Rubin Observatory Scheduler Introduction

Peter Yoachim

University of Washington

For the August 2020 Project and Community Workshop Survey Strategies Session

Rubin Observatory (formerly LSST)

- 8m telescope on Cerro Pachon in Chile
- 3.5 gigapixel camera
- Real-time alert stream
- “making a movie of the sky”

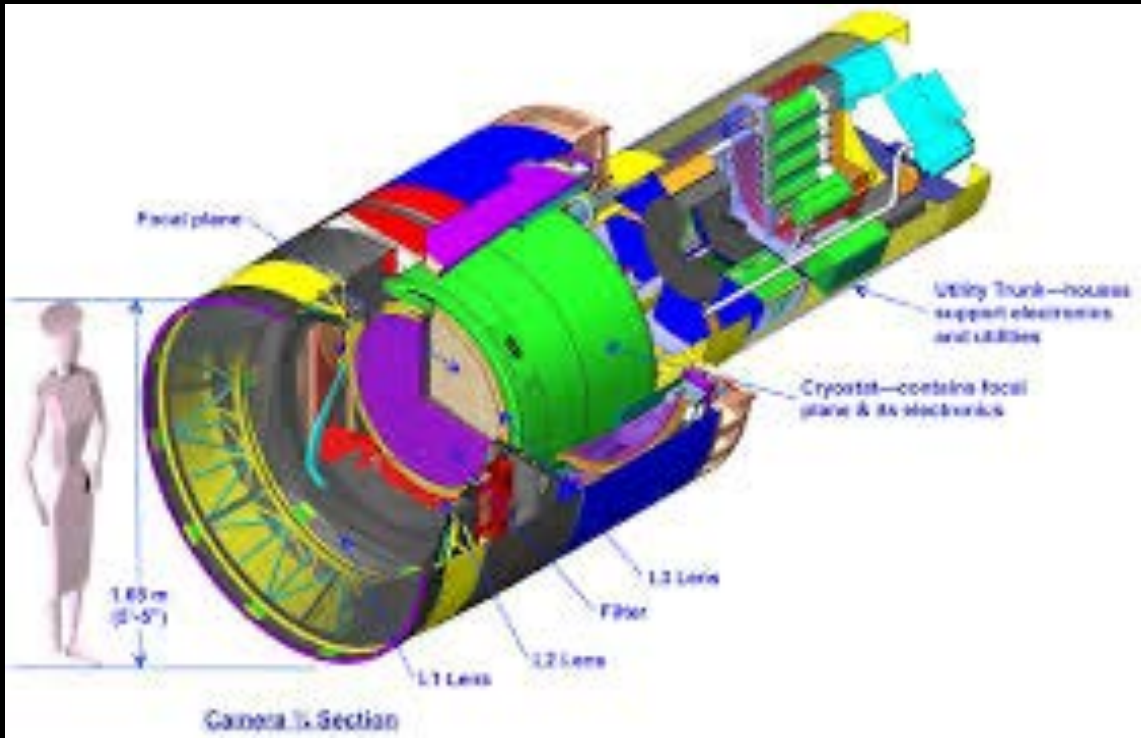

OK, but where exactly should we point the telescope?

Scheduler Motivations

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 years
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

For Rubin, “scheduler” is a misnomer
What we really have is an artificial intelligence that makes real-time decisions about what observations to attempt

“Schedulers” often solve maximization-like problems

Kepler scheduler = maximize the probability of finding a planet in the habitable zone
Gemini scheduler = maximize TAC happiness

We have 4 broad science goals. How many supernova equal an asteroid? ***We can’t cast this as a maximization problem!***

Model observatory

Can give it commands (like, make this observation) and it will simulate slewing to the position and record the conditions of the exposure.

Outputs completed observations

Outputs conditions object

- Map of slewtime
- Map of seeing
- Map of sky brightness
- MJD, LMST
- Next moon rise/set, next twilight start/end

- Kinematic model of telescope, dome, camera
- Historical weather log. We close if 30% of the sky is cloudy, so our weather downtime matches Gemini and SOAR
- Scheduled and unscheduled downtime
- Historical seeing log
- Sky brightness model (ESO model + twilight fit from all sky camera)

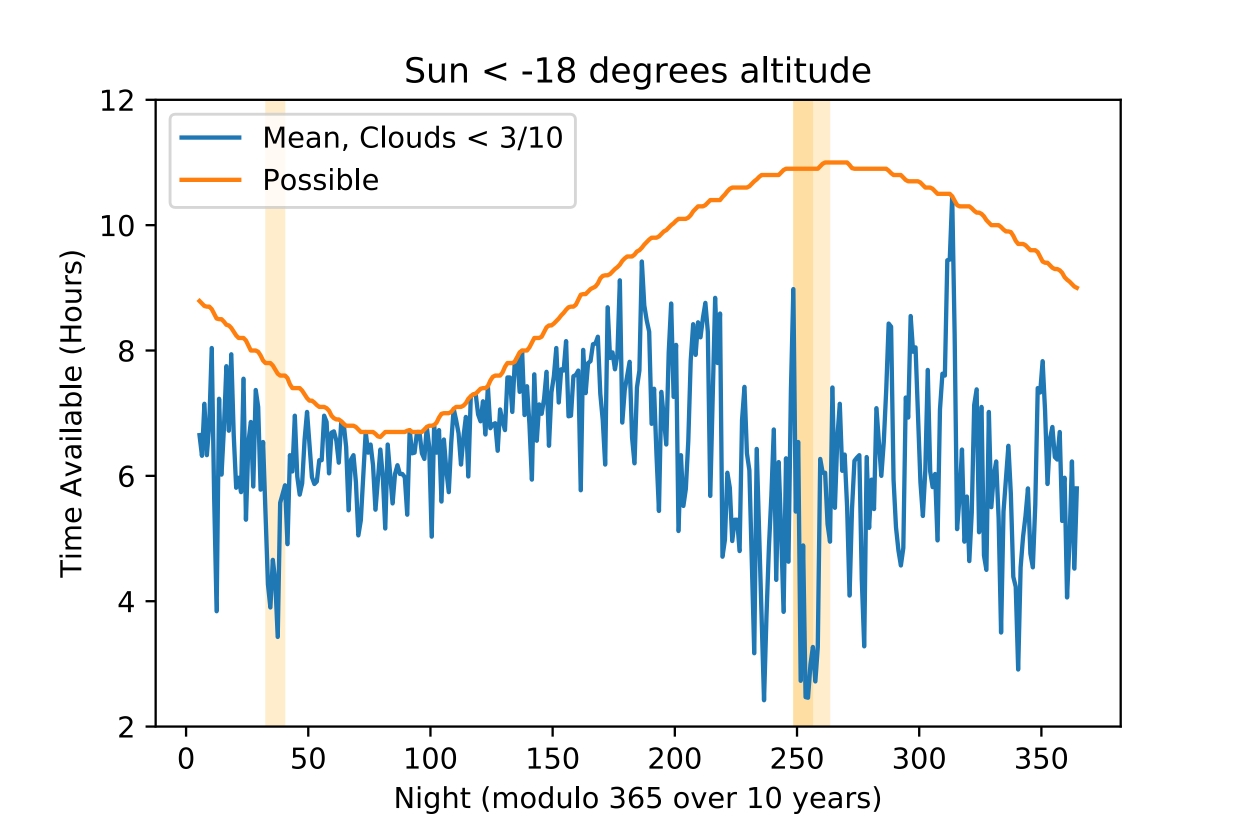

Sky brightness model

- ESO model includes
    - Upper atmosphere
    - Lower atmosphere
    - Airglow continuum
    - Zodiacal light
    - Scattered lunar light
    - Solar twilight

This is expensive to compute, so we pre-compute the sky brightness for the Rubin filters every ~15 minutes and interpolate.

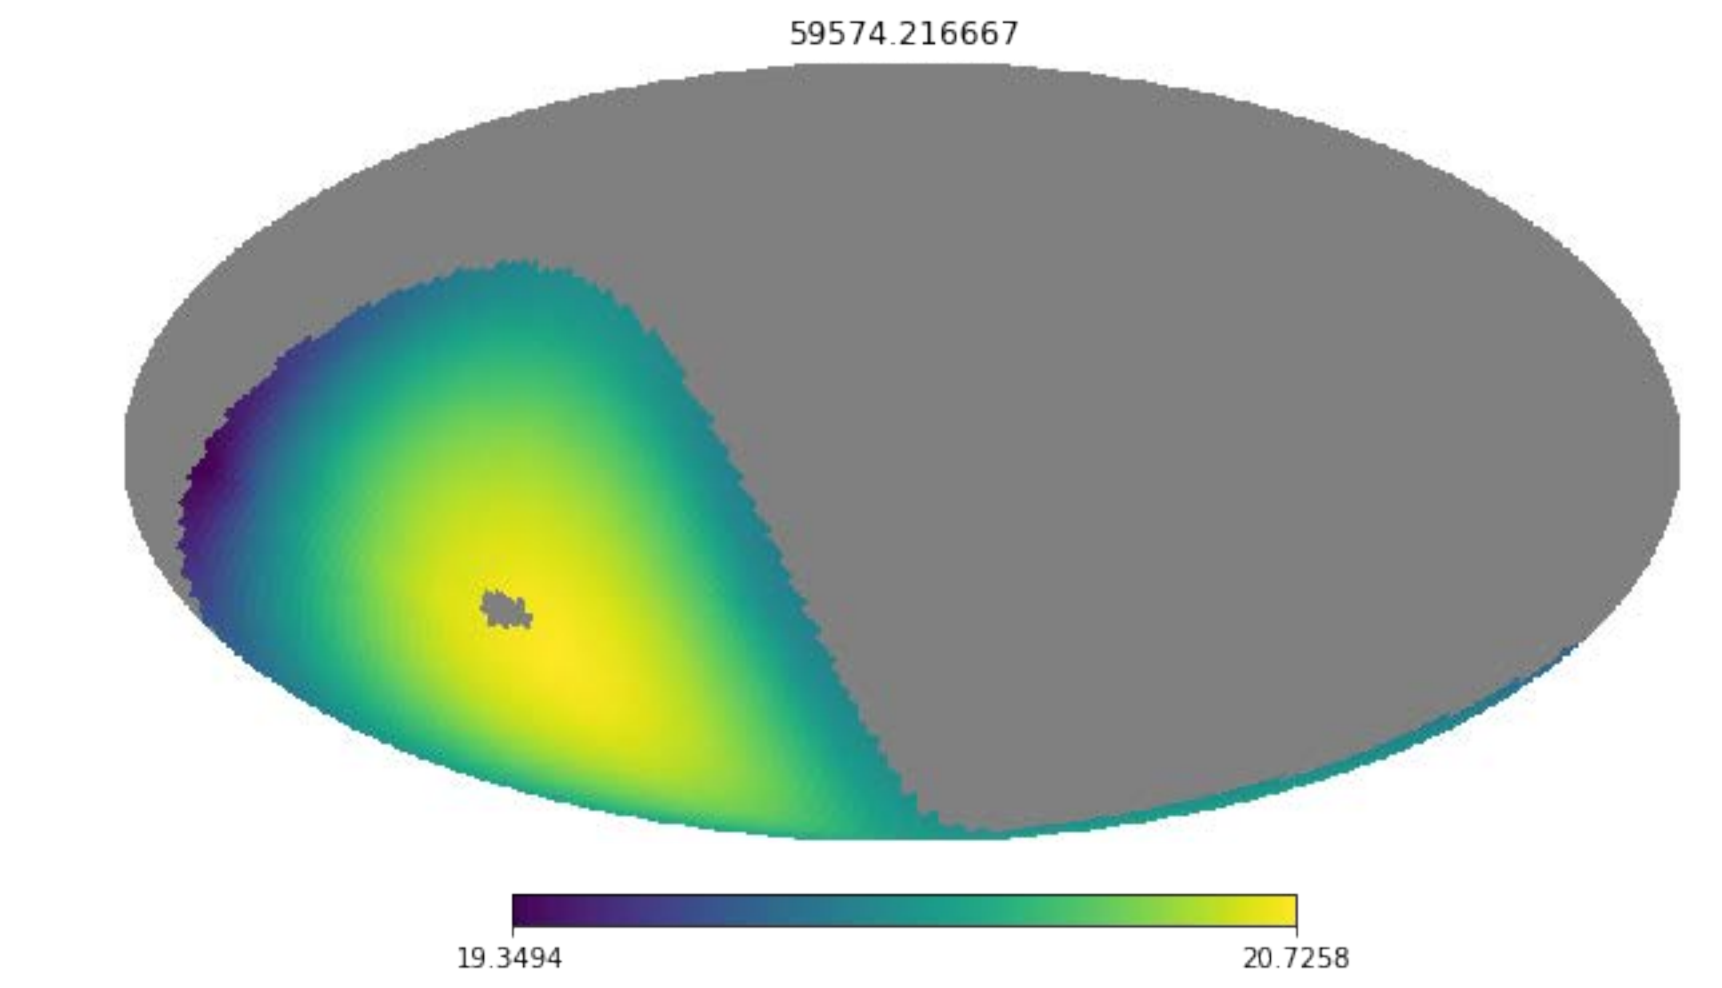

Sky brightness in r (mag/sq arcsec)

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 years
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

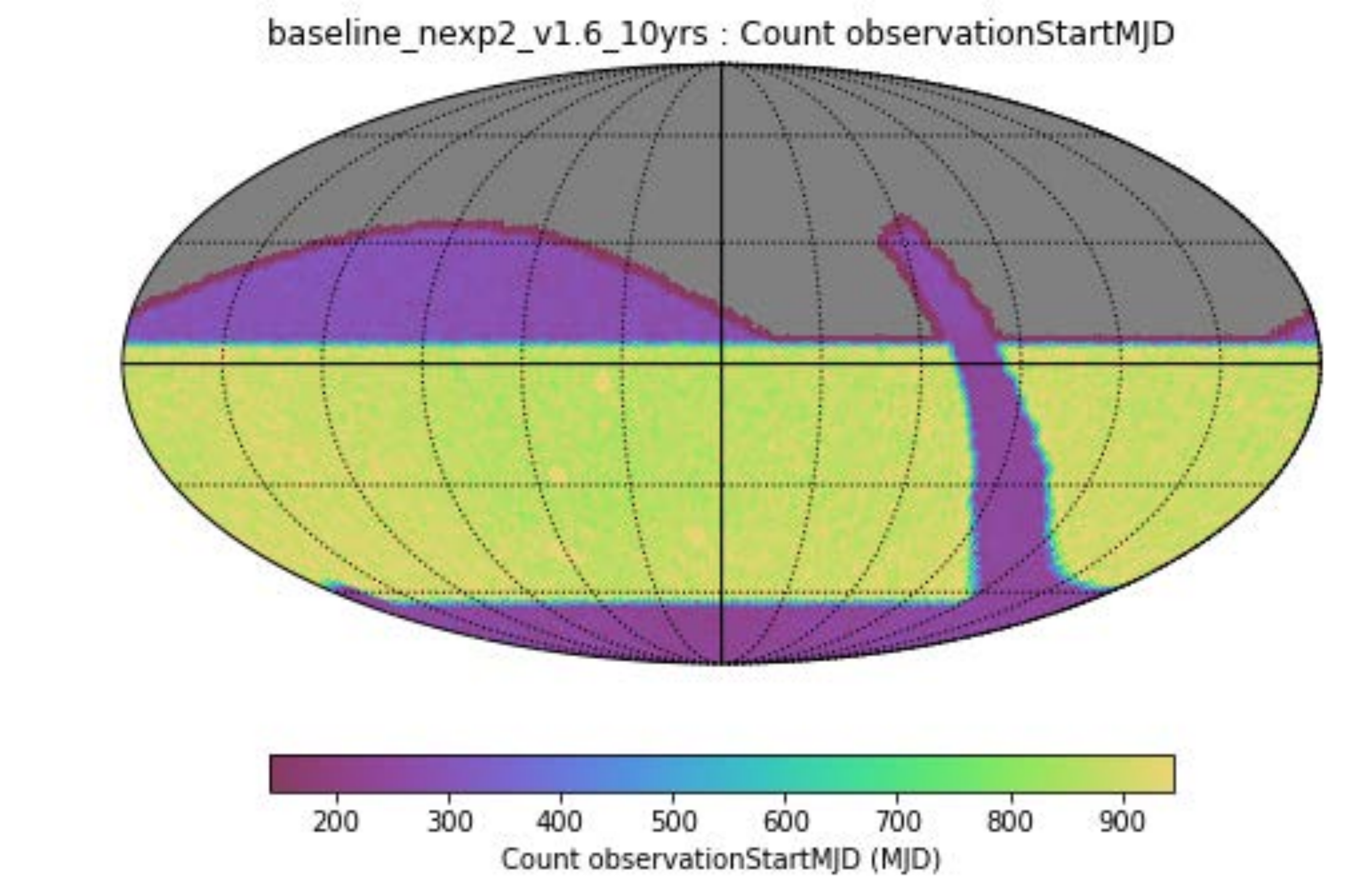

Survey footprint, most of the time is spent in the 18,000 sq degree Wide Fast Deep (WFD) area

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 years
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

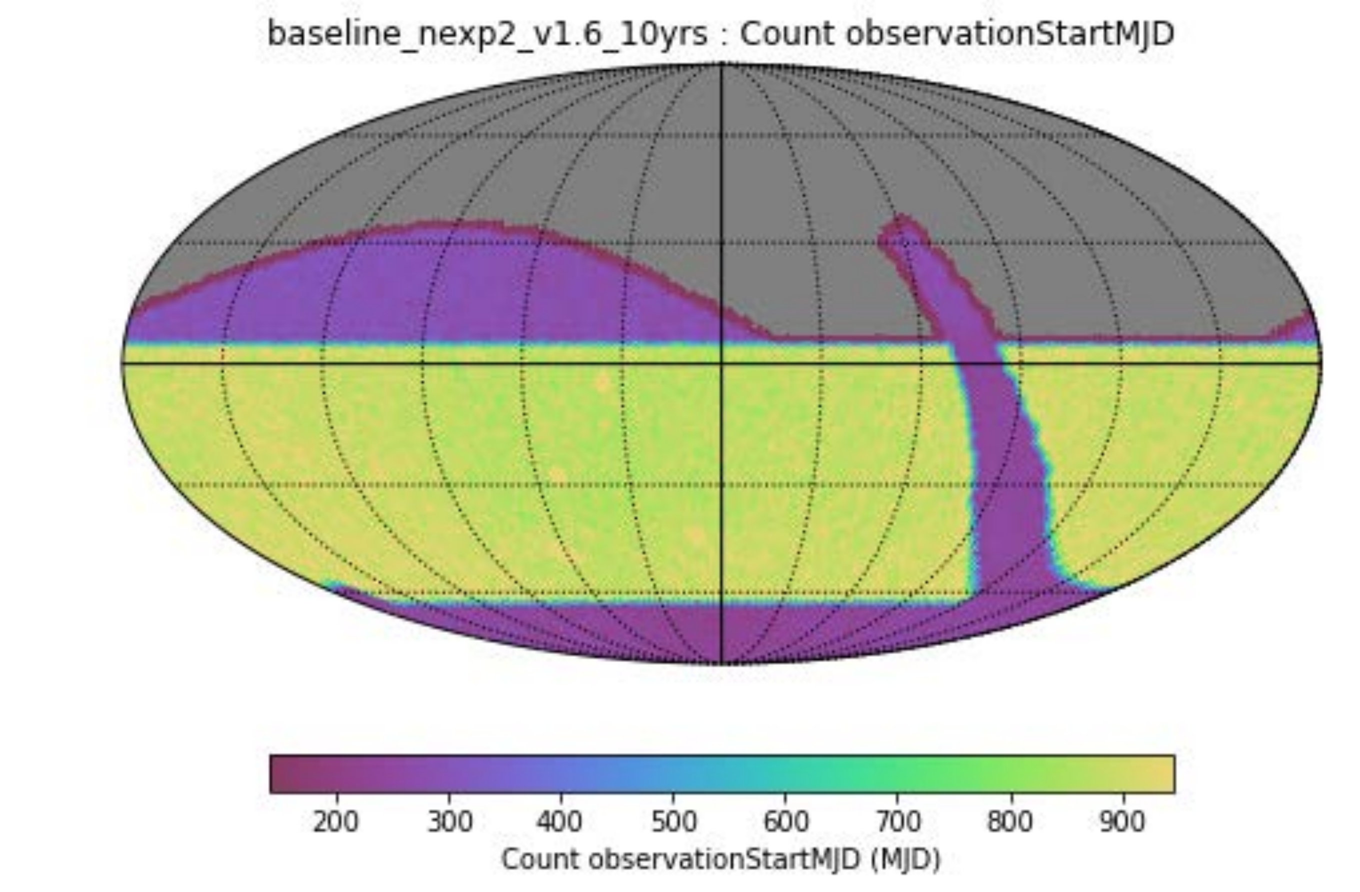

Proper motion depends on temporal baseline, so this requirement basically reduces to we need to survey the entire WFD evenly in year 1 and year 10

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 years
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

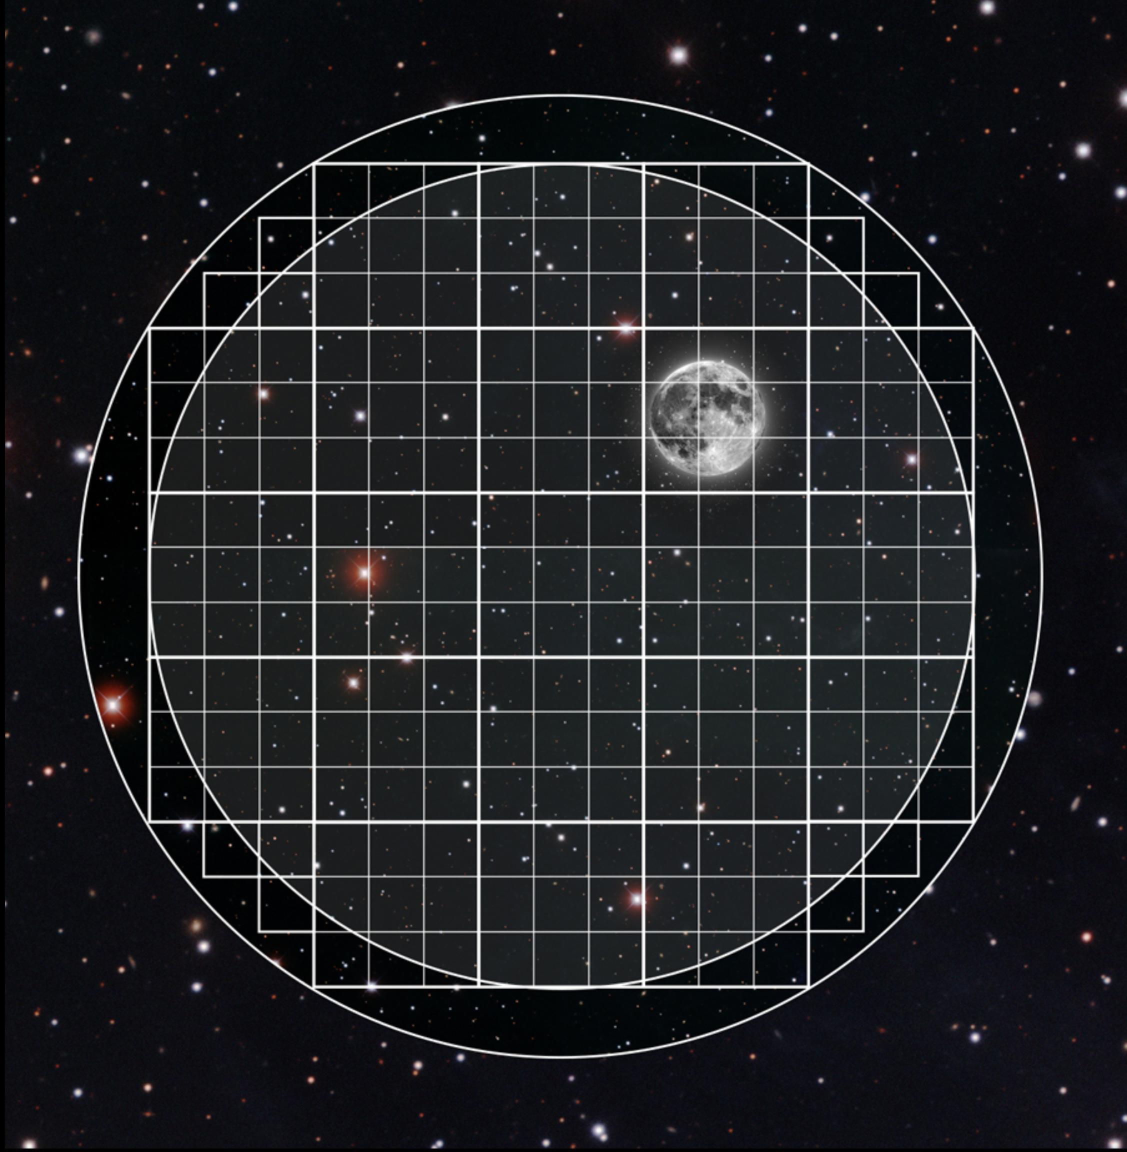

Rubin field of view. Each square is a CCD
Inner circle defines the science FoV

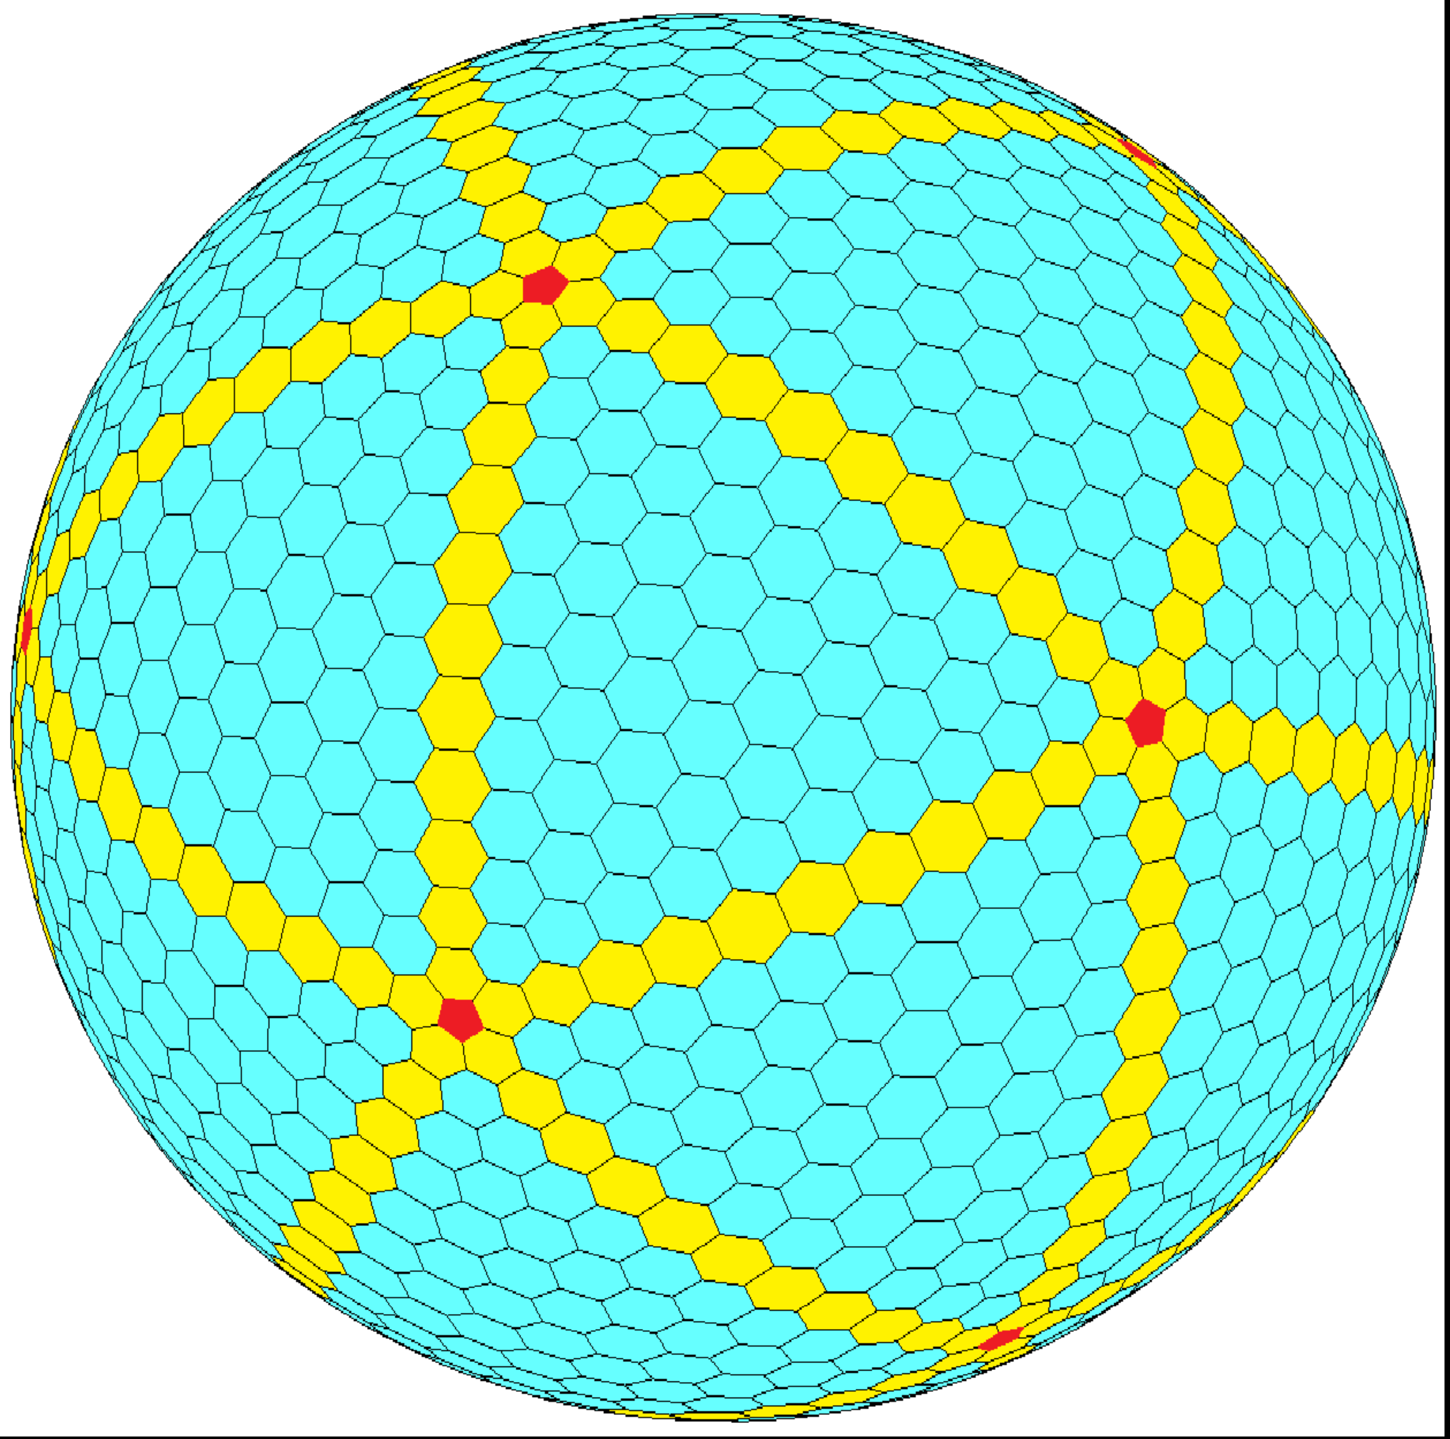

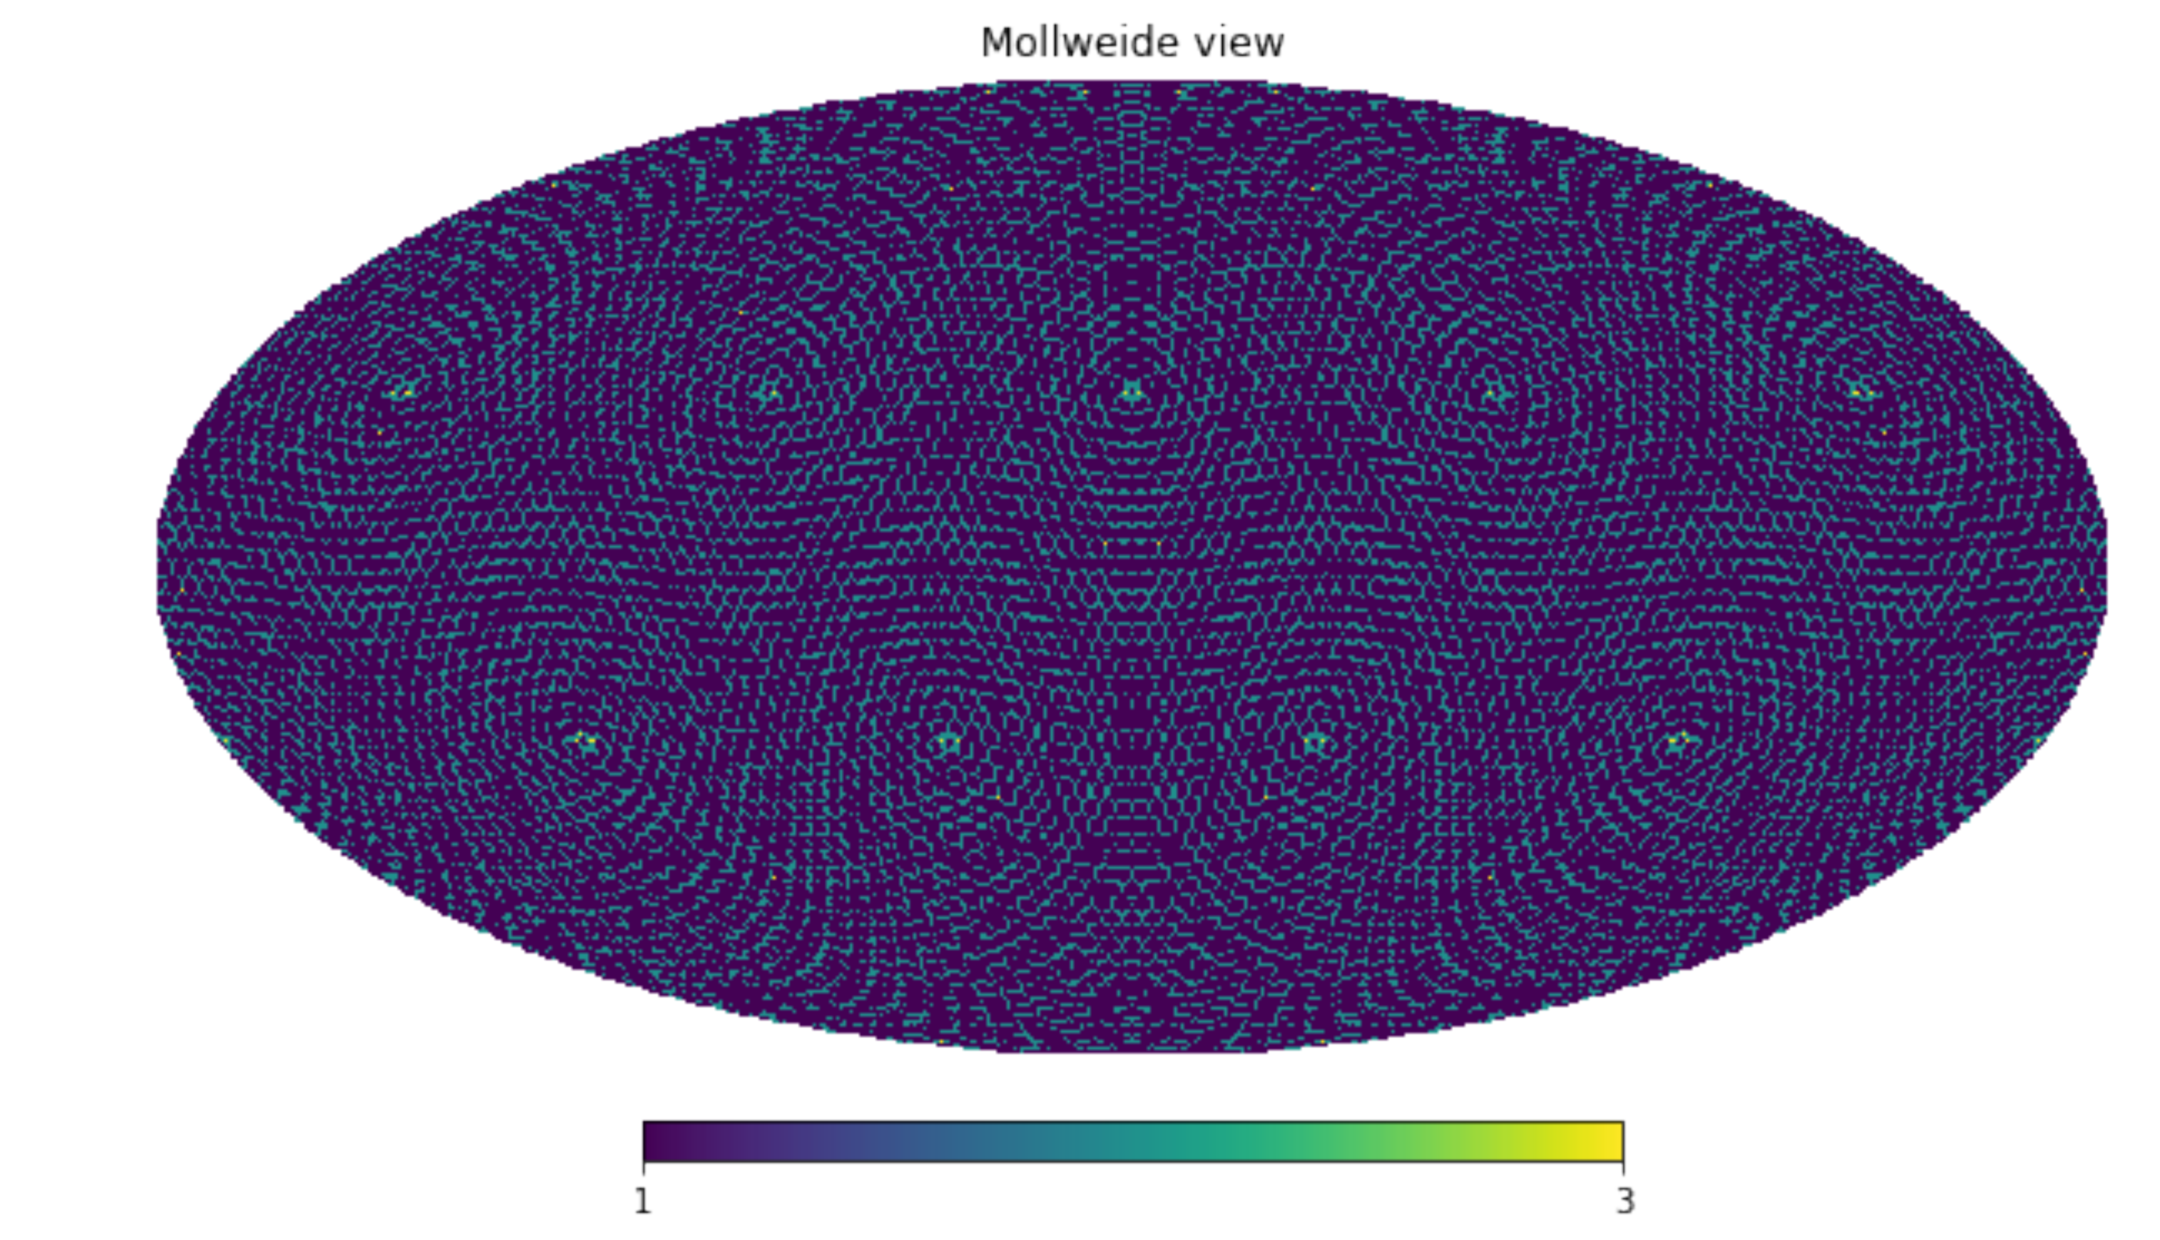

We can cover the full sky with 5292 pointings
The orientation is randomized each night. *There are no fixed fields.*

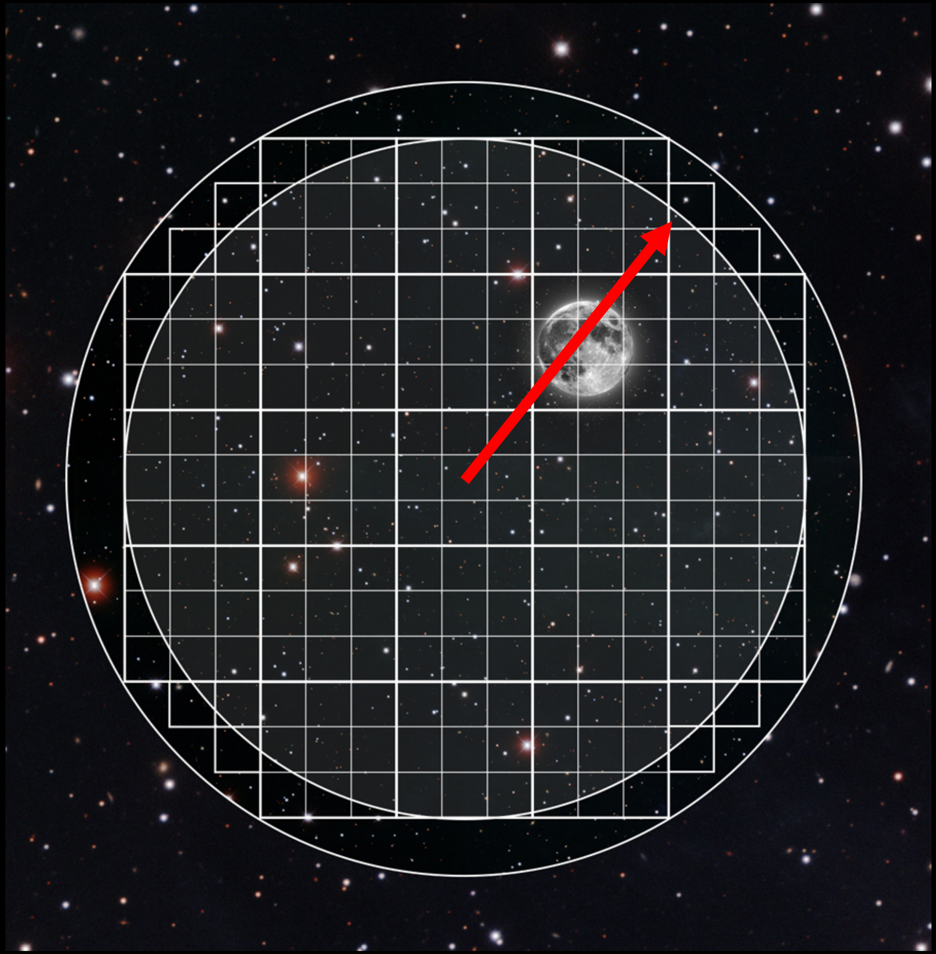

If we observe neighboring pointings, overlap regions will get observed on 0.5-22 minute timescales, meeting the rapid revisit requirement

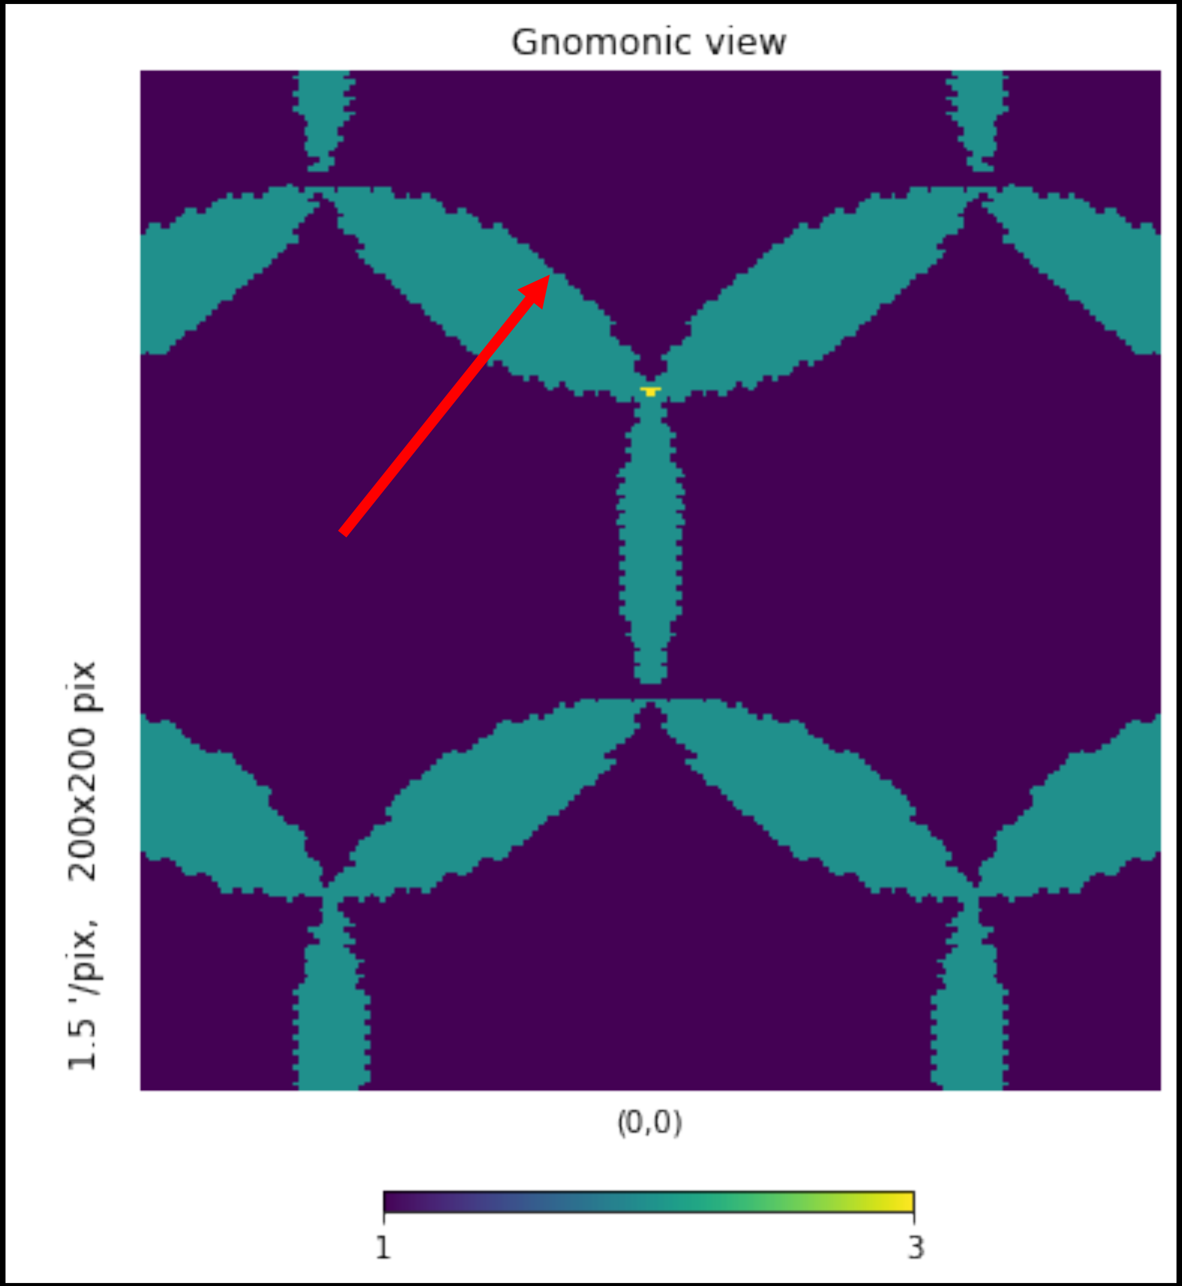

Try to come back to the same position after ~22 minutes in a different filter.

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 years
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

Science Requirement Document

- We need an area of 18,000 square degrees observed 825 times over 10 ye
- Parallax and proper motion precision requirements
- Rapid revisit requirements

The 4 science cases:

- Nature of dark matter and dark energy
- Catalog the solar system
- The variable sky
- Milky Way structure and formation

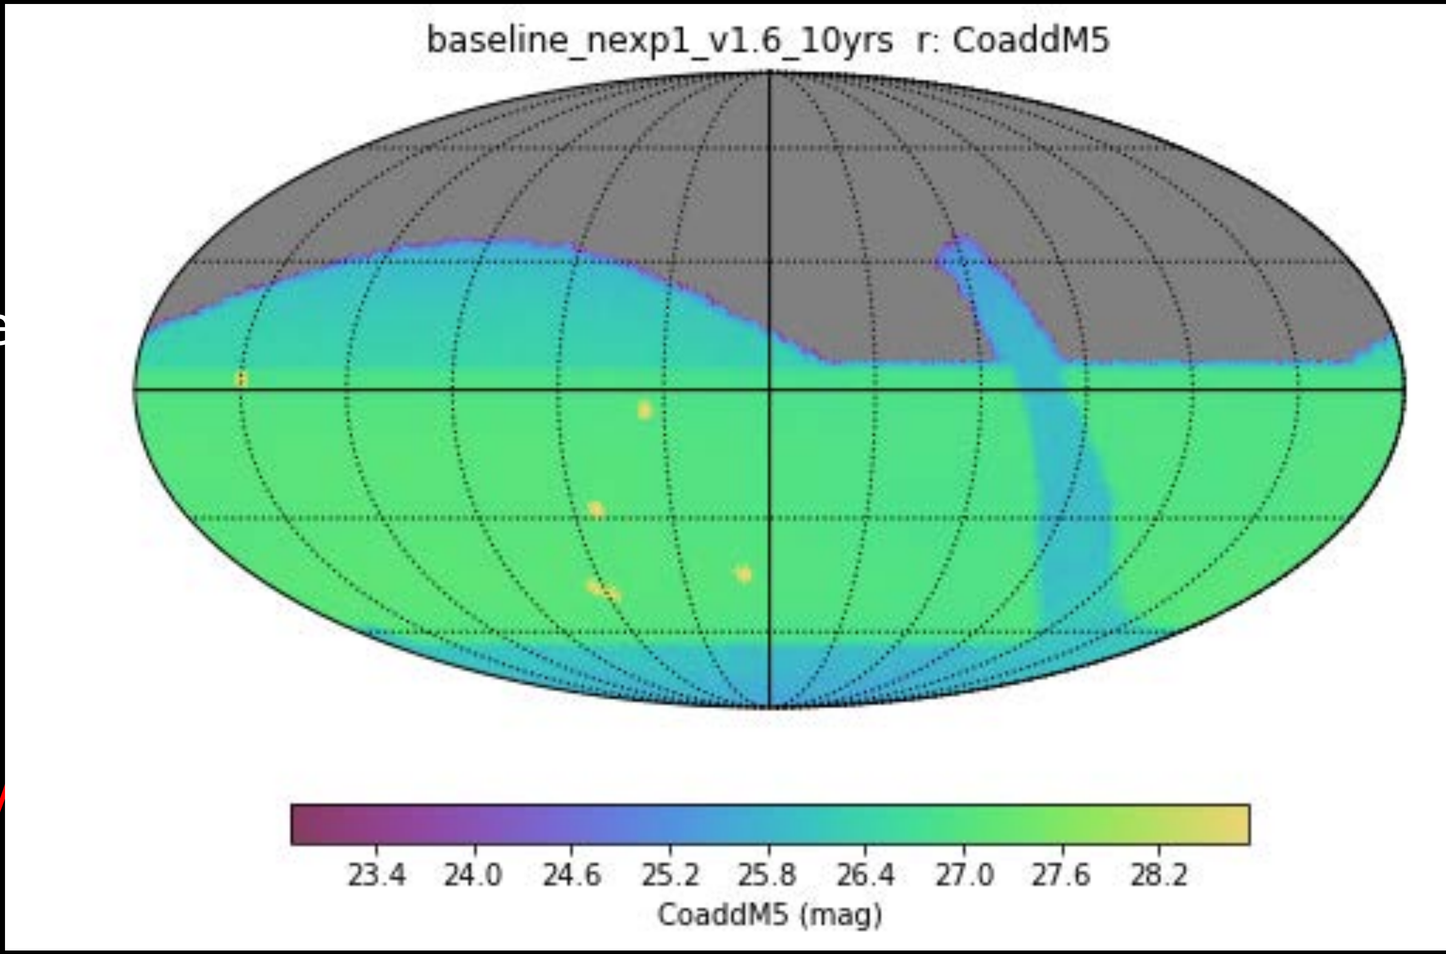

Spatial dithering so WFD area is uniform in depth. Rotational dithering of the camera as well.

The scheduler goes through a decision tree:

1. Is there a deep drilling field up that needs observing? There are 4 announced positions (and probably a 5th position) that will be deep drilling fields.

2. Can I observe a mostly-contiguous sky area twice? (not currently twilight, twilight not starting soon)

3. Pick the best looking spot on the sky for a single observation. These should only happen in and right before twilight (sun altitude -18 degrees to -12 degrees).

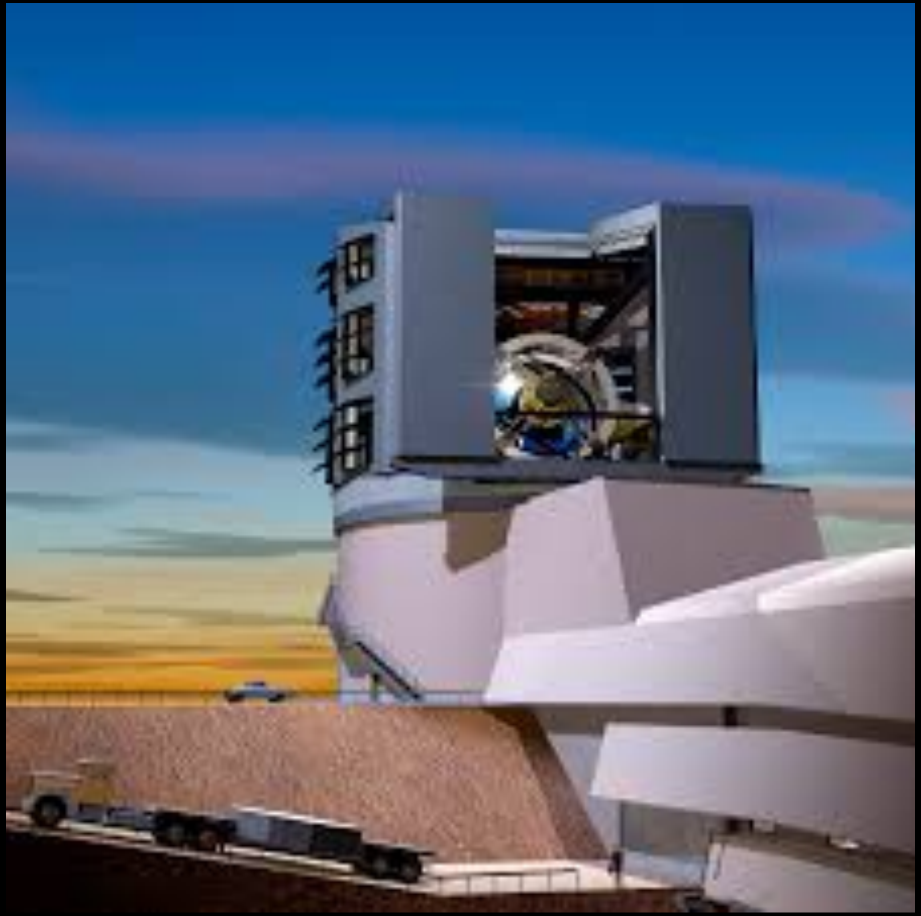

An example night

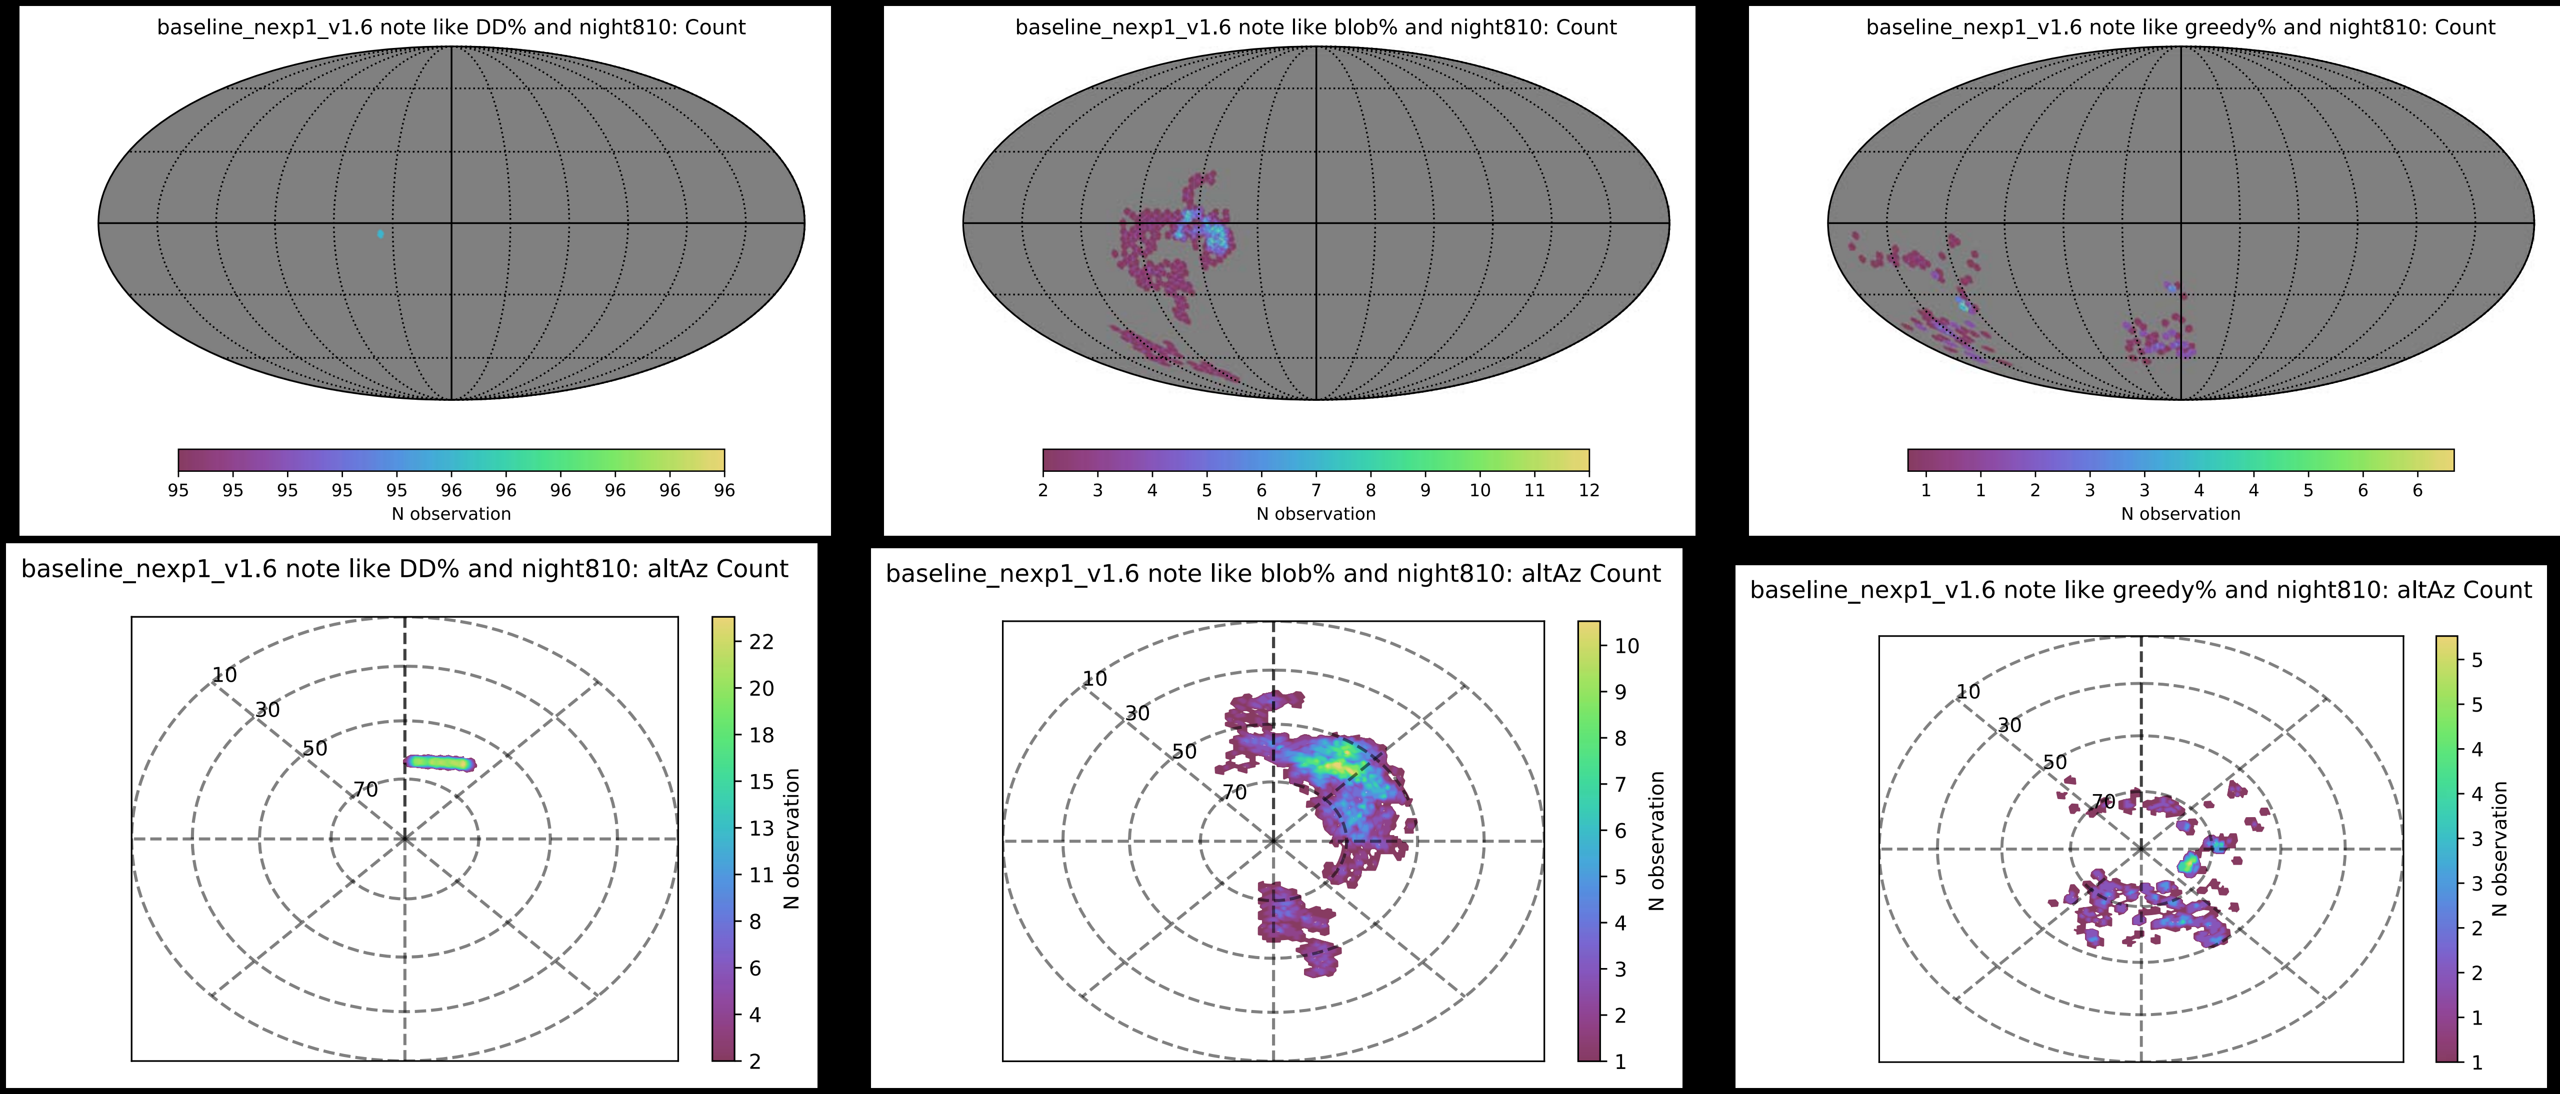

DDF sequence | Paired observations | Twilight observations

Chronologically, the night goes right, middle, left, middle, right.

How do we pick that semi-contiguous blob of sky?

Markov Decision Process (popular in the robotics community) let's us balance different factors

Define the footprint we want the survey to cover in each filter
HEALpix maps, nside=32 (1.8 degree resolution)

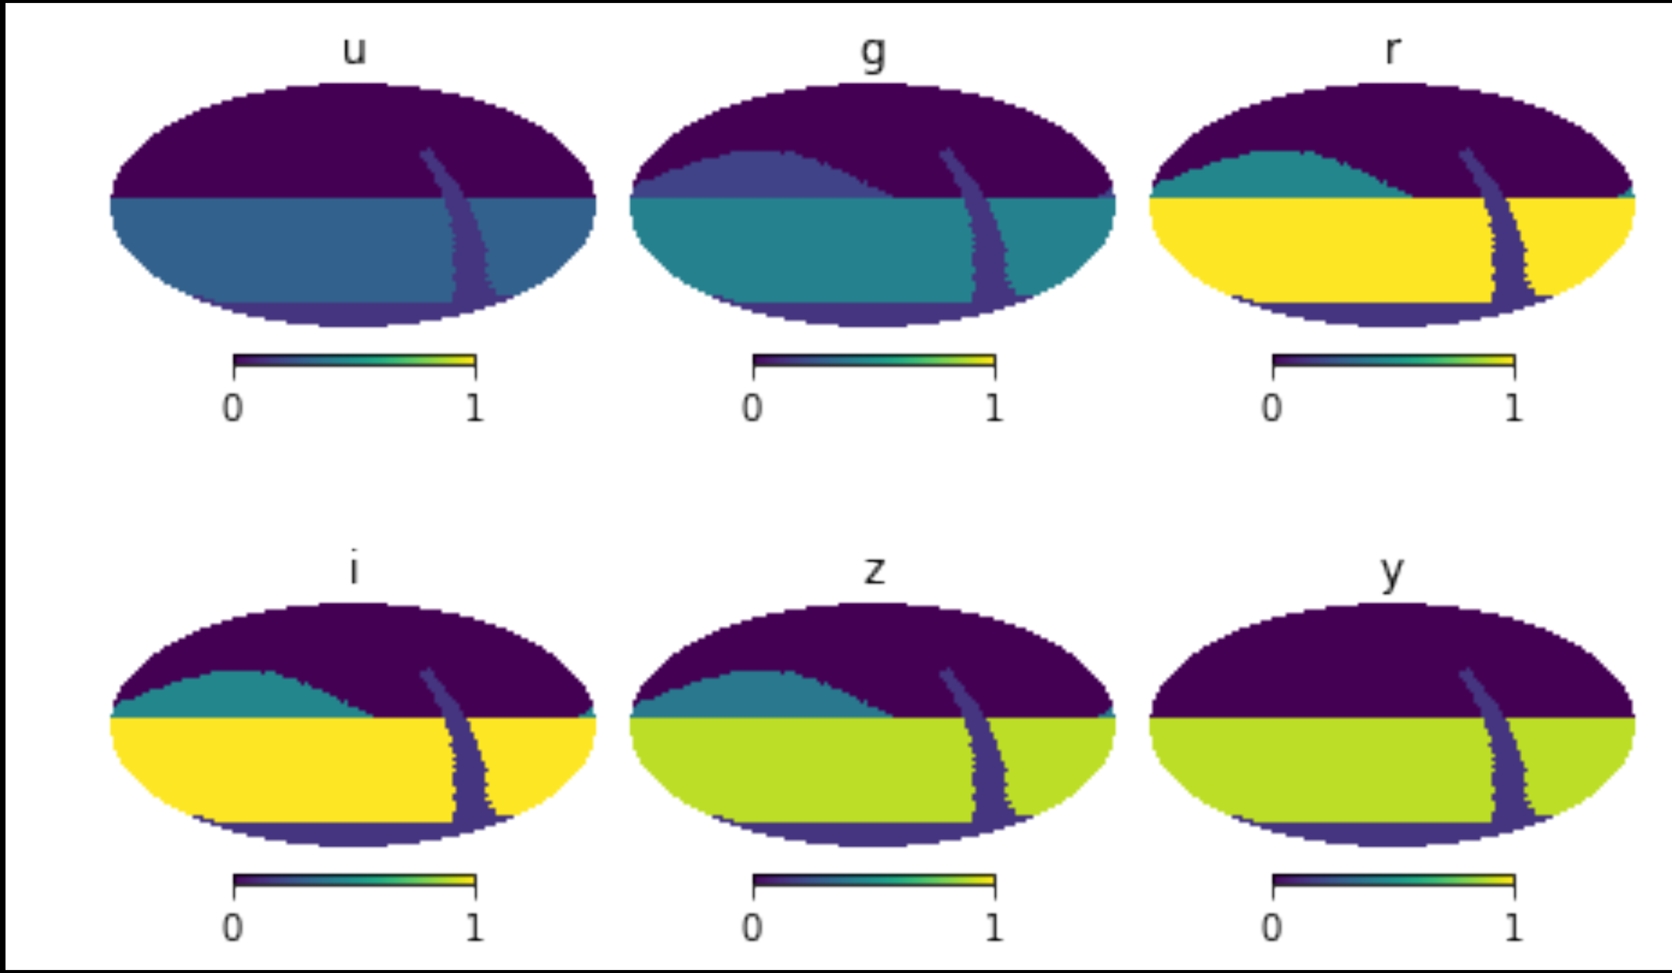

Survey footprint

How does the coverage compare to the desired footprint?

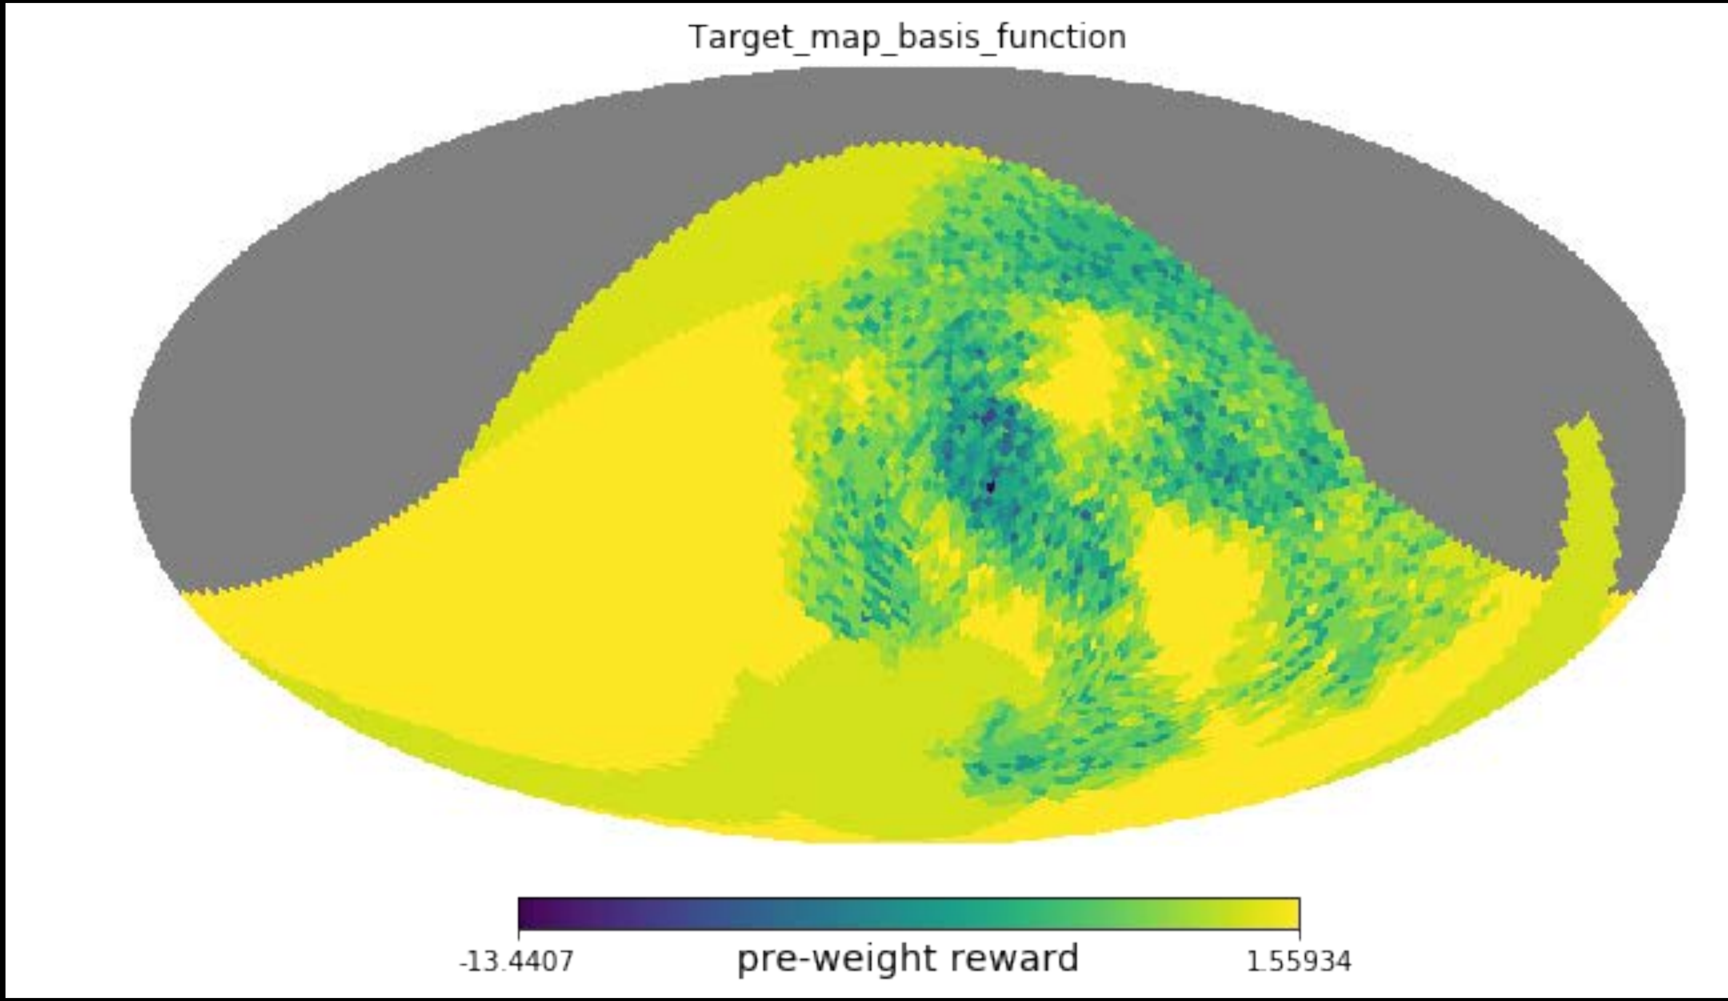

With v1.6, these maps are now time-dependent. No weight on areas that haven’t come into view yet.

5-sigma limiting depth in each filter. Computed from skybrightness, seeing, and airmass. We then look at the difference between the current 5-sigma depth and the best possible 5-sigma depth (for each filter)

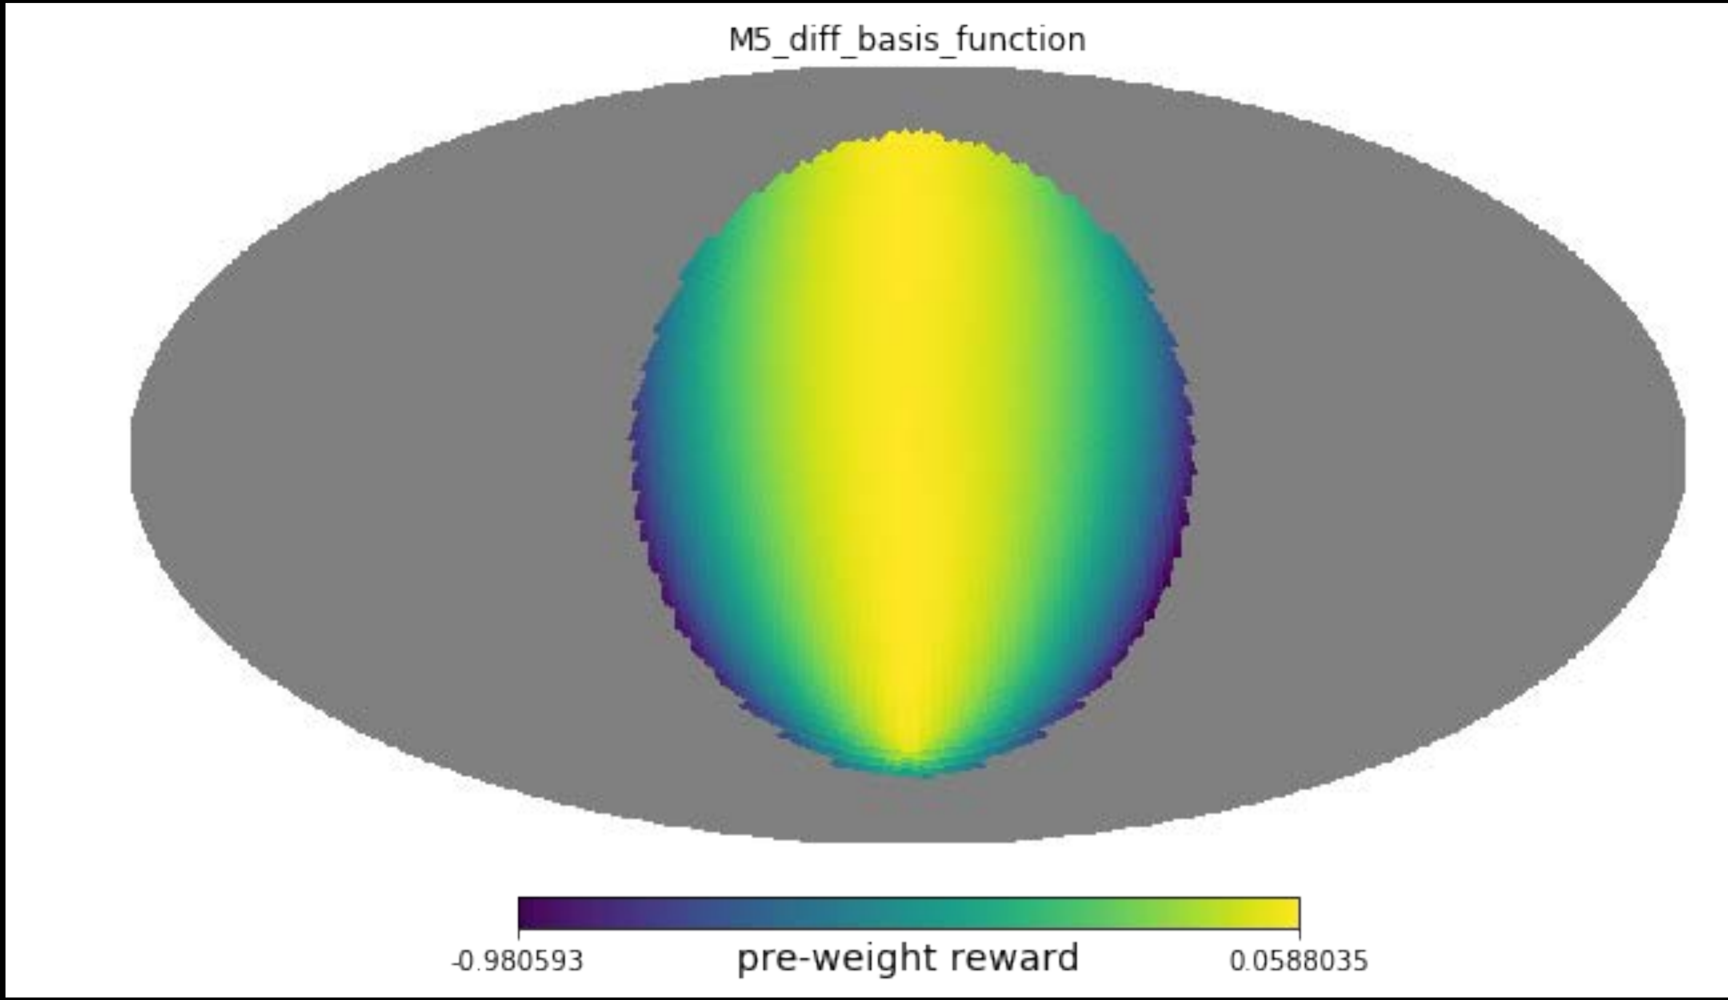

Slewtime: From telescope kinematic model

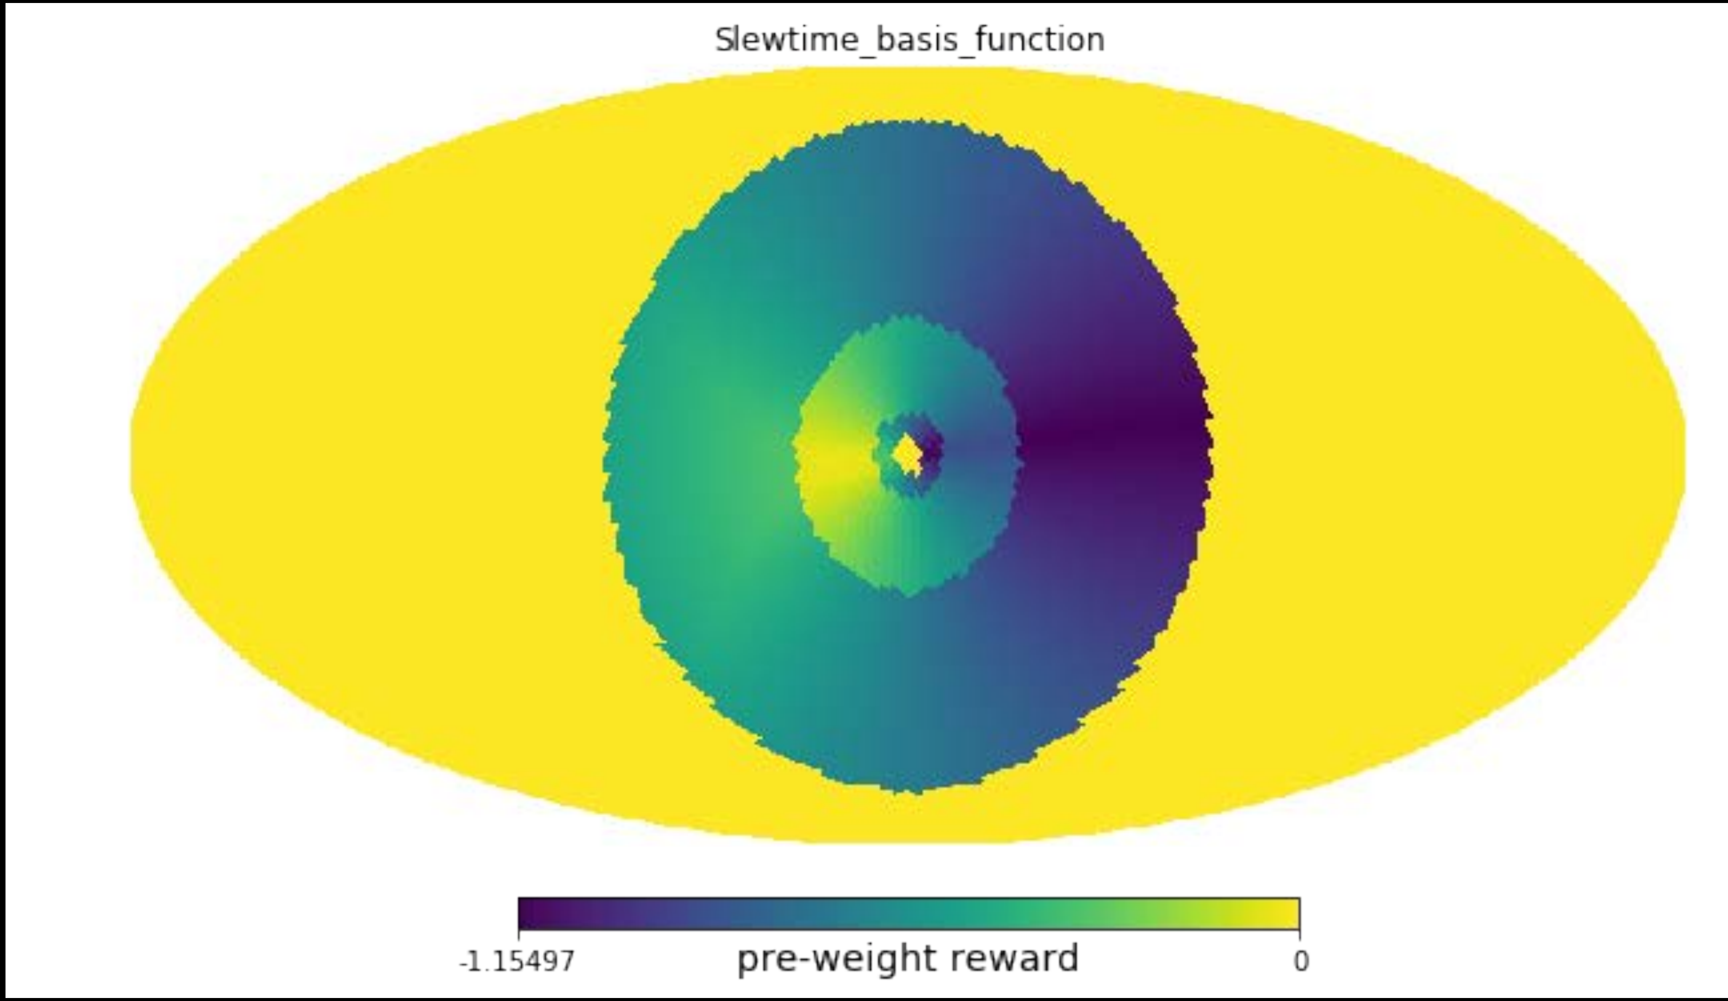


Slew time
Image
Depth
Desired
Footprint

Masks around zenith, moon, bright planets, altitude limit.

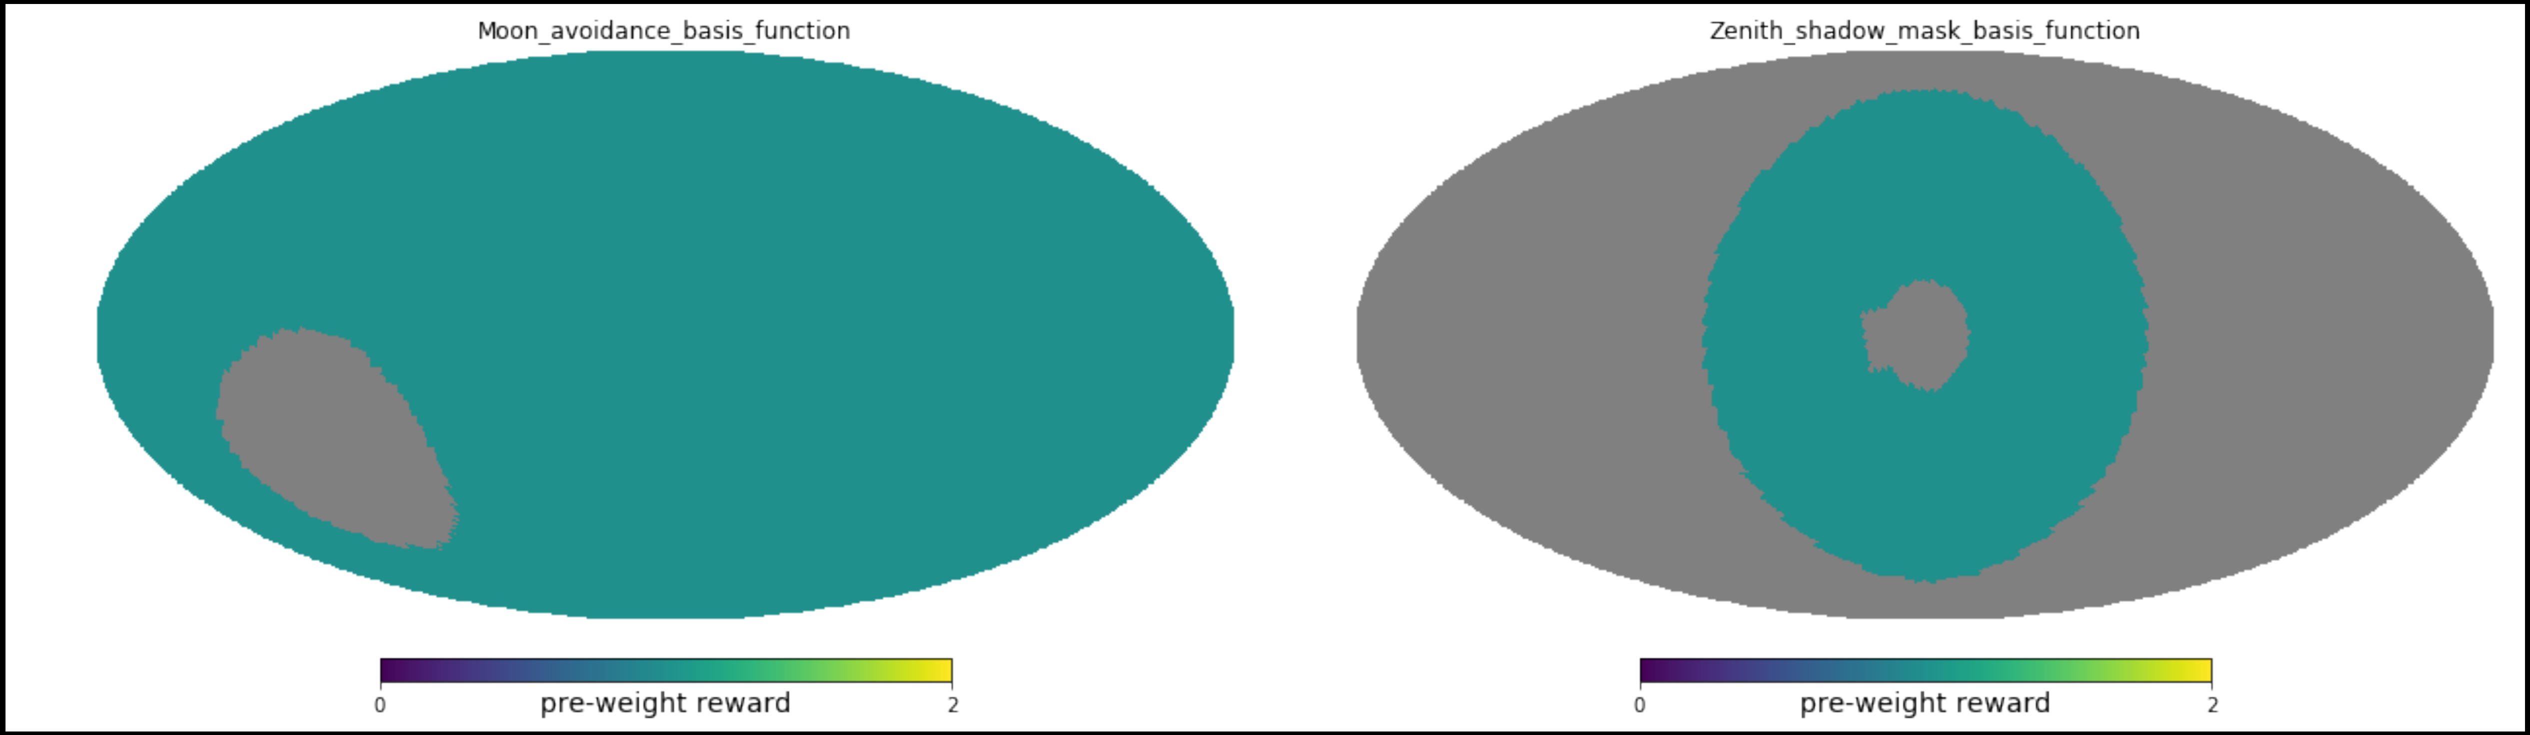

Also a basis function that rewards staying in the same filter (takes 2 minutes to change filter)

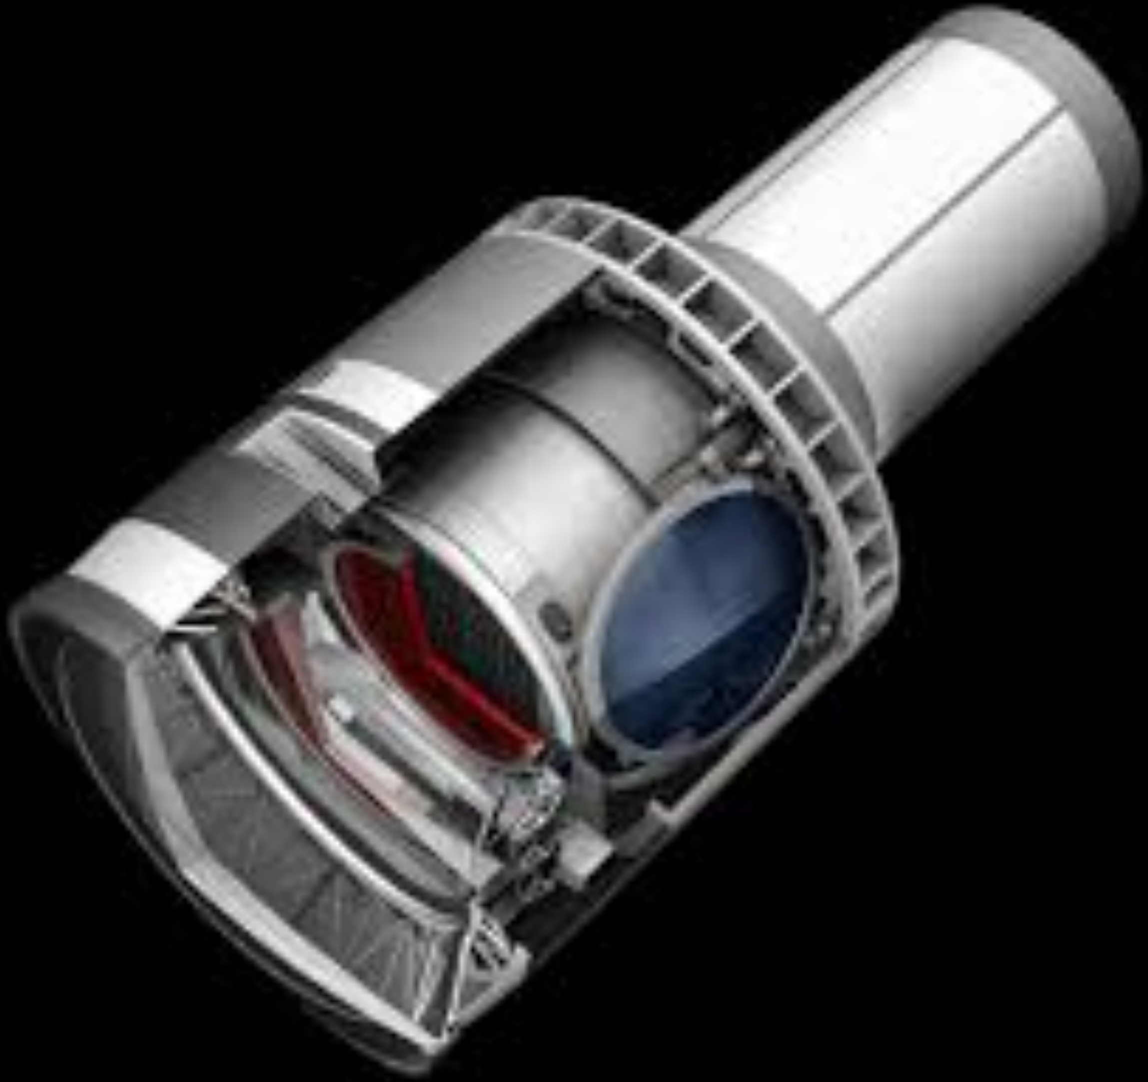

Footprint, depth, slewtime are our basis functions. Each one gets a weight (those are our free parameters), and they are summed to make a final reward function (in 6 filters).

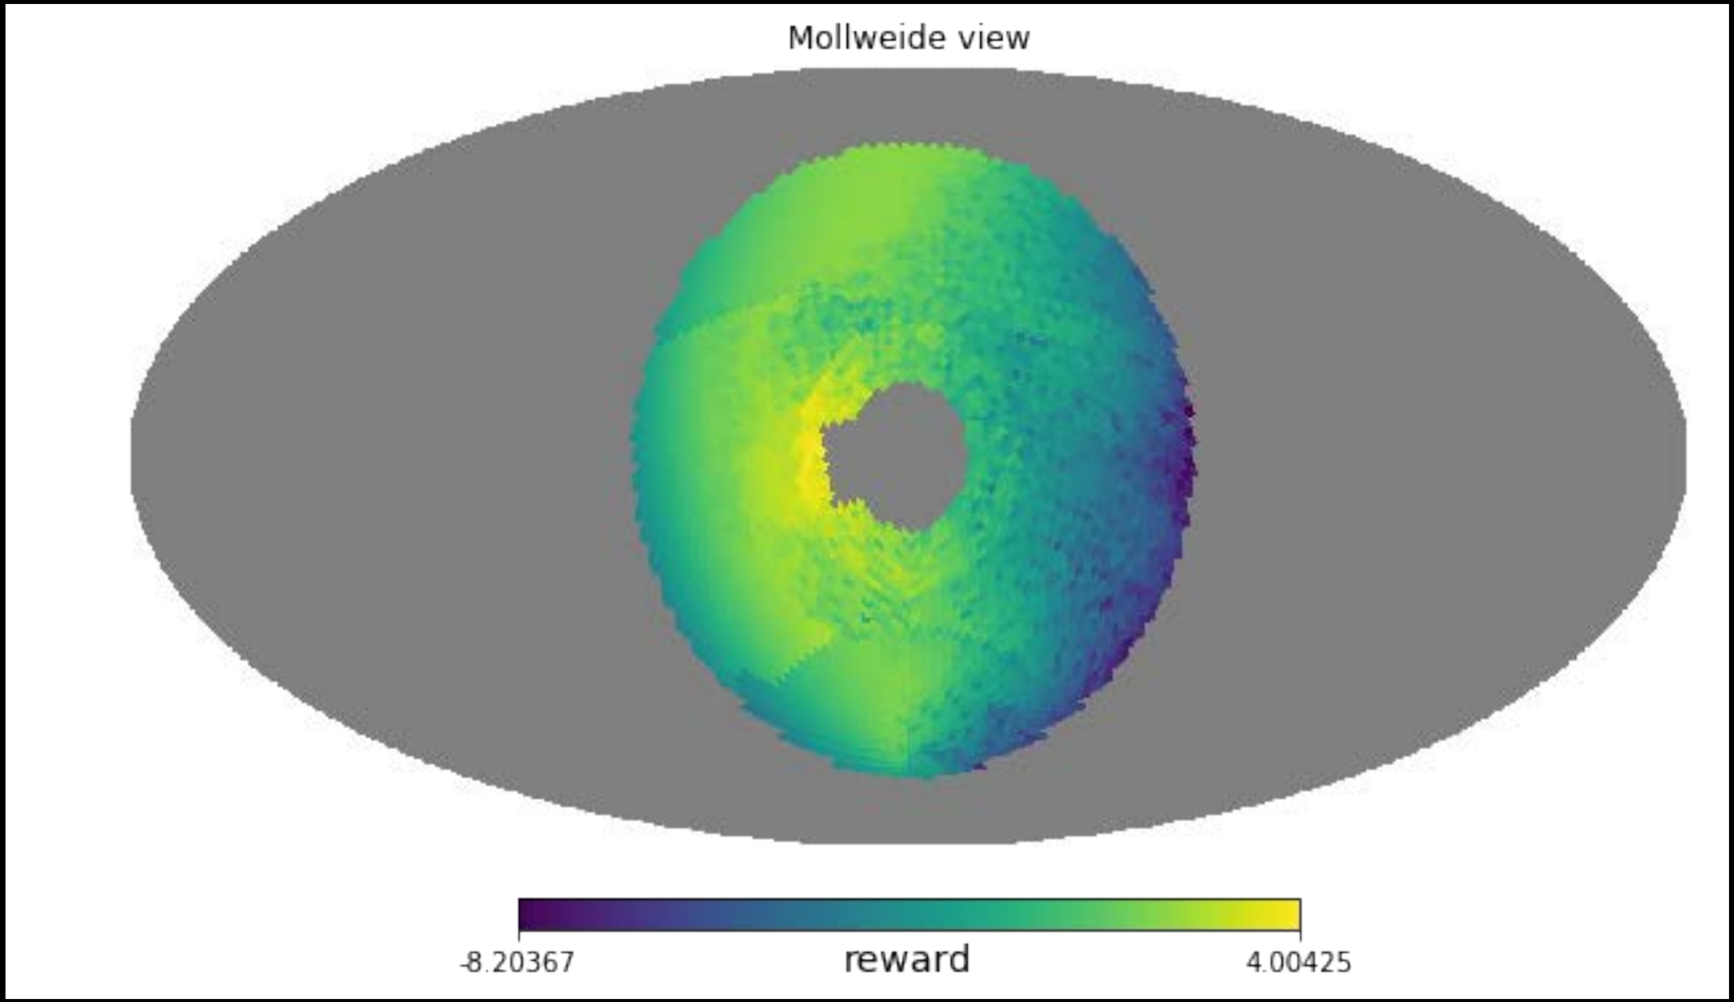

By design, the reward function should be fairly smooth and slowly changing over time

Grow HEALpix blob around the maximum, map to pointings (randomized orientation of tesselation each night)

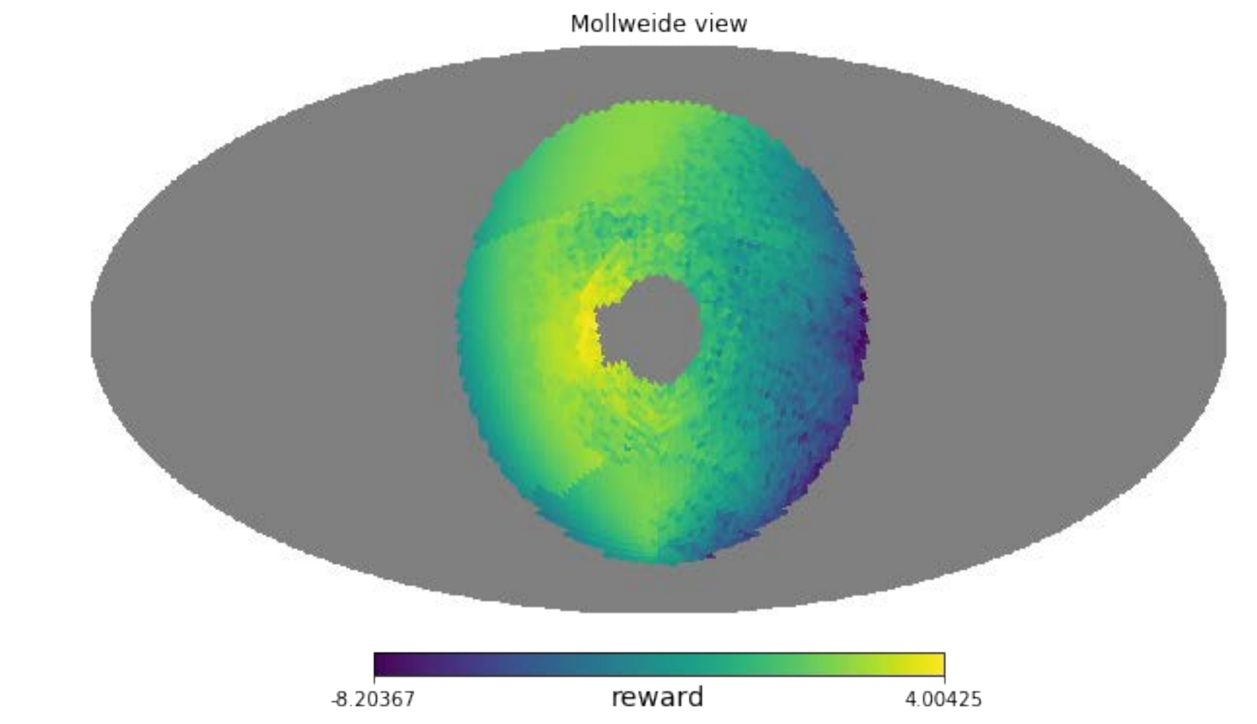

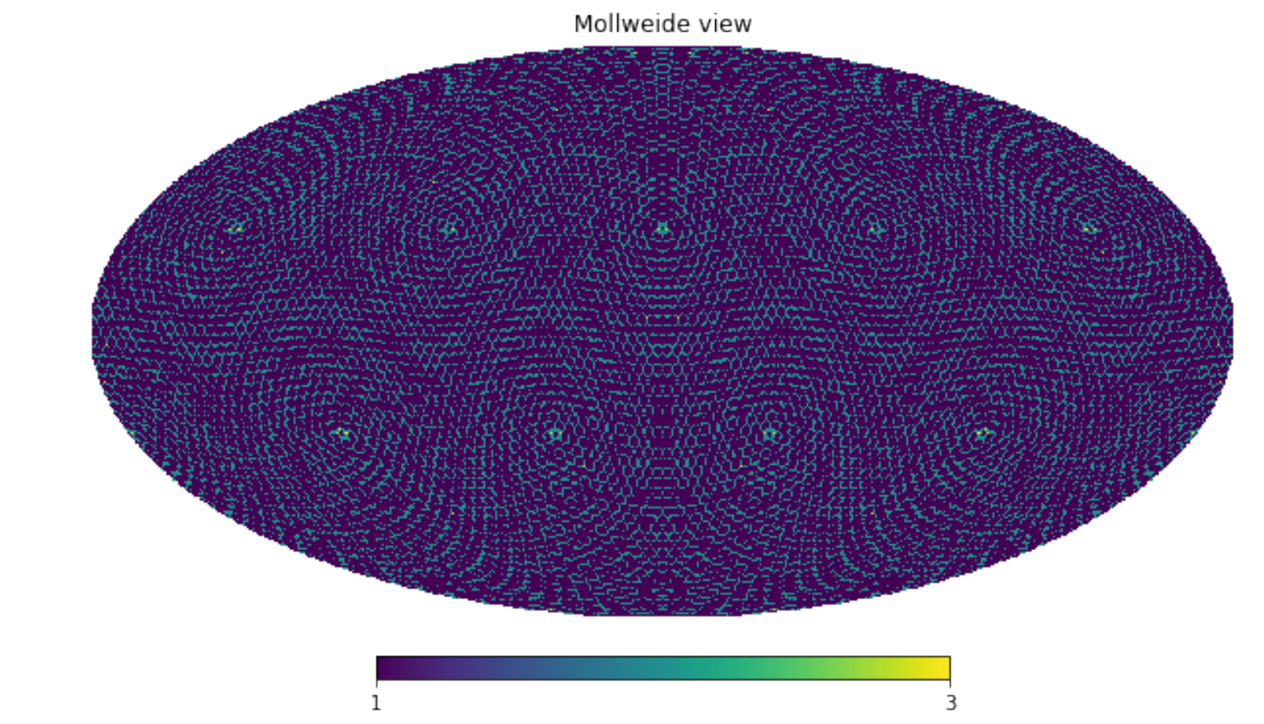

Our list of ~35 pointings

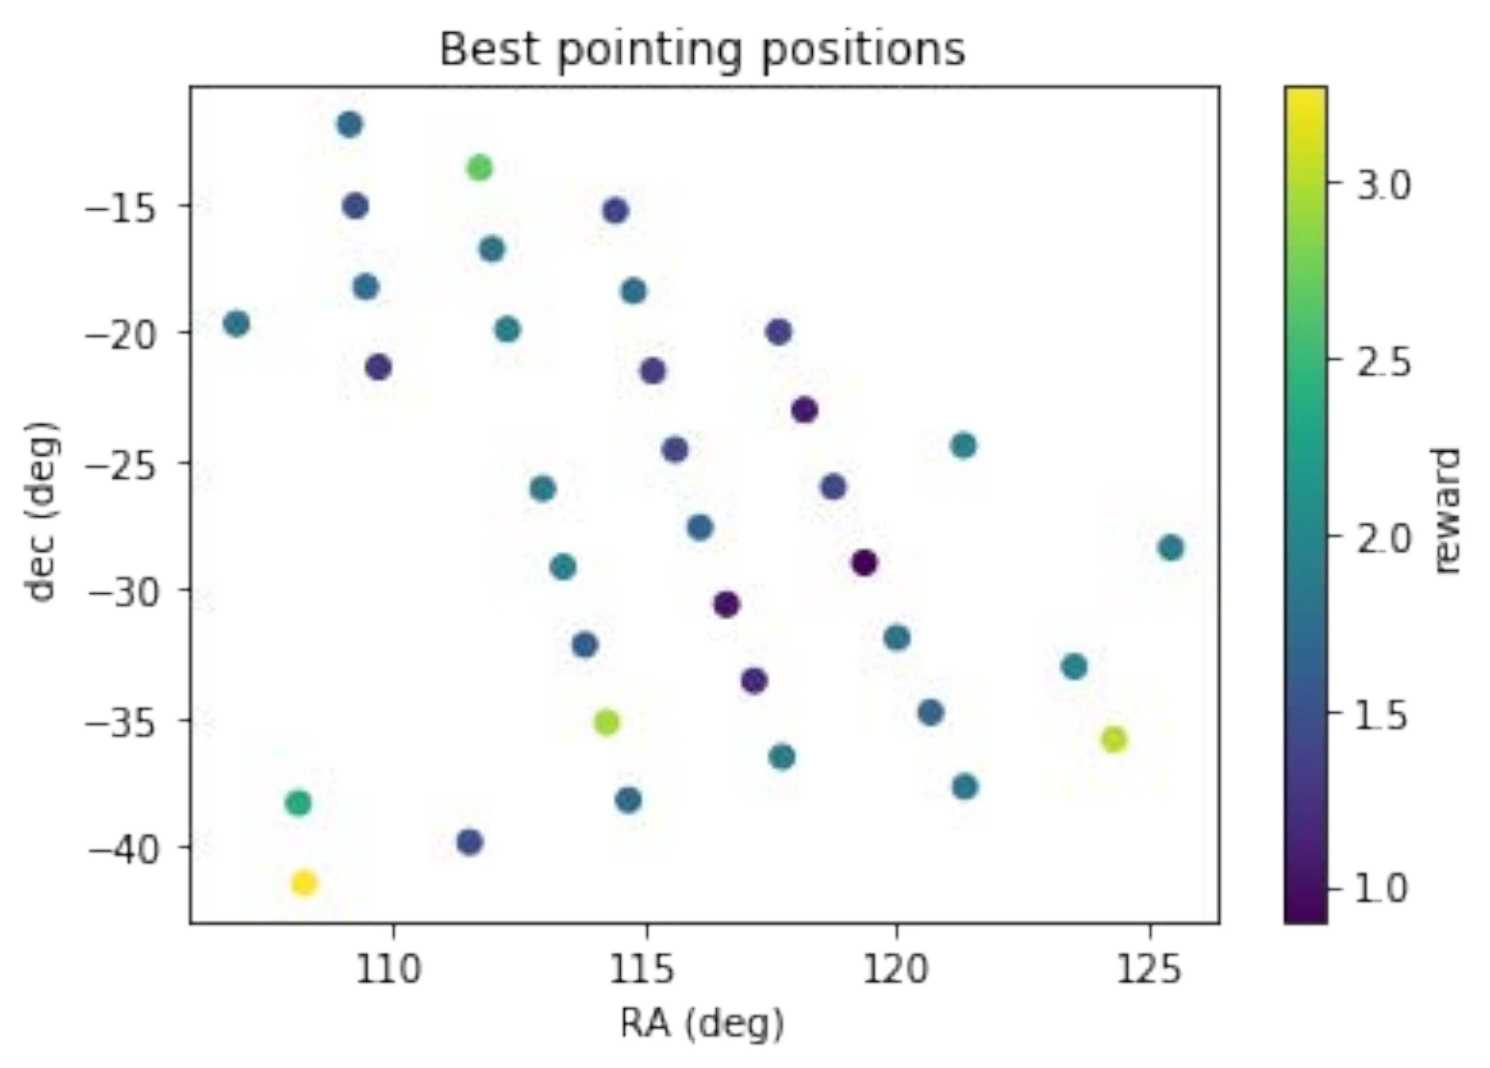

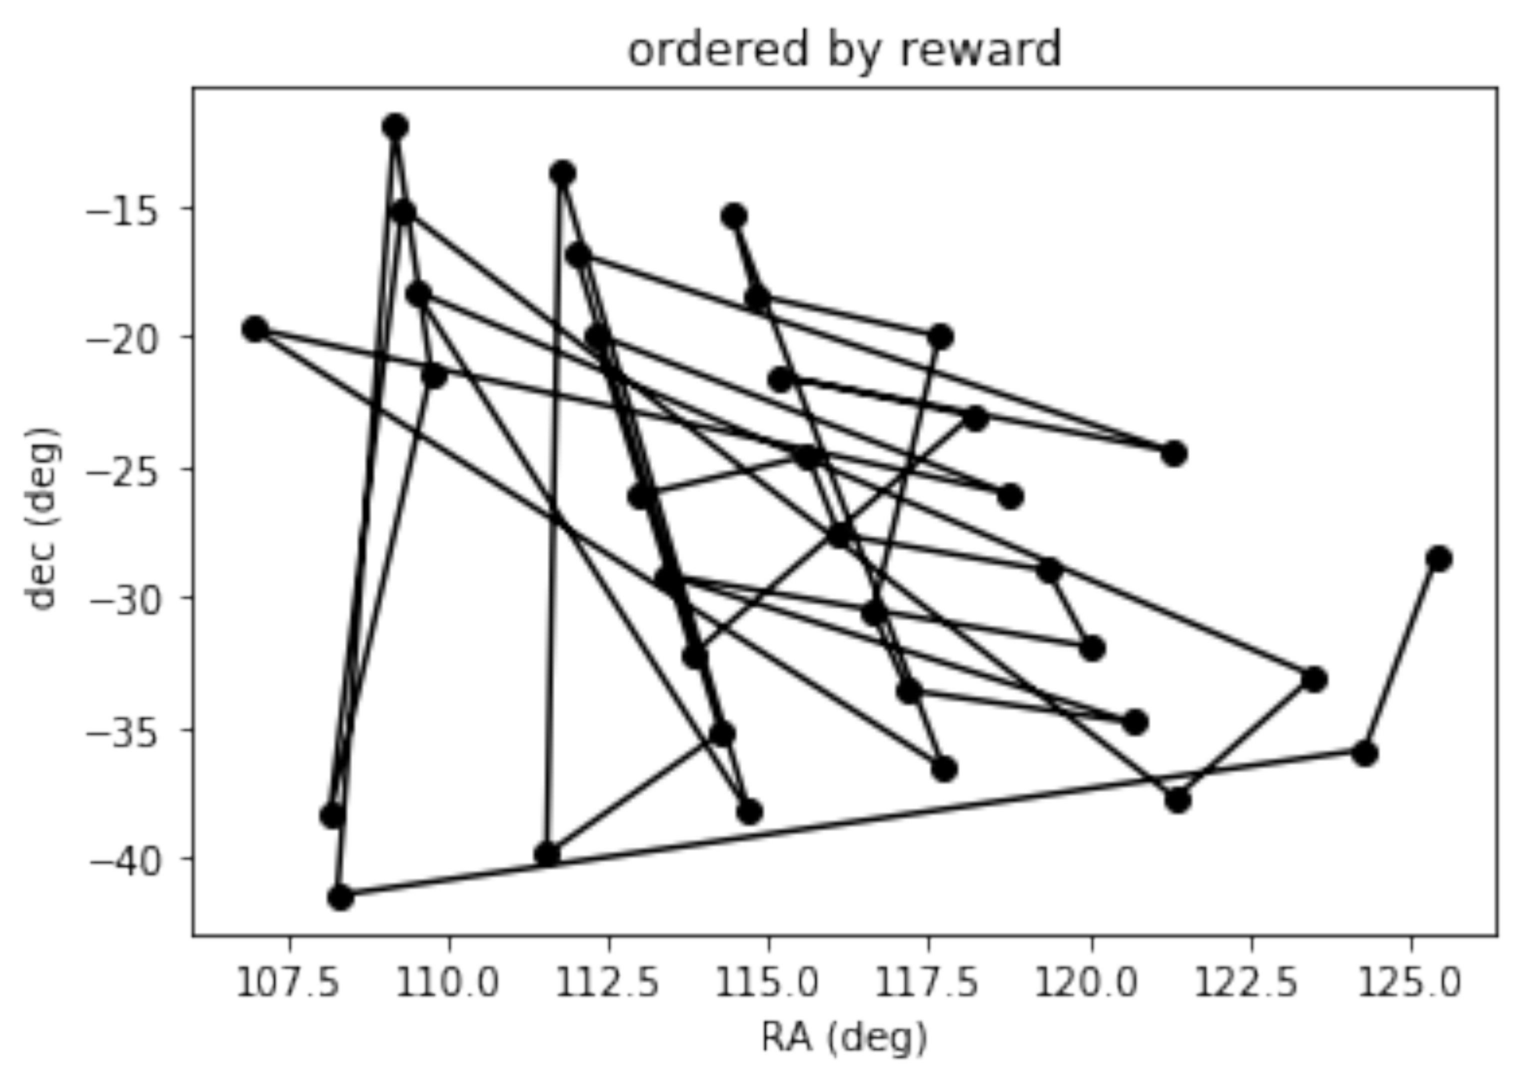

ordered by reward
dec (deg)
RA (deg)
−15
−20
−25
−30
−35
−40
107.5
110.0
112.5
115.0
117.5
120.0
122.5
125.0

Usually, we will observe the path, swap filters, and repeat.

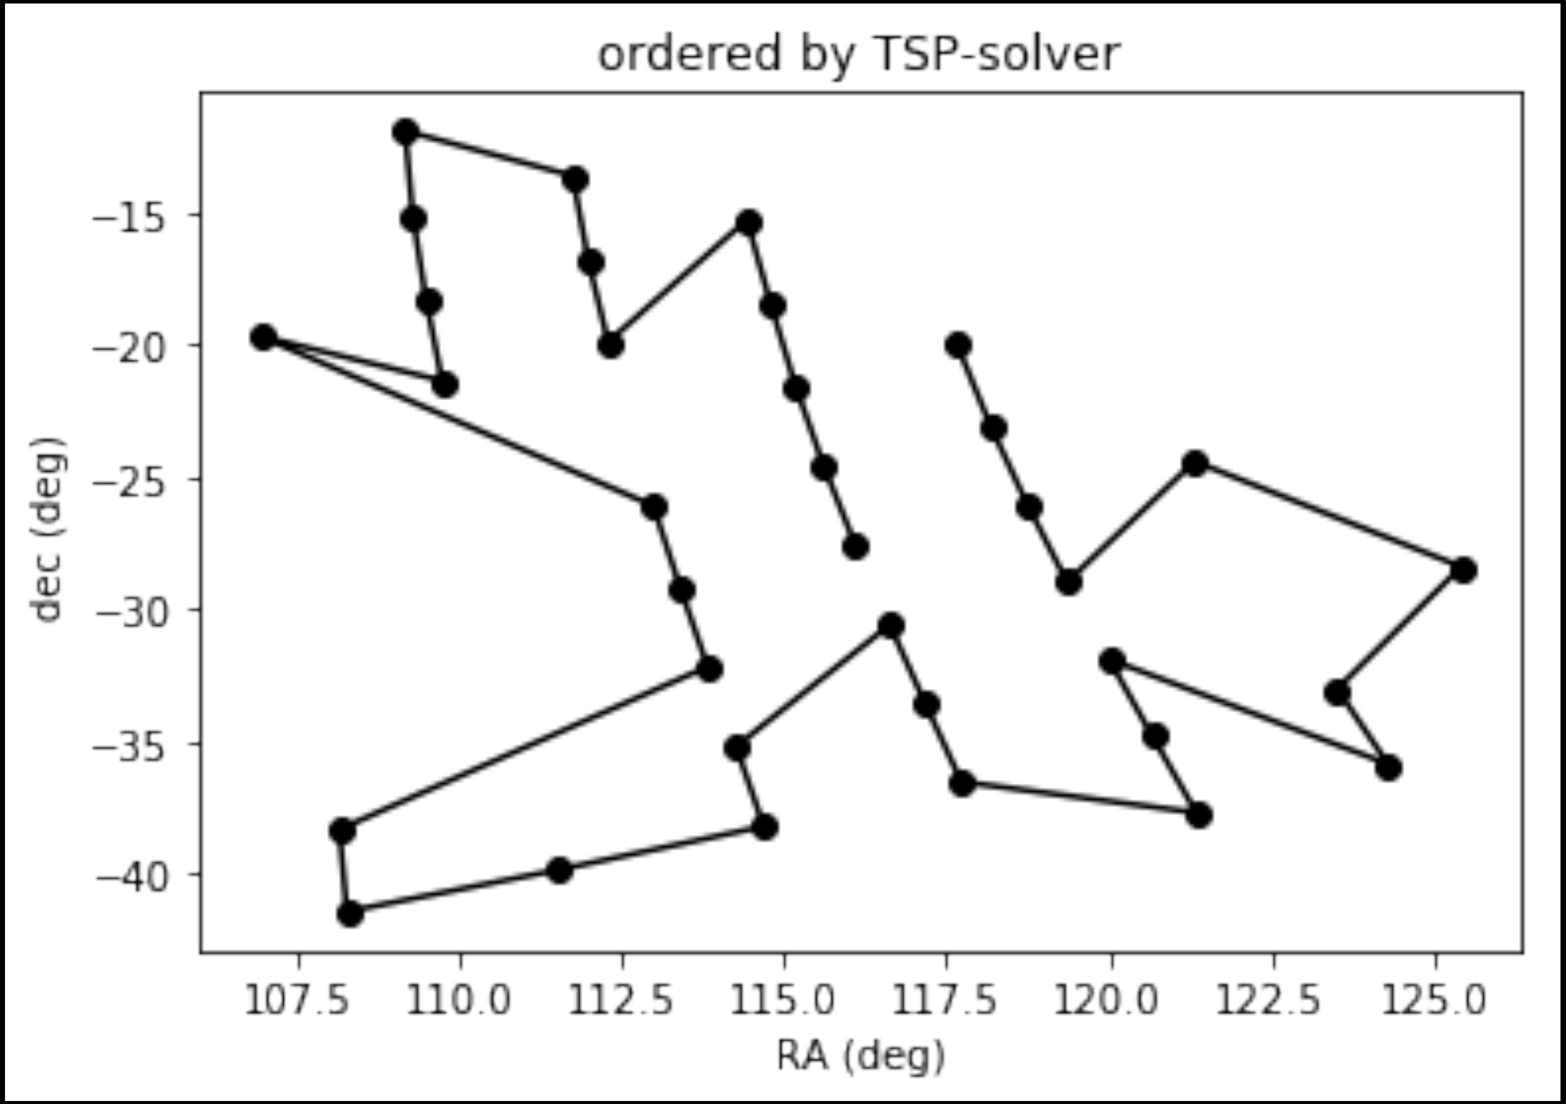

MDP summary

- Compute Basis Functions from the current conditions and observing history
- Compute a Reward Function (sum of weighted Basis Functions). Reward Function in 6 HEALpix maps.
- Pass the Reward Function to a Decision Function (which outputs a list of observations to be executed)
- (optionally) Pass the list of observations to a "detailer", e.g., to specify the camera rotator angle.

Lots of flexibility on changing the telescope behavior

No fixed fields, we track progress at higher resolution
Not much "looking ahead", we are not planning days/weeks/years into the future
No master list of observations to execute

Simulate running the scheduler
for 10 years…

North Ecliptic Spur

Wide Fast Deep (WFD)
18,000 sq degrees

DDFs

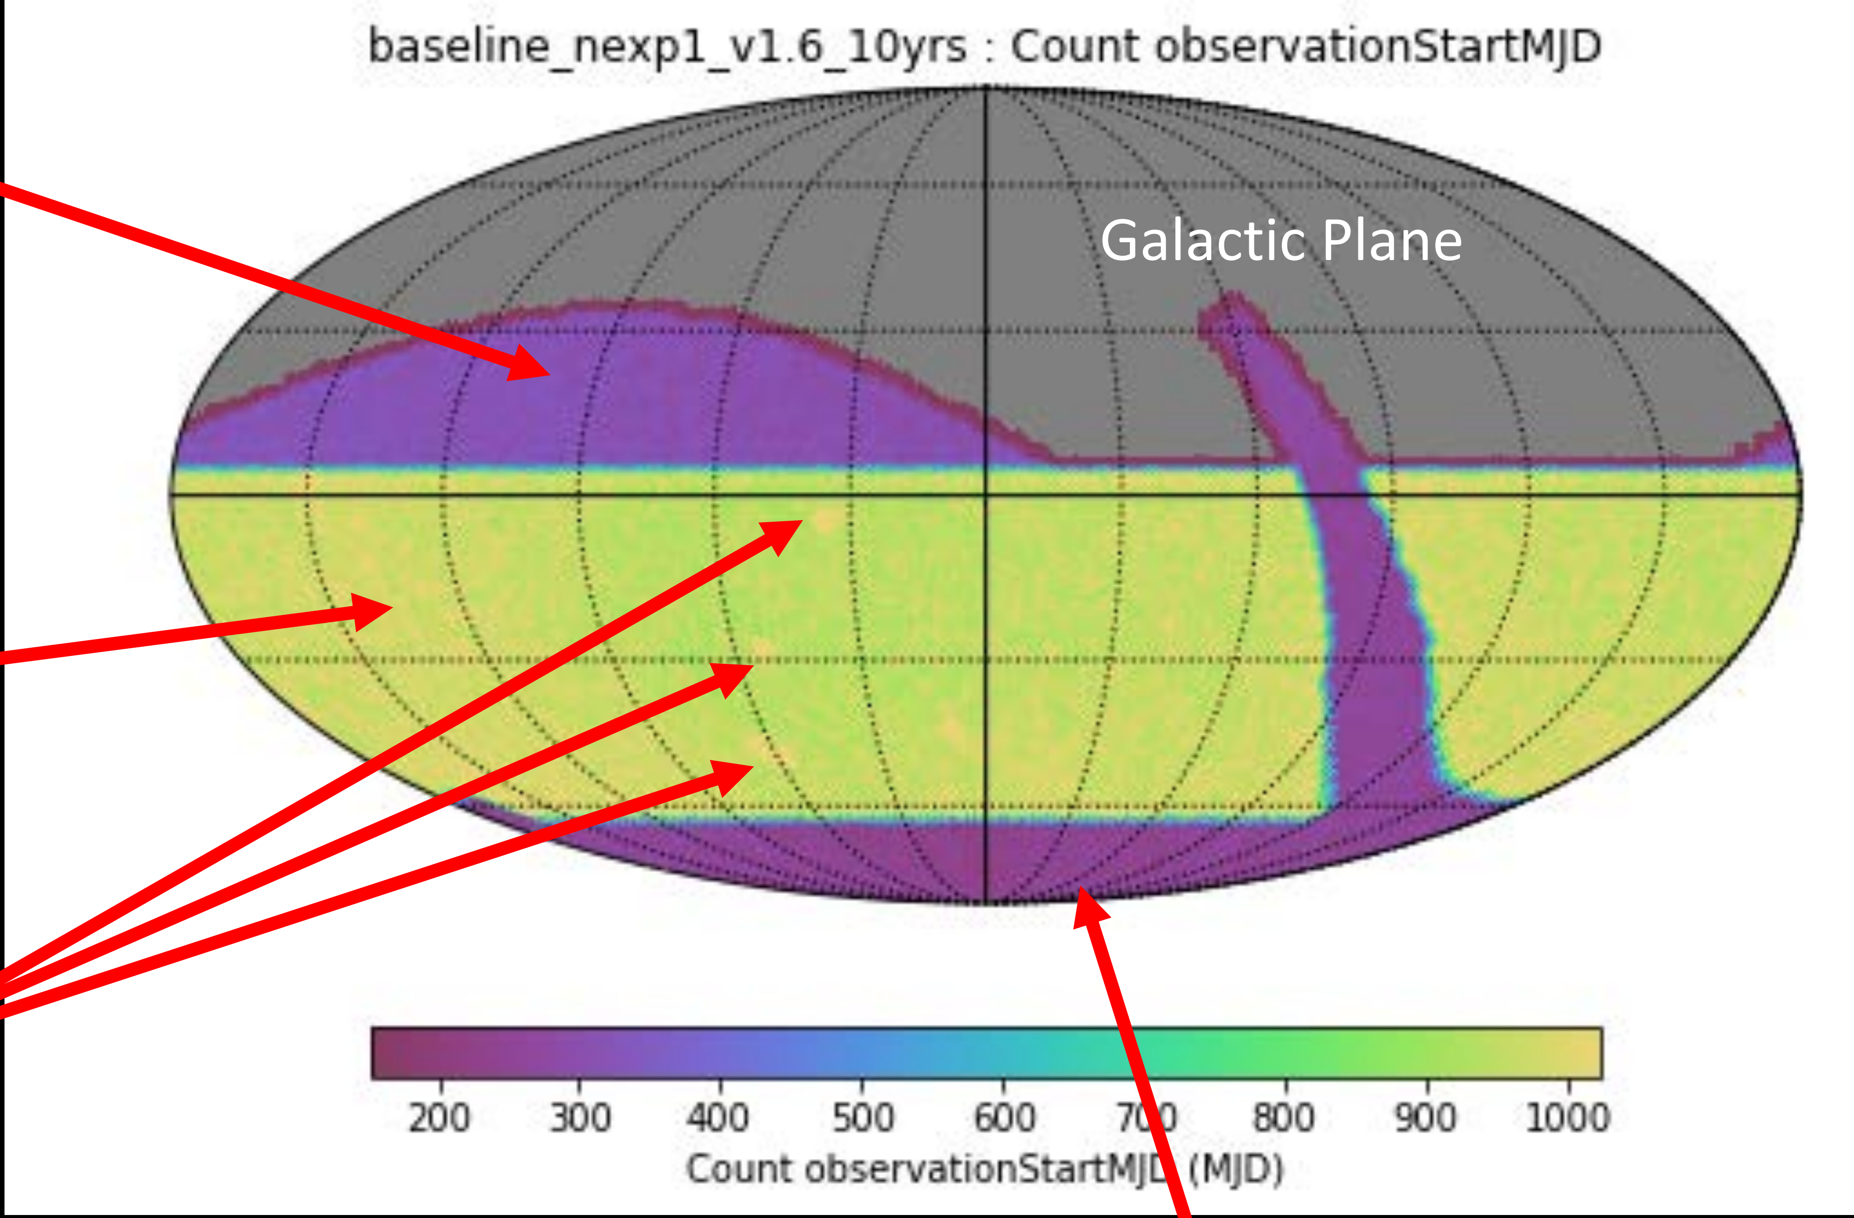

South Celestial Pole

The pointing history in alt/az. Note the log stretch, we are well concentrated at low airmass

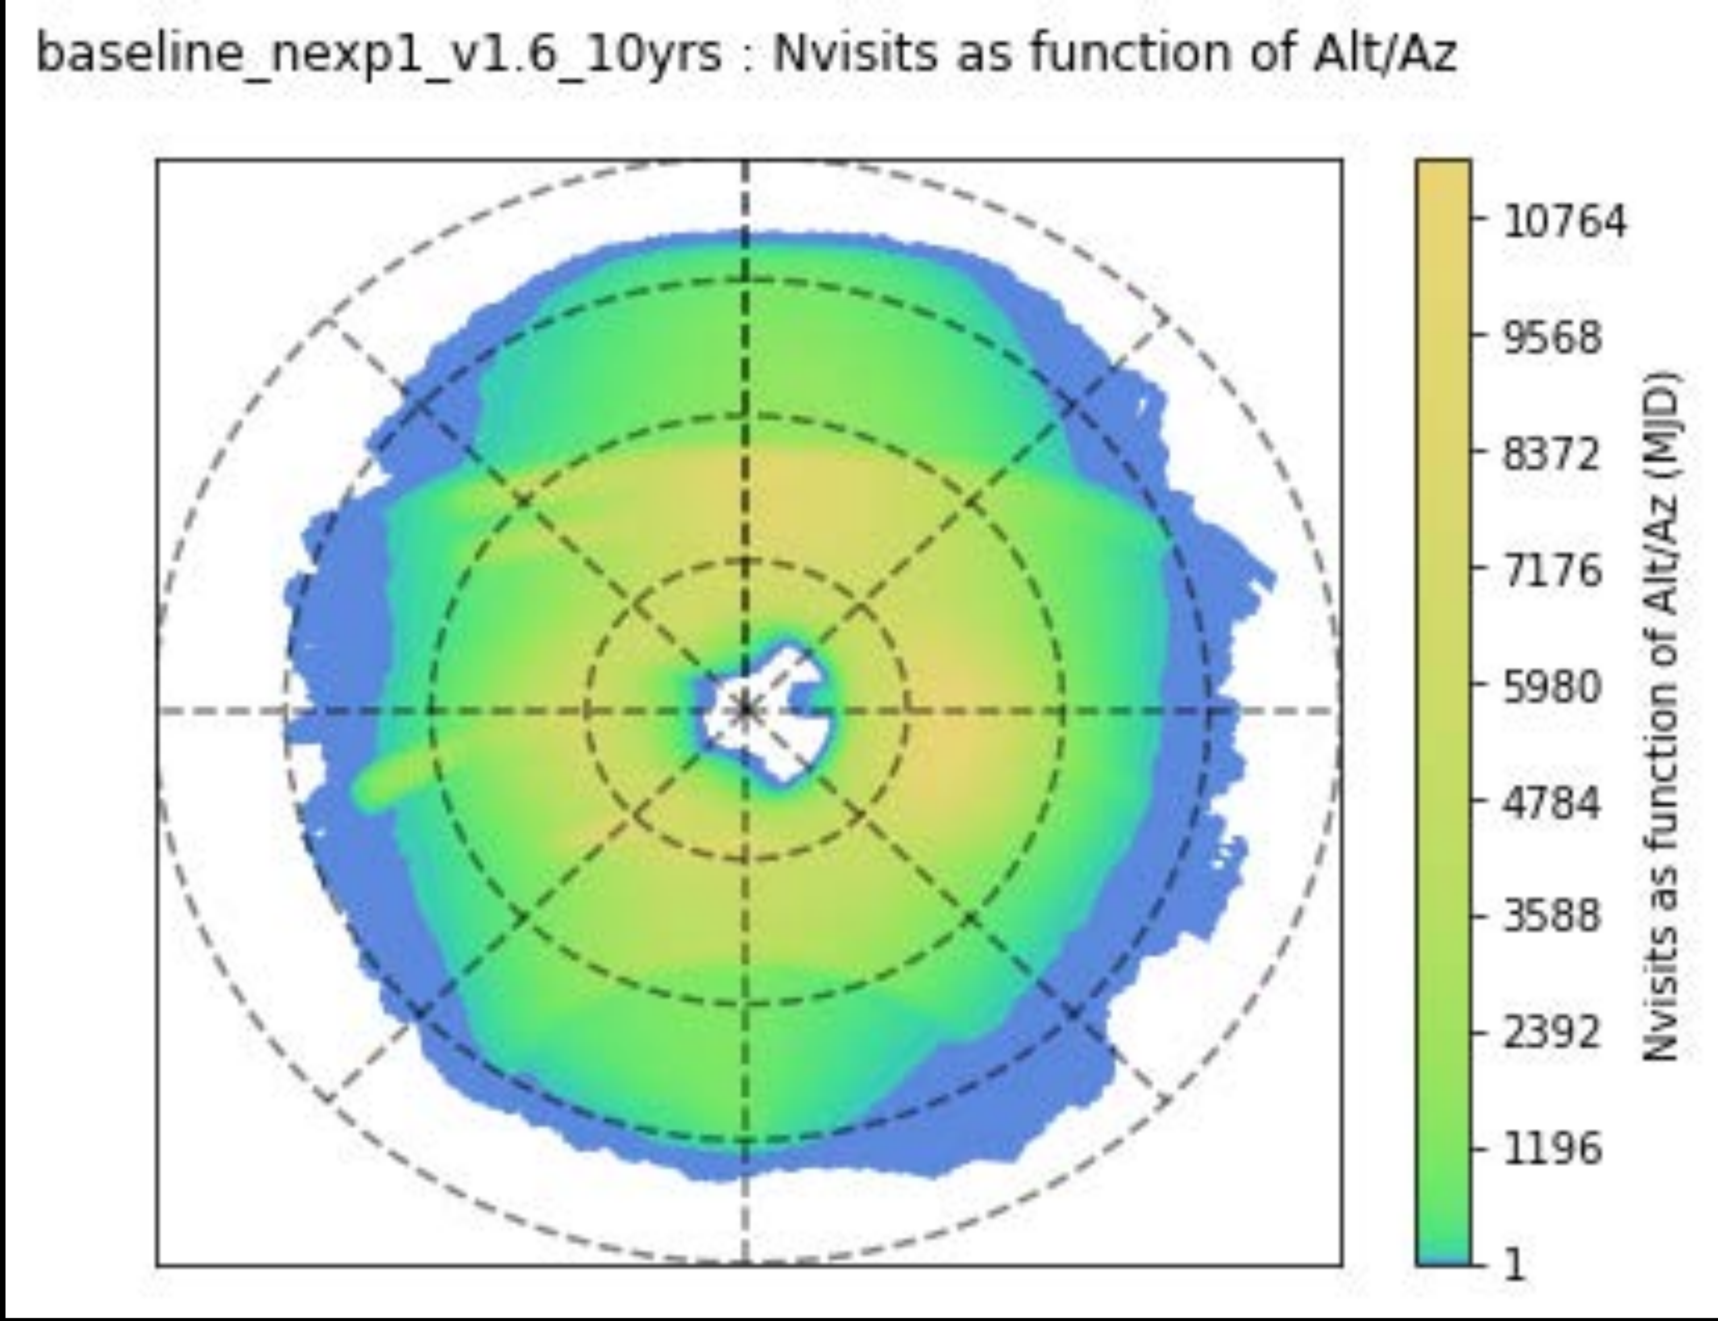

Pointings in Alt/Az with a linear stretch

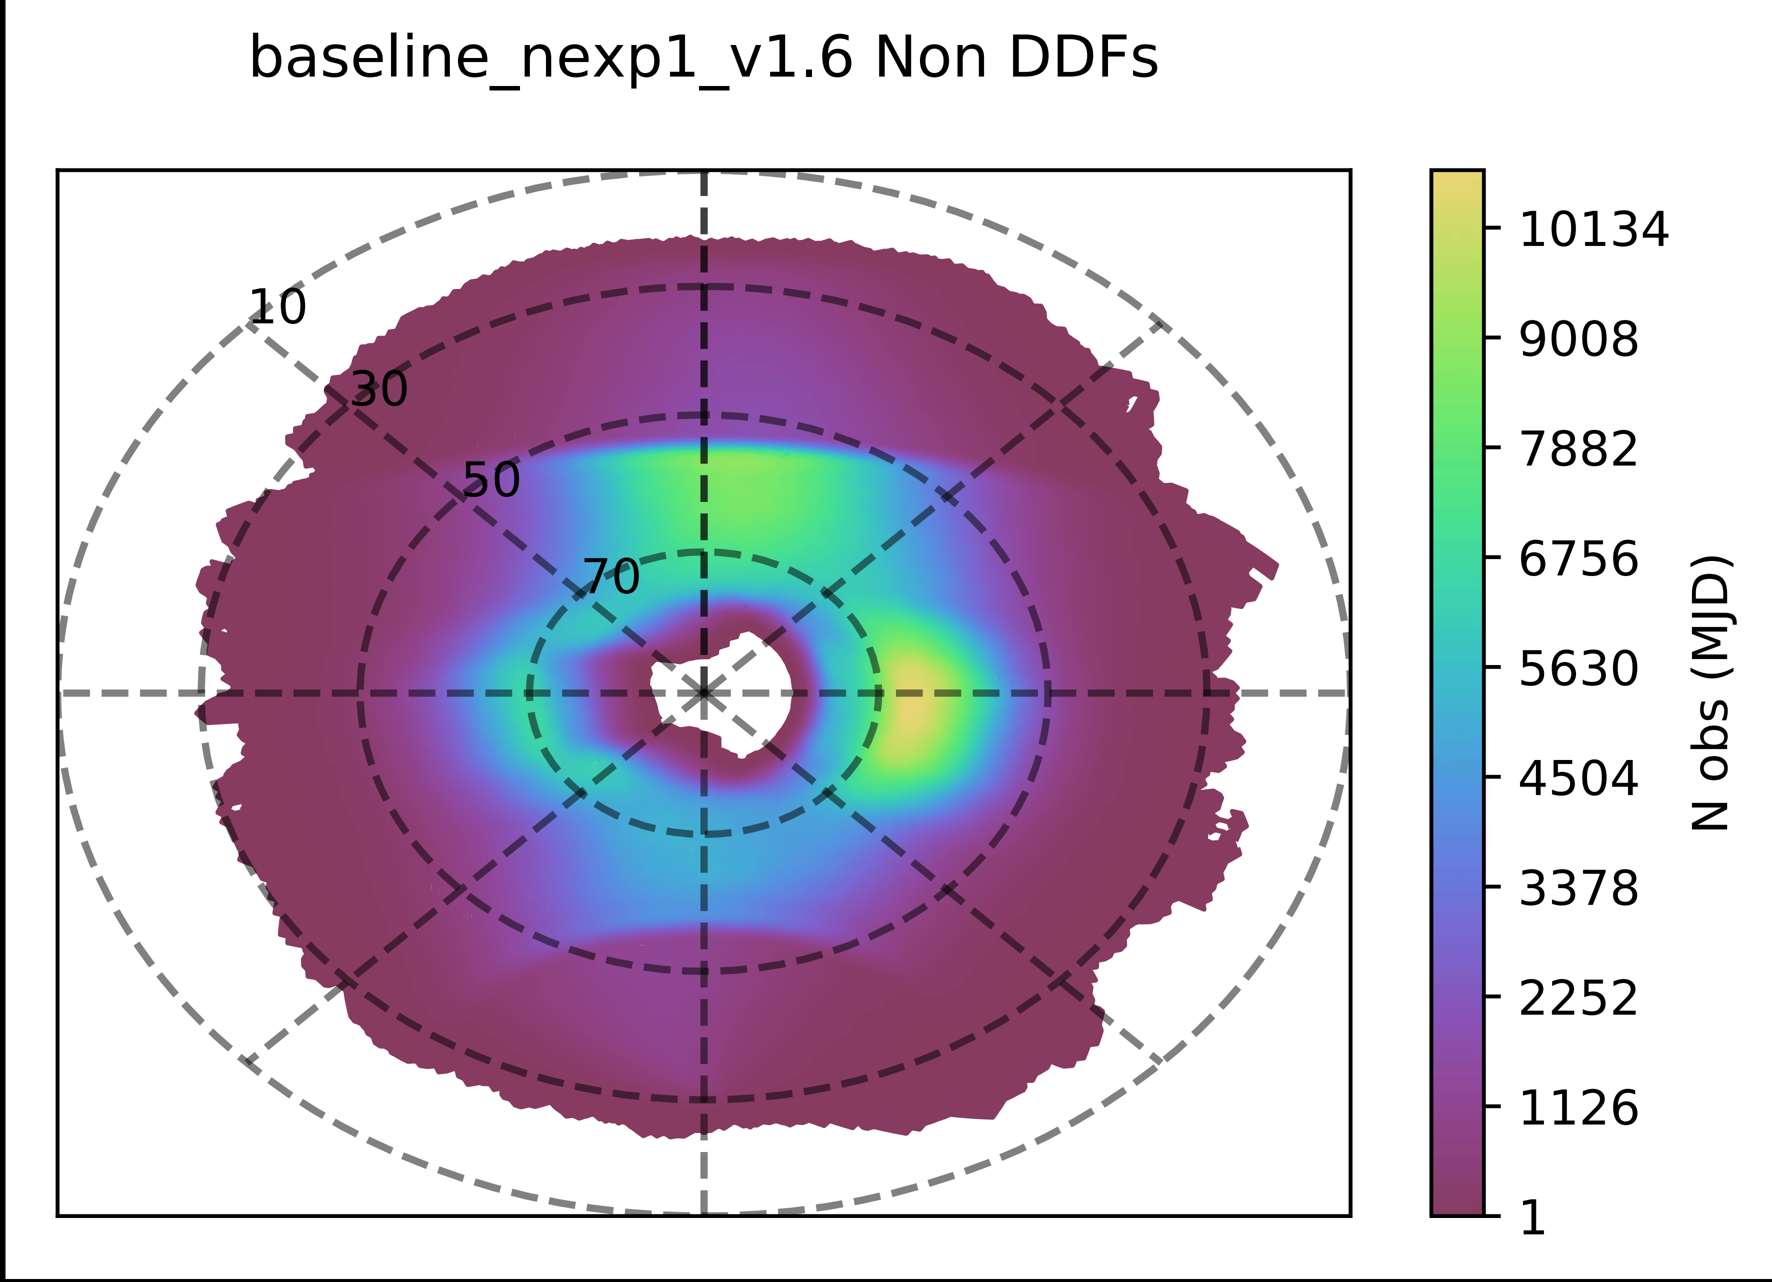

Year 1, color-coded by what filter is being used. Red filters dominate at full moon. White stripes are weather and downtime.

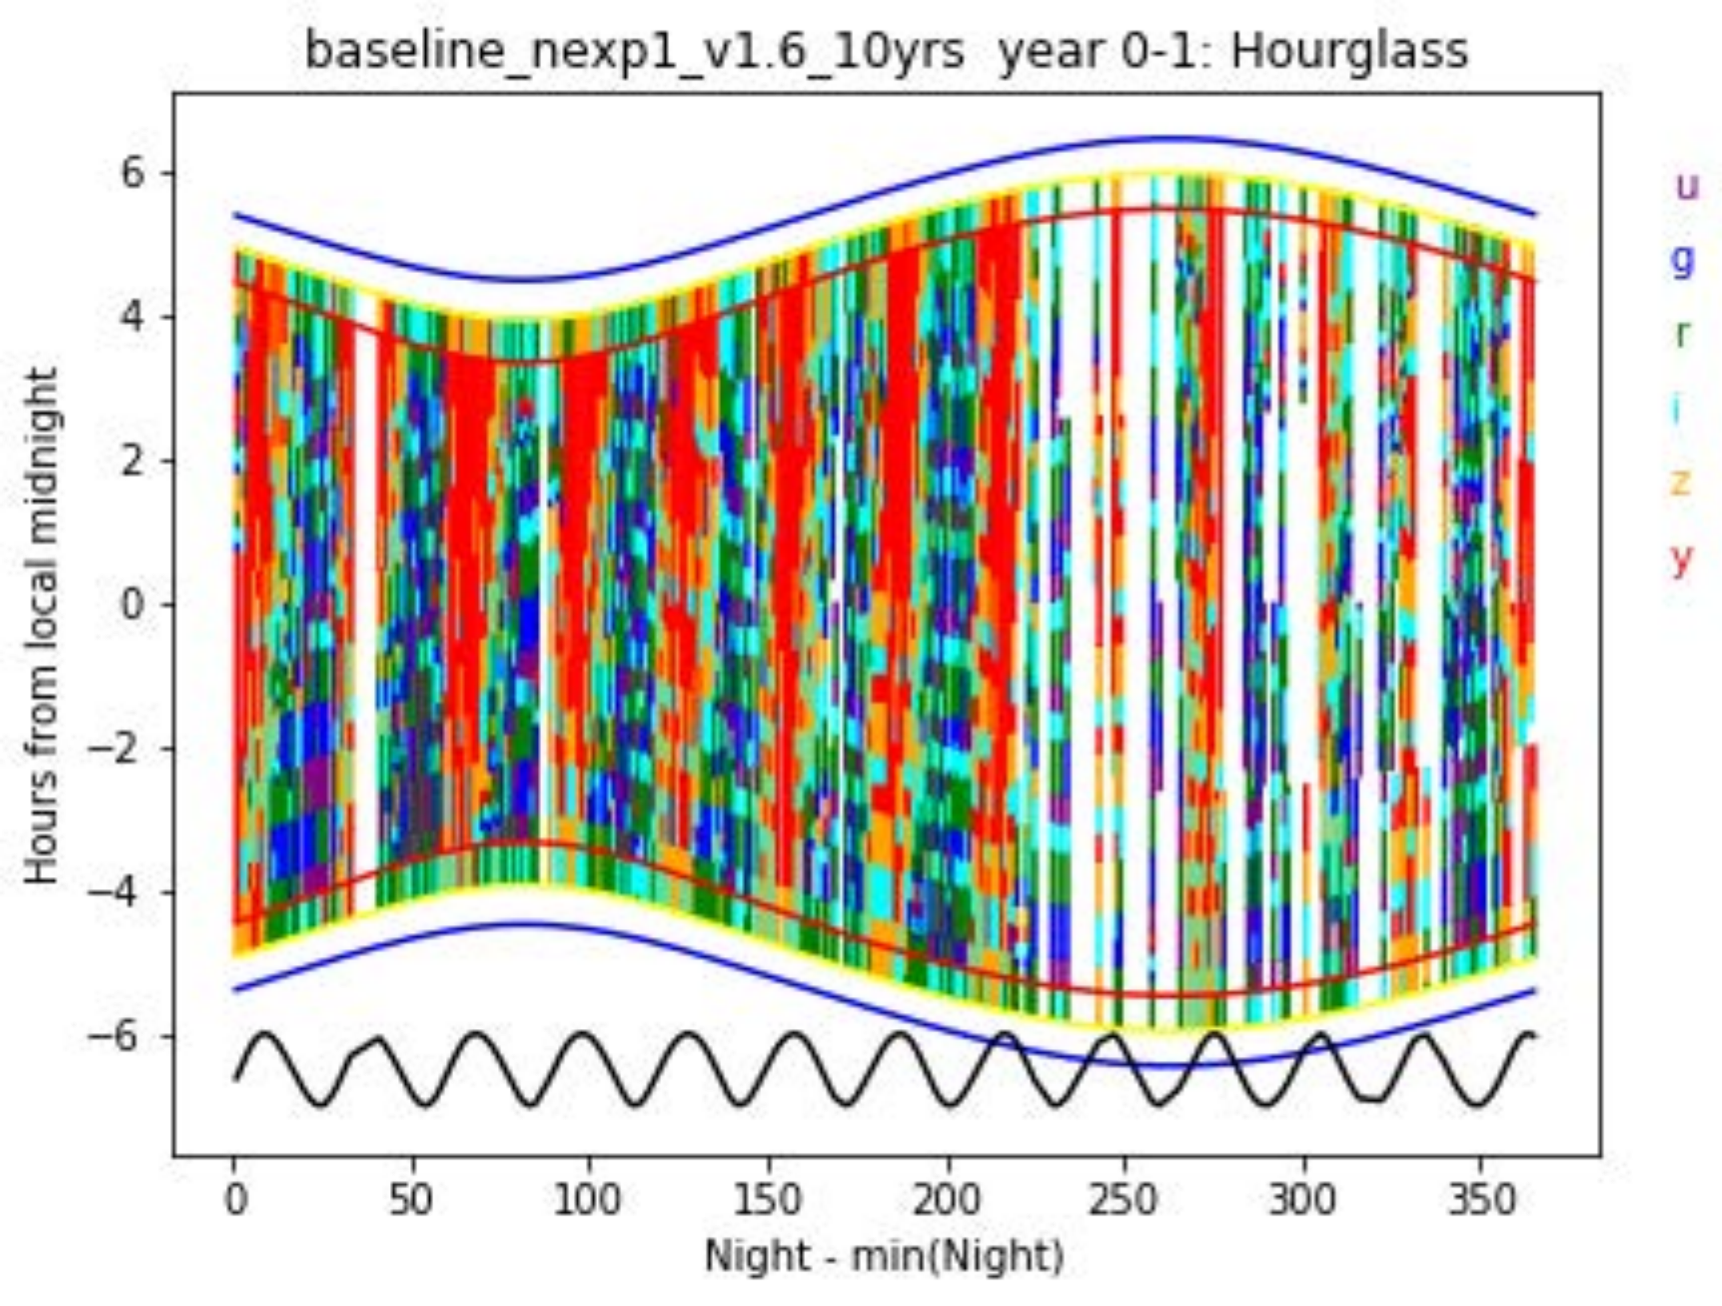

Let’s look at one point in the sky

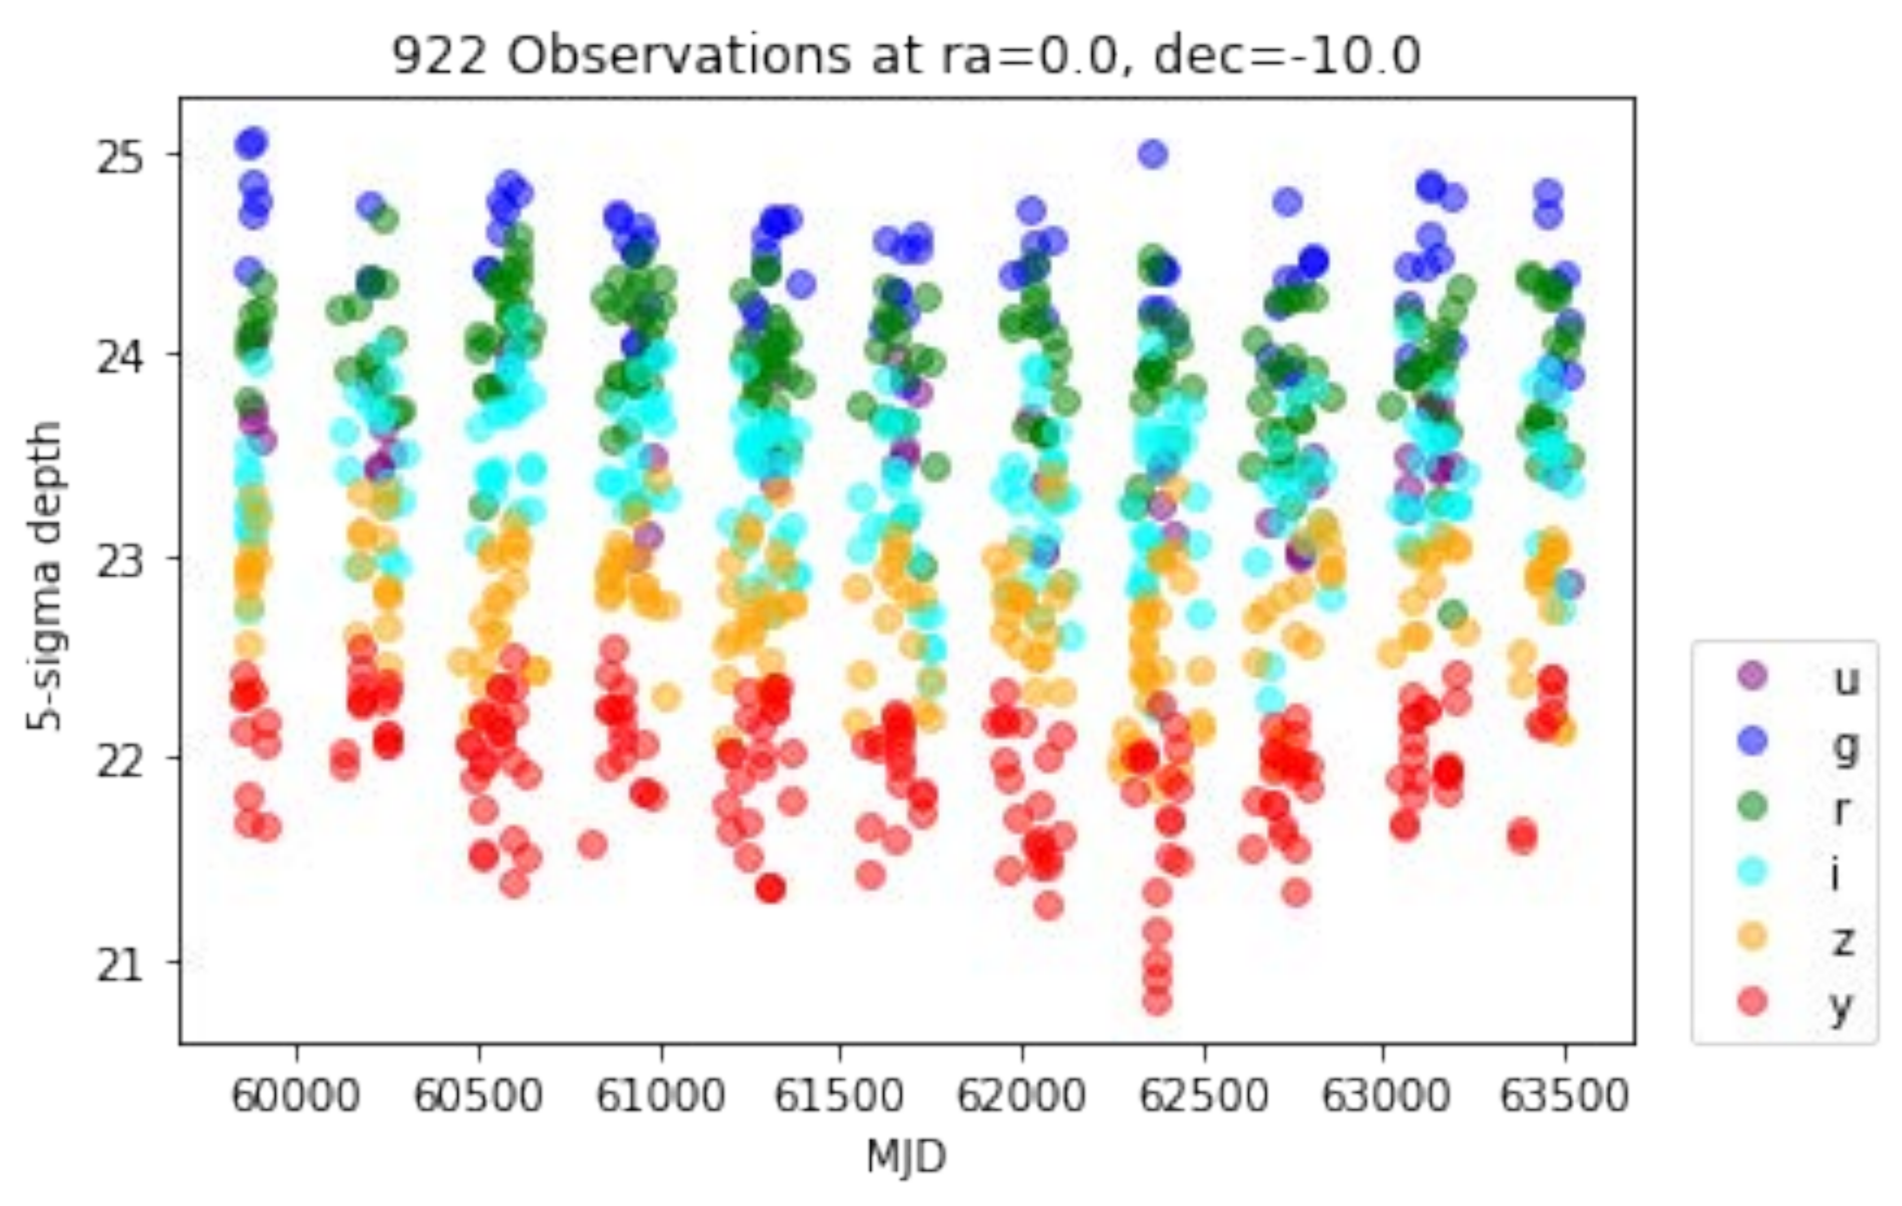

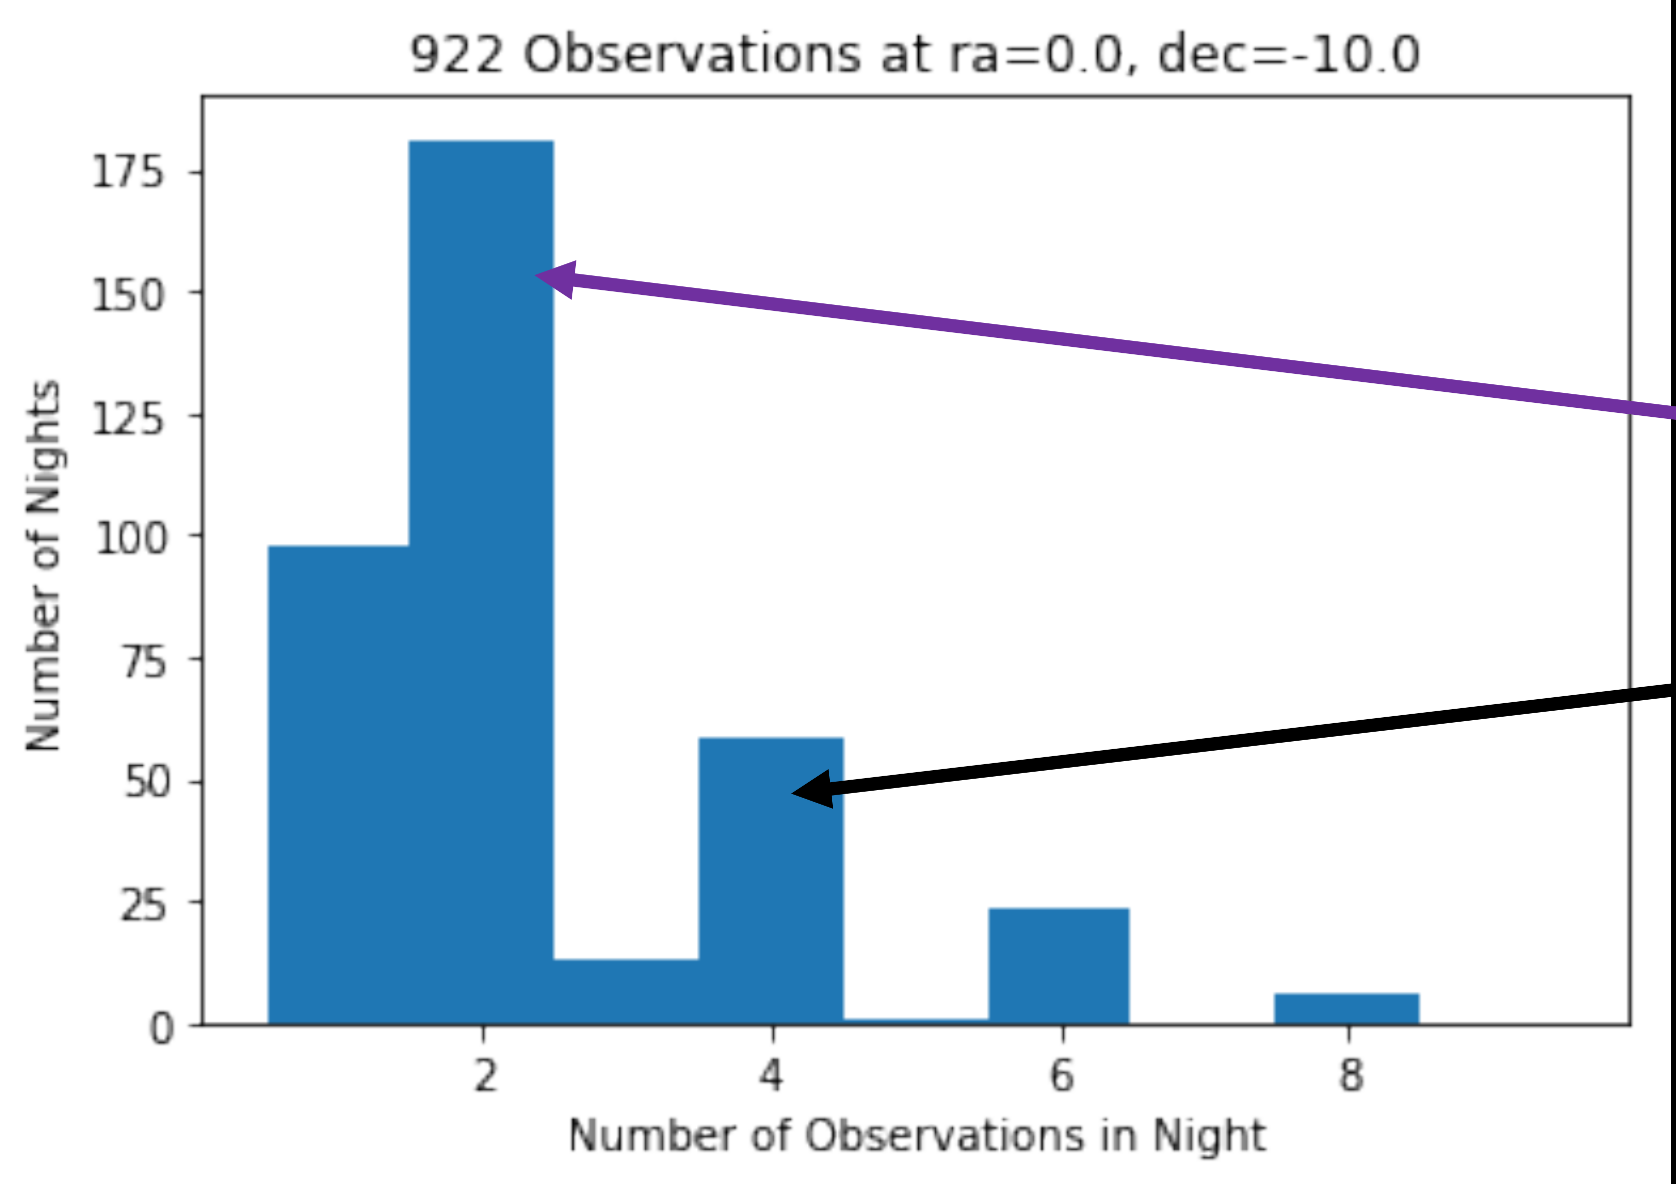

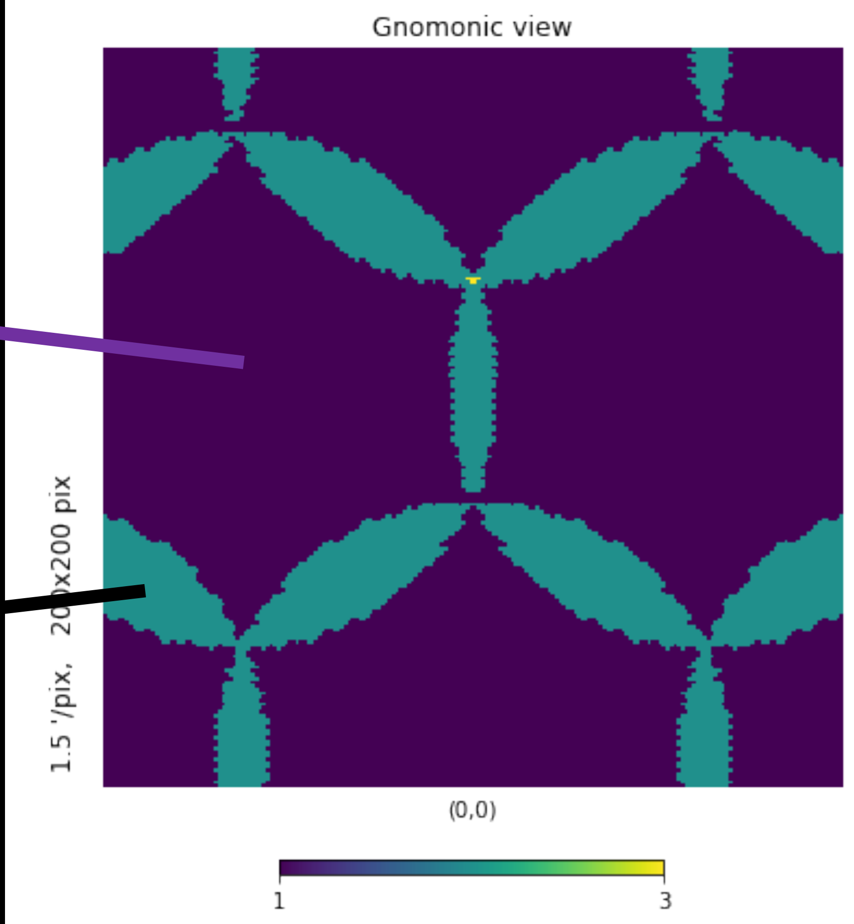

Because we randomize the tessellation and take observations in pairs, most of the time a point in WFD is observed twice in a night. Sometimes it falls in an overlap region and gets 4 observations in a night.

Technical stats

- ~6.5 hours (single core) to simulate 10 years
- 2.2 million observations, recorded to sqlite database (650M)
- Not too tough to install and run, but the pre-computed sky brightness files are 95G for the full 10 years
- “The code is the config”, ~300 lines of python define a scheduler and simulation

Been running on Hyak cluster at UW

Code is on github, help yourself
https://github.com/lsst/sims_featureScheduler

- ~50 Cadence White Papers received in Nov 2018
- Science Advisory Committee gave us a list of suggested runs in April 2019
- We’ve now released a few hundred simulations as we’ve worked through that list

Release Announcements at https://community.lsst.org/
Simulation analysis posted at http://astro-lsst-01.astro.washington.edu:8080/

Huge report on our simulations so far at: https://pstn-051.lsst.io/

Various footprint variations

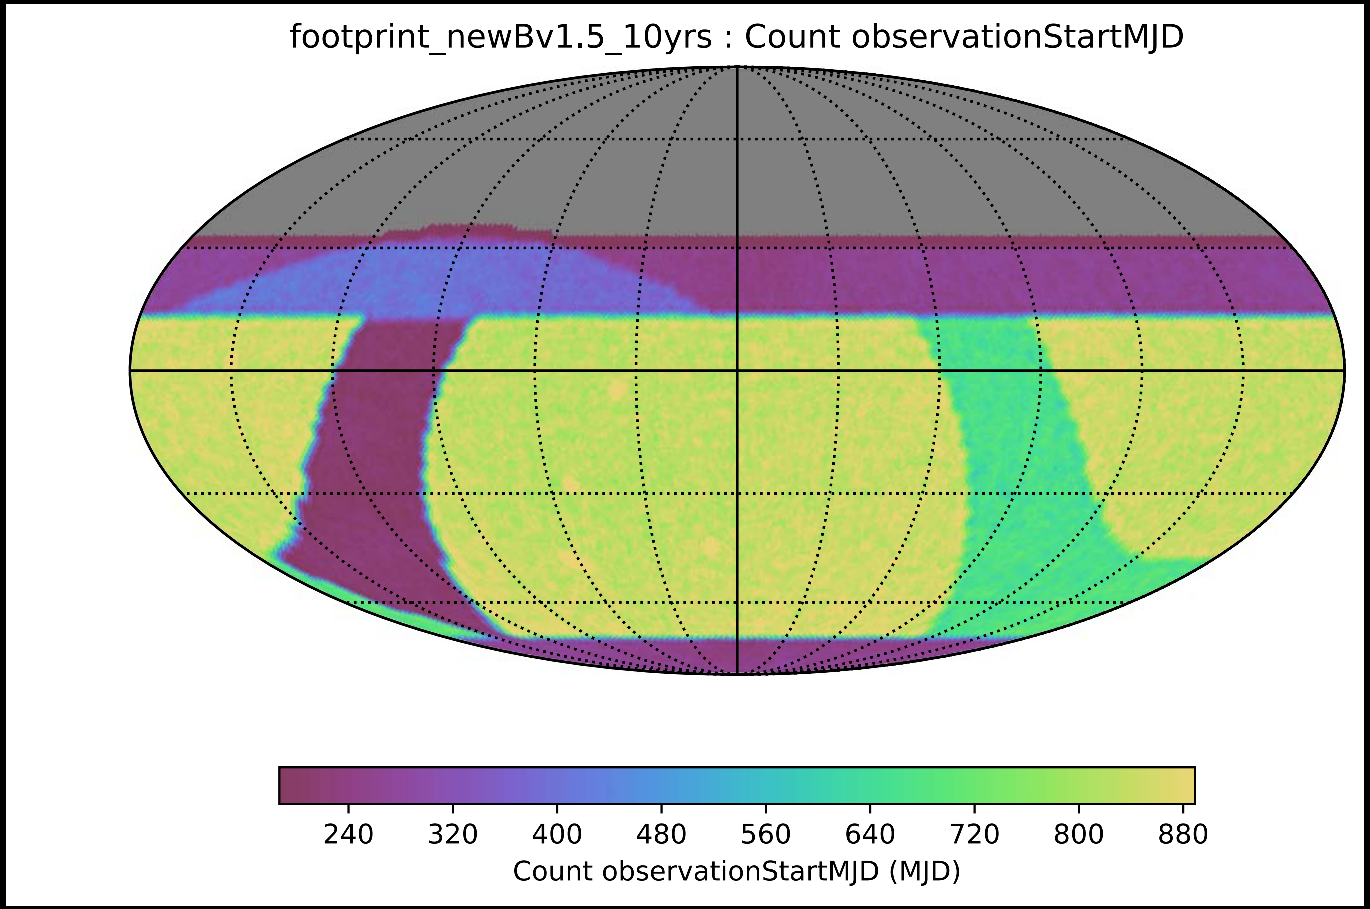

Baseline

Rolling

Number of observations after 10 years (DDF visits removed)

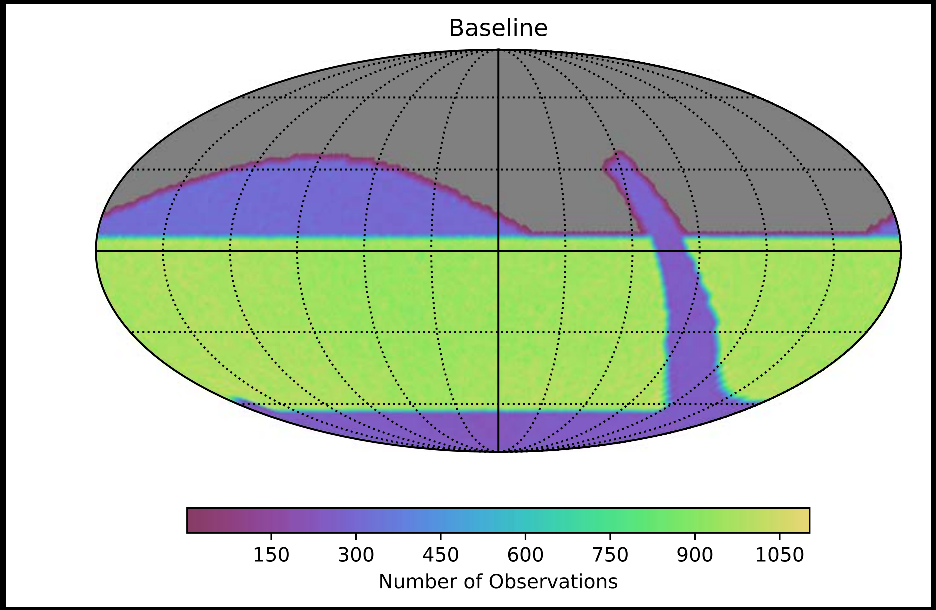

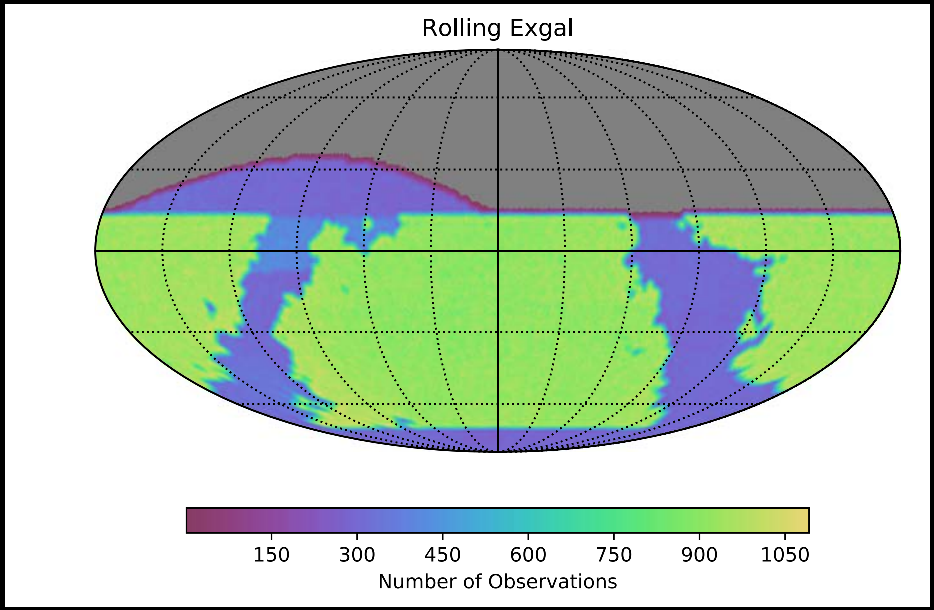

Year 3.5-4.5

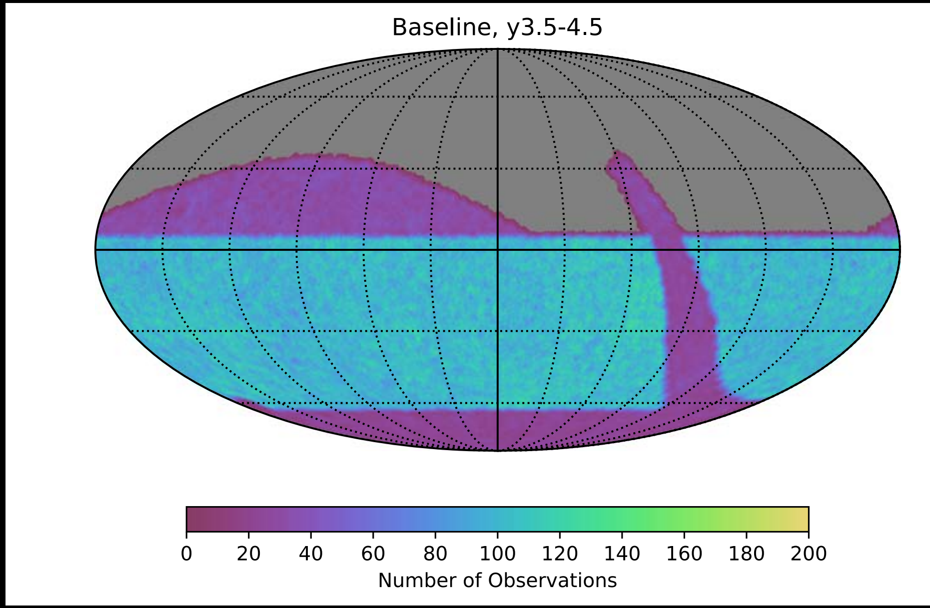

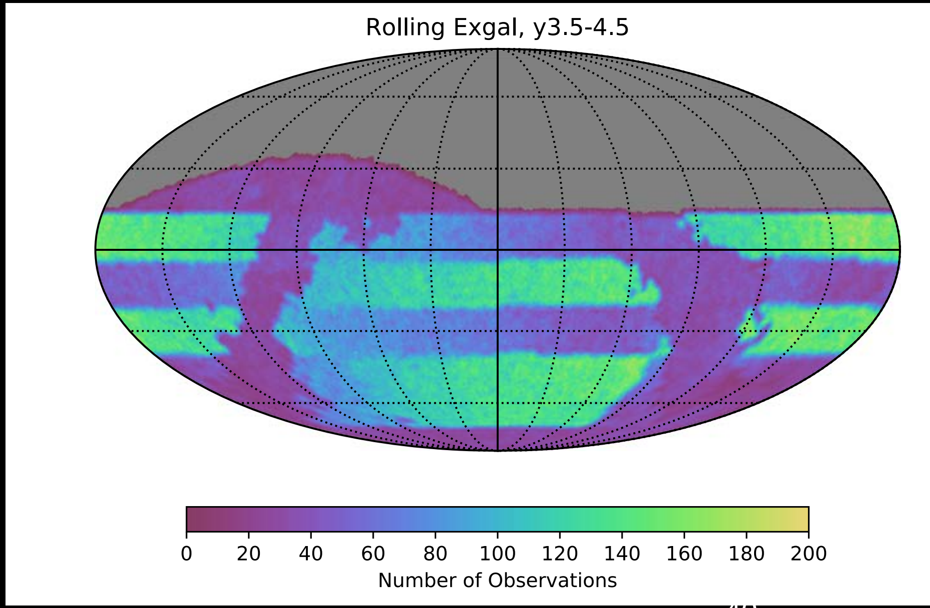

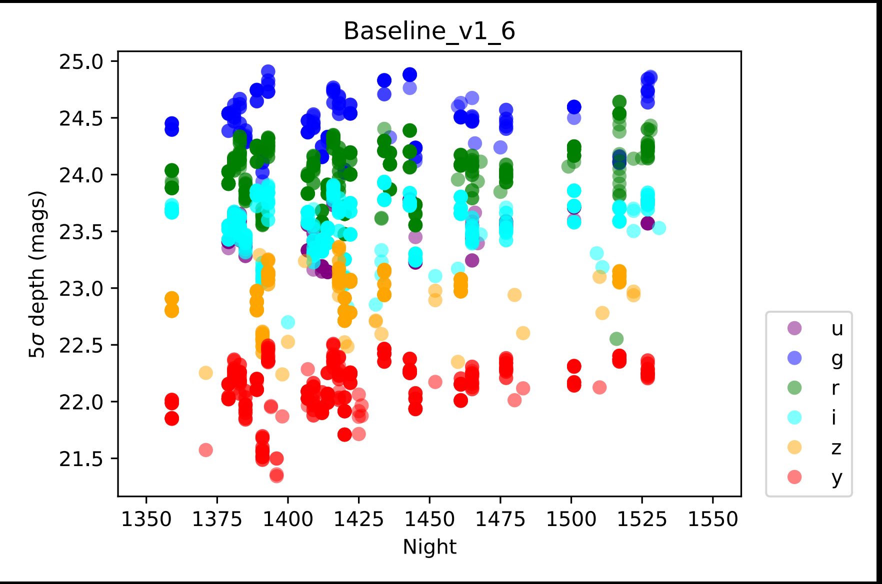

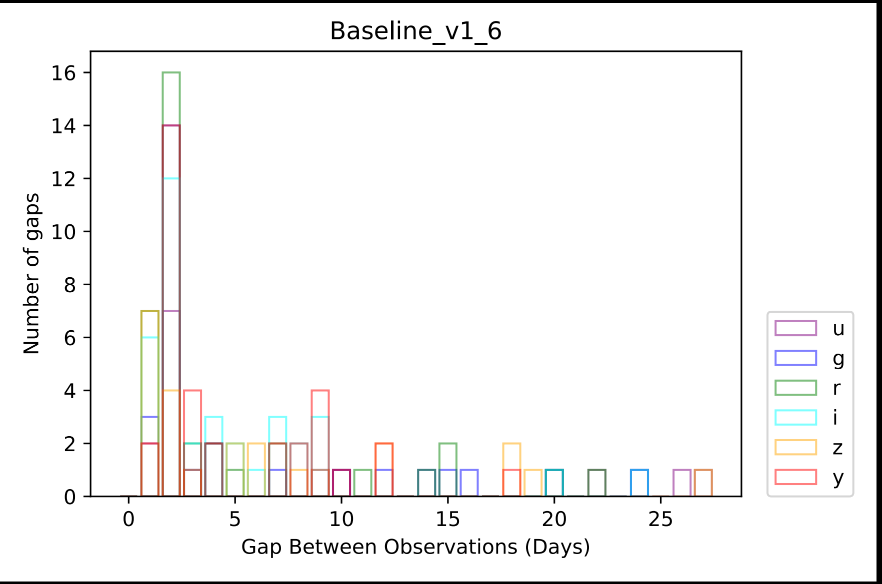

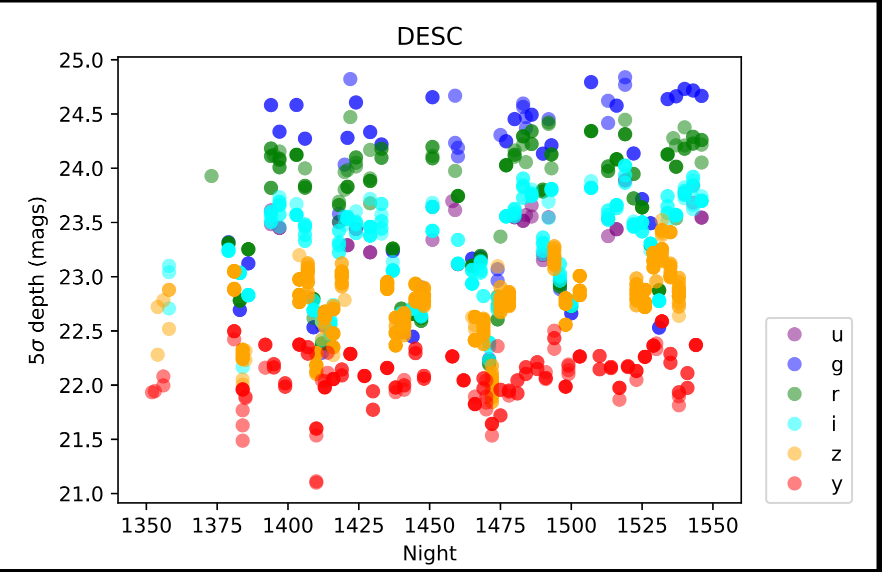

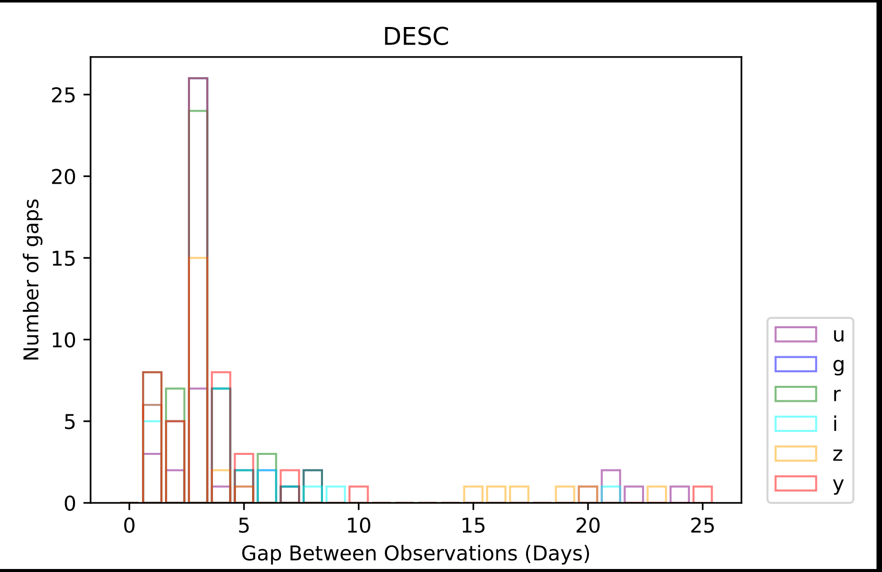

Several variations on how the DDFs execute

Analyzing survey performance

MAF is a package we have written to take in a series of observation metadata (ra, dec, filter, mjd, seeing, airmass, etc)

Just some examples here, see other talk for more.

Final proper motion uncertainty. Assume a fiducial star (20th mag, flat SED), compute expected centroid error from seeing and depth of each image. Error propagation to uncertainty in motion.

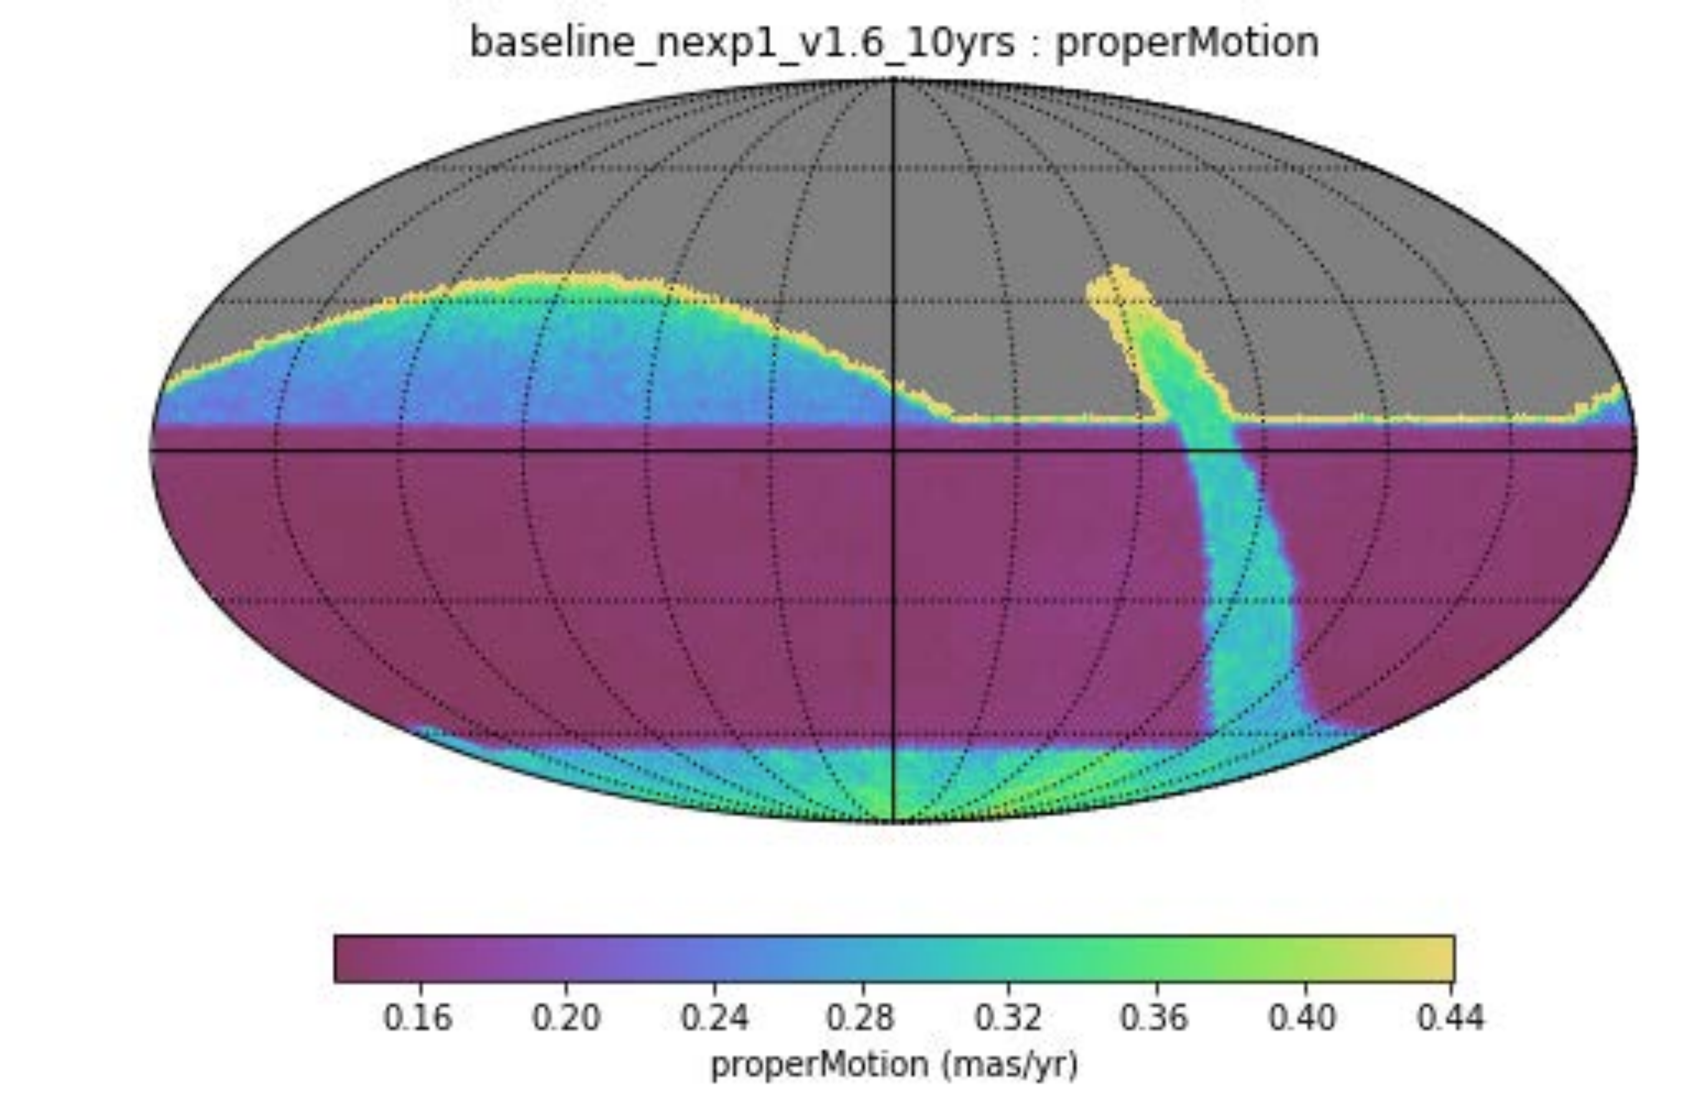

Expected galaxy counts in i, includes dust extinction map in the calculation

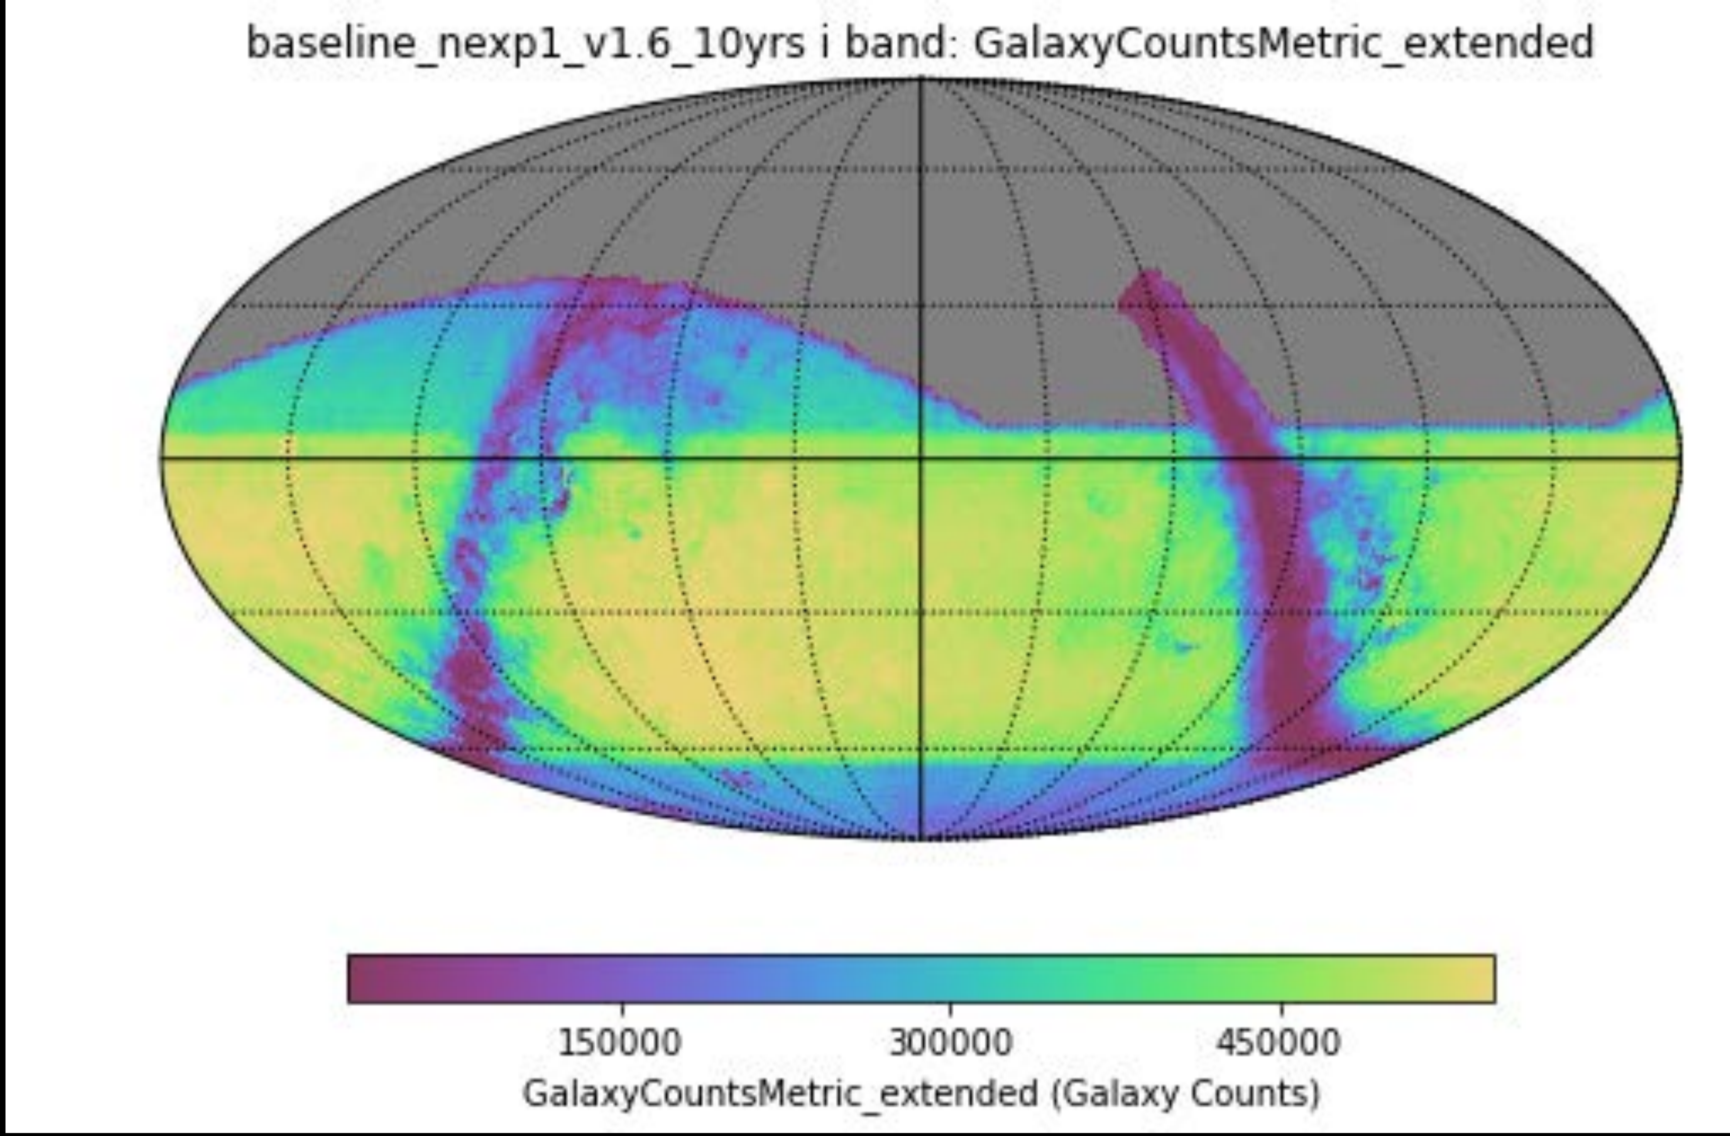

Number of stars. Using TRILEGAL or Galfast galaxy models

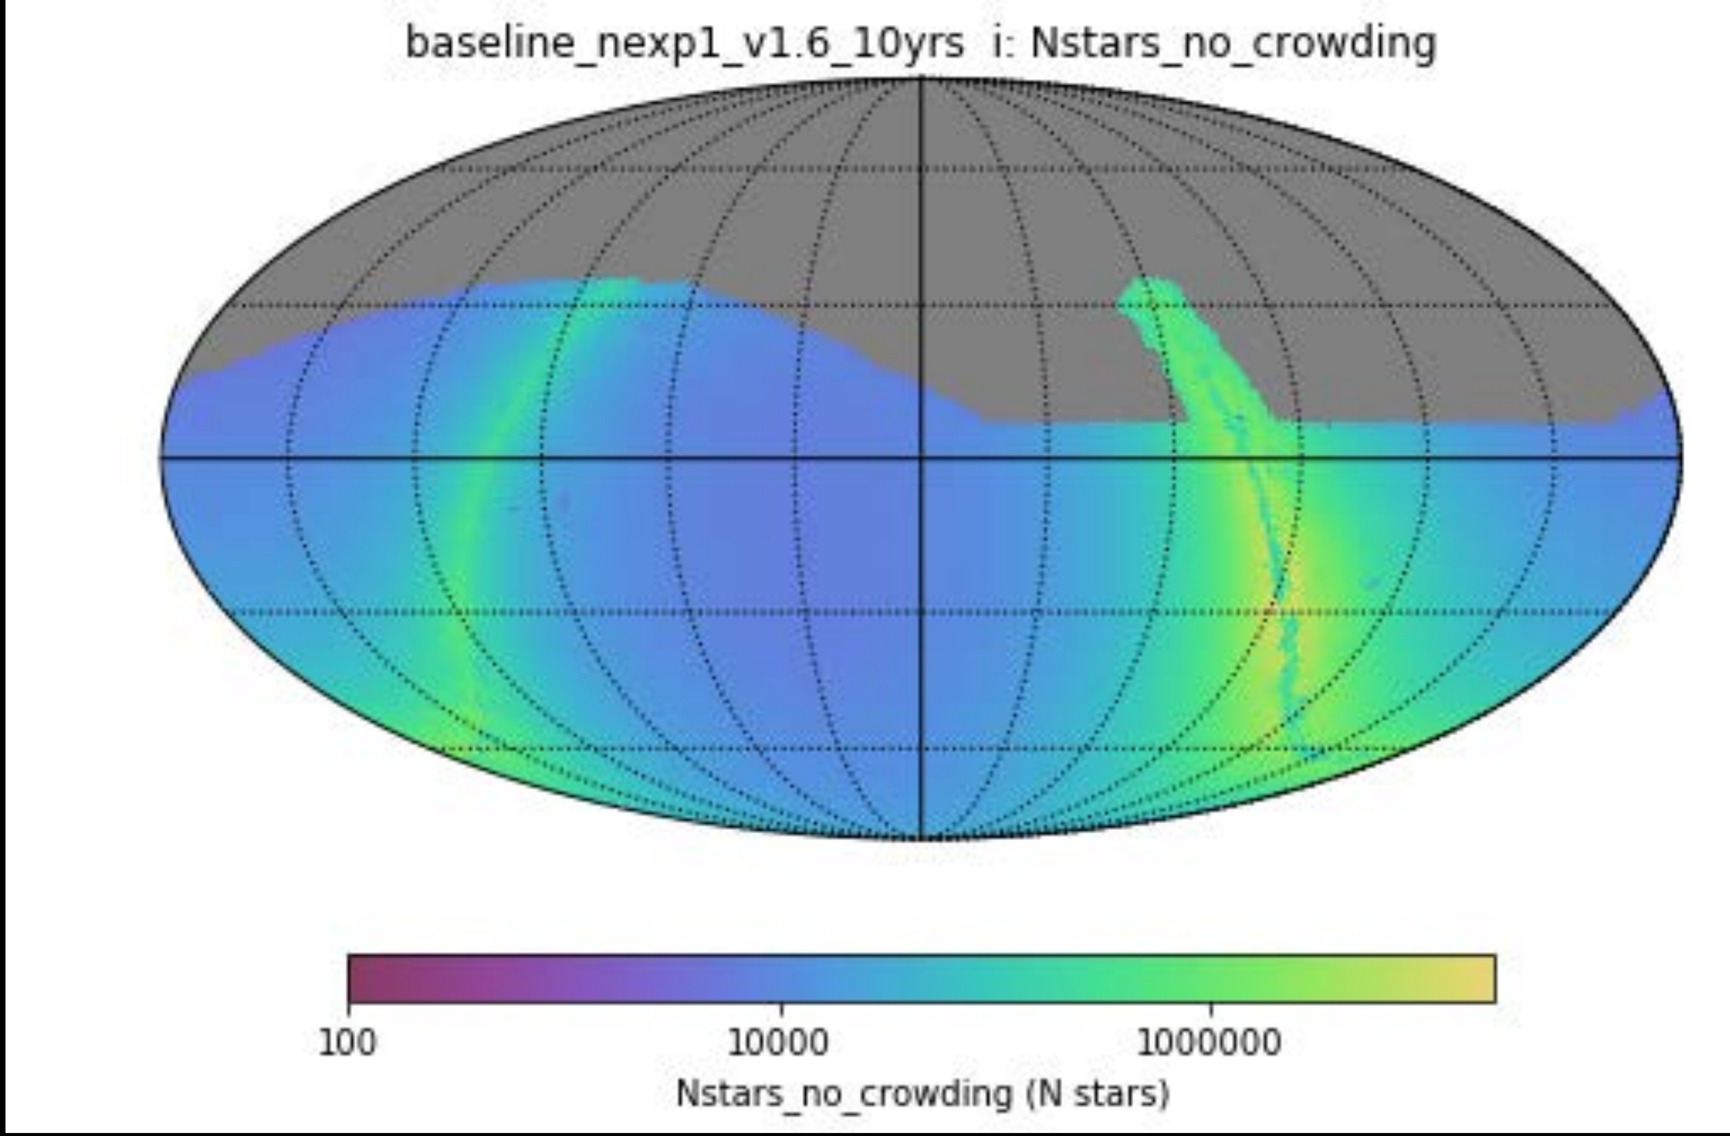

Generate 10k microlensing events and see which ones get recovered

Can see the bulge, LMC, and SMC are in the input distribution.

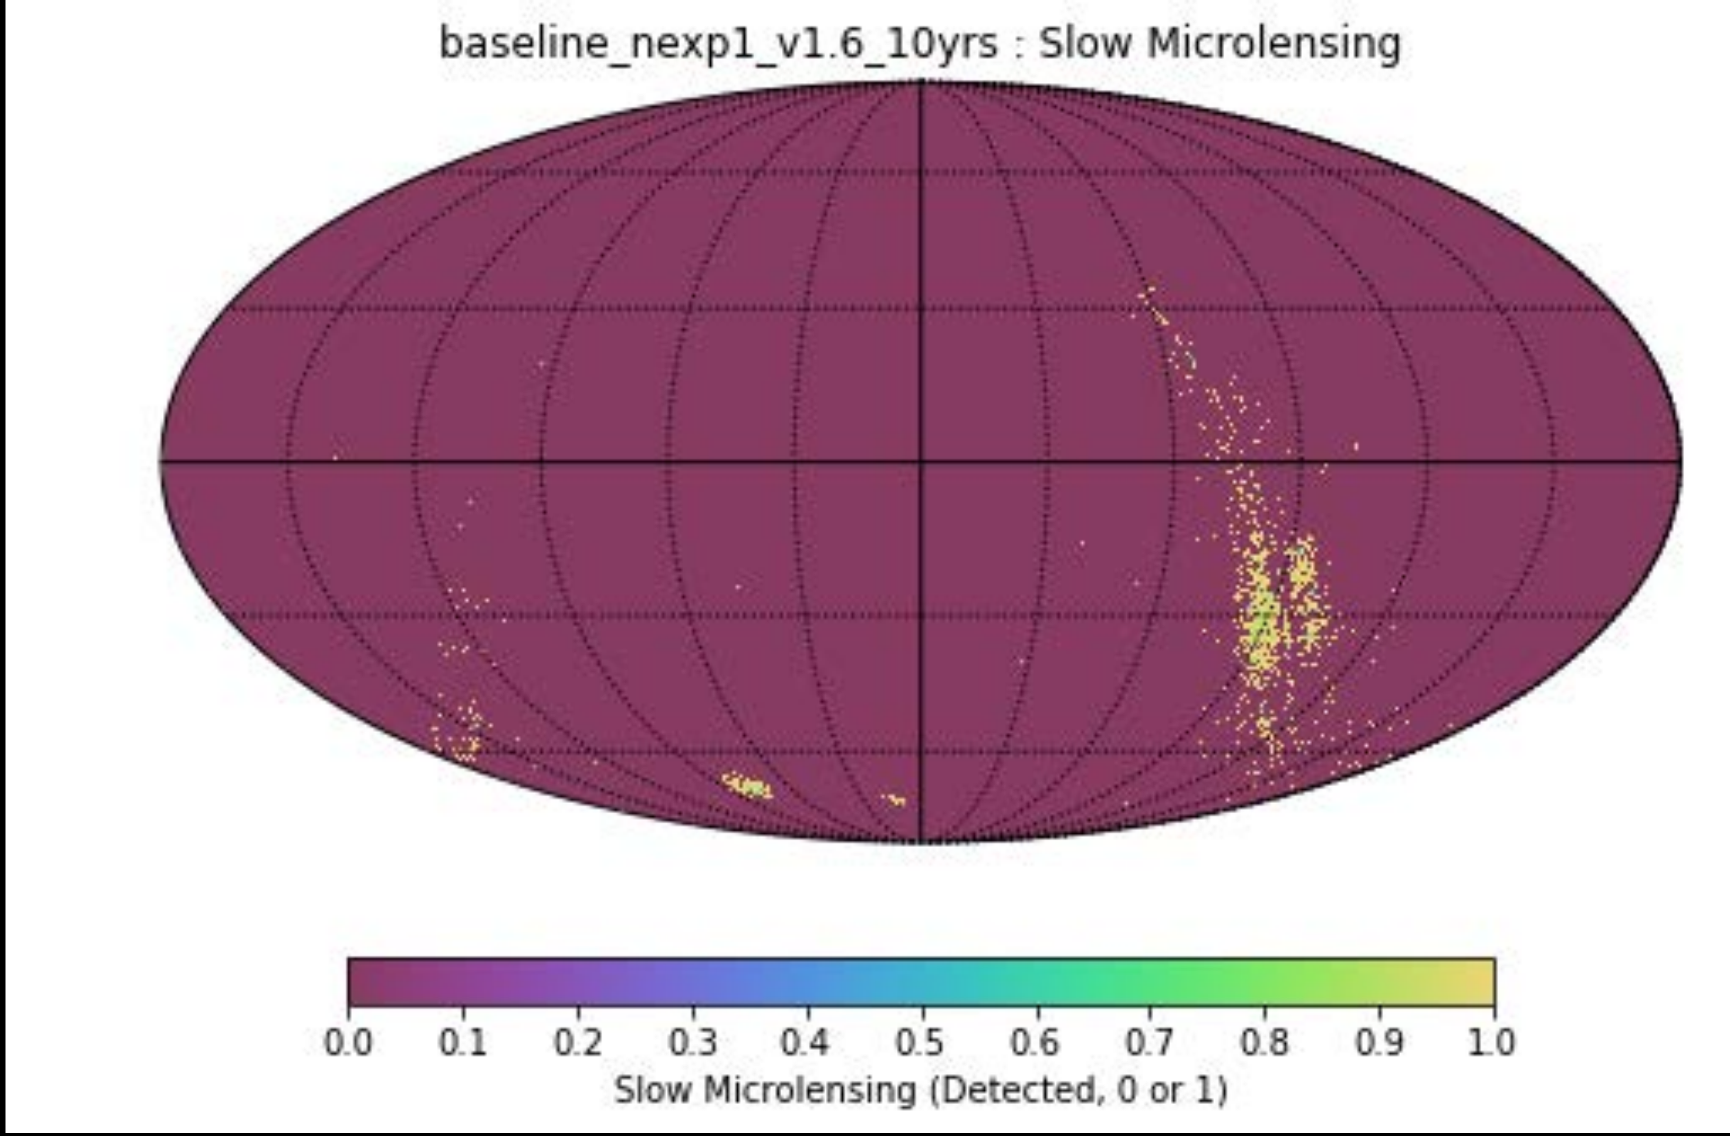

Solar system metrics

Lots of different populations:

NEO

MBA

Trojans

TNO

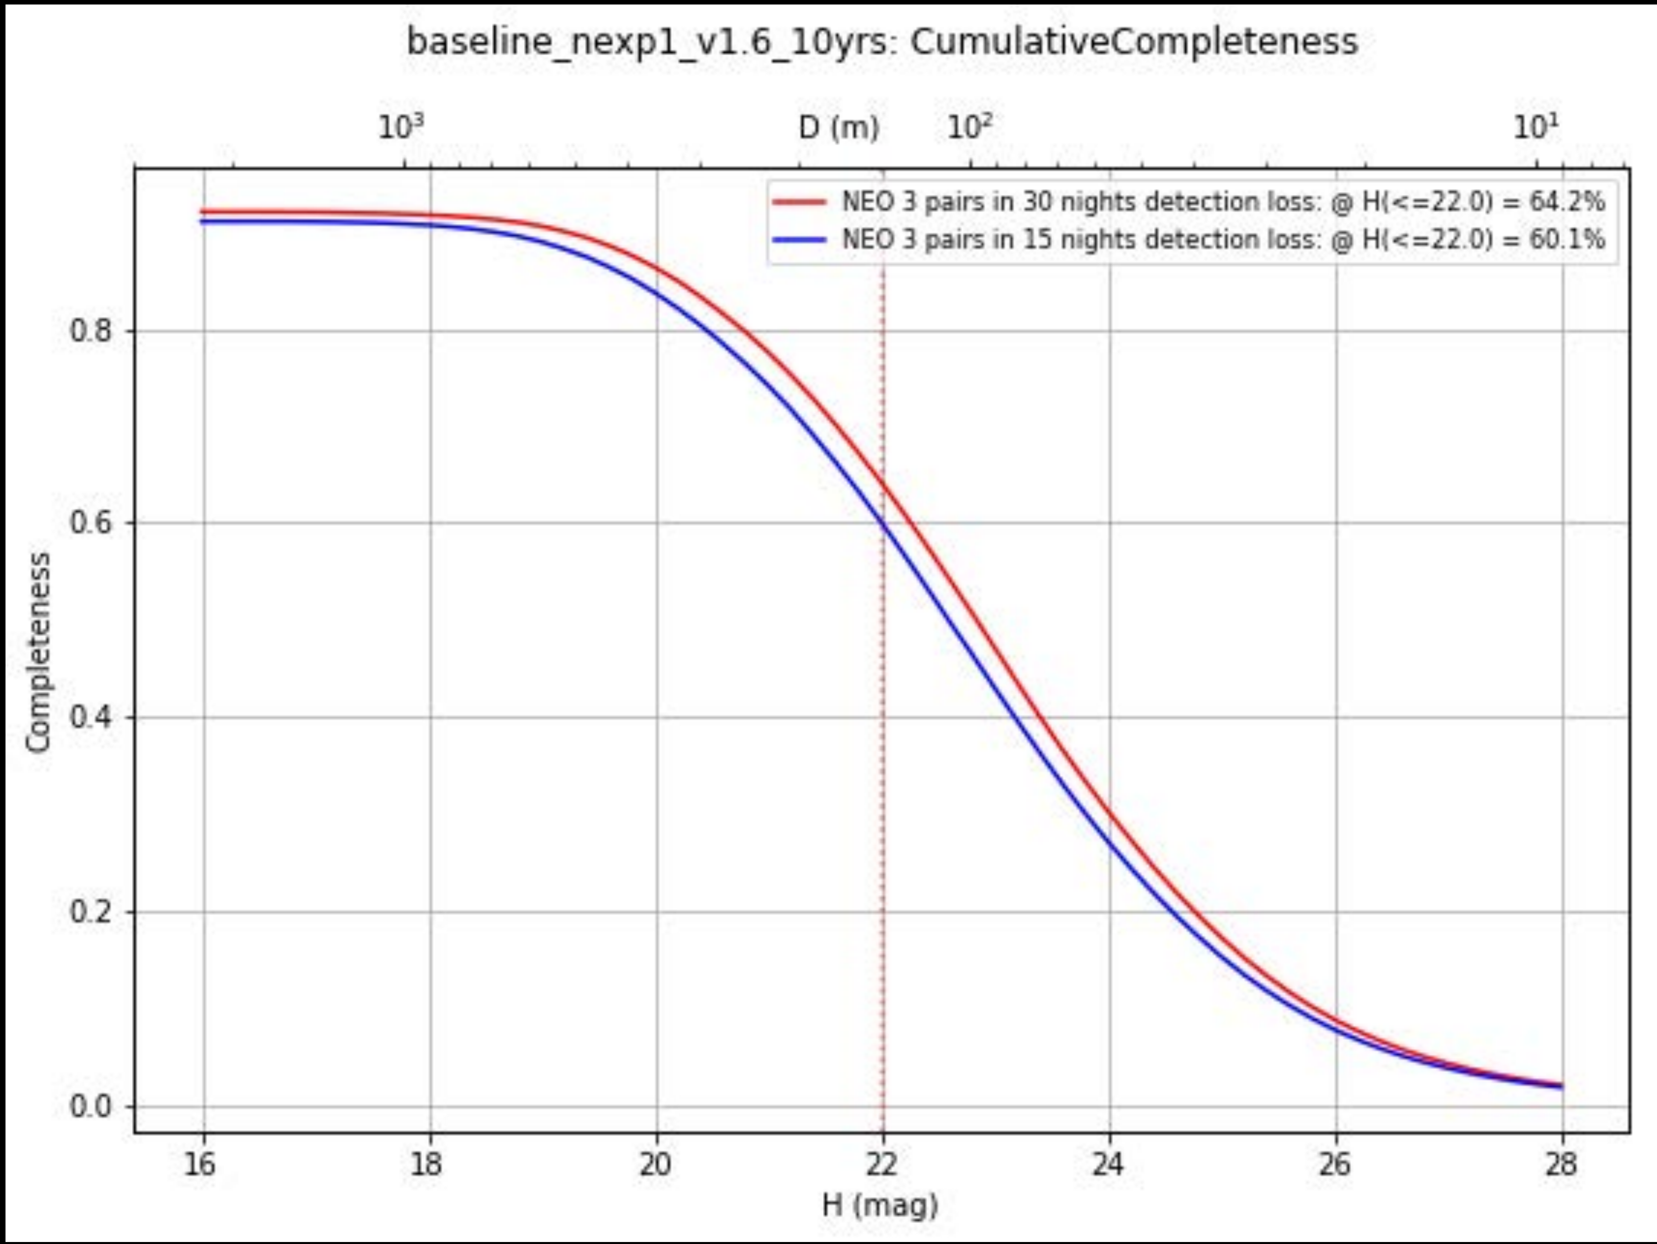

Always looking for better ways to measure survey science performance. We gather contributions from the community at

https://github.com/LSST-nonproject/sims_maf_contrib

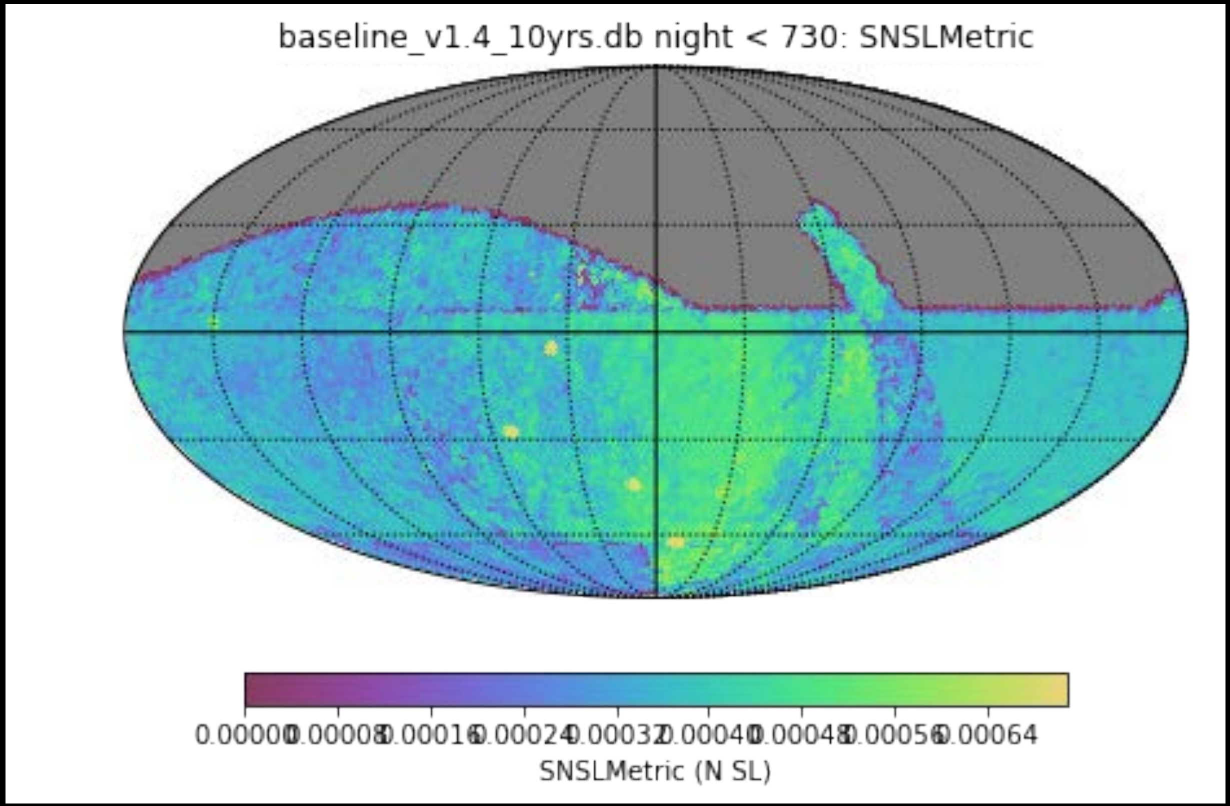

Strongly lensed SNe

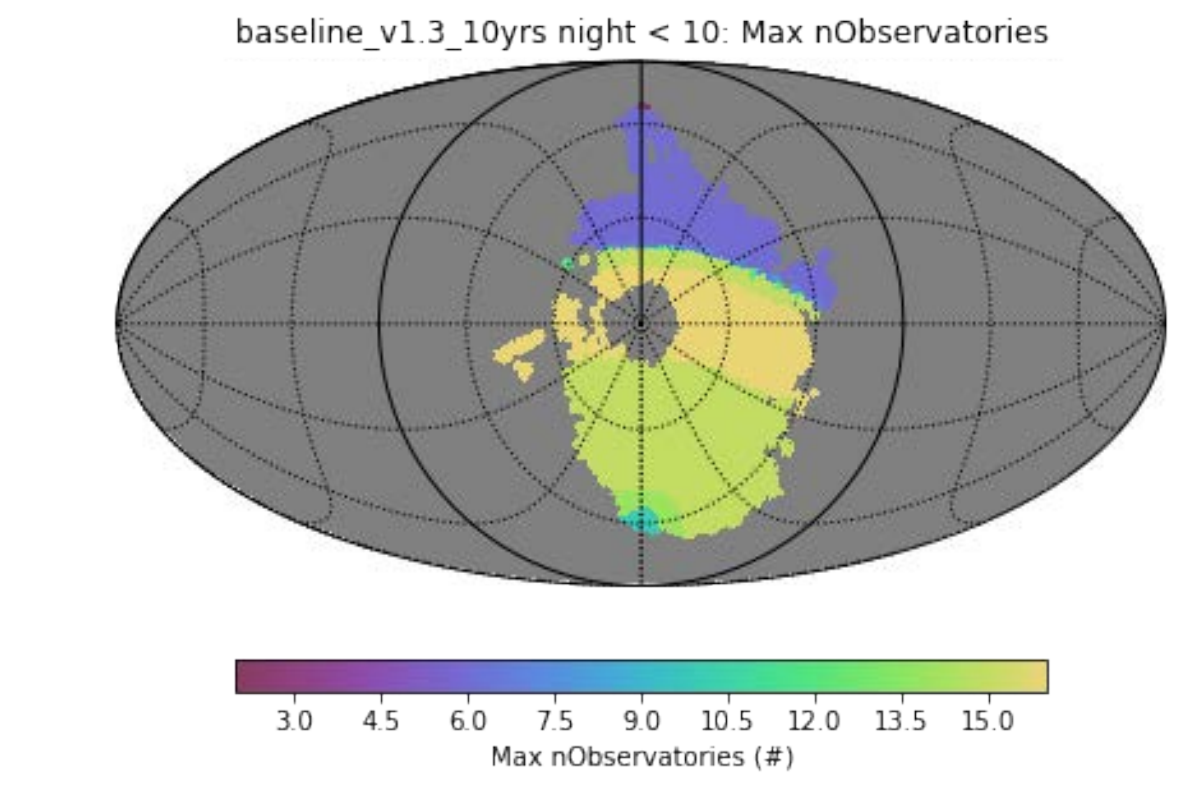

Potential follow up observatories

Awan et al, 2016 Testing LSST Dither Strategies for Survey Uniformity and Large-scale Structure Systematic

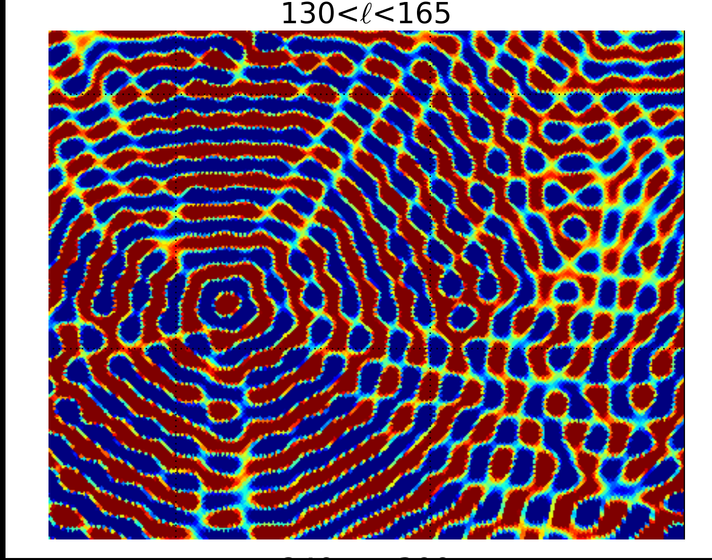

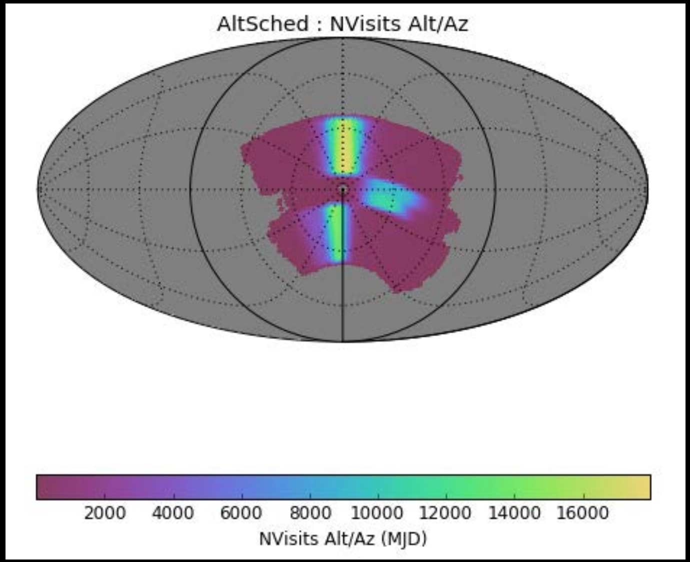

Rothchild et al, 2019 ALTSched: Improved Scheduling for Time-domain Science with LSST

Naghib et al, 2019 A Framework for Telescope Schedulers: With Applications to the Large Synoptic Survey Telescope

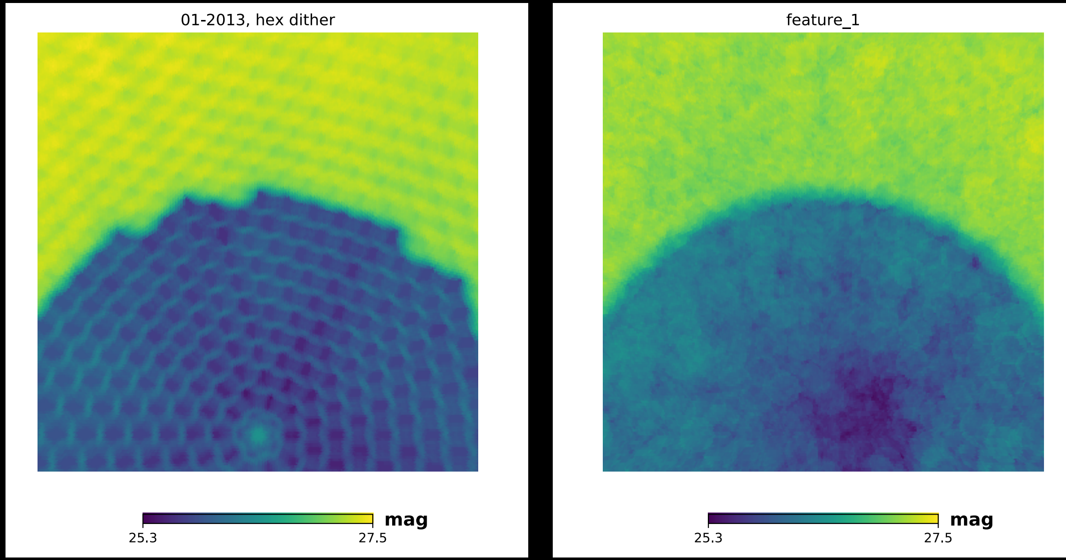

Conclusions

- The Rubin scheduler relies on a decision tree + Markov Decision Process to dynamically select potential observations
- We have a model observatory that includes the kinematics of the telescope and the environmental conditions
- We have simulated a large number of 10 year Rubin surveys
- MAF is our tool for looking at the science performance of different survey simulations

If you have questions, our virtual meeting session is:
Day: Wednesday, Aug 12 2020
Time: 07:30 HST - 10:30 PT - 13:30 EDT - 19:30 CEST - 03:30 AET +1
See https://project.lsst.org/meetings/rubin2020/ for zoom connection details

Or start a discussion on https://community.lsst.org
Feel free to tag @yoachim in your post so I see it