



Master's Thesis

04-02-2026

Sayak Ghosh

***Novae and Optical Transients with Ground-based Gamma-ray
Instruments:***

Optical light curves from H.E.S.S. NSB data and detectability with SWGO

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Novae and optical transients	2
1.2.1	Classical novae in brief	2
1.2.2	What optical light curves encode	2
1.2.3	High-energy emission from novae	3
1.3	Ground-based gamma-ray instruments in time-domain astronomy	4
1.3.1	Imaging atmospheric Cherenkov telescopes	4
1.3.2	Wide-field air-shower arrays	4
1.4	Thesis outline	5
2	Shocks and Particle Acceleration in Novae	6
2.1	Classical novae	6
2.2	Optical emission characteristics and light-curve behaviour	7
2.3	Nova outflows and shock formation	8
2.3.1	Multiple outflow components	9
2.3.2	Radiative shocks and energy dissipation	9
2.3.3	Internal shocks as sites of particle acceleration	9
2.4	Particle acceleration at nova shocks	10
2.4.1	Acceleration timescale and maximum energy	10
2.4.2	Acceleration efficiency and magnetic-field amplification	10
2.5	High-energy emission mechanisms in novae	11
2.5.1	Hadronic emission: inelastic pp collisions and π^0 decay	11
2.5.2	Leptonic emission: inverse Compton and bremsstrahlung	12
2.5.3	Transparency and multiwavelength diagnostics	12
2.6	Observed multiwavelength phenomenology of γ -ray novae	12
3	Air-shower detection physics: IACTs and WCDs	15

3.1	From primary γ rays to extensive air showers	15
3.1.1	Electromagnetic cascades in the atmosphere	15
3.1.2	Cherenkov radiation in air	17
3.2	Imaging atmospheric Cherenkov telescopes: detection physics and the H.E.S.S. implementation	19
3.2.1	Detection principle and reconstruction observables in IACTs . . .	19
3.2.2	HESS array layout and camera geometry	20
3.2.3	H.E.S.S. observing mode and trigger chain	21
3.3	Water Cherenkov detector arrays: detection physics and the SWGO concept	21
3.3.1	Detection principle and reconstruction observables in WCDs . .	21
3.3.2	SWGO: site and array concept	22
3.3.3	SWGO WCD unit design and detector-level physics	23
3.3.4	Readout, timing synchronization, and triggering	24
3.3.5	First-stage deployment: SWGO-A	25
3.3.6	Complementarity of IACTs and WCDs arrays	26
4	Detectability of novae with SWGO	27
4.1	Scope and analysis overview	27
4.2	A parametric nova-shock model: implementation in a custom Nova class	28
4.2.1	Model parameters	28
4.2.2	Shock kinematics	28
4.2.3	Target densities: ejecta and wind	29
4.2.4	Timescales and hadronic interaction efficiency	31
4.2.5	Energetics	31
4.2.6	Acceleration and maximum energy	32
4.3	SWGO instrument response functions	32
4.4	RS Oph as an example: GAMERA SED and SWGO comparison	35
4.5	Integral sensitivity and nova detectability map	36
4.6	Summary and outlook	37
5	Optical light curves from H.E.S.S. Night Sky Background (NSB) data	39
5.1	NSB observables, run metadata, and camera geometry	39
5.1.1	NSB time series format and sampling	40
5.1.2	Run-level metadata and pointing configuration	40

5.1.3	Camera geometry and pixel-ID conventions	40
5.2	Source cluster definition on the camera plane	41
5.2.1	Pointing, wobble and camera coordinates	42
5.2.2	Source localisation on the camera plane	43
5.2.3	Cluster selection strategy	44
5.2.4	Illustration of selected pixels	46
5.3	Step-by-step construction of the NSB derived light curve	47
5.3.1	Representative run and raw cluster light curve	48
5.3.2	Flat-field correction	49
5.3.3	Atmospheric transparency correction	51
5.3.4	Background subtraction	52
5.3.5	Light curve evolution and run continuity	54
5.3.6	Nightly mean NSB proxy and cross-telescope consistency	54
5.4	Run-to-run discontinuities and probable causes	55
5.4.1	Phenomenology of run-to-run discontinuities	55
5.4.2	Probable association with pointing offset transitions	56
5.4.3	Pixel-level origin of the anomalous runs	56
5.4.4	Check of flat-field calibration stability for the dominant pixels	57
5.5	Summary and outlook	59
6	Discussion and conclusion	60
A	Software and repositories	65
A.1	GAMERA	65
A.2	NOVA	65
A.3	nsb-to-nova-lightcurves	66
B	Supplementary plots	67
C	Supplementary figures	75
D	Supplementary tables	76
	Bibliography	80

Chapter 1

Introduction

Novae are among the most frequent explosive transients in the Galaxy, producing rapid optical brightening and rich time variability on timescales of hours to weeks (Bode and Alan Evans, 2008). In recent years, they have also emerged as established high-energy sources following the Fermi Large Area Telescope (Fermi-LAT) detections of GeV emission from multiple Galactic novae (Abdo et al., 2010; Fermi-LAT Collaboration, 2014), demonstrating that particle acceleration in nova-driven shocks can be efficient (Metzger et al., 2015). This makes novae an attractive target for wide-field gamma-ray observatories and, more broadly, for time-domain studies that combine optical and gamma-ray information. In this thesis, we explore two complementary ways in which ground-based gamma-ray instrumentation can contribute to nova science: (i) assessing the detectability of novae with the proposed Southern Wide-field Gamma-ray Observatory (SWGGO) (Conceição and SWGGO Collaboration, 2023), and (ii) extracting optical light curves from High Energy Stereoscopic System (H.E.S.S.) data using Night Sky Background (NSB)-related camera signals to trace optical brightness variations.

1.1 Motivation

Time-domain astronomy studies sources that change with time, where the timing of the emission can reveal the physical processes at work. For explosive transients, the brightness can evolve strongly within hours to days, so observation cadence (time spacing) and duty cycle become crucial. In practice, however, many events are not well covered, especially at early times, due to scheduling constraints, weather, seasonal visibility, and limited sky coverage.

Novae are a useful laboratory for studying explosive outflows and shocks in our Galaxy. They are bright in the optical and evolve on timescales of days to weeks, which makes them accessible to regular monitoring. The detection of high-energy emission from novae also shows that these shocks can accelerate particles to relativistic energies. From an astroparticle point of view, novae therefore provide nearby, time evolving systems in which non-thermal particle populations can be probed through their radiation.

Optical and gamma-ray measurements can provide complementary information. Optical light curves trace how the visible emission evolves as the outflow expands and cools, while gamma rays directly probe non-thermal emission from accelerated particles. Comparing the timing and variability across these bands can therefore help to constrain where and when shocks form and how energy is dissipated during the eruption. In addition, the gamma-ray signal can provide indirect information on the underlying

acceleration and interaction channels (e.g. leptonic versus hadronic scenarios), which connects nova eruptions to broader questions of particle acceleration in astrophysical shocks¹.

Observationally, novae are challenging high-energy targets because their onset is unpredictable and their emission evolves quickly. This motivates instruments with a wide field of view and a high duty cycle. Also, large archival datasets from Imaging Atmospheric Cherenkov Telescopes (IACTs) include camera rate information related to the NSB, recorded continuously while tracking a field. Although these instruments are optimised for gamma-ray observations, the same rate measurements are sensitive to changes in optical sky brightness within the field of view. If this information can be calibrated and corrected for instrumental and atmospheric systematics, it could enable the reconstruction of usable optical light curves for bright transients observed during gamma-ray pointings. Establishing the feasibility of this idea would not only improve temporal coverage for individual events, but could also broaden the scientific use of existing H.E.S.S. archival data as a complementary, high cadence optical monitoring resource. These points motivate the analysis approach adopted in this work, which combines a detectability study for SWGO with an exploration of what can be learned from archival H.E.S.S. observations.

1.2 Novae and optical transients

1.2.1 Classical novae in brief

Classical novae occur in close binary systems in which a white dwarf accretes matter from a companion star. As hydrogen-rich material accumulates on the white dwarf surface, pressure and temperature rise until a thermonuclear runaway ignites in the accreted layer (Bode and Alan Evans, 2008). The resulting explosion ejects material at velocities of order 10^3 km s^{-1} and produces a rapid increase in optical brightness, followed by a decline over days to months. Because the eruption evolves on human-accessible timescales and often reaches naked-eye brightness, novae have long been key targets for optical time-domain astronomy.

In the following, novae and nova-like optical transients are used as a representative class of fast-evolving sources for which dense temporal coverage is particularly valuable. Their optical evolution provides a direct record of the eruption development, while their high-energy emission probes additional non-thermal processes discussed below.

1.2.2 What optical light curves encode

Optical light curves describe how the brightness of a source changes with time. For novae, they capture the rise to maximum light, the subsequent decline, and in some cases additional features such as plateaus, secondary maxima, or dust-related dips which can be observed in Figure 1.1. These features reflect the changing conditions in the expanding ejecta, including its temperature, density, opacity, and geometry. The

¹This is discussed in more detail in Chapter 2.




Figure 1.1: Morphology of a nova light curve, from Bode and Alan Evans (2008).

overall decline rate is often used as a simple empirical descriptor of how quickly the eruption evolves, while shorter-timescale variability can indicate structure in the outflow or interaction with the surrounding environment.

From a practical point of view, optical light curves are also useful because they provide timing information: when the eruption started, when it reached peak brightness, and when major transitions occurred. This timing is important when comparing optical behaviour to emission at other wavelengths, especially at high energies where detections may be limited in time or sensitivity.

1.2.3 High-energy emission from novae

The detection of GeV gamma-ray emission from novae demonstrates that nova eruptions can produce strong shocks and accelerate particles to relativistic energies (Abdo et al., 2010; Fermi-LAT Collaboration, 2014). These shocks may form within the ejecta (for example, if a faster outflow catches up with a slower earlier outflow) or through interaction between the ejecta and surrounding material. In either case, the shock converts kinetic energy into heat and into non-thermal particle populations.

Gamma-ray production can proceed through different channels. In leptonic scenarios, accelerated electrons produce gamma rays through processes such as inverse Compton scattering or bremsstrahlung. In hadronic scenarios, accelerated protons interact with ambient material and produce neutral pions; the subsequent decay ($\pi^0 \rightarrow \gamma\gamma$) generates gamma rays. Distinguishing between these possibilities requires multi-wavelength information and, in particular, constraints on how the high-energy emission evolves relative to the optical light curve (Metzger et al., 2015).

These considerations motivate the observational approaches discussed in the next section: evaluating the detectability of novae and nova-like high-energy emission with SWGO, and assessing the feasibility of reconstructing optical light curves from archival H.E.S.S. observations.

1.3 Ground-based gamma-ray instruments in time-domain astronomy

Ground-based gamma-ray instruments detect very-high-energy photons indirectly, by measuring the secondary particles and Cherenkov light produced when a gamma ray initiates an extensive air shower in the atmosphere. Two major instrument classes are relevant for this thesis: Imaging Atmospheric Cherenkov Telescopes (IACTs), such as H.E.S.S., and wide-field Water Cherenkov Detector (WCD) arrays, such as the proposed SWGO. Their complementary strengths are particularly important in the context of time-domain astronomy, where sky coverage, duty cycle, and observation strategy strongly affect what can be measured.

1.3.1 Imaging atmospheric Cherenkov telescopes

IACT arrays detect nanosecond-scale Cherenkov flashes from air showers using fast optical telescopes and pixelised cameras. This technique provides strong background rejection and good angular resolution, and is most sensitive at energies typically above a few tens of GeV up to multi-TeV energies (instrument-dependent). However, IACT observations are pointed and are limited to clear, dark conditions, which reduces the duty cycle compared to wide-field instruments.

Despite these constraints, IACT archives are valuable for transient studies for two reasons. First, a pointed observation can provide high sensitivity over a short time interval, enabling detailed follow-up once a transient is observed. Second, IACT cameras continuously record rates related to the night-sky background while tracking a field. These rate measurements are primarily used for instrument monitoring and data quality checks, but they are also sensitive to changes in the optical sky brightness within the field of view. This creates the possibility of extracting optical variability information for sufficiently bright transients that occur during gamma-ray observations, provided that instrumental and atmospheric systematics can be controlled.

In this thesis, we explore this idea using archival H.E.S.S. data: we assess whether night-sky-background-related rate measurements can be calibrated to reconstruct usable optical light curves, and discuss the potential implications of such a capability for extending the scientific output of IACT archives.

1.3.2 Wide-field air-shower arrays

Wide-field air-shower arrays or Water Cherenkov Detector (WCD) arrays operate by sampling the secondary particles from extensive air showers at ground level. In contrast to IACTs, they typically observe a large fraction of the overhead sky continuously, leading to a high duty cycle and making them well suited for monitoring unpredictable transients. The trade-off is generally a higher energy threshold and/or lower instantaneous sensitivity compared to pointed IACT observations, depending on energy and source spectrum.

For nova-like emission scenarios that evolve on timescales of days to weeks, continuous monitoring is especially valuable because it reduces reliance on external triggers and

minimises gaps in coverage. The proposed SWGO, as a next-generation wide-field gamma-ray observatory in the Southern hemisphere, is therefore a promising instrument for time domain studies of Galactic transients. In this thesis, we investigate under which conditions SWGO could detect high energy emission from novae and nova-like transients, and how detectability depends on assumptions about the emission timescale and spectrum.

1.4 Thesis outline

Chapter 2 introduces the relevant physical background on novae, optical light curves, and high-energy emission mechanisms. Chapter 3 summarises the instrument concepts and detection physics relevant to this work, with emphasis on the differences between IACTs and WCDs. Chapter 4 presents the SWGO detectability study and discusses the dependence on emission assumptions and timescales. Chapter 5 describes the method to extract optical light curves from archival H.E.S.S. data using NSB related camera rate information, and evaluates its performance on selected cases. Chapter 6 discusses implications, limitations, and prospects for future work, and summarises the main results and conclusions.

Chapter 2

Shocks and Particle Acceleration in Novae

2.1 Classical novae

Classical novae are non-terminal thermonuclear eruptions occurring on the surface of an accreting white dwarf in a compact interacting binary system¹ (Bode and Alan Evans, 2008). Mass transfer from the companion star supplies a hydrogen-rich envelope on the white dwarf. As accretion proceeds, the base of the envelope is compressed and heated under partially degenerate conditions until a thermonuclear runaway ignites. The rapid energy release drives a dramatic increase in luminosity and launches an outflow, while the binary system itself survives the event (Laura Chomiuk et al., 2021). Figure 2.1 summarizes the basic geometry of the interacting binary and the main mass-flow channels (accretion via Roche-lobe overflow and ejection during the eruption) that underpin the discussion in this chapter.

The mass-transfer geometry depends on the binary configuration. In many systems, Roche-lobe overflow forms an accretion disk around the white dwarf, whereas magnetic white dwarfs can channel accretion along field lines. Regardless of the detailed accretion flow, the nova eruption is characterized by (i) a sharp rise to optical maximum, (ii) subsequent decline on timescales of days to months, and (iii) substantial mass loss into the circumbinary environment (Laura Chomiuk et al., 2021). Typical ejecta velocities are of order 10^2 km s^{-1} to 10^3 km s^{-1} , and inferred ejecta masses span several orders of magnitude, depending on the system parameters and the eruption class (Laura Chomiuk et al., 2021). These scales are directly relevant for shock formation and radiative efficiency, which control the emergence of non-thermal emission.

Historically, the optical luminosity of novae was often interpreted primarily in terms of continued nuclear burning and reprocessing in an expanding photosphere. However, the discovery of GeV γ -ray emission from classical novae with Fermi-LAT demonstrates that energetic shocks and efficient particle acceleration occur during the outburst (Fermi-LAT Collaboration, 2014). In some events, coordinated multiwavelength observations provide direct evidence that shocks can contribute substantially to the optical luminosity around maximum light (Aydi et al., 2020). This motivates treating the nova eruption as a time-dependent interaction between multiple outflow components rather than a single homologously expanding shell.

¹In this thesis, “classical nova” generally refers to thermonuclear runaways on accreting white dwarfs and excludes accretion-driven dwarf-nova outbursts in cataclysmic variables.

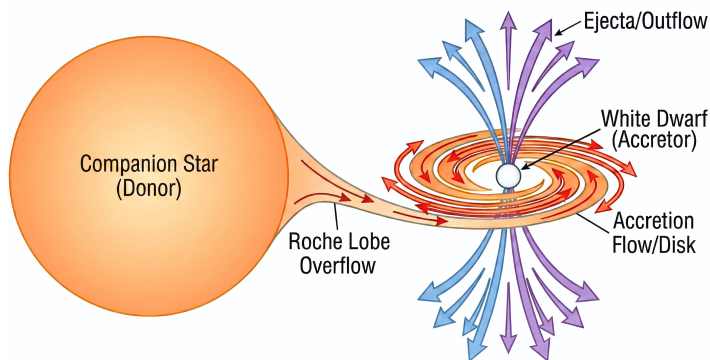


Figure 2.1: A schematic representation of a classical nova system that shows how material is expelled during the nova outburst, how an accretion stream or disk forms around the white dwarf, and how mass is transferred from the companion star through Roche-lobe overflow. The illustration is not to scale and highlights the key aspects relevant to the discussion on nova outflows and shocks in this chapter.

In the following, we focus on the outflow structure expected during nova eruptions and on the formation of internal shocks, which provide the physical basis for particle acceleration and for the multiwavelength signatures exploited later in this thesis.

2.2 Optical emission characteristics and light-curve behaviour

Optical observations provide the most complete temporal coverage for nova eruptions and are therefore a natural reference for defining outburst phases (Laura Chomiuk et al., 2021). In the classical picture, the early optical peak is associated with an optically thick expanding pseudo-photosphere and with reprocessing of higher-energy radiation in the ejecta. As the ejecta expand and thin out, the emergent spectrum and luminosity evolve accordingly. However, multiwavelength campaigns have shown that internal shocks can contribute substantially to the radiated power in some novae, including in the optical band around maximum light (Aydi et al., 2020).

Throughout this thesis, we use standard photometric terminology. Apparent magnitudes are a logarithmic measure of flux. For two measurements with fluxes F_1 and F_2 , the magnitude difference is

$$m_1 - m_2 = -2.5 \log_{10} \left(\frac{F_1}{F_2} \right), \quad (2.1)$$

such that a decline by $\Delta m = 2$ or 3 corresponds to a flux decrease by factors of $10^{0.8}$ and $10^{1.2}$, respectively. This logarithmic scaling is useful when comparing light-curve shapes across a wide dynamic range, while later analysis steps in this work primarily rely on relative brightness variations rather than absolute photometric calibration.

A widely used parameterization of the optical decline is given by the times t_2 and t_3 , defined as the elapsed time after optical maximum for the nova to fade by 2 or 3




Figure 2.2: Schematic optical light curve of a nova illustrating the definition of the decline times t_2 and t_3 , measured from maximum light as the time required for the source to fade by 2 and 3 magnitudes, respectively. The curve also indicates a few other example light-curve features often observed in some novae. This figure is intended for definition and orientation only; it does not represent a specific event.

magnitudes in a given band (Kato and Hachisu, 2023; Payne-Gaposchkin, 1957). The definition of t_2 and t_3 is illustrated schematically in figure 2.2. These timescales serve as simple proxies for the “speed class” of a nova and are commonly used when comparing nova populations and their ejecta properties (e.g. Tappert et al., 2020). While t_2 and t_3 compress complex behaviour into a single number, they provide a reproducible and instrument-independent way to align different events in time.

Beyond a smooth rise and decline, many novae exhibit additional structure such as plateaus, secondary maxima, or dust-formation dips, reflecting changes in opacity, geometry, and energy deposition in the outflow (Laura Chomiuk et al., 2021). Such features are particularly relevant because they can coincide with phases of enhanced shock interaction and thus potentially correlate with high-energy activity (Aydi et al., 2020). These points lead us to describe the eruption as arising from the interaction of several distinct outflow components, which we will see in the next section.

2.3 Nova outflows and shock formation

The GeV γ -ray detections of classical novae imply that strong shocks capable of accelerating particles to relativistic energies develop during the eruption (Laura Chomiuk et al., 2021; Metzger et al., 2015). For most classical novae, interaction with a dense pre-existing circumstellar medium is not expected to dominate the early-time dynamics, motivating scenarios in which the shocks arise predominantly from collisions between distinct components of the nova outflow (Martin et al., 2018). In this picture, the eruption is not adequately described by a single homologously expanding shell; instead, mass loss proceeds in multiple phases whose relative velocities and densities naturally

lead to internal shocks.

2.3.1 Multiple outflow components

Observations across the electromagnetic spectrum support the presence of structured and time-dependent mass ejection in novae, including evidence for an early, relatively slow, and dense outflow followed by a faster wind launched at later times (Laura Chomiuk et al., 2021). When a fast component with velocity v_f catches up to slower material moving at v_s , a two-shock structure develops: a forward shock propagating into the slow outflow and a reverse shock propagating back into the fast wind (Martin et al., 2018). The shock interaction region can compress into a thin cooling layer at high densities, which is also the site where non-thermal particles can interact efficiently with the surrounding gas and radiation fields.

2.3.2 Radiative shocks and energy dissipation

A key distinction is whether the shocks are radiative. At the high post-shock densities expected in nova ejecta, the thermal plasma behind the shock can cool rapidly compared to the expansion timescale, causing the shocked gas to collapse into a dense shell and strongly suppressing the emergent soft X-ray emission due to absorption and reprocessing (Laura Chomiuk et al., 2021; Vurm and Metzger, 2018). In the radiative regime, a large fraction of the shock power can ultimately emerge at lower photon energies, including optical and infrared, providing a physical mechanism for correlations between optical variability and high-energy emission (Aydi et al., 2020).

The instantaneous kinetic power dissipated at a shock can be written schematically as

$$L_{\text{sh}} \sim \frac{1}{2} \dot{M} v_{\text{rel}}^2, \quad (2.2)$$

where \dot{M} characterizes the mass flux through the shock and $v_{\text{rel}} \simeq |v_f - v_s|$ is the relative velocity between colliding outflow components (Metzger et al., 2015). This scaling emphasizes that both the mass-loss history and the velocity stratification of the ejecta control the available shock power and, consequently, the luminosity available for thermal and non-thermal emission.

2.3.3 Internal shocks as sites of particle acceleration

Particle acceleration at nova shocks is commonly described in the framework of diffusive shock acceleration, in which a fraction of particles repeatedly scatter across the shock and gain energy. Internal-shock models have been developed to explain the observed GeV light curves and spectra, with hadronic interactions in the dense cooling layer providing an efficient channel for γ -ray production in many classical novae (Metzger et al., 2015; Martin et al., 2018). Radiative shock models further predict accompanying hard X-ray emission from the same non-thermal particle populations, with the ratio of hard X-ray to γ -ray luminosity providing a diagnostic of the dominant emission channel and microphysical processes at the shock (Vurm and Metzger, 2018).

In the remainder of this chapter, we use this multi-component outflow and internal-shock framework as the physical basis for interpreting optical variability and for motivating time-domain searches for nova-like transients at high energies.

2.4 Particle acceleration at nova shocks

The detection of GeV γ rays from novae implies that a non-thermal particle population is produced efficiently during the outburst (Metzger et al., 2015). A standard framework to describe such non-thermal populations in collisionless astrophysical shocks is Diffusive Shock Acceleration (DSA), in which particles gain energy through repeated shock crossings mediated by scattering off magnetic turbulence (Blandford and Eichler, 1987). In the test-particle limit for strong, non-relativistic shocks, DSA predicts an approximately power-law momentum distribution, $f(p) \propto p^{-4}$, corresponding to a differential energy spectrum close to $dN/dE \propto E^{-2}$ for relativistic particles (Blandford and Eichler, 1987). While the detailed spectral shape depends on shock obliquity, turbulence properties, and non-linear feedback, this scaling provides a useful baseline for interpreting nova γ -ray spectra and for motivating hadronic and leptonic emission scenarios in section 2.5.

2.4.1 Acceleration timescale and maximum energy

A practical way to assess whether nova shocks can accelerate particles to GeV-TeV energies is to compare the acceleration time to the relevant loss and dynamical timescales. In DSA, the characteristic acceleration time to energy E can be written schematically as

$$t_{\text{acc}}(E) \sim \frac{\kappa(E)}{v_{\text{sh}}^2}, \quad (2.3)$$

where v_{sh} is the shock velocity and $\kappa(E)$ is the spatial diffusion coefficient, set by the level of magnetic turbulence near the shock (Blandford and Eichler, 1987). The maximum energy E_{max} is then determined by the condition

$$t_{\text{acc}}(E_{\text{max}}) \lesssim \min(t_{\text{dyn}}, t_{\text{loss}}), \quad (2.4)$$

where t_{dyn} is the characteristic expansion time of the interaction region and t_{loss} includes radiative and collisional losses appropriate to the particle species. In dense nova ejecta, hadrons can lose energy through inelastic proton-proton collisions on a timescale that can be comparable to the expansion time during the early outburst, which simultaneously enhances γ -ray production efficiency and can limit the highest attainable energies (Metzger et al., 2015; Vurm and Metzger, 2018).

2.4.2 Acceleration efficiency and magnetic-field amplification

The luminosity in non-thermal particles is commonly parameterized by an acceleration efficiency ϵ_{nt} , defined as the fraction of the shock power converted into non-thermal particle energy. In nova-shock modelling, ϵ_{nt} enters directly in the conversion from shock power (cf. equation (2.2)) to the expected γ -ray luminosity (Metzger et al., 2015). Kinetic

simulations of non-relativistic shocks show that efficient ion acceleration can occur, particularly for quasi-parallel magnetic-field geometries, and that accelerated particles can drive substantial magnetic-field amplification via streaming instabilities, reducing $\kappa(E)$ and thereby accelerating particles more rapidly (Caprioli and Spitkovsky, 2014). Although the microphysics in nova ejecta is not identical to the interstellar medium, these results provide physical motivation for adopting ϵ_{nt} values at the percent level and for allowing magnetic turbulence levels above the background field in phenomenological nova-shock models (Metzger et al., 2015).

An important complication specific to novae is that the shocks are often radiative at early times (Metzger et al., 2015; Vurm and Metzger, 2018). Rapid post-shock cooling can compress the shocked region into a dense shell, modifying the structure of the shock and potentially affecting particle escape and the relative importance of hadronic versus leptonic emission. In the next section, we therefore treat the γ -ray production mechanisms as model alternatives constrained by the observed spectra and temporal evolution, rather than assuming a single universal acceleration outcome.

2.5 High-energy emission mechanisms in novae

The discovery of transient GeV γ -ray emission from novae with Fermi-LAT established these systems as particle accelerators on timescales of weeks² (Fermi-LAT Collaboration, 2014). In the internal-shock framework discussed above, the γ rays are produced by relativistic particles accelerated at the shocks and interacting with the dense outflow and the intense radiation field. Two broad classes of emission models are commonly considered: hadronic scenarios, in which accelerated ions dominate the γ -ray production, and leptonic scenarios, in which accelerated electrons dominate (Metzger et al., 2015; Martin et al., 2018; Vurm and Metzger, 2018).

2.5.1 Hadronic emission: inelastic pp collisions and π^0 decay

In hadronic models, accelerated protons (and heavier ions) interact inelastically with the ambient gas in the shocked, dense ejecta. These pp collisions produce neutral pions, π^0 , which decay promptly into two γ rays. The efficiency of this channel depends primarily on the target density and the residence time of relativistic hadrons in the interaction region. In the early phases of a nova eruption, the high densities implied by radiative shocks make hadronic interactions efficient and naturally yield GeV emission with durations comparable to the time over which dense shocked material persists (Metzger et al., 2015; Martin et al., 2018). Time-dependent internal-shock calculations further show that hadronic emission can reproduce the main features of observed Fermi-LAT light curves in several novae when interactions occur in the dense cooling layer downstream of the shock (Martin et al., 2018).

²The duration of detectable GeV emission varies between events and analysis choices; see section 2.6.

2.5.2 Leptonic emission: inverse Compton and bremsstrahlung

In leptonic models, accelerated electrons produce γ rays through inverse Compton (IC) scattering of optical/infrared photons and/or through non-thermal bremsstrahlung on ambient ions. The IC channel is sensitive to the photon energy density and to electron losses (including Coulomb losses in dense plasma), while bremsstrahlung scales with target density in a manner analogous to hadronic interactions. Detailed radiative-shock models demonstrate that strong Coulomb cooling of mildly relativistic electrons can suppress the hard X-ray and γ -ray yield expected from purely leptonic scenarios, making combined X-ray and GeV constraints particularly informative (Vurm and Metzger, 2018). More generally, comparisons between γ -ray luminosities and shock power provide constraints on the required electron acceleration efficiency in leptonic interpretations (Metzger et al., 2015).

2.5.3 Transparency and multiwavelength diagnostics

An important requirement for both model classes is that the system is sufficiently transparent for GeV photons to escape. Internal-shock models typically assume that the Fermi-LAT band probes emission produced in regions where the dominant attenuation processes do not fully absorb the GeV photons over the observed timescales (Metzger et al., 2015; Martin et al., 2018). Because the same shock-accelerated particle population generically produces emission extending down to the hard X-ray band, the ratio of hard X-ray to γ -ray luminosity has been proposed as a discriminator between leptonic and hadronic scenarios in radiative shocks (Vurm and Metzger, 2018). In addition, correlations between optical variability and γ -ray activity can arise naturally if a significant fraction of the shock power is reprocessed to optical wavelengths while non-thermal particles produce the GeV component (Metzger et al., 2015).

In the following section, we summarize the key observational phenomenology of GeV-detected novae and the timescales over which the multiwavelength emission is typically observed, providing context for the time-domain analyses pursued later in this thesis.

2.6 Observed multiwavelength phenomenology of γ -ray novae

The first detection of GeV γ -ray emission from a nova was obtained for the symbiotic system V407 Cyg in 2010, where the interaction between the nova ejecta and the dense wind of the red-giant companion provided an early physical interpretation in terms of shocks (Abdo et al., 2010). Subsequent Fermi-LAT observations established classical novae as a distinct class of transient GeV sources, demonstrating that shock-powered particle acceleration is not restricted to symbiotic environments (Fermi-LAT Collaboration, 2014). Systematic searches over the first years of the mission further showed that γ -ray detections are strongly distance-limited and that only a subset of optically discovered Galactic novae are detectable with Fermi-LAT, given typical flux levels and durations (Franckowiak et al., 2018).

The approximate phase ordering of the multiwavelength emission discussed below is summarized schematically in figure 2.3. Interestingly, in figure 2.3, we observe that there is a question mark after the term “*TNR flash*”. This is because, at the time of publication of the article from which the figure was adapted, it wasn’t confirmed yet. But, a more recent publication from König et al. (2022) confirms the detection of an extremely bright and very soft X-ray flash of the classical Galactic nova YZ Reticuli (MGAB-V207) (cf. figure C.1), and this is particularly relevant because this is the same nova about which we will talk more about in chapter 5.

As of late 2025, GeV emission has been reported from more than 20 Galactic novae, spanning classical and recurrent/symbiotic systems, with a wide range of peak fluxes and temporal morphologies (Aydi et al., 2025; Mukai, 2025). The characteristic duration of detectable GeV emission is typically of order days to weeks, with reported extremes ranging from ~ 1 -2 days to several tens of days, depending on the event and analysis time binning (Mukai, 2025; Fermi Science Support Center, 2024). In many cases the GeV activity occurs within a few days of optical maximum and overlaps the early decline, consistent with shock formation during the dense-ejecta phase (Fermi-LAT Collaboration, 2014; Franckowiak et al., 2018).

Multiwavelength follow-up indicates that the shock-powered picture is also reflected outside the GeV band. Hard X-ray emission is frequently weak or absent during the earliest phases, which can be understood as a consequence of absorption and reprocessing in dense ejecta when shocks are radiative (Vurm and Metzger, 2018; Wang et al., 2024). At later times, when the column density drops, thermal and supersoft X-ray emission can become visible, providing complementary diagnostics of the post-outburst evolution (Page et al., 2020). Together, these empirical trends motivate treating the optical and GeV light curves as time-domain tracers of the same underlying shock evolution, while recognizing that detectability depends on distance, cadence, and wavelength-dependent opacity.

The physical concepts discussed in this chapter motivates two complementary directions pursued in the remainder of this thesis. First, the multi-component outflow scenario and the resulting internal shocks provide a framework in which optical variability and GeV γ -ray activity can trace the same underlying shock evolution (sections 2.3 and 2.6). This motivates constructing dense optical time series with well-defined systematics control, as optical data provide the most complete temporal coverage for many eruptions. In particular, even when absolute photometric calibration is not available, relative brightness variations can be used to identify outburst phases and to search for temporal correlations with high-energy activity.

Second, the observational context summarized in section 2.6 shows that novae constitute transient particle accelerators on timescales of days to weeks, with event-to-event diversity in peak flux and duration. This motivates assessing the detectability of nova-like transients with current and future ground-based γ -ray instruments, especially in regimes where dense time-domain coverage is achievable and where complementary optical information can constrain the emission phase.

Accordingly, in the next chapters we develop and validate a method to extract optical light curves from H.E.S.S. camera NSB data, including a dedicated treatment of instrumental and atmospheric systematics, and evaluate the prospects for detecting




Figure 2.3: Schematic timeline of a nova eruption illustrating the approximate phase ordering of key multiwavelength tracers relative to the outflow evolution (slow and fast components and their shock interaction). Curves indicate qualitative detectability rather than absolute flux, and the timing and duration of each phase vary between events. The schematic is intended as a qualitative guide to the phase ordering discussed in section 2.6. Adapted from Laura Chomiuk et al. (2021).

novae and nova-like optical transients at very high energies with SWGO, using the physical and observational timescales established here as guiding inputs, as discussed earlier in section 1.1.

Chapter 3

Air-shower detection physics: IACTs and WCDs

Chapter 2 motivated novae as time dependent particle accelerators, in which shocks can channel kinetic power into non-thermal particle populations and high energy photons. From the ground, however, VHE γ rays are not observed directly. Instead, the atmosphere converts the primary particle (incoming Cosmic ray or γ ray) into an Extensive Air Shower (EAS), and ground-based instruments infer the primary γ ray properties from the secondary cascade and its associated radiation fields.

This chapter summarizes the detector technologies used for air-shower detection and the physical principles that connect air showers to measurable signals. As discussed above, we focus on two complementary approaches: imaging the Cherenkov light emitted by shower particles in air using High Energy Stereoscopic System (H.E.S.S.) as a representative IACT system, and measuring the shower particle footprint at ground using WCD units as planned by the Southern Wide-field Gamma-ray Observatory (SWGGO) collaboration. The goal is to establish the measurement chain, starting with shower development, light production, sensor response, and reconstruction observables needed for the analysis methods introduced in the following chapters.

3.1 From primary γ rays to extensive air showers

Ground-based instruments infer the properties of primary VHE γ rays through the EAS they initiate in the atmosphere. In the energy range relevant for modern ground-based observatories, the shower development is governed by well-understood electromagnetic interactions, while the dominant backgrounds are hadron induced showers with additional hadronic and muonic components. In this section, we summarize the key features of EAS development that determine the production and space-time structure of Cherenkov light. (Particle Data Group, 2025; Particle Data Group, 2024)

3.1.1 Electromagnetic cascades in the atmosphere

A primary γ ray entering the atmosphere typically converts via pair production in the Coulomb field of a nucleus, producing an e^+e^- pair. The electrons and positrons radiate bremsstrahlung photons, which can again convert by pair production, leading to a multiplicative electromagnetic cascade. The longitudinal development is often




Figure 3.1: Left: Heitler-model schematic of an electromagnetic cascade initiated by a γ ray in the atmosphere. The primary photon converts in the Coulomb field of an atomic nucleus, producing an e^+e^- pair; subsequent bremsstrahlung emission and repeated pair production generate a multiplicative shower. The cascade is shown in discrete steps of one radiation length X_0 , illustrating the rapid growth of particle multiplicity in the early shower development. Right: Monte Carlo visualization of a 100 GeV photon-induced air shower as a function of altitude, for a fixed altitude of the first interaction. The simulated track density highlights the concentration of secondary charged particles along the shower axis and the longitudinal spread of the cascade through the atmosphere. Illustration taken from Steppa (2022).

characterized in terms of the slant depth X (in g cm^{-2}) and the depth of shower maximum, X_{max} , at which the number of charged particles reaches its peak. (Particle Data Group, 2024)

A useful conceptual description is provided by the Heitler model (cf. figure 3.1), in which the cascade develops through successive particle splittings with approximately equal energy sharing. In this picture, the number of particles grows as $N \sim 2^t$, where t is the depth expressed in units of radiation lengths, until the typical particle energy falls to the critical energy E_c where ionization losses become comparable to radiative losses. The shower maximum then occurs after approximately

$$t_{\text{max}} \simeq \frac{\ln(E_\gamma/E_c)}{\ln 2}, \quad (3.1)$$

and the maximum number of charged particles scales as $N_{\text{max}} \sim E_\gamma/E_c$. While quantitative predictions require Monte Carlo simulations, these scalings capture the essential dependence on primary energy that drives instrument design: higher energy primaries generate larger cascades that develop deeper and produce more Cherenkov emitting charged particles. (Matthews, 2005)

Hadron-induced showers differ qualitatively because hadronic interactions produce secondary mesons and a substantial muonic component. The resulting shower images and ground footprints exhibit larger event-to-event fluctuations and additional substructure relative to purely electromagnetic cascades. Modern analyses exploit these morphological differences statistically for γ /hadron separation, using either image-based observables

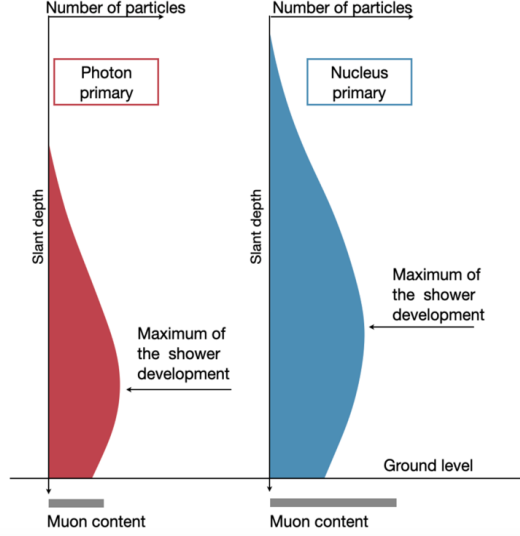


Figure 3.2: Schematic comparison of the longitudinal development of a γ -ray induced electromagnetic shower (left) and a hadron induced shower (right) as a function of slant depth X . The shower maximum X_{\max} marks the depth where the charged-particle content peaks. Hadron-induced showers exhibit a substantial muonic component and larger fluctuations, motivating γ /hadron discrimination in ground-based air-shower experiments. Adapted from Pierre Auger Collaboration (2023).

in IACTs or footprint-based observables in ground arrays. (Particle Data Group, 2025)

A schematic comparison of electromagnetic and hadron-induced shower development is shown in figure 3.2. Figure 3.3 summarizes the measurement chain from a primary γ ray to the observables recorded by IACTs and WCDs.

3.1.2 Cherenkov radiation in air

Charged shower particles emit Cherenkov radiation when their speed exceeds the phase velocity of light in the medium,

$$\beta n(\lambda) > 1, \quad (3.2)$$

where $\beta = v/c$ and $n(\lambda)$ is the wavelength-dependent refractive index of air. The radiation is concentrated on a cone with opening angle θ_c given by

$$\cos \theta_c = \frac{1}{\beta n(\lambda)}. \quad (3.3)$$

The Cherenkov emission geometry is illustrated in figure 3.4. The diagram also indicates the wavefront opening angle η ¹; in dispersive media, $\theta_c + \eta \neq 90^\circ$ because phase and group velocities differ (Particle Data Group, 2024).

¹Not to be confused with refractive index n .





Figure 3.3: Schematic of ground-based γ ray detection with extensive air showers. A primary γ ray initiates an air shower in the atmosphere; IACTs (left) record the brief Cherenkov light flash produced in air, while WCDs arrays (center/right) sample the secondary particles at ground. The inset illustrates the water-Cherenkov technique, where charged shower particles generate Cherenkov light in a water tank that is measured by PMTs. Taken from SWGO Collaboration (2019)



(a) Wavefront geometry in a medium of refractive index n (Cherenkov angle θ_c). (b) Illustration of coherent emission and the resulting Cherenkov cone for $v > c/n$.

Figure 3.4: Cherenkov emission in a dielectric medium. (a) A charged particle in the medium emits coherent radiation when $\beta n > 1$, producing emission at the Cherenkov angle θ_c with $\cos \theta_c = (\beta n)^{-1}$; the wavefront opening angle η is indicated. (b) Schematic view of the coherent wavefront build-up and the corresponding Cherenkov cone. Illustration from Vikas Joshi (2019)

In air, this geometry leads to a short (nanosecond-scale) flash at ground distributed over a characteristic “light pool” with a lateral extent of order $\sim 10^2$ m. (Particle Data Group, 2024)

The emitted photon yield is strongly weighted toward short wavelengths. In the Frank-Tamm description, the spectral yield per path length scales approximately as

$$\frac{d^2N}{dx d\lambda} \simeq \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right), \quad (3.4)$$

where α is the fine-structure constant. In practice, the detectable spectrum is shaped by atmospheric transmission², mirror reflectivity, light concentrators, and the sensor quantum efficiency, so that IACTs are typically most sensitive in the near-UV/blue optical band. (Particle Data Group, 2024; Bernlöhr et al., 2003)

These features motivate the two complementary detector concepts discussed in the remainder of this chapter. IACTs form an image of the Cherenkov light emitted in air, enabling stereoscopic reconstruction of shower geometry and image based background rejection. Ground arrays based on WCDs, in contrast, sample the shower particle footprint at ground, using Cherenkov emission in water to reconstruct the shower-front timing and lateral signal distribution with high duty cycle and wide field of view.


3.2 Imaging atmospheric Cherenkov telescopes: detection physics and the H.E.S.S. implementation

3.2.1 Detection principle and reconstruction observables in IACTs

In the IACT technique, the VHE primary is derived indirectly from the Cherenkov light emitted by relativistic charged particles in the electromagnetic cascade (cf. section 3.1.2). The Cherenkov flash lasts only a few nanoseconds at ground and is spatially confined to a light pool with a lateral scale of order $\sim 10^2$ m, so the measurement is intrinsically photon-starved and requires fast optics and fast readout (Hinton and Hofmann, 2009). A tessellated mirror collects the Cherenkov photons and images the angular light distribution onto a pixelated focal-plane camera. Each triggered telescope thus records a two-dimensional intensity pattern that encodes the shower geometry projected onto the camera plane.

After calibration and image cleaning, the cleaned image is commonly summarized by a small set of shape and orientation parameters (e.g. Hillas-type moments), which provide estimators for the shower axis direction and the impact parameter with respect to the telescope (Hillas, 1985). Stereoscopic observations with multiple telescopes viewing the same cascade from different positions enable a geometric reconstruction of the shower axis in three dimensions and reduce degeneracies present in single-telescope images (Hinton and Hofmann, 2009). Event classification exploits that γ ray induced electromagnetic showers produce comparatively compact and regular images, while

²Such as Rayleigh and aerosol scattering, ozone absorption.



(a) Array layout of H.E.S.S.: the four H.E.S.S. I telescopes (CT1–CT4) form an approximate square with side length 120 m, and the H.E.S.S. II telescope (CT5) is located near the centre. (b) Photograph of the H.E.S.S. site in Namibia showing CT1–CT4 and the central CT5.

Figure 3.5: The H.E.S.S. instrument and its array geometry. The schematic highlights the characteristic telescope positions, while the photograph provides a visual impression of the site and the relative scale of CT5 compared to CT1–CT4. Photo credit: Clementina Medina, (H.E.S.S. Collaboration, 2025).

hadron-induced showers tend to exhibit increased substructure and larger fluctuations due to hadronic multiplicity and a muonic component (cf. section 3.1.1).

3.2.2 H.E.S.S. array layout and camera geometry

The High Energy Stereoscopic System (H.E.S.S.) instrument is located in Namibia on the Khomas Highland at an altitude of about 1800 m above sea level. Figure 3.5 shows a view of the array with the four HESS I telescopes and the central HESS II telescope along with a schematic diagram. In its first phase (HESS I), four telescopes (CT1–CT4) were constructed in 2002–2003 and arranged on the corners of a square with side length 120 m (cf. figure 3.5a) to enable stereoscopic imaging over the Cherenkov light pool. Each of the HESS I telescopes has a segmented reflector composed of 382 round mirror facets, a focal length of 15 m, and a total mirror area of 108 m^2 . The corresponding cameras have a field of view of 5° and consist of 960 PMTs as light sensors with a pixel size of 0.16° (H.E.S.S. Collaboration, 2025). The camera pixels are arranged on a hexagonal grid to provide approximately uniform angular sampling across the field of view; the pixelization defines the intrinsic image sampling scale that enters image parametrization and reconstruction (Hinton and Hofmann, 2009).

Since 2012, the second phase (HESS II) adds a fifth, larger telescope (CT5) located at the centre of the original array to extend the sensitivity towards lower energies. CT5 uses 875 hexagonal mirror facets, has a focal length of 36 m, and a total mirror area of 614 m^2 . The CT5 camera comprises 2048 PMTs and provides a 3.2° field of view (H.E.S.S. Collaboration, 2025). The original HESS I cameras were replaced in the 2015–2016 upgrade campaign (often denoted HESS IU), primarily to modernize

the trigger/readout electronics and reduce dead time while retaining the 960-PMT photosensor plane (Giavitto et al., 2016).

3.2.3 H.E.S.S. observing mode and trigger chain

H.E.S.S. observations are typically taken in wobble mode, where the source is offset from the camera centre by $\pm 0.5^\circ$ and the offset direction is alternated in right ascension and declination between runs of about 28 min. This strategy enables simultaneous background control regions within the same field of view and is commonly paired with reflected-region or ring-based background models (Aharonian and H.E.S.S. Collaboration, 2006; Berge et al., 2007).

Because the Cherenkov signal is both faint and fast, H.E.S.S. employs a staged trigger to suppress NSB fluctuations. A camera trigger begins from a single-pixel threshold of order $\sim 4 \text{ p.e.}^3$ within a $\sim 1.5 \text{ ns}$ window at the PMT, followed by a multiplicity condition at camera level: within a sector (a square group of 64 pixels), three pixels must satisfy the coincidence requirement (Funk et al., 2004). In stereoscopic operation, events are typically recorded only if at least two telescopes trigger within a coincidence window of 80 ns (Funk et al., 2004). In contrast, CT5 can be read out in a single-telescope configuration (mono mode), while observations using all five telescopes are often referred to as hybrid observations (H.E.S.S. Collaboration, 2025).

3.3 Water Cherenkov detector arrays: detection physics and the SWGO concept

In contrast to IACTs, which image the air-Cherenkov light emitted during shower development, Water Cherenkov Detector (WCD) arrays infer the primary particle properties by sampling the secondary shower footprint at ground (cf. figure 3.3). In this approach, relativistic charged particles reaching the detector plane generate Cherenkov light inside water volumes instrumented with fast photosensors. The measured signal amplitudes and arrival times across the array encode the shower-front geometry, the lateral particle distribution, and, through composition-dependent muon content, information useful for γ /hadron separation. (SWGO Collaboration, 2025)

3.3.1 Detection principle and reconstruction observables in WCDs

A WCD registers Cherenkov photons produced when a charged particle traverses water with $\beta n_w(\lambda) > 1$ (cf. section 3.1.2). Since refractive index of water, $n_w \simeq 1.33$ in the optical, the Cherenkov angle for ultrarelativistic particles is large ($\theta_c \approx 41^\circ$), and a significant fraction of the light is emitted toward the tank walls and floor, where it is redistributed by reflections and recorded by photomultiplier tubes (PMTs). The fundamental observables are (i) the integrated charge (or calibrated amplitude) per

³The pixel threshold is expressed in photoelectrons (p.e.). Following Funk et al. (2004), 1 p.e. is defined as the mean amplitude of a single-photoelectron signal at the pixel comparator (about 21 mV), providing a convenient calibration unit for the discriminator threshold.

detector, which tracks the local particle content and deposited energy, and (ii) the signal time, which traces the shower-front arrival time at that detector position. (SWGOCollaboration, 2025)

At array level, direction reconstruction is typically obtained by fitting a curved shower-front model to the measured hit times, accounting for per-channel timing offsets and signal propagation delays. The shower core location and energy proxy are inferred from the lateral distribution of signal amplitudes (e.g. a fitted lateral distribution function or a size parameter defined at a reference distance). The same amplitude pattern also provides background-discriminating information: hadron-induced showers exhibit larger intrinsic fluctuations and, crucially, a higher muon content at ground than purely electromagnetic cascades (cf. section 3.1.1). Since muons tend to deposit energy in compact, track-like patterns through the detector volume, they can be exploited as handles for γ /hadron separation when the detector design provides sensitivity to penetrating particles. (SWGOCollaboration, 2025)

3.3.2 SWGO: site and array concept

The Southern Wide-field Gamma-ray Observatory (SWGOCollaboration, 2025) is a planned southern-hemisphere, wide-field γ -ray observatory based on a large array of WCDs. Its baseline site is at Pampa La Bola in the Atacama Astronomical Park in Chile (cf. figure 3.6), selected to provide a high-altitude environment and logistical access suitable for a $\sim \text{km}^2$ -scale deployment. The design is explicitly optimized for continuous monitoring with large instantaneous field of view and high duty cycle, complementing pointed IACT observations. (SWGOCollaboration, 2025)

The baseline array layout adopts a three-zone configuration in which the tank density decreases with radius (cf. figure 3.6). The central, high-fill-factor region provides the low-energy reach (from hundreds of GeV upward), while the intermediate and outer zones extend the instrumented area and effective area at multi-TeV energies. In the current baseline, the central zone fill factor⁴ is 70%, surrounded by zones with fill factors of 4% and 1.7%, respectively, yielding a total instrumented area of approximately 1 km^2 and a total of 3763 WCDs (SWGOCollaboration, 2025). The staged density profile is therefore a detector-level implementation of the physics requirement that higher-energy showers illuminate larger footprints and benefit from larger sampling area.

Since the publication of SWGOCollaboration (2025), the Analysis & Simulations working group has settled on an updated reference configuration for Milestone 7 (M7), denoted D8, based on a double-layer “D” tank design and a three-zone layout (Conceição, Glombitza, et al., 2025), which is what this thesis will mainly focus on. In D8 the zone fill factors are 70%, 4.0%, and 1.7% with outer radii of 156 m, 400 m, and 560 m, respectively, corresponding to a total of 3763 tanks (approximately 1 km^2 coverage) (Conceição, Glombitza, et al., 2025). Compared to earlier layouts, the zone 1 density is slightly reduced, while zones 2 and 3 are slightly increased to preserve direction reconstruction performance, which depends on tank spacing and not only on the fill

⁴The fill factor denotes the fraction of the geometric array area covered by detector units; a high fill factor corresponds to dense sampling that improves low-energy triggering and reconstruction, while lower fill factors extend the instrumented area for efficient detection of large high-energy shower footprints.

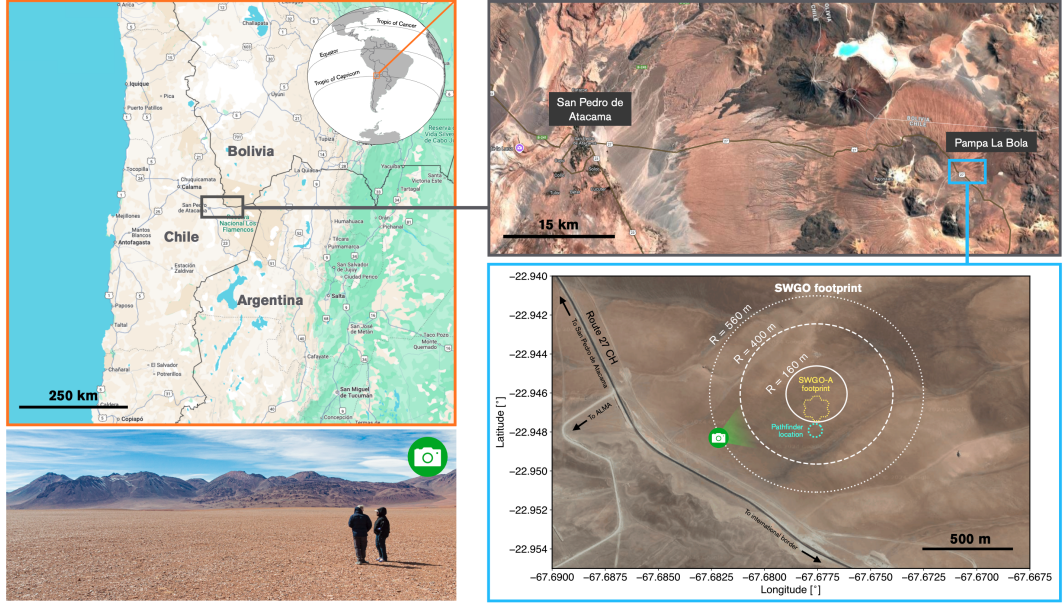


Figure 3.6: Proposed SWGO site at Pampa La Bola (Atacama Astronomical Park, Chile) and indicative footprints for the staged deployment and the full array concept. The baseline design adopts a dense core with lower-density outer regions to sample the increasing lateral extent of the ground particle footprint with primary energy. Adapted from SWGO Collaboration (2025).

factor (Conceição, Glombitza, et al., 2025).

3.3.3 SWGO WCD unit design and detector-level physics

The baseline SWGO central-array detector unit is a cylindrical steel tank (radius 5.2 m, height 4.1 m) assembled on site and lined with a bladder that contains a total water height of 4.0 m. The bladder is divided into two optically separated volumes: an upper water volume for efficient collection of shower-electromagnetic signals and a lower volume of height 0.8 m designed to enhance sensitivity to penetrating muons (cf. figure 3.7a). The inner surfaces employ highly reflective materials (e.g. Tyvek) to increase light collection and reduce position-dependent response variations. (SWGO Collaboration, 2025)

In the M7 reference configuration, both the D8 array and its pathfinder-scale counterpart D9 employ a double-layer D-tank unit design (Conceição, Glombitza, et al., 2025). Deliverable D7.2 specifies this dual-layer station with a diameter of 5.20 m and a total height of 4.1 m (cf. figure 3.7b), with an 8 inch high-quantum-efficiency PMT in the upper chamber and an 8 inch Hamamatsu R5912 PMT in the lower chamber (Conceição, Glombitza, et al., 2025). The upper chamber is described with white walls and a black base, while the lower chamber is entirely white, reflecting an optimization of light collection and background rejection in the two channels (Conceição, Glombitza, et al., 2025).

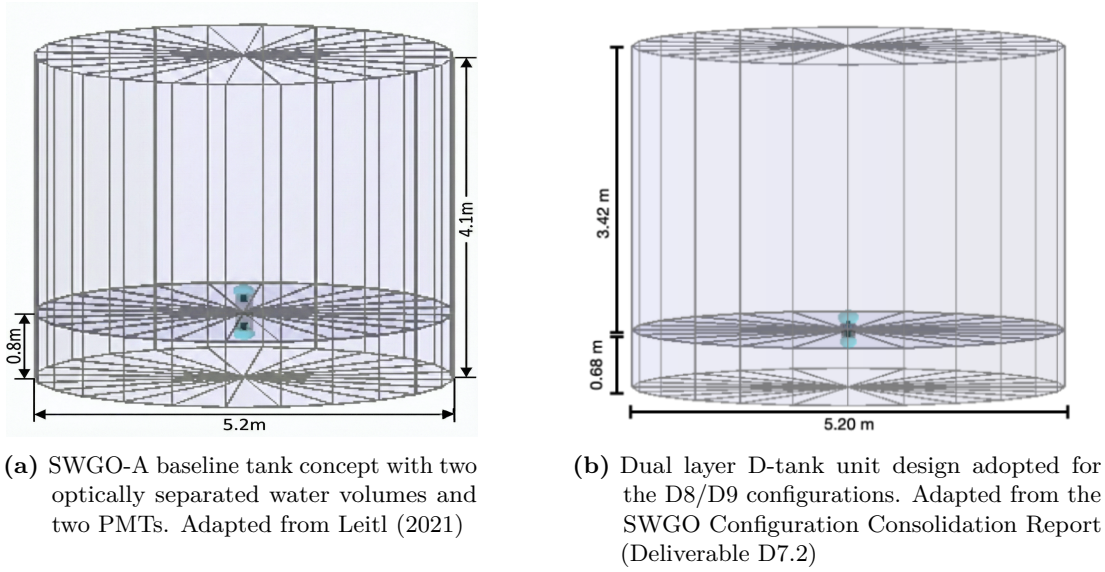


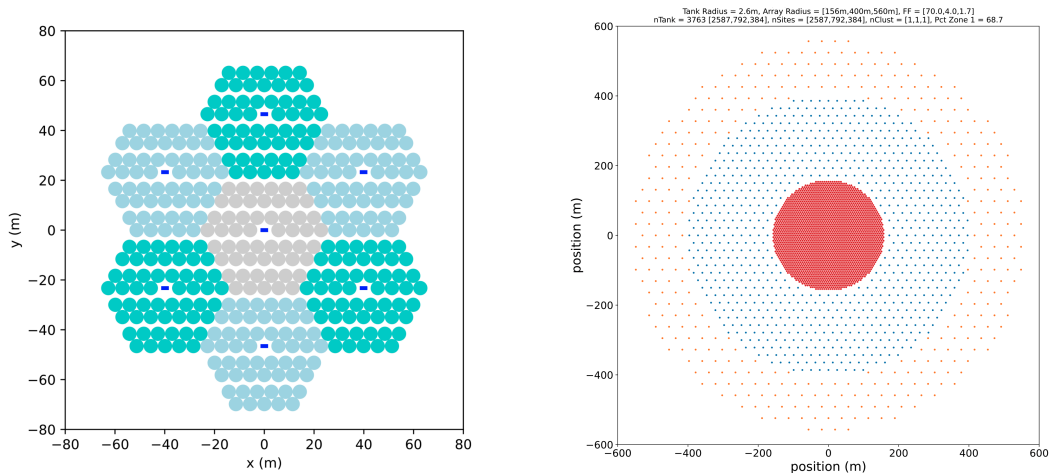
Figure 3.7: Comparison of SWGO WCD unit concepts. Figure (a) shows the SWGO-A pathfinder oriented A-tank concept. Figure (b) shows the dual-layer D-tank unit design used as the baseline for the D8/D9 configurations, featuring an upper chamber optimized for the electromagnetic shower component and a lower chamber intended to enhance the response to penetrating muons for γ /hadron separation.

A double PMT module is deployed per tank, with one large PMT viewing the upper volume and a second PMT instrumenting the lower volume, facing downwards. The lower chamber provides muon-tagging capability, which is a direct detector level response to the physical origin of the dominant background: hadron-induced air showers contain more muons at ground than γ induced cascades. By separating a muon enriched response channel from the electromagnetic dominated response, the design enables γ /hadron separation using observables built from the relative light yield and timing between the two compartments. (SWGO Collaboration, 2025)

3.3.4 Readout, timing synchronization, and triggering

The SWGO baseline readout architecture routes PMT signals from each WCD to nearby field nodes that provide digitization, local control, and timing distribution. In the baseline design, PMT waveforms are digitized using FlashCam electronics operated at 250 M Sa s^{-1} with deep buffering, enabling waveform-level extraction of arrival times and integrated charges. A first-level trigger is formed at the channel level using a threshold at the fraction-of-photoelectron scale, placing requirements on electronic noise and gain stability. (SWGO Collaboration, 2025)

To achieve the sub-nanosecond relative timing accuracy needed for shower front fits, the baseline design uses a White Rabbit timing system to distribute a common time and frequency reference across the array. In configurations where high voltage is delivered over shared cabling, dedicated hardware measures per-channel signal propagation delays along the cable (time-domain reflectometry) to correct for channel-to-channel transit times and temperature-dependent variations (SWGO Collaboration, 2025). These



(a) SWGO-A planned layout (cluster-based concept). Adapted from SWGO Collaboration (2025). (b) Baseline D8 layout with three radial zones and varying fill factor. Adapted from the SWGO Configuration Consolidation Report (Deliverable D7.2)

Figure 3.8: Comparison of SWGO array-layout concepts. Figure (a) shows the planned SWGO-A layout based on a cluster and field-node architecture. Figure (b) shows the baseline D8 configuration, implemented as a multi-zone array with decreasing fill factor at increasing radius. Adapted from SWGO collaboration design documents.

components implement a practical measurement chain from the fast, nanosecond-scale Cherenkov emission to calibrated, array-synchronized hit times used in direction reconstruction.

3.3.5 First-stage deployment: SWGO-A

The first construction stage, SWGO-A, is conceived as a dense segment of the inner array to validate deployment procedures, integration of readout and timing, and simulation-to-data consistency (cf. figure 3.8a). In the SWGO white paper, this stage is described as 385 WCDs arranged in seven clusters of 55 tanks with an overall fill factor of 65%, each cluster connected to a single field node (SWGO Collaboration, 2025).

In the M7 consolidation, SWGO-A is instead represented by the dedicated pathfinder configuration D9, defined as a reduced version of D8 with 361 stations and an array diameter of 58 m (Conceição, Glombitza, et al., 2025). D9 is intended to cover part of the dense inner zone and to prioritize low-energy performance while maintaining the same tank concept and design assumptions as the baseline configuration (Conceição, Glombitza, et al., 2025). Tank spacing and service aisles are incorporated in the layout to enable installation and maintenance while preserving dense sampling close to threshold (SWGO Collaboration, 2025).

3.3.6 Complementarity of IACTs and WCDs arrays

The IACT and WCD approaches discussed in this chapter provide complementary access to the same underlying physical object: the secondary radiation fields and particle footprint produced by a Extensive Air Shower (EAS) in the atmosphere. IACTs form a short-exposure image of the Cherenkov light emitted during shower development, yielding a compact set of image-shape and timing observables per triggered telescope and enabling stereoscopic reconstruction when multiple telescopes view the same event (cf. section 3.2). WCDs, in contrast, sample the shower at the detector plane, recording charge and timing information across an extended instrumented area and providing direct sensitivity to the particle content at ground, including the muonic component (cf. section 3.3). (SWG0 Collaboration, 2025)

These different measurement strategies naturally emphasize different regions of parameter space. Imaging of the air Cherenkov flash offers event-wise morphological information that supports precise direction reconstruction and powerful γ /hadron separation, at the cost of a limited instantaneous field of view and an observational duty cycle constrained by night time and atmospheric conditions. Ground-based WCD arrays, by construction, operate with a wide field of view and near continuous duty cycle, making them well suited for time-domain monitoring, unbiased surveys, and the detection of transient or long-duration emission. (SWG0 Collaboration, 2025)

For the transient phenomena considered in this thesis, these trade-offs motivate a two-pronged strategy. Wide field monitoring can identify and follow evolving emission on timescales comparable to the source variability, while pointed IACT observations provide complementary measurements with higher per-event reconstruction fidelity for targeted follow-up. The remainder of the thesis therefore uses the detector physics framework established in this chapter as the basis for the analysis methods and datasets introduced next.

Chapter 4

Detectability of novae with SWGO

Chapter 2 established novae as time-dependent particle accelerators in which shocks can power non-thermal emission, while section 3.3 summarised the WCD measurement concept and the design goals of SWGO. Building on this framework, we now quantify under which assumptions nova-like VHE γ -ray emission would be detectable with SWGO in a forward-modelling approach.

The analysis presented in this chapter combines (i) a parameterised nova-emission model, implemented as a dedicated `Nova` class¹ to explore how various physical parameters may or may not depend directly on each other and evolve with time, with (ii) SWGO Instrument Response Functions (IRFs) to eventually calculate the expected integral sensitivity of SWGO to novae. Since nova outbursts evolve on timescales of days to weeks, we focus on integration timescales relevant for transient monitoring and adopt a reference exposure of 30 days for the baseline sensitivity calculations. We first validate the pipeline using a worked example based on the symbiotic nova RS Oph, where spectral energy distributions are computed with GAMERA (Hahn, 2016) (see section A.1), and then generalise to sensitivity/detectability plots that include the dependence on observing conditions, in particular zenith angle and shower-core distance in the IRFs².

4.1 Scope and analysis overview

The goal of this chapter is to provide a quantitative detectability assessment for representative nova-like emission scenarios under the current SWGO design and analysis assumptions. This chapter does not aim to predict the number of novae detectable by SWGO. Instead, we evaluate the detectability of representative nova-like emission scenarios by folding a parameterised source model through SWGO IRFs and identifying which physical parameters (e.g. distance, shock velocity etc.) most strongly affect whether a detection criterion is met. Unless stated otherwise, the instrument concept and baseline performance assumptions follow the SWGO white paper (SWGO Collaboration, 2025).

¹Written in Python (cf. section A.2).

²Throughout this chapter, we use the term “detectability” to denote the comparison of model-predicted flux levels to a sensitivity criterion defined in section 4.5.

4.2 A parametric nova-shock model: implementation in a custom Nova class

To explore the detectability of novae and nova-like transients with ground-based γ -ray instruments, we require a model that maps a compact set of physical parameters (ejecta mass, shock velocity, circumstellar density, and microphysical acceleration parameters) onto derived quantities such as the target density for hadronic interactions, characteristic timescales, and the maximum particle energy. We therefore implemented a lightweight, parametric nova model as a Python class (`Nova`). The baseline prescriptions for shock expansion, densities, and energetics follow the RS Oph modelling in MAGIC Collaboration (2022), while the broader physical context (radiative shocks, dense environments, and the connection between shocks and optical emission) is summarized in the review by Laura Chomiuk et al. (2021).

4.2.1 Model parameters

The `Nova` class is parameterized by the time since outburst t , a characteristic shock velocity v_{sh} , the ejecta mass M_{ej} , and (where applicable) a red-giant wind described by a mass-loss rate \dot{M} and wind speed v_{wind} . The model also includes a geometric shell-thickness parameter h (we assume an ejecta layer of thickness $\Delta R = hR_{\text{sh}}$), and microphysical parameters entering particle acceleration and losses (magnetic field B , acceleration parameter ξ , and the inelastic pp cross section σ_{pp}). Unless stated otherwise, the default normalization corresponds to RS Oph-like values (e.g. $v_{\text{sh}} \sim 4500 \text{ km s}^{-1}$ at early times and $t \sim 3 \text{ d}$; see MAGIC Collaboration 2022 for the adopted reference parameters).

4.2.2 Shock kinematics

Following MAGIC Collaboration (2022), the shock radius is related to the shock velocity by

$$R_{\text{sh}}(t) \simeq 1.2 \times 10^{14} \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}} \right) \left(\frac{t}{3 \text{ d}} \right) \text{ cm}, \quad (4.1)$$

which corresponds to near-linear expansion with approximately constant v_{sh} at early times.

In addition, we implemented an optional phenomenological deceleration model, parameterized as a broken power law in time (constant velocity up to t_{dec} , followed by $v_{\text{sh}} \propto t^{-m_{\text{dec}}}$):

$$v_{\text{sh}}(t) = \begin{cases} v_0, & t \leq t_{\text{dec}}, \\ v_0 \left(\frac{t}{t_{\text{dec}}} \right)^{-m_{\text{dec}}}, & t > t_{\text{dec}}, \end{cases} \quad (4.2)$$

The resulting time evolution of the shock velocity for representative parameter choices is shown in Figure 4.1.

The corresponding shock radius is obtained from $R_{\text{sh}}(t) = \int_0^t v_{\text{sh}}(t') dt'$, which we evaluate analytically. This extension is used to test the sensitivity of derived quantities

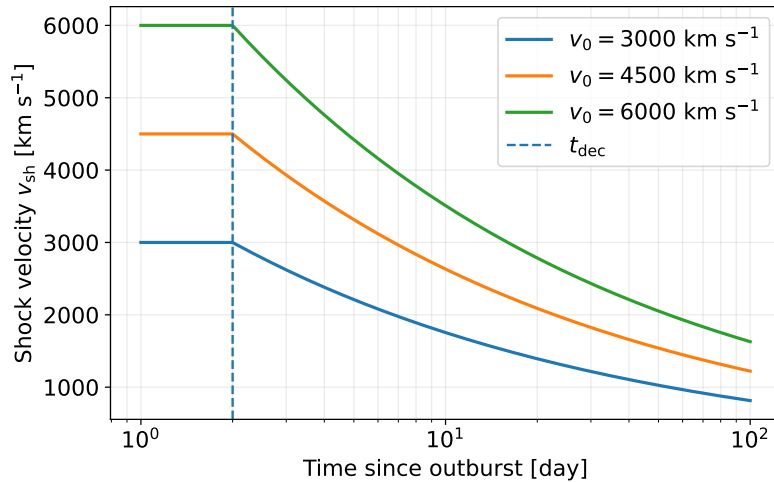



Figure 4.1: Shock velocity evolution $v_{\text{sh}}(t)$ in the Nova class for a constant-velocity model and for the deceleration method in Equation (4.2). In the decelerating case, the transition occurs at $t_{\text{dec}} = 2$ d (with slope set by m_{dec}).

to deceleration while keeping the RS Oph-normalized scalings as the baseline. Figure 4.2 illustrates the resulting shock-radius evolution. The constant-velocity case yields $R_{\text{sh}} \propto t$, while enabling deceleration produces a slower late-time growth for $t > t_{\text{dec}}$.



(a) Constant-velocity expansion for different v_0 . (b) Same v_0 values with deceleration enabled ($t_{\text{dec}} = 2$ d).

Figure 4.2: Shock radius evolution $R_{\text{sh}}(t)$ in the Nova class. Left: constant-velocity expansion consistent with the RS Oph-normalized scaling in MAGIC Collaboration (2022). Right: effect of enabling the broken power-law deceleration model (see Equation (4.2)).

4.2.3 Target densities: ejecta and wind

Hadronic γ -ray production depends on the density of target material encountered by accelerated protons. We include two components: (i) the expanding nova ejecta and (ii) the circumstellar wind of the companion (relevant for symbiotic systems such as RS Oph).

Assuming the ejecta occupy a thin shell of thickness h at radius R_{sh} , the ejecta number density is

$$n_{\text{ej}}(t) = \frac{M_{\text{ej}}}{4\pi h R_{\text{sh}}^3 m_p} \simeq 6.0 \times 10^8 \left(\frac{M_{\text{ej}}}{10^{-6} M_{\odot}} \right) \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}} \right)^{-3} \left(\frac{t}{3 \text{ d}} \right)^{-3} \left(\frac{h}{0.1} \right)^{-1} \text{ cm}^{-3}, \quad (4.3)$$

and the wind density of the secondary star is

$$n_{\text{wind}}(t) = \frac{\dot{M}}{4\pi R_{\text{sh}}^2 v_{\text{wind}} m_p} \simeq 1.1 \times 10^8 \left(\frac{\dot{M}}{5 \times 10^{-7} M_{\odot} \text{ yr}^{-1}} \right) \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}} \right)^{-2} \left(\frac{t}{3 \text{ d}} \right)^{-2} \left(\frac{v_{\text{wind}}}{10 \text{ km s}^{-1}} \right)^{-1} \text{ cm}^{-3}. \quad (4.4)$$

These expressions and their normalizations follow MAGIC Collaboration (2022). As a sanity check, upon plotting the number density for both the nova ejecta and the secondary star wind density, we observe that the slope for the $n_{\text{ej}}(t)$ is much steeper and drops quicker than that of $n_{\text{wind}}(t)$, indicating that the number density in the expanding ejecta shell declines much quicker, as expected.




Figure 4.3: Time evolution of the number densities used in the Nova class for an RS Oph-like default normalization. The ejecta shell density $n_{\text{ej}}(t)$ follows Equation (4.3) and decreases as $n_{\text{ej}} \propto R_{\text{sh}}^{-3} \propto t^{-3}$ for constant-velocity expansion, while the red-giant wind density $n_{\text{wind}}(t)$ follows Equation (4.4) and decreases as $n_{\text{wind}} \propto R_{\text{sh}}^{-2} \propto t^{-2}$. The different slopes reflect geometric dilution of the expanding ejecta shell compared to the $1/r^2$ density profile of a steady wind.

In the Nova class, we also define a simple “target density” proxy

$$n_{\text{target}}(t) = \max [n_{\text{ej}}(t), n_{\text{wind}}(t)], \quad (4.5)$$

which avoids discontinuities when one component dominates, while still reflecting the controlling interaction scale.

4.2.4 Timescales and hadronic interaction efficiency

The inelastic pp cooling timescale is

$$t_{pp} = \frac{1}{n_{\text{target}} c \sigma_{pp}}, \quad (4.6)$$

consistent with the RS Oph scaling used in MAGIC Collaboration (2022). In the ballistic regime, the dynamical timescale is $t_{\text{dyn}} \equiv t$. We use the bounded proxy

$$f_{pp} = 1 - \exp\left(-\frac{t_{\text{dyn}}}{t_{pp}}\right), \quad (4.7)$$

which interpolates smoothly between $f_{pp} \approx t_{\text{dyn}}/t_{pp}$ for inefficient interactions and $f_{pp} \rightarrow 1$ for calorimetric conditions.

4.2.5 Energetics

For consistency with the RS Oph modelling, we adopt the ejecta kinetic energy scaling

$$E_k \simeq \frac{1}{2} M_{\text{ej}} v_{\text{sh}}^2 \simeq 2.0 \times 10^{44} \left(\frac{M_{\text{ej}}}{10^{-6} M_{\odot}}\right) \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}}\right)^2 \text{ erg}, \quad (4.8)$$

and the RS Oph-normalized proton energy requirement (from the fit-based normalization in MAGIC Collaboration 2022).

$$E_{p,\text{nova}} \simeq 0.44 \times 10^{44} \left(\frac{M_{\text{ej}}}{10^{-6} M_{\odot}}\right)^{-1} \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}}\right)^3 \left(\frac{d}{2.45 \text{ kpc}}\right)^{-2} \left(\frac{h}{0.1}\right) \text{ erg}. \quad (4.9)$$

This implies a conversion efficiency

$$\epsilon_p \equiv \frac{E_{p,\text{nova}}}{E_k} \simeq 0.22 \left(\frac{M_{\text{ej}}}{10^{-6} M_{\odot}}\right)^{-2} \left(\frac{v_{\text{sh}}}{4500 \text{ km s}^{-1}}\right) \left(\frac{d}{2.45 \text{ kpc}}\right)^{-2} \left(\frac{h}{0.1}\right), \quad (4.10)$$

which we use as a convenient normalization when comparing different nova parameter choices. Figure 4.4 visualizes equation (4.8) in the $(M_{\text{ej}}, v_{\text{sh}})$ plane and indicates where representative novae fall relative to the RS Oph-like normalization.

While figure 4.4 highlights the location of representative novae in the $(M_{\text{ej}}, v_{\text{sh}})$ plane and the implied kinetic-energy scale, we additionally use the RS Oph-based normalization in equations (4.9) and (4.10) to evaluate the kinetic shock energy to proton energy conversion efficiency. As a consistency check of the implementation, figure 4.5 shows that the RS Oph reference point lies at $\epsilon_p \simeq 0.2$ ($\sim 20\%$), consistent with the RS Oph calibrated value adopted from MAGIC Collaboration (2022).

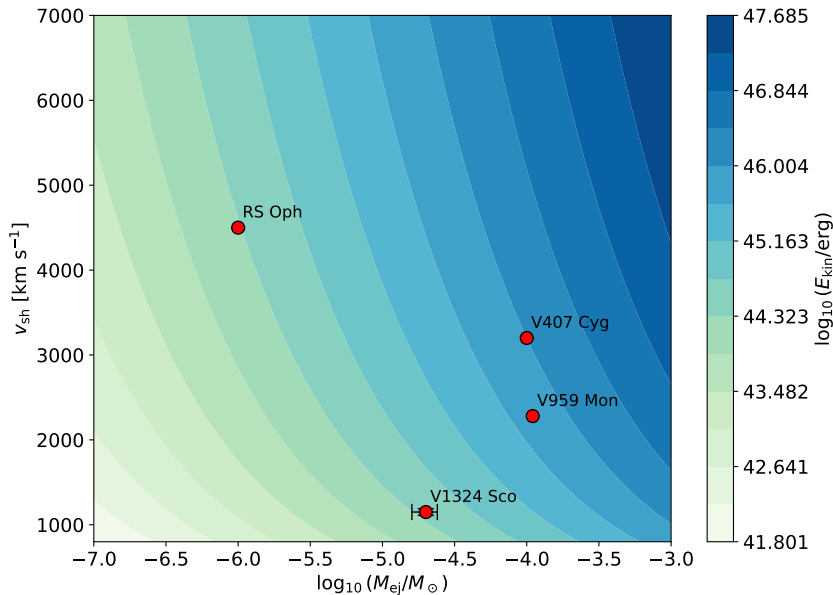


Figure 4.4: Contour plot of the outflow kinetic energy $E_k = \frac{1}{2}M_{\text{ej}}v_{\text{sh}}^2$ from Equation (4.8) in the $(M_{\text{ej}}, v_{\text{sh}})$ plane. The overplotted markers indicate representative literature values for selected novae (listed in table D.3). At fixed E_k , the degeneracy follows $M_{\text{ej}} \propto v_{\text{sh}}^{-2}$.

4.2.6 Acceleration and maximum energy

We parametrize particle acceleration using the acceleration rate

$$\left(\frac{dE}{dt}\right)_{\text{acc}} = \xi c \frac{E}{R_L(E)}, \quad (4.11)$$

with corresponding timescale

$$t_{\text{acc}} \simeq 3.9 \left(\frac{E}{300 \text{ GeV}}\right) \left(\frac{\xi B}{10^{-7} \text{ G}}\right)^{-1} \text{ d}, \quad (4.12)$$

as in MAGIC Collaboration (2022). In the RS Oph parameter setting, the maximum proton energy is typically limited by the available time since outburst rather than by pp losses, motivating an “escape/age-limited” interpretation for E_{max} , appropriately implemented in the Nova class. Figure 4.6 summarizes the maximum proton energy E_{max} predicted by the acceleration prescription implemented in the Nova class as a function of shock velocity and mass loss rate by wind velocity \dot{M}/v_{wind} .

4.3 SWGO instrument response functions

To translate a physical source model into an expected detection significance, we require an instrument-level description of how γ -ray showers are detected and reconstructed. In




Figure 4.5: Kinetic energy of the outflow (equation (4.8)) and RS Oph calibrated proton conversion efficiency (contours; $\epsilon_p \equiv E_{p,nova}/E_k$ from equations (4.9) and (4.10)) in the (M_{ej}, v_{sh}) plane. The RS Oph reference point is plotted for comparison with the adopted RS Oph-based normalization (MAGIC Collaboration, 2022).

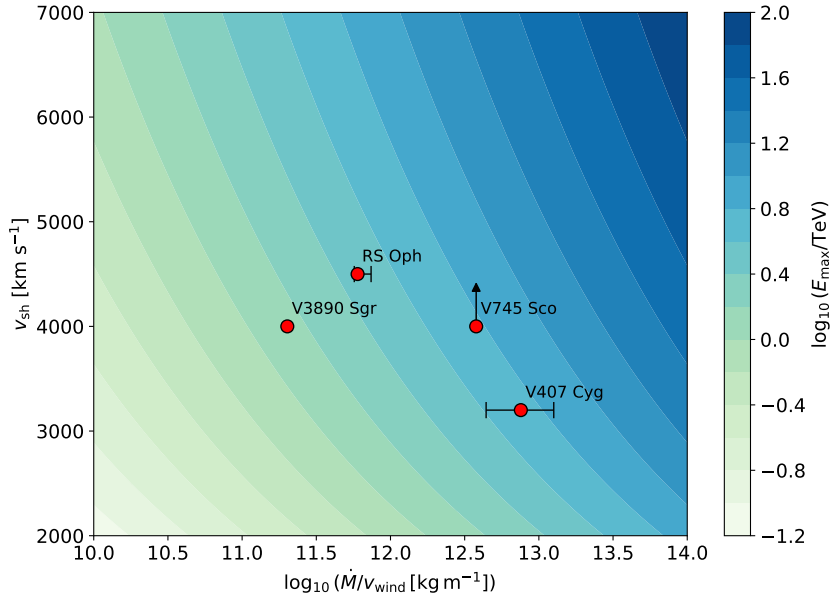


Figure 4.6: Contour map of the maximum proton energy E_{max} predicted by the Nova class as a function of shock velocity v_{sh} and mass loss rate by wind velocity \dot{M}/v_{wind} . The colour scale shows $\log_{10}(E_{max}/\text{TeV})$, computed using the analytic escape-limited scaling implemented in Nova. **E_max:** $E_{max} \propto (\dot{M}/v_{wind})^{1/2} v_{sh}^2$, with default parameters $Z = 1$ and $\xi_{esc} = 0.01$. Red markers indicate representative literature values for selected novae (cf. table D.2); horizontal error bars denote ranges in \dot{M}/v_{wind} where available, and arrows denote one-sided limits.

the standard high-energy analysis framework this information is captured by Instrument Response Functions (IRFs), which provide the conditional mapping from true photon properties to reconstructed event distributions.

For a ground-based air-shower array such as SWGO, the relevant IRF components are:

(i) the effective area $A_{\text{eff}}(E, \theta, \dots)$, which converts incident photon flux into an expected event rate as a function of true energy E and observing conditions (e.g. zenith angle θ);

(ii) the Point Spread Function (PSF), describing the distribution of reconstructed directions around the true arrival direction;

(iii) the energy dispersion, describing the probability to reconstruct an event with energy E_{reco} given true E ; and

(iv) the residual background rate after γ /hadron separation cuts. In practice these IRFs are derived from Monte-Carlo simulations of showers and detector response, followed by the same reconstruction and event-selection chain used for analysis.

Because SWGO observes continuously with a wide field of view, the exposure for a fixed celestial position is accumulated while that position transits through the local sky. This implies that the response relevant for a given source is not described by a single zenith angle, but by an effective response integrated along the zenith-angle trajectory over the chosen time window. In this work we adopt a reference exposure of 30 days, motivated by the typical days-to-weeks evolution timescales of nova high-energy emission (as discussed earlier in chapter 1).

We follow the `pyswgo` IRF-production chain (as described in SWGO Collaboration (2026)) to see how analysis-ready IRFs can be obtained in two steps. First, reconstructed event samples are converted to event-level quantities and classified to define energy-dependent quality and γ /hadron-separation cuts. Second, these selections are used to build *table IRFs*, i.e. zenith-binned response tables together with the associated GADF-style index files for high-level analysis. Finally, to obtain responses appropriate for a custom integration time, the table IRFs are projected into the sky and integrated over the selected time interval to produce *map IRFs* (effective exposure, background maps, PSF and energy-dispersion maps) for a given sky position and observation period. The resulting map IRFs are then used to compute differential sensitivity curves using `pyswgo-sensitivity` (usage shown below) for the 30-day exposure, which we use below to evaluate nova detectability.

```

1 usage: pyswgo-sensitivity [-h] --config CONFIG [--binning BINNING]
2                             [--input_irf_path INPUT_IRF_PATH]
3                             [--output_path OUTPUT_PATH] [--log-level LOG_LEVEL]
4                             [--overwrite] [--dataset-wise] [--ra RA] [--dec DEC]
5                             [--n-transits N_TRANSITS]

```

4.4 RS Oph as an example: GAMERA SED and SWGO comparison

In the previous section, we derived the SWGO differential sensitivity curves from the 30 d IRFs. These sensitivity tables provide the minimum detectable energy flux (expressed as $E^2 dN/dE$) as a function of reconstructed energy under the analysis assumptions used in the IRF production. In the following, we use these results to perform a simple end-to-end detectability check using a representative nova scenario.

Model setup and photon production

As a worked example, we consider parameters motivated by the recurrent nova RS Oph adapted mainly from MAGIC Collaboration (2022), H.E.S.S. Collaboration (2022), and Laura Chomiuk et al. (2021). The non-thermal proton population is constructed using our custom nova model described in section 4.2, which provides (i) the total energy injected into relativistic protons, E_p^{nova} , and (ii) an effective maximum proton energy E_{max} . We assume a power-law proton spectrum with an exponential cutoff

$$\frac{dN_p}{dE_p} \propto \left(\frac{E_p}{E_{\text{ref}}}\right)^{-\alpha} \exp\left(-\frac{E_p}{E_{\text{max}}}\right), \quad (4.13)$$

with spectral index α . The normalization is chosen such that the total energy in protons equals the value returned by our nova model

$$\int E_p \frac{dN_p}{dE_p} dE_p = E_p^{\text{nova}}, \quad (4.14)$$

i.e. the spectrum is renormalized numerically to satisfy the energy limit predicted for the RS Oph like parameter set.

We then use GAMERA to compute the resulting γ ray emission from inelastic proton-proton interactions in an ambient target medium. In this calculation, the distance to the source is set by the RS Oph-like configuration, such that GAMERA returns the spectral energy distribution (SED) at Earth. Since we only inject protons, the component shown in the comparison below corresponds to the hadronic p - p channel.

Comparison to SWGO differential sensitivity

Figure 4.7 shows the resulting hadronic SED from GAMERA (solid blue) together with the SWGO differential sensitivity curve derived for an assumed source index of -2 (black dashed)³, both plotted in terms of $E^2 dN/dE$.

The key point is that, in the SWGO-relevant energy range (roughly \gtrsim TeV), the RS Oph like hadronic SED lies below the 30 d differential sensitivity curve. Within this simplified steady-state setup, this implies that an RS Oph-like nova at the assumed

³See figure B.1 for a similar plot with both indices -2 and -3 .

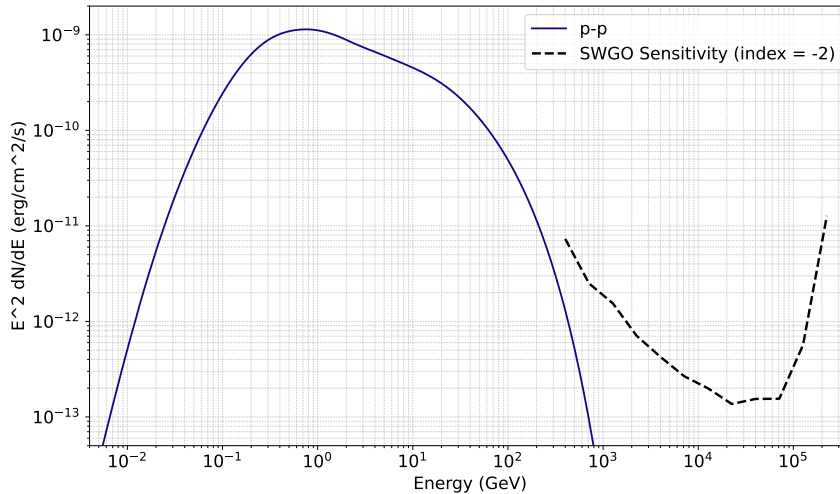


Figure 4.7: Hadronic γ -ray spectral energy distribution (SED) from p - p interactions computed with GAMERA for an RS Oph-like nova parameter set (solid blue), compared to the SWGO 30 d differential sensitivity for an assumed spectral index of -2 (black dashed). Both curves are shown as $E^2 dN/dE$ at Earth.

distance and energy limit would not be detectable by SWGO at the 30 d sensitivity level in the hadronic channel. Detectability would require a substantially larger γ -ray energy flux in the TeV range, which could in principle be achieved by combinations of a larger non-thermal proton energy limit, a closer distance, a harder/less-curved spectrum at high energies (e.g. larger E_{\max}), and/or environmental changes that enhance hadronic emission (e.g. higher effective target density and interaction efficiency).

4.5 Integral sensitivity and nova detectability map

To summarize nova detectability in a compact way, we convert the 30-day *differential* sensitivity obtained from the SWGO IRFs into an integral (energy-integrated) sensitivity threshold over 0.5-100 TeV, and compare this to the predicted energy flux from our hadronic (p - p) nova model across parameter space. Concretely, the differential sensitivity table/dataset is produced from the D8 IRFs (effective area, PSF, energy dispersion, and residual background), and provides the minimum flux per energy bin required to satisfy the adopted detection criterion. By integrating this threshold over 0.5-100 TeV (consistent with the same binning and spectral assumption used in the sensitivity calculation), we obtain a single reference energy-flux limit for a 30 day exposure.

Figure 4.8 shows the resulting detectability map in the (d, v_{sh}) plane. The colour scale indicates the model-predicted integrated p - p energy flux at Earth (in $\text{erg cm}^{-2} \text{s}^{-1}$) for each parameter combination, where the distance d is responsible for the geometric dilution and the shock velocity v_{sh} controls the available shock power and maximum particle energies in the model. The dotted white line marks the D8 integral sensitivity threshold, such that model events above (toward higher fluxes) would be expected to be detectable within the assumed 30 day exposure, while events below are likely to remain

undetected.

For perspective, we also plot representative literature estimates for selected novae (listed in table D.4). Their position with respect to the D8 threshold provides an order-of-magnitude estimate of whether similar systems would be detectable by SWGO under comparable conditions. The same integral sensitivity map produced with the D9 configuration is provided in figure B.2, with noticeably much worse sensitivity.

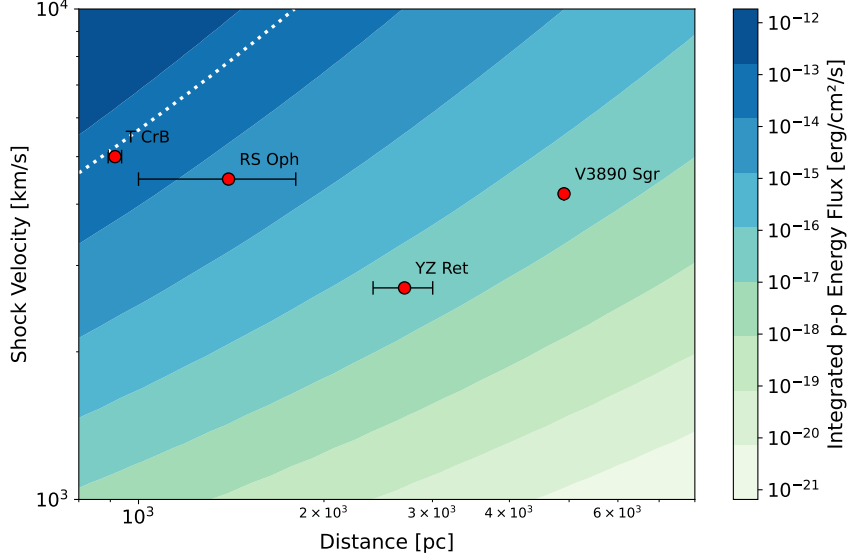


Figure 4.8: Nova detectability contour map for the SWGO D8 configuration based on the 30 day IRF-derived sensitivity. The colour scale shows the model predicted integrated p - p energy flux at Earth in the 0.5-100 TeV band as a function of distance d and shock velocity v_{sh} . The dotted white curve indicates the D8 integral sensitivity threshold for a 30 day exposure, $F_{\text{sens}}^{(\text{D8})} \simeq 7.366 \times 10^{-14} \text{ erg cm}^{-2} \text{ s}^{-1}$ (integrated over 0.5-100 TeV); parameter combinations above this line correspond to fluxes expected to be detectable under the adopted detection criterion. Red markers denote representative literature estimates for selected novae.

4.6 Summary and outlook

In this chapter we quantified the detectability of representative nova-like TeV emission scenarios with SWGO in a forward modelling approach. We combined a parametric nova shock model (implemented as a dedicated Python Nova class; see section A.2) with SWGO IRFs to (i) validate the analysis chain on a worked RS Oph example, and (ii) generalise to an integral sensitivity based detectability criterion for a 30 day exposure.

Using the D8 IRFs, we derived a reference integral sensitivity threshold by integrating the IRF based differential sensitivity over 0.5 TeV to 100 TeV and used this threshold to construct a detectability contour map in the (d, v_{sh}) plane (figure 4.8). The map provides a compact summary of how distance and the shock properties encoded by v_{sh} control whether a given source is expected to lie above or below the adopted detection criterion. Comparing representative literature estimates for selected novae to the D8

boundary indicates that, under the assumptions adopted here, only the most nearby and energetic systems are expected to approach the SWGO discovery space on 30 day timescales. Repeating the same procedure for the D9 configuration yields a higher detectability threshold (figure B.2). These results establish a working pipeline for translating a time dependent nova emission model into an SWGO specific detectability diagnostic.

Chapter 5

Optical light curves from H.E.S.S. Night Sky Background (NSB) data

The cameras of H.E.S.S. continuously record signals induced by the Night Sky Background (NSB) in each pixel¹. In standard H.E.S.S. analyses, NSB related observables (e.g. pedestal widths and pixel currents) are monitored to assess data quality, to flag problematic pixels, and to support calibration and stability studies (Aharonian et al., 2004). In this thesis, we repurpose these camera monitoring products to obtain a proxy for optical flux in order to extract time-resolved light curves of classical novae, with the nova MGAB-V207 (commonly known as YZ-Reticuli) as the primary case study. Figure 5.1 shows the optical lightcurve and the location of the nova in the constellation Reticulum.

This chapter describes the complete analysis chain from raw NSB data to calibrated light curves, and outlines the major instrumental and analysis-induced systematics that limit run-to-run continuity. The detector concept and H.E.S.S. instrumentation are introduced earlier in chapter 3; Here we focus on (i) building an optical proxy from pixel-level NSB information, (ii) applying corrections for differences between pixels and for atmospheric transparency, and (iii) using a data-driven quality-control method to find potentially faulty pixels and any offsets that cause sudden jumps in the data. The scripts used in this chapter can be found in section A.3.

We begin by defining the NSB derived observables used in this work and the associated run metadata (telescope ID, timestamps, pointing, offset etc.), together with the camera geometry needed to map pixel IDs to focal-plane coordinates.

5.1 NSB observables, run metadata, and camera geometry

This work is based on per-pixel monitoring outputs that quantify the NSB level in the H.E.S.S. cameras, together with run-level pointing information and a camera geometry description that maps pixel IDs to focal-plane coordinates. In the following, we define these analysis inputs and the conventions used throughout this chapter.

¹Sometimes obtained indirectly, for example, for CT1-CT4 it can be a derived value from the pedestal distribution

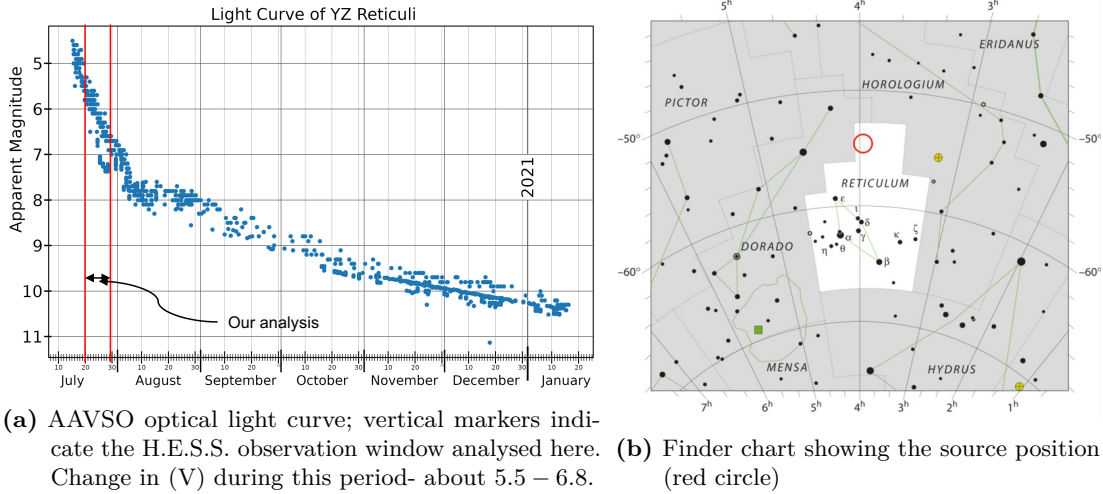


Figure 5.1: Context for the MGAB-V207 (YZ Ret) dataset used in this chapter: optical brightness evolution from AAVSO data, showing a decline in the overall apparent magnitude (V) from around 4.5 to 10.5 (left) and the source’s location on the sky (right). The location panel is adapted from IAU et al. (2025).

5.1.1 NSB time series format and sampling


For all telescopes considered here (CT1-CT5), the raw NSB time series are stored as a table with three columns: (i) a pixel identifier `pix_id` (unitless), (ii) an NSB rate in units of MHz, and (iii) a timestamp in UTC. Each row therefore represents the NSB estimate of a single pixel at a given time. The cadence differs between telescopes due to differences in the underlying monitoring configuration: for CT5 the time increment is 0.1 s, while for CT1-CT4 it is typically in the range 10 s to 20 s.

5.1.2 Run-level metadata and pointing configuration

To relate the per-pixel NSB measurements to a sky position and to identify changes in the camera illumination pattern between runs, we combine the NSB time series with run-level metadata obtained from an independent query script. The metadata include, at minimum, the run identifier, observation start time, duration, target coordinates, and the pointing offset configuration used for wobble observations (cf. section 5.2.1). For MGAB-V207, the relevant fields are shown in table D.1 and include `Target_RA`, `Target_Dec`, `Offset_Mode` and the offset components `Offset_x` and `Offset_y`.

5.1.3 Camera geometry and pixel-ID conventions

The reconstruction of nova light curves requires identifying and summing a compact cluster of pixels around the source position on the camera plane. For this purpose, we use a camera geometry file that maps each pixel identifier to a pair of focal-plane coordinates (x, y) . The pixel-ID to focal-plane coordinate mapping used throughout this chapter is illustrated for CT1-CT4 and CT5 in figure 5.2. For CT5, this mapping is provided as a plain-text table with columns `pix_id`, `x_coord`, and `y_coord`. For



(a) CT1–CT4 camera geometry. Each of the four 12m diameter telescopes has a camera with 960 pixels or PMTs. (b) CT5 FlashCam geometry. The single, larger 28m diameter telescope has a camera with 2,048 pixels.

Figure 5.2: Pixel layouts of the H.E.S.S. cameras used in this work (see figure 3.5 for the full telescope layout). Each hexagon corresponds to a single camera pixel; coordinates are shown in the focal-plane frame defined by the respective geometry files.

CT1-CT4, the available geometry information is distributed in the HAP convention, which uses a different pixel-ID numbering scheme than the NSB files used here. We therefore generated an equivalent CT1-CT4 geometry table in the same `pix_id;x;y` format by converting the HAP pixel IDs to the pixel-ID convention used in our NSB time series. This ensures that a given `pix_id` consistently refers to the same physical camera pixel across all analysis steps (geometry visualisation, cluster selection, and light curve extraction). Figure 5.2 illustrates how the pixel layout looks for both type of telescopes.

In section 5.2 we describe how the source pixel cluster is defined on the camera plane using this geometry information, and how the cluster definition differs between the CT5 exploratory analysis and the CT1-CT4 production pipeline.

5.2 Source cluster definition on the camera plane

The NSB derived optical proxy is constructed by integrating the NSB time series over a fixed set of camera pixels associated with the nova position (as we will discuss in section 5.3). In this work, the source localisation and pixel selection are performed directly in the camera plane by combining (i) run-wise pointing metadata (target coordinates and wobble offsets) as a guide, (ii) reference NSB maps from internal monitoring products (NSB maps from the HESS-wiki), and (iii) the pixel geometry mapping `pix_id` \mapsto (x, y) introduced in subsection 5.1.3. In the following, we describe the strategy used to define a reproducible pixel cluster for CT5 and CT1-CT4, together with the systems used to store the selected pixel lists for downstream corrections and

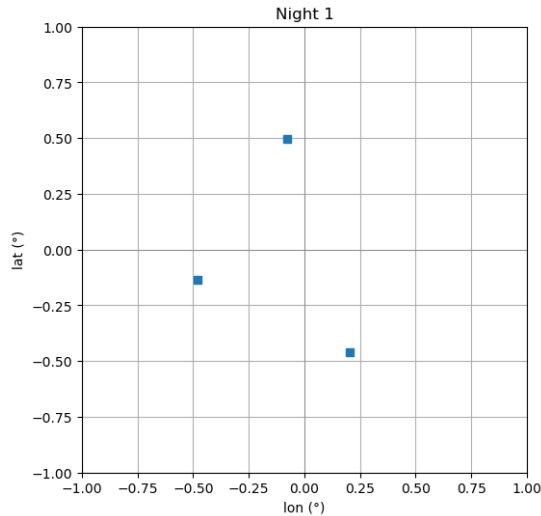


Figure 5.3: Example of run-wise source positions in camera coordinates for the first observation night (Run numbers: 161096, 161097, 161098) of MGAB-V207 (See table 5.1 for the pointing metadata and figure B.3 for similar position plots for all runs across all nights). Each marker corresponds to one run and represents the effective source displacement in the camera plane derived from the run metadata. The clustering of points reflects the discrete wobble configurations used during observations. Adapted from the HESS Wiki.

light curve construction.

5.2.1 Pointing, wobble and camera coordinates

H.E.S.S. observations of MGAB-V207 were performed in wobble mode, in which the source is placed at a desired offset from the camera centre to enable background control regions within the same field of view. For each run, we can get the typical pointing parameters from the `Monitor Run Data` table accessed via HAP tools, including the target coordinates $(\alpha_{\text{src}}, \delta_{\text{src}})$ and the commanded offsets $(\Delta\alpha, \Delta\delta)$ in the RA/Dec convention (fields `Target_RA`, `Target_Dec`, `Offset_Mode`, `Offset_x`, and `Offset_y`). Throughout this chapter, we use these values as reference to locate the pixel cluster contributing to the source (in this case, the nova MGAB-V207) (see section 5.2).

The numerical values of the offsets depend on the convention in which they are specified. In the RA/Dec convention, the longitudinal offset is given as a right-ascension difference $\Delta\alpha$, whereas the corresponding small-angle displacement in the tangent plane scales approximately as $\Delta\alpha \cos \delta$ at source declination δ . For MGAB-V207, $\delta_{\text{src}} \simeq -54.78^\circ$, such that $\cos \delta_{\text{src}} \approx 0.58$. Consequently, a nominal offset of $\Delta\alpha \simeq 0.87^\circ$ corresponds to an effective camera-plane displacement of $\Delta\alpha \cos \delta_{\text{src}} \simeq 0.50^\circ$, consistent with the expected values (cf. table 5.1). Small run-to-run deviations around the nominal wobble locations are expected from finite pointing accuracy and metadata conventions, and can be seen in figure 5.3.

Table 5.1: Run metadata (pointing related) for the MGAB-V207 ToO observations obtained using `Monitor_Run_Data` table accessed via HAP tools

run_id	Target_Name	Target_RA	Target_Dec	Offset_Mode	Offset_x	Offset_y	Run_Start_Time	Duration
161096	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-19 02:58:28	1683.0
161097	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-19 03:28:39	1683.0
161098	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-19 03:57:52	951.0
161131	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-20 02:48:50	1682.0
161132	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-20 03:18:04	1683.0
161133	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-20 03:47:17	1587.0
161163	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-21 02:38:46	1683.0
161164	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-21 03:08:50	1682.0
161165	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-21 03:38:09	2043.0
161195	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-22 02:08:42	1683.0
161196	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-22 02:38:40	1682.0
161197	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-22 03:07:53	1682.0
161198	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-22 03:37:04	1682.0
161224	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-23 02:10:13	1683.0
161225	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-23 02:40:13	1683.0
161226	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-23 03:09:25	1682.0
161227	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-23 03:38:38	2071.0
161268	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-24 02:08:40	1682.0
161269	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-24 02:38:28	1682.0
161270	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-24 03:07:40	1682.0
161271	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-24 03:36:55	2102.0
161306	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-25 02:39:17	1682.0
161307	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-25 03:09:21	1682.0
161308	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-25 03:38:34	2042.0
161334	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-26 02:37:45	1682.0
161335	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-26 03:07:01	1683.0
161336	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-26 03:36:17	2162.0
161366	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-27 02:40:55	1682.0
161367	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-27 03:10:05	1682.0
161368	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	-0.87	0.0	2020-07-27 03:39:16	1215.0
161426	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-28 02:28:21	1682.0
161427	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-28 02:58:26	1682.0
161428	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-28 03:27:40	1682.0
161476	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	0.5	2020-07-29 02:25:54	1682.0
161477	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.87	0.0	2020-07-29 02:55:09	1682.0
161478	ToO MGAB-V207	59.623125	-54.778111	RA/Dec	0.0	-0.5	2020-07-29 03:24:23	2823.0

5.2.2 Source localisation on the camera plane

For each run, we determine an approximate source location in the camera field of view from the run metadata obtained as described in the previous subsection, and validate this location against the NSB maps available via the H.E.S.S. internal web-summary tools. Within this constrained region, we define a seed pixel as the pixel with the maximum NSB rate in the run-wise NSB map. To verify that the selected seed pixel is consistent with a transient optical source rather than a static background feature, we inspect the time evolution of the NSB camera maps across all timestamps within the run. Based on this inspection, we define a fixed pixel cluster that encloses the apparent source footprint on the camera plane, augmented by an additional one-pixel-wide boundary layer of surrounding pixels. This augmented cluster is used in the subsequent light curve extraction to capture possible source flux and will be described in detail in section 5.2.3.

The detailed cluster-definition procedure differs between CT5 and CT1-CT4. For CT5, the pixel selection is obtained with an automated procedure adapted from Thomas' satellite-trail detection code (Lang, 2025) by modifying the selection cuts and optimisation parameters for nova-like variability. For CT1-CT4, where the automated selection did not generate stable tracks, the pixel selection was performed manually following the same seed pixel and boundary layer rule by making a frame-by-frame animation of each run using `Plotly` in Python. Both procedures and their associated limitations are

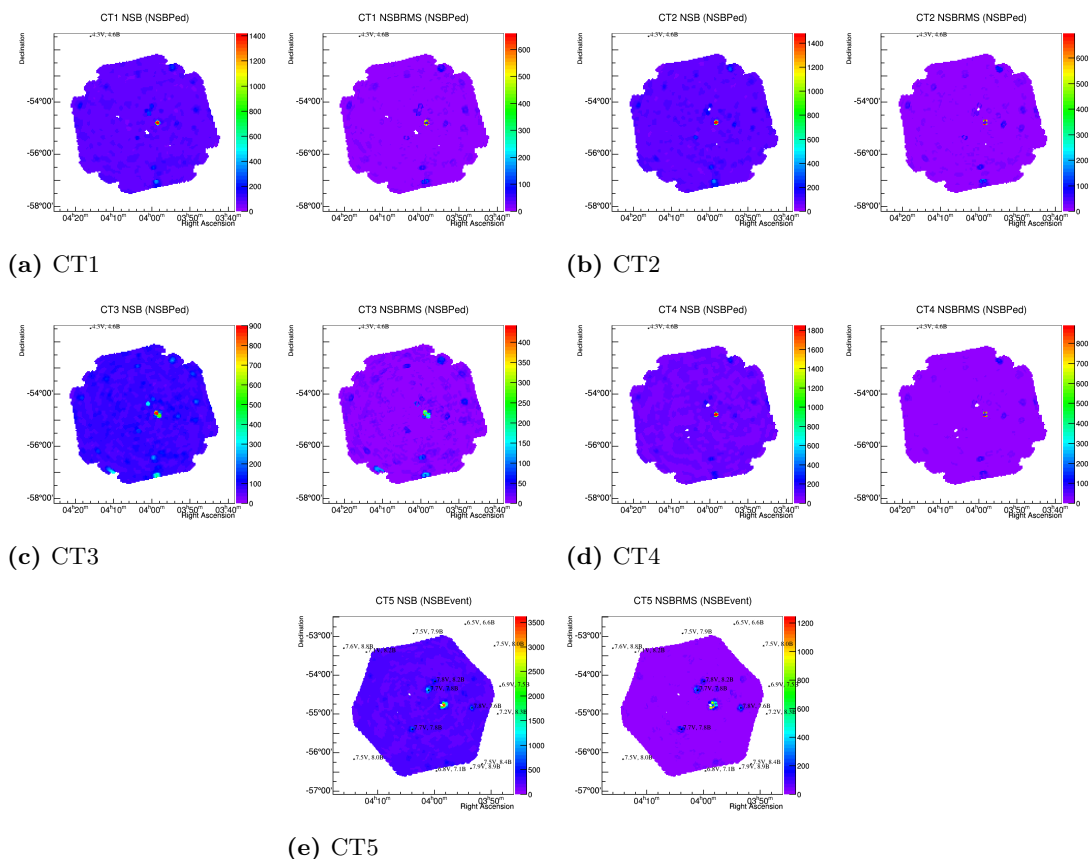


Figure 5.4: NSB maps for CT1-CT5. For each telescope, the left panel shows the NSB map and the right panel shows the corresponding NSB RMS map (as indicated in the panel titles). These correspond to run number 161097 and were taken from the HESS web summary for identifying the source location on the camera FOV, as described in section 5.2.3. Here, the target source (MGAB-V207) is the unmarked bright patch in the middle.

described in detail in section 5.2.3.

5.2.3 Cluster selection strategy

For each run, we first estimate an approximate source location in the camera field of view using the pointing metadata described in subsection 5.2.1 and cross check this location against reference NSB maps from the internal web summary (cf. figure 5.4). Within this constrained region, we define a seed pixel as the pixel with the maximum run wise NSB rate value. We then inspect the time evolution of the camera frames constructed using the raw NSB data, across all timestamps within the run to ensure that the candidate signal is consistent with transient source related variability rather than a static illumination feature². Based on this inspection, we define a connected pixel

²We also checked other catalogues like Vizieir (Tycho-2 catalogue) etc., for any such static illumination features nearby or background stars in the close neighbourhood of the MGAB-V207 region on the camera field of view (mainly sources brighter than $\text{mag}(V) \sim 12$) but found none.

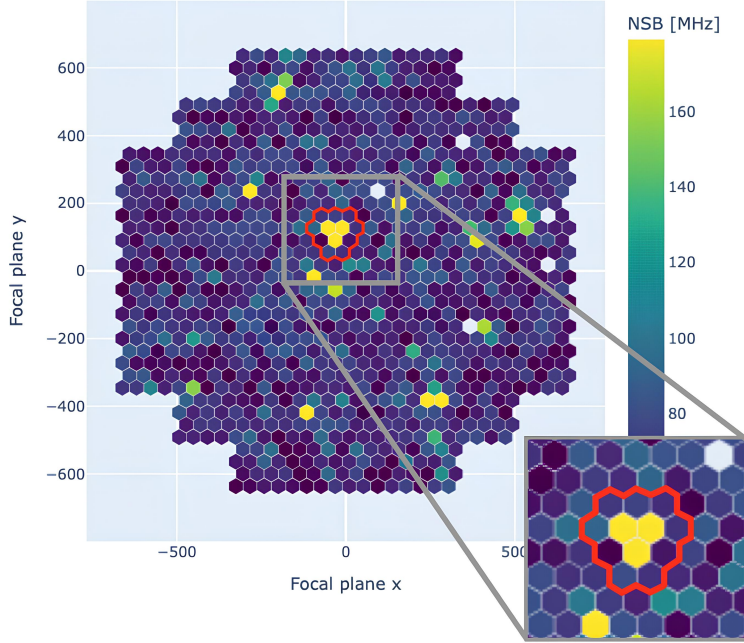


Figure 5.5: One frame of the CT1 FOV for run 161096 showing the cluster of pixels representing the source (MGAB-V207). The zoomed in frame shows clearly the extra 1-pixel layer convention described in section 5.2.3. Hence, all the pixels within the red boundary are considered in the cluster.

cluster that encloses the apparent source footprint on the camera plane and augment it by a one pixel wide boundary layer (a one ring wider in the hexagonal pixel grid, see figure 5.5) to capture possible leakage due to pointing jitter and to enable robustness checks in the subsequent light curve construction.

Nova-track identification (CT5). For CT5, the cluster is obtained with an automated procedure adapted from the satellite-trail detection pipeline developed by Thomas (Lang, 2025) (Specifically the `track_sorter` function). The track sorter can return multiple candidate tracks per run, including satellite-like transients and localised illumination features. To identify the nova candidate, we apply a simple ranking based on (i) how long it lasts through the duration of the run, (ii) location consistency, and (iii) brightness (NSB rate). First, we require the track duration $\Delta t_{\text{tr}} = t_{\text{last}} - t_{\text{first}}$ to be consistent with a source that remains in the camera for the full run, i.e. Δt_{tr} must exceed a fixed fraction of the run duration ($> 95\%$). Among the surviving tracks, we select the one whose pixel cluster centroid lies closest to the expected source region on the camera plane derived from the nominal wobble configuration. If multiple tracks satisfy these criteria, we choose the track with the largest mean NSB level within its pixel set (computed from the track brightness samples).

CT1-CT4: manual cluster definition with run wise animation: For CT1-CT4, the automated track procedure did not yield stable clusters, and the cluster definition

Table 5.2: Parameters used for CT5 track construction and filtering in the adapted `track_sorter` (For more details on the cuts or parameters refer to the functions in `functions_import_essential.py` from Lang (2025)). Tracks are built by associating spatially adjacent pixel hits within a fixed temporal window and are retained only if they satisfy the filtering cuts listed below.

Parameter / cut	Value / criterion
Time association window	$\Delta t_{\text{assoc}} = 60 \text{ s}$
Neighbour requirement	at least one neighbour overlap with recent track hits
Neighbour lookback	last $N_{\text{lb}} = 20$ pixel entries in the track
Spatial extent	$N_{\text{pix}} \leq 10$ (unique pixels per track)
Temporal sampling	$N_t > 1$ and $\langle \Delta t \rangle \leq 5 \text{ s}$
Apparent velocity	$v \leq 0.002 \text{ deg s}^{-1}$

was therefore performed manually following the same seed pixel and boundary layer rule discussed earlier. For each run, we use the metadata based approximate source location as discussed earlier (cf. table 5.1) and the internal web summary NSB maps (cf. figure 5.4) as reference, identify the brightest pixel in the constrained region, and generate a frame-by-frame animation of the camera FOV using the raw NSB data over the full run using `Plotly` in Python, as we talked about in section 5.2.2. The source region is then defined as the surrounding pixel group exhibiting similar variability consistent with the nova location; the final cluster is obtained by adding a one pixel wide boundary layer around this region³ (cf. figure 5.5). This approach prioritises robustness against under selection at the possibility of increased background contribution, which is taken into account for corrections in later sections of this chapter.

Cluster data storage for further processing: For reproducibility, the selected pixel lists are stored in run wise JSON files and used as fixed inputs for all subsequent steps (pixel level corrections, run aggregation, and light curve binning). Each JSON entry records, at minimum, the run identifier, telescope identifier, zenith angle (necessary for performing AOD corrections; this aspect will be discussed later in detail) and the final cluster pixel IDs for each time frame.

5.2.4 Illustration of selected pixels

Figure 5.6 illustrates the pixel-cluster definition introduced in subsection 5.2.3 using a representative CT5 run. Both panels show the same run and the same selected pixel cluster, displayed at two different timestamps. The only visible change is that one pixel within the cluster randomly disappears when its NSB rate rises to very large values (typically $\gtrsim 2500 \text{ MHz}$ in this dataset). This on/off behaviour is not a property of the cluster selection procedure, rather, it is an instrumental effect that can imprint artificial step-like features in the derived light curve if the affected pixel contributes to the signal in one time interval but not in the next.

³Initially, the light curves were obtained without taking into account the extra single layer boundary condition, which led to some inconsistencies in the overall expected structure of the curve due to possible light leakage into neighbouring pixels, hence the approach was discontinued.



(a) “Pix-OFF”: the affected pixel is absent at this timestamp. (b) “Pix-ON”: the same pixel is present at this timestamp.

Figure 5.6: Example of random CT5 pixel saturation within a fixed source pixel cluster. Both panels show the same run (run ID 161096) and the same selected cluster at two different UTC timestamps (indicated above each frame). The zoomed areas highlight the source region: one pixel within the cluster is present in one frame but absent in the other when its NSB rate becomes very large (typically $\gtrsim 2500$ MHz in this dataset). Such abrupt pixel on-offs can introduce artificial jumps in the lightcurve.

The impact of such occasional pixel drop outs is shown in figure 5.7 for the same CT5 run (161096) used in the “Pix-ON”/“Pix-OFF” example (figure 5.6). We plot the mean NSB rate of the selected track cluster versus run time together with the number of pixels contributing to the cluster at each timestamp. Whenever the number of active pixels changes abruptly, the cluster-mean light curve exhibits a corresponding step-like variation. This behaviour is consistent with a saturation-related on/off response of one (or sometimes more) pixels within the cluster and demonstrates that, for small CT5 clusters, the inferred light curve can be dominated by instrumental effects rather than source variability. For this reason, we drop the CT5 approach and move forward with CT1-CT4, since such pixel saturation effects are typically not observed there (at least in the runs analysed for the purpose of this work), and the light curves in the following sections of the chapter are therefore derived from CT1-CT4.

5.3 Step-by-step construction of the NSB derived light curve

To demonstrate how the NSB-derived optical light curve is constructed in practice, we illustrate the full correction chain using a single representative CT1 run (run 161096)




Figure 5.7: CT5 track-based light curve for MGAB-V207 in run 161096. Top: mean NSB rate of the selected track cluster as a function of run time. Bottom: number of pixels contributing to the cluster at each timestamp (as returned by the track-based selection). Discrete changes in the number of active pixels coincide with step-like changes in the cluster-mean NSB level, indicating that occasional pixel drop-outs within the small CT5 cluster can imprint artificial discontinuities in the NSB-derived optical light curve.

as a worked example⁴. We start from the raw cluster light curve extracted from the calibrated NSB time series and then apply, step by step, the instrumental and atmospheric corrections and the run-wise cleaning procedures introduced in this chapter. Comparing the intermediate results for the same run makes the impact of each processing stage directly apparent and motivates the final light curves for MGAB-V207. The corresponding light curves for all analysed runs and telescopes are provided in the appendix (chapter B) to document the full dataset without overcrowding the main text.

5.3.1 Representative run and raw cluster light curve

Following the CT1-CT4 cluster identification procedure described in section 5.2.3, we store the run wise selection in a compact JSON configuration file. The file is organized as a dictionary keyed by run ID, where each entry contains (i) the list of camera pixel identifiers that define the adopted cluster and (ii) the corresponding observing zenith angle:

$$\text{config}[\text{run_id}] = \{ \text{"pixels"} : [p_1, p_2, \dots], \text{"zenith"} : z \}. \quad (5.1)$$

The "pixels" field therefore encodes the spatial extraction region used to build the

⁴The run from CT1 is used for illustration; the same procedure is applied identically to all runs from telescopes, and the full set of light curves is provided in the appendix (cf. chapter B).

NSB derived light curve for a given run, while "zenith" records the zenith angle in degrees. The zenith angle is used specifically in the AOD correction method discussed in section 5.3.3, where the optical depth is associated to an attenuation factor along the slant path through the atmosphere.

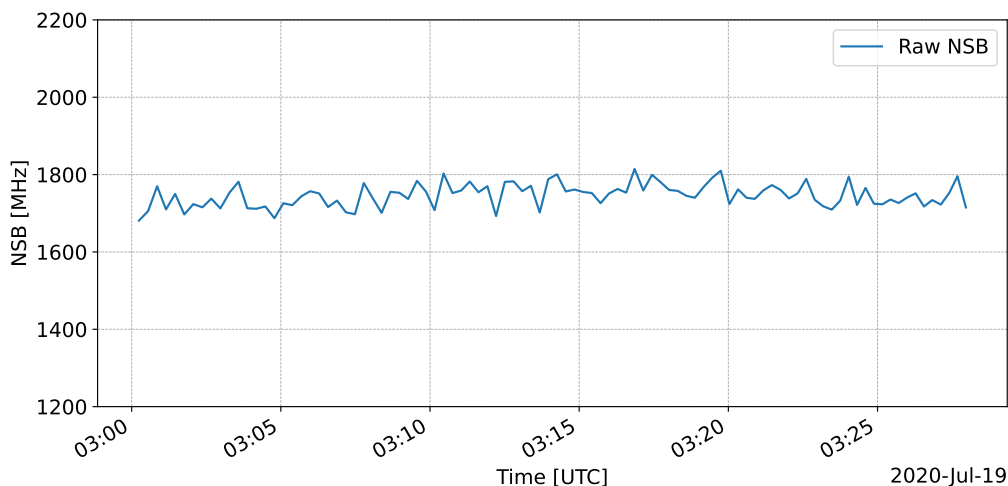


Figure 5.8: Raw NSB-derived light curve for CT1 run 161096 (MGAB-V207 field) prior to any instrumental or atmospheric corrections. The curve shows the summed NSB rate over all pixels in the selected cluster (cf. section 5.3.1). This time series serves as the baseline input for the correction chain presented in section 5.3.

Figure 5.8 shows the raw cluster time series extracted for run 161096, which we use as the input to the following correction steps discussed below.

5.3.2 Flat-field correction

Pixel-to-pixel variations in the camera response (e.g., due to differences in light collection efficiency and photodetector gain) imprint a fixed pattern structure on the measured NSB rates. To place all pixels on a consistent relative scale before forming the cluster sum, we apply a run wise flat-field correction based on the standard H.E.S.S. flat-field calibration model (Aharonian et al., 2004).

For each run and telescope, we query a flat-field coefficient table using the run start time as reference. Concretely, the run time stamp is extracted from the raw per pixel NSB file as the earliest recorded time in the run. We then obtain the corresponding flat-field coefficients via `get_FlatField(date=run_dt, telId=tel_id)` (cf. section A.3), which returns a coefficient per camera pixel⁵. An example output looks like this:

Pixel	Coefficient_hg	Coefficient_lg	charge_mean_hg	charge_mean_lg
0	0.977048	0.977048	129.147	-1.0
1	0.991012	0.991012	131.184	-1.0

⁵Flat-field coefficients are queried from the H.E.S.S. calibration database using a `MySQLdb` connection (credentials in `~/dbtoolsrc`); for each run date and telescope we select the latest valid set in `Calib3_FlatField_Set` and retrieve per-pixel coefficients from `Calib3_FlatField`.

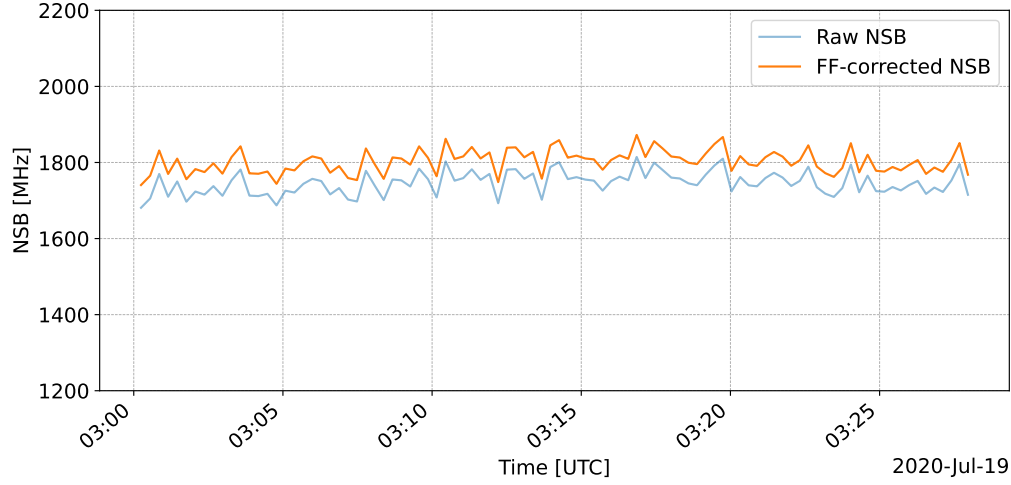


Figure 5.9: Raw and flat-field corrected NSB derived light curves for CT1 run 161096. In both cases the light curve is defined as the sum of NSB rates over the selected pixel cluster; the flat-field correction applies the pixel wise factors f_i from equation (5.2) before summation (cf. equation (5.4)).

5	2	0.891101	0.891101	117.937	-1.0
6	3	0.912502	0.912502	120.954	-1.0
7	4	0.845686	0.845686	111.651	-1.0

In this work, we use the high-gain coefficients `Coefficient_hg`. To make the correction dimensionless and preserve the overall normalization of the summed light curve, the coefficients are normalized by their median value:

$$f_i = \frac{C_i^{\text{hg}}}{\text{median}(C^{\text{hg}})}, \quad (5.2)$$

where C_i^{hg} denotes the tabulated high gain coefficient for pixel i . Let $R_i(t)$ denote the raw NSB rate for pixel i at time t . The flat-field corrected pixel series is computed by a multiplicative scaling,

$$R_i^{\text{ff}}(t) = f_i R_i(t). \quad (5.3)$$

After this correction, the NSB derived light curve for a given run is constructed as the sum over the selected cluster pixels \mathcal{C} (cf. section 5.3.1),

$$L_{\text{ff}}(t) = \sum_{i \in \mathcal{C}} R_i^{\text{ff}}(t) = \sum_{i \in \mathcal{C}} f_i R_i(t). \quad (5.4)$$

Figure 5.9 illustrates the impact of the flat-field correction on the representative run 161096, comparing the raw light curve to the corrected series $L_{\text{ff}}(t)$.

For run 161096, the flat-field correction corresponds to a median rescaling of $L_{\text{ff}}/L_{\text{raw}} = 1.031985$ (i.e. +3.198%), equivalent to a median absolute shift of 56.265 MHz

for this run. This indicates that the correction primarily adjusts the relative pixel contributions in the cluster sum. A similar trend is observed for all runs across all telescopes.

This establishes a consistent relative pixel response within the cluster prior to the AOD correction, which depends on the observing geometry (in particular the zenith angle stored in the run wise configuration; cf. section 5.3.3).

5.3.3 Atmospheric transparency correction

The measured NSB rates are affected by time dependent atmospheric extinction, primarily due to aerosols. To correct the NSB derived light curves for changes in atmospheric transparency, we use the AOD provided by the AERONET sun photometer network (Holben et al., 1998). The AOD is a dimensionless measure of the vertical aerosol optical depth and quantifies the attenuation of light by aerosols along the line of sight.

We extract the AOD from the downloaded AERONET data file, selecting the column at the target wavelength (here 380 nm). Since AERONET measurements are taken during daytime, whereas the H.E.S.S. observations relevant for this work were performed at night, the AOD value used for a given observing night is obtained by bracketing the night with the last AOD measurement of the preceding day and the first measurement of the following day, and taking their mean. This provides a stable proxy for the aerosol load during the corresponding night, while avoiding extrapolation over multiple days.

For an observation at zenith angle θ_z , the attenuation along the slant path is larger than the vertical extinction. We approximate this with an airmass factor $m(\theta_z)$ using the Kasten-Young parameterization (Kasten and Young, 1989),

$$m(\theta_z) = \frac{1}{\cos(\theta_z) + 0.50572 (96.07995^\circ - \theta_z)^{-1.6364}}, \quad (5.5)$$

with θ_z in degrees. For each run, the zenith angle stored in the run wise configuration (cf. section 5.3.1) is converted to $m(\theta_z)$, and the AOD is converted to a multiplicative attenuation factor via the Beer-Lambert law,

$$T(\theta_z, t) = \exp[-\tau_a(t) m(\theta_z)], \quad (5.6)$$

where $\tau_a(t)$ denotes the vertical AOD at the chosen wavelength. Since the measured signal is attenuated by T , the corrected NSB rate is obtained by multiplication with T^{-1} :

$$R_i^{\text{corr}}(t) = R_i^{\text{ff}}(t) \exp[\tau_a(t) m(\theta_z)], \quad (5.7)$$

applied identically to all pixels i in the camera at a given time stamp. The NSB derived light curve used in the following is then constructed by summing $R_i^{\text{corr}}(t)$ over the cluster pixels. This procedure assumes that the aerosol extinction is approximately stable over the observing night and that a single bracketing AOD value is adequate.

Figure 5.10 compares the raw, flat-field corrected, and flat-field plus AOD corrected NSB-derived light curves for the representative run 161096.

Applying the AOD based correction shifts the flat-field corrected light curve for run 161096 upward by a median factor of $L_{\text{ff+aod}}/L_{\text{ff}} = 1.062597$ (i.e. +6.260%), because




Figure 5.10: Stepwise impact of the flat-field and atmospheric transparency corrections on the NSB derived light curve for CT1 run 161096. Shown are the raw light curve, the flat-field corrected series (cf. section 5.3.2), and the series after additionally applying the AOD based transparency correction from section 5.3.3.

it accounts for atmospheric extinction by aerosols along the line of sight. Together with the flat-field correction ($L_{\text{ff}}/L_{\text{raw}} = 1.031985$), this results in a total median upward shift of $L_{\text{ff+aod}}/L_{\text{raw}} = 1.096585$ (i.e. +9.658%) relative to the raw light curve for this run. A qualitatively similar behaviour is observed for all analysed runs and telescopes as in the previous subsection.

5.3.4 Background subtraction

After applying the flat-field and AOD corrections, the NSB derived light curve still contains a large contribution from the diffuse night-sky background and instrumental offsets that are common across the camera. To isolate relative variations associated with the selected cluster, we construct a run wise background estimate from pixels outside the signal cluster and subtract an empirical baseline.

For each run, the signal pixel set \mathcal{C} is taken from the run wise JSON configuration (cf. section 5.3.1). Using the corrected per pixel time series $R_i^{\text{corr}}(t)$ (cf. section 5.3.3), we form the cluster-summed signal series

$$S(t) = \sum_{i \in \mathcal{C}} R_i^{\text{corr}}(t). \quad (5.8)$$

The background pixel set \mathcal{B} is defined as all remaining camera pixels present in the corrected file, $\mathcal{B} = \{\text{all pixels}\} - \mathcal{C}$, and the corresponding summed series is

$$B_{\text{cam}}(t) = \sum_{j \in \mathcal{B}} R_j^{\text{corr}}(t). \quad (5.9)$$

Since both $S(t)$ and $B_{\text{cam}}(t)$ are sums over different numbers of pixels, we scale the background to the size of the signal cluster by using the camera average per-pixel

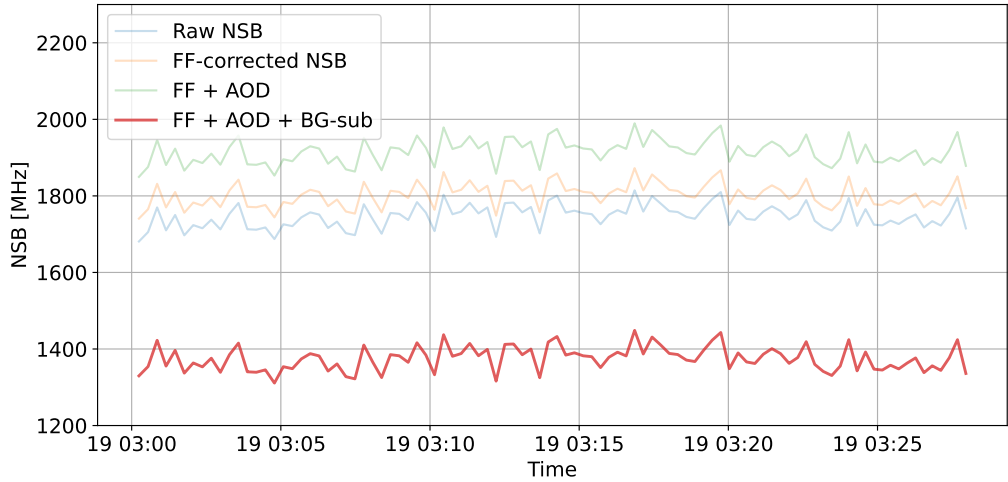


Figure 5.11: Step by step evolution of the NSB derived light curve for CT1 run 161096. Shown are the raw light curve, the flat-field corrected series, the additional AOD based transparency correction, and the baseline subtracted curve obtained by subtracting a rolling low percentile baseline estimated from pixels outside the signal cluster (cf. section 5.3.4).

background level,

$$B(t) = \frac{B_{\text{cam}}(t)}{|\mathcal{B}|} |\mathcal{C}|. \quad (5.10)$$

In this work we subtract a slowly varying baseline derived from the scaled background using a rolling low percentile estimator. More specifically, we compute a rolling q -quantile $b(t)$ of $B(t)$ within a time window of width Δt and define the background subtracted light curve as

$$L_{\text{bgsub}}(t) = S(t) - b(t), \quad (5.11)$$

with $q = 0.05$ and $\Delta t = 5$ min for the results shown in this work, unless stated otherwise. This procedure subtracts a slowly varying baseline that captures camera wide background levels and their gradual evolution, while the use of a rolling low percentile estimator reduces sensitivity to short lived upward fluctuations in the background estimate⁶.

Figure 5.11 summarizes the cumulative effect of the successive corrections for the representative run 161096.

The background subtraction step shifts the FF+AOD light curve downward by removing a median baseline of 541.543 MHz, which corresponds to a median fractional removal of 28.160%. This is expected because the subtraction accounts for the camera wide NSB level (and its slow evolution) estimated from pixels outside the signal cluster. The Median Absolute Deviation (MAD) derived scatter decreases from 36.208 MHz to 33.882 MHz (ratio 0.9358), which shows that large camera wide variations are reduced while the short timescale structure of the signal is kept.

⁶This choice is motivated by runs taken close to twilight or during changes in lunar illumination, where the NSB level may (or may not) be affected due to evolving sky brightness.

5.3.5 Light curve evolution and run continuity

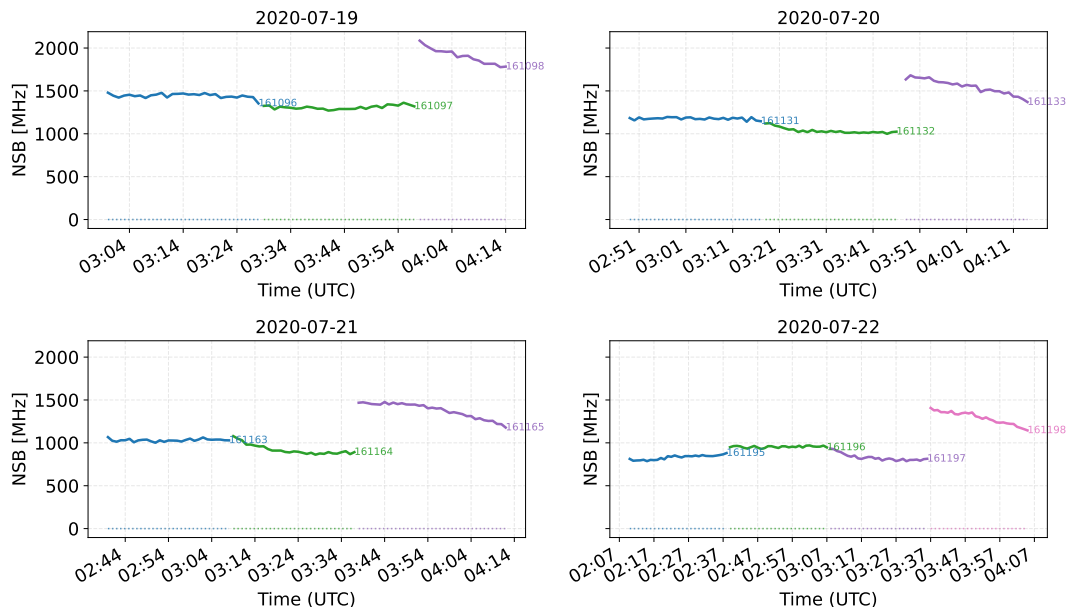


Figure 5.12: Background subtracted NSB derived light curves for the first four nights of MGAB-V207 observation as recorded by CT2. Each panel shows the light curves for the individual runs within the night; the time series are binned to 1 min. While the corrected curves are expected to connect smoothly between consecutive runs within a night, run-to-run discontinuities of up to $\gtrsim 500$ MHz are observed in some cases. Corresponding light curves for CT1-CT4 over the full observation period for all telescopes are provided in chapter B.

After applying the full correction chain described above (FF, AOD corrections, and baseline background subtraction), we next inspect the run-to-run behaviour of the NSB derived light curves on nightly time scales. Figure 5.12 shows the light curves for the first four observing nights of observation as recorded by CT2. CT2 is used here as a representative example for visualization; the corresponding light curves for all telescopes (CT1-CT4) and all analysed nights are provided in chapter B.

5.3.6 Nightly mean NSB proxy and cross-telescope consistency

To provide a compact overview of the night-to-night evolution, we compute for each observing night the mean NSB rate of the selected source cluster after applying the full correction chain and the baseline subtraction described above. We then average this nightly mean across CT1-CT4 to obtain a single proxy for the relative optical brightness evolution over the analysed time window. Figure 5.13 shows that this proxy declines from 2020-07-19 to 2020-07-29, consistent with a fading source over the observation period. An exception is a clear upward deviation on 2020-07-23, which motivates a check of telescope-to-telescope consistency.

Figure 5.13 shows the corresponding nightly mean time series separately for CT1-CT4.

The overall decline, is expected of the nova and the upward deviation on 2020-07-23 are present in all four telescopes, the exact reason for which remains unknown at the moment⁷. In the following sections, we move from nightly averages to run-resolved time series in order to assess intra-night continuity and to identify systematic effects that can bias the reconstructed evolution on shorter time scales.




Figure 5.13: Nightly mean NSB rate of the selected source cluster, averaged across CT1-CT4, used as a proxy for the relative optical brightness evolution during the MGAB-V207 observation window (for the data, see table D.5). Points show the nightly mean after applying FF and AOD corrections and the baseline subtraction from section 5.3.4; error bars denote the standard error of the mean across runs contributing to each night. The proxy declines over the analysed period, with an upward deviation on 2020-07-23 that is examined for consistency across individual telescopes in figure 5.14.

5.4 Run-to-run discontinuities and probable causes

5.4.1 Phenomenology of run-to-run discontinuities

As shown in figure 5.12, the background subtracted NSB derived light curves do not always join smoothly between consecutive runs within the same night. In several nights, the final run exhibits a clear offset relative to the preceding run, with discontinuities that reach $\gtrsim 500$ MHz in the CT2 example. Such jumps exceed the level expected from run-to-run variations in observing conditions alone. For instance, changes in effective light collection can occur when the illumination pattern shifts across the Winston cone gaps, which reduces the effective aperture; however, this geometric effect is expected to be at the $\mathcal{O}(10\%)$ level and does not readily explain the larger offsets observed here. Similar behaviour is visible across multiple nights and for almost all telescopes, while the full set of nightly light curves for CT1-CT4 is provided in chapter B.

⁷Though, it is consistent within the error ranges in the AAVSO light curve.

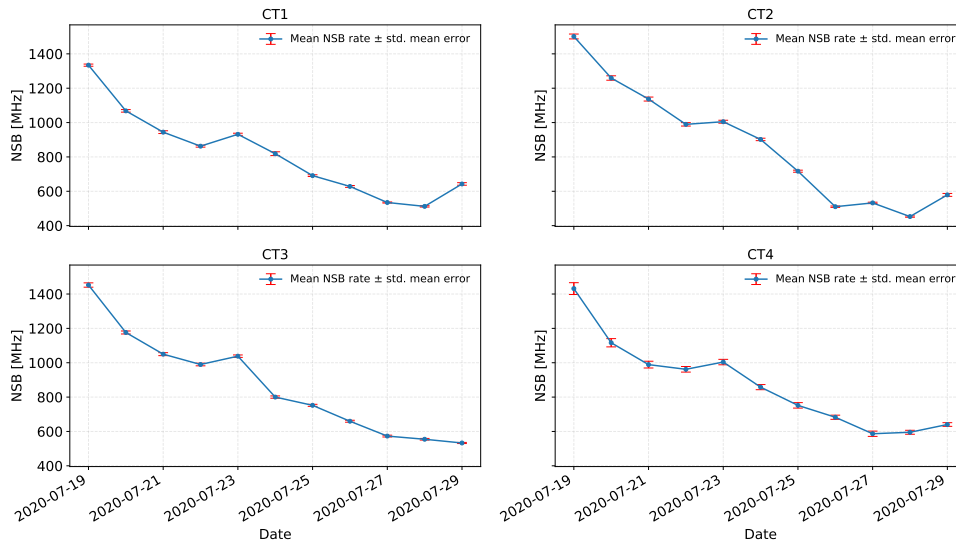


Figure 5.14: Nightly mean NSB proxy light curves shown separately for CT1-CT4, computed as in figure 5.13. The shared downward trend and the upward deviation on 2020-07-23 appear consistently across all telescopes. (For the data, see table D.5)

5.4.2 Probable association with pointing offset transitions

Inspecting figure 5.12 we clearly see that the most noticeable discontinuities coincide with specific changes in the pointing offset configuration between runs, when cross-checked with web summary metadata. For the first night in CT2, the anomalous run occurs when the pointing offset transitions from a $+0.5$ deg declination offset ($+0.5y$) in the preceding run to a -0.5 deg right-ascension offset ($-0.5x$) in the subsequent run⁸. This repeated occurrence suggests that the discontinuities may be linked to run dependent conditions that change together with the offset position though, at the same time, we find no concrete evidence that the effect is directly caused by the pointing offset itself. We therefore treat the offset transition as an empirical indicator for when the discontinuity occurs, rather than as a demonstrated physical cause.

5.4.3 Pixel-level origin of the anomalous runs

To localize the origin of the run-to-run offsets, we inspected the corrected time series at the level of individual pixels within the signal cluster. Figure 5.15 shows the baseline subtracted per-pixel light curves for CT2 for the first four observing nights, with the runs of each night overlaid. In each night, one run exhibits a pronounced offset in the light curve (cf. figure 5.12); at the per-pixel level, this behaviour is consistently driven by a single pixel whose baseline-subtracted NSB level is substantially higher than the remaining cluster pixels. For CT2, the same pixel (pix_id 477) is repeatedly associated with the anomalous runs across all four nights, indicating that the discontinuities may be caused by a localized, pixel-specific contribution rather than a camera wide effect.

⁸Here we use the internal web summary notation for the offset pattern used in the MGAB-V207 observations.

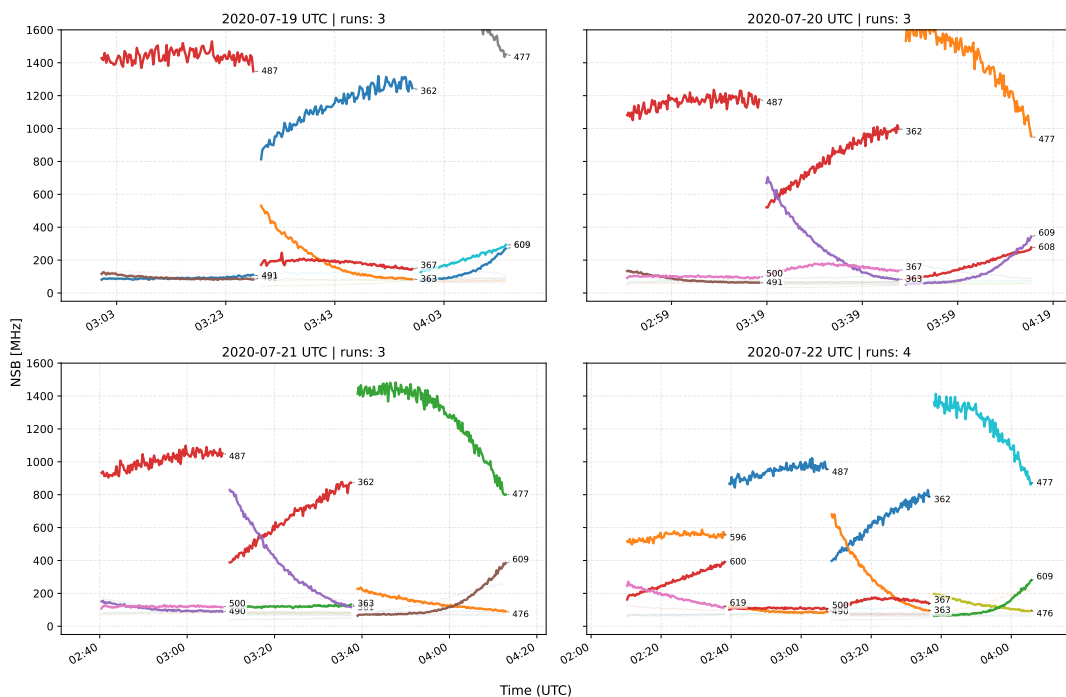


Figure 5.15: Baseline subtracted per-pixel NSB time series for CT2 for the first four observing nights of the MGAB-V207 observation. Each panel corresponds to one night; curves show the individual cluster pixels for all runs of that night overlaid (after applying FF and AOD corrections and the baseline subtraction from section 5.3.4). The three most prominent pixels in each run is highlighted with a darker colour. The anomalous run-to-run offsets seen in the light curves (cf. figure 5.12) are consistently dominated by a single pixel (here, pixel 477) with an abnormally higher baseline subtracted NSB rate.

A similar single-pixel dominance is observed in other telescopes, with the anomalous runs repeatedly involving CT1 pixel 609, CT4 pixel 487, and CT2 pixel 477; for CT3 no equally stable single pixel association could be established at the time of writing the thesis.

A potential concern is that the manual cluster selection procedure could bias the cluster composition and thereby induce discontinuities. We tested the selection criteria on multiple runs and telescopes and verified that the anomaly persists under reasonable variations of the selected cluster, indicating that the discontinuity is not an artefact of the manual selection step. The repeated association with a specific pixel instead points to a run-dependent effect that alters the relative scaling of that pixel in the corrected series, motivating a closer inspection of the calibration factors applied during these runs.

5.4.4 Check of flat-field calibration stability for the dominant pixels

The repeated dominance of a single pixel in the anomalous runs motivates a targeted check of the calibration factors applied to that pixel. In our correction chain, pixel wise




Figure 5.16: Temporal evolution of the H.E.S.S. FF coefficients for the pixels most strongly associated with the run-to-run discontinuities (CT1 pixel 609, CT2 pixel 477, CT4 pixel 487) over the MGAB-V207 observation window. The vertical dashed lines mark the nights in which the observations took place (2020-07-19 to 2020-07-29). Over this period, the coefficients remain approximately constant. A longer-baseline version of this trend plot is shown in figure B.4.

gain variations are accounted for by the FF coefficient; a run-dependent miscalibration or a sudden change in the applied coefficient could therefore introduce an artificial offset in the corrected NSB time series. We therefore inspected the temporal evolution of the H.E.S.S. FF coefficients for the pixels that repeatedly dominate the discontinuities (CT1 pixel 609, CT2 pixel 477, and CT4 pixel 487) over a time window covering the MGAB-V207 observations.

Figure 5.16 shows that the FF coefficients for these pixels remain approximately constant over the relevant period, with no step-like change coincident with the observing nights highlighted in section 5.4.2. Within this window, the coefficient variations do not exhibit a discrete transition that could naturally account for the observed run-to-run offsets in the baseline subtracted light curves. A longer trend study over the full available time range is provided in figure B.4, which does show significant changes in the FF coefficients at times, but none of those lie within our analysis time frame, hence possibly couldn't contribute to the noticed anomaly.

5.5 Summary and outlook

In this chapter we established a reproducible analysis pipeline to extract NSB derived optical proxy light curves from H.E.S.S. CT1–CT4 data for the MGAB-V207 observations. The procedure combines pixel selection, FF and AOD corrections, and baseline background subtraction (section 5.3.4) to construct run-resolved and nightly time series (section 5.3.5). Applying this chain demonstrates that light curve extraction from NSB data is feasible in practice and yields stable large scale trends across multiple nights and telescopes.

At the same time, we identify run-to-run discontinuities that prevent the corrected light curves from joining smoothly between consecutive runs within a night (section 5.4). These offsets can reach values that are difficult to explain by expected run-to-run changes in observing conditions alone. A pixel-level breakdown shows that the most prominent discontinuities are repeatedly driven by a single pixel within the selected cluster (section 5.4.3). While several of the strongest discontinuities coincide with transitions in the pointing offset configuration (section 5.4.2), this coincidence is treated as an indirect indicator rather than a demonstrated physical cause, as of now. Inspecting the temporal behaviour of the applied FF calibration factors for the dominant pixels did not reveal a discrete change over the observation window that would trivially explain the observed jumps (section 5.4.4). The origin of the discontinuities therefore remains unresolved at the time of writing, and we treat them as a first-order limitation of the present NSB-based light curve results.

The dominant limitation of the present pipeline is therefore most likely systematic in nature: run-to-run discontinuities persist despite the applied correction chain and can be driven by localized pixel level behaviour. A consolidated discussion of plausible explanations and proposed follow-up tests, together with the implications for the broader thesis goals, is deferred to chapter 6.

Chapter 6

Discussion and conclusion

This thesis was motivated by a practical question: to what extent can ground based γ ray instruments contribute to nova time-domain science beyond standard event level analyses? We address this in two ways: by directly probing VHE emission with SWGO, and by trying to extract secondary/proxy optical information from NSB data using H.E.S.S.. The results are therefore best interpreted as instrument-focused feasibility studies. They create analysis workflows that can be reliably reproduced and clearly specify the assumptions and constraints that limit how the results can be interpreted.

In the SWGO study (chapter 4), we combined a parametric nova shock model with IRF-based sensitivity products to derive detectability diagnostics in the 0.5 – 100 TeV band for a 30 d exposure. For the baseline SWGO D8 configuration, the resulting integral sensitivity is of order $7 \times 10^{-14} \text{ erg cm}^{-2} \text{ s}^{-1}$. The corresponding detectability contours show that D8 is substantially more promising than the pathfinder-like D9 configuration, whose integral sensitivity is nearly an order of magnitude higher (cf. figure 4.8 and the corresponding D9 figure in the appendix: figure B.2). This difference is also visible when placing representative literature novae on the D8 detectability map: for the parameter combinations shown, most candidates do not cross the D8 threshold. However, the most optimistic cases lie close enough that moderate improvements in sensitivity (or, equivalently, modest shifts in model realisations that increase the TeV-band energy flux at Earth) could plausibly bring systems such as T CrB into the discovery space, and potentially even RS Oph under favourable assumptions. In contrast, the D9 threshold is sufficiently higher that comparable scenarios remain clearly under it, making D9 less compelling for nova detectability within the framework explored here.

A further outcome of the SWGO study is that the forward-modelling approach makes the key modelling assumptions explicit. The current nova class intentionally follows simple analytic prescriptions from the literature it is based on. As a consequence, several time dependencies are implemented as unbroken power laws, without additional cutoffs or transition terms. For example, the shock-radius evolution is parameterised by a simple power-law form, whereas more complete descriptions may require additional scales or drop-off parameters to capture changes in the outflow structure, radiative losses, or interactions with a non-uniform environment. Extending the model to include such features, ideally benchmarked against hydrodynamic simulations or well constrained multiwavelength nova datasets, would be a natural next step. This would reduce model-driven ambiguity in the inferred TeV-band flux evolution and strengthen the link between astrophysical scenarios and SWGO detectability diagnostics.

In the H.E.S.S. study (chapter 5), we showed that camera monitoring products can

be repurposed into time-resolved NSB light curves that serve as a relative proxy for optical brightness, and we established a working correction and analysis pipeline to do so. At the same time, the end-to-end results reveal systematic limitations that currently dominate the uncertainty.

A further practical extension of the NSB-derived light curves would be to express them in a differential magnitude scale rather than in raw NSB rate. This would require extracting NSB time series for one or more non variable reference sources in the same field of view and forming a relative measurement with respect to this reference ensemble. In this way, the nova proxy could be expressed as an instrumental differential magnitude, enabling comparisons between different nights and, to a first approximation, under varying observing conditions, while not relying on any absolute photometric calibration.

The most important limitation is run-to-run discontinuities that persist after applying FF and AOD corrections, and which can be driven by localized pixel-level behaviour within the selected cluster (sections 5.3.5, 5.4.3 and 5.5). Two concrete improvement directions follow from the discussions in chapter 5 and from feedback in one of the H.E.S.S. analysis call meetings. First, the FF calibration may not fully capture the effective spectral response relevant for a nova field: the nova spectrum and the wavelength dependent PMT response can introduce uncertainty that is not represented by a single global correction factor. Addressing this would require either a spectral-response-aware correction or a dedicated uncertainty term that propagates calibration uncertainty into the final light curves. Second, targeted validation tests can help determine whether the dominant pixel-level fluctuations reflect intrinsic pixel behaviour or arise specifically in the transient illumination context. For example, analysing the same pixel cluster (including the identified problematic pixel) in observations without a nova contribution provides a control sample for per-pixel NSB behaviour. This can clarify whether the discontinuities are caused by a persistent pixel abnormality (e.g. an incorrect FF coefficient or a gain/voltage calibration issue) or by a transient-dependent effect. Together, these steps directly test the leading hypotheses for the discontinuities and provide a route toward a more robust optical proxy product.

Overall, this thesis established two complementary approaches to nova time-domain science with ground based γ ray observatories. For SWGO, we implemented a lightweight nova-shock model and demonstrated a forward-modelling pipeline that yields quantitative, IRF-based detectability diagnostics, so that sensitivity and detectability limits can be stated within clearly defined assumptions. For H.E.S.S., we developed and validated a working pipeline to extract optical proxy light curves from NSB monitoring data, and we identified the dominant systematic effects that must be controlled before stronger physical inferences are justified. The ability to derive light curves from NSB data is valuable because it can improve temporal coverage for individual events and broaden the scientific use of H.E.S.S. archival data as a complementary, high-cadence optical monitoring resource, as discussed in section 1.1, provided that the pixel-level and calibration driven systematics highlighted in this work are quantitatively mitigated.

Acronyms

AERONET Aerosol Robotic Network.

AOD Atmospheric Optical Depth.

DSA Diffusive Shock Acceleration.

EAS Extensive Air Shower.

Fermi-LAT Fermi Large Area Telescope.

FF Flat-fielding.

GADF Gamma Astro Data Format.

H.E.S.S. High Energy Stereoscopic System.

HAP H.E.S.S. Analysis Package.

IACT Imaging Atmospheric Cherenkov Telescope.

IRF Instrument Response Function.

MAD Median Absolute Deviation.

NSB Night Sky Background.

PMT Photo Multiplier Tube.

PSF Point Spread Function.

SWGO Southern Wide-field Gamma-ray Observatory.

UTC Coordinated Universal Time.

VHE Very High Energy.

WCD Water Cherenkov Detector.

Appendix A

Software and repositories

This appendix lists the software packages and code repositories used to produce the results in this thesis, together with version access dates to support reproducibility.

A.1 GAMERA

GAMERA is an open-source C++/Python framework for spectral modelling of non-thermal astrophysical sources. It provides a modular interface to define time-dependent leptonic and hadronic particle populations and to compute the associated photon emission for a range of source classes. In this thesis, GAMERA is used to generate model spectral energy distributions (SEDs) for the RS Oph worked example and for representative nova-like emission scenarios in section 4.4.

Repository version: The source code is available online:<https://github.com/libgamera/GAMERA>. All results in this thesis were produced with the repository version as on 2025-02-15.

License: GAMERA is distributed under the GNU Lesser General Public License (LGPL), version 2.1 or later.

A.2 NOVA

The module `novae_custom_class_final.py` in the repository defines a single class, `Nova`, that implements a compact, RS Oph motivated, parametric nova-shock model for quick calculations and parameter scans. It is intended for reproducible figure generation and for exploring how shock kinematics and ambient densities affect basic non-thermal proxies such as hadronic interaction timescales and an order of magnitude gamma-ray flux proxy.

The implementation uses Astropy quantities throughout, so units are enforced at runtime.

Repository version: The source code is available online:<https://github.com/c0smic-65/novae-vhe-modeling.git>. All results in this thesis were produced with the repository version as on 2026-01-30.

A.3 nsb-to-nova-lightcurves

The repository `nsb-to-nova-lightcurves` collects the analysis scripts used to convert corrected H.E.S.S. NSB time series into an optical light curve proxy for nova outbursts (see chapter 5).

Repository version: The source code is available online:<https://github.com/c0smic-65/nsb-to-nova-lightcurves>. All results in this thesis were produced with the repository version as on 2026-02-03.

Appendix B

Supplementary plots




Figure B.1: Similar plot to figure 4.7, but with index = -3 as well. Interesting point to note here is that ideally one would expect the curve for index = -3 to be above the one for index = -2, but this is not the case. The reason for that is still under investigation at the time of writing this thesis




Figure B.2: Nova detectability map for the SWGO D9 configuration based on the 30 day IRF derived sensitivity. The colour scale shows the model predicted integrated p - p energy flux at Earth in the 0.5–100 TeV band as a function of distance d and shock velocity v_{sh} . The dotted white curve indicates the D9 integral sensitivity threshold for a 30 day exposure, $F_{\text{sens}}^{(\text{D9})} \simeq 7.56 \times 10^{-13} \text{ erg cm}^{-2} \text{ s}^{-1}$ (integrated over 0.5–100 TeV); As expected, the integral sensitivity value is much worse than D8.

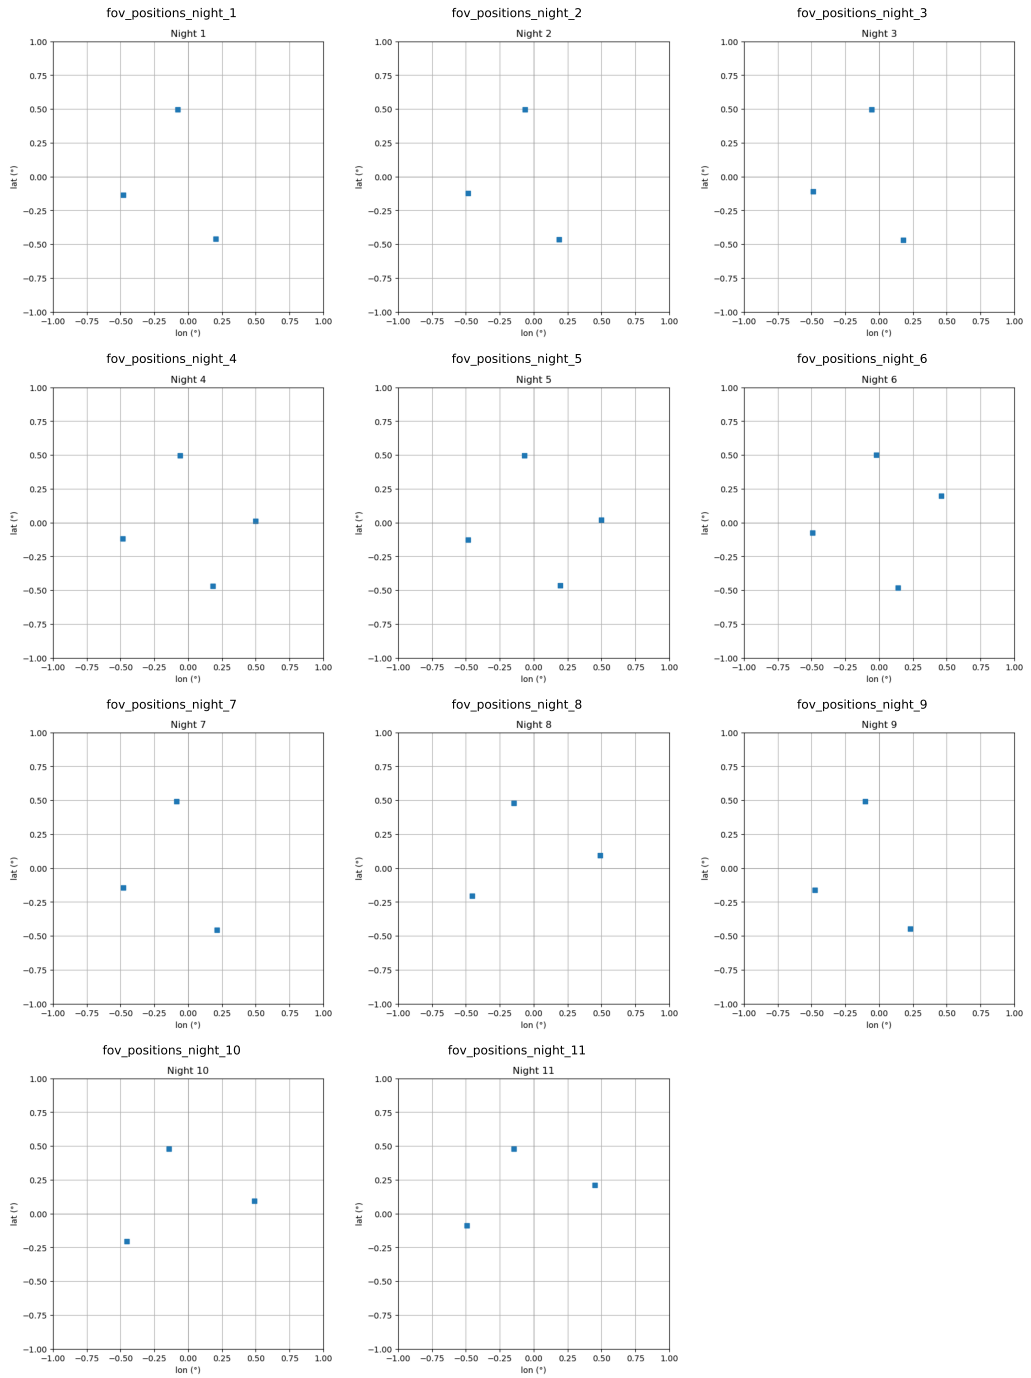


Figure B.3: Sky-coordinate offsets of the reconstructed source position (MGAB-V207) on the camera field of view for each observing night. The panels show the relative longitude and latitude offsets (in degrees); each marker corresponds to the nightly cluster position used for the corresponding pixel-region selection discussed in section 5.2.1. Adapted from the HESS Wiki.




Figure B.4: Flat fielding coefficient trends of possibly problematic pixels, similar to figure 5.16, but for a timescale of about a year. The red lines mark the period of observation of MGAB-V207.

Nightly NSB light curves - all telescopes

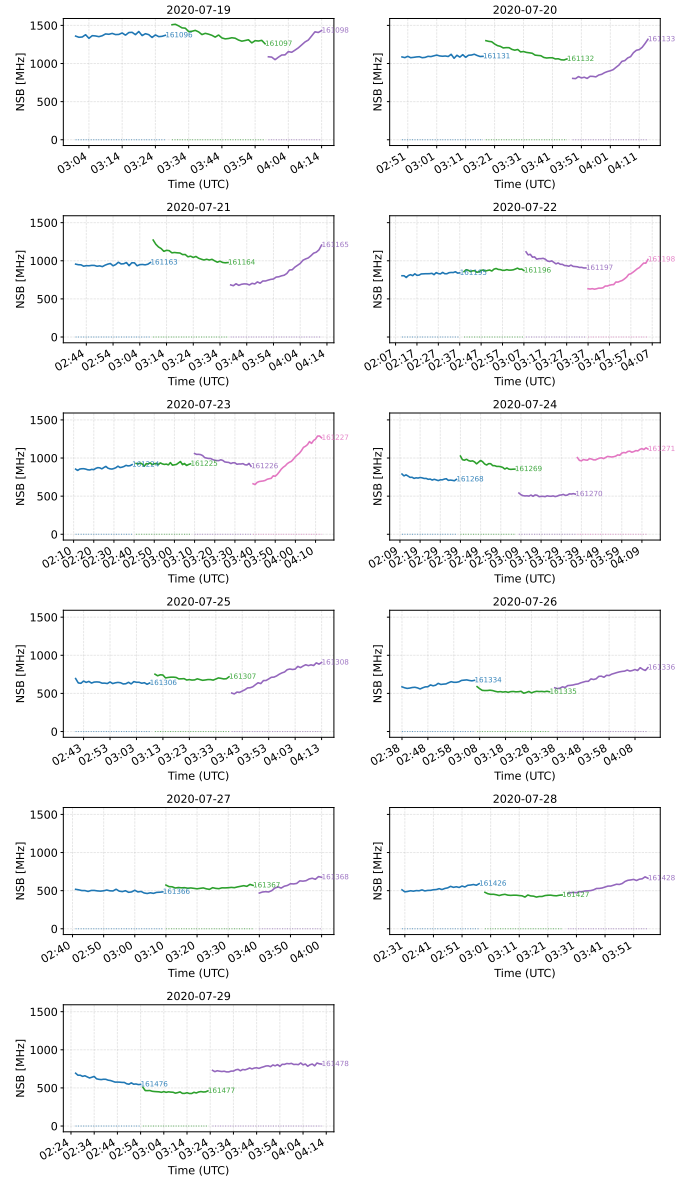


Figure B.5: Nightly NSB light curves for CT1. Each coloured segment corresponds to an individual run.

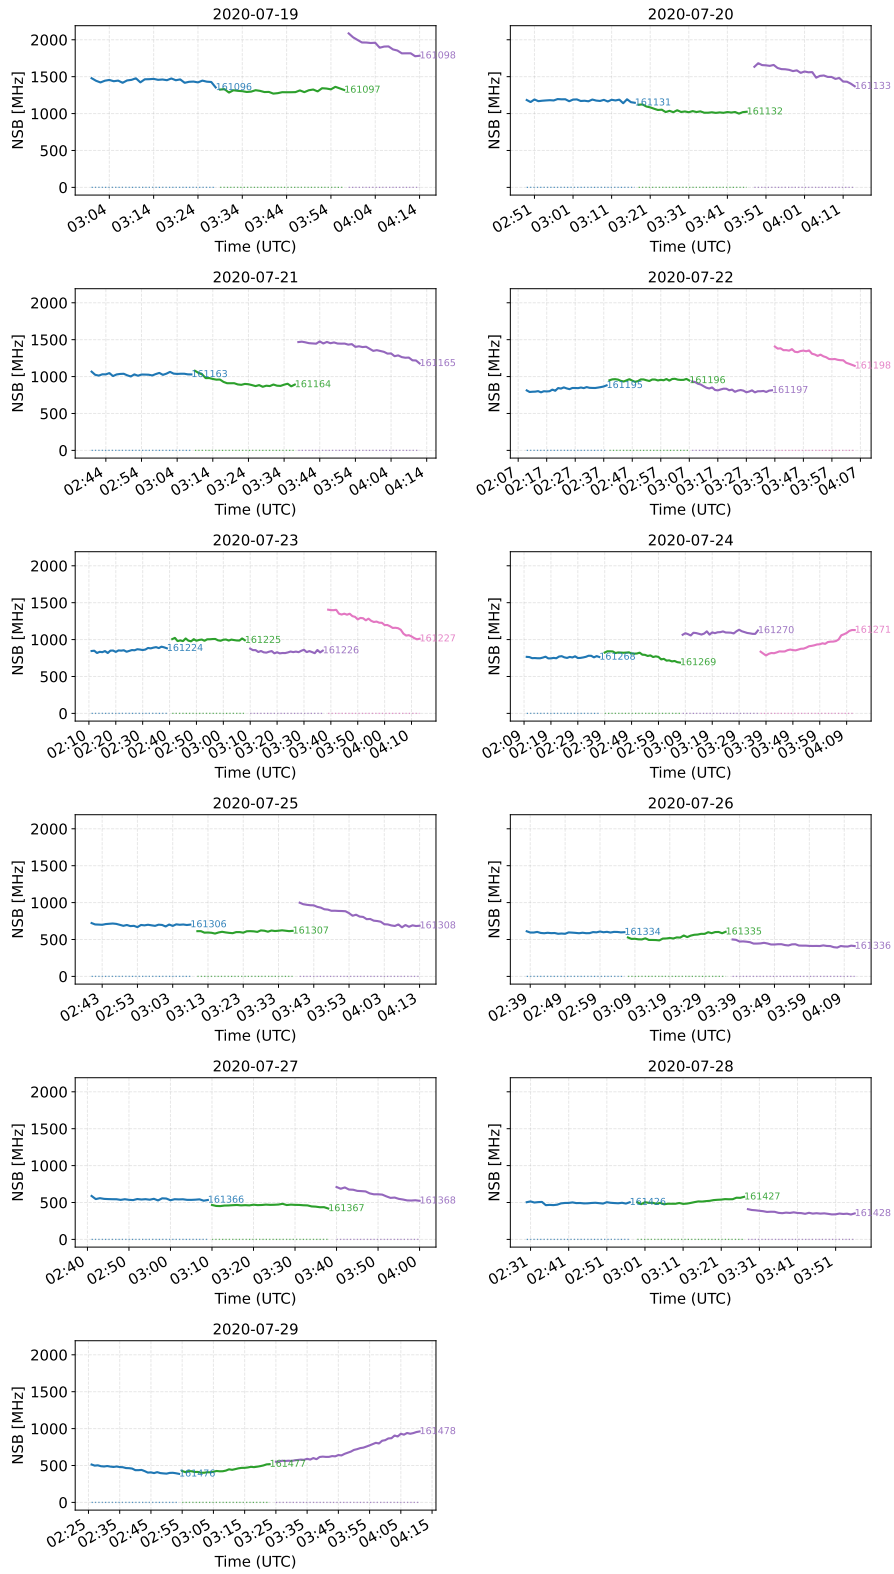


Figure B.6: Nightly NSB light curves for CT2. Each coloured segment corresponds to an individual run.

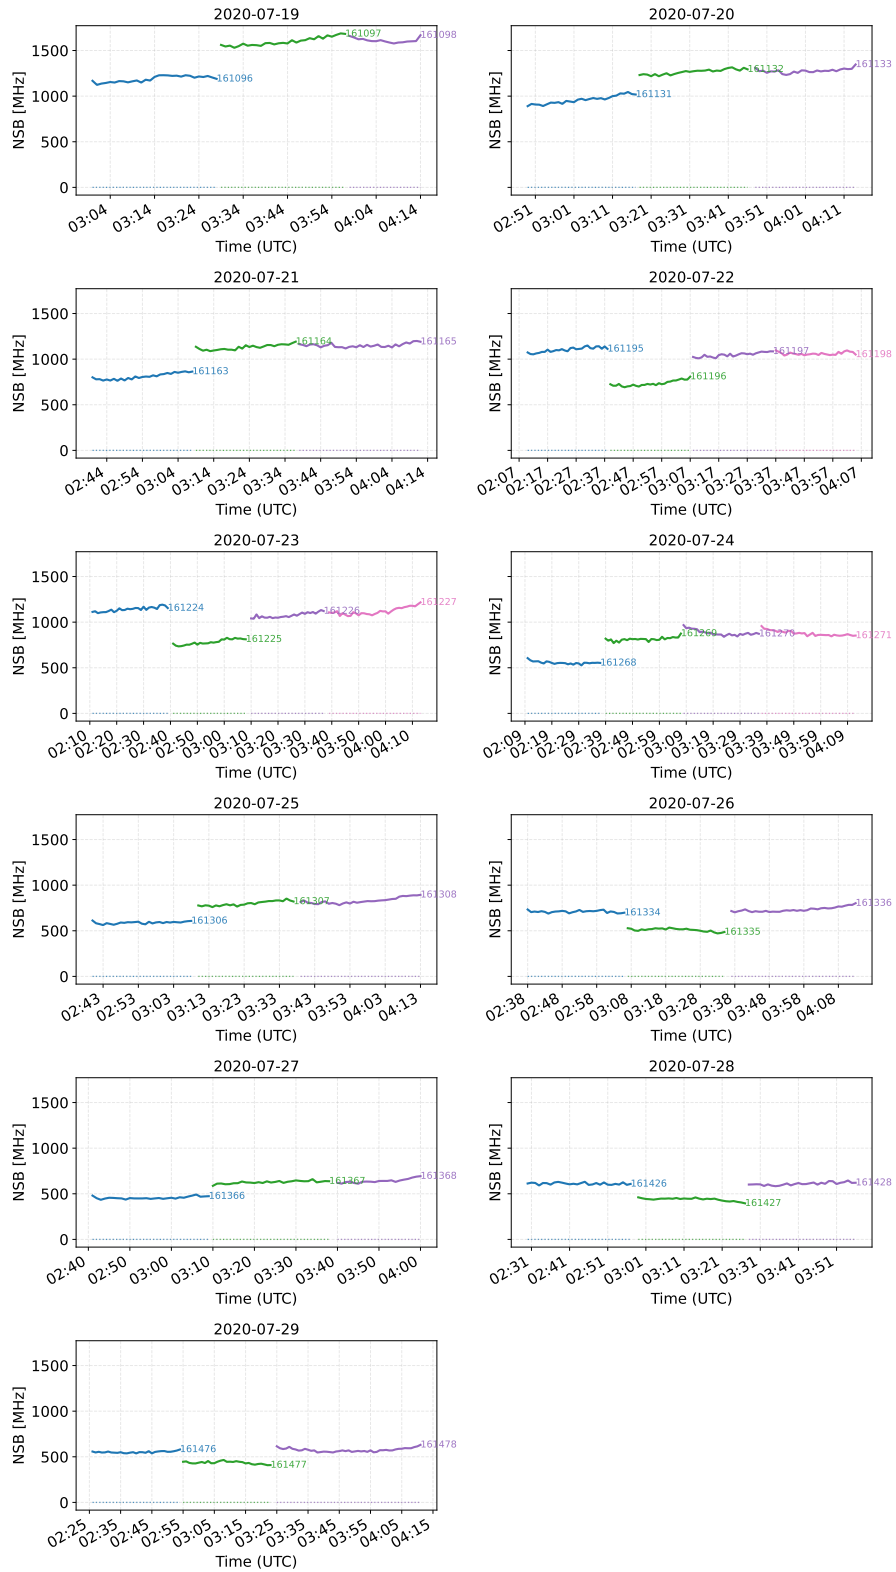


Figure B.7: Nightly NSB light curves for CT3. Each coloured segment corresponds to an individual run.

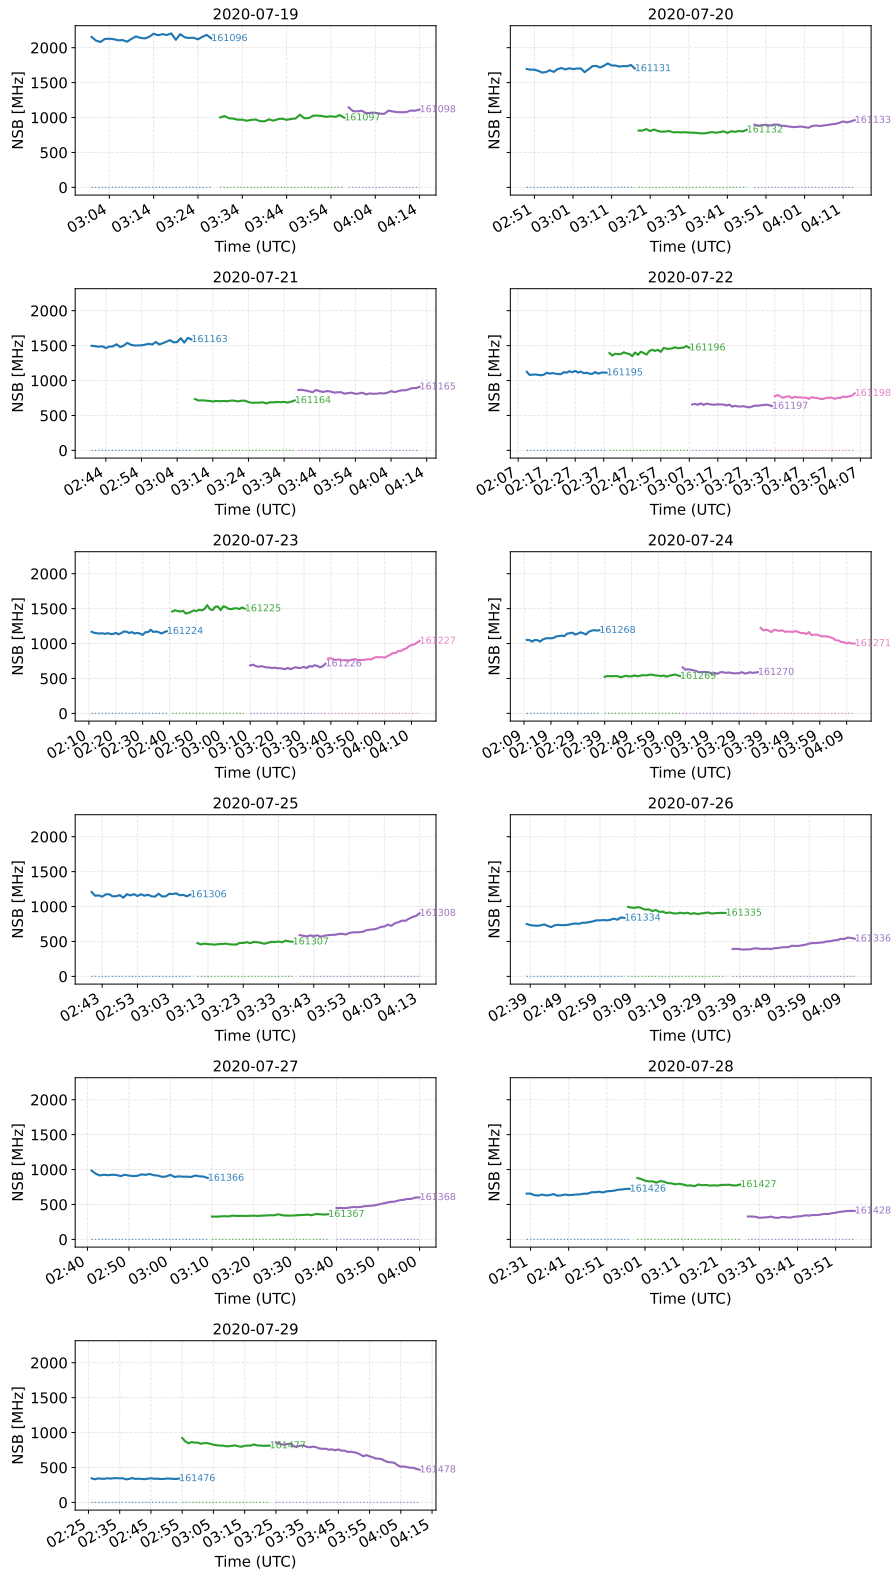


Figure B.8: Nightly NSB light curves for CT4. Each coloured segment corresponds to an individual run.

Appendix C

Supplementary figures




Figure C.1: Combined sky image from all seven eROSITA cameras in the 0.2–0.6 keV band. At $t_0 = 2020-07-07 \ 16:47:20.64$ (TT), during the second all-sky survey, eROSITA detected a bright, previously unknown soft X-ray flash that was strongly affected by pile-up. No source is visible in the scans obtained four hours before or after the event. Adapted from König et al. (2022)

Appendix D

Supplementary tables

Table D.1: Example run metadata used in this work for run 161096 (MGAB-V207), as returned by the run-metadata query script.

Field	Value
Entry	396528
WhenEntered	2022-08-17 07:19:07
Run	161096
RunType	ObservationRun
Target_Name	ToO MGAB-V207
Offset_Mode	RA/Dec
Offset_x	0.87
Offset_y	0.0
Target_RA	59.623125
Target_Dec	-54.778111
Number_Of_Telescopes	5
Telescope_Pattern	62
Datasets_Pattern	46
Run_Start_Time	2020-07-19 02:58:28
First_Event_Time	2020-07-19 03:00:14
Duration	1683.0
Events	4453566

Table D.2: Literature values for systems overplotted on the escape-limited E_{\max} contour map (cf. figure 4.6). Only parameters with an explicit source are listed. The wind velocity was assumed to be $\sim 10 \text{ km s}^{-1}$ for similar novae wherever a value was not found in literature

Nova	v_{sh} [km s^{-1}]	\dot{M} [$M_{\odot} \text{ yr}^{-1}$]	v_{wind} [km s^{-1}]	\dot{M}/v_{wind} [kg m^{-1}]
RS Oph	4500 (MAGIC Collaboration, 2022)	–	~ 10	$(5.7\text{--}7.4) \times 10^{11}$ (fiducial 6.0×10^{11}) (H.E.S.S. Collaboration, 2022)
V407 Cyg	~ 3200 (L. Chomiuk et al., 2012)	$\sim (0.7\text{--}2) \times 10^{-6}$ (L. Chomiuk et al., 2012)	–	$\sim (4.4 \times 10^{12}\text{--}1.3 \times 10^{13})$
V3890 Sgr	$v(t) = (1840 \pm 490)t^{-0.75 \pm 0.13}$ (with t in days) (A. Evans et al., 2022)	3.2×10^{-8} (Nyamai et al., 2023)	–	$\sim 2.0 \times 10^{11}$
V745 Sco	$\gtrsim 4000$ (Banerjee et al., 2014)	$\lesssim 6 \times 10^{-7}$ (Molina et al., 2024)	–	$\lesssim 3.8 \times 10^{12}$

Table D.3: Literature values for novae plotted in figure 4.4 for ejecta mass M_{ej} and a characteristic velocity (v_{sh} taken here as a shock speed when available, otherwise an outer/maximum ejecta velocity). Only entries with an explicit source are listed; uncertainties/ranges are included only when stated in the cited reference.

Nova	v_{sh} [km s^{-1}]	M_{ej} [M_{\odot}]	Literature
RS Oph (2021)	4000–5000	$(3.54\text{--}3.83) \times 10^{-6}$	MAGIC Collaboration (2022) and H.E.S.S. Collaboration (2022)
V407 Cyg (2010)	~ 3200	$\sim 10^{-4}$	L. Chomiuk et al. (2012)
V959 Mon (2012)	2280	1.1×10^{-4}	Healy et al. (2017)
V1324 Sco (2012)	1150 ± 40	$(2.0 \pm 0.4) \times 10^{-5}$	Martin et al. (2018)

Table D.4: Literature values for the novae plotted in figure 4.8

Nova Name	Distance and Error (pc)	Velocity (km/s)
RS Oph	1400 ± 400 (Barry et al., 2008)	4500 (H.E.S.S. Collaboration, 2022)
T CrB	915.82 ± 23.00 (Wikipedia)	5000 (arXiv:2405.01257)
V3890 Sgr	4905.62 ± 0 (Wikipedia)	4200 (ATel 13047)
YZ Ret	2703 ± 300 (Wikipedia)	2700 (ATel 13867)

Table D.5: Nightly mean NSB rates (with standard errors) for CT1–CT4 and the combined average over CT1–CT4.

Date (UTC)	Night	CT1	CT2	CT3	CT4	CT1–4 mean
2020-07-19	1	(1334.191 ± 6.309) MHz	(1501.162 ± 14.166) MHz	(1452.356 ± 12.584) MHz	(1431.633 ± 34.369) MHz	(1431.318 ± 9.745) MHz
2020-07-20	2	(1068.187 ± 7.469) MHz	(1259.366 ± 12.834) MHz	(1176.101 ± 8.690) MHz	(1116.337 ± 24.072) MHz	(1157.723 ± 7.513) MHz
2020-07-21	3	(944.226 ± 8.072) MHz	(1136.866 ± 11.571) MHz	(1049.939 ± 8.639) MHz	(989.071 ± 19.716) MHz	(1032.774 ± 6.633) MHz
2020-07-22	4	(862.602 ± 5.639) MHz	(989.682 ± 9.938) MHz	(989.726 ± 7.663) MHz	(961.798 ± 16.252) MHz	(951.563 ± 5.400) MHz
2020-07-23	5	(931.713 ± 6.335) MHz	(1004.968 ± 8.514) MHz	(1038.184 ± 7.123) MHz	(1003.876 ± 15.832) MHz	(994.679 ± 5.065) MHz
2020-07-24	6	(818.794 ± 11.036) MHz	(901.859 ± 7.059) MHz	(800.111 ± 6.712) MHz	(858.023 ± 14.977) MHz	(845.918 ± 5.225) MHz
2020-07-25	7	(691.054 ± 5.040) MHz	(716.947 ± 5.956) MHz	(752.029 ± 5.983) MHz	(751.878 ± 15.873) MHz	(727.399 ± 4.549) MHz
2020-07-26	8	(628.248 ± 5.409) MHz	(510.297 ± 4.041) MHz	(659.480 ± 5.679) MHz	(682.709 ± 12.089) MHz	(616.907 ± 4.074) MHz
2020-07-27	9	(534.734 ± 3.161) MHz	(532.348 ± 4.159) MHz	(573.077 ± 5.279) MHz	(586.654 ± 15.394) MHz	(555.752 ± 4.151) MHz
2020-07-28	10	(512.155 ± 3.940) MHz	(452.826 ± 3.928) MHz	(554.575 ± 4.776) MHz	(595.231 ± 11.536) MHz	(526.014 ± 3.618) MHz
2020-07-29	11	(642.687 ± 7.422) MHz	(578.950 ± 8.473) MHz	(532.888 ± 3.279) MHz	(640.384 ± 10.670) MHz	(597.864 ± 4.120) MHz

Bibliography

- Abdo, A. A. et al. (2010). “Gamma-Ray Emission Concurrent with the Nova in the Symbiotic Binary V407 Cygni”. In: *Science* 329.5993, pp. 817–821. DOI: 10.1126/science.1192537.
- Aharonian, F. et al. (2004). “Calibration of cameras of the H.E.S.S. detector”. In: *Astroparticle Physics* 22.2, pp. 109–125. DOI: 10.1016/j.astropartphys.2004.06.006.
- Aharonian, F. and H.E.S.S. Collaboration (Oct. 2006). “Observations of the Crab nebula with H.E.S.S.” In: *Astronomy & Astrophysics* 457.3, pp. 899–915. DOI: 10.1051/0004-6361:20065351. arXiv: astro-ph/0607333.
- Aydi, Elias et al. (2025). “Multiple outflows and delayed ejections revealed by early observations of a nova eruption”. In: *Nature Astronomy*. DOI: 10.1038/s41550-025-02725-1. URL: <https://www.nature.com/articles/s41550-025-02725-1>.
- Aydi, Elias, Kirill V. Sokolovsky, Laura Chomiuk, Elad Steinberg, et al. (2020). “Direct evidence for shock-powered optical emission in a nova”. In: *Nature Astronomy* 4, pp. 776–780. DOI: 10.1038/s41550-020-1070-y. URL: <https://www.nature.com/articles/s41550-020-1070-y>.
- Banerjee, D. P. K., V. Joshi, V. Venkataraman, N. M. Ashok, et al. (2014). “Near-IR Studies of Recurrent Nova V745 Scorpii during its 2014 Outburst”. In: *The Astrophysical Journal Letters* 785.1, p. L11. DOI: 10.1088/2041-8205/785/1/L11. arXiv: 1403.0651 [astro-ph.HE].
- Berge, D., S. Funk, and J. Hinton (2007). “Background modelling in very-high-energy γ -ray astronomy”. In: *Astronomy & Astrophysics* 466.3, pp. 1219–1229. DOI: 10.1051/0004-6361:20066704. arXiv: astro-ph/0610959.
- Bernlöhr, K. et al. (2003). “The optical system of the H.E.S.S. imaging atmospheric Cherenkov telescopes. Part I: Layout and components of the system”. In: *Astroparticle Physics* 20.2, pp. 111–128. DOI: 10.1016/S0927-6505(03)00171-3. arXiv: astro-ph/0308246.
- Blandford, Roger and David Eichler (1987). “Particle Acceleration at Astrophysical Shocks: A Theory of Cosmic Ray Origin”. In: *Physics Reports* 154.1, pp. 1–75. DOI: 10.1016/0370-1573(87)90134-7. URL: <https://ui.adsabs.harvard.edu/abs/1987PhR...154....1B/abstract>.
- Bode, Michael F. and Alan Evans, eds. (2008). *Classical Novae*. 2nd. Cambridge: Cambridge University Press.

- Caprioli, Damiano and Anatoly Spitkovsky (2014). “Simulations of Ion Acceleration at Non-relativistic Shocks. I. Acceleration Efficiency”. In: *The Astrophysical Journal* 783.2, 91, p. 91. DOI: 10.1088/0004-637X/783/2/91. URL: <https://ui.adsabs.harvard.edu/abs/2014ApJ...783...91C/abstract>.
- Chomiuk, L., M. I. Krauss, M. P. Rupen, et al. (2012). “The Radio Light Curve of the Gamma-Ray Nova in V407 Cyg: Thermal Emission from the Ionized Symbiotic Envelope, Devoured from within by the Nova Blast”. In: *The Astrophysical Journal* 761.2, p. 173. DOI: 10.1088/0004-637X/761/2/173.
- Chomiuk, Laura, Brian D. Metzger, and Ken J. Shen (2021). “New Insights into Classical Novae”. In: *Annual Review of Astronomy and Astrophysics* 59. DOI: 10.1146/annurev-astro-112420-114502. URL: <https://ui.adsabs.harvard.edu/abs/2021ARA%26A...59..391C>.
- Conceição, Ruben, Jonas Glombitza, and Andrew Smith (Jan. 28, 2025). *The Southern Wide-field Gamma-ray Observatory: Configuration Consolidation Report (SWG0 Deliverable D7.2)*. Tech. rep. SWGO-AS/Site-24-004. Accessed: 2026-01-30. SWGO Collaboration. URL: https://www.swgo.org/SWGOwiki/lib/exe/fetch.php?media=simulations:configuration_consolidation_report_d7_2.pdf.
- Conceição, Ruben and SWGO Collaboration (2023). *The Southern Wide-field Gamma-ray Observatory*. arXiv: 2309.04577 [astro-ph.HE].
- Evans, A., D. P. K. Banerjee, R. D. Gehrz, et al. (2022). “The 2019 outburst of the recurrent nova V3890 Sagittarii”. In: *Monthly Notices of the Royal Astronomical Society* 515, pp. 1278–1297.
- Fermi Science Support Center (2024). *Fermi Cycle 17 Approved Programs*. Mission document (PDF). Accessed: 2026-01-15. URL: <https://fermi.gsfc.nasa.gov/ssc/proposals/cycle17/Cycle17ApprovedPrograms.pdf>.
- Fermi-LAT Collaboration (2014). “Fermi Establishes Classical Novae as a Distinct Class of Gamma-Ray Sources”. In: *Science* 345.6196, pp. 554–558. DOI: 10.1126/science.1253947.
- Franckowiak, Anna, Pierre Jean, Matthew Wood, C. C. Cheung, et al. (2018). “Search for gamma-ray emission from Galactic novae with the *Fermi*-LAT”. In: *Astronomy & Astrophysics* 609, A120, A120. DOI: 10.1051/0004-6361/201731516. URL: https://www.aanda.org/articles/aa/full_html/2018/01/aa31516-17/aa31516-17.html.
- Funk, S., G. Hermann, J. Hinton, D. Berge, et al. (2004). “The trigger system of the H.E.S.S. telescope array”. In: *Astroparticle Physics* 22.3–4, pp. 285–296. DOI: 10.1016/j.astropartphys.2004.08.001. arXiv: astro-ph/0408375.
- Giavitto, G. et al. (2016). “Upgraded cameras for the H.E.S.S. imaging atmospheric Cherenkov telescopes”. In: *Proc. SPIE 9908, Ground-based and Airborne Instrumentation for Astronomy VI*. DOI: 10.1117/12.2231865.
- H.E.S.S. Collaboration (2022). “Time-resolved hadronic particle acceleration in the recurrent nova RS Ophiuchi”. In: *Science* 376.6588, pp. 77–80. DOI: 10.1126/science.abn0567.

- H.E.S.S. Collaboration (2025). *H.E.S.S. – High Energy Stereoscopic System*. URL: <https://www.mpi-hd.mpg.de/hfm/HESS/> (visited on 12/12/2025).
- Hahn, Joachim (2016). “GAMERA – A Modular Framework for Spectral Modeling in VHE Astronomy”. In: *Proceedings of Science*. Vol. ICRC2015. PoS. Accessed: 2026-01-25, p. 917. DOI: 10.22323/1.236.0917. URL: <https://pos.sissa.it/236/917/pdf>.
- Healy, A. A. et al. (2017). “Multi-epoch radio imaging of γ -ray Nova V959 Mon”. In: *Monthly Notices of the Royal Astronomical Society* 469.4, pp. 3976–3994.
- Hillas, A. M. (1985). “Cerenkov light images of EAS produced by primary gamma rays and by nuclei”. In: *Proceedings of the 19th International Cosmic Ray Conference (ICRC)* 3, pp. 445–448.
- Hinton, J. A. and W. Hofmann (2009). “Teraelectronvolt Astronomy”. In: *Annual Review of Astronomy and Astrophysics* 47, pp. 523–565. DOI: 10.1146/annurev-astro-082708-101816.
- Holben, B. N. et al. (1998). “AERONET—A Federated Instrument Network and Data Archive for Aerosol Characterization”. In: *Remote Sensing of Environment* 66.1, pp. 1–16. DOI: 10.1016/S0034-4257(98)00031-5.
- IAU, Sky & Telescope magazine, Sinnott, Roger, and Fienberg, Rick (2025). *Milky Way with constellation outlines (IAU / Sky & Telescope)*. Wikimedia Commons. CC BY 3.0. Wikimedia Commons file page (curid=15412184). URL: <https://commons.wikimedia.org/w/index.php?curid=15412184> (visited on 12/28/2025).
- Joshi, Vikas (2019). “Reconstruction and Analysis of Highest Energy γ -Rays and its Application to Pulsar Wind Nebulae”. heiDOK document 26062; thesis defense: 2019-02-06; deposited: 2019-02-15. Dissertation. Heidelberg University (Universität Heidelberg). DOI: 10.11588/heidok.00026062. URL: <https://archiv.ub.uni-heidelberg.de/volltextserver/26062/>.
- Kasten, F. and A. T. Young (1989). “Revised optical air mass tables and approximation formula”. In: *Applied Optics* 28.22, pp. 4735–4738. DOI: 10.1364/AO.28.004735.
- Kato, Mariko and Izumi Hachisu (2023). “Theoretical Light-curve Models of the Symbiotic Nova CN Cha: Optical Flat Peak for 3 Yr”. In: *The Astrophysical Journal* 951.2, 128, p. 128. DOI: 10.3847/1538-4357/acdb4c. URL: <https://doi.org/10.3847/1538-4357/acdb4c>.
- König, Ole, Jörn Wilms, Riccardo Arcodia, Thomas Dauser, et al. (2022). “X-ray detection of a nova in the fireball phase”. In: *Nature* 605.7909, pp. 248–250. DOI: 10.1038/s41586-022-04635-y. arXiv: 2209.05125 [astro-ph.HE]. URL: <https://arxiv.org/abs/2209.05125>.
- Lang, Thomas (2025). *Satellite-trails*. GitHub repository. URL: <https://github.com/Tho-Lang/Satellite-trails> (visited on 02/01/2025).
- Leitl, Franziska (June 2021). “Core and Energy Estimations for SWGO Using a Template-based Shower Reconstruction Method”. Erlangen Centre for Astroparticle Physics (ECAP). Supervisor: Prof. Dr. Christopher van Eldik. Master’s Thesis in Physics. Er-

- langen, Germany: Friedrich-Alexander-Universität Erlangen-Nürnberg. URL: https://ecap.nat.fau.de/wp-content/uploads/2021/07/2021_MA_Leit1.pdf (visited on 01/24/2026).
- MAGIC Collaboration (2022). “Proton acceleration in thermonuclear nova explosions revealed by gamma rays”. In: arXiv:2202.07681 (v3, 10 Nov 2022). DOI: 10.48550/arXiv.2202.07681. arXiv: 2202.07681 [astro-ph.HE].
- Martin, Pierrick, Guillaume Dubus, Pierre Jean, Vincent Tatischeff, and Cyril Dosne (2018). “Gamma-ray emission from internal shocks in novae”. In: *Astronomy & Astrophysics* 612, A38. DOI: 10.1051/0004-6361/201731692. arXiv: 1710.05515 [astro-ph.HE]. URL: https://www.aanda.org/articles/aa/full_html/2018/04/aa31692-17/aa31692-17.html.
- Matthews, J. (2005). “A Heitler model of extensive air showers”. In: *Astroparticle Physics* 22.5–6, pp. 387–397. DOI: 10.1016/j.astropartphys.2004.09.003.
- Metzger, Brian D., Thomas Finzell, Indrek Vurm, Romain Hascoët, et al. (2015). “Gamma-ray novae as probes of relativistic particle acceleration at non-relativistic shocks”. In: *Monthly Notices of the Royal Astronomical Society* 450.3, pp. 2739–2748. DOI: 10.1093/mnras/stv742.
- Molina, I. et al. (2024). “The symbiotic recurrent nova V745 Sco at radio wavelengths”. In: *Monthly Notices of the Royal Astronomical Society*. DOI: 10.1093/mnras/stae2093.
- Mukai, Koji (2025). *Fermi/LAT Detected Novae*. NASA Goddard Space Flight Center web page. Last updated: 2025-09-27, accessed: 2026-01-15. URL: <https://asd.gsfc.nasa.gov/Koji.Mukai/novae/latnovae.html>.
- Nyamai, M. M., J. D. Linford, J. R. Allison, et al. (2023). “Synchrotron emission from double-peaked radio light curves of the symbiotic recurrent nova V3890 Sgr”. In: *Monthly Notices of the Royal Astronomical Society* 523.2, pp. 1661–1682. DOI: 10.1093/mnras/stad1534.
- Page, K. L., A. P. Beardmore, J. P. Osborne, et al. (2020). “Neil Gehrels *Swift* Observatory studies of supersoft novae”. In: *Advances in Space Research* 66.5, pp. 1169–1180. DOI: 10.1016/j.asr.2019.11.017. URL: <https://ui.adsabs.harvard.edu/abs/2020AdSpR..66.1169P/abstract>.
- Particle Data Group (2024). *Passage of Particles Through Matter*. Review of Particle Physics (2024), PDG review chapter. Accessed: 2026-01-22. URL: <https://pdg.lbl.gov/2024/reviews/rpp2024-rev-passage-particles-matter.pdf>.
- (2025). *Cosmic Rays*. Review of Particle Physics (2025), PDG review chapter. Accessed: 2026-01-22. URL: <https://pdg.lbl.gov/2025/reviews/rpp2025-rev-cosmic-rays.pdf>.
- Payne-Gaposchkin, Cecilia H. (1957). *The Galactic Novae*. Interscience Monographs and Texts in Physics and Astronomy. Amsterdam: North-Holland Publishing Company.
- Pierre Auger Collaboration (2023). “Search for photons above 10^{19} eV with the surface detector of the Pierre Auger Observatory”. In: *Journal of Cosmology and Astroparticle*

- Physics* 2023.05, p. 021. DOI: 10.1088/1475-7516/2023/05/021. arXiv: 2209.05926 [astro-ph.HE].
- Steppa, Constantin (Jan. 2022). “Modelling the galactic population of very-high-energy gamma-ray sources”. PhD thesis. DOI: 10.25932/publishup-54947.
- SWG0 Collaboration (July 2019). *Exploring the Extreme Universe: International Collaboration for a New Gamma-ray Observatory Launched*. https://www.swgo.org/SWG0Wiki/doku.php?id=press_release_en. Press release (July 1, 2019), accessed 2026-01-29.
- (2025). “Ground-based gamma-ray astronomy with the Southern Wide-field Gamma-ray Observatory (SWG0)”. In: White paper (arXiv:2506.01786v2). arXiv: 2506.01786 [astro-ph.HE].
- (2026). *Users Guide — pyswgo (IRF production documentation)*. pyswgo internal documentation. Version 0.1.dev61+g75e9e97fb.d20260125, last updated 2026-01-25. URL: https://swgo-collaboration.gitlab.io/irf-production/user_guide.html#producing-instrument-response-functions-irfs (visited on 01/30/2026).
- Tappert, C., N. Vogt, A. Ederoclite, L. Schmidtobreick, et al. (2020). “The luminosity evolution of nova shells. I. A new analysis of old data”. In: *Astronomy & Astrophysics* 641, A122. DOI: 10.1051/0004-6361/202037913. arXiv: 2007.10940 [astro-ph.SR].
- Vurm, Indrek and Brian D. Metzger (2018). “High-energy Emission from Nonrelativistic Radiative Shocks: Application to Gamma-ray Novae”. In: *The Astrophysical Journal* 852.2, 62, p. 62. DOI: 10.3847/1538-4357/aa9c4a. arXiv: 1611.04532 [astro-ph.HE]. URL: <https://ui.adsabs.harvard.edu/abs/2018ApJ...852...62V/abstract>.
- Wang, H.-H. et al. (2024). “Evidence of the γ -ray counterpart of nova FM Cir from *Fermi* and *Swift*”. In: *Monthly Notices of the Royal Astronomical Society: Letters* 531.1, pp. L63–L68. DOI: 10.1093/mnrasl/slae033. URL: <https://academic.oup.com/mnrasl/article/531/1/L63/7645456>.

Acknowledgements

I would like to express my sincere gratitude to Dr. Alison Mitchell for her supervision, guidance, and continued support throughout this thesis, and for giving me this wonderful opportunity. I also thank Sam for his valuable guidance during the early phases of this work, which helped shape the direction of the project from the outset. I am grateful to Martin for his help with SWGO IRFs in particular and for his careful review and constructive feedback on the thesis as a whole. I further thank Gerrit for his support with the H.E.S.S. analysis, as well as for detailed feedback and helpful discussions.

I am also deeply thankful to my friends for their constant encouragement, understanding, and companionship throughout this journey.

Finally, I would like to thank my parents for their constant encouragement and unwavering support throughout my studies.

For language polishing, I used the tools Writefull and Grammarly for grammar, punctuation, and language refinement, and Canva to create selected schematic figures; responsibility for the scientific content, analysis, and conclusions remains entirely my own.

Declaration of Originality

I, Sayak Ghosh, student registration number: 23268734, hereby confirm that I completed the submitted work independently and without the unauthorized assistance of third parties and without the use of undisclosed and, in particular, unauthorized aids. This work has not been previously submitted in its current form or in a similar form to any other examination authorities and has not been accepted as part of an examination by any other examination authority.

Where the wording has been taken from other people's work or ideas, this has been properly acknowledged and referenced. This also applies to drawings, sketches, diagrams and sources from the Internet.

In particular, I am aware that the use of artificial intelligence is forbidden unless its use as an aid has been expressly permitted by the examiner. This applies in particular to chatbots (especially ChatGPT) and such programs in general that can complete the tasks of the examination or parts thereof on my behalf.

Any infringements of the above rules constitute fraud or attempted fraud and shall lead to the examination being graded "fail" ("nicht bestanden").

Place, Date

Signature