/19/2025).

Schipani, P., S. Campana, R. Claudi, H. U. Käufl, M. Accardo, M. Aliverti, A. Baruffolo, S. Ben Ami, F. Biondi, A. Brucalassi, G. Capasso, R. Cosentino, F. D'Alessio, P. D'Avanzo, O. Hershko, D. Gardiol, H. Kuncarayacti, M. Munari, A. Rubin, S. Scuderi, F. Vitali, J. Achrén, J. Antonio Araiza-Duran, I. Arcavi, A. Bianco, E. Cappellaro, M. Colapietro, M. Della Valle, O. Diner, S. D'Orsi, D. Fantinel, J. Fynbo, A. Gal-Yam, M. Genoni, M. Hirvonen, J. Kotilainen, T. Kumar, M. Landoni, J. Lehti, G. Li Causi, D. Loreggia, L. Marafatto, S. Mattila, G. Pariani, G. Pignata, M. Rappaport, D. Ricci, M. Riva, B. Salasnich, R. Zanmar Sanchez, S. Smartt, and M. Turatto (July 2018). "SOXS: A Wide Band Spectrograph to Follow up Transients". In: *Ground-Based and Airborne Instrumentation for Astronomy VII*. Vol. 10702. SPIE, pp. 96–107. DOI: `10.1117/12.2307349`. (Visited on 08/10/2025).

Schlekat, Donovan, Megan Dubay, Dylan Dutton, Joshua Haislip, and Daniel Reichart (Mar. 2025). "The Future of Optical Astronomy with the Skynet Robotic Telescope Network". en. In: *Bulletin of the AAS* 57.2. ISSN: 2330-9458, 0002-7537. URL: `https://baas.aas.org/pub/2025n2i302p14/release/1` (visited on 07/19/2025).

Serra, Jean Paul (1982). *Image analysis and mathematical morphology*. en. London ; New York: Academic Press. ISBN: 978-0-12-637240-3.

Shapiro, Linda G. and George C. Stockman (2001). *Computer vision*. eng. OCLC: 247860912. Upper Saddle River, NJ: Prentice Hall. ISBN: 978-0-13-030796-5.

Smartt, S. J. et al. (Nov. 2017). "A Kilonova as the Electromagnetic Counterpart to a Gravitational-Wave Source". In: *Nature* 551.7678, pp. 75–79. ISSN: 1476-4687. DOI: `10.1038/nature24303`. (Visited on 09/10/2025).

Smith, K. W., R. D. Williams, D. R. Young, A. Ibsen, S. J. Smartt, A. Lawrence, D. Morris, S. Voutsinas, and M. Nicholl (Jan. 2019). "Lasair: The Transient Alert Broker for LSST:UK". In: *Research Notes of the AAS* 3.1, p. 26. DOI: `10.3847/2515-5172/ab020f`.

Smith, Ken (Oct. 2019). “Lasair: The transient alert broker for LSST:UK”. In: ADS Bibcode: 2019eeu..confE..51S, p. 51. DOI: `10.5281/zenodo.3478098`. URL: `https://ui.adsabs.harvard.edu/abs/2019eeu..confE..51S` (visited on 07/19/2025).

Smith, Robert and Iain Steele (Feb. 2017). “LT Telescope and IO:O Throughput”. en. In: Publisher: figshare. DOI: `10.6084/m9.figshare.4659421.v1`. URL: `https://figshare.com/articles/journal_contribution/LT_Telescope_and_IO_O_Throughput/4659421/1` (visited on 10/18/2024).

Smith, Roger, David Walker, and Hugo E Schwarz (2004). “The Tololo all Sky Camera: Tasca”. In: *Scientific Detectors for Astronomy: The Beginning of a New Era*. Springer, pp. 379–384.

Steele, Iain A., Robert J. Smith, Paul C. Rees, Ian P. Baker, S. D. Bates, Michael F. Bode, Mark K. Bowman, Dave Carter, Jason Etherton, Martyn J. Ford, Stephen N. Fraser, A. Gomboc, Robert D. J. Lett, Anthony G. Mansfield, Jonathon M. Marchant, Gustavo A. Medrano-Cerda, Christopher J. Mottram, D. Raback, A. B. Scott, M. D. Tomlinson, and R. Zamanov (Sept. 2004). “The Liverpool Telescope: Performance and First Results”. In: *SPIE Astronomical Telescopes + Instrumentation*. Ed. by Jacobus M. Oschmann Jr. USA, p. 679. DOI: `10.1117/12.551456`. (Visited on 05/26/2020).

Sterken, Chr. and J. Manfroid (1992). *Astronomical Photometry, a Guide*. en. Vol. 175. Astrophysics and Space Science Library. Dordrecht: Springer Netherlands. ISBN: 978-94-011-2476-8. DOI: `10.1007/978-94-011-2476-8`. URL: `http://link.springer.com/10.1007/978-94-011-2476-8` (visited on 08/23/2024).

Street, R. A., A. Adamson, J. P. Blakeslee, R. D. Blum, A. S. Bolton, T. Boroson, M. Bowman, C. Briceño, J. H. Elias, E. Gomez, S. Heathcote, E. Heinrich-Josties, A. Hopkinson, C.-H. Lee, B. W. Miller, J. Nation, S. Ridgway, D. R. Silva, and L. J. Storrie-Lombardi (Dec. 2020). “The Astrophysical Events Observatories Network (AEON)”. In: *Observatory Operations: Strategies, Processes, and Systems VIII*. Vol. 11449. SPIE, pp. 471–484. DOI: `10.1117/12.2559986`. (Visited on 08/10/2025).

Stubbs, Christopher W. and John L. Tonry (Aug. 2006). “Toward 1% Photometry: End-to-End Calibration of Astronomical Telescopes and Detectors”. en. In: *The Astrophysical Journal* 646.2. Publisher: American Astronomical Society, pp. 1436–1444. ISSN: 0004-637X, 1538-4357. DOI: `10.1086/505138`. URL: `https://iopscience.iop.org/article/10.1086/505138` (visited on 07/19/2025).

Talens, G. J. J., J. F. P. Spronck, A.-L. Lesage, G. P. P. L. Otten, R. Stuik, D. Pollacco, and I. A. G. Snellen (May 2017). “The Multi-site All-Sky CAmeRA (MASCARA): Finding transiting exoplanets around bright ($m_V < 8$) stars”. en. In: *Astronomy & Astrophysics* 601. Publisher: EDP Sciences, A11. ISSN: 0004-6361, 1432-0746. DOI: `10.1051/0004-6361/201630319`. URL: `http://www.aanda.org/10.1051/0004-6361/201630319` (visited on 07/18/2025).

Tanvir, N. R., A. J. Levan, A. S. Fruchter, J. Hjorth, R. A. Hounsell, K. Wiersema, and R. Tunnicliffe (Aug. 2013). “A "Kilonova" Associated with Short-Duration Gamma-Ray Burst 130603B”. In: *Nature* 500.7464, pp. 547–549. ISSN: 0028-0836, 1476-4687. DOI: `10.1038/nature12505`. arXiv: `1306.4971 [astro-ph]`. (Visited on 09/28/2025).

Tonry, J. L., L. Denneau, A. N. Heinze, B. Stalder, K. W. Smith, S. J. Smartt, C. W. Stubbs, H. J. Weiland, and A. Rest (May 2018). “ATLAS: A High-cadence All-sky Survey System”. In: *Publications of the Astronomical Society of the Pacific* 130.988, p. 064505. ISSN: 1538-3873. DOI: `10.1088/1538-3873/aabadf`. (Visited on 08/08/2025).

Torres, Carlos, Emilio Cuevas Agulló, and Juan Carlos Guerra García (2003). “Estudio de la procedencia de las masas de aire en la región subtropical: episodios de polvo sahariano”. In.

Torres, Carlos, Emilio Cuevas Agulló, Juan Carlos Guerra García, and Virgilio Carreño Corbella (2001). “Caracterización de las masas de aire en la región subtropical”. In.

Tüg, H., N. M. White, and G. W. Lockwood (Dec. 1977). “Absolute Energy Distributions of Alpha Lyrae and 109 Virginis from 3295 Å to 9040 Å.” In: *Astronomy and Astrophysics* 61, pp. 679–684. ISSN: 0004-6361. (Visited on 05/14/2025).

van Velzen, Sjoert, Suvi Gezari, Erica Hammerstein, Nathaniel Roth, Sara Frederick, Charlotte Ward, Tiara Hung, S. Bradley Cenko, Robert Stein, Daniel A. Perley, Kirsty Taggart, Ryan J. Foley, Jesper Sollerman, Nadejda Blagorodnova, Igor Andreoni, Eric C. Bellm, Valery Brinnel, Kishalay De, Richard Dekany, Michael Feeney, Christoffer Fremling, Matteo Giomi, V. Zach Golkhou, Matthew J. Graham, Anna. Y. Q. Ho, Mansi M. Kasliwal, Charles D. Kilpatrick, Shrinivas R. Kulkarni, Thomas Kupfer, Russ R. Laher, Ashish Mahabal, Frank J. Masci, Adam A. Miller, Jakob Nordin, Reed Riddle, Ben Rusholme, Jakob van Santen, Yashvi Sharma, David L. Shupe, and Maayane T. Soumagnac (Feb. 2021). “Seventeen Tidal Disruption Events from the First Half of ZTF Survey Observations: Entering a New Era of Population Studies”. In: *The Astrophysical Journal* 908.1, p. 4. ISSN: 0004-637X. DOI: `10.3847/1538-4357/abc258`. (Visited on 09/10/2025).

Varela, A. M., C. Bertolin, C. Muñoz-Tuñón, S. Ortolani, and J. J. Fuensalida (Dec. 2008). “Astronomical Site Selection: On the Use of Satellite Data for Aerosol Content Monitoring”. In: *Monthly Notices of the Royal Astronomical Society* 391.2, pp. 507–520. ISSN: 00358711, 13652966. DOI: `10.1111/j.1365-2966.2008.13803.x`. (Visited on 12/11/2023).

Vestrand, W. T., Konstantin N. Borozdin, Steven P. Brumby, Donald E. Casperson, Edward E. Fenimore, Mark C. Galassi, Katherine McGowan, Simon J. Perkins, William C. Priedhorsky, Daniel Starr, Robert White, Przemek Wozniak, and James A. Wren (Nov. 2002). “The RAPTOR experiment: a system for monitoring the optical sky in real time”. In: ed. by Robert I. Kibrick. Waikoloa, HI, pp. 126–136. DOI: `10.1117/12.`

459515. URL: `http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=874881` (visited on 08/07/2025).

Vogt, Steven S., Matthew Radovan, Robert Kibrick, R. Paul Butler, Barry Alcott, Steve Allen, Pamela Arriagada, Mike Bolte, Jennifer Burt, Jerry Cabak, Kostas Chloros, David Cowley, William Deich, Brian Dupraw, Wayne Earthman, Harland Epps, Sandra Faber, Debra Fischer, Elinor Gates, David Hilyard, Brad Holden, Ken Johnston, Sandy Keiser, Dick Kanto, Myra Katsuki, Lee Laiterman, Kyle Lanclos, Greg Laughlin, Jeff Lewis, Chris Lockwood, Paul Lynam, Geoffrey Marcy, Maureen McLean, Joe Miller, Tony Misch, Michael Peck, Terry Pfister, Andrew Phillips, Eugenio Rivera, Dale Sandford, Mike Saylor, Richard Stover, Matthew Thompson, Bernie Walp, James Ward, John Wareham, Mingzhi Wei, and Chris Wright (Apr. 2014). "APF—The Lick Observatory Automated Planet Finder". In: *Publications of the Astronomical Society of the Pacific* 126.938, pp. 359–379. ISSN: 1538-3873. DOI: `10.1086/676120`. URL: `http://dx.doi.org/10.1086/676120`.

Ward, Alan (2005). "The Importance of Understanding Clouds". In: *NASA*.

Wells, D. C., E. W. Greisen, and R. H. Harten (June 1981). "FITS - a Flexible Image Transport System". In: *Astronomy and Astrophysics Supplement Series* 44. Publisher: EDP ADS Bibcode: 1981A&AS...44..363W, p. 363. ISSN: 0365-0138. URL: `https://ui.adsabs.harvard.edu/abs/1981A&AS...44..363W` (visited on 02/28/2025).

Williams, Roy D, Gareth P Francis, Andy Lawrence, Terence M Sloan, Stephen J Smartt, Ken W Smith, and David R Young (Jan. 2024). "Enabling Science from the Rubin Alert Stream with Lasair". In: *RAS Techniques and Instruments* 3.1, pp. 362–371. ISSN: 2752-8200. DOI: `10.1093/rasti/rzae024`. (Visited on 09/28/2025).

Woo, Mason, Jackie Neider, Tom Davis, and Dave Shreiner (1999). *OpenGL programming guide: the official guide to learning OpenGL, version 1.2.* Addison-Wesley Longman Publishing Co., Inc.

Ye, Liang, Yufeng Wang, Zhiguo Cao, Zhibiao Yang, and Huasong Min (2022). "A Self Training Mechanism With Scanty and Incompletely Annotated Samples for Learning-Based Cloud Detection in Whole Sky Images". In: *Earth and Space Science* 9.6, e2022EA002220. ISSN: 2333-5084. DOI: `10.1029/2022EA002220`. (Visited on 12/26/2023).

Young, Andrew T. and William M. Irvine (Oct. 1967). "Multicolor Photoelectric Photometry of the Brighter Planets. I. Program and Procedure". In: *The Astronomical Journal* 72, p. 945. ISSN: 0004-6256. DOI: `10.1086/110366`. (Visited on 05/14/2025).

Zaher, Ali, Stéphane Thil, Julien Nou, Adama Traoré, and Stéphane Grieu (July 2017). "Comparative Study of Algorithms for Cloud Motion Estimation Using Sky-Imaging Data". In: *IFAC-PapersOnLine* 50.1, pp. 5934–5939. ISSN: 24058963. DOI: `10.1016/j.ifacol.2017.08.1488`. (Visited on 03/06/2022).

Zhang, Songiie, Zhekang Dong, Xinyi Yang, Songiian Chai, Zhao Xu, and Donglian Qi (July 2019). "Intrahour Cloud Tracking Based on Optical Flow". In: *2019 Chinese*

Control Conference (CCC). Guangzhou, China: IEEE, pp. 3023–3028. ISBN: 978-988-15639-7-2. DOI: `10.23919/ChiCC.2019.8865296`. (Visited on 03/06/2022).

Zhu, Xianyi, Yi Xiao, Guanghua Tan, Shizhe Zhou, Chi-Sing Leung, and Yan Zheng (Aug. 2020). "GPU-accelerated 2D OTSU and 2D Entropy-Based Thresholding". In: *Journal of Real-Time Image Processing* 17.4, pp. 993–1005. ISSN: 1861-8200, 1861-8219. DOI: `10.1007/s11554-018-00848-5`. (Visited on 10/04/2020).