A COMPLETE GUIDE TO THE SOFTWARE

AstroTracker

Everything it does, and why it does it

AstroTracker is a program that keeps track of astrophotography — the hobby of taking photographs of galaxies, nebulae and star clusters from your own back garden. This document explains what the program does, feature by feature, from the ground up. It assumes you have never pointed a telescope at anything.

Written for two readers: someone who has never heard of any of this, and an astrophotographer deciding whether the program is worth their time.
Covers version 0.23.3.

How to read this

This document is built in layers. You do not have to read it in order, but it is written so that you can.

Part One explains the hobby. Not the software — the hobby. What astrophotographers actually do on a clear night, what the equipment is, what the words mean, and why the whole business generates a mountain of files and decisions that nobody can hold in their head. If you already own a telescope, skip it; you will learn nothing. If you do not, read it, because nothing in Part Three onward will make sense without it.

Part Two says what AstroTracker is in a page, states the handful of principles it is built on, and covers installing it and what it costs. The principles explain most of the design decisions further in.

Part Three is the foundation: how the program reads your hard drive and works out, without being told, which photographs belong together. Almost everything else depends on this, so it gets its own part.

Part Four is the bulk of the document — every screen in the program, with what it shows and what you can do with it.

Parts Five to Eight cover the things that happen while you are asleep, the phone access, the numbers the program computes and what they mean, and an honest account of what leaves your computer and what the program does *not* do.

Technical terms are explained the first time they appear and collected in a glossary at the end. Where a section gets into the reasoning behind a design, it is marked like this:

> **WHY IT WORKS THIS WAY**
>
> These boxes explain a decision — usually one that was made after something simpler turned out to be wrong. They can be skipped without losing the thread.

> **IN PLAIN TERMS**
>
> These boxes restate something technical in everyday language. If a paragraph lost you, the plain-terms box after it is the same thing said again more simply.

Contents

What astrophotography actually involves

Before any of the software makes sense, you need to know what the person using it was doing the night before.

Astrophotography is photographing things in deep space — galaxies, nebulae, star clusters — usually from an ordinary back garden. The objects are not small. The Andromeda Galaxy is six times wider than the full moon in our sky. The problem is not size. The problem is that they are unimaginably *faint*. A camera that photographs a wedding in a thirtieth of a second needs, for these, to sit with its shutter open for several minutes, and then to do that again a hundred times.

A single photograph is not a photograph

Here is the part that surprises everyone. The finished picture you have seen of the Orion Nebula is not one exposure. It is typically between fifty and a thousand separate exposures, each three to ten minutes long, added together by a computer afterwards.

The reason is noise. Every digital camera sensor produces a faint random fizz — some from the electronics, some from the sky itself, which is never truly black because of moonlight, town lights and the atmosphere glowing. On a normal daytime photograph this fizz is invisible because the subject is a million times brighter. On a galaxy it is not; the fizz and the galaxy are about equally faint.

Randomness has a useful property: it partly cancels itself out when you average it. The signal from the galaxy is in the same place in every frame. The noise is in a different place every time. Average a hundred frames together and the galaxy stays put while the fizz shrinks. It shrinks slowly — by the square root of the number of frames, so four frames are twice as good as one, and a hundred frames are ten times as good — which is exactly why people shoot so many.

> **IN PLAIN TERMS**
>
> Imagine trying to hear a whisper in a room full of static. One recording is useless. But if you record the same whisper a hundred times and play them all at once, the whisper adds up and the static, being different each time, mostly cancels. That is stacking, and it is the entire basis of the hobby.

Each of those individual exposures is called a **sub** — short for sub-exposure, because it is one part of the whole. The total time of all the subs added together is the **integration**. When an astrophotographer says "I have twelve hours on the Horsehead," they mean twelve hours of total shutter-open time, accumulated over however many nights it took.

It takes many nights

Twelve hours of darkness does not happen in one night. Real nights are short, partly cloudy, interrupted by the moon rising, and stopped by dew, wind and equipment failure. A serious target is assembled across weeks or months. A project started in September may not be finished until January — and by then the object has moved out of the evening sky and you have to wait a year.

So the hobby has an unusual shape: long-running projects, each made of many nights, each night made of many files, all of which must eventually be found again and added together correctly. Lose track of which files belong to which project and the work is wasted.

What the equipment is

A working setup — the hobby calls it a **rig** — is four things bolted together:

The mount

> A motorised tripod head that turns slowly westward to cancel the Earth's rotation. Without it every star becomes a streak. It is the most important and most expensive part, and the one that fails most interestingly.

The telescope

> Here it is just a lens. What matters is its focal length (how magnified the view is) and its aperture (how much light it gathers). Many astrophotographers own several and swap between them for different targets.

The camera

> A dedicated astronomy camera, cooled to around −10 °C to reduce that electronic fizz. Usually monochrome — it sees only in grey.

The filter wheel

> If the camera is monochrome, colour comes from photographing the same object repeatedly through red, green and blue filters and combining them. There are also **narrowband** filters that admit only one precise colour of light emitted by glowing gas — hydrogen, oxygen, sulphur — and block everything else, including most moonlight and streetlight. This is how people photograph nebulae from cities.

On top of that sits the software: a program on a laptop that runs the whole night unattended — slewing to the target, checking it is centred, focusing, and working through a list of exposures. Two programs matter for this document: **Sequence Generator Pro** (usually "SGP"), which runs the night, and **PHD2**, which does guiding.

Guiding, focus and seeing

Three things go wrong constantly, and understanding them explains half of AstroTracker.

Guiding. No mount tracks perfectly. Over a five-minute exposure a small error turns round stars into short dashes. So a second, small camera watches one star and, many times a minute, tells the mount to nudge. How well this is going is measured as **RMS error** in arc-seconds — a measure of angle so small that the full moon is 1,800 arc-seconds across. Good guiding is under about one arc-second. When guiding fails — a cloud eats the guide star — the mount drifts and everything shot in that period is ruined.

Focus. A telescope tube shrinks as it cools through the night, which shifts the focus. Sharp at 10 p.m., soft by 2 a.m. So the software re-focuses every hour or two, or whenever the temperature drops a degree.

Seeing. The atmosphere is turbulent. Starlight passing through it wobbles, and a star that ought to be a point is smeared into a small blurry disc. How big that disc is — measured as **FWHM**, roughly "how wide the star is halfway down its brightness" — is the seeing. It is weather, not equipment: nothing you own can fix a night of bad seeing, and on a good night the same telescope produces visibly sharper images. It changes hour to hour.

> **IN PLAIN TERMS**
>
> Seeing is the twinkle. Stars twinkle because the air is moving. The more they twinkle, the more smeared your photographs are. Some nights are steady and some are not, and there is nothing you can do about it except know which kind of night you had.

The other three kinds of frame

The real photographs of the target are called **lights**. But a raw light frame is not usable on its own, because the camera and telescope each add their own fixed defects: hot pixels that always read bright, dust shadows on the sensor window, and a natural darkening toward the corners of any lens. Three other kinds of frame measure those defects so they can be subtracted out. Together they are called **calibration frames**.

KIND	HOW IT IS TAKEN	WHAT IT REMOVES
Dark	With the lens cap on, same exposure length and same temperature as the lights.	The sensor's own heat glow and its hot pixels.
Bias	Lens cap on, shortest possible exposure.	The fixed electrical offset the camera adds to every reading.
Flat	Pointed at an evenly lit surface — dawn sky, or a lightbox.	Dust shadows and corner darkening. Must be re-shot whenever anything is unscrewed.

The rules around these are fiddly and they are the source of a great deal of wasted work. Darks must match the exposure length, the temperature and the camera's gain settings. Flats must match the filter

and the exact optical arrangement — take the camera off to clean it and every old flat is void. Realise in March that you never shot flats for a set of lights in November and there is nothing to be done: the data is compromised and you cannot go back, because the dust has moved.

> **WHY IT WORKS THIS WAY**
>
> This is the single most common way a night's work is lost, and it is entirely a bookkeeping failure rather than a skill failure. It is also, not coincidentally, one of the first things AstroTracker was built to watch for.

The words, and what they mean

A short reference you can come back to. The full glossary is at the end of the document; this is the minimum needed to read the rest.

WORD	MEANING
Sub	One exposure. The building block. A night produces dozens.
Light	A sub of the actual target.
Dark, bias, flat	Calibration frames — see *What astrophotography actually involves*.
Integration	Total exposure time added up across every sub.
Stacking	Combining all the subs into the final image.
Rig	One particular combination of mount, telescope and camera.
Target	The object — a galaxy, a nebula. Usually a catalogue number: M31, NGC 7000, IC 1396.
Project	One target shot on one rig — the unit of work that gets finished.
FITS	The file format astronomy uses. Each file carries a text *header* describing how it was taken.
Header	The label inside every FITS file: target, filter, exposure, temperature, date, telescope, camera.
Filter	Coloured glass in the light path. L, R, G, B are broadband; Ha, OIII, SII are narrowband.
Arc-second	An angle: 1/3600 of a degree. The full moon is about 1,800 across.
FWHM	How wide a star looks. Smaller is sharper. The main quality number.
Eccentricity	How oval a star is. Round is 0. Ovals mean tracking or tilt problems.
Guiding / RMS	Keeping the mount on target, and how badly it is wandering.
Seeing	Atmospheric blurring. Weather, not equipment.
Transparency	How much light the sky is letting through — reduced by haze and thin cloud.
Meridian flip	Mid-night, the mount must swing to the other side of the pier. A common point of failure.
Bortle class	How dark your sky is, 1 (remote desert) to 9 (city centre).
Plate solve	Identifying exactly where a photograph points by matching its star pattern to a catalogue.

Gain / offset	Two camera settings: gain is its sensitivity, offset a small fixed amount added to every reading. Darks are only valid for lights with matching values.

Why this becomes a filing problem

Everything so far is about one night. The difficulty is what happens after two hundred of them.

The arithmetic

A single night on one telescope produces perhaps sixty files, each between twenty and fifty megabytes. Somebody with three telescopes who images whenever it is clear will produce ten to twenty thousand files a year. The library this document was written alongside holds **143,479 light frames**, taken over just under four years with about nine telescopes and lenses — several terabytes, spread over a fast internal drive for current work and a network drive for the archive.

No one remembers 143,479 of anything. And the questions that matter are exactly the ones that require remembering:

- How many hours do I actually have on this target, and how many through each filter?
- Did I ever shoot flats for those November nights?
- Are my darks still valid, or was that before I changed the camera's gain?
- Which of last night's sixty frames are ruined, and by what?
- Is this project finished, or is it two hours short of usable?
- Why does that one night look worse than the others?
- Is tonight worth going out for, given the moon?
- Did last night's session actually run, or did it stop at 1 a.m.?

What people do instead

Almost everyone builds the same three things by hand, and all three break.

A folder naming scheme. Something like `2026\MyScope\M31\Night_3\`. It works until the capture software saves somewhere unexpected, or two targets get shot in one night and land in one folder, or a folder is renamed and now the drive disagrees with the spreadsheet. Folder names are a promise the computer never checks.

A spreadsheet. A row per project, totals per filter, notes on what is missing. It is accurate the week you make it. Then you image at midnight, go to bed, and forget to update it, and within a month it is quietly wrong — which is worse than not having it.

Memory. Which works for the current project and fails completely for the one you set aside eight months ago when it dropped below the trees.

WHY IT WORKS THIS WAY

Notice that every one of these fails for the same reason: they are all a *second* record of something the files themselves already know. Every FITS file already states its target, filter, exposure, temperature and time in its header. The information is not missing. It is simply spread across a hundred thousand files where no human can see it. That observation is the whole idea behind AstroTracker.

And the decisions

Beyond bookkeeping there is a second class of problem: judgement calls that need more information than one person can hold.

Is tonight worth setting up for? That depends on what is above the horizon after dark, how high it gets, when the moon rises, how bright it is, how close it is to your target, which of your filters survive that much moonlight, and whether it is forecast to cloud over at one in the morning. Every one of those is knowable. Putting them together, at six in the evening, while cooking dinner, is not.

Or: you have eight hours on a galaxy. Is it worth another four? Hours alone cannot say. What decides it is how faint the part you care about is — the dim outer arms, a wisp of gas at the edge — against the noise in your subs, and that depends on your sky, your optics, the length of each sub and the target. Four hours under a bright moon are not worth four under a dark one; sixty hours on a faint shell of gas can still be noisy while five on a bright nebula are clean; and the improvement from more data slows down the more you already have. There is a real answer, and it is already in the subs you have. Nobody works it out by hand.

Or, at two in the morning: is it still clear? Did the sequence stop? Is the mount still tracking? Most people answer these by getting out of bed and looking.

AstroTracker exists because all of this is answerable from information the computer already has — the file headers and the subs themselves, the log files the capture software writes anyway, a free weather forecast, and an ordinary camera pointed at the sky.

The short answer

AstroTracker reads the astrophotography files already on your computer, works out what they are, and then answers the questions you would otherwise answer from memory, a spreadsheet, or not at all.

You install it on the Windows PC that holds your images. You tell it which folders your data lives in. It reads every file's header — the text label inside each one — and from that alone it builds a picture of your entire library: which telescope took what, which photographs belong to which target, how many hours you have on each, and what calibration is missing. Tell it where you observe from and it adds what is worth doing tonight.

A label cannot tell a sharp photograph from a cloudy one, so for that one question it looks at the picture itself. It measures each sub: the size and shape of its stars, how bright the sky behind them was, the faint glow thin cloud spreads around bright stars, how much the sky changed against the subs taken either side of it, and how likely it is that cloud passed through. Subs that stand out from the rest of their night are flagged for you to look at (see *Review — judging your photographs*). Measuring reads a file and never changes it, and deciding which subs to throw out stays with you.

From then on it watches those folders. New files appearing at one in the morning are picked up within seconds, added to the right project without being told, and measured shortly after. If the flow of files stops while it is still dark, it says so. If you have an all-sky camera and the sky clouds over, it says so. In the morning there is a summary of what the night produced.

It runs in three places at once, all the same program:

- **A tray icon** on the PC, which is how it starts and stays running.
- **A web page** at `http://127.0.0.1:8765` — this is the actual interface, opened in whatever browser you already use. There is nothing to log in to.
- **Your phone**: the same pages over your home network or a VPN, plus push notifications that reach you anywhere. Both are off until you switch them on.

> **IN PLAIN TERMS**
>
> Think of a warehouse manager who walks the aisles every night, reads the label on every box, keeps a perfect running inventory, checks each new delivery for damage, notices when a delivery stops arriving, and leaves a note on your desk each morning about what came in and what is missing. Nothing inside a box is ever changed and nothing is thrown away.

Five principles it is built on

These are worth stating early, because nearly every design decision in the rest of the document follows from one of them.

1. Your files stay where they are, and stay yours

AstroTracker never modifies an image file. Not to add a keyword, not to fix a header, not ever. It reads them and records what it found in its own database; the thumbnails and measurements it makes are kept in its own folder too, never beside your files. Your photographs are the product of years of work and a program that rewrites them is a program that can destroy them.

When it does move files — and it can, because filing is half the problem — it never does so on its own initiative. Every move is proposed on screen first, confirmed by you, logged, and undoable afterwards. There is also a **Read-only mode**, under **Settings › Safety**, that disables file movement entirely, for people who would rather it never touched anything.

2. Headers are the truth; filenames are a hint

A folder called `M31` might contain anything. The header inside each file states what the telescope was actually pointing at, through which filter, for how long, at what temperature. AstroTracker believes the header. Where a filename or folder name disagrees, the header wins, and the discrepancy is usually itself worth reporting to you.

3. The scan reads labels, not pictures

During a scan, AstroTracker reads only the header of each file, never the image data. This is the difference between reading a few kilobytes and reading fifty megabytes. On a local drive a scan gets through a few thousand files a minute; over a network share about thirty a second, so the first scan of a 100,000-file archive takes roughly an hour. You can carry on using the program while it runs, and the wait happens once: after that, a file whose size and date have not changed is not read again.

Image data is read only for the jobs that need it, never as part of a scan. One of those runs on its own: by default each new sub that lands in a current-work folder is measured shortly after the scan finds it, so its star size and shape, background, thumbnail and signs of cloud are ready in Review by the time you look. **Measure new subs on current-work folders automatically**, under **Settings › Watching and rescans**, turns that off. Everything else waits for you to ask: measuring or re-measuring subs from Review or the Library page, opening a sub at full size in Review, and **How much more?**, which reads a few subs per filter together with their darks and flats.

4. Everything stays on your computer

There is no account, no cloud service, no upload. The database, thumbnails and log live in one folder on your PC (see *Installing it, and where things live*), and backups, if you turn them on, go only to the folder you choose. Your subs and calibration frames are never copied anywhere.

The program does reach the internet for a short list of things, set out in full in *What leaves your computer*: chiefly a weather forecast for your site, a picture of the sky region you are framing, target catalogues, a check for program updates, and push notifications if you turn them on. None of them carries an image of yours.

5. Measure, do not assume

Where the program makes a claim about your data, it is computed from your data. The seeing estimate comes from the stars in your own subs. The advice about which filters survive tonight's moon starts from textbook values and is then re-fitted from your own measured sky brightness. The focuser's temperature coefficient is derived from your own autofocus runs, and until there are enough of them the program says what is still missing rather than inventing a number.

How much more? judges a project not by its hours but by how far the faint parts of the target stand above the noise in a few of your own calibrated subs, and says how many more hours each filter needs to reach the level you count as clean, which is a setting. The **star glow** Review shows on each sub, the haze thin cloud spreads around bright stars, is judged against the clearest this target has looked through that filter, not against a fixed number.

> **WHY IT WORKS THIS WAY**
>
> A tool that tells you what it does not yet know is more useful than one that always has an answer. The focuser coefficient on the Rigs page will say **Collecting** for an entire season rather than fit a confident-looking line through four data points.

Installing it, and where things live

The installation

A single installer, `AstroTracker-Setup-0.23.3.exe`, about 55 MB. It installs for your user account rather than system-wide, which means no administrator prompt — and, more usefully, means the program can later update itself without one. It offers a desktop shortcut, a "start when I sign in" option, and an optional install of Syncthing, unticked to begin with (a free file-syncing tool, used later if you need to bring subs or log files over from a separate observatory PC).

The installer and the program are digitally signed, so Windows shows a verified publisher rather than a warning about an unknown one. Later versions arrive without a trip back to a website: the program checks for new releases itself and can download and install them, and it installs one only when it carries that same signature, as described in *Backups, updates and staying alive*.

Where your data lives

Everything AstroTracker knows lives in one folder: `%LOCALAPPDATA%\AstroTracker`. Inside it are the database (a single file, `astrotracker.db`), the log, and a few folders, mostly of pictures the program has made, fetched or been given:

FOLDER	WHAT IT HOLDS
`thumbs`	Small previews of the subs it has measured. They are made again whenever they are needed, so they are safe to delete.
`skycache`	Sky survey pictures of your targets, fetched once and kept. With an all-sky camera, also up to eight nights of its frames for the timelapse, roughly 50 MB a night, and what the program has learned about which parts of the camera's picture are sky.
`finals`	The finished pictures you upload to your projects, each kept as a JPEG to view and a small one for lists. The file you uploaded stays wherever you exported it.
`updates`	The newest installer the program has downloaded for itself, and a record of the last install.

Back up that one folder and you have backed up everything the program knows; uninstalling asks whether you want it deleted. **Settings › Data on this PC** shows where the folder is and how large the database, the thumbnails and the log have grown, with a button that clears the thumbnails.

Your own subs are never copied there, so the folder stays modest: a library of about 195,000 files, 143,479 of them light frames, produces a database of about 220 MB.

The first run

The first time it opens, a setup wizard walks through six steps after a short welcome, which also lists what the program never does. You can move past any of them, each can be changed later, and each has its own page in the program, so nothing here is a one-time decision.

1. **Folders.** Which folders hold your data, and what role each plays — described in *Telling it where your files are*, because it is the single most important thing you tell the program.
2. **Scan.** Press **Start scan** and it reads the header of every file in those folders, showing as it goes how many files it has seen and read, how many are new, changed or gone, and how many it could not read. On a first run it reads a few thousand files a minute from a local drive and about thirty a second over a network share, so a 100,000-file archive takes roughly an hour. You need not wait: **Continue while it scans** moves on to the next step, and a scan that is interrupted picks up where it left off.
3. **Rigs.** The telescope-and-camera combinations the scan found. It works these out by itself from the files — see *Recognising your telescopes without being told* — so this step has something to show only once the scan has found light frames. What you add is what the headers cannot say: a name you recognise, whether the camera is a colour one (a mono and a colour version of the same camera often look identical in the headers), the folder names that belong to it, an owner if the data is a friend's, and details such as which filters it carries.
4. **Sites.** Where you observe from: a name, latitude, longitude and, if you like, elevation. The numbers are decimal degrees with north and east positive: Greenwich is `51.48, 0.00` and Denver about `39.7, -105.0`. If your capture software writes the site into its headers, the wizard fills them in for you. A telescope kept somewhere else, at a remote observatory say, can have a site of its own with its own time zone. This drives every calculation about when it gets dark, what is above the horizon, and how much the moon will hurt.
5. **Logs** (optional). Where your capture software's log files are: PHD2's for guiding and SGP's for the night's sequence. They live on the PC that runs the telescope, so the step walks through copying them over with Syncthing, and a **Check folder** button counts the logs it finds.
6. **Phone** (optional). Two tickboxes, **Allow phone access on this network** and **Send push notifications through ntfy**, each with the steps to set up your phone — see *Your phone*.

A few things are not in the wizard. Under **Settings › Observing sites** are your home site's time zone (blank means the same as this PC), how dark your sky is (the Bortle class, taken as 6, a bright suburban sky, until you set it) and the shape of your horizon. An all-sky camera is set under **Settings › Weather and sky**; see *The all-sky camera and finding cloud*.

The last screen's **Open AstroTracker** button closes the wizard. If the scan is still running, its progress moves to a banner across the top of every page, and you can use the program while it works.

What it costs, and the trial

AstroTracker is paid software with a 30-night trial. Nothing is crippled during it and nothing is asked of you to start.

The trial

Download it, run it, and you have **30 nights** of the whole program. Every feature, your whole library, no limits on how many telescopes or how many frames. There is no sign-up, no email address, and no payment details — the installer does not ask and there is nothing to create an account for.

It is counted in *nights* rather than clock time, and that is a deliberate detail. The trial begins at midday and ends at midday, because a night in this program runs from noon to noon, so it can never run out in the middle of a session. Installing at eleven at night does not cost you a night either: whatever time you start, you get all thirty. For the last week a banner across the top of every page counts the nights down and names the morning the trial ends.

> **WHY IT WORKS THIS WAY**
>
> AstroTracker is partly a safety monitor — it is what tells you at 2 a.m. that your telescope has stopped producing frames. Software like that must not choose the middle of the night to announce a commercial deadline. If the trial does run out while a rig is imaging, it keeps watching that session and stops in the morning instead.

When the trial ends

Nothing you have collected goes anywhere. Every project, every night, every grade and measurement stays exactly where it is and stays visible. Your photographs were never touched in the first place.

What stops is new work, which means anything that would add to or change what the program knows. In the background, folders are no longer watched or rescanned, new photographs are not measured, and the morning summary, the stall and rain alerts and the season reminders are no longer sent. On the pages, every change is turned down: grading a photograph, writing a note, setting a goal, changing a project's status, saving settings, starting a scan, moving files. Every page still opens and everything on it can still be read, and a banner across the top says the trial has ended, with an **Enter a key** button.

A few things carry on, because they look after what you already have rather than adding to it: the daily backup and **Back up now**, checking for and installing updates, and saving a problem report from Settings › Support. Settings › Licence stays open as well, since that is where the key goes. Entering a key restores the rest immediately, with nothing to re-scan and nothing lost, because nothing was ever taken away.

The key, and what it does not do

A licence key is a short piece of text that arrives by email. You paste it into Settings › Licence, press **Use this key**, and that is the whole procedure.

It is verified **on your own computer, offline**. There is no activation server, no licence server, no account to log into, and no periodic check-in. Nothing is transmitted when you enter a key, when it is checked, or ever afterwards. An observatory PC with no internet connection at all behaves exactly like one online.

> **IN PLAIN TERMS**
>
> Most paid software phones home to check you are still allowed to use it. This does not. The key proves itself by mathematics rather than by asking permission, so the program keeps working whether or not the internet does — and whether or not the company that sold it to you still exists.

The key covers your own machines, so the PC in the observatory and the one in the house are both fine. Moving to a new computer is a matter of pressing **Remove this key** on the old one and pasting the key into the new one.

Telling it where your files are

This is the one thing the program cannot work out for itself, and the one thing it asks you to get right.

You give AstroTracker a list of top-level folders — it calls them **library roots** — and for each one you say what kind of folder it is. You do this in the **Folders** step of the setup wizard the first time, and on the **Library** page whenever you like after that. Add the folder that holds your rig folders rather than each target folder inside it, since a root is read all the way down. The role matters, because it changes how often the folder is looked at and what the program expects to find inside. There are three to choose from.

ROLE	WHAT IT MEANS	HOW IT IS WATCHED
Current work	Where tonight's files land and where you process. Add one for every place new subs land: a folder per rig, such as `D:\Astro\MyScope` and `D:\Astro\SecondScope`, is typical, and several are normal.	Watched live. New files are noticed within seconds.
Archive	Finished data, usually on a NAS or large external drive. Nothing new appears here except when you move a completed project in.	Scanned on a schedule — every six hours by default — or on demand.
Calibration library	A dark and bias library shared across projects. No targets live here.	Scanned on the same schedule as an archive.

A root can be a local drive, an external disk, or a network path such as `\\NAS\Astro\Archive`. The distinction between current and archive is mostly about cost: watching a network archive continuously would hammer it for no benefit, since nothing changes there on its own.

Never one root inside another

What the program will not accept is a root inside another root, or one that contains a root you already have. A root reads everything beneath it, so the inner folder's files would belong to both, and one of the two scans could only fail. It refuses the folder and tells you why, in one of two ways:

- "that folder is inside `D:\Astro`, which is already a library root and scans everything under it" — there is nothing to do: that folder is already covered.
- "that folder contains the library root `D:\Astro\MyScope`; remove that one first" — remove the smaller root on the Library page, then add the larger one. Removing a root makes the program forget the grades you gave its subs, apart from `BAD_` names and `bad` folders, as *Library* explains.

A rig's own calibration folder, set on the Rigs page, follows the same rule; see *Calibration: darks, bias and flats*.

The rejected-subs folder

The Library list can show one more role, `rejected`, which is not among the three you pick from. That root is made for you when you fill in **Folder for rejected subs** under Settings › Rejected subs: the folder the subs you grade bad are moved to (see *Review — judging your photographs*). The folder is created if it does not exist and follows the setting: change the folder and the root moves with it, clear it and the root goes. Everything in it counts as bad, and adds nothing to your totals. Choose a folder outside your other roots: one inside them would be a root inside a root, and its scan would be refused.

Stage folders

Within a working drive, most people separate new data from data they are working on. The default names are `New Data` and `Processing`, and moving a project's folder from one to the other is how you signal that you have stopped shooting and started editing. AstroTracker notices the move and updates the project's status accordingly. If your drives use different names — "Incoming", "Stacking" — you give it yours under Settings › Library layout, separated by commas if there are several.

What it expects to find, and what it ignores

Beyond the roots and stage names, the program does not require any particular folder layout. It reads headers, so a folder can be called anything. But two things are worth knowing.

It **ignores the output folders of stacking software** — the `calibrated`, `registered`, `master` and `logs` folders that WBPP and similar tools generate. Those contain processed copies of frames you already have. Counting them would double or triple your integration totals, which is worse than useless.

It notices **files whose names begin with** `BAD_` and treats that as a human judgement to be preserved rather than a naming accident. If you have been marking bad frames that way for years, the program keeps that opinion. A sub sitting in a folder called `bad`, `rejected`, `reject`, `quarantine` or `trash` is read the same way.

The scan: headers, not pictures

What actually happens when AstroTracker walks your drives.

For every FITS file it finds, the program opens just far enough to read the header block and then closes it again. A header is a few kilobytes of plain text at the front of a file that may be fifty megabytes. Reading the headers of 195,000 files is a manageable job; reading their contents would be several terabytes of disk traffic.

From that header it takes, where present:

WHAT	USED FOR
Target name	Grouping subs into projects.
Frame type	Light, dark, flat or bias.
Filter	Per-filter totals, goals, moonlight advice, flat matching.
Exposure length	Integration totals and dark matching.
Date and time	Which night a sub belongs to; ordering; the live view.
Sensor temperature	Dark matching, by the cooler's set point where the header records one.
Gain and offset	Dark and bias validity; drift detection.
Focal length, pixel size, image dimensions, binning	Identifying which telescope took it — see the next chapter.
Coordinates and rotation	Matching subs to targets by where they point; framing; flat rotation matching.

Which night does a sub belong to?

An observing night crosses midnight, so a frame taken at 1 a.m. on the 14th belongs to the night of the 13th. AstroTracker dates every night by its evening, which is how astrophotographers already speak, and keeps the frame's own clock time separately. A single sub, though, is labelled with the calendar date and time it was taken: one from seven minutes past midnight reads `2026-07-14 00:07`, and hovering over it shows the night it belongs to, "Night of 2026-07-13".

Times are handled in local clock time throughout, because that is how you experience a night. Sequence Generator Pro writes both the local time and a UTC timestamp into every header, and the program reads the local one; where a capture program recorded only UTC, an approximate local time is worked out from the longitude in the header. A telescope in another time zone can be given that zone, and its subs, its logs and its darkness times are then all read on that clock rather than yours.

After the first scan

The first scan is the only slow one. Afterwards the program keeps track of what it has already seen, so a re-scan only looks at what changed. Working drives are watched continuously: a file appearing during the night is picked up within seconds and appears in the live view without you having to press **Scan**; *The watcher* describes how.

Files that turn out to be unreadable — a truncated download, a drive that dropped out mid-write — are counted and reported rather than silently skipped, because a file that cannot be read is usually a file that needs attention.

Recognising your telescopes without being told

A surprisingly hard problem with a clean solution, and one of the parts of the program most worth understanding.

If you own three telescopes, every question about your library depends on knowing which one took each photograph. Six hours on a galaxy through a big scope is a different thing from six hours through a small one, and darks from one camera are useless for another.

The obvious approach is to read it from the header, which has fields named `TELESCOP` and `INSTRUME` for exactly this. Unfortunately those fields are, in practice, worthless. Sequence Generator Pro writes the *mount driver's* name into `TELESCOP` — so every rig on the same mount brand claims to be the same telescope — and `INSTRUME` frequently reads `ASI Camera (1)`, which identifies nothing at all.

> **WHY IT WORKS THIS WAY**
>
> Rather than asking you to fix your headers, AstroTracker identifies a rig by a **fingerprint** — a combination of things that *are* reliably recorded and that together are unique to one optical setup:
>
> **focal length + pixel size + image dimensions + binning**
>
> Two different telescopes essentially never share all four. The same telescope, shot the same way, always does. Binning, the one of the four that people do change between nights, is allowed for, as the next section explains. It needs no cooperation from your capture software and works retroactively on a decade of files.

So the program discovers your telescopes rather than being told about them. After the first scan, the Rigs page lists what it found. You give each one a name — "80 mm refractor", "8-inch Newtonian" — and can add the things a header cannot know: which filters it carries, where its log files are, its guiding limit.

One rig, more than one fingerprint

Strictly, the fingerprint changes when you change binning. Switch the camera to bin 2 for a night (four pixels read out as one, so the picture is half as wide and half as tall and each pixel twice the size) and those frames carry a different fingerprint from the same telescope's bin 1 frames. Taken literally, that would give you a second rig for one night's work.

So before it creates a rig, the scanner checks whether the newcomer is one you already have: a focal length within 2 per cent of an existing rig's, and the same sensor once the binning is undone. If so, the frames join that rig and it simply gains a second fingerprint. The same check absorbs a focal length written slightly differently in the header, as happens when a capture profile is edited. The Rigs page

lists the fingerprints each rig has collected and the binnings it has been seen at; its **Binning** setting is just the one you normally shoot at, and frames at another binning still belong to the rig.

What the check cannot explain, such as a few frames shot with the camera set to a cropped readout by mistake, still becomes a rig of its own. But a new rig holding fewer than ten lights is flagged for review until you confirm it; *Alerts* explains how to fold it into the rig it resembles or drop it.

Cameras are identified separately

Calibration frames create a complication. A dark frame has no focal length in its header, because the lens cap was on and often no telescope driver was even connected. So darks cannot carry the full fingerprint.

The solution follows the physics. A dark frame is a property of the *camera body*, not of the telescope — one camera used on three telescopes needs only one set of darks. So AstroTracker matches calibration frames on a shorter **camera fingerprint**: pixel size, image dimensions and binning, plus the folder they were found in. Darks are then organised by camera body rather than by rig, and a light frame matches the darks of its rig's camera.

Two identical camera models are told apart where it matters — by which telescope they were shot on, by a folder named after the rig, or by you when it genuinely cannot be determined. This also keeps a monochrome and a colour version of the same sensor from being confused.

> **A DELIBERATE RESTRICTION**
>
> Only light frames can create a rig. A folder of darks discovered before any lights must not be allowed to invent a telescope that does not exist; it attaches to a camera instead. This sounds like a detail and is the difference between a clean rig list and a list full of phantoms.

Rigs you do not own

A rig can be marked as belonging to someone else — data from a friend, or from a collaboration. Those rigs are excluded from calibration alerts and from planning, because nagging you about missing flats for somebody else's telescope is noise. A rig can also be marked **retired**: it sinks to the bottom of every list, raises no alerts and sits out tonight's planning, while its projects and data remain exactly as they were.

A friend's setup can be identical to yours, the same telescope model with the same camera, and then it shares your fingerprint. There the folder decides. **Split by folder…** on the Rigs page gives the frames under a folder name you choose (the folder their data arrives in, say) a rig of their own with your friend as owner, and later frames in that folder follow it. Frames with no such folder in their path always stay with your own rig.

Turning photographs into projects

A **project** is one target shot on one rig. It is the unit that gets planned, tracked and finished, and everything on the Active page is organised around it.

The same nebula shot on two different telescopes is two projects, because they are genuinely two pieces of work with their own hours, their own calibration and their own framing. (If you want them viewed together, there is a campaign view for exactly that, described in *Finished, and the campaign view*.)

How a sub finds its project

Grouping happens one rig at a time: a sub is only ever compared with projects on the rig that took it. (Lights shot with no telescope driver connected have no focal length in their headers. When the camera alone does not settle which rig took them, the folder they are filed under becomes a rig of its own, named after that folder, so they do not sit unassigned and invisible.) Among that rig's projects, three rules decide, in this order:

1. **The same folder, under the same name.** A sub joins a project that already has subs in its target folder, if its target name matches, wherever the telescope was pointing. Night, filter and flats subfolders count as part of the target folder. This keeps a mosaic's panels, or a target you re-centred halfway through, together as one project when they share a folder.
2. **Where the telescope was pointing.** Otherwise the coordinates in the header decide, whatever the target name says. A sub joins the nearest project within 35 per cent of the narrower side of the rig's field of view (about a third of a degree for a field 1.4 by 0.9 degrees; half a degree when the field is not known), or up to three times as far when the target name matches as well. This is what catches a name left blank, misspelt or still set to last night's target, and the same object written two ways: M31 one night, Andromeda the next.
3. **The name, when there is no pointing.** Subs whose headers carry no coordinates, usually because no mount was connected, are matched by target name alone.

Names are compared ignoring case, spaces and punctuation, so `NGC 7000`, `ngc7000` and `NGC_7000` are one name. Where the header has no name, or only a placeholder such as `Target`, the name of the target folder stands in. A sub that none of the rules place starts a new project, named after its target; the project page's **Rename** changes that. A project you created yourself before shooting takes part too, provided you gave it a rig and the program knew where the target is, so the first night's subs find it by pointing.

> **WHY IT WORKS THIS WAY**
>
> The target name is typed by a person, or carried over from last night's sequence because nobody changed it. The coordinates come from the mount. Two subs aimed at the same patch of sky through the same telescope are the same piece of work whatever they are called, so the name decides only within one folder, or when there is nothing better to go on.

The last two rules only ever lead to a project whose status fits where the sub lives: subs in a Current work folder never join a Complete or Cancelled project, and subs in an Archive join only Complete ones, so an old archive is never pulled into this season's work. A target you return to after finishing its project therefore starts a new project, unless the new subs are saved, under the same name, into the very folder where the finished project's own subs still sit. Archived subs that match nothing finished start a new project that is already Complete.

When it groups something wrongly

The program never stops to ask: every sub ends up in a project, an existing one or a new one. When the result is not what you meant (a mosaic's panels kept in separate folders under different names, say, split into projects of their own), you put it right on the project page: **Move to…** for chosen subs, **Merge into…** for a whole project (see *Inside one project*). Subs sent with **Move to…** are marked as placed by you: no rescan moves them again, and neither does **Regroup projects** on the Library page, which rebuilds all the grouping from the files. A merge holds through rescans, but merged subs are not marked, so **Regroup projects** can split them apart again; a merge also deletes the emptied project's goals and cannot be undone in the app. For a correction that must survive everything, use **Move to…**.

The two-targets-in-one-folder problem

Sequence Generator Pro creates one folder per sequence. If you shoot a second target the same night, its frames land in the first target's folder — and now the folder name is a lie.

AstroTracker catches this because it reads headers rather than folder names. The morning summary reports how many frames of which project are sitting in another target's folder, and the Organize page offers each of them a folder of its own, created beside the one they landed in and numbered by that project's own nights. Nothing moves until you confirm.

Flats are handled carefully here: they follow only when something identifies which target they belong to — the target name in their header, or their rotation angle if the rig is set to require matching rotation. Most flats are simply called "flats" and were shot once for the whole night, in which case they stay where they are and legitimately serve both projects.

The same care applies before anything moves. Flats, darks and bias count towards a project whose lights share their folder. Where two projects share it, a frame whose target name picks one out, such as `flats_M101`, counts for that one; everything else counts for the folder's owner: the project the folder is

named after, or failing that, the one whose subs were there first. A frame you placed by hand stays where you put it.

The life of a project

Projects move through a small set of states. The important design rule is that **status only ever moves forward**: where its files sit can move a project on to Processing or Complete, never back, and a project you mark Complete or Cancelled is left as it is.

STATUS	MEANING	HOW IT IS REACHED
Planning	Intended, nothing shot yet.	Created by you.
Acquiring	Actively collecting data.	Automatically, when light frames start a new project. A project you planned yourself is not moved on when its first subs arrive: change its status yourself.
Processing	Done shooting, now editing.	Automatically once all its light frames are in a Processing stage folder.
Complete	Finished.	By you, or automatically once all its light frames are in the archive.
Postponed	Set aside, will resume.	By you, with a reason.
Cancelled	Abandoned.	By you.

Postponed, and the season reminder

Most projects are not abandoned; they are interrupted. The target sinks below the trees in October and does not come back until March, or the weather simply never cooperated. Marking a project **Postponed** with a reason — out of season, weather, gear, other — takes it off the active list without losing anything: the files stay where they are, the grades and totals are kept, and it sits in its own group with an estimate of when the target returns, worked out from its rig's site and your own horizon profile.

Then the useful part. Every day AstroTracker works out how long each postponed target will be up in the coming night's astronomical darkness at its rig's site, "up" meaning what it means on the tonight chart: above your altitude limit or the horizon you drew, whichever is higher. The first night the target is up for at least half the dark hours, AstroTracker tells you — in the app and on your phone — "You can continue M106", after three in the afternoon so there is time to plan the night.

It tells you once per season, and re-arms after the target has been out of range for a month. It is a season reminder, not a weather one: if the target already clears that bar on the day you postpone the project, the reminder waits until it has gone and come back. The switch is under Settings › Weather

and sky. If new light frames for a postponed project simply appear, it puts itself back to Acquiring without being asked.

> **IN PLAIN TERMS**
>
> You are not expected to remember that the thing you gave up on in October is worth restarting in March. The program remembers, and tells you on the first night it is actually worth doing.

Finishing, and archiving

Marking a project Complete moves no files — that is deliberate, because "I have finished shooting this" and "I have tidied it away" are different decisions taken at different times.

When you are ready to tidy, **Archive finished projects** on the Organize page moves a completed project's lights and flats into your archive, night folders and all, as one logged, undoable batch (see *Organize — moving files safely*).

Alternatively, move the folder to your archive by hand. The project follows on its own. Its notes, goals, every sub's grade and reasons, and all the quality measurements carry over to the new location, and the status moves to Complete at the archive's next scan. Nothing is re-measured and nothing is lost.

Moving files by hand

Files get moved around in Explorer all the time: a night dragged into the wrong folder and back, a finished project sent to the archive, an old one brought back to a working drive. AstroTracker follows the exposure, not the path it happens to sit at. When a sub disappears from one folder and the same exposure turns up in another, within one drive or between a working drive and the archive, it keeps its grade, its reject reasons and its measurements. Moved within one drive or into the archive, it keeps its project too; brought back to a working drive, it is grouped afresh (see *Bringing old nights back into a working project* below). To give a sub a different project, use **Move to…** on the project page.

It knows the exposure by its header — the camera, the moment the shutter opened, the exposure length and the frame type — rather than by its file name. When two copies of the same sub move together, say a night folder and a copy of it, each copy keeps its own grade. A sub whose file has vanished is remembered for 120 days; if the file turns up anywhere in your library within that time, it picks up where it left off.

> **WHY IT WORKS THIS WAY**
>
> A grade belongs to the exposure, not to the file name. That is why a sub you rename out of its `BAD_` prefix stays rejected: removing the prefix is not read as a change of mind. To take a sub back, grade it again on the project's Review page (see *Review — judging your photographs*).

Bringing old nights back into a working project

Sometimes a target you finished and archived deserves more: a second season to go deeper, or a filter you did not have the first time. The nights you already shot belong with the new ones. Move (or copy) the old night folders in beside the new ones, then press **Rescan this project's folders** on the project page (see *Inside one project*) rather than wait for the next full scan.

The old subs appear in the project as soon as that rescan finishes, ungraded at first. They get their grades, reasons and measurements back once the archive notices they have left: at its next scan, every six hours by default, or sooner if you press **Scan** on the archive under Library. A grade you give one of them in the meantime is kept. If you copied rather than moved, the originals are still in the archive and Organize lists each pair under Duplicates; the copies take over the grades only once you delete the originals.

> **A KNOWN LIMIT**
>
> Because the working drive is read before the archive notices the files have gone, the moved subs can land in a second project of their own instead of the one you meant. If that happens, open one of the two and use **Merge into…** on the project page to put them back together.

Active

The landing page: everything you are working on, telescope by telescope.

The page is called **Active** in the menu because it holds only work in progress: projects being planned, shot or processed. Postponed projects are shown apart, and finished or abandoned ones have their own page (*Finished, and the campaign view*).

Along the top are a few totals: how many projects are active, the integration in the whole library, how many rigs and how many files the program knows, and, when there are any, an **Unassigned** count of files that belong to no project. The tile calls them calibration frames, and most are: darks and bias kept in a calibration library belong to a camera rather than a project, so they always count here. If the latest night produced light frames, a **Last night** card follows with what each project gained and buttons to deal with the flagged subs (*The morning summary* describes it). A **+ Plan a target** link leads to the Planning page, which is where tonight's chart, the recommendations and the catalogue checklists are.

Then comes a card for each telescope with work in progress: its name, the telescope and camera on it, the hours and light frames its active projects add up to, and a row of horizontal bars — one per filter — showing how those hours are distributed. Below the bars are the names of its biggest projects, and any postponed ones with when their target is back in the sky. A friend's telescope carries its owner's name, and retired telescopes come last, under their own heading.

Click a telescope and you get its projects, one card each. Tabs along the top narrow the list: **Active** to begin with, one tab for each status the telescope has projects in, and **All**. Every card shows the target, its first and last night, its status, its hours, nights and light frames, and the same bars for that project alone.

Those bars are the quickest read on the page. If you have set goals per filter, each bar is scaled to its own goal with a marker at 100%, so a glance tells you that red is finished and blue is halfway; a total goal adds one bar of its own, showing how much of it is done. Otherwise the bars scale to the largest filter, which still shows imbalance — a common and easily missed problem, since colour data shot in unequal amounts limits the whole image to whichever channel is weakest.

The badges on a project card

Along the bottom of each project card, small badges say what needs attention and what tonight is good for:

BADGE	MEANING
Flats paired	Every filter on every night has flats to go with it. Otherwise the badge counts the filter-nights still missing them.
… without darks	How many of the project's exposure settings no dark run covers (*Calibration: darks, bias and flats*).
… marked bad	How many subs are graded bad.
Revisit	You have marked the project to come back to, with the filters you want to add.
tonight: …	Which filters suit the target tonight, given the moon: **tonight: Ha, SII**, **tonight: OIII (ok)**, **tonight: any filter** or **tonight: wait for dark**. It appears only when the target is up tonight; hold the mouse over it for the full verdict, explained under *Best filters tonight* below.
Back …	On a postponed project, roughly when its target returns ("Back mid March"), or **In season now**, with the reason it was postponed.

Goals

You can set a total goal for a project, or a goal per filter — six hours of luminance, two each of red, green and blue. Set them with **Set goals** on the project's own page, or give a total when you plan a target; per-filter goals, when there are any, take the place of the total. Goals drive the bars, on the cards and on the project's page, and the **Last call** panel on the Planning page, which forecasts whether each project will reach its goal before its target leaves for the season (*Planning*); with per-filter goals it adds up what each filter still lacks, so plenty of luminance cannot hide missing red. There is no alert when a goal is reached: a filter that has met its goal shows the goal figure beside its bar in green, and **Last call** reads **Goal reached**.

> **WHY IT WORKS THIS WAY**
> A goal is a number of hours you chose, not a verdict on the picture. Whether a filter is clean enough yet is what **How much more?** on the project's page measures, and it ignores your goals (*Inside one project*).

Tonight's chart, on the Planning page

The chart of the coming night is not on this page but on Planning, in the panel headed **Tonight**, and it is the part of the program most used at six in the evening. It covers every project being planned, shot or postponed whose place in the sky is known, except those on retired telescopes.

The horizontal axis is the night, from evening through to dawn. The dark window — true astronomical darkness, not merely sunset — is shaded. Each target you are working on gets an altitude curve showing how high it climbs and when, drawn solid only where it is genuinely usable: in the dark, above

your minimum altitude or your local horizon, whichever is higher. The moon is drawn as a dashed line so you can see at once whether your target is up when the moon is not.

Each target has a checkbox to show or hide its curve, with **Show all** and **Hide all**, and the choice is remembered in your browser. With a dozen projects the chart is unreadable until you can turn some off.

Your actual horizon

Nobody images from an open plain. There are trees, a neighbour's roof, a hill. AstroTracker lets you describe your real horizon as a list of compass bearing and altitude pairs — `0:20, 90:35, 180:15, 270:25` means twenty degrees of trees to the north, thirty-five degrees of something to the east, and so on — or a single number for a flat horizon.

That profile, typed under Settings › Observing sites, is then used everywhere it matters: the tonight chart, everything else on the Planning page that asks when a target is up, and the season reminder for postponed projects. A target that technically rises at nine but does not clear your oak tree until eleven is shown as rising at eleven, because that is the truth for you.

Best filters tonight

The tonight chart goes further and ranks the *pairs* — which project, through which filter, gives the most useful data tonight given where the moon is and how bright it is: the **Best use of tonight** line under the chart. You do not have to tell it which filters each rig carries; it uses the filters it has seen in that rig's light frames. Each target's row on the chart, and the **Best tonight** tile on a project's own page, might read:

```
best Ha, SII · ok OIII 48% · wait L, R, G, B
```

Meaning: the narrowband filters will barely notice tonight's moon, oxygen will suffer moderately, and there is no point shooting broadband colour at all. On the project cards the same verdict is shortened to the **tonight** badge. *The moon, and which filters survive it* explains how that is computed and why it is more trustworthy than a rule of thumb.

Inside one project

Open a project and you get everything known about that target on that telescope.

Nights

A list of every night the project was shot, with how many subs came from each, through which filters, and how long they ran. This is the record that is otherwise impossible to reconstruct: it is how you discover that the four hours you thought you had on a target were actually ninety minutes across three frustrated attempts.

Calibration coverage

The **Sub-exposures and darks** panel takes each group of subs (exposure, gain, offset and temperature) and names the dark run that would be used — the nearest in time from the right camera body — while the **Nights** table shows the flats for each filter on each night. Where something is missing or questionable it says so specifically rather than generically: not "check your calibration" but **Ha: none** on a night with no hydrogen-alpha flats, **none** where no darks match, **unverified** where the darks' headers carry no offset, or **body?** where a dark run's camera body has not been settled.

Flats that came from the rig's own calibration folder rather than from the project's night folder are marked **set**, and holding the pointer over one names the set, so you can see at a glance whether a project is relying on a shared flat set from three weeks ago.

How much more?

The most-used part of the page. It answers: given what I already have, and how good it was, is another session worth it?

Hours alone cannot answer that, so the program measures the pictures instead of counting the clock. Press **Work it out** (the button reads **Measure again** once there is a result, and **Measuring…** while it runs). It calibrates a few of your own subs of each filter with the project's darks and flats, and works out the **signal-to-noise ratio**, or **SNR**, that a stack of all your subs would have: how clearly the faint outskirts of the target, always the last parts to come clean, stand out from the grain. Nothing is stacked, but real files are read, so it takes a little while, typically about half a minute a filter.

The answer is a table with one row per filter:

COLUMN	WHAT IT TELLS YOU
Hours	The integration you have in that filter.
SNR	How clearly the faint structure would stand out in a stack of all of it.
Verdict	**needs more**, **almost there** or **clean**. Out of the box, an SNR below 4 needs more, 4 to 6 is almost there, and 6 and up is clean.
Next band	The total hours at which the filter moves up to the next verdict, if more of the same subs come in.
Clean	The total hours at which it reaches clean, or **reached** if it already has.

Both hour figures are totals, not extra hours on top. A **>** in front of one means more than four times the hours you have now, beyond the point where the forecast is worth much. A filter whose files could not be reached, on an archive drive that is not plugged in for instance, shows **not measured**.

The bands, and the patch of pixels the noise is judged over (3 × 3 by default), are yours to set under **How much more?** on the Settings page, because what counts as clean is one imager's eye rather than a textbook figure.

The notes behind the i

Where the measurement noticed something about a filter, an **i** at the end of its row opens a short note: that the filter was measured without darks or flats, so its number is rough; whether your subs are long enough (**sky-limited: longer subs would gain nothing**, or **read noise matters: longer subs would help**); or how much fainter the target is in that filter than in another, and how many times the hours that costs.

Exposure groups, and where it is judged

If you shot one filter at more than one sub length (five-minute subs one season, ten-minute the next), **Show exposure groups** lists each on its own row, with its subs and hours, the SNR it alone would give, and whether it was calibrated with darks and flats.

Once there is a result, **Where to judge it** shows one of the project's subs: drag a box over the part of the target you judge the picture by, press **Use this box**, and measure again. Without a box the whole frame is judged, which is right for most targets; the box is for the exception, such as a galaxy whose faint halo you are happy to leave noisier than its bright body. **Back to the whole frame** removes it.

Subs you have rejected are left out. The result is kept on the project with the date and time it was measured, so it is there the next time you open the page; press **Measure again** after a few more nights. *Signal, noise, and when a filter is done* goes through the arithmetic.

> **IN PLAIN TERMS**
>
> Hours are not the finish line; a clean picture is. The program looks at the faintest parts of your target in a handful of your own subs and tells you, filter by filter, whether it is done, nearly done or needs more, and roughly how many hours in total would get it there.

Framing

Framing preview, in the **Rig fit** panel, shows how the target sits in your camera's field of view. The program fetches a real photograph of that patch of sky from a survey archive and draws your sensor's rectangle over it. The rectangle stays fixed and horizontal while the sky rotates underneath, so you can set a camera angle and see exactly what will fit — and drag the sky to re-centre the frame. **Save framing** keeps the angle and centre. When other rigs have shot the same target, the panel lists them with a **View campaign** link (*Finished, and the campaign view*).

The preview answers the question people otherwise answer by shooting a test frame and looking: will the whole of the Veil Nebula fit, or do I need to rotate the camera, or is this the wrong telescope for this object entirely? The panel's first line gives the answer in words, such as **comfortable**, **tight fit** or **mosaic, about 4 panels**, with the share of the frame's short side the target spans, from its size in the catalogue or one you enter under **Target size, arcmin**. Until you save a framing, the preview opens at the angle your frames were actually shot at.

Sub review, and moving bad frames

The **Sub quality** tile at the top of the page shows how many of the project's lights have been measured, how many were flagged and how many are still waiting; clicking it opens the sub-by-sub grading screen described in *Review — judging your photographs*. Once there are rejected subs still sitting among the good ones, a button such as **Move 12 bad frames** appears in the page header. It asks first, then moves them out of the way in one batch you can undo from the Organize page, so that when you stack, only the good frames are there.

Folders, and nights copied in by hand

The **Folders** panel lists every folder the project's files live in, with how many files each holds and an **Open** button that opens it in Explorer on the computer running AstroTracker. (**Open containing folder** in the header does the same for the folder holding the most files.) It is a quick way to see that a project is spread across a working drive and an archive, or that one night landed somewhere unexpected.

Rescan this project's folders reads just this project's folders again, Night subfolders included, without scanning the whole drive: for nights you copied or moved in by hand (*The life of a project* shows how to bring an earlier season back). The page refreshes when the scan finishes; if another scan is already running, it says so and you try again once it is done.

Renaming, merging, and moving files between projects

Grouping is automatic, but you have the last word. **Rename** in the header, or a double-click on the project's name, changes what it is called. **Merge into…** folds this project into another you pick from the list: every file moves across and this one is gone. It asks first, and it cannot be undone from the page.

The **Files** list at the bottom shows each file with its type, filter, exposure, night, time (in UTC), gain, temperature, rotation and grade, and a **Show** button that opens its folder with the file selected. Tick files (the box at the top ticks them all) and **Move to…** sends them to another project, or to a **New project…** on the same rig. Nothing moves on disk, and files placed this way stay where you put them when the folders are scanned again (*Turning photographs into projects*).

Delete removes the project but never its subs: they stay on disk and in the index, and the next scan groups them into projects again. What goes with it is what belonged only to the project, such as its goals and the copy of the finished picture it kept for you (your own image file is not touched). A planning project with no files simply goes.

Exports

AstroBin CSV in the page header writes the project's acquisition details in the layout AstroBin — the astrophotography sharing site — imports when you upload an image: exposure counts per filter, dates and calibration frame counts. That is tedious to assemble by hand and the program already knows all of it. **Summary** gives a plain-text version for a caption or a forum post, and a CSV of every project is on the Organize page.

Imaging — the live view

What is happening right now. This is the page that gets opened on a phone at two in the morning.

Each telescope that is running gets a card. The card is built from files arriving on disk and from the log files your capture software writes anyway, so it needs no special connection to anything.

Above the cards sits the **night arc**, the noon-to-noon bar described under *Weather*, with a line for now, so you know how much usable night is left; with telescopes at more than one site, each site gets its own.

What a rig card shows

- **Tonight** — the hours and number of subs since local noon, and how many came through each filter.
- **Time since the last sub.** The single most informative number on the page. If your exposures are five minutes and it has been twenty-two minutes, something is wrong, and the card says so: once the next sub is later than the margin set in Settings, the time turns into **STALLED** and a note says how many minutes late the next sub is.
- **The latest sub's measurements** — its filter and exposure, star size (FWHM), shape (eccentricity, shown as *e*) and the number of stars found, then any warning flags it earned, or "looks fine" — so you can see focus drifting or cloud arriving without opening anything. The flags are explained in *Review — judging your photographs*.
- **Trends across the night** for sky brightness and star size, as two small lines that appear once three subs have been measured, so a gradual decline is visible as a decline rather than as one bad frame. When the stars late in the night are more than a quarter larger than they were early on, the star-size line turns amber and the card warns "focus drifting: stars are getting bigger".
- **The seeing estimate**, in arc-seconds with a plain-language band from excellent to bad.
- **Cloud over the target.** With an all-sky camera at your home site, a warning such as **clouds over M31: 45%** appears once the camera sees 30 percent cover or more in the direction of the rig's target.
- **A guide graph** in the style of PHD2's own: the last ten minutes of guiding error, RA in blue and declination in red, with your limit marked as a dashed line. A rig guiding badly is obvious at a glance.
- **The latest picture** from that rig, with buttons to grade it right there without opening the project: **clouds**, **trailing**, **focus**, **airplane / satellite** (for a streak one left across the frame) and **other**

each reject it for that one reason, replacing any it had; **good** passes it and **clear** removes the grade. Several reasons at once are given on the Review page.

Tonight's sequence as a chain

The card also draws the night as a chain of steps: sequence started, target acquired, centred, guiding, focused, imaging, ended. The step the rig is currently on is highlighted with its time. Steps where something went wrong turn amber or red, and underneath is a count of the night's incidents — guide star lost, recovery attempts, meridian flips, warnings.

This is assembled entirely from the SGP and PHD2 log files, and it turns a wall of log text into something you can read in one second from bed. At the foot of the card, **Show session log** lists the night's latest events. **Clear errors** tidies away the warnings once you have dealt with them and **Clear** hides the lot; anything new still appears, and past sessions and the morning summary keep every event.

> **WHY IT WORKS THIS WAY**
>
> The log also reveals something no amount of file-watching can: whether a sequence *finished normally*. Without it, a rig that completed its plan at 4 a.m. and a rig that crashed at 4 a.m. look identical — both simply stop producing files — and you would get a "no new subs" alarm either way. With the log, the card says "sequence finished" and no alert is sent.

Past sessions

A dropdown at the top opens any previous night, with its sub charts and guiding charts intact. This is where you go when an image did not come out and you want to know why: you can see that the guiding degraded at 1 a.m., or that the star sizes grew steadily from midnight, or that half the frames were shot through cloud.

Live hardware, and control

Optionally — and only for a rig on the same network — the card can show what the hardware is doing this second. Sequence Generator Pro runs a small service on the observatory PC; enter that PC's address as the rig's **SGP API host** on the Rigs page and the card gains a **Rig now** panel showing the camera's state and temperature, which filter is in the beam, where the focuser is, and whether the mount is parked, slewing or tracking. It recognises an autofocus run while it is happening. A rig with an address keeps its card all day, so the panel is there before the night's first sub.

Filters are shown by name rather than slot number. The service will only ever report "slot 2", so the names are read out of the rig's own log file — which means they also work for a telescope whose logs you sync but whose PC you cannot reach. Failing that, type them into the rig's **Filter wheel** box on the Rigs page.

Beyond watching, you can optionally allow **control**. Switched on per rig with **Allow control** on the Rigs page, and off by default, it adds a **Rig actions** menu that can park and unpark the mount, abort the frame being exposed, set the cooler, move the filter wheel, move the focuser to a position or step it in or out, connect or disconnect the equipment, start a sequence, and set the meridian flip delay (how many minutes past the meridian the mount waits before it flips). Every action except the flip delay asks for confirmation naming exactly what will be sent, and every command, accepted or refused, is written into that rig's session log. It works from a paired phone, which is the entire point of it — "it is raining and I am in bed" is the case it exists for.

> **READ THIS BEFORE ENABLING CONTROL**
>
> Sequence Generator Pro's service has **no password of any kind**. Anyone who can reach that port on your network can park your mount. Windows blocks the port until you allow it, and the guide inside the program gives you a firewall rule that restricts it to the one PC running AstroTracker. Leave control switched off and the program is strictly read-only.
>
> One thing the service cannot do at any price is *stop a sequence* — there is only a start command. Aborting a frame stops that frame and SGP begins the next one. If clouds arrive, aborting and parking is the closest you can get remotely.
>
> Nor can it switch the camera's cooler on or off. Setting the cooler only changes its target temperature, so if the cooler is off, SGP accepts the number and nothing happens; the panel's camera line says whether the camera is actually cooling.

Weather

Everything about the sky itself, for every site you observe from.

The forecast

A week of hourly cloud cover, from Open-Meteo — a free service needing no account or key — under a line giving tonight's verdict and its best clear stretch. It scrolls sideways, each night's dark hours are outlined so you can see at a glance which nights are actually usable, likely rain is marked with a blue dot, and a row of chips, one per night with its clear hours, jumps to any night. Planning Thursday from Monday morning works.

The rain alert

While a telescope is imaging, or while the site is in darkness, the program watches the same forecast and warns you once — tray, in the app, and on your phone — when rain becomes likely within your chosen lead time. You set both the probability threshold and how much warning you want, under Settings › Weather and sky; out of the box it warns at a 50 percent chance of rain within three hours.

The Weather page states plainly what the alert is currently watching for and names the first rainy hour within your lead time, or says there is none, so you are never guessing whether it is armed.

> **AN HONEST LIMITATION**
> This is a forecast, not radar. It will not catch a shower that was not predicted. It is there to stop you leaving a telescope out overnight when rain was always coming, not to react to weather already overhead.

The night arc, and multiple sites

The page draws the night itself as a bar from noon to noon: sunset, the start of astronomical darkness, dawn (the moment that darkness ends) and sunrise, with the stages of twilight in between shaded and a line marking the present moment. The moon is not drawn here; its phase, rise and set are on the Planning page. If you observe from more than one place — a home site and a dark-sky site, or a remote telescope in another state — each is a site of its own with its own coordinates, horizon, Bortle class and time zone, and everything is computed per site. On the Weather page each site gets a panel of its own with its own arc and forecast, noting whether it is dark there now and how far its clock is from yours.

The all-sky camera

If you have an all-sky camera and have told the program where its pictures land, this page shows its newest picture too, with how old it is, the cloud reading taken from it through the night, a chart of cloud cover you can click to bring back the picture from any moment, a button that plays the night as a

timelapse, and **Notify me when it clears** for a night you are waiting on. *The all-sky camera and finding cloud* explains how all of it works. Without a camera set up, this part of the page simply does not appear.

Seeing

The current seeing estimate, in arc-seconds and as a word, derived from the stars in the subs landing right now. *Seeing, guessed from your own stars* explains how. It also appears on the Imaging page, on each rig card, in past sessions and in the morning summary, because "how good was the sky" is the question you ask about a night afterwards as often as during.

Planning

What to point at tonight, what to point at next, and whether the targets you are already working on will be finished before they leave the sky for the year.

The Planning page sits between Weather and Active. It is made of panels, and **Arrange** puts them in the order you want and hides the ones you do not use. The layout is kept in the program, so the PC and the phone show the same page. Above them, **Find a target** opens any catalogue object's details by name, and **+ Plan a target** creates a planning project from a name it looks up, an optional size and goal, the rig, and a framing preview.

Catalogues, downloaded once

Planning works from the standard lists of deep-sky objects: OpenNGC (the NGC and IC catalogues and Messier's list) and the Sharpless, LBN, van den Bergh and Barnard catalogues from the astronomical data centre in Strasbourg. The first time, press **Download catalogues**; it takes under a minute and never needs repeating. They are fetched onto your computer rather than shipped inside the program, because their publishers keep the say over redistributing them.

The night, and the moon calendar

How many hours of real darkness tonight has, where the moon is, and what that means: a night for broadband (colour, galaxies), for narrowband (emission nebulae through hydrogen-alpha, oxygen and sulphur filters), or for a mixture. Below it are five weeks of nights, each showing how much of its darkness is free of the moon. Pick any of them, or any date in the **Night** box at the top of the page, and the rest of the page plans that night instead, except **Tonight's chart**, which always shows tonight, and **Last call**, which weighs the whole season; **Back to tonight** returns.

Will you finish in time?

Every target has a season. For each project on your own telescopes that you are shooting or planning, the **Last call** panel estimates the hours you are likely to add before the target leaves for the year and, if the project has a goal, whether that reaches it: **On track**, **Tight** or **Unlikely this season** (without a goal it reads **No goal set**). The estimate comes from your own record rather than from a perfect sky: how long the target is up each night, how often that telescope has actually imaged in that month before, how much of the time a target was up you have turned into subs, and how many projects share the telescope. Postponed and out-of-season projects follow under **Coming back**, with when their target returns.

> **IN PLAIN TERMS**
>
> It answers "if I keep going the way I usually do, will this be done before it disappears behind the house until next year?" — early enough to do something about it.

Recommended targets

For each of your own telescopes, three targets you have never imaged, ranked for the night: how long they are up and how high, how well they fit that telescope's field of view, how much the moon will cost them, how well known they are, and how much they resemble what you usually shoot. Each comes with its reasons. **All · NB · BB** beside a telescope's name switches between every kind of target, narrowband targets only, and broadband ones only. **Not for me** stops a target coming back, and **Plan it** opens the planning form filled in for it.

Worth another visit

Finished is not always finished. **Revisit later** on a project marks it to come back to, with the filters you want to add — hydrogen-alpha for a galaxy, say — and a note. Those projects are listed first here, with whether the target is in season, when it returns if not, and how the moon suits each filter you picked, so a bright moon that rules out colour can still say "Ha good". Below them the page suggests finished projects with under six hours whose target is back in the sky.

Checklists, and each target's details

The **Catalogues** panel holds the checklists — Messier, Caldwell, the Herschel 400, Sharpless, van den Bergh, Barnard, LBN, NGC and IC — ticked from where your frames actually pointed rather than from what your folders are called. A galaxy caught in the corner of another target's frame counts, and so does data from friends' telescopes; a solid tile was a project's target, a dashed one was caught in another frame. Click any object for its details: when it is up, how it fits each telescope, which of your subs caught it, and an **AstroBin** link that opens other people's photographs of it in your browser. There you can also **Plan it**, star it as a **Favourite** (favourites are listed under **Last call** with their season), or overrule the tick with **Count it as imaged** or **Doesn't count**.

The all-sky camera and finding cloud

Optional, and the most involved piece of measurement in the program.

An all-sky camera is a cheap camera with a fisheye lens pointed straight up, taking a picture every minute. Many astrophotographers already have one. AstroTracker does not talk to the camera; it reads whatever folder the camera's software leaves pictures in — AllSkEye on Windows, indi-allsky on Linux, or any script that saves a JPEG. It can also read a single file or a URL. The Weather page shows the picture with its age, and says so plainly when it has gone stale for more than fifteen minutes, because a cloud watch that has silently stopped is worse than none.

Reading cloud from a picture

Once a source is set, the newest picture is read once a minute from half an hour before astronomical darkness begins to half an hour after it ends, which the program works out from your home site's latitude and longitude, so those must be set. Each reading is kept for two weeks.

The method rests on an observation that is counter-intuitive until you have seen it: at a light-polluted site, **clouds are brighter than the sky**. They reflect the town back down. So the detector builds a reference of what a clear sky looks like from the last hour of frames — anything that stays put, such as the mount, houses and a neighbour's lamp, cancels out — and anything that lights up against that reference is cloud. Cover is the fraction of sky lit that way; haze is then measured from the density of stars still visible in the parts that remain clear.

The reading is broken down by compass wedge and altitude ring rather than given as one number, because "cloud in the north-west, coming this way" is actionable and "40% cloud" is not. A trend, set against the reading about twelve minutes earlier, says whether it is clearing, holding or closing in. On the page the reading is one line, such as **sky now: cloudy · 62% cover**, then where the cloud lies or which way it is moving in from, the trend, and how many stars the camera sees against a clear sky's count. The words run clear, thin cloud, cloudy and overcast.

The camera works out which way it is pointing

For any of that to be reported as real compass directions, the program has to know how the camera is oriented — and nobody installs one level and square to north.

It works it out from the sky itself. The stars rotate about the celestial pole, and the direction of that rotation between frames is something only the sky can tell you; it settles which way up the image is. Then the bright-star catalogue pins down the exact geometry. This takes about forty minutes of a clear night, and is redone from scratch on a clear night about once a week, because nothing about the mounting is assumed. If you move or replace the camera, **Forget** under Settings › Weather and sky

clears the old answer until the next clear night, and **Find north now** beside it solves at once when forty minutes of frames are in. Two nights solved independently agreed to under two degrees.

Once north is known, compass letters are drawn on the picture with a ring on Polaris, and the cloud map beside it (dark for clear sky, light for cloud, grey where there is no sky) is drawn as the sky looks when you stand under it and look up — north at the top, east on the *left*, which is the opposite of a map of the ground and catches everyone out once. Until then, directions follow the picture, and Review cannot say how much cloud lay over a sub's target.

> **WHY IT WORKS THIS WAY**
>
> The moon is deliberately excluded from the cloud reading. It is a light, not weather; it moves; and it lights up the equipment in the frame. Left in, the sky around a bright moon reads as lit cloud and the detector cries wolf every month.
>
> Similarly, where the camera software writes both a display JPEG and a raw FITS frame, the FITS is what gets measured. The JPEG is re-stretched by the camera software as the scene changes — a yard light going out is enough — and one such re-stretch was once read as an overcast sky. The page still *shows* you the JPEG, which is in colour.

What the cloud reading is used for

- **An alert while imaging:** one notification per episode when cover stays at or above your threshold (60% unless you change it) for three minutes, saying which direction it is coming from.
- **An all-clear**, if you want one — see *Every alert it can send*.
- **On the Review page:** each sub is annotated with what the camera saw in the direction of its target at the moment it was taken, which turns "this frame looks bad" into "this frame was shot through cloud", and its **All-sky** button opens the camera's timelapse of that night at that minute.
- **On the Imaging card:** a chip with the cloud over the rig's target once it reaches 30%, and, when a rig has stalled and the camera sees 50% cover or more, "clouds" beside the stall instead of leaving you to guess.
- **A night history chart** above the picture, with a line at your alert threshold: hover or drag for any minute's reading, and click a minute to see the actual all-sky photograph and cloud map from that moment (**Back to live** returns).
- **A timelapse of the night**, played from the pictures the reading sampled — described next.

Playing the night back

▶ **Play the night as a timelapse**, in the All-sky camera panel on the Weather page, runs the camera's pictures one after another, the way Review plays a set of subs. They are the pictures the cloud reading sampled, one a minute through darkness, so every frame carries its own reading; the button

therefore needs **Read cloud cover from the camera during darkness** (Settings › Weather and sky, on by default). It opens on the most recent night, and a **Night** list at the top offers the others still available, up to about two weeks back. Where the camera writes colour JPEGs, those are what you see; a camera that saves only raw frames plays in grey.

Above the scrubber — the slider you drag through the night — the night's cloud reading is drawn as a line on the same time scale, labelled "cloud cover above, time below", so a bank of cloud is easy to find and drag straight to. The frame on screen shows its time and its percentage of cloud.

CONTROL	WHAT IT DOES
▶ Play / ⏸ Pause	Starts and stops the sequence; the Space bar does the same.
◀ ▶	One frame back or forward; so do the arrow keys.
10/s 20/s 30/s	How many frames a second.
all *N* frames	A quick play shows an even sample of no more than 300 frames. On a night with more than that, this tick box appears and loads every one.

Esc closes the player. A minute the camera has no picture for shows the previous frame rather than a blank, so the playback does not flicker.

The program keeps each sampled picture on disk as the night goes, up to eight nights of about 50 MB each, so the morning's play starts straight away. An older night is read from the camera's archive the first time you play it, which is slower — a progress bar shows how far it has got — and is kept after that.

KNOWN LIMITS, STATED PLAINLY

- Thin high cloud dims stars before it removes them, so it reads late.
- Dew or frost on the dome reads as overcast — which is arguably correct, since you cannot image through it either.
- In the last half hour before the watch stops for the morning, a brightening eastern sky can read as cloud.

Review — judging your photographs

Sixty frames arrived last night. Some of them are ruined. This is where you find out which, and it is where most of the program's measurement work surfaces.

Frames go wrong in ordinary ways: a cloud drifted over, the guiding slipped and the stars are short dashes, the focus wandered as the tube cooled, a satellite crossed the field, someone turned on a floodlight. Stacking a bad frame with good ones does not average the badness away — it actively degrades the result. So every serious workflow includes throwing frames out, and doing it by eye across a thousand subs is punishing.

What is measured

AstroTracker reads the actual pixels of each sub and computes the numbers below. New subs in your current-work folders are measured on their own soon after they arrive (see *The watcher*); anything else when you ask, with **Measure N unmeasured** or **Re-measure all** at the top of a project's Review page, or **Re-measure N subs** on the Library page for subs an older version measured. The line under the page title keeps count: measured, flagged, rejected and unmeasured.

MEASUREMENT	WHAT IT TELLS YOU
Star count	How many stars were detected. A sharp drop means cloud or haze took them away.
FWHM	How wide the stars are — sharpness. Rises with bad seeing or bad focus.
HFR	A second sharpness measure, less sensitive to the exact star profile.
Eccentricity	How oval the stars are. Round is good; oval means tracking error, wind, or optical tilt.
Background	How bright the sky was. Rises with moonlight, cloud lit from below, or a neighbour's light.
Noise	The random fizz on top of the sky. Shown for reference, and one of the things the cloud likelihood weighs.
Gradient	How uneven the background is across the frame. A ramp means light coming in from one side.
Altitude	How high the target stood when the sub was taken. Low down, the sky is brighter and the stars fewer with no cloud at all.
Star glow	Light spread around the brightest stars by thin cloud, as a multiple of clear sky. Explained below.
Sky change	How much the background moved against the subs either side of it. Explained below.

Cloud likely	A trained model's estimate, in per cent, of how likely the sub is cloudy. Explained below.
Score	A single combined number, so a list can be sorted worst-first.

In the detail pane each number stands beside the same number for the session's best sub, so you can see at once how far this one is off.

> **WHY IT WORKS THIS WAY**
>
> Star size is measured on a fixed sample of the brightest 200 stars, with saturated ones skipped. Both choices matter. If the sample floated with however many stars happened to be in the frame, a rich field would reach further down toward the faint end, where stars measure artificially tight — so the same optics would report different sharpness on different targets. And a star whose core is at full well is flat-topped, so measuring it reports it as wider than it is.
>
> The *count* is separate and uncapped. On a good frame from a large telescope it runs to seven or ten thousand.

Flags: comparing a sub to its own night

Absolute thresholds are useless here. A star size of 4.5 pixels is excellent on one telescope and poor on another, and the same telescope varies by night. So each sub is compared against the *median of its own session* — the same target, same filter, same night — and flagged when it stands out from its neighbours.

That is what produces flags like **clouds** (stars vanished and the background rose relative to the rest of the night), **trailing** (stars markedly more oval), **few stars**, and **bright sky**. A night that was fine until 1 a.m. lights up from 1 a.m. onward. A night that was uniformly mediocre produces few flags, because nothing in it stands out from the rest of it; the star glow and the cloud likelihood, below, look beyond the one night for that.

Where an all-sky camera is set up, each sub also carries the most cloud the camera saw in the direction of that target while the sub was being taken, which converts a suspicion into a fact. When that reached 30%, the card says so (**cloud over target: 45%**, with **check** added from 40%); otherwise, when the sky as a whole read 30% or more, the card gives that instead (**camera: 45% cloud**), which is all it can say before the camera knows where north is. Any sub with a camera reading also gets an **All-sky** button in the detail pane, which plays the camera's pictures of that night from the minute the sub was taken, with a line on the scrubber marking it (see *The all-sky camera and finding cloud*). The readings are kept for two weeks, so older subs show neither.

Three numbers for cloud

Thin cloud is the hardest fault to see by eye, and the one that does most damage once it is in a stack. The flags catch cloud that comes and goes; three more numbers go after the rest. The first two are

measurements and the third a suggestion; none of them rejects anything on its own. What the card and the detail pane show is below; *Finding cloud in your own subs* explains how each is worked out and how far to trust it.

Star glow

Thin cloud spreads a bright star's light into a faint halo, while its core, and so its measured size, hardly changes. **Star glow** measures that halo around the brightest stars in the frame, as a multiple of clear sky: **1.00×** is as clear as this target has looked in that filter on any night (on a project's first night, that night alone), and higher means more glow. Hover over the number to see it against its own night only.

From 1.4× a chip appears on the card, **glow 1.4×**, in amber, turning red from 1.6×. Below 1.4 the number is only in the detail pane, because clear subs wander a little on their own. The sort **Star glow, strongest first** puts the worst at the top.

It reads **no bright star** when the field holds no star bright enough to show a halo, or the session has fewer than four subs to compare, and **re-measure** on a sub measured by an older version. It needs the focal length and the pixel size in the sub's header.

Sky change

Watch a cloudy night as a slideshow and the background visibly shifts from frame to frame; on a clear night it sits still. **Sky change** puts a number on that: how unevenly a sub's background differs from the subs just before and after it in the same filter, taken within half an hour, as a percentage of the sky's brightness, with smooth changes such as the moon's glow taken out. Cloud drifting through is lumpy, so it raises the number.

There is no chip and no fixed level that counts as bad, because what is normal depends on how you shoot, above all on how long a filter wheel leaves between two subs of the same filter. Sort by **Sky change, largest first** and look at the top of the list. It reads **no neighbour** when no sub in that filter was taken within half an hour either side, and **re-measure** on a sub measured by an older version. It is a different thing from the **Sky** view described below, which shows you the background as a picture.

Cloud likely

Cloud likely is a trained model's estimate, in per cent, of how likely the sub is cloudy. It weighs the measurements together, star glow and sky change included, each against the rest of the night, the project's other nights and the neighbouring subs, and it allows for where the target, the moon and the sun stood (seen from the rig's own site), so a bright sky low down or near twilight is not called cloud. It catches thick cloud too, where the stars and their glow simply vanish. From 50% the card shows a chip, **cloud 64%**, in amber, red from 90%, and the sort **Cloud likely first** brings the doubtful subs to the top. Subs measured by an older version get a likelihood straight away; the star glow and sky change it also uses arrive after a re-measure.

A SUGGESTION, NOT A VERDICT

The model ships inside the program, learned from one library's hand-graded subs, and does not retrain itself on the grades you give. How well it carries over to other equipment and other skies is not yet proven. Treat it as a second opinion: it never rejects or moves anything on its own.

Altitude, and subs below your horizon

The **Altitude** row is worked out from the pointing and time in the sub's header and the coordinates of the rig's site, and does not appear for a sub without them.

If you have drawn your horizon for the rig's site under Settings, in **Observing sites** (see "Your actual horizon" in *Active*), a sub taken while the target was below it gets a red **below horizon · 19°** mark on its thumbnail, and the row reads, for example, **19° · below your 25° horizon**. The telescope was looking at a roof, a tree or the glow over a town. In the numbers that can look very like cloud, but it is not, and the cure is different: end the sequence earlier rather than wait for a clearer night. Only the drawn horizon sets off this mark; the altitude limit does not.

Grading

You grade a sub good or bad, and a bad one can be given a reason: **clouds**, **trailing**, **focus**, **airplane / satellite** (a bright line drawn across the frame), or **other**. A sub can carry more than one, since a frame can be trailed, out of focus and crossed by an aeroplane all at once. In the detail pane each reason is a button, lit when the sub has that reason; clicking it adds the reason or takes it off. The thumbnail shows the first reason and how many more (✕ **trailing +2**; hover for the full list), and the project page's sub list reads **bad · trailing, focus**.

Everything is on the keyboard, and most grades move you straight on to the next sub, so a hundred frames can be worked through quickly:

KEY	WHAT IT DOES
Arrow keys	Move to the previous or next sub.
X	Reject the sub, or take the rejection back, and move on.
C, T, F, A, O	Reject with that one reason (clouds, trailing, focus, airplane / satellite, other) and move on. The fast path.
Shift + reason	Add that reason to the sub, or take it off, and stay on it, so a second and a third can follow.
G or Space	Mark the sub good, or take that back, and move on.
U	Clear the grade.

B	Blink against the session's best sub.
V	Cycle the view: Thumb, Full, Sky.
Esc	Back to the thumbnail view.

Beyond one at a time:

- **Multiple selection.** Shift-click a range, Ctrl-click individually, or use a Select toggle on a touchscreen, then grade them all at once. A reason button then gives every marked sub exactly that one reason.
- **A whole session.** Each night-and-filter row of the sessions table has a **Reject all as…** list that rejects every sub of that night and filter with one reason — for the night where the clouds came in at midnight and never left. Subs you already graded by hand are left alone.
- **The rest of a session, good.** Beside it, **Rest good** marks every sub of that night and filter you have not graded as good, but leaves the subs the statistics flagged ungraded, for you to look at; subs you already graded are untouched. It asks first, then says what it did, for example "58 marked good, 4 flagged left ungraded".
- **Everything flagged last night.** That one is not on this page: the **Last night** panel on the Active page has a **Reject N flagged subs** button, which rejects every flagged sub of that night you have not graded, using the flag itself as the reason (see *The morning summary*).

Auto-grade, a test feature

Off by default, and switched on under **Sub quality** in Settings (**Auto-grade on the Review page (test feature)**), it adds an **Auto-grade (test)** button to the top of every project's Review page, which lets the cloud model reject the subs it is surest about. The button opens a panel that counts the subs nobody has graded that have a cloud likelihood, and offers three levels. Each takes a sub when either number reaches it, and shows how many subs that would be before anything is touched:

LEVEL	TAKES A SUB AT	ON NIGHTS THE MODEL HAD NEVER SEEN
Cautious	cloud likely 90% or more, or star glow 1.8× or more	caught about 2 in 3 cloudy subs, and wrongly took about 1 good sub in 40
Balanced	cloud likely 70% or more, or star glow 1.6× or more	caught about 5 in 6, and wrongly took about 1 good sub in 20
Eager	cloud likely 50% or more, or star glow 1.4× or more	caught about 9 in 10, and wrongly took about 1 good sub in 10

Pick a level (Cautious to begin with), tick **also move them with my rejected subs** if you want them moved as well (not available in read-only mode), and press **Reject N as clouds**, which asks once more. Those subs become rejected as clouds, and their cards say • **auto**. A sub you grade by hand afterwards is your grade again. **Undo auto grades (N)** clears every grade the model gave in the project; files it already moved stay where they are until you undo that move on the Organize page.

> **TREAT IT AS A TEST**
> It is there to be checked against a stack, and it will sometimes be wrong. Auto grades are never used to train the model: the training export marks them as the model's own verdicts, not yours.

Looking at them properly

Thumbnails are not enough to judge a frame, so each sub can be opened three ways:

Thumb

> The small cached preview, for fast scanning.

Full

> A large, stretched rendering where star shape and faint detail are visible.

Sky

> The background only, with stars removed and the remainder stretched hard. Thin cloud, haze and gradients that are invisible in the normal view stand out as blotches and ramps. This is the view that makes "reject: clouds" an easy call.

There is also **Blink**, which alternates between the current sub and the session's best one — the classic way to see a difference that is real but too subtle to hold in memory.

Full and Sky open over the page rather than replacing it, so the grid keeps its place and closing returns you to the sub you were looking at.

Timelapse

Pick a filter and a sort order, press **Timelapse**, and every sub you are looking at is loaded with a progress bar and then played in sequence.

This shows what no per-frame number can. Cloud drifting across a night, focus gradually softening, the target slowly walking across the frame because the mount is not quite polar aligned, a satellite trail, a gradient rotating as the sky moves — these are obvious when the frames run one after another and nearly invisible one at a time.

It plays at 2, 5, 10 or 20 frames a second. The space bar plays and pauses, the arrow keys step one sub at a time, and a scrubber jumps anywhere in the sequence. The reject buttons and all five reasons,

airplane / satellite included, stay on screen throughout, and the grading keys (X, C, T, F, A, O, G and U) work there too; only the space bar means play and pause instead of good. Grading pauses playback — because grading is a decision about the frame in front of you, and it should not scroll away while you make it.

Everything is fetched before playing so that nothing stutters: about fifteen seconds for a 147-frame set. Very large sets play through a moving window rather than loading whole.

Sorting and filtering

Chips above the grid show one filter at a time across every night, so red can be compared with red — which is what you want, since filters differ systematically and mixing them hides real outliers. The sessions table narrows to a single night and filter, and the two work together. Tabs switch between all subs, the flagged ones, the rejected ones and the ones not yet measured. A sort box orders by date (oldest or newest first), by filter then date, worst score first, star glow strongest first, sky change largest first, or cloud likely first. Each thumbnail carries the calendar date and time it was taken, such as **2026-09-15 04:16**, and its filter; for a sub taken after midnight, hovering over the date shows the observing night it belongs to (**Night of 2026-09-14**).

Quarantine

Rejected subs stay in the database with their grade and reason no matter where the files go — that record is valuable. Grades and their reasons are the labels a cloud model learns from (the one in the program learned from subs graded this way), and the training export in Settings can write yours out. But you do not want them in the folder you are about to stack.

Move rejected subs (N) at the top of the Review page (**Move N bad frames** on the project page) moves them out in one undoable batch; in read-only mode it is greyed out. If you set a rejected-subs folder they land in `<folder>\<rig>\<project>\`; without one they go to a `bad` folder beside their siblings. Either way they remain tracked, marked as moved, and out of your integration totals.

Organize — moving files safely

The program's answer to a decade of accumulated filing drift. It is also the place where the most care has been taken, because this is where a bug costs you data.

Every section on this page folds away and stays folded, because on a large library the page is otherwise twenty-odd screens long.

Proposed moves

AstroTracker compares where files are against where their headers say they belong, and proposes moves in groups. It only ever proposes moves within your current-work folders: an archive is finished work, and is left as it is.

Misfiled subs

> Lights belonging to a project whose folder is elsewhere, identified by pointing and rig. Flats whose name or rotation says they belong go with them.

Stray subs

> Lights with no target folder at all, sitting loose in a root or stage folder.

Darks and bias in project folders

> Dark and bias frames saved inside a project folder, to be filed into your Calibration library folder under the rig's name and their settings, such as `MyScope\Darks_300s_100g_50o_-10C`. Shown only when you tick **Include darks and bias found in project folders**, and only if you have a folder with the Calibration role.

Two targets in one folder

> The SGP problem described in *Turning photographs into projects* — each target gets its own folder beside the one they landed in. A folder belongs to the target it is named after, however many subs of another target end up in it.

For each proposed move you see the file, its project, its type, its night, where it would go and why. You tick what you want. Nothing whatever happens until you confirm.

The safety rules

- Files are **never overwritten**. A collision is reported, not resolved silently.
- Every batch is **logged** with each individual move.
- Every batch can be **undone**, putting everything back where it came from.

- A move outside your library roots is **refused** — the program will not write to a destination it does not manage.
- **Read-only mode: never move files**, under Settings › Safety, switches all of it off, undo included.

Archive finished projects

A list of completed projects whose files still sit on a working drive, with the destination each would go to in your archive. Tick, confirm, and the lights and flats move together, night folders and all, as one undoable batch. Darks stay in the dark library where they belong. It needs a library folder with the Archive role; until you add one on the Library page, the section says so.

The archive layout is a template, `<year>/<rig>/<target>/<night>` by default, and the section shows the one in use. You change it under Settings › Library layout, where **Learn from my archive** reads the pattern off the projects already there, so nothing about your folder structure is assumed. The year used is that of the project's last night, so a project spanning New Year stays in one folder. Each destination can be changed with **Edit** before you confirm, and **Reset** puts back the proposal.

Move history

The last fifty batches, newest first, each with what it was, how many files moved, how many failed, a details view listing every individual move, and an Undo button; older batches stay in the log. This is the record that makes the whole feature safe to use: nothing is unrecoverable, and you can always see what happened.

Moves started on other pages are batches like any other, listed and undone here too: **Move bad frames** on a project page, **Move rejected subs** and the auto-grade's **also move them with my rejected subs** on the Review page, and the button on the **Last night** panel that moves the night's rejects. Undoing a batch of rejected subs puts them back where they were and takes the rejection away, so they come back ungraded.

Duplicates

The same capture — same camera, same exposure, same timestamp — indexed in more than one place. This happens through years of backups, half-finished reorganisations and copies made "just in case". The page groups them by folder pair, so you see the shape of the problem: "these 820 captures exist in both of these two folders."

AstroTracker never deletes anything. It shows you the copies, with file sizes, and a **Compare pictures** button that puts their thumbnails side by side so you can confirm they really are the same capture; a group whose copies differ in size is marked **sizes differ**, which is worth a closer look before you delete either. A Show button opens the folder. The decision, and the deletion, are yours.

A grade belongs to the exposure, not to the file, so you need not work out which copy you graded before choosing which to delete. Once a copy is gone, the next scan of its folder notices, and the copy you kept takes over whatever it lacked: the grade, the quality measurements, and the project if it had none. Nothing it already has is replaced, and reject reasons come across only with the grade they explain: a copy you passed as good never picks up a reason from a copy you rejected. The same rule lets a sub you move to another folder by hand keep its grade and its measurements.

Exports

The last section holds one button, **All projects as CSV**: a spreadsheet with a row per project giving its target, rig, status, hours in total and per filter, subs, nights, first and last night, and coordinates. A single project's AstroBin file and text summary are on its own page.

Library

The folders themselves: the library roots described in *Telling it where your files are*, which also explains the roles and why a folder inside another root is refused. To add one, type or paste its path, choose **Current work**, **Archive** or **Calibration library**, and press **Add folder**; a line under the form explains the role you picked. The rejected-subs folder is not added here: it joins the list once you set it under Settings › Rejected subs.

Each root gets one row:

COLUMN	WHAT IT SHOWS
The dot	Green when the folder can be reached, red when it cannot: an unplugged drive, or a network share that is switched off.
Folder, Role	The path, and the role you gave it.
Files	How many files the program has indexed there.
Last scan	When the folder was last scanned, or "never".
Watching	**live** for a working folder watched for new files as they arrive; **scheduled** for one rescanned on a timer (archives, calibration libraries, and working folders when live watching is off in Settings); **error** when the watch could not be set up, most often because the folder cannot be reached (hover for the reason); **off** when the watcher is not running. After that come how many file changes it has noticed, "pending rescan" while it waits for the folder to go quiet, and when the next scheduled rescan is due, in UTC rather than local time.

Each row has a **Scan** button, to scan that folder now rather than wait for the schedule, and a **Remove** button. With more than one folder, **Scan all** scans them one after another.

When a drive is unplugged

Nothing needs switching off. The folder's dot turns red and its Scan button is greyed out. Scheduled rescans pass it by, and **Scan all** skips it without touching anything. Its files stay in the index, grades and all, and everything carries on once the drive is back.

> **REMOVE FORGETS**
>
> **Remove** touches no file on disk, but the program forgets the folder: its files leave the index, and with them the grades and reject reasons you gave them. Adding it again starts from scratch, apart from `BAD_` names and `bad` folders, which are read afresh. A folder's role cannot be changed once it is added, and removing and re-adding it is no free way round that, so choose the role with care.

Re-measure

This is also where the **Re-measure** button lives. When the program's measurement method improves, existing measurements become stale — computed a different way from today's. The button appears only when some of your subs were measured by an older version, and says how many: **Re-measure 1,234 subs**. It is deliberately scoped to subs that already have measurements rather than to every file you own: measuring a library for the first time is a much longer job and a different decision.

Pressing it asks first. A sub measured by an older version lacks readings newer versions take, such as the star glow and the sky change that the cloud likelihood on the Review page is built from (see *Finding cloud in your own subs*). The work runs in the background while you carry on: an hour or more on a large library, longest for subs on a network drive, with a progress bar, a count, how many subs could not be read and the file it is on. Moving or deleting subs meanwhile does not hold it up. To bring a single project up to date instead, use **Re-measure all** on that project's Review page.

Regroup projects, and the dark library

Regroup projects rebuilds the grouping of files into projects from scratch, from the files themselves. Subs you placed in a project by hand keep it, and planning projects, goals and notes are kept; an automatically grouped project you renamed may get its default name back, which is why it asks first. It cannot run while a scan is running.

At the bottom of the page is the **Dark and bias library**: every dark and bias run found anywhere in your folders, by camera body, described in *Calibration: darks, bias and flats*.

Rigs

Each telescope setup, grouped into Identity, Optics, Camera, Capture PC and Calibration.

Most of it is discovered from your files. What you add is the part no header can tell it:

- **A name** — "ED127" rather than a fingerprint.
- **What its filters are called**, and how wide the narrowband ones are. There is no list to type: every name the `FILTER` header has carried on the rig's frames is read through a built-in table, so "Lum" counts as L and "H-alpha" as Ha. **Show filter names** lists them, each with a **Shown as** name you can correct and, for a narrowband filter, a **Bandpass** in nanometres; **Save filter names** stores both. A corrected name applies to every frame of the rig, old and new, and its per-filter totals follow. This drives all the moonlight advice; a narrowband filter with no width entered counts as 6 nm.
- **Which site** it observes from, which settles its darkness times, forecast and time zone.
- **Where its log files are** (**Log folder**), so guiding and sequence events can be read.
- **Its capture PC's address** (**SGP API host**), for live camera, focuser and mount state on the Imaging page, and whether **Allow control** lets the program send it commands.
- **Its own guiding limit**, if a rig with a very different focal length or mount needs a different threshold from the global one.
- **A calibration folder**, if the rig keeps its own — see *Calibration: darks, bias and flats*.
- **Retired** or **owned by someone else**, which quiet it down as described in *Recognising your telescopes without being told*.

At the foot of each rig, **Confirm** (**Save** once it is confirmed) stores your changes. **Delete**, **Merge into…**, **Split by folder…** and the **Ignored configurations** list below the rigs are described in *Recognising your telescopes without being told*.

The focuser temperature model

A worked example of the measure-do-not-assume principle. Telescopes shift focus as they cool, and SGP can compensate automatically if you give it a coefficient in steps per degree. Finding that number is normally a manual experiment most people never do.

AstroTracker derives it from your own nights. Each time a rig finishes an autofocus run, it records where the focuser settled and the temperature the focuser reported. That needs the rig's **SGP API host**: the log files mention the temperature too rarely, so a rig whose logs are only copied across records nothing and shows no focus panel. Once there are ten runs spread over at least eight degrees that actually track the temperature, the rig's **focus vs temperature** line states its own steps per

degree and where the focuser should sit at the temperature it reads now, give or take the scatter of the runs behind it.

Until then it says what is still missing rather than a number. Four runs on one mild night would fit a slope through nothing but scatter. Expect it to say "collecting" for a season. A focuser with no temperature probe cannot take part, and nothing says so: SGP reports a placeholder that is thrown away, so the panel keeps saying "Collecting: no autofocus runs recorded yet". If it still says that after a few autofocus runs, check that the focuser reports a temperature.

Hardware

Cameras and telescopes entered once as things in their own right, then picked from dropdowns when you build a rig — rather than typing the same sensor dimensions and pixel size into every rig that uses that camera. An "Import from my rigs" button seeds the list from what has already been discovered.

This matters because equipment gets recombined. The same camera moves between three telescopes; a telescope gets a different camera for a season. Keeping cameras and optics separate from the combinations they form means the history stays coherent when you rearrange. Editing a number here updates every rig that uses it, and a reducer or Barlow is a second optic entry at its own focal length, since the focal length is what the headers carry. The cameras entered here are also the camera bodies the dark library is organised by.

Calibration: darks, bias and flats

The bookkeeping that quietly ruins people's data, tracked properly.

Dark runs

Darks are organised by **camera body** first, then by **run** — one folder shot on one night. The **Dark and bias library** at the foot of the Library page lists each run's sets, dates and frame counts. A light frame matches darks from its own rig's camera with the same exposure, gain, offset and binning, a temperature within 2 °C and shot within a year of the lights (both under Settings › Calibration matching), choosing the nearest run in time, and the project page names which run it would use for each exposure group.

Runs can be **retired** with the **Retire** button beside each one, which is essential after a sensor is cleaned or serviced, or the darks are re-shot: a retired run's darks stop matching every light, older ones too, rather than lingering and quietly being applied to data they no longer describe.

Where two cameras on the Hardware page share a sensor model, a run's body may be genuinely ambiguous. The program works it out from the telescope it was shot on or a folder named after the rig, and where it truly cannot, it raises a **Dark runs** alert and you pick the body in the dark library; until then the run matches either body, marked unverified. People with one camera are never asked.

> **WHY IT WORKS THIS WAY**
>
> Older versions of SGP did not write the camera's `OFFSET` into headers. Rather than treating a missing offset as zero — which would silently match darks that may be wrong — such frames match on the other keys and are counted as **unverified**. The program tells you the match is uncertain instead of pretending it is not.

Flats

Flats are matched to lights in the same night folder first; without one, flats shot within 14 days of the session count (Settings › Calibration matching). If a rig keeps a single calibration folder that you re-shoot periodically, tell the program where it is: everything in it from that rig's camera belongs to the rig, and the flats become dated sets, filter by filter, read from the headers — so the subfolders can be named anything at all.

The folder is scanned for you: outside your library roots it becomes a Calibration library root of its own, named after the rig; inside one, that root's scans already cover it. Either way its frames belong to the rig because of where they sit, not because of which root found them. Only flats under a Calibration library root become dated sets, though, so keep such a folder inside a Calibration root, not a Current work or Archive one.

A set applies to every project of that rig from the night it was shot until you shoot a newer one, and a set shot up to two days after the lights still counts for them. The rig page lists the set currently in force per filter, its age and rotation, the older sets, and lets you retire one. Per rig you can also set:

- **Flats must match rotation**, off by default: only flats shot within 3° of the lights' rotation count — essential if you rotate the camera between targets, since a rotated flat corrects dust in the wrong place;
- **Flats valid for, days**: how long a set applies (blank: until a newer one);
- **Remind me when flats are older than, days**: an alert when the set in force gets old (blank: never);
- **Darks valid for, days**: how old darks may be (blank: the Settings value).

Flats in a project's own night folder always win over the shared set, and the project page marks flats that came from the rig's folder and names the set, so you always know what is being relied on.

Alerts

A single list of things that are actually wrong and actionable.

Alerts are computed from the current state of the library rather than stored, which means they disappear on their own when the situation is fixed. They cover:

- Missing flats for a given project, filter and night.
- Missing darks; darks in use that are older than the **Dark-age reminder** in Settings, if you switch it on; and dark runs needing a camera body assigned.
- Flats whose rotation does not match the lights, with how far off they are, on a rig whose flats must match rotation.
- Gain or offset drift — a setting changed between sessions, which invalidates calibration.
- Flats in force for a rig that are older than its **Remind me when flats are older than, days** on the Rigs page (off until you give it a number).
- A rig the scan has found that you have not confirmed yet.
- A rig the scan created from only a handful of photographs, which is usually a mistake.
- A filter name in the headers that matches none of the standard names.
- Files whose headers cannot be read, per library folder: still being written, or not FITS.
- Postponed projects whose target is back in season.

An alert about a project links to it; one about a rig has a **Rigs** button. The one-time nudge for a rig imaging without a log folder is a notification rather than an alert, described in *Every alert it can send*.

Rigs to confirm, and rigs that look like mistakes

A rig the scan discovers arrives unconfirmed, with a **Rig** alert, "New rig detected", under the name the program made up for it: give it a proper name, say whether its camera is mono or colour, and press **Confirm** on the Rigs page. One created from fewer than ten lights (Settings › Rig review) gets a **Rig review** alert instead, because that is far more often a camera setting changed by mistake for part of a night, such as a cropped readout, than a new telescope. The alert names the existing rig it resembles, if there is one. Confirm the rig if it is real; otherwise **Merge into…** or **Delete** it, as *Recognising your telescopes without being told* describes. Nothing on disk is touched either way.

Filter names

Capture software writes whatever the filter wheel was told to call each position, so one hydrogen-alpha filter can arrive as "Ha", "H-alpha" or "Ha 3nm"; common spellings like these already count as one filter. A header value that matches none of the standard names — L, R, G, B, Ha, OIII, SII, and None

for no filter — raises a **Filter name** alert with the number of frames carrying it. Set what it should be shown as under **Show filter names** on the Rigs page (see *Rigs*). A filter with no standard equivalent, such as a dual-band filter, can keep its own name; dismiss its alert.

Dismissing, and old nights

Dismissing an alert hides it until the underlying situation changes — each alert has a stable identity, so dismissing "no Ha flats for 2026-09-06" does not hide a different night's problem. Each kind has its own tab with a Dismiss-all button, and dismissed alerts can be shown and restored at the foot of the page.

A date under Settings › Sub quality, **Hide calibration alerts for nights before**, hides missing-flat and gain-or-offset drift alerts about nights before it, which is how you adopt the program mid-library without drowning in complaints about 2019. Missing-dark alerts belong to a project rather than a night, so the date leaves them alone. Missing flats, missing darks and drift stop being raised once a project is marked complete or cancelled, and a rig that belongs to someone else, or has been retired, raises no calibration or drift alerts at all.

Finished, and the campaign view

Finished lists completed and cancelled projects, grouped by telescope: a card per rig with its hours and projects, your own rigs first, then each friend's, with retired rigs at the end of each group. Pick one to see its projects, newest first, with a box to filter by target. Cancelled ones are labelled as such, so a target you gave up on stays visible as a record of what was tried rather than vanishing.

The **campaign view** takes one target across every rig and every owner and shows it as a single body of work. This is for the case where you have shot the same nebula on two of your own telescopes over three years, or where several people are contributing to one image: the individual projects stay separate, with their own calibration and framing, but the campaign totals them, filter by filter, with each project's share. Open it from **View campaign →** on a project page, which appears when another project has the same target.

The finished picture

When you mark a project complete, the program offers to keep the processed image you made from it (a JPEG, PNG, WebP or TIFF). It appears on the project's card in Finished and in the **Final image** panel on the project's page, where a click opens it full screen, a further click shows full resolution and Esc closes it. The program keeps two JPEG copies of its own, one to view, at most 4096 pixels on its longer side, and a 720-pixel thumbnail, in a `finals` folder inside its data folder, so even a very large TIFF costs only a few megabytes; the file you chose is never touched, and the daily backup copies the `finals` folder too. A complete project without a picture shows an empty frame with an **Upload the final image** button; **Replace** and **Remove** change it later.

Settings

Grouped by subject, in the order they appear on the page:

GROUP	WHAT IS IN IT
Observing sites	Your home site: name, latitude, longitude, elevation, time zone (blank means this PC's), **Altitude limit** (30°), **Bortle class** (6) and horizon. Below it, any other sites, each with its own horizon, time zone and Bortle class. Feeds every sky calculation.
Calibration matching	How old flats may be when there is no Night folder to pair them by (14 days), how far a dark's temperature may differ from the lights' (2 °C), and how old a dark may be (365 days).
Dark-age reminder	**Remind me to retake darks**: an alert when a dark set your active projects rely on is older than a number of days (90). Off by default: plenty of imagers never need to retake darks.
Safety	**Read-only mode: never move files** (off). Switched on, it stops every file move: Organize and Quarantine are switched off; everything else keeps working.
Library layout	The archive template, stage folder names, and **Learn from my archive**.
Rejected subs	Where quarantined frames go. Left empty, they go into a `bad` folder beside the rest.
Phone access	**Allow phones and other computers on this network** (off), an optional second code for Tailscale, and the pairing QR codes with their links (*Your phone*).
Push notifications	**Send push notifications through ntfy** (off), the per-phone setup, a test push, and how late a rig's next sub must be before it counts as stalled (10 minutes).
Weather and sky	The rain warning (on: a 50% chance within 3 hours), the season reminder (on), the all-sky picture folder or URL, reading cloud from it (on), the cloud warning (on, at 60% cover), the all-clear (off, at 15% or less), the sky circle, and where north lies in the all-sky picture.
Guiding and seeing	The guiding warning (on): a limit of 1.5″ total RMS, held for 5 minutes.
How much more?	**Needs more below, SNR** (4), **Clean from, SNR** (6) and **Judged per patch of** (3 × 3 pixels): where a filter's verdict changes, as *Signal, noise, and when a filter is done* explains.

Morning summary	Whether it is sent (on), the time it arrives (08:00), and **Send a test summary now**.
Sub quality	**Hide calibration alerts for nights before** (a date), the trailing thresholds (eccentricity 0.75, and 0.15 above the session's median), **Re-flag measured subs**, and **Auto-grade on the Review page (test feature)**, off by default (see *Review — judging your photographs*).
Rig review	How few lights make a rig the scanner created worth a second look (10).
Watching and rescans	Whether current-work folders are watched (on) and new subs measured automatically (on), the settle time (8 seconds), and how often everything is rescanned: current work every 15 minutes, archives every 6 hours.
Target catalogues	Downloads the object catalogues the Planning page and its checklists use, or imports an OpenNGC file you downloaded yourself.
Data on this PC	Where the program keeps its data, how much room the database, thumbnails, sky cache and log take, and **Clear thumbnails**. Your FITS files are never copied here.
Backups	**Back up the database daily** (off), the folder, the time (09:00), copies to keep (7), **Back up now**.
Updates	**Check for updates** (on, every 6 hours), automatic installs (off), the channel (stable), **Check now**.
Export training data	Writes your graded subs, the ones the program flagged and a sample of ordinary ones to a folder, as training material for the sub-quality model on another machine. Running it again adds only what is new.
Licence	How many trial nights are left, or who the copy is licensed to; the box to paste a key into; links to the licence terms, the privacy statement and the open-source licences.
Support	Saving a problem report to email to the developer, and **What's new** for the full version history.

The groups from Observing sites to Watching and rescans, and the Backups and Updates choices further down, are kept when you press **Save settings**, which sits just below Watching and rescans. The other sections below that button work through buttons of their own, and so does the list of other observing sites, which keeps each change as you make it.

Guides

Reached from the top right of any page. How-to articles, grouped by subject. Most are for the parts of the setup that involve other machines — getting a rig's log files onto this PC (with Syncthing, if you like), opening the one network port a rig needs, the all-sky camera, phone access and push notifications, backups. The rest explain things people trip over in the program itself: what each library folder's role means, two targets that ended up in one folder, flats and rotation, controlling a rig from the Imaging page, and how the focuser's response to temperature is learned.

Two details make these more useful than ordinary documentation. Commands come with *this* PC's own address already filled in and a Copy button, so they can be run as they stand rather than adapted. And a small **i** beside individual fields opens the guide for that specific field — so the explanation of "SGP API host" is next to the box asking for it, rather than in a manual you would have to know to go and read.

The watcher

Most of AstroTracker's value is delivered while nobody is looking at it.

Working drives are watched continuously. A file appearing at 1:14 a.m. is noticed within seconds: its header is read, it is attached to the right project, the rig's live card updates, and the count of subs for the night goes up. Nothing has to be scanned and nobody has to press anything.

It is within seconds rather than instantly, and on purpose. The program waits until the folder has been quiet for a few seconds — eight by default, the **Settle time** under **Watching and rescans** in Settings — before reading it, so a file that the capture program is still writing, or that is still arriving from another PC, is never read half-finished.

Alongside the live watch, a reconciling scan runs every fifteen minutes by default to catch what a file-system watcher can miss — files that arrived while the program was closed, a drive that was disconnected, a folder moved by hand in Explorer. Archive and calibration roots are scanned on a schedule instead, every six hours by default, because nothing changes there unprompted.

New subs are measured as they arrive

Whichever way a new sub on a working drive is found, filing it is only half the job. Once the scan has filed it, the program measures it in the background and makes its thumbnail for Review; when a batch is done, the project's flags are worked out again, each sub against the rest of its session. So by the morning Review already has last night's numbers and flags, and there was nothing to start by hand. What the numbers mean is in *How a photograph is measured*.

It is on by default: **Measure new subs on current-work folders automatically**, under **Watching and rescans** in Settings. Only light frames on working drives are measured this way, including older ones in a working folder you have just added. Subs in an archive are measured only when you ask, because measuring years of old data is hours of work you should choose to start.

> **IN PLAIN TERMS**
>
> The scan itself still reads only headers. Measuring is a separate, slower job that opens the picture, one sub at a time, after the scan has filed it — and like everything else here, it only reads the file and never changes it.

When a folder cannot be reached

Drives get unplugged and network shares go to sleep. A folder that cannot be reached at its rescan, or whose scan fails part-way through, is noted in the program's log and simply tried again at its next turn — fifteen minutes later for a working drive, six hours for an archive, with the default settings — rather

than every few seconds. The first rescan after it comes back picks up everything that arrived meanwhile. The *Library* page shows which folders can be reached right now.

Pages that keep themselves up to date

You do not have to press reload either. When a scan adds, changes or removes a sub, or a batch of measurements finishes, the Active, Imaging, Weather, Planning, Review and Finished pages refresh themselves, and so does the count on the Alerts tab. A scan that found nothing leaves the page alone, so a Review page is not rebuilt under you while you grade.

The other watches

Watching folders is one of several background jobs, each on its own timer. The alerts in *Every alert it can send* come from three of them: the night checks, once a minute (the morning summary, the stall and rain alerts, the log-folder reminder, the daily backup, the update check and the season reminders); the log reader, every thirty seconds (guiding and sequence trouble); and the all-sky watch, once a minute (cloud, the all-clear and a stopped camera). A long archive scan never makes a cloud warning wait, and a check that fails is written to the log while the rest carry on. The night checks take turns, though: if the daily backup is writing to a network share that has stopped answering, the stall and rain alerts wait until the copy finishes or gives up.

Every alert it can send

Notifications go to three places at once: a Windows tray notification, a banner across the top of the app when a page is open, and — if you set it up — a push to your phone that reaches you anywhere. Together they mean you find out whether you are at the computer, in the house, or away. A few that only matter when you are sitting at the computer stay off the phone; they are pointed out below.

The stall alert

The most valuable one. While it is dark and a telescope has been producing subs, the program works out when the next one is due: the sub being taken has to finish, and the next takes its own exposure, so twenty-minute subs arriving forty minutes apart is not a problem. When the next sub is more than ten minutes late (the margin is yours to change), something has gone wrong — the mount lost the plot, the sequence errored, the camera disconnected, clouds arrived — and you get one notification.

It is careful about false alarms. A rig whose log shows the sequence has ended — finished, or aborted, which has an alert of its own below — does not raise a stall, and none is raised in the last minutes before dawn, when a night ending is expected. The notification itself asks the question — sequence stopped, clouds, or a crash? — and says how many subs the rig took tonight and when the last one started. With an all-sky camera, the rig's card on the Imaging page answers it: at fifty percent cloud cover or more, its stall warning adds "clouds" and the cover, so you can tell weather from a failure at a glance.

> **IN PLAIN TERMS**
>
> The alternative is setting an alarm for 3 a.m. to go and check. This is the feature that lets people sleep.

Guiding

While a rig is guiding, its total RMS error is read from the PHD2 log every thirty seconds. When it stays above the limit — 1.5 arc-seconds by default, or a per-rig value — for five closed minutes in a row (the number of minutes is yours to change too), you get one warning with the RA and declination split, which usually identifies the cause. A quiet note follows in the session log when it settles again.

Five closed minutes matters: guiding spikes constantly for a few seconds at a time, and an alert on every spike would be noise. Five sustained minutes is a real problem.

Cloud, and the all-clear

With an all-sky camera, one notification per episode when cover stays at or above your threshold (60 percent by default) for three minutes while a rig is imaging, saying which direction it is arriving from.

There is a matching all-clear, off by default. It only ever follows a warning — so it cannot arrive out of nowhere — and it waits for a genuinely open sky, fifteen percent cover or less held for about twelve minutes. A brief gap in the cloud will not fire it. That asymmetry is deliberate: the warning is worth a glance, the all-clear is worth getting out of bed for. The message says how long the cloudy spell lasted and what each rig is currently doing, so you know immediately whether anything stopped and needs restarting.

That covers the case where you were already imaging. For the other case — sitting up waiting for cloud to clear *before* starting — there is a **Notify me when it clears** button on the Weather page. It does not care whether anything is imaging, it is armed for tonight only, it can be cancelled, and if the sky is already clear it tells you so rather than making you wait for nothing.

When the all-sky camera itself stops

The cloud watch is only as good as the pictures it is given. If the camera stops producing them while a rig is imaging — its software has stopped, say, or its computer has gone to sleep — you get one notification once the newest picture is fifteen minutes old, saying when it was taken and which rigs are imaging. If the folder or web address the pictures come from cannot be read at all, that notification comes at once and says so. Without it, the first sign would be the guiding alerts when the cloud nobody warned you about arrives. The watch carries on by itself when pictures come back.

Rain

While a rig is imaging or the site is dark, one warning when the forecast shows rain within your lead time. A forecast, not radar.

Sequence trouble

Read from the SGP and PHD2 logs: guide star lost, failed recovery attempts, and SGP's own error notifications. These are the events that turn a productive night into four hours of the mount pointing at nothing.

The serious ones — a lost guide star, a failed recovery, an error, an aborted sequence — reach your phone. Lesser warnings, such as SGP starting a recovery, stay in the tray and the app. And the same trouble on the same rig within ten minutes arrives as one notification, not five.

Season reminders

The first night a postponed target is up for at least half the dark hours: "You can continue M106." It arrives that afternoon, once per season, and shows on the Alerts page too.

Housekeeping

The rest are about the program itself rather than the night.

- **A backup that failed**, usually because the folder it copies to cannot be reached. It asks whether that folder is reachable and tries again an hour later.
- **A damaged database**, found by the check that runs every time the program starts. The notification sends you to the app, where a banner offers to restore the newest backup; see *Backups, updates and staying alive*.
- **Updates**: one notification when a new version is available, and one when it has been installed. With automatic installs on, you are also told when a download was refused because it is not signed by the AstroTracker publisher, or when this copy cannot install itself (because it sits in Program Files, say), with what to do about it. The brief notice that an install is starting stays on the computer.
- **A rig imaging without a log folder**: a one-time nudge per rig, so guiding and sequence events are not silently missing from your records. It goes to the tray and the app only, never to your phone.

The morning summary

One message, waiting when you get up, about the night that just ended.

It reports:

- How much integration the night added: the total, then each project's hours filter by filter, for up to four projects, the biggest first.
- The seeing over the whole night, when its stars allow an estimate, as a word and a number — "Seeing good, about 2.2″" — worked out as *Seeing, guessed from your own stars* explains.
- How many subs were flagged — for cloud, for trailing, or for showing no stars at all.
- Anything that needs acting on: flats not taken, a gain or offset that drifted, subs of one project sitting in another project's folder.

The point is that a night either needs your attention or it does not, and you should be able to tell which in ten seconds over coffee: anything worth acting on is gathered on one last line that begins "Check:", so a night without that line needs nothing from you. It deliberately says nothing about a night more than two days old, whether on a first launch or after the computer has been off for a few days: by then it is no longer news.

It goes out once a day, after 8 a.m. unless you choose another time under Settings › Morning summary, where it can also be switched off. A night is dated by the evening it began, so its summary never goes out while the night is still running; if the computer was off at that hour it arrives when the program next starts, and a night that produced no subs sends nothing. It travels like every other notification: the Windows tray, a banner at the top of the app, and your phone if push is set up (*Your phone*). **Send a test summary now**, in the same place, sends the latest night's summary at once, whatever the time.

The **Last night** panel near the top of the Active page shows the same night in more detail — the number of subs, every project rather than the first four, and what each project's flags were for — though not the seeing or the note about subs in another project's folder. It has two buttons: **Reject N flagged subs** rejects every flagged sub you have not already graded, using the flag as the reason, and **Move N rejected to dump** moves all rejected subs of the night to the rejected folder across every project, as a single undoable batch. For a cloudy night, that is the entire cleanup in two clicks.

Backups, updates and staying alive

Backups

The database holds things that exist nowhere else: every sub's grade and the reason for it, your dark runs and their camera bodies, hardware, sites, horizons, notes and goals. Your image files are safe on disk regardless — but the judgements are not.

So: a folder (a NAS share works), a time, and how many copies to keep. Backups are off until you tick **Back up the database daily** and name the folder. Once a day the database is copied using SQLite's online backup, which is safe while the program is running, then verified, and older copies beyond your limit are removed. The same run copies the pictures you added as projects' finished images into a `finals` folder beside those copies, because they live in the data folder rather than in the database. If the copy fails because the share is offline you get a notification and it retries an hour later. There is a **Back up now** button, and with backups set up one is written automatically before any update installs.

Moving to a new PC, or restoring by hand, is: install, quit from the tray, copy the newest backup (a file named like `astrotracker-20260924-0900.db`) over `astrotracker.db` in `%LOCALAPPDATA%\AstroTracker`, copy the backup's `finals` folder into that same folder, and start. Leave out the `finals` folder and the projects' finished pictures are missing, since the database only names them. Thumbnails rebuild on demand, and your folders, rigs and grades come back as they were, provided each folder has the same path on the new PC: give the drives the same letters. Removing a folder on the Library page and adding it again under a new letter starts it from scratch, without its grades (see *Library*).

The database is checked every start

If the file is damaged — it happens, usually because of an unclean shutdown or a failing drive — a red banner says so and offers to restore the newest backup with one click, rather than letting pages fail one at a time with mysterious errors. The damaged file is kept beside the restored one rather than thrown away. The button is **Restore newest backup and close**: the program puts the backup in place and closes, and when you start it again from the Start menu it opens on the restored copy. Anything changed since that backup was made is lost, and with no backup folder set there is nothing to restore, which is the best argument for switching backups on.

Updates

A few times a day the program asks GitHub for the newest release. When one exists you get a banner and one notification.

- **Download** fetches the installer and checks two things: that it is signed by the AstroTracker publisher, and that it matches the checksum published with the release. One that fails is deleted,

never installed, and the signature is checked once more just before an installer runs. The releases always come from AstroTracker's own page on GitHub; that is not a setting anyone can change.

- **Install now** closes the program, runs the installer silently, and starts the new version. The page reloads itself when the server is back and shows What's new.
- **Automatic installs**, off by default, do the same unattended — but only outside darkness, never while a rig is imaging, and not within ten minutes of a start.
- **Channel**: stable, or beta to also receive pre-releases.

After an update, the first page load shows **What's new** — the features and fixes since the version you had, not since the beginning. The full history is behind the **What's new** button under Support in Settings.

Keeping up with a large library

Libraries grow into the hundreds of thousands of files, and the program is built for that: on a library of about 191,000 files the alerts take under a second to work out, and a change you make during a big scan, such as grading a sub, waits its turn at the database rather than failing.

When something goes wrong

A page that hits an error shows the error where the page was, with **Reload**, **Copy error** and **Save problem report…**, instead of going blank. The rest of the program is unaffected, and opening another page from the menu clears it. The same **Save problem report…** is under **Support** in Settings.

What it writes down about itself

Most problems are understood afterwards, from the log: `astrotracker.log` in the data folder. It records when each scan and each measuring run started and finished and what it did — how many files it looked at, added, updated and removed, how many subs it measured, and how long it all took.

When something hangs, the program notices. Every thirty seconds a watchdog looks for a request from a page that has run for more than twenty seconds, and for a scan or measuring run that has made no progress for ten minutes, and writes into the log where every part of the program was at that moment, which names what it is stuck on. The pages report their own errors, slow requests and freezes to the same log, and a page that froze and was then closed is written up by the next one to open, with the last click or key before it.

> **IN PLAIN TERMS**
>
> If AstroTracker ever freezes, leave it a minute before you close it. By then it has written down what it was stuck on, and the report you save once it is running again carries that.

The log looks after itself: at 5 MB it starts a fresh file and keeps the three before it. If another program holds the file open so that it cannot be set aside, that is written into the log and tried again ten minutes later.

The problem report

A problem report is a zip file: the last 2 MB of each log file, and a snapshot of the installation — the version, the Windows version and where the program is installed, your settings, library folders and rigs, how many files of each kind the library holds, the current alerts, what each part of the program was doing at that moment, and up to fifty files whose headers could not be read. There are no images in it and no database.

Save problem report… downloads it like any file from a web page. Email it to the developer with a line about what you were doing; AstroTracker never sends it anywhere by itself. Before it is written, what is nobody else's business is taken out: your licence key, the phone pairing code, your push topic, any user name and password in the all-sky camera's or push server's address, and any access code the log recorded. Your site's position is kept only to the whole degree, which still shows how dark the nights get there without pointing at your house. It is an ordinary zip, so you can open it and look before you send it.

Your phone

Two separate things: seeing the pages, and being told without looking.

Seeing the pages

Under Settings › Phone access, tick **Allow phones and other computers on this network**, press **Save settings** and restart AstroTracker; Windows then asks once whether to allow it on private networks, and you allow it. Then scan a QR code in Settings with your phone's camera: it opens the Imaging page and pairs the phone at the same time. Any device without the pairing code is refused.

Add it to your home screen and it opens like an app. Away from home, connect to your home VPN first; if you use Tailscale, tick **Also show a Tailscale code when Tailscale is running** (off by default) for a second QR code with that address.

The phone gets the real interface, not a cut-down one. Tables drop their least useful columns at phone width so the information you need to make a decision fits without sideways swiping, and text wraps normally rather than collapsing into a column of single letters. The one exception is the **How much more?** table on a project page, which needs every column: it scrolls sideways, with the filter column held in place so each row keeps its name.

Nearly everything works from a phone, including moving files, undoing a batch, archiving and quarantining, and — if you enabled it — rig control. A few things stay on the PC. Opening a folder and revealing a file both open Explorer on the PC's own screen, so from a phone they would do nothing you could see: there the rule is *where the action lands*. The rest are kept to the PC for trust instead, so that somebody who picks up your unlocked phone cannot re-pair themselves, read files off the PC or copy your database somewhere of theirs: making a **New pairing code**, restoring a backup, adding or removing library folders, and changing any setting that names a folder or a web address (the rejected-subs, backup and training-export folders, the all-sky picture's source and the push server). A phone can still save the rest of the Settings page.

Another computer on your network

Any browser on your home network can open AstroTracker the same way: a laptop, a Mac, another PC in the house. Under each QR code Settings writes out its link, headed "For another computer, open this link there once:", with a **Copy** button. Open it once in the other computer's browser and it lands on the Imaging page, paired, with the same pages and the same three exceptions.

The pairing is then kept in that browser as a cookie (a small note a browser stores for one website) for a year, after which you open the link or scan the code again; until then a bookmark to `http://<this PC's address>:8765/` is all you need. Your router hands the PC that address and may hand out a different

one after a restart. Give the PC a fixed address in the router's settings (usually called a DHCP reservation) and bookmarks and home-screen icons keep working.

> **IN PLAIN TERMS**
>
> The QR code and the link are the same thing: the code is only that address drawn as a picture, pairing code included. A phone reads it with its camera; a computer simply opens the address. Either way, treat the link like a house key. If it ever goes astray, **New pairing code** in Settings withdraws it from every device, and each one pairs again.

Push notifications

Push uses ntfy, a free open-source service. There is no account and no open port: the program publishes to a private topic and your phone subscribes to it. Tick **Send push notifications through ntfy** under Settings › Push notifications; it uses the public ntfy.sh server unless you give your own under **Server**. Which alerts reach the phone is set out in *Every alert it can send*.

The setup asks which phone you are holding, because the two genuinely differ. On Android you scan a code and it subscribes in one step. On an iPhone you cannot — the ntfy app has no link handler on iOS — so the topic is pasted instead: open Settings on the phone itself, through phone access, and the **Copy** button beside the topic puts it on the phone's clipboard. If you would rather not install an app at all, ntfy's own web page can be added to your home screen and delivers real push notifications from there (on an iPhone, iOS 16.4 or newer).

Once the settings are saved, **Send a test push** checks the whole chain; if it reports the message sent and nothing arrives, the subscription on the phone is what to look at. The topic name is a secret — anyone who knows it can read your notifications — and **New topic** replaces it if it ever leaks, after which the phone subscribes to the new one.

How a photograph is measured

Four numbers do most of the work. This chapter says what each one means physically, so that the flags on the Review page (*Review — judging your photographs*) are not magic.

Star size — FWHM and HFR

A star is a point of light infinitely far away, so in a perfect telescope on a perfect night it would be a single pixel. It never is. The atmosphere smears it, the optics spread it slightly, and the result is a small fuzzy blob a few pixels across.

FWHM — full width at half maximum — measures that blob's width at the level where it is half as bright as its centre. **HFR**, half-flux radius, measures the radius containing half the star's total light. They measure the same thing differently; HFR is less sensitive to the exact shape of the star's profile, which is why focusing routines prefer it.

Smaller is sharper. A rising FWHM through a night means focus is drifting as the tube cools, or the seeing is deteriorating, and the two can be told apart because focus drift is gradual and one-directional while seeing wanders.

Eccentricity — how round the stars are

Zero is a perfect circle; higher numbers mean more oval. Stars should be round. When they are not, the cause is almost always mechanical:

- **Ovals all in the same direction** across the frame: the mount moved during the exposure — guiding error, wind, a cable snag.
- **Ovals pointing in different directions in different corners**: optical, usually tilt in the camera or focuser.

Eccentricity is the measurement behind the **trailing** flag, and it is the one people most often miss by eye on a small preview.

Trailing also has a fixed bar, not only a comparison. On real optics the stars of perfectly good subs measure 0.4 to 0.55, because field curvature, tilt and coarse pixels all make them slightly oval, so a sub is only called trailing when its eccentricity reaches **0.75** *and*, on a night with five or more subs in that filter, sits at least **0.15** above that night's median. Both numbers are under Settings › Sub quality: lower the bar to be stricter, press **Save settings**, and **Re-flag measured subs** then applies it to everything already measured.

Background and noise

The sky is never black. **Background** is how bright it was in a given sub — rising with moonlight, with cloud lit from below by a town, with dawn, or with a neighbour's floodlight. **Noise** is the random fizz on top of it.

These two matter beyond flagging bad frames. Noise is one half of signal-to-noise, how far the faint parts of your target stand clear of that fizz, which decides whether a filter is finished; the **How much more?** panel measures it more carefully, on calibrated subs (*Signal, noise, and when a filter is done*). And because moonlight makes a sky noisier, each filter's noise on moonlit nights against dark ones shows how much of the moon it lets through (*The moon, and which filters survive it*).

Gradient

How uneven the background is across the frame — a ramp from one side to the other. Caused by light entering from one direction: the moon just out of frame, a streetlight, dawn, or thin cloud catching town light on one side of the sky. Gradients can be removed in processing, but steep ones eat into your signal and are worth knowing about.

The flags

Review turns these numbers into the flags on a sub's card. Most ask whether a number *stands out*: lies more than three times the session's usual scatter from its median. Why a sub is judged against its own session is the next section.

FLAG	WHEN
no stars	No star was found at all.
trailing	Eccentricity at the fixed bar above, and clear of the session's median by the margin.
clouds	The background stands out above the session's and the stars drop: under half the session's typical count, or clearly below it.
few stars	Under half the session's typical star count, without the brighter background that would make it clouds.
bright sky	The background stands out and is at least 10% brighter than the session's, with the stars still there.
soft	Star size stands out and is at least 15% larger than the session's.
elongated	Eccentricity stands out by at least 0.06, short of trailing.
gradient	The gradient stands out, by at least 2% of the background.

Why comparison beats thresholds

None of these numbers has a universally good value. Four-pixel stars are excellent through one telescope and mediocre through another. A background of 700 is dark under one sky and bright under another.

So AstroTracker compares each sub against the median of its own session — same target, same filter, same night — and flags what stands out from its neighbours. This makes the flags meaningful on any equipment, at any site, without calibration or configuration. It also has a consequence worth knowing: a night that was uniformly poor produces few flags, because nothing in it is anomalous. The flags find the *changes*. They need five subs of a filter in a night to compare; a shorter session gets only the fixed checks: a sub with no stars at all, and the trailing bar.

How good a night was overall is answered elsewhere: how sharp, by the seeing estimate (*Seeing, guessed from your own stars*); whether a project is deep enough yet, by How much more?; and a night that was hazy from start to finish, by the **star glow** and the **cloud likelihood**, which also compare a sub with the project's other nights in that filter (*Finding cloud in your own subs*).

Finding cloud in your own subs

Review puts three numbers on every measured sub that are about cloud and nothing else: **Star glow**, **Sky change** and **Cloud likely**. This chapter says what each one reads, why it is built that way, and how far to trust it.

Thick cloud is easy to find: the stars vanish, the background jumps, and the **clouds** flag in *Review — judging your photographs* catches it. Thin cloud is the problem. The frame looks normal and the stars are no wider than on a clear night, but in the finished stack it becomes a soft bloom around the bright stars that is very hard to process away. People who grade by eye spot it two ways: a halo around the bright stars, and a background that swims when the subs are played one after another. The first two numbers measure exactly those; the third comes from a model that weighs everything, those two included. All three work from the subs alone, with or without an all-sky camera.

Star glow: the halo around the bright stars

Thin cloud is a haze of ice crystals or droplets that scatters starlight slightly off its path, so a bright star seen through it picks up a faint halo — an **aureole** — spreading several arc-minutes out. (An arc-minute is a sixtieth of a degree; the full moon is about thirty across.) The star's core hardly changes, which is why star size (FWHM) cannot see thin cloud.

So in each sub the program takes the four brightest stars — saturated ones included, because here the bright ones are the point — and reads the typical brightness in rings around each:

RING	DISTANCE FROM THE STAR	WHAT IT HOLDS
Close in	0.5 to 1.5 arc-minutes	The telescope's own scattered light: always the same share of the star, whatever the sky does.
Near ring	3 to 9 arc-minutes	Where a cloud's halo is strongest through a long focal length.
Far ring	9 to 25 arc-minutes	Where it peaks on a wide field. It does not fit a small field, and is left out there.
Sky	30 to 40 arc-minutes	The sky beyond the halo. Where it fits in the frame the rings are read above it, not above the whole frame's sky, which varies across a wide field.

The distances are angles on the sky, turned into pixels from the focal length and pixel size in the header, because the halo belongs to the cloud and not to the telescope. A header without both numbers gives no glow.

> **WHY IT WORKS THIS WAY**
>
> A ring holds two lights: the telescope's own scatter, always the same share of the star, and the cloud's halo. Read on its own, a ring gets *darker* under thick cloud, which dims the star, and brighter as a rising target brightens it, and neither of those is haze. Divided by the light close in, which is all telescope scatter, it leaves the *shape* of the halo, and only the air changes that.

A shape means something only against the same star under a clear sky. The program finds that star in every sub by its position (mirrored after a meridian flip, when the frame turns round) and takes its clear-sky shape from the low end of what it showed in subs that kept at least 80% of their night's typical star count, twice: over that night, and over every night of the project in that filter. Where a star was seen over a wide enough span of altitude, its clear-sky halo is allowed to grow towards the horizon, because clear air scatters light too and low down the light crosses more of it. A sub's glow is its stars' outer light over what clear sky predicts, added up so the brightest star carries it and three faint ones cannot outvote it, and Review shows the largest answer: either ring, against either comparison. **1.0 is as clear as that target has looked; 1.5× is half as much light again in the halo.**

> **WHY IT WORKS THIS WAY**
>
> A night compared only with itself cannot see haze that lasted all night: its least hazy subs become its idea of clear. A bright star's halo on the good subs of other nights holds steady, which is why the project's history is the second comparison.

A star whose clear-sky halo is fainter than 1% of the sky is reading background and is ignored. A field with no bright enough star, common through a small field of view, shows **no bright star**.

Sky change: the background that moves

While each sub is measured, the program shrinks the frame four pixels by four, cuts it into a grid of 24 by 36 blocks and takes the **median** — the middle value — of each. A star covers only a few pixels of a block, so the stars vanish and a small map of the sky's brightness remains, divided by its own median so only the pattern is left.

That map is divided by the map of the sub before and of the sub after in the same filter, if they were taken within 30 minutes. A steady sky gives a flat result. The best-fitting smooth tilt is removed, because the moon moving or a town's glow turning with the sky changes it as an even ramp, not in lumps. The photon noise goes too: light arrives in random packets, so any two frames differ slightly, and how much is worked out from the sub's own noise and brightness, which puts cameras and filters on one scale. What remains is how lumpy the change was, as a percentage of the sky. The sub's figure is the *smaller* of its two comparisons: a cloudy sub between clear ones differs from both, a clear sub beside a cloudy one from only one.

With one filter's subs back to back, a clear sky reads a few hundredths of a per cent, and cloud drifting through from about a tenth up. But no level counts as bad: a filter wheel cycling four filters can leave thirteen minutes between subs of one filter, the sky drifts in that time, and one project's good subs can read as high as another's cloudy ones. So it is for sorting, and it feeds the cloud likelihood. The neighbours are the comparison, not the night's typical map, because on a mostly cloudy night the typical map is cloudy and the clean subs would become the odd ones out.

Cloud likely: a model that weighs everything

The third number comes from a machine-learning model, a program that learns a rule from examples instead of being given one: 300 small decision trees, each asking up to four yes-or-no questions about a sub's numbers, their answers added into one score. It ships inside the program and runs on your computer with no internet and no extra software. It reads no pixels, only numbers already stored, so a whole library is scored in seconds. Its inputs are nearly all comparisons:

INPUT	WHAT IT ASKS
Against expectation	Star count, star size, background, noise and gradient, against what the night, and the project's subs of that filter and sub length, say they should have been for where the target, moon and sun stood.
Against neighbours	How star count, background, noise and gradient changed from the sub before and after, and how restless they were over the five around it, each on its own filter's scale so a cycling filter wheel is not read as change.
Glow and sky change	Both, as above.

"What it should have been" matters. Low down, the light crosses more air and more of a town's glow, so the sky is brighter and the stars fewer; the moon and twilight brighten it too. None of that is cloud. So each measurement is fitted, through the night's subs and through all the project's subs of that filter and sub length, against how high the target stood, the moon's height while it is up and how close the sun is to rising or setting. Each line may slope only the physical way, and subs that stray far from it are set aside. On a night when the target climbed from 25° to 81° with the moon down, the first hour looked like cloud against the night's median and sat within 2% of expected against the fit. This needs the latitude and longitude of the rig's site (Settings › Observing sites) and the target's position in the headers; without them, or with fewer than eight subs to fit, the plain median is used.

> **WHY IT WORKS THIS WAY**
>
> The target's altitude is not an input of its own. Given it, the model learned a habit of the person who graded its examples, who seldom passed a low sub as good, and read "low" as "cloud". Altitude belongs in what a sub should have shown, not in the verdict.

The model learned from 3,355 subs graded by hand in the developer's own library, 2,431 good and 924 rejected for cloud, from four telescopes. Your own grades do not retrain it. It was tested on whole nights held back from training, so its figures below are for nights it had never seen. The number **ranks** subs well but is not a calibrated probability: cloudy subs were weighted up in training because there were fewer of them, so 73% does not mean 73 in 100 such subs were cloudy. It sorts and marks subs in Review and never grades, moves or alerts on its own. After an update every measured project is rescored at the first start, in seconds, without reading pixels or touching your grades.

> **IN PLAIN TERMS**
>
> The likelihood is a second opinion from something that has seen thousands of graded subs and knows a low target or a bright moon is not cloud. It is good at putting the doubtful subs first. It is not a judge, and it decides nothing unless you ask it to.

Why two numbers, and why auto-grade takes either

The model already uses the glow; the two stay apart because they fail differently. **The glow measures the sky** but cannot see thick cloud, where the star and its halo dim together, or any field without a bright star. **The model copies a person**, and is only as good as the grades it learned from: five subs once passed as good stood out on the glow, while the model, having learned subs like them as good, put none of them as high as 70%; looked at again, all five were bloated.

So the auto-grade test feature takes an ungraded sub when *either* number reaches the level picked, each likelihood paired with a glow that good subs rarely reach; the sky change counts through the model. Its three levels, and how it is used and undone, are in *Review — judging your photographs*.

What the numbers caught

Measured on the library the model came from. **Caught** is the share of subs rejected for cloud that reach the mark; **false alarms** the share of good subs that do.

MARK	ON THE CARD	CAUGHT	FALSE ALARMS
Star glow 1.4×	chip	54%	5.6%
Star glow 1.6×	red chip	49%	1.6%
Star glow 1.8×	red chip	44%	0.8%
Cloud likely 50%	chip	86%	6%
Cloud likely 90%	red chip	70%	1.8%

The glow catches about half of all cloud with few false alarms: thin cloud round a bright star. The other half — thick cloud, fields with no bright star, cloud that left the halos clean but moved the background

— is what the likelihood and the sky change are for.

STATED HONESTLY

- The subs beside a cloudy one, and a night's first sub, can score higher than they deserve: the model reads change from the neighbours.
- A project's first night has no history, so haze that lasted all night can pass.
- Dew or frost on the camera's window is not told apart from cloud, and neither is a lit building or tree in the way; the **below horizon** mark in Review, from the horizon you draw under Settings › Observing sites, helps with the second.
- Old projects with no grades have been seen to score high in bulk, and nobody has yet checked why.
- How well a model trained on one library carries over to other equipment is unproven. That is why the likelihood only sorts and marks, and auto-grade is a test feature, off by default: try it on a project you can stack with and without its rejects before trusting it.

Signal, noise, and when a filter is done

The arithmetic behind "is another night worth it?", which is the most consequential question the program answers.

Not all hours are equal

Signal-to-noise improves with the square root of integration time. Four hours are twice as good as one; sixteen hours are four times as good as one. This is why the second half of a project always feels less rewarding than the first.

But hours on their own say nothing about whether a picture is finished. Sixty hours on a faint shell of OIII can still look noisy, while five hours on a bright reflection nebula is already smooth. The target, the filter, the telescope, the sub length and the sky — moon, haze, light pollution — all change what an hour is worth. Two hours under a bright moon are genuinely not two hours.

What decides whether a filter is done

When you look at a finished picture and call it noisy or clean, you are judging the faintest parts you care about — the outer wisps of a nebula, the dust around it, the halo of a galaxy. The bright middle is smooth long before they are. What you are seeing is their **signal-to-noise ratio**, or SNR: how far the faint structure stands above the random fizz on top of it. At an SNR of one it is lost in the fizz; the higher it climbs, the smoother it looks.

The scale you look at matters too. Nobody judges a picture one pixel at a time: the eye takes in small patches, and averaging a patch of 3 × 3 pixels cuts the fizz to a third. So what decides whether a filter is done is the SNR of the faint structure, at the scale you look at it. That is what the **How much more?** panel on a project page measures.

Read from a handful of subs

Measuring the SNR of a stack sounds as if it needs a stack. It does not. Press **Work it out** and the program takes five single subs from each of a filter's **exposure groups** (its subs at one sub length, gain and offset), spread across the project's nights, and calibrates each with darks, bias and flats that match it. From each calibrated sub it reads two numbers:

- **The signal**: how far the faint parts of the target stand above the darkest sky in the frame, in that one sub.
- **The noise**: how much the fizz in that one sub scatters from pixel to pixel.

Stacking adds the signal up faster than the noise, by exactly the square-root law above, so the stack's SNR follows from those two numbers and a count of the subs, with the noise taken over the patch you

judge by:

```
SNR of the stack = signal per sub × √(number of subs) ÷ noise per sub
```

The useful property is what this leaves out. The size and speed of the telescope, how dark your sky is, the moon, haze, the filter and the sub length all shape those two numbers, so they are all inside the measurement and none of them has to be modelled. You enter no camera or telescope details, and a brighter sky counts the way it shows in the picture: as more noise.

"The faint parts" has a precise meaning. The stars are set aside, the frame is smoothed so that structure is read rather than grain, and the darkest sky in the frame is taken as the floor. Of what rises above that floor, the program takes the level that nine-tenths of the judged area reaches: the faint outskirts of the target rather than its bright middle.

> **IN PLAIN TERMS**
>
> Five photographs are enough to know how bright the faint parts of a target are and how grainy a single photograph is. From those two facts and a count of how many photographs you have, the program can say how grainy the faint parts of the finished stack will be, without stacking anything.

> **WHY IT WORKS THIS WAY**
>
> It was checked against finished work before it was trusted: forecasts from five subs came within a few percent of the noise the finished masters (the stacked results) actually measured. Over the whole frame, the default bands below sorted five of six projects the way their imager had already judged them by eye. The sixth was the Andromeda galaxy, which its imager calls done while accepting a faint halo noisier than the galaxy's body; over the whole frame the measurement called it unfinished. That is why you can draw a box.

Where it looks

By default the whole frame is judged, less a thin border at the edges where flats correct least well, because the whole frame is what you look at. A box you draw with **Where to judge it** (see *Inside one project*) takes the frame's place from the next measurement and stays with the project until you go back to the whole frame.

It deliberately does not use a catalogue's size for the target. Catalogues often list only a nebula's bright core, and a box on the core reads every filter as clean when the picture around it is not.

Three bands, set to your eye

Each filter gets a verdict from its SNR:

VERDICT	SNR OF THE FAINT PARTS	WHAT IT MEANS
needs more	below 4	The master still looks noisy in the faint parts.
almost there	4 to 6	Close, but not yet clean.
clean	6 and up	Clean with little noise reduction.

The two numbers and the size of the patch are settings, under **Settings › How much more?**. The default patch, 3 × 3 pixels, is close to how you see a drizzled master (drizzling is a way of stacking that enlarges the result) after light noise reduction; a single pixel is the strictest choice. They are settings because they are one imager's eye, not a textbook's: the defaults come from one imager's finished and unfinished projects. The surest way to set yours is to measure a few projects you have already judged and move the numbers until the verdicts agree with you.

When to stop

Because SNR grows with the square root of time, the program can say what more of the same subs would buy. The total hours at which a filter reaches a band are the hours you have now multiplied by the square of (band ÷ SNR now). With the default bands:

- A filter at SNR 3 after five hours reaches 4, **almost there**, at $5 \times (4 \div 3)^2$: about nine hours in total.
- It reaches 6, **clean**, at $5 \times (6 \div 3)^2$ = twenty hours. Doubling the SNR takes four times the time.
- One at SNR 1.5 after five hours would need eighty hours to reach clean, which is further out than the program will stand behind.

These are the panel's **Next band** and **Clean** figures (see *Inside one project*). The forecast assumes more subs like the ones you have. When it comes to more than four times the hours you already have, it is shown with a "greater than" sign, as in `> 80 h`: that far beyond what was measured, the number is a direction rather than a promise.

The program gives its advice in these terms: a band, and the hours to the next one. It is not telling you to stop; it is telling you what the next band would cost, so the decision is informed rather than a feeling.

What else it notices

When there is more to say about a filter, its row has an **i** button with notes: how faint the target is in that filter compared with another, whether its subs are long enough, and whether it was measured without its darks or flats.

Some filters see less of the target

Filters differ not only in their noise but in the faint signal itself. Dust that shows three times fainter in green than in luminance (which passes every colour at once) needs nine times the hours in green for

the same result: the square-root law again. The program compares how bright the faint parts are per second of exposure in each filter, and when one is at least one and a half times fainter than the brightest, the note says how many times fainter and how many times the hours it needs. When almost nothing shows in a filter, twenty times fainter or more, it says the hours forecast is not a useful target for that filter.

Whether your subs are long enough

Every time the camera reads out a picture it adds a fixed amount of electronic noise, the **read noise**, however long the exposure. The sky's own noise grows with the exposure, so in a long enough sub it swamps the read noise, and the sub is called **sky-limited**. The program measures the read noise on a bias frame (or a dark, when there is no bias) and works out its share of the noise in the sub length carrying most of the filter's hours. Noises combine by their squares, so the share is of the squared noise:

- A tenth or less: "sky-limited: longer subs would gain nothing".
- Three-tenths or more: "read noise matters: longer subs would help".
- In between: "read noise is a small part of the noise; sub length is fine".

Several sub lengths in one filter

If you shot one filter at more than one sub length, gain or offset — 300-second subs one season and 600-second the next — each exposure group is measured on its own, and the groups are combined the way stacking software combines them, weighting each sub by how clean it is. A noisier group can never make the result worse, only add less. **Show exposure groups** on the panel lists each group's own figures (see *Inside one project*).

What it needs, and what it cannot tell you

It reads real pixels, so it runs only when you ask, in the background, and most of its time goes on reading files. It needs:

- **The subs themselves**, on a drive the computer can reach: at least two of each group, or that group is not measured. Subs you have marked bad are left out, and a sub stored in two folders counts once.
- **Calibration that matches.** Darks come from anywhere in the library, matched on camera, exposure, gain and offset and closest in temperature; bias on camera, gain and offset; flats through the same filter from the project's own folders first and the rig's otherwise, closest in date. Four of each are combined into a quick master. If the sub length carrying most of a filter's hours had no darks or no flats, the note says so, as in "measured without flats: treat the number as rough".

STATED HONESTLY

Until the bands are set to your eye, its verdicts are someone else's. Its forecasts assume more subs like the ones you have. And it measures the faint parts; it does not judge the picture. A halo you are happy to leave noisy is your call, which is what the box is for.

The moon, and which filters survive it

The most common planning question in the hobby, usually answered with folklore.

The folklore is: narrowband under the moon, broadband when it is dark. It is not wrong, but it is coarse. It does not tell you whether tonight's particular moon, at that particular separation from your particular target, is bad enough to matter — or which of your three narrowband filters suffers most.

How the sky brightness is computed

AstroTracker uses an established model of moonlit sky brightness (Krisciunas and Schaefer, 1991) which takes the moon's phase, how high it is, how far it is from your target, how much atmosphere each is seen through, and a typical figure for how much the air dims light. That yields how bright the sky is at your target's position, at each moment of the night.

Your dark sky comes from the Bortle class you set for the site. The ratio of the two is how many times brighter the moon makes your sky in that direction, right now — and it changes through the night as the moon rises and sets and as your target climbs and sinks. The site is your home site, except on the Planning page's Tonight panel, where a telescope at another site with a Bortle class of its own is judged under that one.

Each filter feels a different share

Moonlight is reflected sunlight — a continuous spread of all colours. A filter only admits the colours in its window, so how much moonlight it lets in depends on how wide that window is and where it sits.

- **Broadband** filters admit most of it. Blue suffers worst, because the atmosphere scatters blue light most; red suffers least.
- **Narrowband** filters admit far less, because a 3 nm window is a tiny slice of a continuum: about a twentieth of what luminance takes in for Ha and SII, about a seventh for OIII.
- Among the narrowband three, **OIII** fares worst — it sits at 500 nm, where scattering is strongest — and Ha and SII, further into the red, fare better. A wider band admits more.

Those are the starting values. They are then **refitted from your own subs**: every measured sub carries its background noise rate and the moon conditions at the time it was taken, so the program can work out the share each of *your* filters actually feels at *your* site. Filters differ between manufacturers and light pollution differs between gardens, and after enough subs under enough moons the numbers are yours rather than a textbook's.

Enough means at least twenty measured subs through a filter taken with the moon adding under 5% to the sky, and twenty with it adding more than 30%; until a filter has both, it keeps its starting value. The

fit is redone once a day and kept between a fifth of and five times the starting value, so a few odd subs cannot turn it upside down.

The program knows a filter by its name alone: luminance (a clear or UV/IR-cut filter counts as luminance), red, green, blue, Ha, OIII and SII. Any other gets no verdict, which usually includes a dual-band filter under its brand name; one named, say, Ha-OIII is judged as Ha.

What you see

For each filter the result is a **yield**: what an hour through that filter tonight is worth, as a share of the same hour under a moonless sky, averaged over the time your target is up in the dark. At 100% the moon costs you nothing; at 50%, two hours tonight gather what one would on a dark night. The yield becomes a verdict per filter — **best**, **ok** with its percentage, or **wait** — at these lines:

FILTER	BEST AT	OK DOWN TO
Narrowband (Ha, OIII, SII)	60% or more	35%
Red, green, blue	60% or more	40%
Luminance	80% or more	None: below 80% it says wait

> **WHY IT WORKS THIS WAY**
>
> Luminance is held to a higher bar because moonlight leaves a gradient across it that costs more to clean up afterwards than the hours it saves. It is either a good night for luminance or a night to wait for a darker one.

You meet the verdicts on the **tonight** badge of the Active page's project cards, on the **Best tonight** tile of a project's page (hold the pointer over it for each filter's percentage, and which filters have been learned from your subs), and on each target's line in the Planning page's **Tonight** panel (see *Active*). When the moon stays below the horizon while the target is up in the dark, the line reads `moon down: any filter`.

The same yields, worked out for whichever night you pick, are behind how much the moon costs each recommended target and what **Worth another visit** says about each filter you asked for (see *Planning*).

Best use of tonight

Under the Tonight panel's list, **Best use of tonight** puts your projects in order of what the night can give them. For each, it takes the filter with the highest yield (and any others within five points of it) and multiplies the hours the target is up in the dark by that yield, giving **effective hours**. Four hours up at a 50% yield is two effective hours: as much data as two hours on a moonless night. Up to six projects

are listed, highest first, and a project whose every filter says wait is left off, so the top of the list is simply the best use of the night.

Effective hours rank tonight's choices; they do not say whether a project needs more data at all. **How much more?** on the project's page answers that by measuring the picture itself (see *Signal, noise, and when a filter is done*).

Seeing, guessed from your own stars

Proper seeing monitors are dedicated instruments costing thousands. AstroTracker estimates it from data you are already producing.

The logic: a star's measured width is the combination of several things added together — the atmosphere's blurring, the telescope's own optical spread, the coarseness of the pixel grid, half of the guiding error (guiding smears stars too), and the extra atmosphere you look through when pointing low rather than overhead.

Every one of those except the atmosphere is either known or measured. The focal length and pixel size give the pixel scale — how many arc-seconds of sky each pixel covers — and with it how much the pixel grid itself widens a star. The optical limit (the diffraction limit) comes from the telescope's aperture, the width of its main lens or mirror: the wider it is, the smaller the finest point of light it can form. For a 127 mm telescope that is about 0.9 arc-seconds, worked out for the green light at the middle of the visible range.

The aperture is the one number your photographs do not carry. It comes from the telescope you entered on the Hardware page, or from the **Aperture mm** box when you edit a rig on the Rigs page. A rig without one has nothing taken off for its optics, so its estimate reads worse than it should — most noticeably on a small telescope, whose optical limit is the largest.

The PHD2 log gives the guiding error. Only half of it is charged, because part of PHD2's number is the atmosphere making the guide star dance, and that is already in the star width. It is also never allowed to account for more than 30% of the total: a guiding figure that large beside small stars is noise from a faint guide star, not real tracking error, which would show in the stars. The target's altitude is computed. Subtract them all from the measured star width and what remains is the atmosphere — converted to arc-seconds and corrected to what it would be overhead, which is how seeing is conventionally quoted.

Subs flagged for cloud or trailing, or with too few stars, and subs graded bad are left out, since their star size says nothing about the sky. While you image, the estimate follows each telescope's latest subs; for a past night it uses the whole night. Where several telescopes are imaging at once their estimates are combined from the better end, because everything not taken out can only make a telescope read worse, never better. A telescope with coarse pixels, over 2 arc-seconds each, is left out when a finer one is imaging; when only such telescopes are imaging, the estimate is marked rough.

It is reported as a number and a band — excellent up to 1.5 arc-seconds, good to 2.5, average to 3.5, poor to 4.5, bad beyond — on the Weather page, the Imaging header and each rig card, in the Planning page's Tonight panel while imaging, in past sessions, and in the morning summary.

STATED HONESTLY

This is an estimate, not a measurement. Focus and filter are not removed from it, so a soft focus reads as worse seeing than it really is, and a narrowband filter reads as slightly better. It is reliable for comparing one night against another on the same equipment, which is what it is for. It is not a calibrated instrument reading.

What leaves your computer

A complete list, because "local-first" should be verifiable rather than a slogan.

WHAT	WHERE	CARRIES	OPTIONAL?
Weather forecast	Open-Meteo	Your coordinates. No account, no key.	Works without it; the page says it is offline.
Sky survey images	CDS hips2fits	The coordinates of the patch you are framing. Cached on disk afterwards.	Only while a framing preview is on screen: on a project once you press **Framing preview**, on the New project form once you choose a target, and in a target's details on the Planning page when you have a rig to frame it.
Target catalogues	GitHub (OpenNGC) and CDS VizieR, or VizieR's Harvard mirror	Nothing but a request for the published catalogue files, which names the version you run.	Only when you press Download catalogues.
Target name lookup	CDS Sesame	The name you typed, when the downloaded catalogues do not know it.	Only for a name the catalogues do not have.
Update check	GitHub	A request for the latest release number, which names the version you run, and, when you choose to install one, the installer and its checksum.	Can be switched off. Installing waits for you unless you tick Install updates automatically.
Push notifications	ntfy	Short text messages, to a private topic. No images.	Off until you enable it.

That is the entire list of what goes out to the internet. The pages load nothing from elsewhere: their typefaces ship with the program. The astronomy library inside it never fetches anything by itself either; it works from the tables it was installed with. And a problem report is not on the list because the program never sends one: you save it, and email it if you want help. A link such as **AstroBin** opens

another website in your browser only when you click it. Your images are never uploaded. There is no account, no telemetry and no analytics. The database never leaves your PC except to the backup folder you nominate, which is usually your own NAS.

The licence check is not on that list, because it does not leave your computer either. A key is verified offline against a public key built into the program. The trial's bookkeeping stays on your machine too: the date you first ran AstroTracker, in the database and again in the Windows registry at `HKCU\Software\AstroTracker` so that reinstalling does not silently restart a trial; the latest time the program has seen on your PC's clock, so that winding the clock back cannot lengthen the trial; and, if the trial runs out while a rig is imaging, a note to hold the lock off until midday. None of it identifies you.

What a problem report carries

A problem report is a zip file: the end of each log file, which includes the pages' own notes of errors, slow requests and freezes, and a snapshot of the program — your settings, library folders, rigs and alerts, and what each part was doing. No images and no database. **Save problem report…** in Settings › Support writes it, and you email it yourself; the program never sends it. *Backups, updates and staying alive* lists everything in it.

> **WHAT IS TAKEN OUT FIRST**
> Your licence key, the phone pairing code, your push topic, any user name and password in a camera's or server's address, and any access code in the log. Your site's position stays only to the whole degree. Your site's name and your folder names do stay, because they are what make a report readable; open the zip and look before you send it.

On your own network

A few more connections normally stay inside your own network, and each exists only if you set it up. When the **All-sky picture folder or URL** setting holds a web address rather than a folder, AstroTracker fetches the camera's newest picture from that address. A rig with an **SGP API host** on the Rigs page is asked every five seconds what its camera, filter wheel, focuser and mount are doing, and is sent a command only when that rig's **Allow control** is on and you send one yourself (see *Imaging — the live view*). And with phone access switched on, your phone talks straight to this PC over your own Wi-Fi, or through Tailscale if you use it (see *Your phone*).

Going the other way, only AstroTracker's own pages can use it. Any other website open in your browser is refused if it tries to send the program a command or read its settings, and so is a site that tries to pass itself off as this PC.

What it will not do to your files

- It never modifies a FITS file. Not one byte, not ever.
- It never writes image data into your library: no cleaned copies, no new image files, only the moves you confirm.
- It never deletes anything — not even duplicates it has identified for you.
- It never moves a file without showing you first, and every move is logged and undoable.
- It does not lose your work when you move files yourself: a sub moved in Explorer keeps its grade, reject reasons and measurements if it turns up in your library again within 120 days (see *The life of a project*).
- It will not move a file outside the library roots you configured.
- Read-only mode disables file movement entirely.

What it does not do

An honest boundary, so nobody arrives expecting the wrong tool.

It does not process images

AstroTracker does not stack, align or edit anything, and it writes no calibrated subs. That is the job of PixInsight, Siril, APP or DeepSkyStacker, and they do it well. AstroTracker tells you which frames to feed them and whether you have the calibration to do it properly. The only calibrating it does is in memory, on the few subs **How much more?** measures.

Nor does it clean subs: it never removes a satellite or aeroplane trail or saves a corrected copy of anything. Your stacking software's pixel rejection takes out most trails.

It does not run your telescope

It does not plan sequences, slew, focus or guide. SGP and PHD2 do that. AstroTracker reads what they did. The optional rig actions are deliberate one-off commands — park, abort, set the cooler — for when something has gone wrong and you are not in the same building. It is not an observatory control system, and it cannot even stop a sequence, because the interface it talks to has no such command.

It is Windows-only, on a PC that can open the files

It reads your files through Windows, from the PC's own drives or from network shares such as a NAS, so it runs on a PC that can open them. Phones and other computers reach it through a browser on your network, which covers most of what people actually want.

It is one person's software

Written for one library and then generalised. It has been tested hard against a real 141,000-frame collection, four telescopes and a network archive — which is more than many hobby tools get — but it is not a company product with a support desk. There is a problem-report button that reaches the developer directly.

It does not grade your subs for you

Review puts numbers on every measured light frame — how sharp and round the stars are, how bright and even the background is, **Star glow**, **Sky change** and **Cloud likely** — and can sort the doubtful subs first, but none of them rejects, moves or grades anything by itself. The one thing that acts on them is a test feature, off until you switch it on under Settings › Sub quality: an **Auto-grade (test)** button on a project's Review page that rejects as clouds the ungraded subs the cloud model is surest about. It acts only when you press it, never touches a sub you graded, and one click takes its grades back (see *Review — judging your photographs*).

TRAINED ON ONE LIBRARY

The cloud model learned from the graded subs of the library AstroTracker was written for. Your own grades are kept, with every reason, but nothing inside the app learns from them. On other cameras, telescopes and skies, treat its number as a suggestion to check against the subs themselves, and try the auto-grade on a project you can stack with and without its rejects before you rely on it.

Things it does not yet know how to do

- **Judge subs the way you would.** Only cloud has a model behind it. Trailing, soft and oval stars are flagged against the rest of their night, but nothing grades them for you. Nothing looks for aeroplane or satellite trails at all: you mark those yourself with the **airplane / satellite** reason. And it cannot tell dew or frost on the camera's sensor window from cloud; dew leaves a blotch that stays put in the frame and grows through the night, and nothing measures that yet.
- **Analyse guiding in depth.** It reads the guide logs for events and RMS, but the deeper analysis a dedicated log viewer offers — periodic error, backlash, polar alignment drift, whether a mount is behaving differently than it did a year ago — is designed and waiting on data from more than one kind of mount.
- **See weather it was not told about.** The rain alert is a forecast; the cloud detector needs a camera and reads thin high cloud late.

WHY THIS LIST IS HERE

Every one of these is a thing the program could plausibly claim to do badly. Saying "collecting" for a season, or "this is an estimate, not a measurement", is the same principle as refusing to fit a focuser slope through four data points. A tool you can trust about small things is a tool you can trust about large ones.

Glossary

TERM	MEANING
Airmass	How much atmosphere you look through, counted in thicknesses of the air straight overhead: 1 overhead, about 2 at 30° up. Low down the sky is brighter and fewer stars show even with no cloud, and the program allows for that when it looks for cloud.
All-sky camera	A cheap camera with a fisheye lens pointed straight up, photographing the whole sky about once a minute. AstroTracker reads its pictures to tell cloud from clear sky.
Arc-second	1/3600 of a degree. The full moon is about 1,800 arc-seconds wide. Guiding errors and seeing are measured in these.
Aureole	The halo of light that thin cloud scatters around a star, spreading well beyond the star itself. It is the bloom that is hard to process out of a stack, and what star glow measures.
Auto-grade	A test feature, off until you switch it on under Settings › Sub quality, that rejects as clouds your ungraded subs whose cloud likelihood or star glow reaches the level you pick. Its grades are marked "auto" and one click takes them all back.
Bias frame	A zero-length exposure with the lens capped, measuring the camera's fixed electrical offset.
Binning	Combining neighbouring pixels into one, trading resolution for sensitivity. Part of a rig's fingerprint.
Bortle class	A 1–9 scale of sky darkness. 1 is a remote desert, 9 a city centre.
Calibrated sub	A sub with its dark subtracted and divided by its flat, which takes out the sensor's heat glow, dust shadows and dark corners and leaves the sky and the target. **How much more?** measures calibrated subs.
Calibration frames	Darks, bias and flats — the frames that measure the equipment's defects so they can be removed.
Cloud likelihood	A trained model's estimate that a sub is cloudy, shown in per cent on Review as **Cloud likely**. A suggestion: it sorts and marks subs, and never rejects or moves one by itself.
Dark frame	A capped exposure of the same length and temperature as the lights, measuring sensor heat glow and hot pixels.
Dark run	One batch of darks shot on one night, by one camera body. The unit AstroTracker organises darks into.

Eccentricity	How oval a star is. 0 is a circle. High values mean tracking error or optical tilt.
Effective hours	Clock hours multiplied by the share of their worth a filter keeps under tonight's moon (its yield): under a bright moon an hour of narrowband keeps most of its worth, an hour of luminance little of it. Shown on the Planning page's **Best use of tonight** line.
Exposure group	A filter's subs taken at one exposure length, gain and offset. **How much more?** measures each group and combines them the way stacking software would.
Faint structure	The dim outer parts of a target: the edges of a nebula, the outer halo of a galaxy. Their noise decides whether a finished stack looks clean, so they are what **How much more?** measures.
Fingerprint	Focal length + pixel size + image dimensions + binning: how AstroTracker identifies which rig took a frame.
FITS	The astronomy image format. Every file carries a text header describing how it was taken.
Flat frame	An exposure of an evenly lit surface, measuring dust shadows and corner darkening. Must match the optical setup exactly.
FWHM	Full width at half maximum — how wide a star appears. The main sharpness measure. Smaller is better.
Gain / offset	Camera sensitivity settings. Calibration frames are only valid for lights with matching values.
Guiding	A second camera watching one star and correcting the mount many times a minute.
HFR	Half-flux radius — a second sharpness measure, less sensitive to star profile shape.
Header	The text label inside a FITS file: target, filter, exposure, temperature, date, optics.
Horizon profile	Your real horizon of trees, roofs and hills, entered under Settings › Observing sites as bearing:altitude pairs (180:15 is 15° up due south) or one number all the way round. A target counts as up only above it, and Review marks subs taken below it.
Integration	Total exposure time across every sub of a project.
Library root	A top-level folder you have told AstroTracker about, with a role.
Light frame	A sub of the actual target, as opposed to a calibration frame.
Meridian flip	The mount swinging to the other side of the pier as the target crosses the meridian: due south from the northern hemisphere, due north from the southern. A common failure point.

Narrowband	Filters admitting only one emission line — Ha, OIII, SII — blocking almost all moonlight and streetlight.
Observing night	A night dated by its evening, so 1 a.m. on the 14th belongs to the night of the 13th.
PHD2	The standard free guiding program. Its log is one of AstroTracker's two main log sources.
Plate solve	Identifying exactly where an image points by matching its star pattern against a catalogue.
Project	One target on one rig. The unit of work AstroTracker tracks.
Read noise	The camera's own electronic noise, added once each time a sub is read out, whatever its length. For the same total time, fewer and longer subs carry less of it.
Rig	One combination of mount, telescope and camera.
RMS	The measure of guiding error, in arc-seconds. Under about 1″ is good.
Seeing	Atmospheric blurring — the twinkle. Weather, not equipment.
SGP	Sequence Generator Pro, which runs the night. Its log is AstroTracker's other main source.
Sky change	How much a sub's background moved against the subs either side of it in the same filter, in per cent of the sky, with smooth gradients such as the moon's taken out. Cloud drifting through is lumpy and raises it.
Sky-limited	A sub whose noise comes almost entirely from the sky's own light rather than from read noise. Longer subs gain nothing then; **How much more?** says when a filter's subs have got there.
SNR	Signal-to-noise ratio: how far the signal stands above the noise. It grows with the square root of exposure time, so four times the hours doubles it. **How much more?** reads it in a target's faint structure.
Stacking	Combining many subs into one image so that noise cancels and signal accumulates.
Star glow	How much light thin cloud has spread around a sub's brightest stars (the aureole), as a multiple of the clearest this target has looked in that filter; 1.0 is clear.
Sub	One exposure. Short for sub-exposure.
Timelapse	A set of pictures played one after another like a film. Review plays a project's subs this way, and the all-sky camera's night can be played the same way, with the cloud reading drawn under it.

Transparency	How much light the sky lets through. Reduced by haze and thin cloud.

This document describes AstroTracker 0.23.3. Features are added regularly; the version history inside the program, behind the **What's new** button under Settings › Support, is always current.