VERA C. RUBIN OBSERVATORY

U.S. National Science Foundation

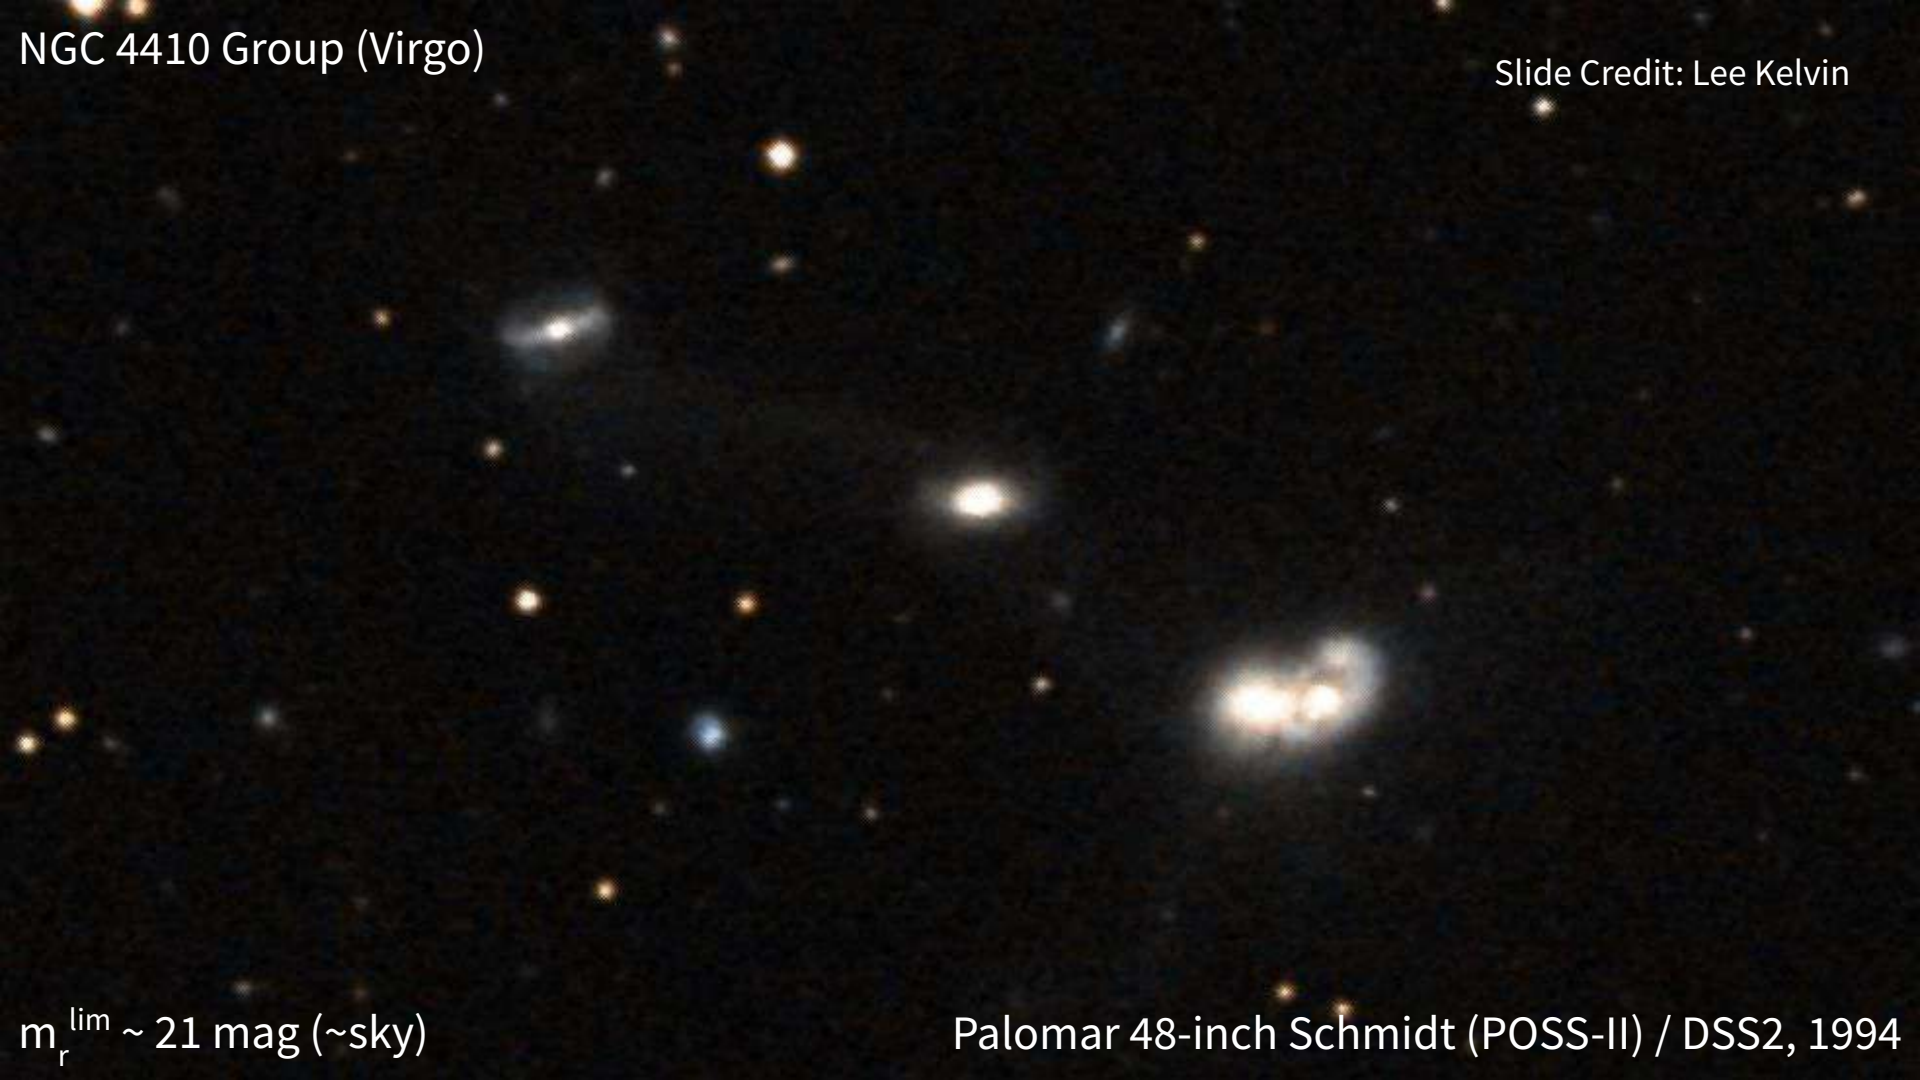
U.S. DEPARTMENT of ENERGY | Office of Science

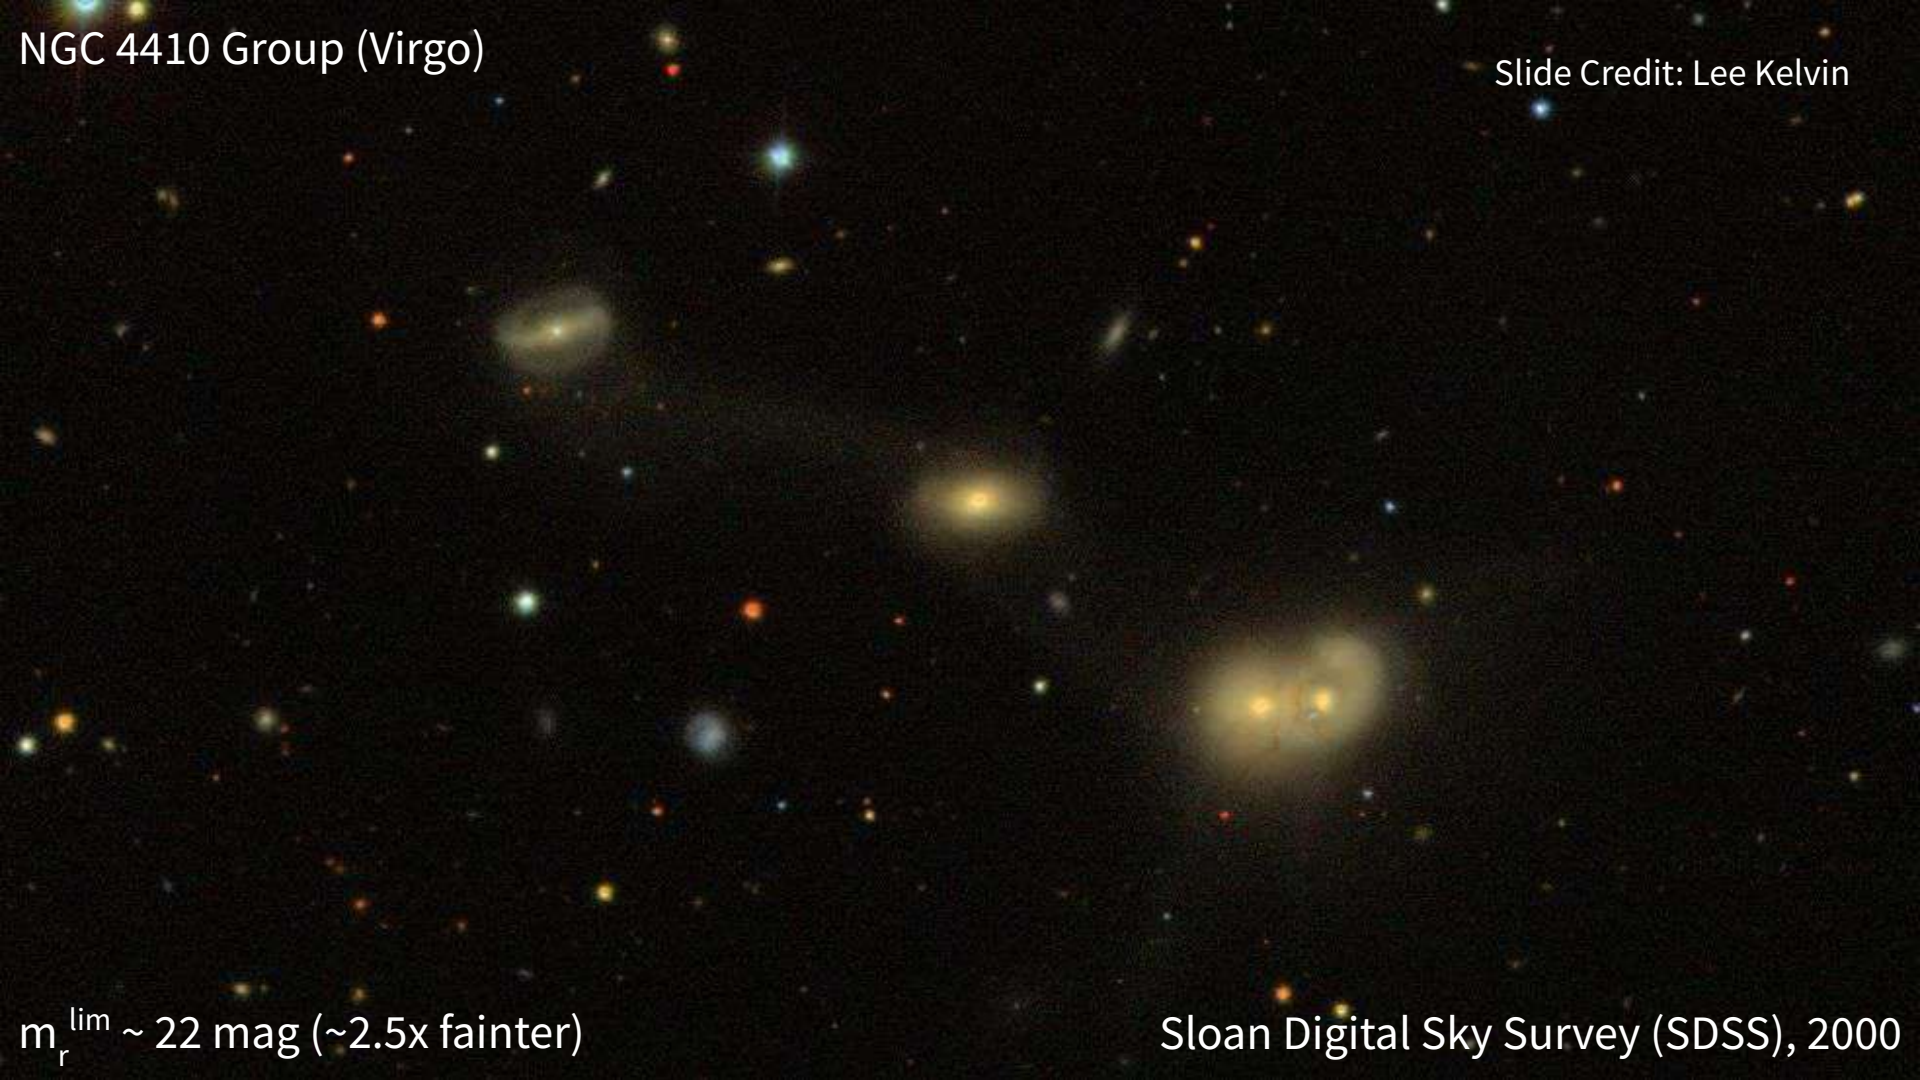
Data Management and Pipelines

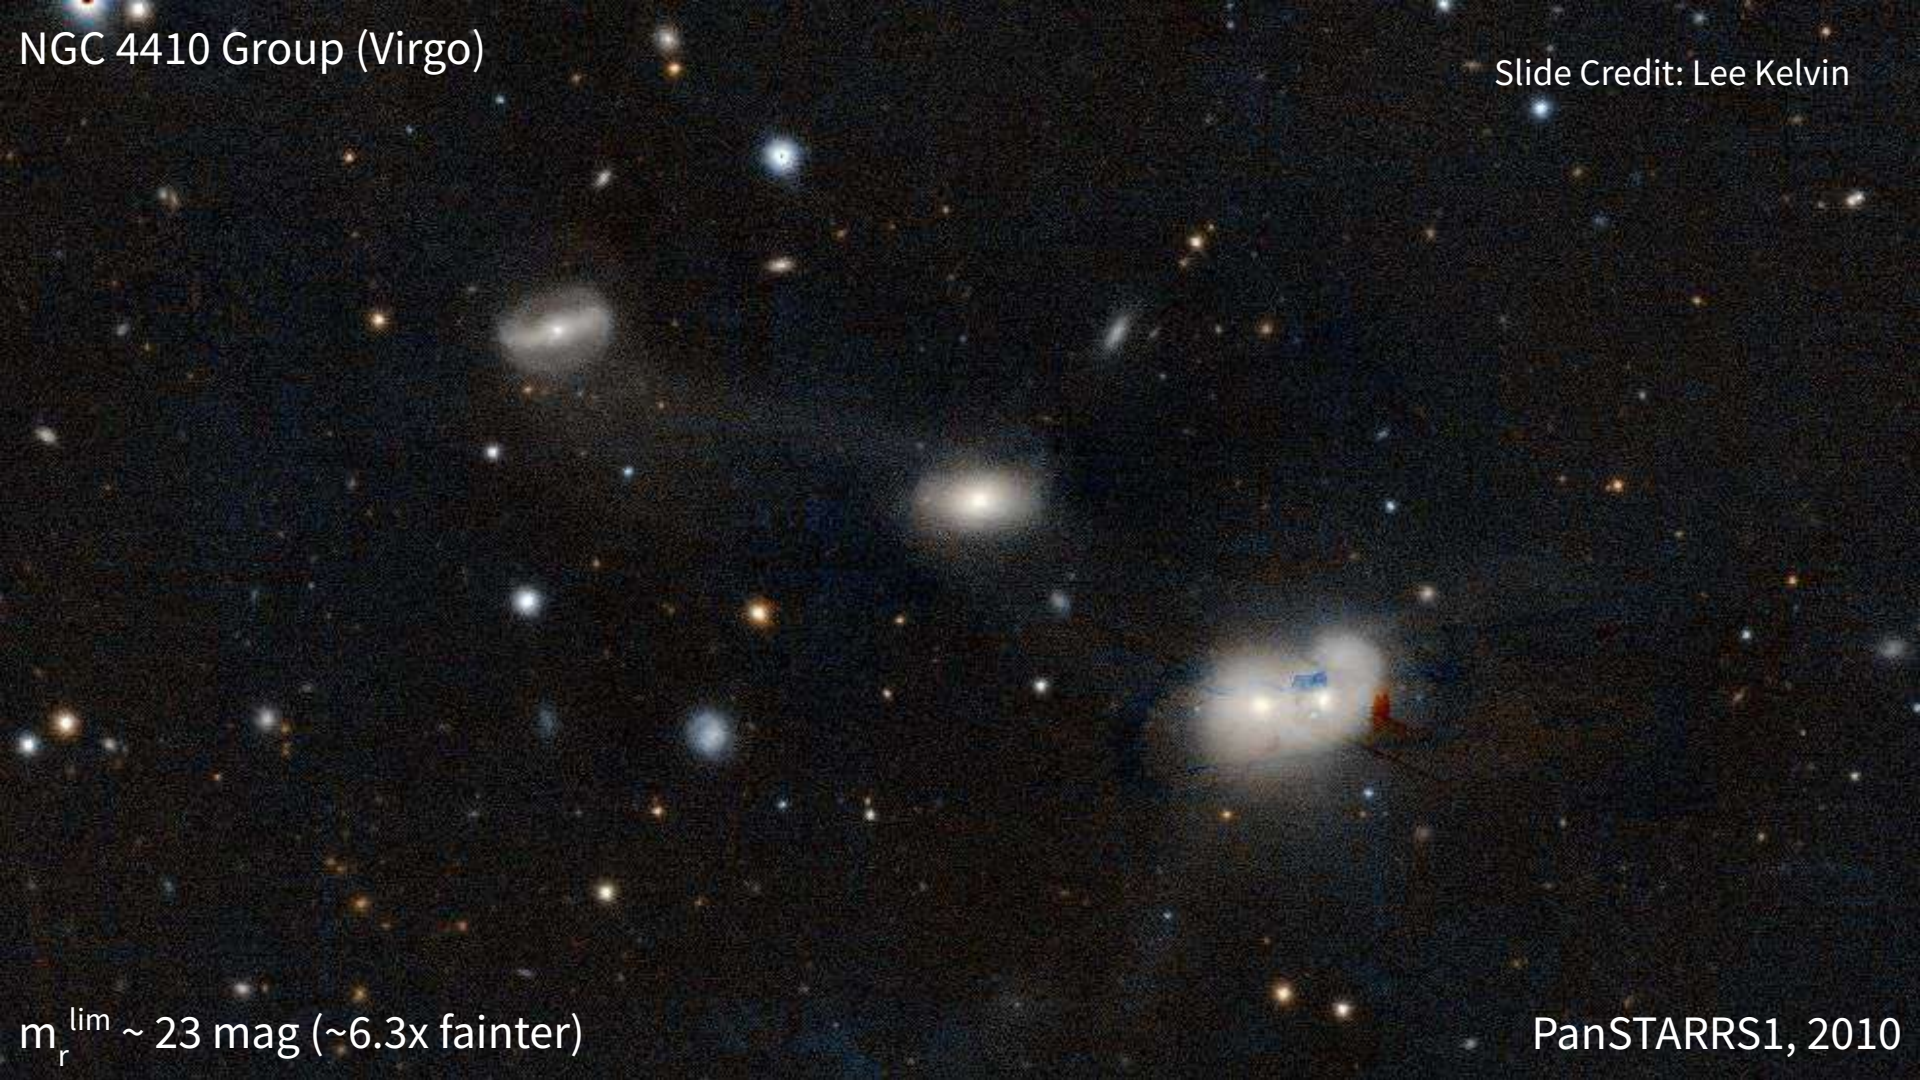
Yusra AlSayyad

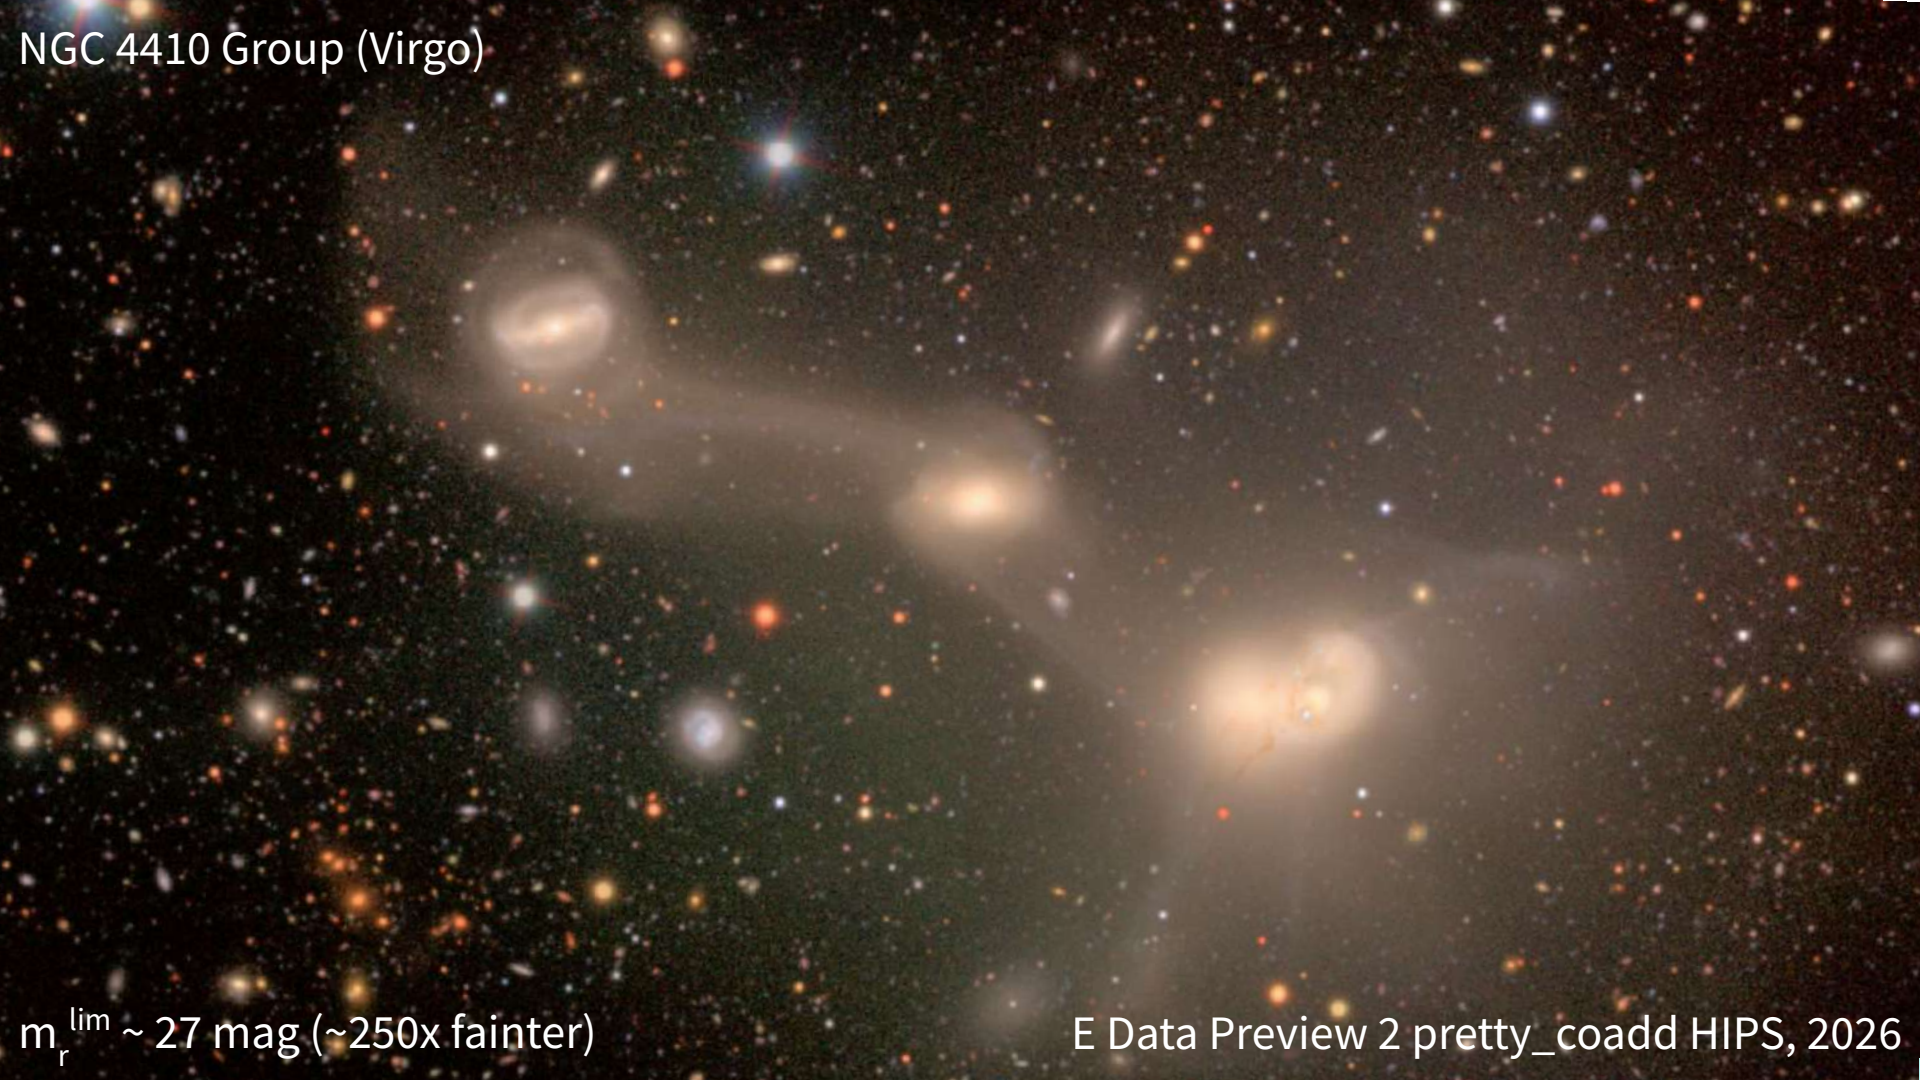
Deputy Manager for Rubin Data Management

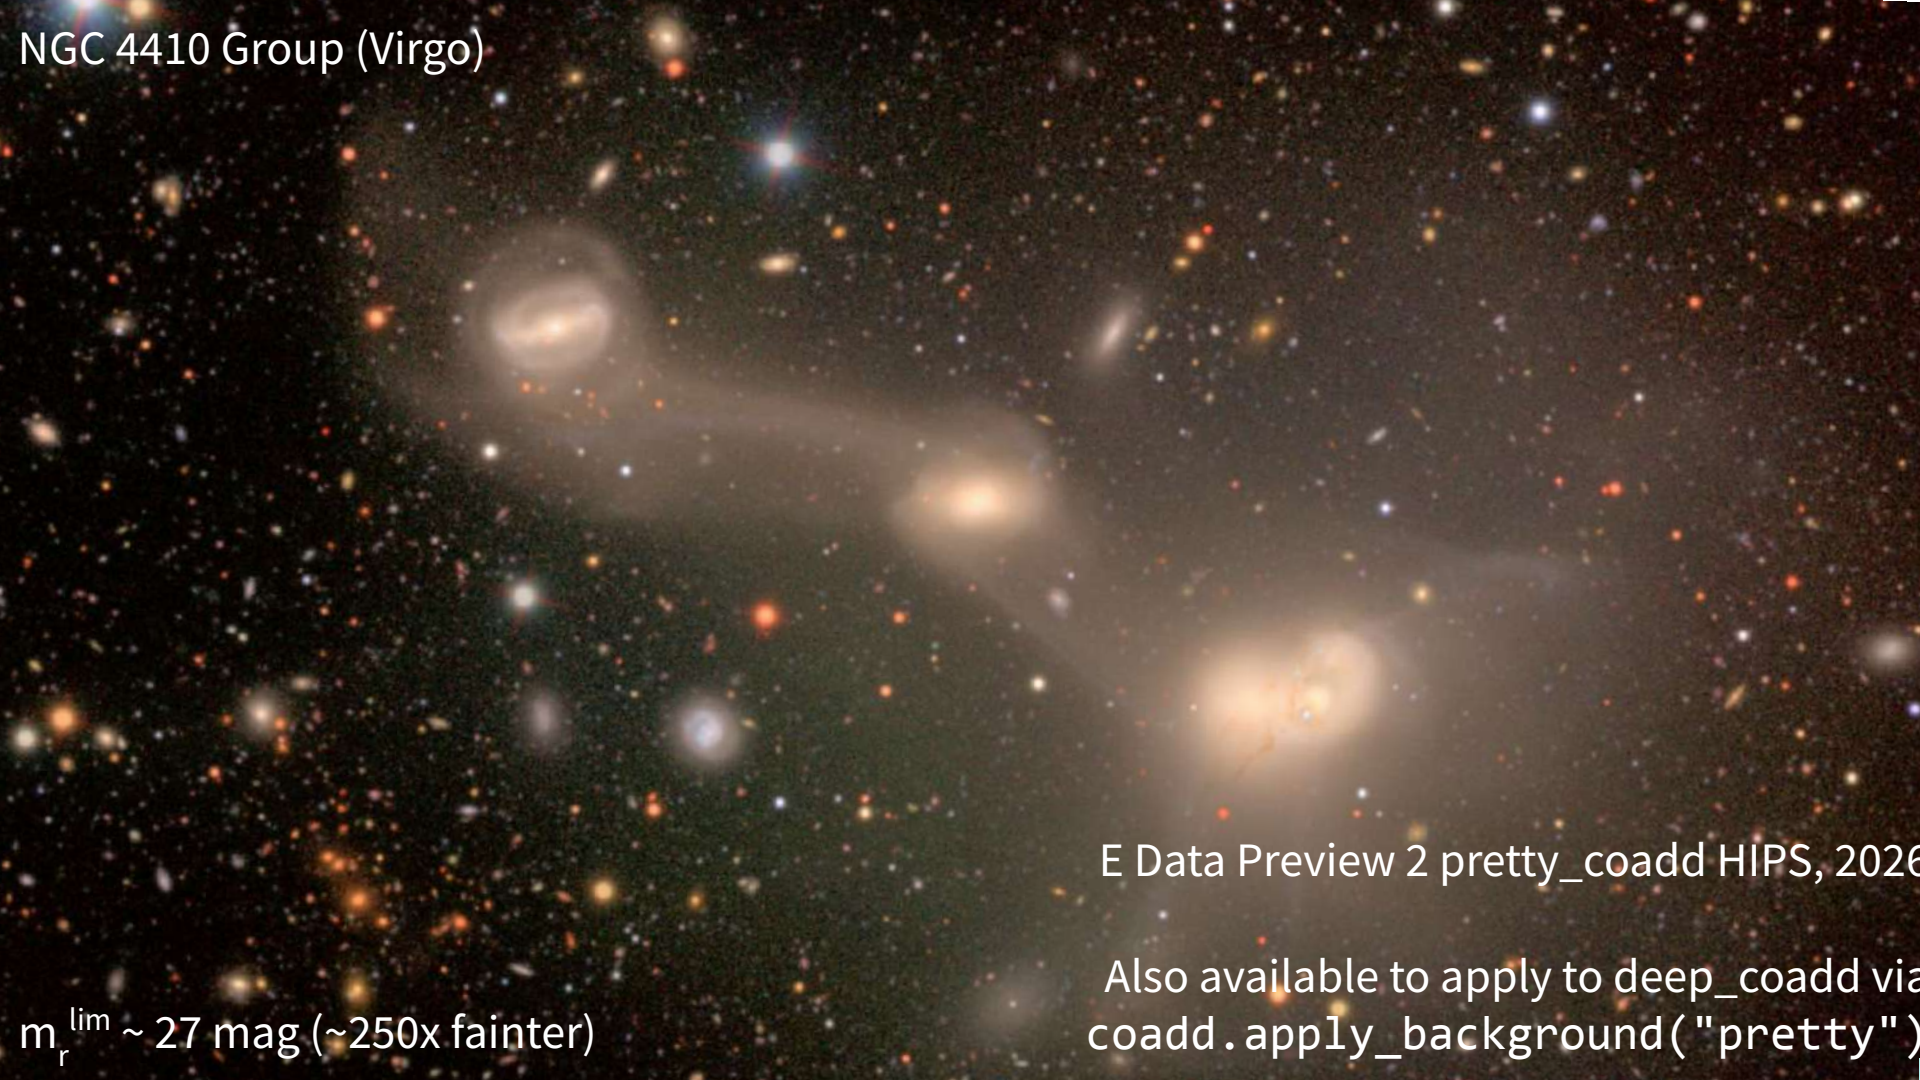



NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
m_r^{lim} ~ 21 mag (~sky)
Palomar 48-inch Schmidt (POSS-II) / DSS2, 1994

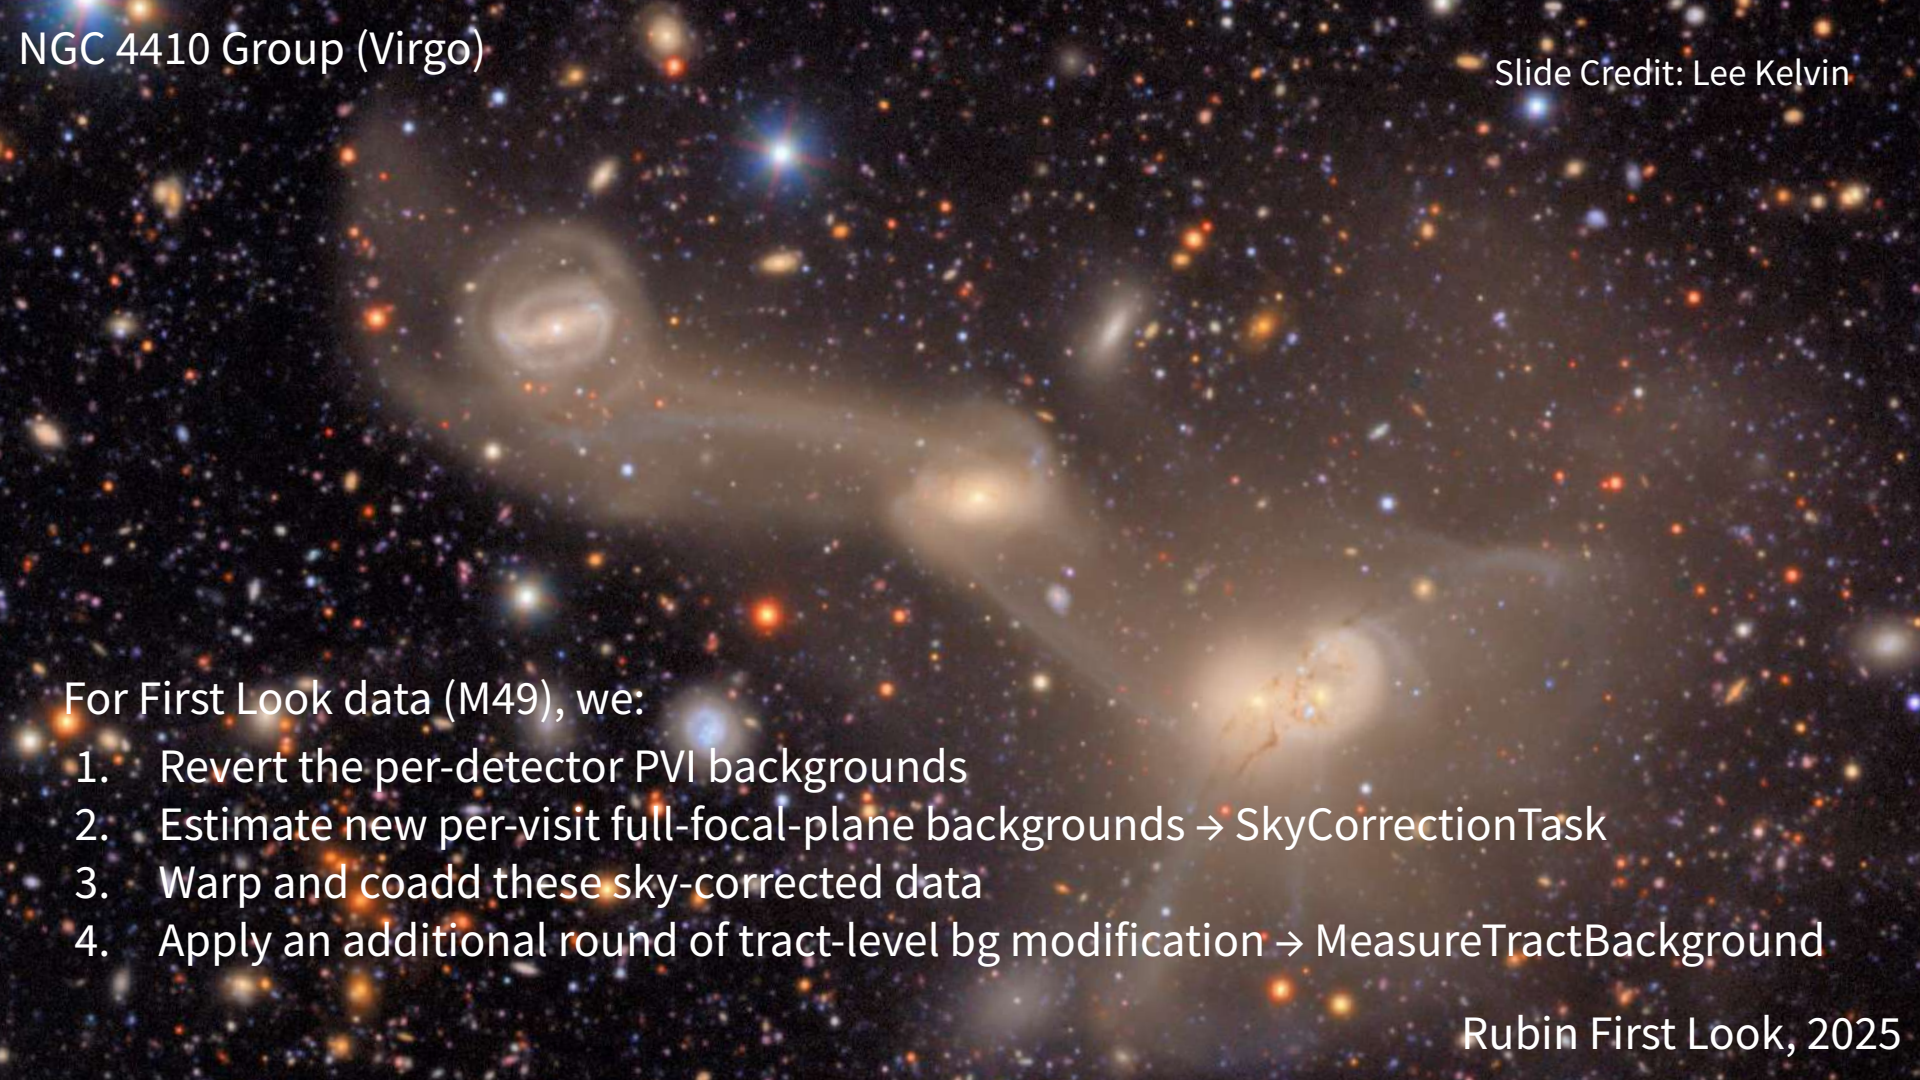

NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
m_r^{lim} ~ 22 mag (~2.5x fainter)
Sloan Digital Sky Survey (SDSS), 2000


NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
m_r^{lim} ~ 23 mag (~6.3x fainter)
PanSTARRS1, 2010

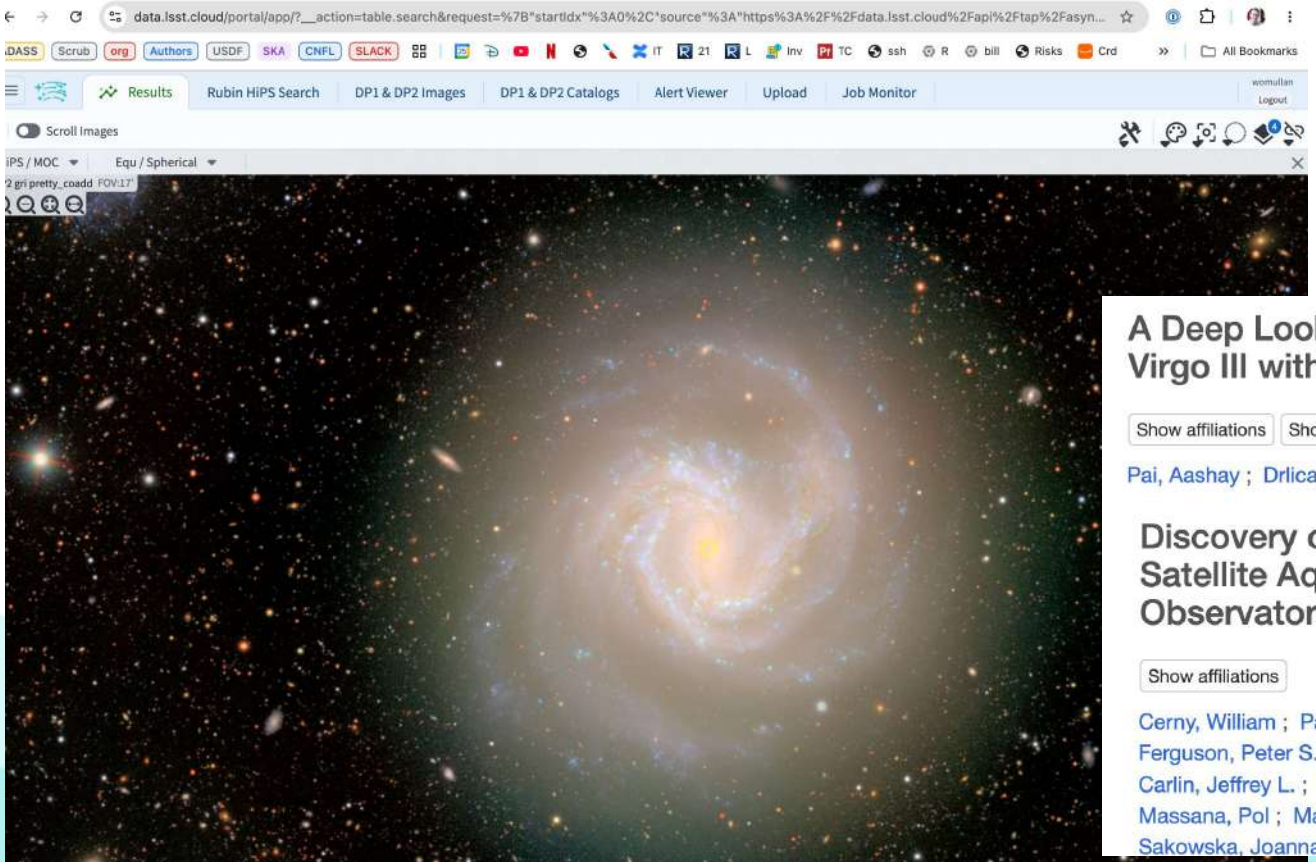

NGC 4410 Group (Virgo)
m_r^{lim} ~ 27 mag (~250x fainter)
E Data Preview 2 pretty_coadd HIPS, 2026


NGC 4410 Group (Virgo)
E Data Preview 2 pretty_coadd HIPS, 2026
Also available to apply to deep_coadd via
coadd.apply_background("pretty")
m_r^{lim} ~ 27 mag (~250x fainter)


NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
m_r^{lim} ~ 27 mag (~250x fainter)
Rubin First Look, 2025


NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
For First Look data (M49), we:
1. Revert the per-detector PVI backgrounds
2. Estimate new per-visit full-focal-plane backgrounds → SkyCorrectionTask
3. Warp and coadd these sky-corrected data
4. Apply an additional round of tract-level bg modification → MeasureTractBackground
Rubin First Look, 2025

Data Preview 2 (DP2)
Released Early DP2 (EDP2) **on time for RCW** July 27



EDP2 = catalogs and coadds
Full DP2 produced and release planned for December

First publication during Rubin Community Workshop
Second soon after

A Deep Look at the Ultra-Faint Milky Way Satellite Virgo III with Rubin Observatory Data Preview 2

Show affiliations | Show 9 more authors

Pai, Aashay ; Drlica-Wagner, Alex ; Ferguson, Peter S. ; ...

Discovery of the Distant, Ultra-Faint Milky Way Satellite Aquarius IV with the Vera C. Rubin Observatory Early Data Preview 2

Show affiliations

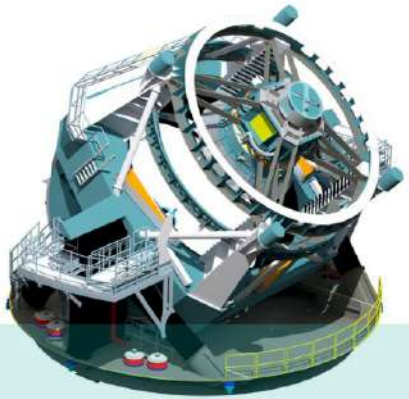
Cerny, William ; Pai, Aashay ; Drlica-Wagner, Alex ; Pace, Andrew B. ; Ferguson, Peter S. ; Geha, Marla ; Tan, Chin Yi ; Campana, Sasha ; Carlin, Jeffrey L. ; Crnojević, Denija ; Ji, Alexander P. ; Limberg, Guilherme ; Massana, Pol ; Mau, Sidney ; Medina, Gustavo E. ; Mutlu-Pakdil, Burçin ; Sakowska, Joanna D. ; Shipp, Nora ; Stringfellow, Guy S.

Both Prompt and Data Release Data Products are already available

Raw Data: 20TB/night



Sequential 30s images covering the entire visible sky every few days



LSST SCIENCE PIPELINES

Prompt Data Products

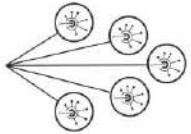
- Alerts incl. science, template and difference image cutouts — 60s — via Alert Streams
- Catalogs of detections incl. difference images, transient, variable & solar system sources — 24h — via Prompt Products
- Raw & processed visit images (PVIs), difference images — 80h — via Image Services

Data Release Data Products

Final 10yr Data Release:

- Images: 5.5 million x 3.2 Gpixels
- Catalog: 15PB, 37 billion objects

via Data Releases



Community Brokers

Rubin Data Access Centres (DACs)

USA (USDF)
Chile (CLDF)
France (FRDF)
United Kingdom (UKDF)

Independent Data Access Centers (IDACs)

Access to proprietary data and the Science Platform require Rubin data rights

Rubin Science Platform

Provides access to LSST Data Products and services for all science users and project staff.

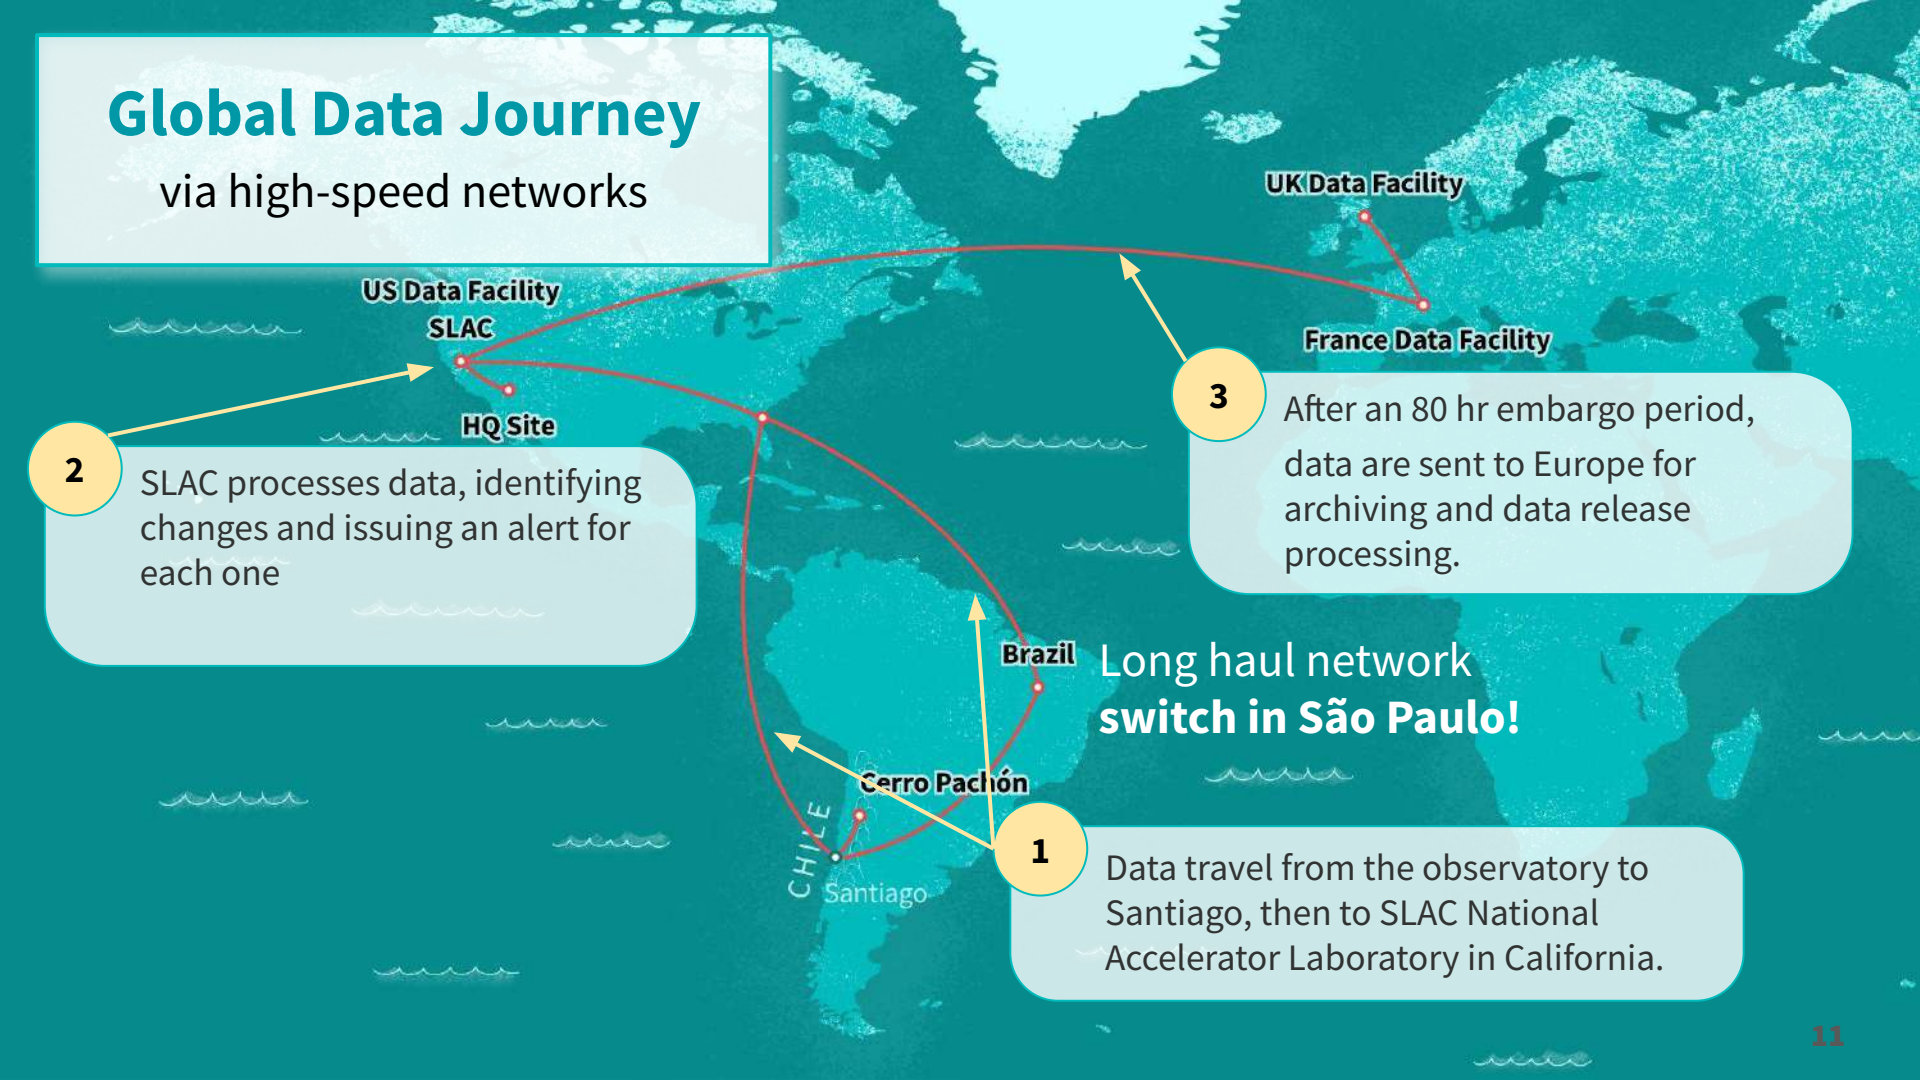




Credit: Leanne Guy

Global Data Journey

via high-speed networks



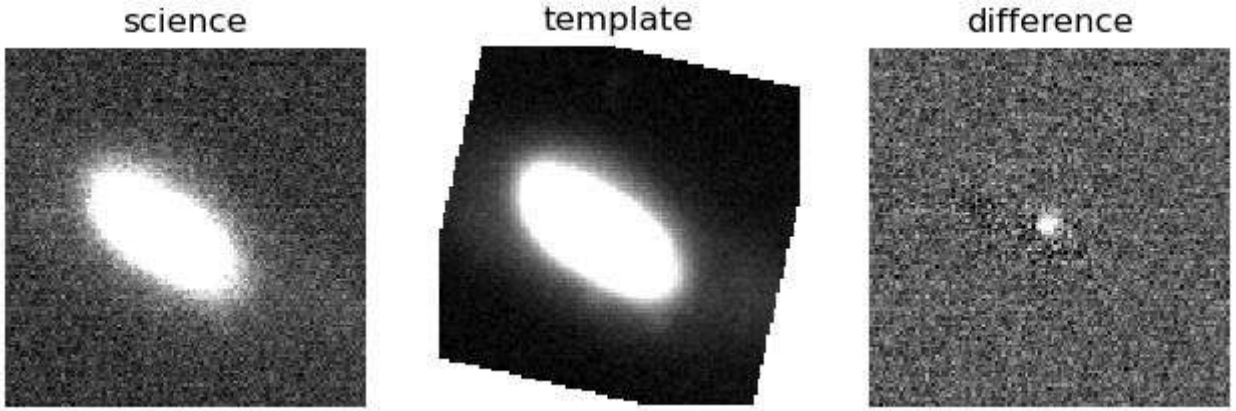
2 SLAC processes data, identifying changes and issuing an alert for each one

3 After an 80 hr embargo period, data are sent to Europe for archiving and data release processing.

Long haul network
switch in São Paulo!

1 Data travel from the observatory to Santiago, then to SLAC National Accelerator Laboratory in California.

How “changes” are identified in real time



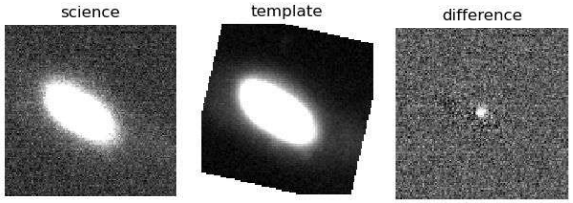
Above: a “detection triplet” with Rubin commissioning data for an unclassified time-domain event (an extragalactic transient) which was first discovered by the Zwicky Transient Facility.

Image Subtraction

Immediately after a new image is obtained and processed, an older (template) image is subtracted from it, to create a difference image.

Detection algorithm run on difference image to detect sources – *positive or negative* – with a signal-to-noise ratio >5.

What is an “alert”, exactly?



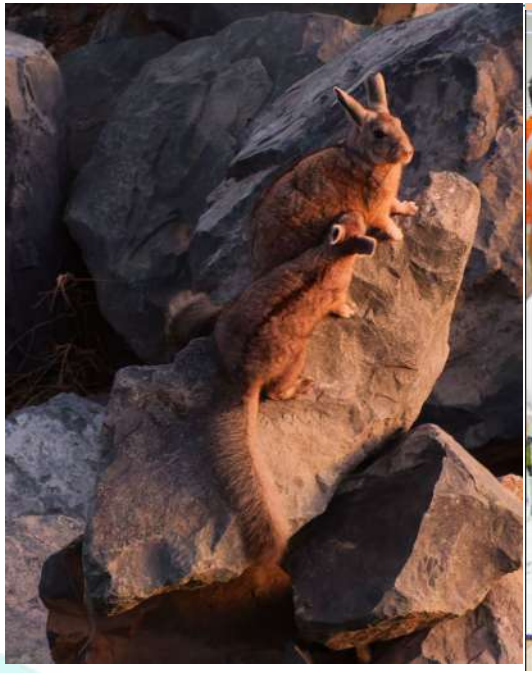
One “alert” is issued for each difference image source detected with signal-to-noise ratio >5.

An “alert” is a machine-readable file that contains data for the detected source:

- the sky position (coordinates)
- measurements (magnitude)
- observational metadata (date, time, filter)
- 12-month history of past detections
- very small image cutouts


```
{diaSource:
 {"diaSourceId":611255759837069401,
  "visit": 202411080532,
  "detector": 42,
  "MJD": 60623.261598,
  "band": "r",
  "radec": {"ra": 53.1249,
            "dec": -27.7399},
  "raVar": 5.1,
  "decVar": 10.1,
  "ra_dec_Cov": 1.241,
  "x": 112.1,
  "y": 121.1,
  "xVar": 1.2,
  "yVar": 12.1,
  "x_y_Cov": 11.2,
  "snr": 41.1,
  "psfFlux": 1241.0},
 history: …
}
```

Rubin summit 24 Feb 2026 first public alerts

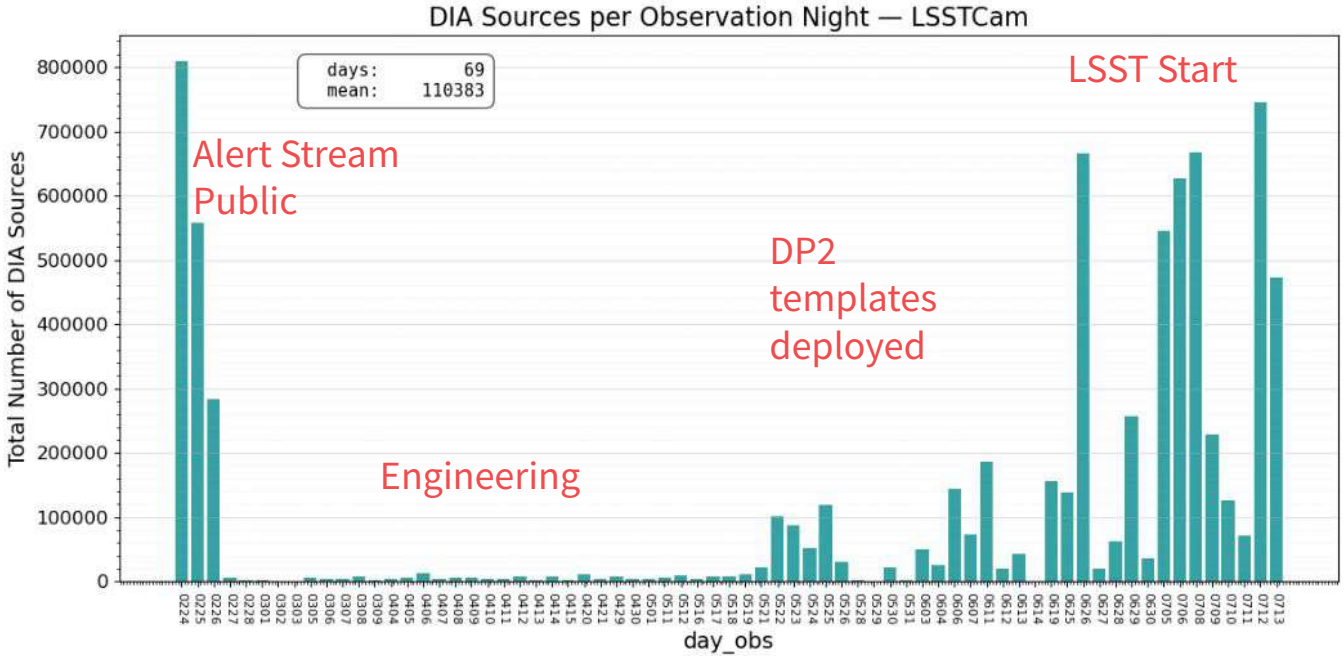


Vizcacha were waiting at sunset



A little after 21:35 Chile, time first alerts from the night appeared on the brokers' websites. On zoom with the brokers.

Alerts started flowing 24 February 2026

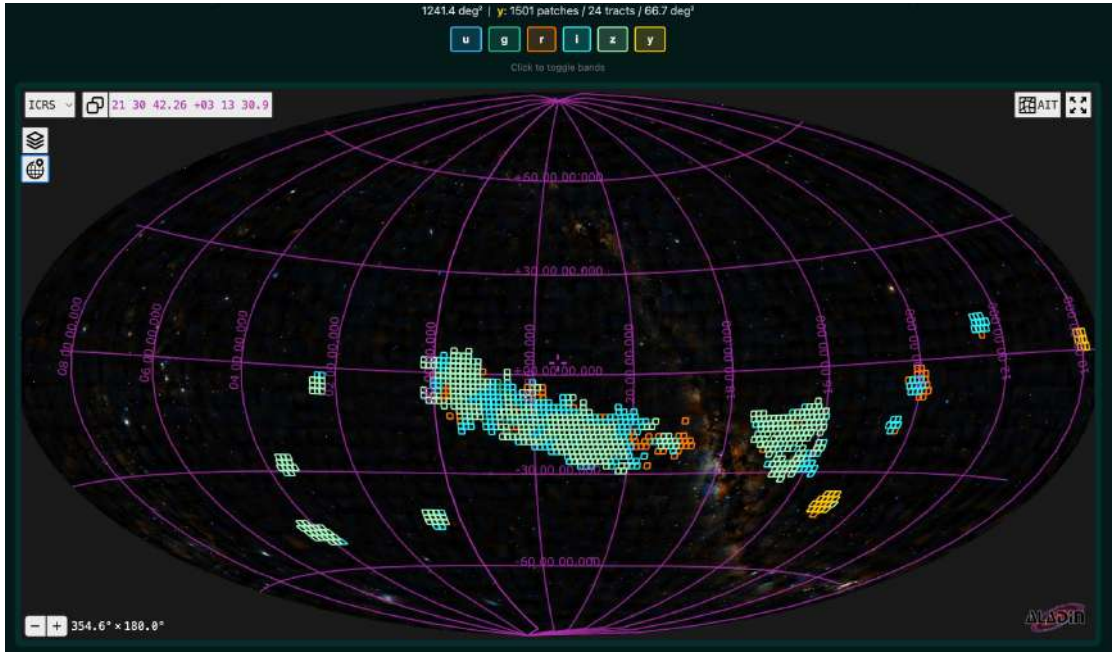


ls.st/alert-stream-status

Alert rate varies night to night depending on what is observed.

Learn more about alerts and brokers:
rubinobservatory.org/for-scientists/data-products/alerts-and-brokers

DP2 template images were deployed May 21-Jun 11
Just as Earth was coming back around

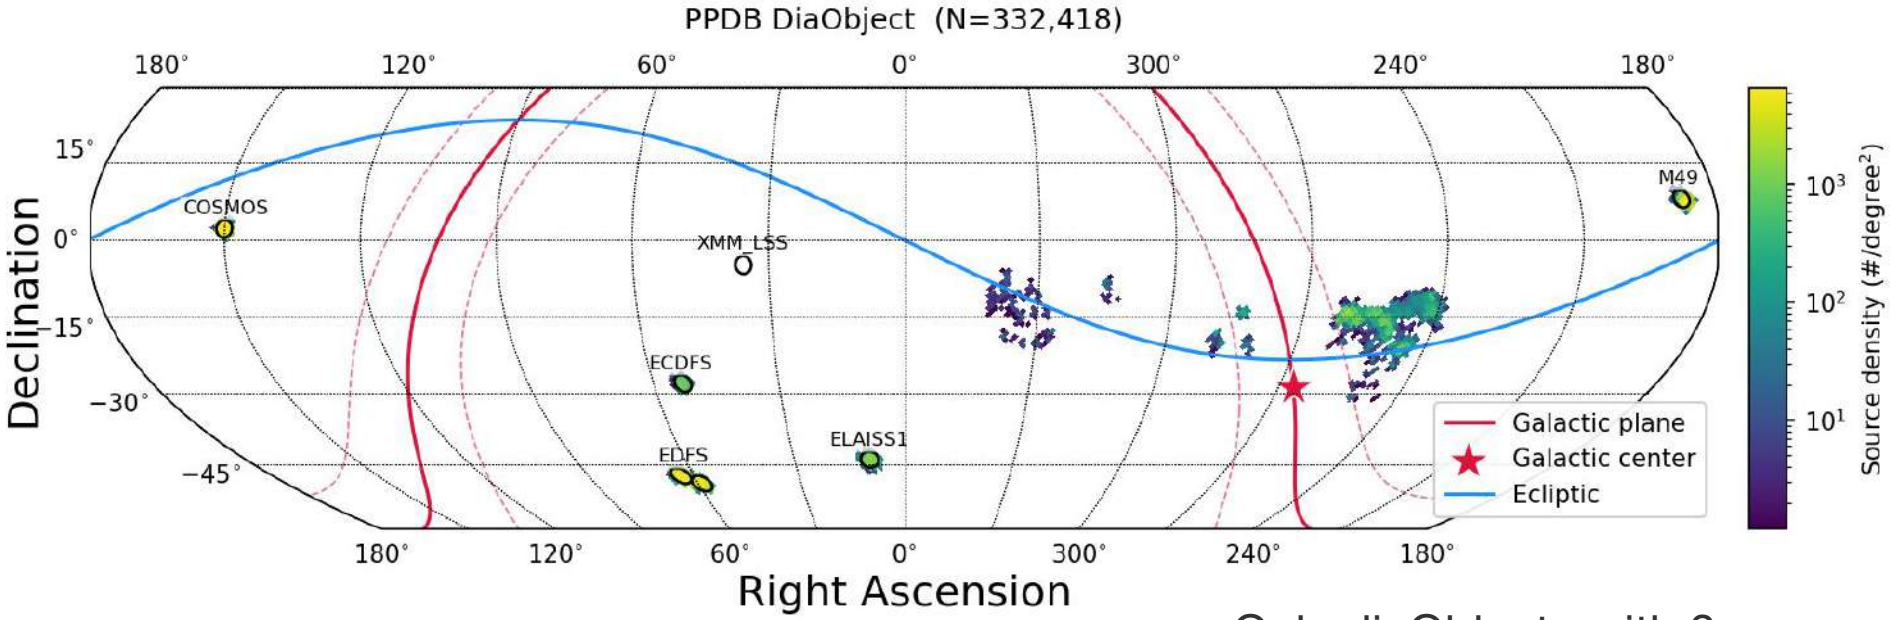


ls.st/template-coverage

Current coverage (2026-07-18)

- u: 94 deg^2
- g: 302 deg^2
- r: 1033 deg^2
- i: 1530 deg^2
- z: 1241 deg^2
- y: 67 deg^2

DiaObjects from alerts so far



Only diaObjects with 3+ diaSources

Figure credit: Ian Sullivan

DiaSources from alerts so far

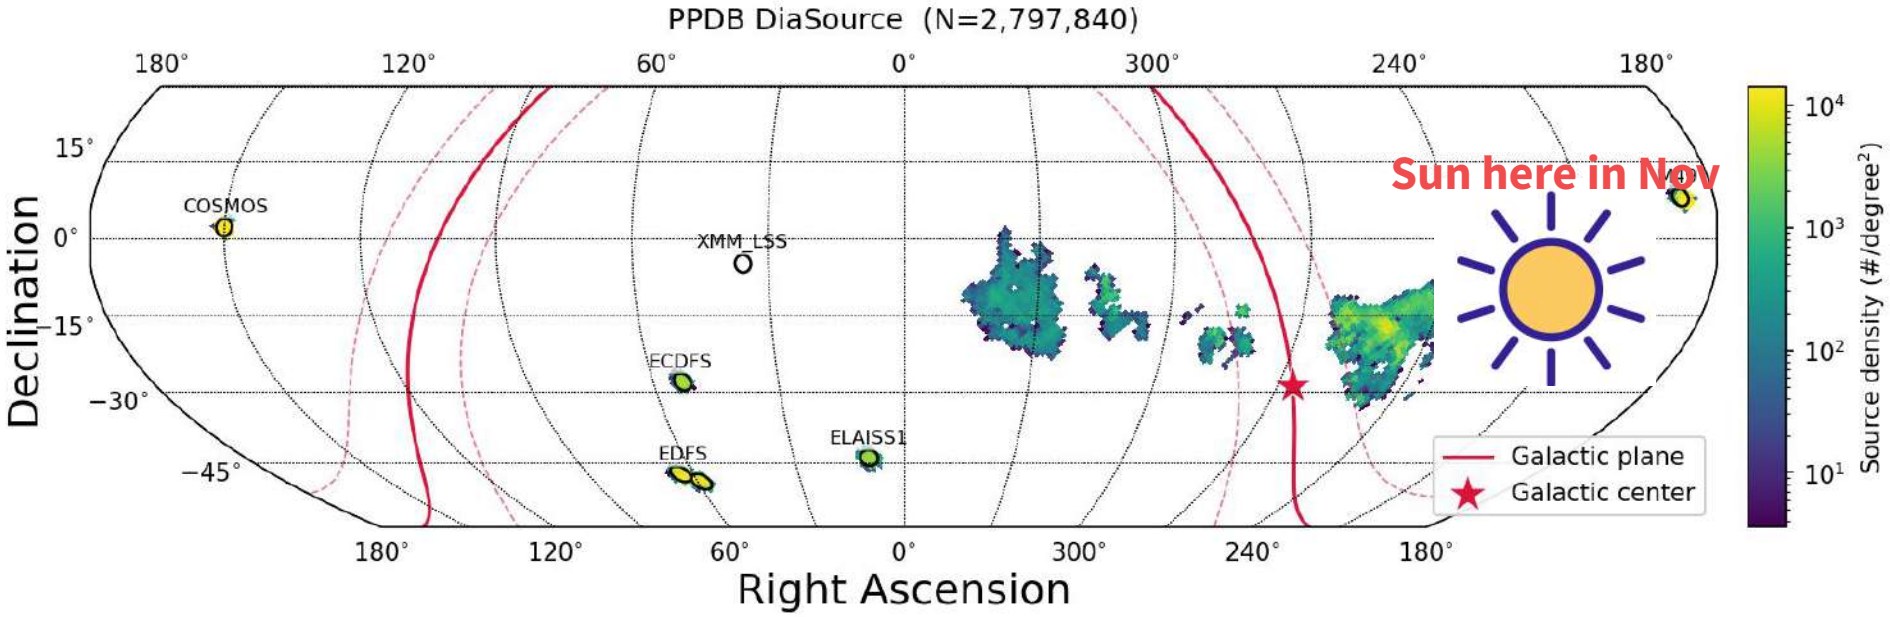
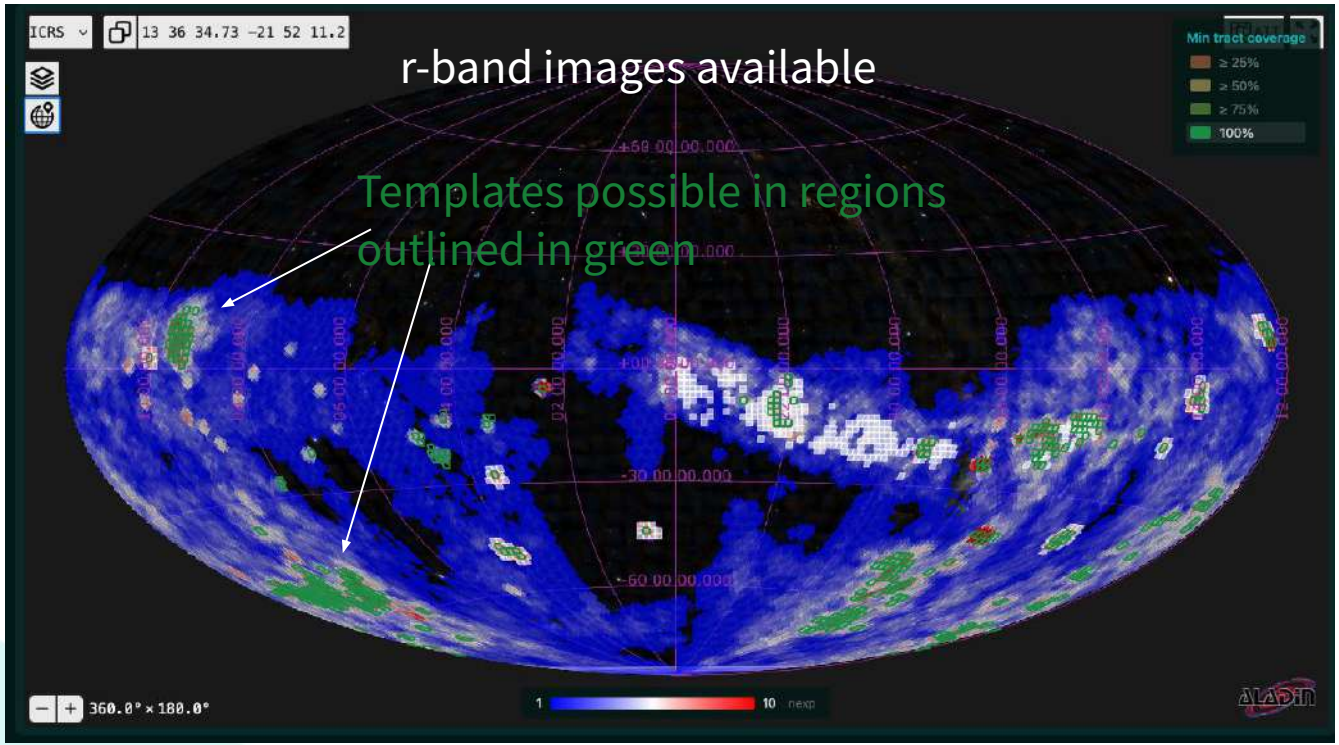
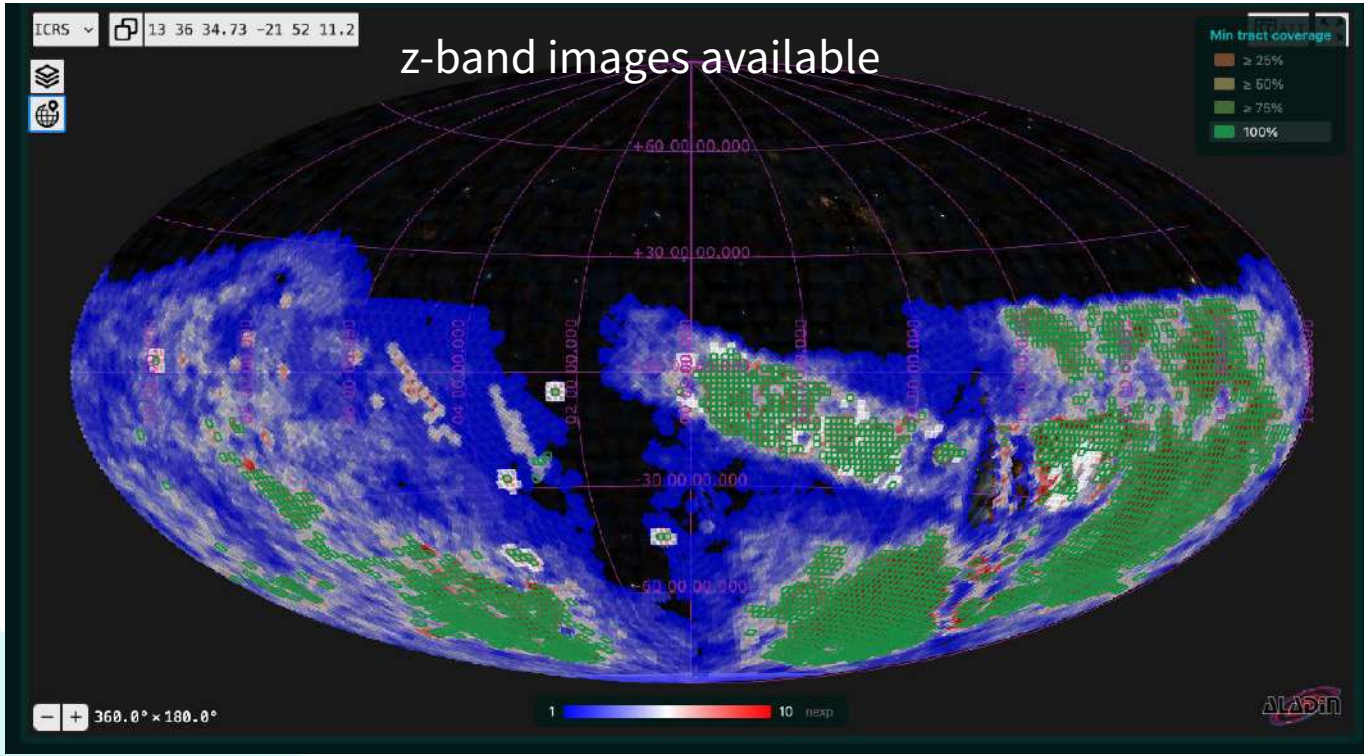


Figure credit: Ian Sullivan

Only diaSources with 0.75+ reliability

Data productions **next** processing will produce incremental templates





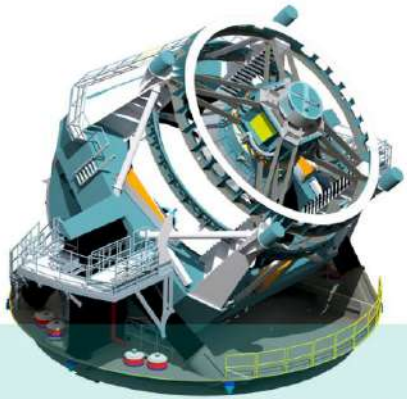
Driven by the gravitational wave observing run starting in January, collaborating with the Dark Energy Science Collaboration to produce some targeted DECam templates

Both Prompt and Data Release Data Products are already available

Raw Data: 20TB/night



Sequential 30s images covering the entire visible sky every few days



LSST SCIENCE PIPELINES

Prompt Data Products

- Alerts incl. science, template and difference image cutouts — 60s — via Alert Streams
- Catalogs of detections incl. difference images, transient, variable & solar system sources — 24h — via Prompt Products
- Raw & processed visit images (PVIs), difference images — 80h — via Image Services

Data Release Data Products

Final 10yr Data Release:

- Images: 5.5 million x 3.2 Gpixels
- Catalog: 15PB, 37 billion objects

via Data Releases

Public/online since February 25 2026

Brokers

Expected **Nov 30 2026**

Expected **Jan-Mar 2027**

France (FRDF)
United Kingdom (UKDF)

Early DP2 online Jul 2026
Full DP2 expected Dec 2026

Access to proprietary data and the Science Platform require Rubin data rights

Rubin Science Platform

Provides access to LSST Data Products and services for all science users and project staff.



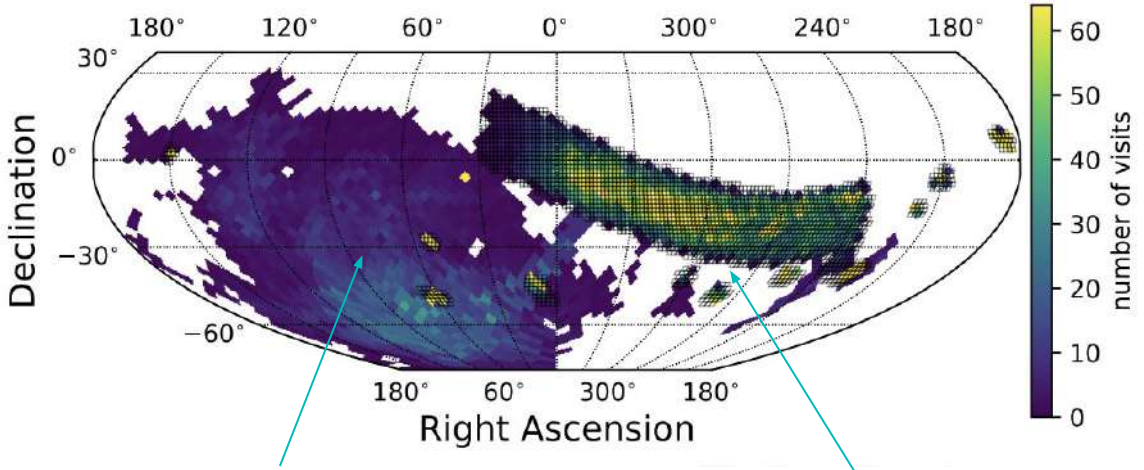


Credit: Leanne Guy

EDP2 (data taken 2025-04-25 –> 2026-01-6)

DP2 Production in numbers:

- 30,000 observations (“visits”) x 180 detectors
- = 5.4 million images (400x DP1)
- 3 million CPU-hours
- 786,420,720 Objects
- 3000 sq deg coadded
- 15000 sq deg w/ at least 1-band coverage. That is ⅓ of the entire (not visible) sky
- 400TB of “final” data products
- 7PB of “intermediate” data products



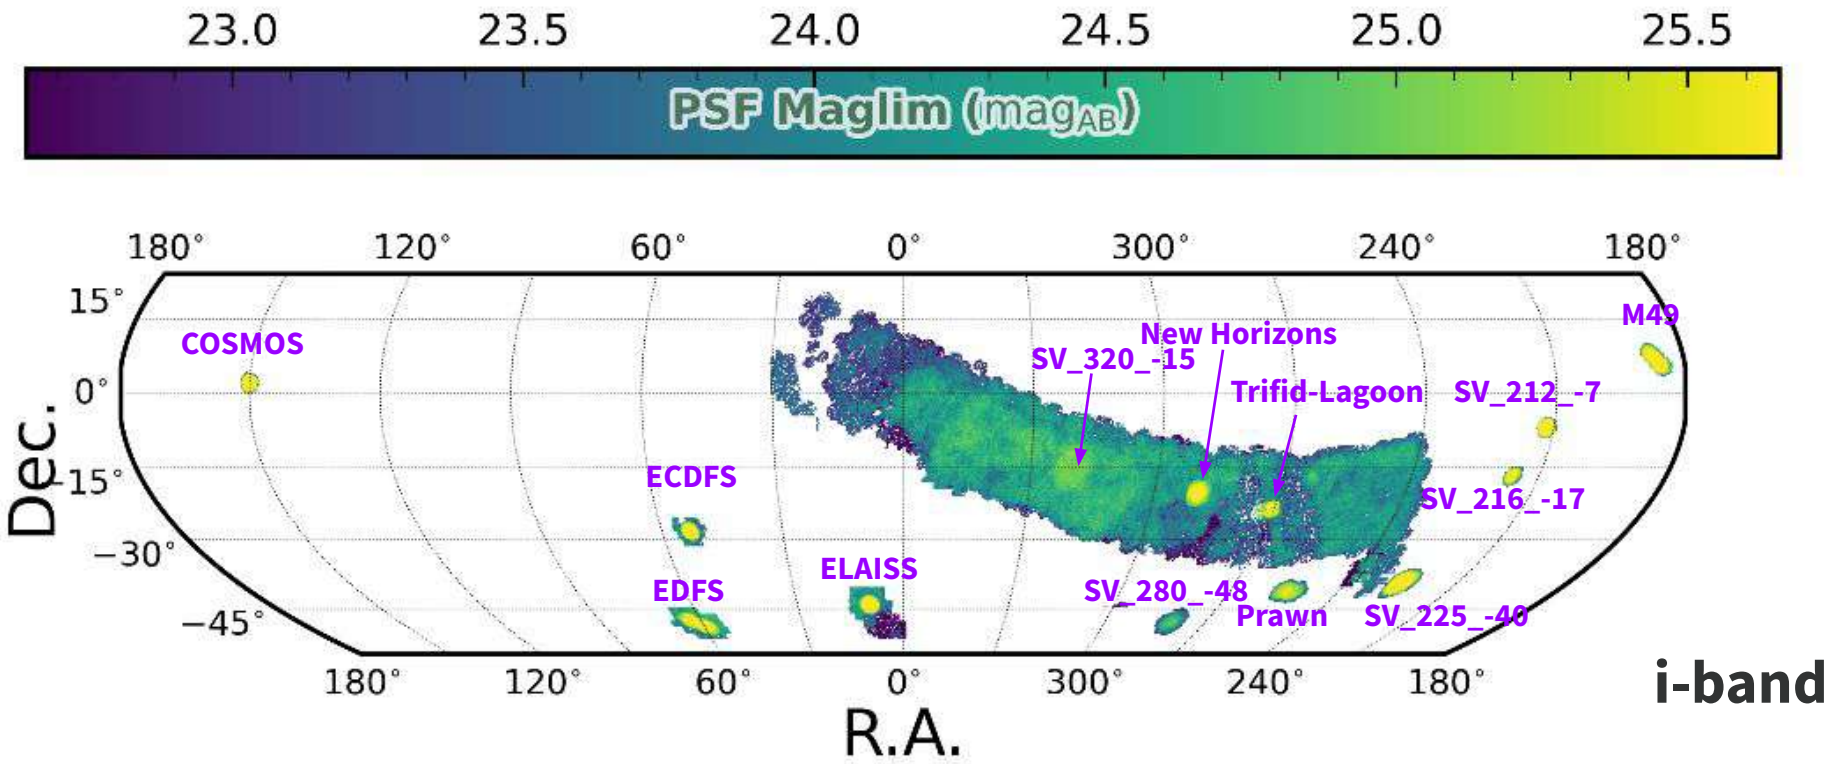
Single visit-level data products only:
Visit images + Source Catalogs

Multi-epoch data products:

- Coadded images, difference images
- Objects, DiaSources, DiaObjects, ForcedSources (lightcurves)

DR1 will be at least 10x more data

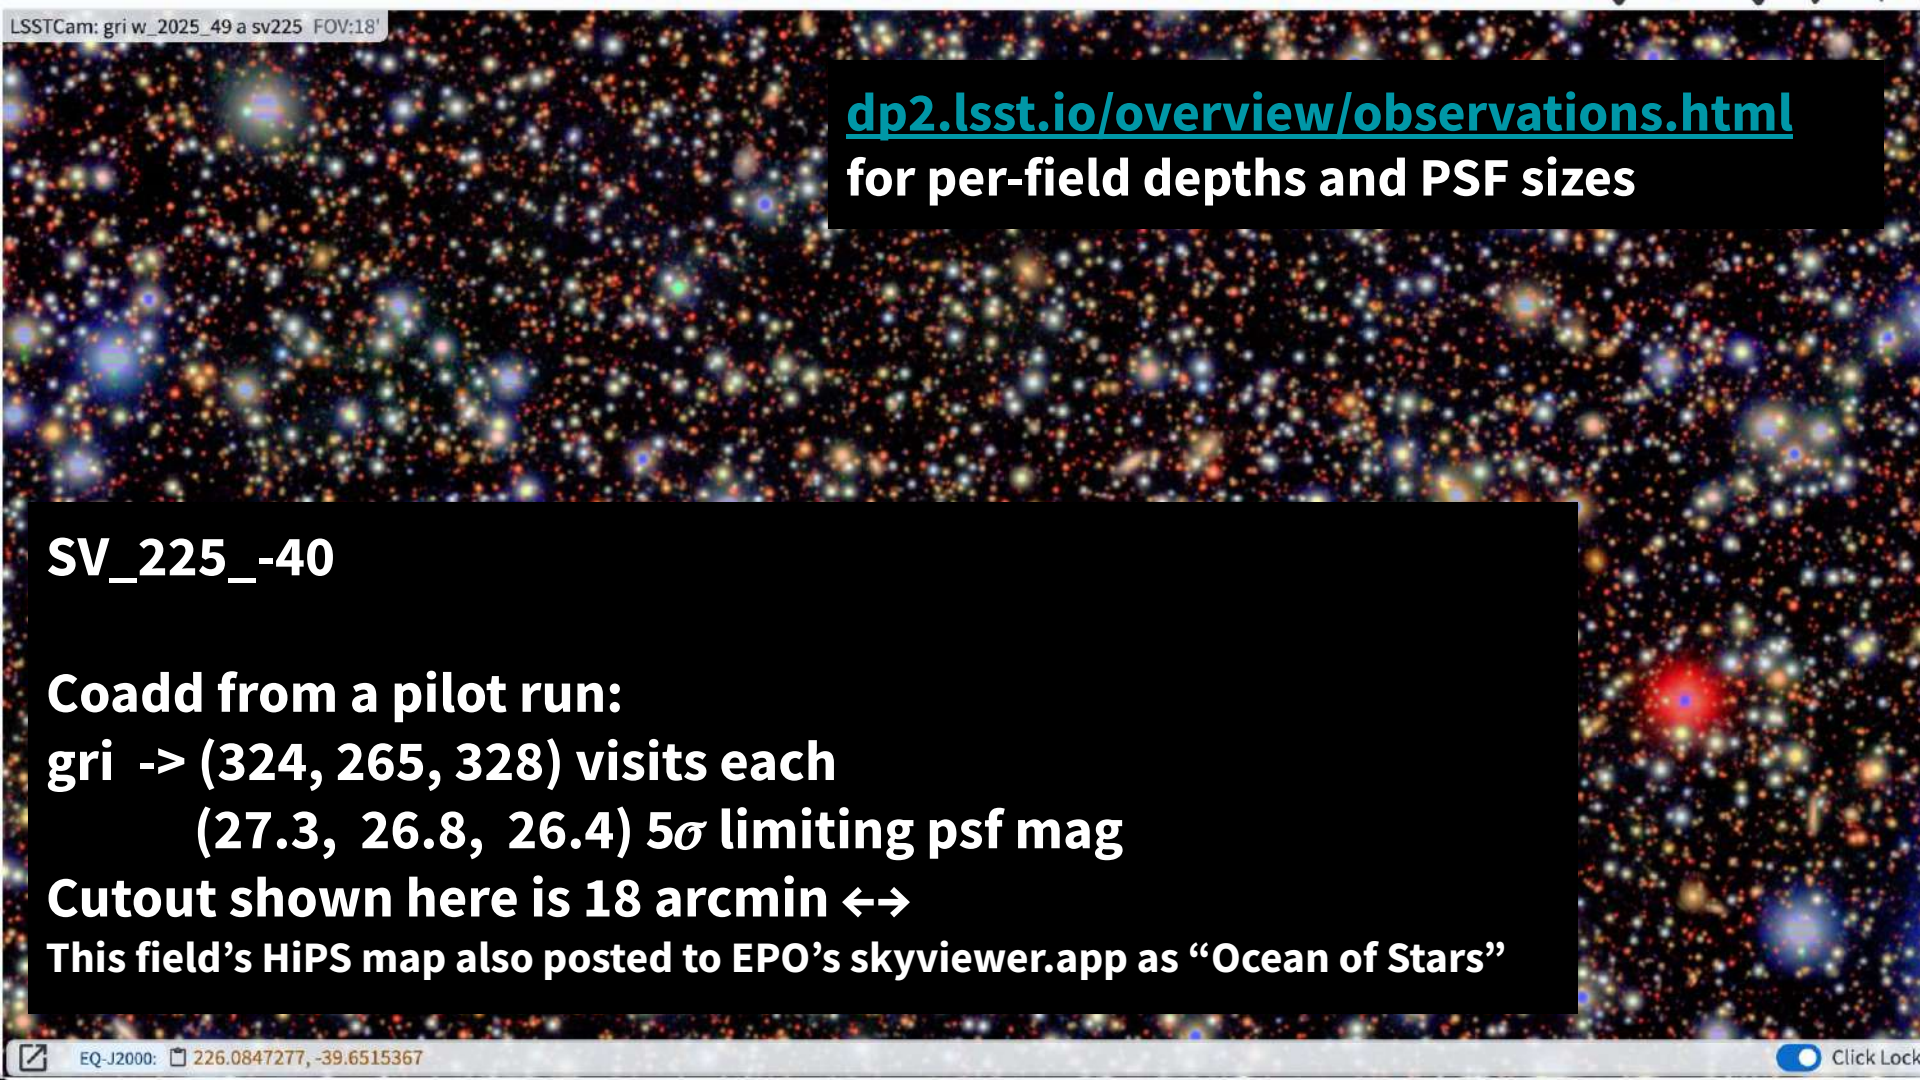
DP2 - Deep coadd depth varies across footprint



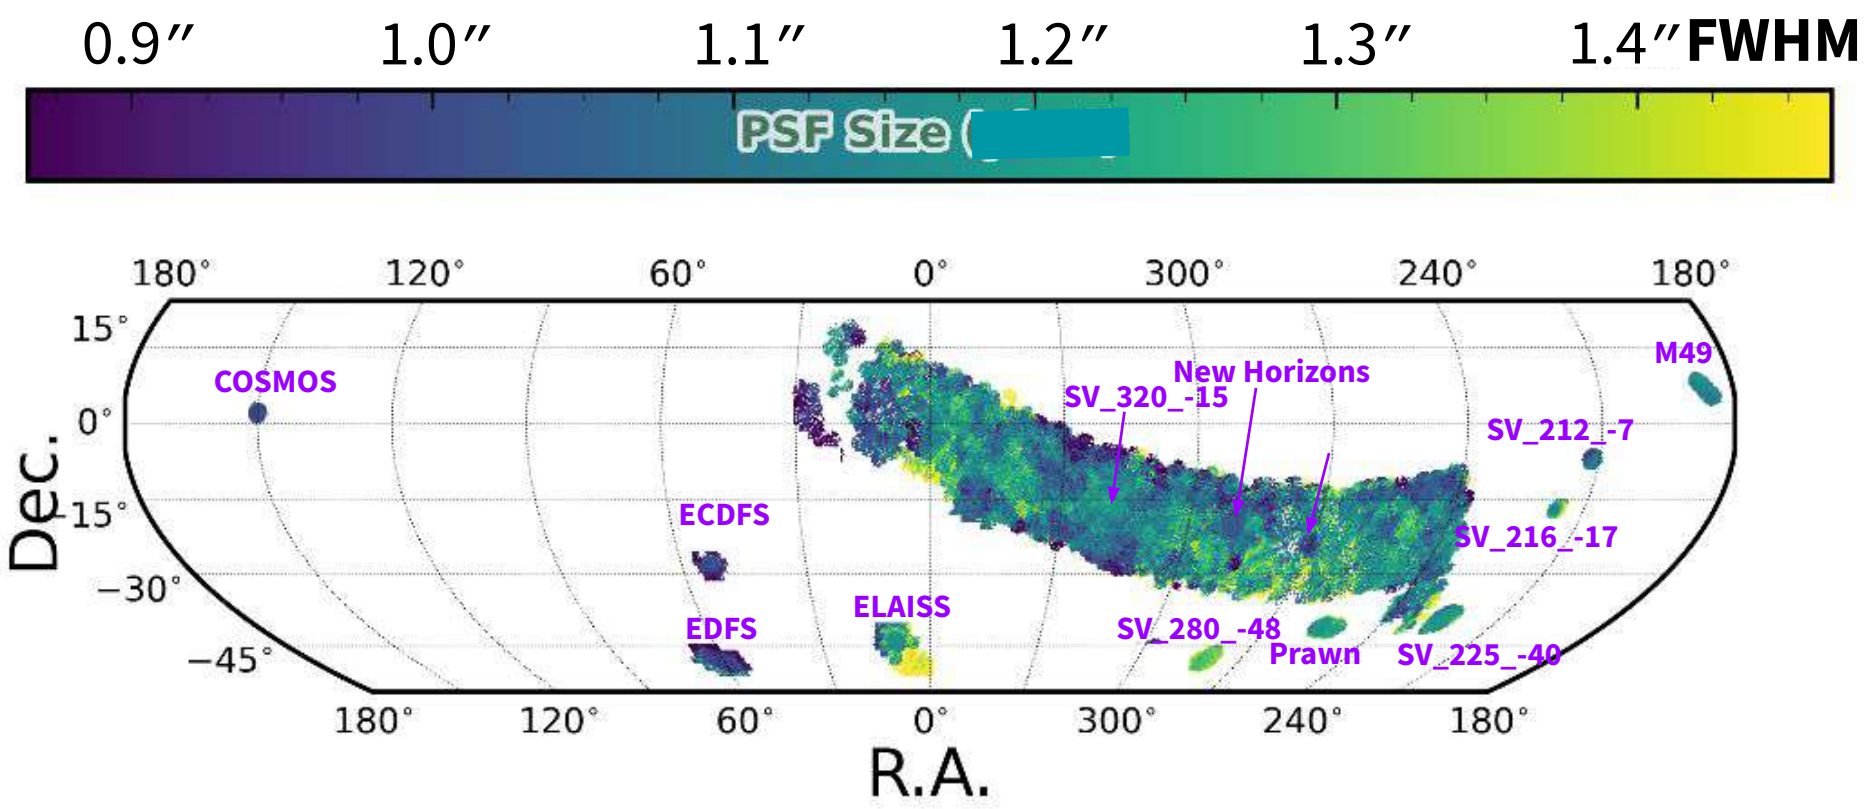
i-band


DP2 gri pretty_coadd FOV:9.6'
COSMOS
gri (67, 133, 121) visits each
(26.5, 26.2, 26.8) 5σ limiting psf mag each
Cutout shown here is 9.6 arcmin ←→
This field's HiPS map also posted to EPO's skyviewer.app
EQ-J2000: 148.9525431, 1.5537083

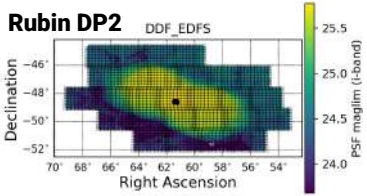

DP2 gri pretty_coadd FOV:18'
EDF-S
gri -> (31, 33, 57) visits
(25.6, 25.5, 25.7) 5σ limiting psf mag
Cutout shown here is 18 arcmin ↔
EQ-J2000: 64.1990033, -47.8156305

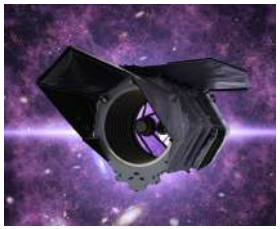

LSSTCam: gri w_2025_49 a sv225 FOV:18'
dp2.lsst.io/overview/observations.html
for per-field depths and PSF sizes
SV_225_-40
Coadd from a pilot run:
gri -> (324, 265, 328) visits each
(27.3, 26.8, 26.4) 5σ limiting psf mag
Cutout shown here is 18 arcmin ↔
This field's HiPS map also posted to EPO's skyviewer.app as "Ocean of Stars"
EQ-J2000: 226.0847277, -39.6515367
Click Lock

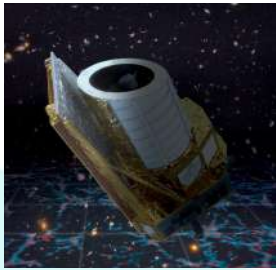
DP2 - Deep coadd i-band average PSF (image quality)



A Public Euclid Deep Field South (EDF-S) Dataset for Building Rubin/Roman/Euclid Joint Processing

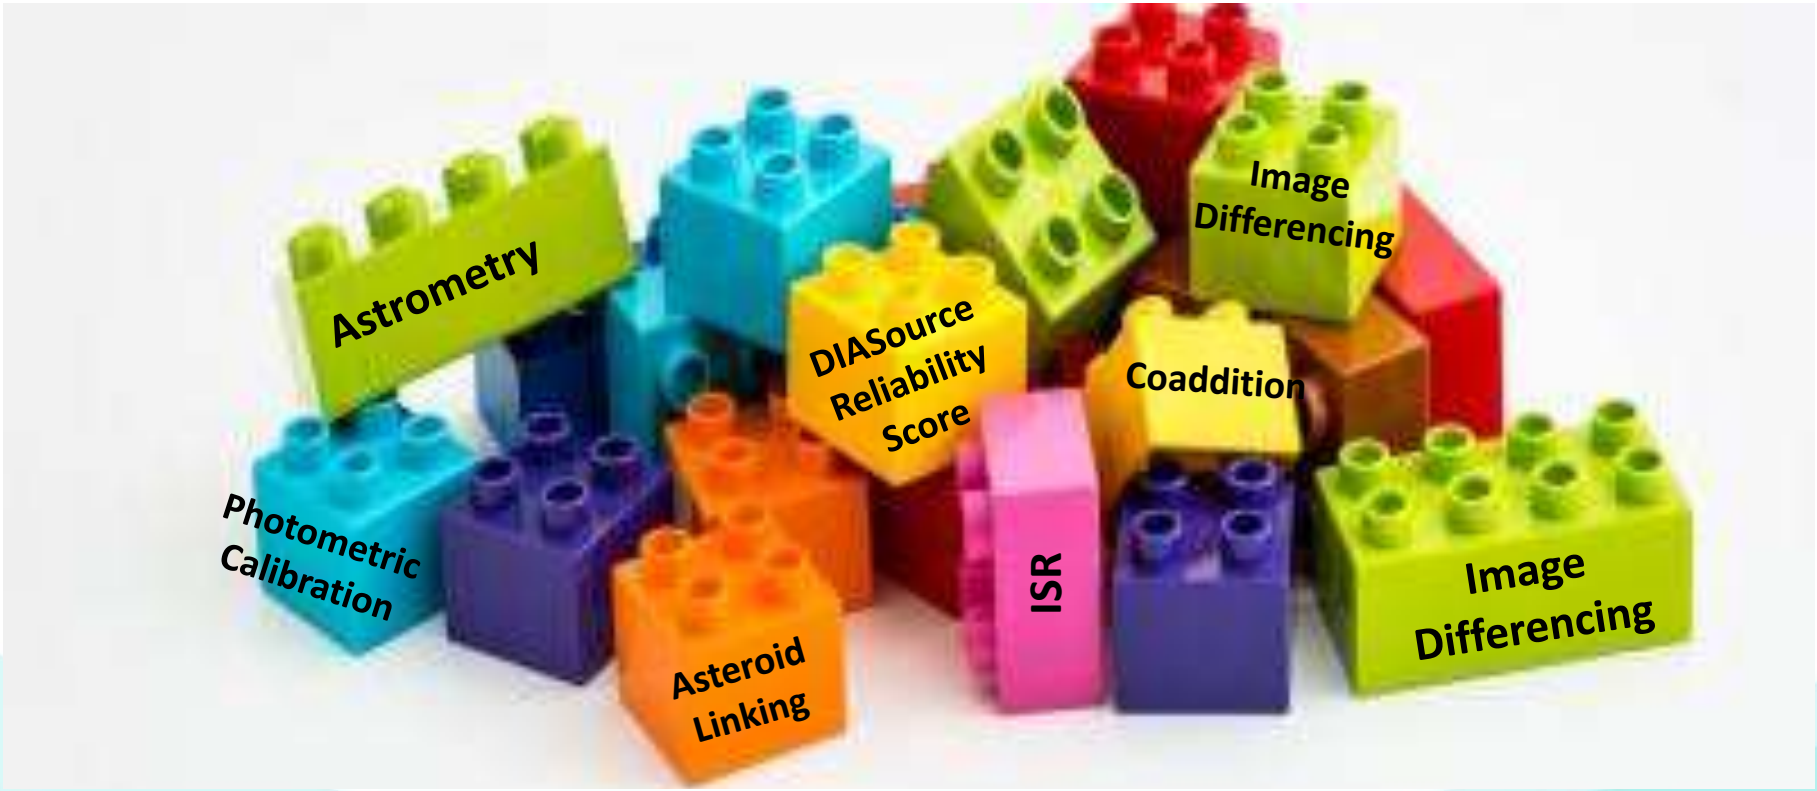


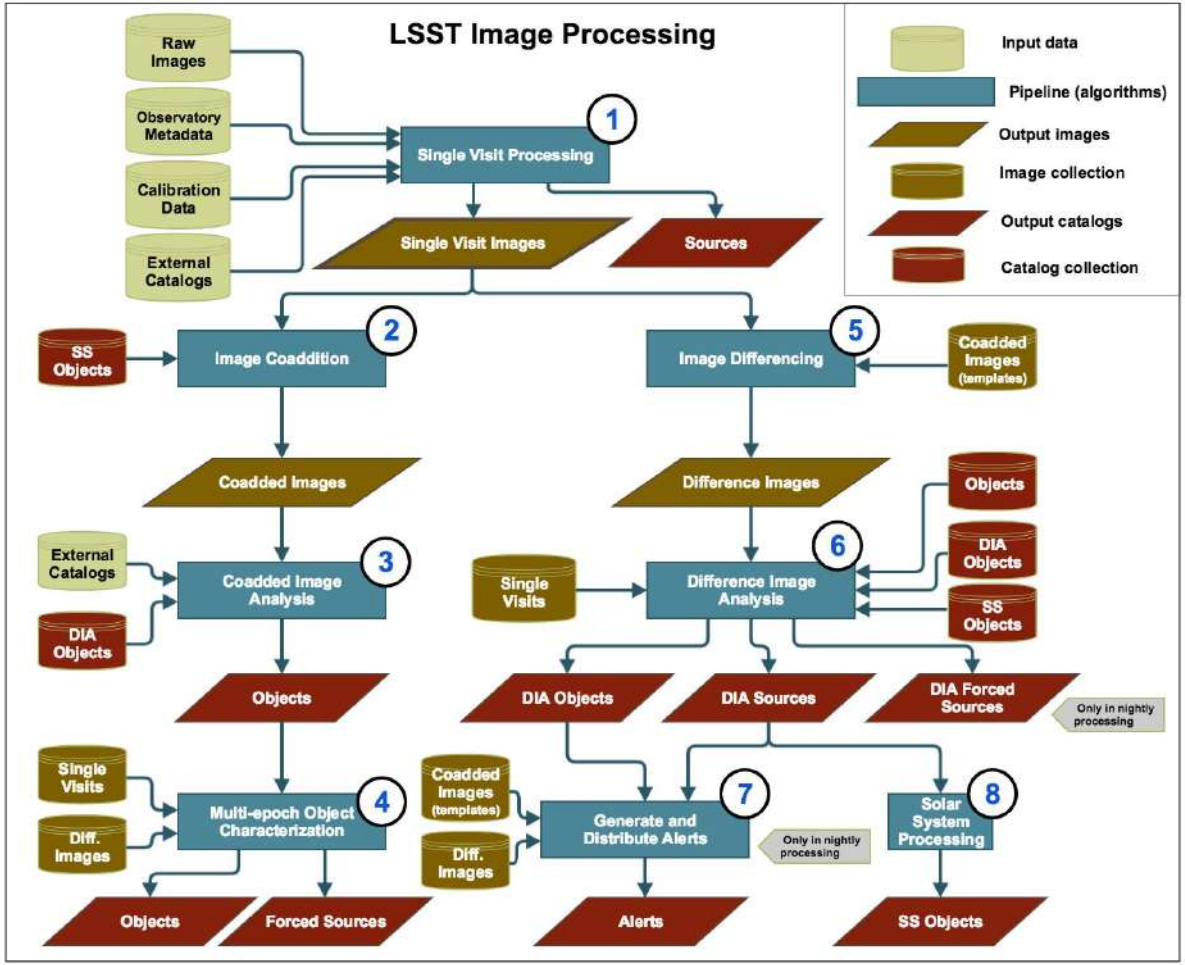




- By 2027, Rubin, Euclid & Roman will be producing overlapping petabyte-scale datasets. Co-processing unlocks science no survey can do alone.
- Yet each runs different data access policies, hindering the overlapping community from building the shared infrastructure and analyses needed to realize this science.
- **Our proposal: make the 23 sq. deg. Rubin EDF-S data fully public to provide one common open dataset across all three surveys to build infrastructure for joint processing (rtn-047.lsst.io)**
- COSMOS did this 20 yrs ago and remains a cornerstone field today

Alert Prod (AP), Data Release Prod (DRP) Pipelines constructed from the same algorithmic components called Tasks




LSST Image Processing
Input data
Pipeline (algorithms)
Output images
Image collection
Output catalogs
Catalog collection
Raw Images
Observatory Metadata
Calibration Data
External Catalogs
1
Single Visit Processing
Single Visit Images
Sources
SS Objects
2
Image Coaddition
Coadded Images
External Catalogs
DIA Objects
3
Coadded Image Analysis
Objects
Single Visits
Diff. Images
4
Multi-epoch Object Characterization
Objects
Forced Sources
5
Image Differencing
Coadded Images (templates)
Difference Images
Single Visits
6
Difference Image Analysis
Objects
DIA Objects
SS Objects
DIA Objects
DIA Sources
DIA Forced Sources
Only in nightly processing
Coadded Images (templates)
Diff. Images
7
Generate and Distribute Alerts
Only in nightly processing
Alerts
8
Solar System Processing
SS Objects

Optimizing the Algorithms & Pipelines for LSSTCam

First look coadds in 3mo, public alerts in 9mo, Data Release in 15mo after first photon.

Only possible because we developed the pipelines on precursor datasets like Hyper Suprime-Cam.

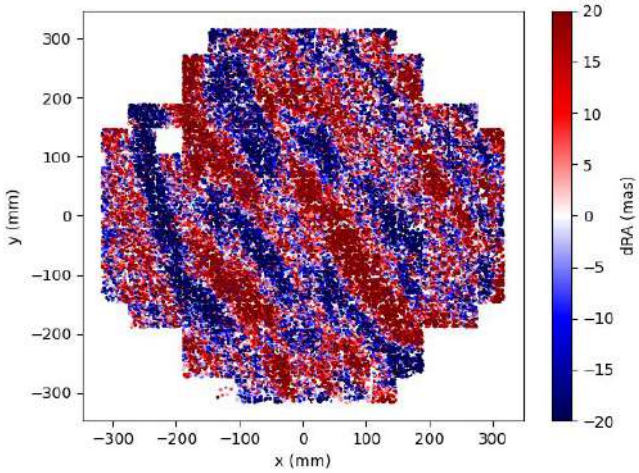
But, every camera + survey is unique. Aimed to optimize pipelines for LSSTCam to release commissioning data ASAP. Mantra was: DP2 fast. DR1 good.

More pipeline work for DR1/DR2 to achieve performance needed.

DES took 6 years of pipeline development for cosmology results.

Examples of LSSTCam specials:

- **2x ccd vendors** => 2x the detector effects and focal plane discontinuities.
- **Extra-large focal plane** => structure in photometric transmission (clouds)
- **Short exposure times** => atmospheric turbulence in astrometry



Example of Astrometric Residuals in visit 2025121600090 due to **atmospheric turbulence.** Developed new Task to model it, which reduced **median astrometric repeatability** from ~13 mas to **~8mas.**

Science Performance on DP2 data

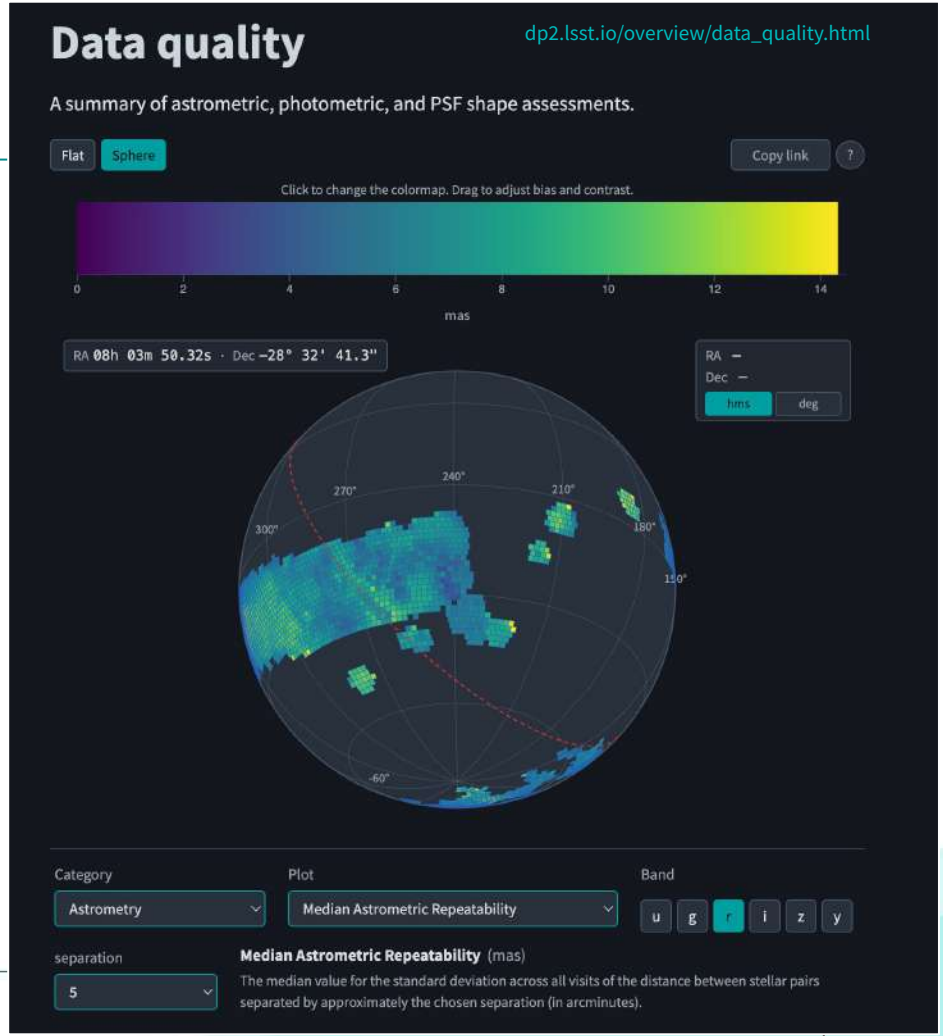
dp2.lsst.io/overview/data_quality.html→
DP2 Paper

Astrometry:

- Relative repeatability ~8 mas →
  - design spec 10mas
- Median absolute match to Gaia DR3 ~1mas
  - (design spec 50mas)

Photometry:

- Relative repeatability 3-4 mmag (*griz*)
  - (design spec 5 mmag)



In summary

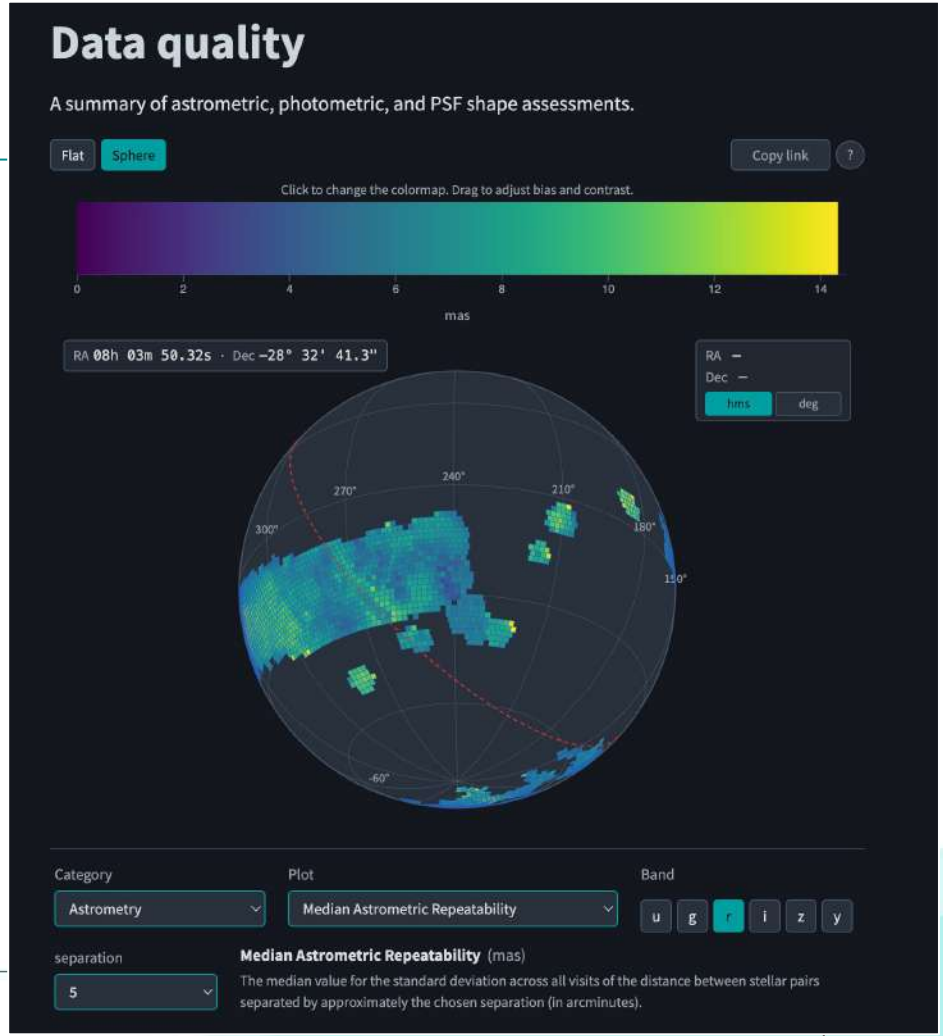
Rubin adapted the LSST Science Pipelines for LSSTCam as fast as possible so that we could get feedback from YOU before DR1.

=> rubin.community is there to get your feedback

If you see something weird, say something

Check against known issues:
dp2.lsst.io/overview/known_issues.html

Docs at
prompt-products.lsst.io
dp2.lsst.io



Q&A Menu

What do we know about DR1?

Performance of: Astrometric Calibration, Photometric Calibration, Image subtraction, background subtraction

Expected pipeline changes between DP2 and next DR.

Image format change between DP1 and DP2

Happy to talk about these during the breaks.

Also, for more detail of each algo see pipelines talk (video, slides) from July 2026 at the RCW talk.

Appendix

New Pipeline Features for the next data release

These are in progress, and the later we start processing DR1, the more of ths features in (addition to data!) will be included:

- Wavelet-based, multi-band coadd background subtraction and detection
- Chromatically corrected flat-fielding. Sky colors. Multi-LED Flats.
- Chromatic PSFs
- Automatic exclusion from coadds: images with unmodelable PSFs, clouds, ice-clouds
- Deblending on difference images
- Add a couple photo-z's to the object table
- Releasing the preliminary isolated star association tables that are used as **input** to astrometric and photometric calibration (in addition to the parallax and proper motion tables that were included in DP2)

New Pipeline Features for the next data release

Potential table schema changes:

- Removing the forced photometry on the direct science images in ForcedSource Tables?
- Removing unnecessary columns from the Source table

Fixes for the various Known Issues (for example)

- SPIKE mask sizes
- Add systematics/noise floors to error columns
- 2-D Gaussian process interpolation for pixel interpolation in ISR and no-data regions in coadds.
- n_image HIPS, per-cell table of coadd inputs, to make it easier way to find out how many visits have gone into a coadd.

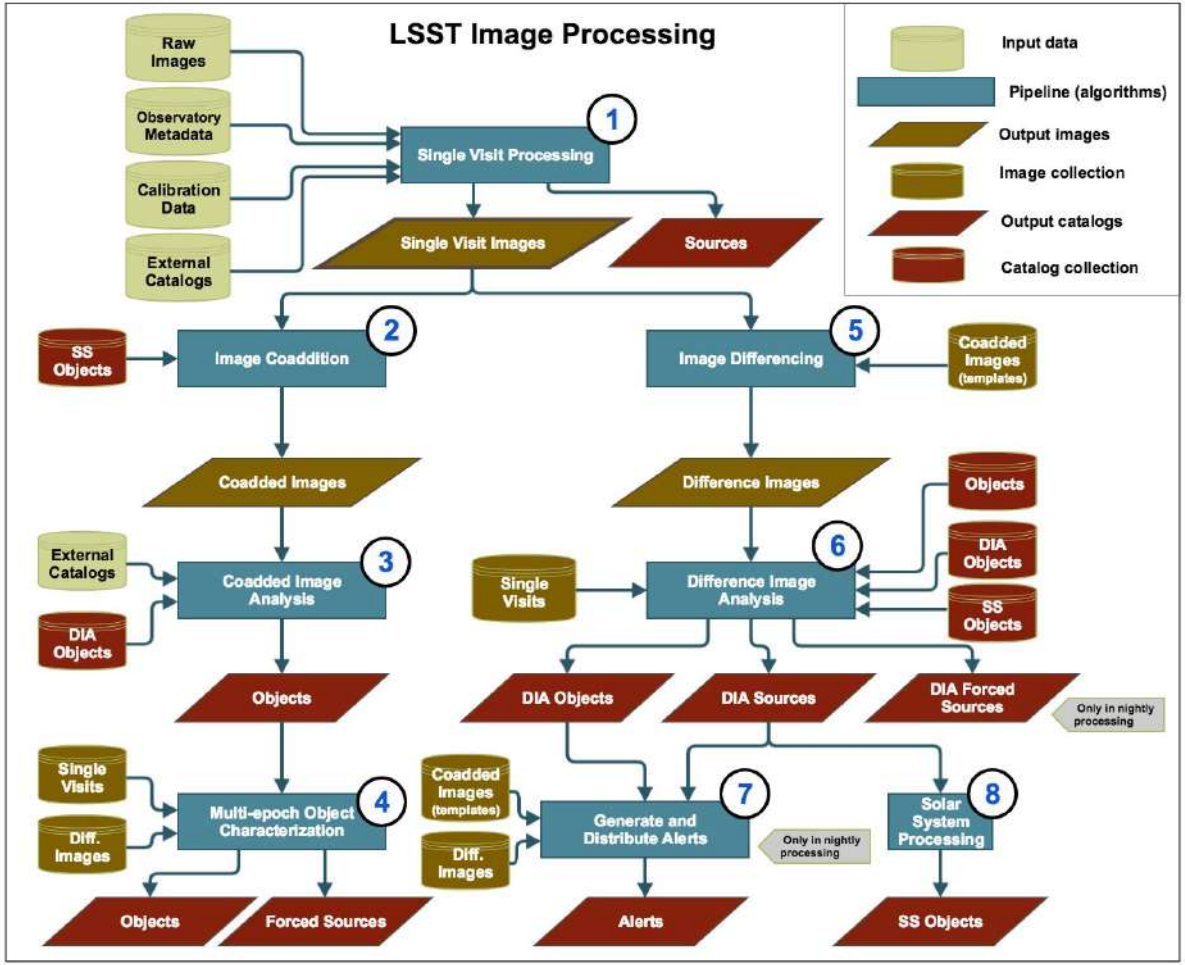

LSST Image Processing
Input data
Pipeline (algorithms)
Output images
Image collection
Output catalogs
Catalog collection
Raw Images
Observatory Metadata
Calibration Data
External Catalogs
1
Single Visit Processing
Single Visit Images
Sources
SS Objects
2
Image Coaddition
Coadded Images
External Catalogs
DIA Objects
3
Coadded Image Analysis
Objects
Single Visits
Diff. Images
4
Multi-epoch Object Characterization
Objects
Forced Sources
5
Image Differencing
Coadded Images (templates)
Difference Images
Single Visits
6
Difference Image Analysis
Objects
DIA Objects
SS Objects
DIA Objects
DIA Sources
DIA Forced Sources
Only in nightly processing
Coadded Images (templates)
Diff. Images
7
Generate and Distribute Alerts
Only in nightly processing
Alerts
8
Solar System Processing
SS Objects

Image processing, source detection & measurement

Instrument signature removal (ISR)
- image reduction for **raw** images

Calibration
- astrometric solution (WCS)
- photometric calibration (flux units nJy)
- PSF estimation
- result: **visit_image** (available DP2)

Coaddition
- mean stacking, weighted by inverse variance
- all images with PSF FWHM < 1.7 → **deep_coadd**
- the 1/3 best-seeing images → **template_coadd** (DP2)

Difference image analysis
- subtract templates from visit images
- result: **difference_image** (DP2)

Detection
- “detected” means a signal-to-noise ratio > 5

Deblending
- Scarlet Lite deblender algorithm

Measurement
- centroid, size, shape, and flux measurements
- fit with, e.g., PSF, Gaussian, or extended models

Forced photometry
- a measurement is made at a fixed coordinate
- can be sub-threshold (SNR<5)

Moving object linking
- individual detections are linked into moving objects

Catalog descriptions

Direct and coadd image analysis Tables

Source
- measurements for SNR>5 detections in **visit_images**

Object
- multi-band forced photometry at the positions of all single-band detections in the **deep_coadd** images

ForcedSource (on Object)
- Multi-epoch forced photometry at the postions of all **objects** in every **visit_image** and every **difference_image**
- lightcurves for **objects**

Difference image analysis (DIA) tables

DiaSource
- measurements for SNR>5 detections in **difference_images**

DiaObject
- positions of associated DiaSources (radius 1”)
- variability characterization parameters

ForcedSourceOnDiaObject
- Mulit-epoch forced photometry at the positions of all **DiaObjects** in every **visit_image** and every **difference_image**
- lightcurves for **DiaObjects**

Catalog descriptions

Moving objects Tables

DiaSource
- measurements for SNR>5 detections in **difference_images**

SSObject
- associations of DiaSources into moving objects

SSSource
- instantaneous physical properties of an SSObject at the time of every detection as a DiaSource

MPCORB
- derived orbital parameters from the Minor Planets Center (MPC), not limited to LSST detections

Observational metadata tables

Visit
- individual observations
- refers to full focal plane / boresight
- date, time, band, coordinates, airmass, etc.

VisitDetector
- individual observations with each detector
- one row per detector per visit
- airmass, PSF FWHM, zeropoint, magnitude limit

Browse the table schemas:
sdm-schemas.lsst.io/dp2.html

Data Preview 2 Processing:

Took place over ~3-4 months, used 2.8 million CPU hours

DR1 expected to be 10 times the size

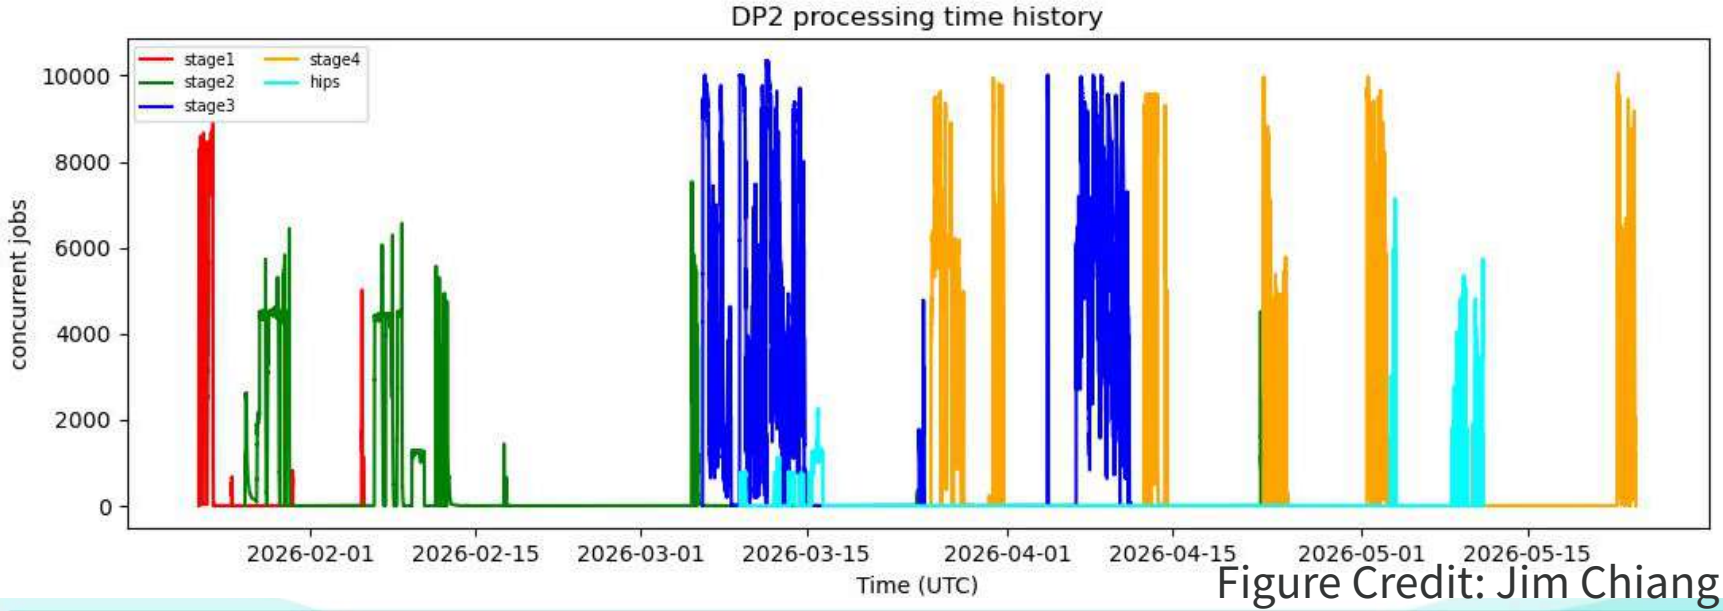
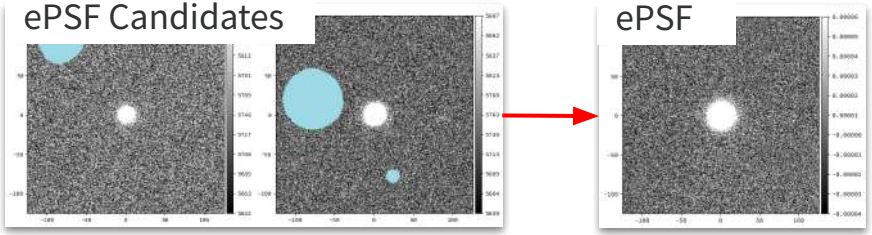


Figure Credit: Jim Chiang

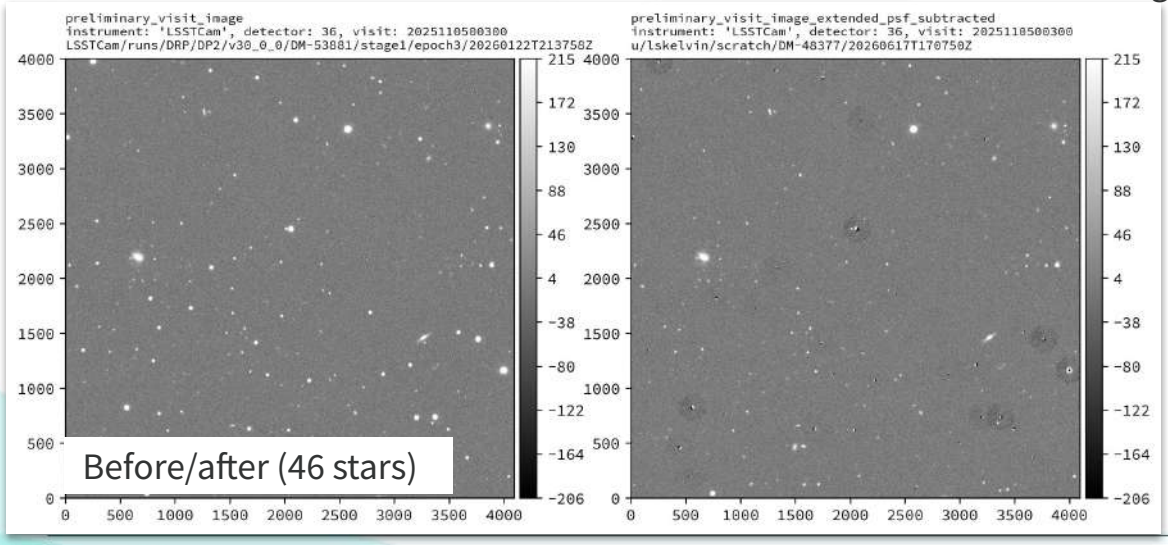
Extended PSF Generation (DM-48377)

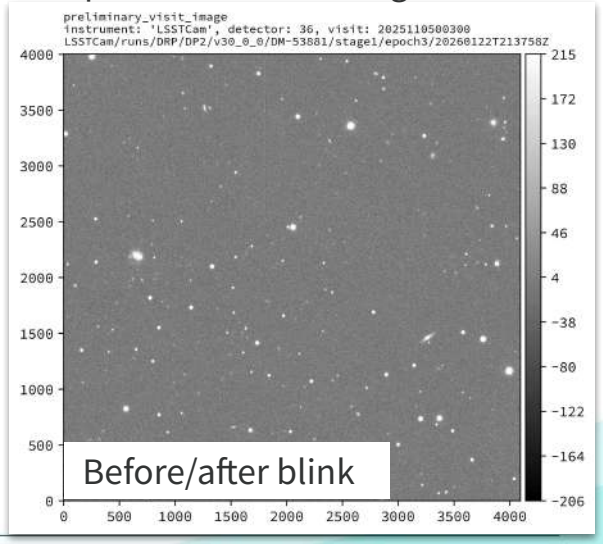


Final ePSF subtraction task in-review on DM-48377;

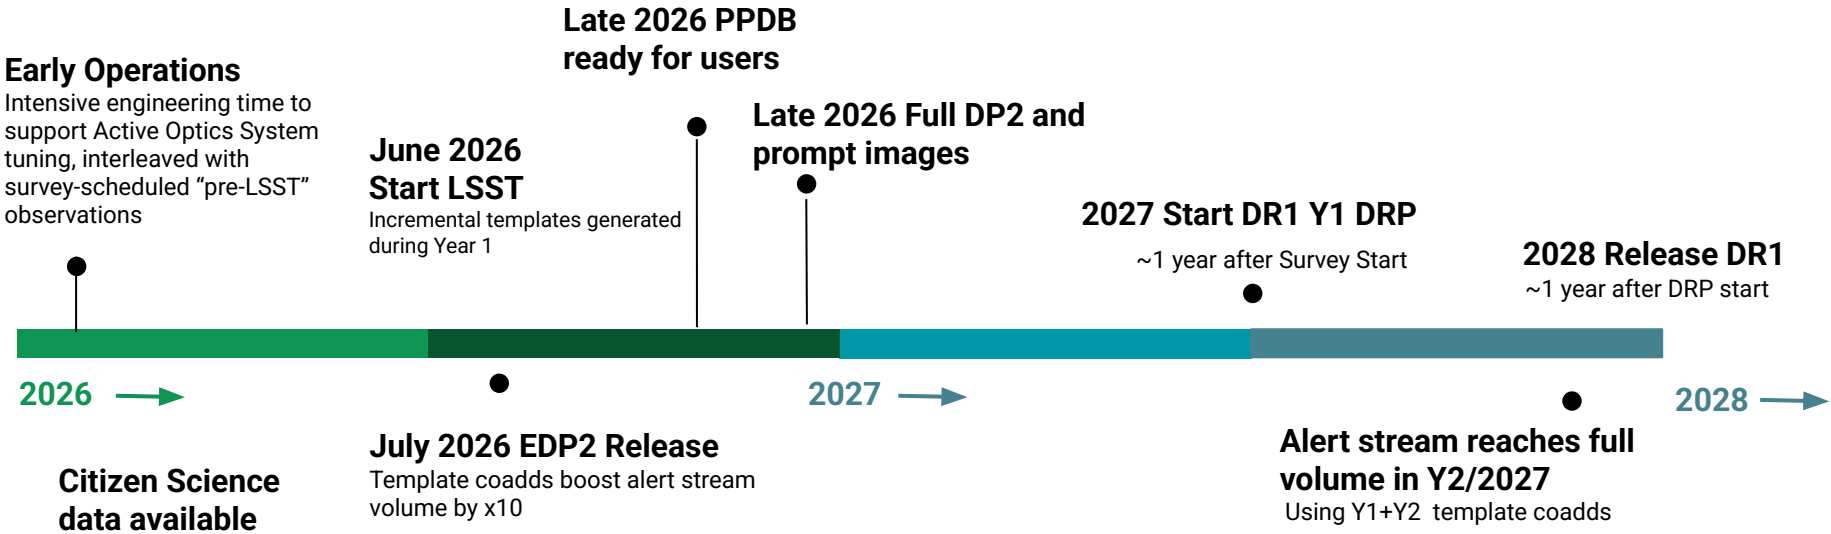
This completes basic ePSF pipeline task framework → decisions to be made now on how/where to implement.

Need to experiment increasing S/N of ePSF candidates to resolve wing noise amplification → stacking multi-visit.





Rubin Y1-Y2 Timeline



The DR1 clock has started, and a pattern of regular survey data releases will follow

Speaker: Jim Bosch
3 min

Introducing `lsst.images`

DM is rolling out new image classes and file formats to (mostly) replace those in `afw`.

- The old file format was too opaque - nearly impossible to document.
- The old file format was bad for partial reads over network I/O.
- The `afw` types have a ton of footguns, and removing them without breaking pipeline code is difficult.
- The split between C++ and Python makes `afw` hard to document, distribute (no `pip install`), and maintain.

The new types and formats will be a lot easier to maintain and support over the course of the survey. We should have done this long ago, but better now than never.

Speaker: Jim Bosch
3 min

More information about `lsst.images`

Documentation is at https://images.lsst.io, with more on the rationale in DMTN-339.

Docs are under heavy development right now, but we've already got:

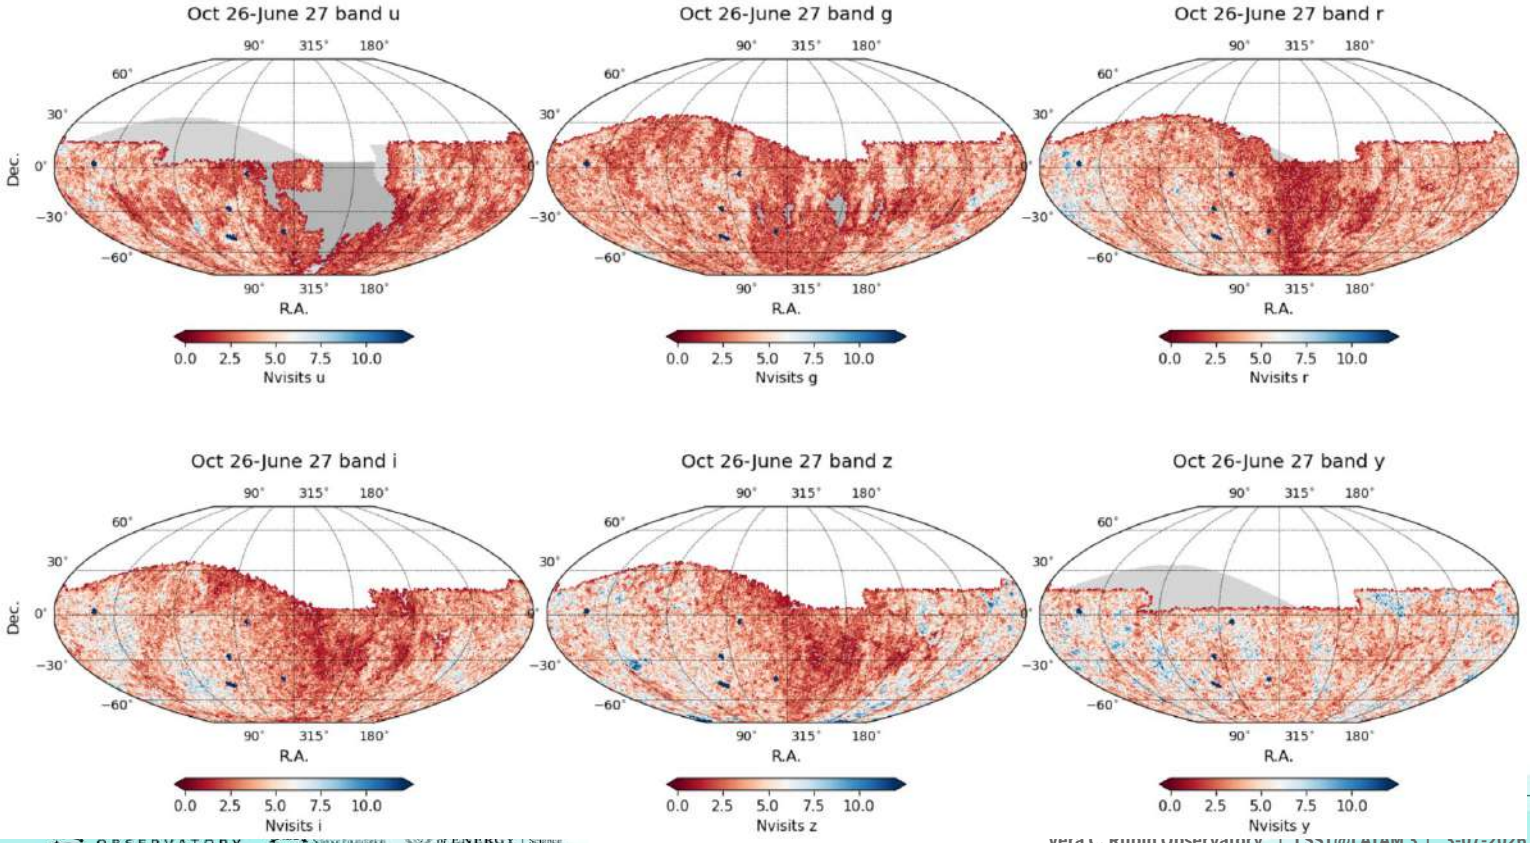
- a "for afw users" guide
- pretty good API reference documentation

with a lot more on the way.

The new types can't do everything the old types can yet:

- image display still goes through `afw`
- essentially all `Tasks` still use `afw`.

Updated simulations show expected number of visits between June ‘26 -’27 not sufficient to coadd area more than DDFs

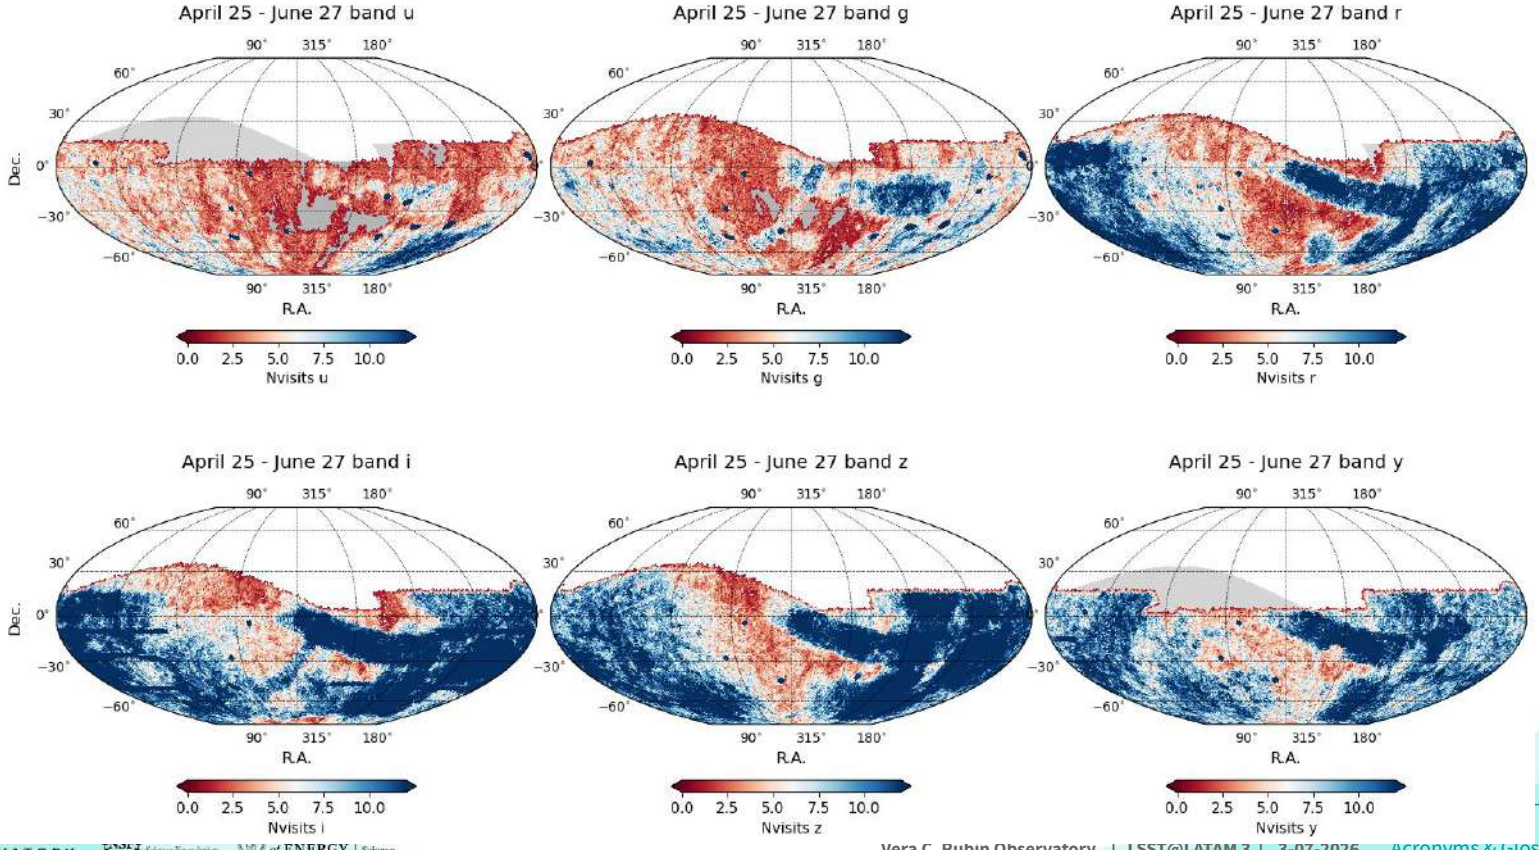


White to blue: 6+ visits (minimum raw visits for coaddition after quality cuts)

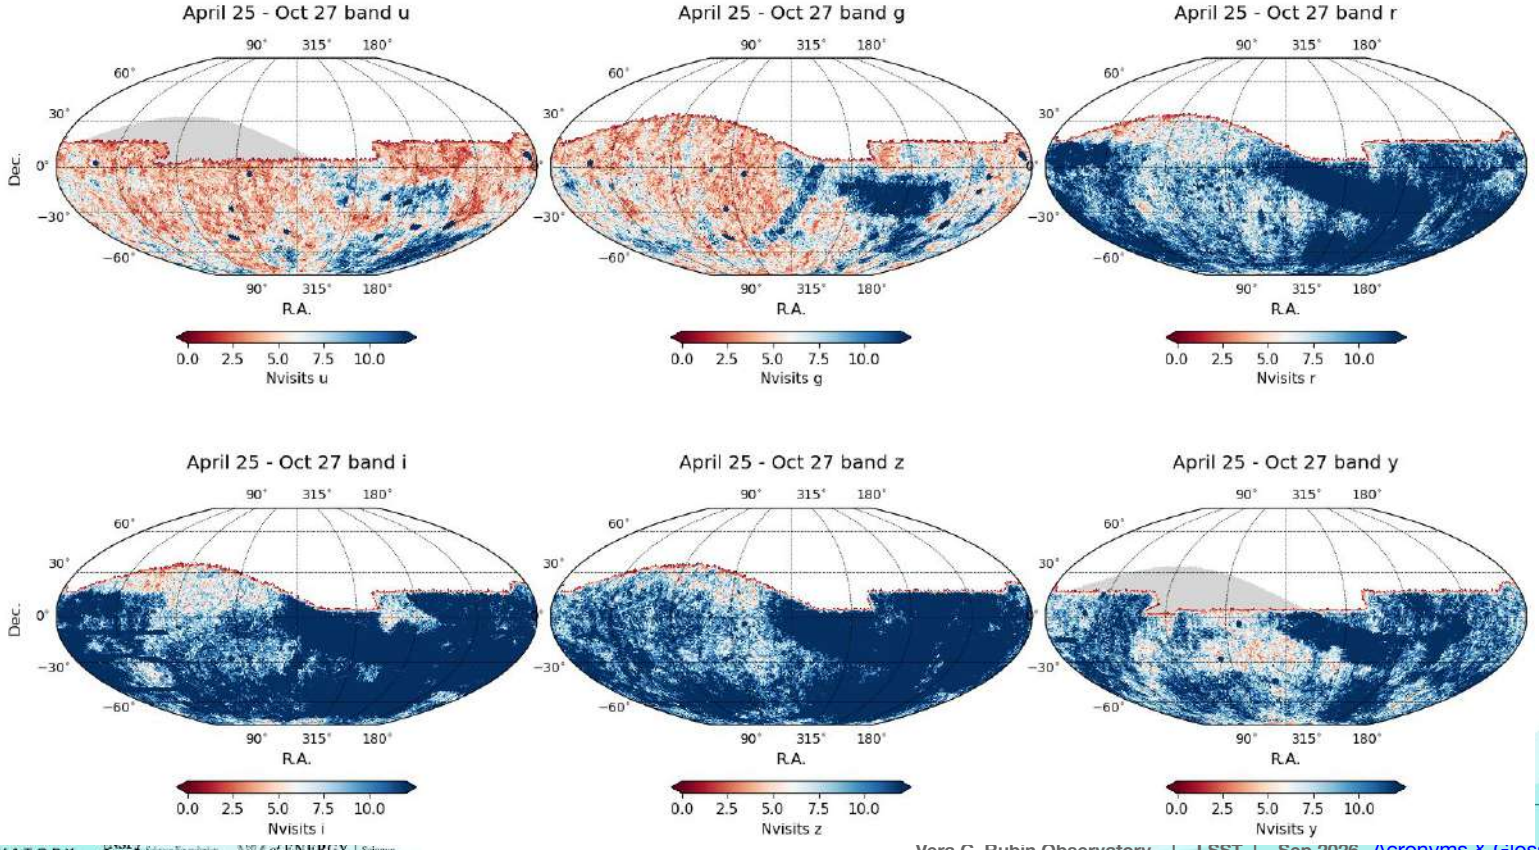
White to red: <6 visits (uncoaddable)

Simulations from Lynne Jones

Existing visits + 8 months



Existing visits + 1 year



Data Release 1 (DR1) and the near term strategy for the LSST is our primary challenge to address with the community

Anticipate one year of processing for the first year of LSST.

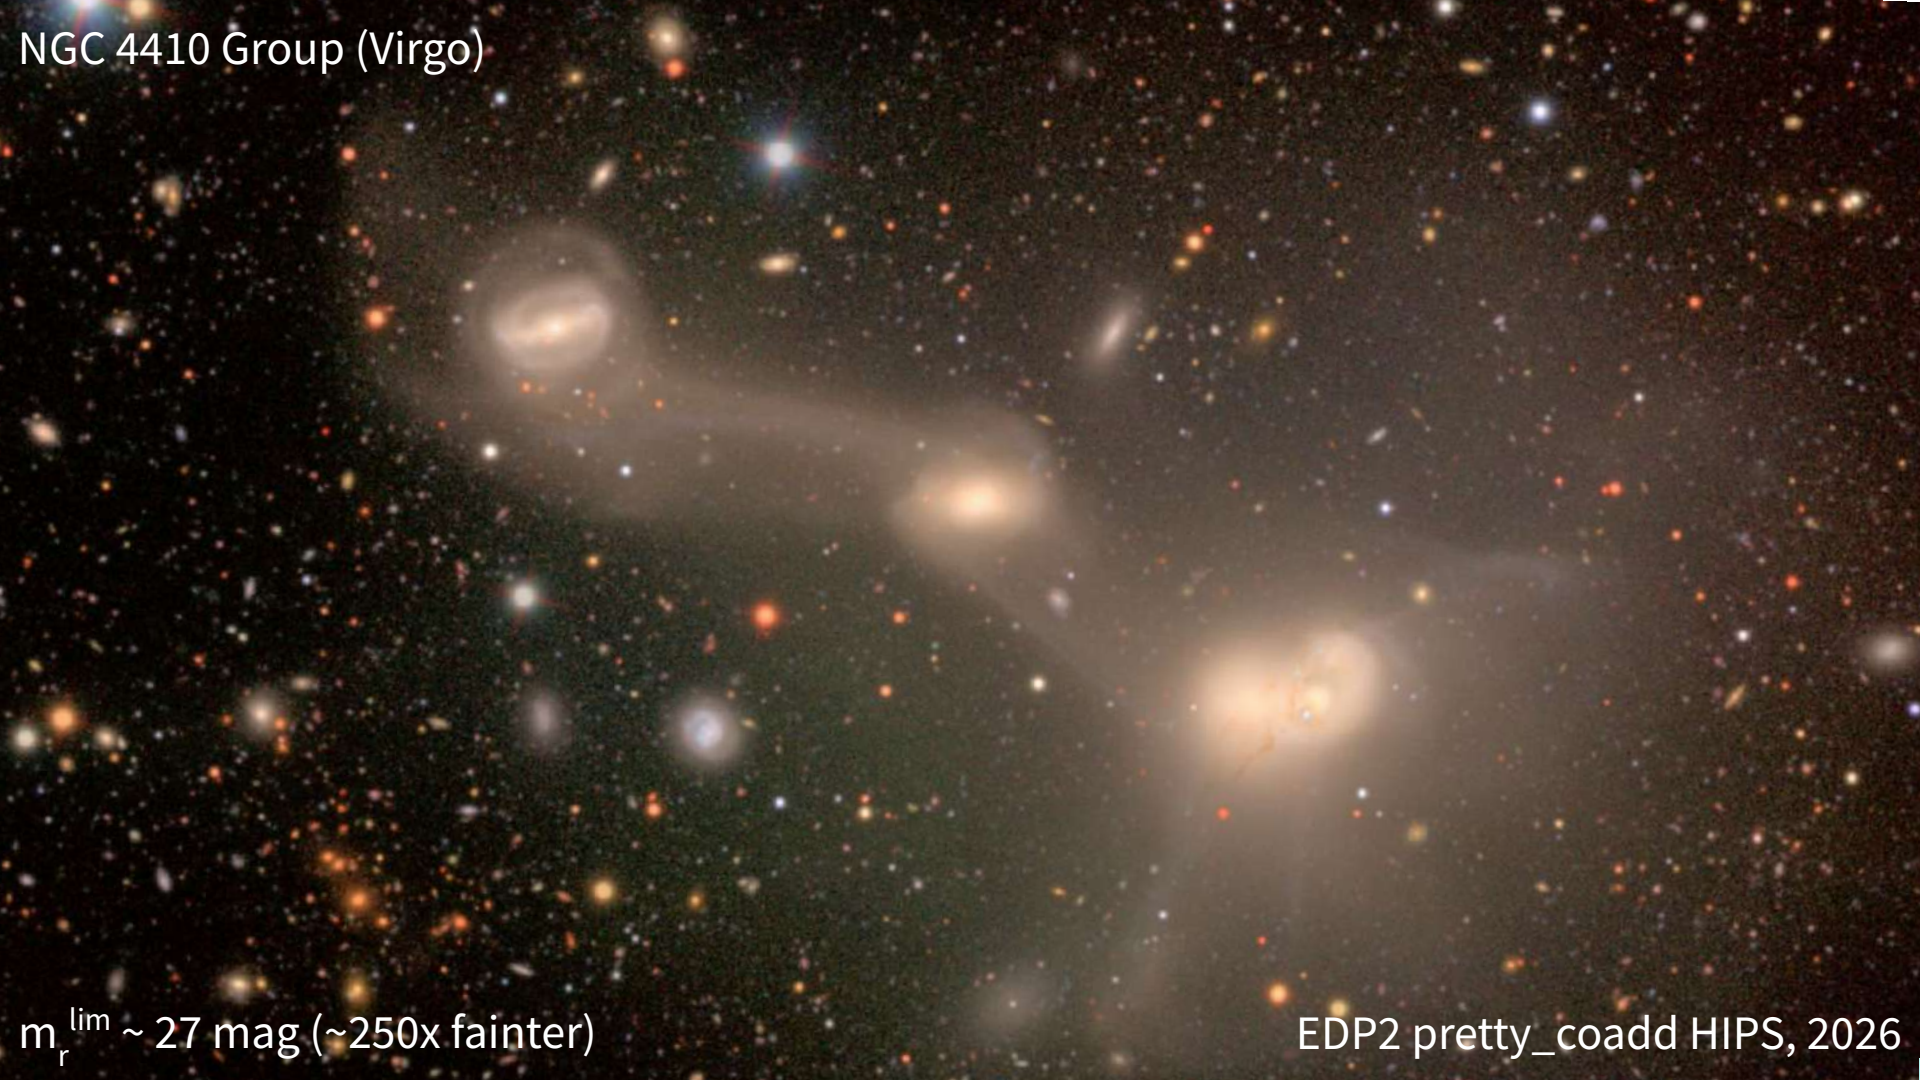
DR1 nominally expected mid 2028

- We will need to balance getting data to the community as soon as possible with the need to reach a uniform Wide Fast Deep data release early in the survey.
- We will need to build templates in all filters as soon as possible.
- We will need to balance science goals with Rubin resources.

Data Release Board will make a recommendation this fall.

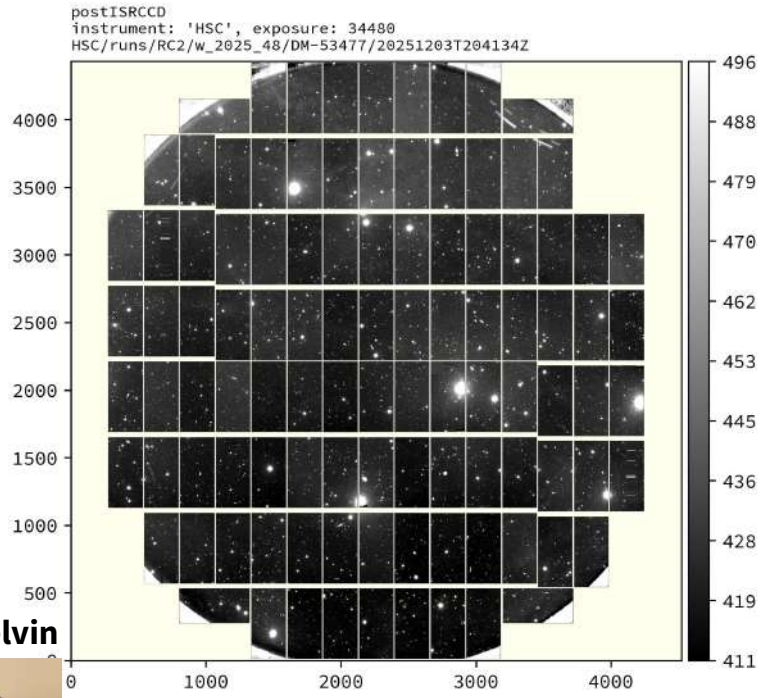
Survey Cadence Optimization Committee **considering tradeoffs** for which of footprint area, bands, including pre-LSST data, delaying cutoff until Oct ‘27 can yield

Pipeline Performance and Known Issues

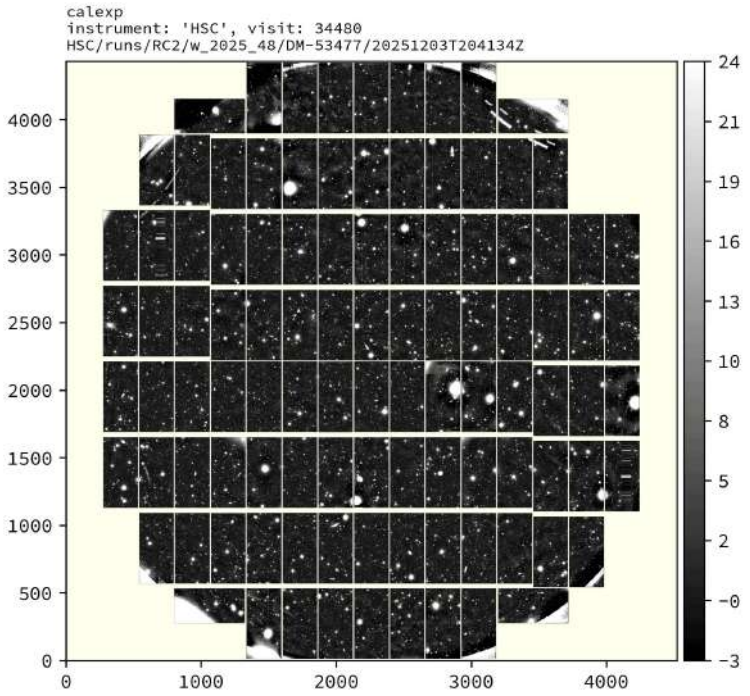

NGC 4410 Group (Virgo)
m_r^{lim} ~ 27 mag (~250x fainter)
EDP2 pretty_coadd HIPS, 2026


NGC 4410 Group (Virgo)
Slide Credit: Lee Kelvin
m_r^{lim} ~ 27 mag (~250x fainter)
Rubin First Look, 2025

SkyCorrectionTask in HSC (g): tried and true



Input: post-ISR

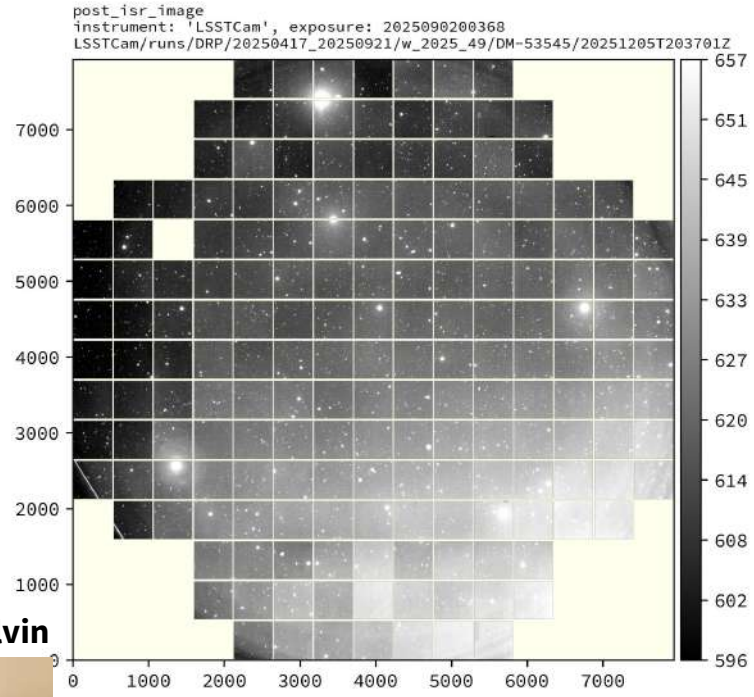


Output: with/without sky correction

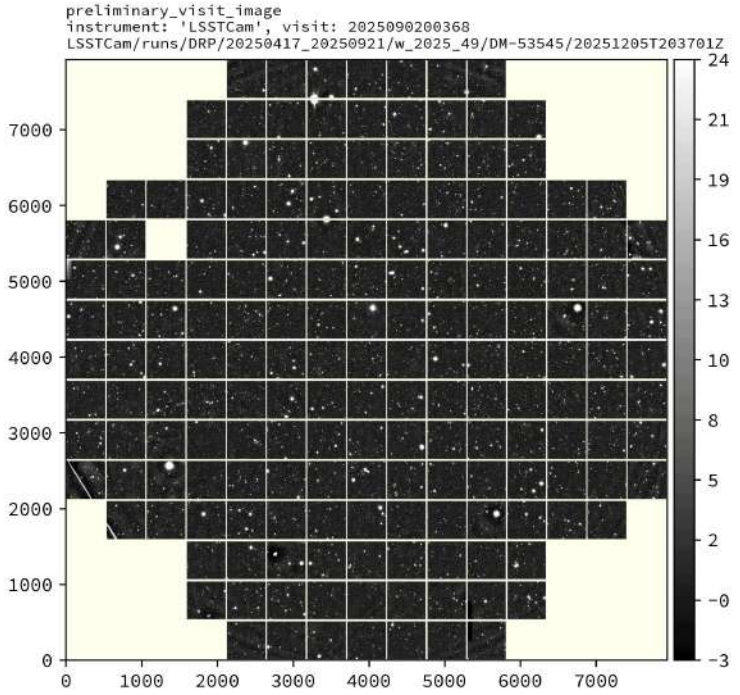
Lee Kelvin



SkyCorrectionTask in LSSTCam (g): works as expected…



Input: post-ISR

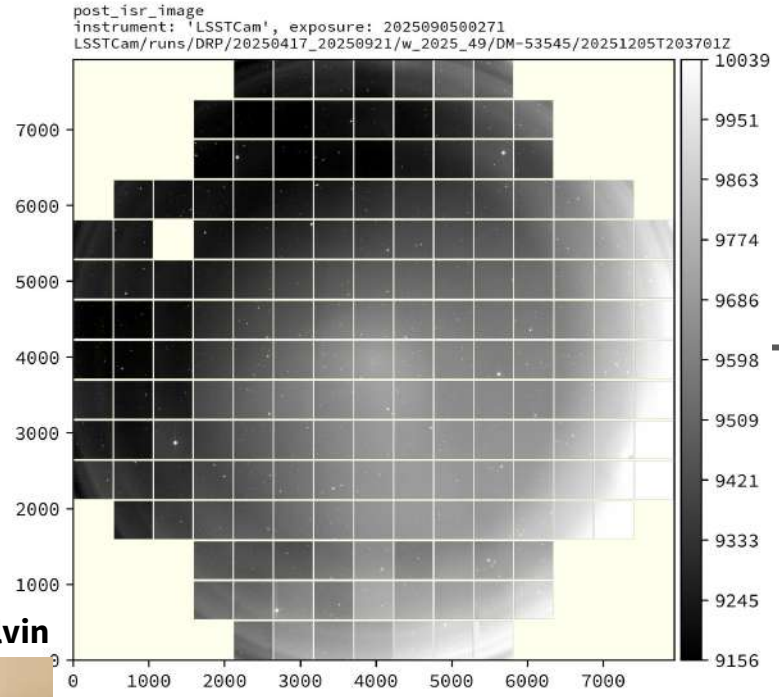


Output: with/without sky correction

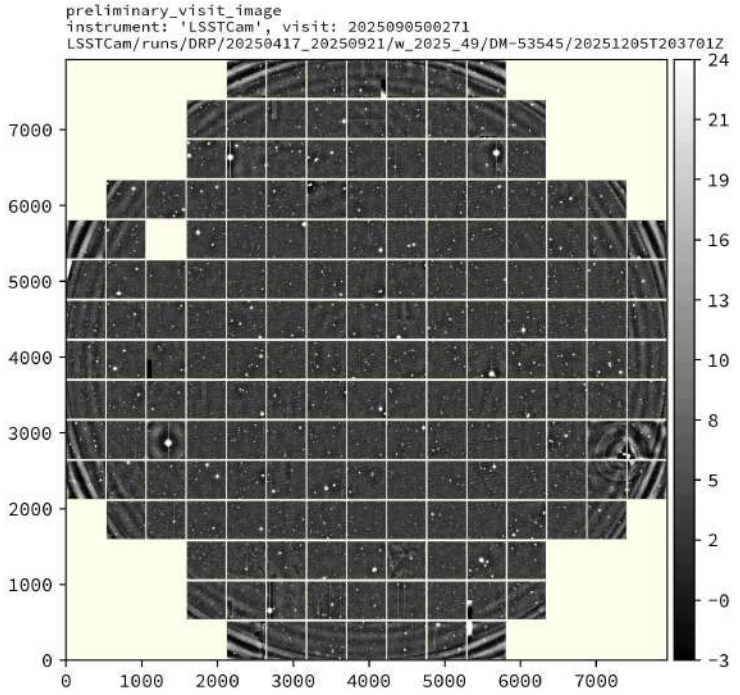
Lee Kelvin



SkyCorrectionTask in LSSTCam (g): … until it doesn't



Input: post-ISR



Output: with/without sky correction

Lee Kelvin

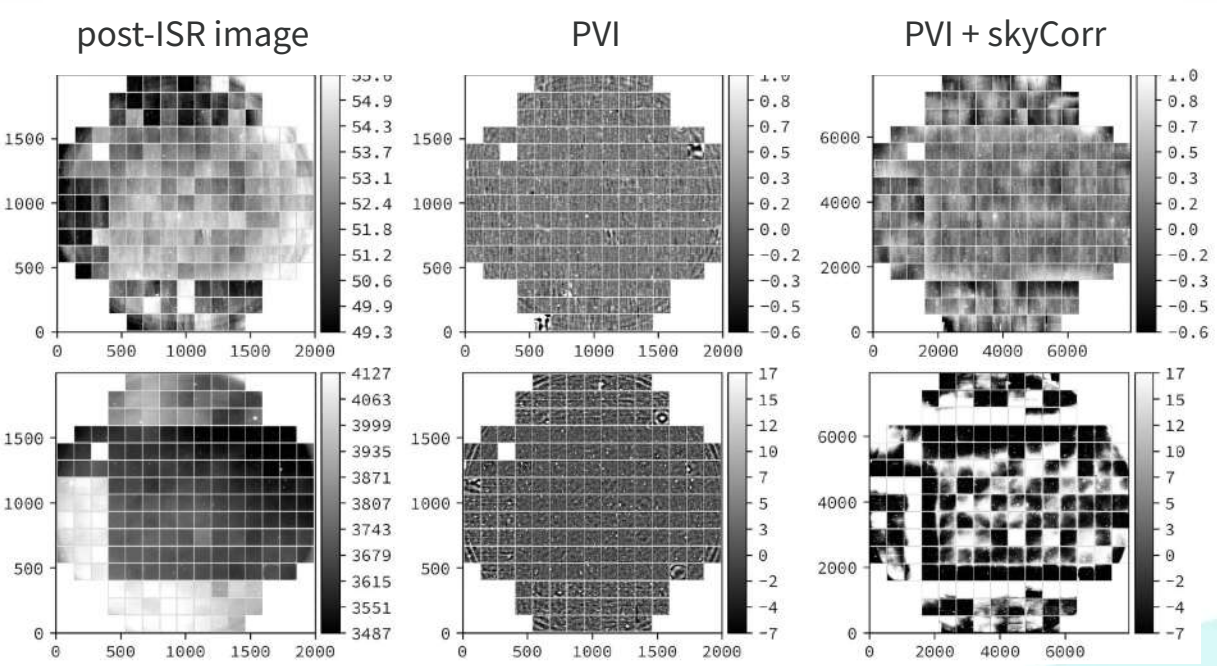


SkyCorrectionTask: worse as step-function offsets grow

SkyCorrectionTask is not designed to accommodate detector-to-detector offsets in the input data (~post-ISR).

Primary sources:

- e2v/ITL
- offset detectors
- offset REBs
- detectors with different chromatic response



Still wanted to provide a large-scale background with the coadds for LSB science

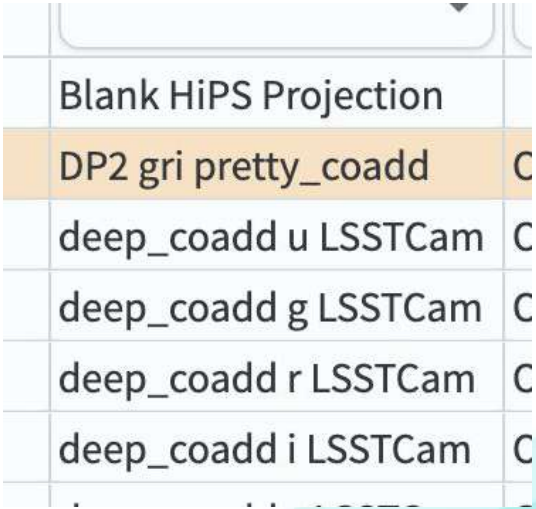
Original plan was to use skyCorr on the visits that go into THE deep_coadd, then subtract off the coadd background before detecting and measuring objects.

But, commissioning skyCorr never reached the robustness needed for the coadd we were trusting our Objects on.

Therefore, after processing DP2, we made a second coadd called `pretty_coadd` for the gri HIPS map.

Looks OK in some regions.

So we present it to you as a second background attached to the `deep_coadd` files for you to assess.



Instrumental Signature Removal (ISR)

Alex Broughton

- DP1 ISR: developed from a detailed physical model of the detector and readout hardware
- DP2 ISR: refined through calibration system improvements
- Three calibration epochs identified during LSSTCam commissioning



Credit: Agnès Ferté (modified by Alex Broughton)

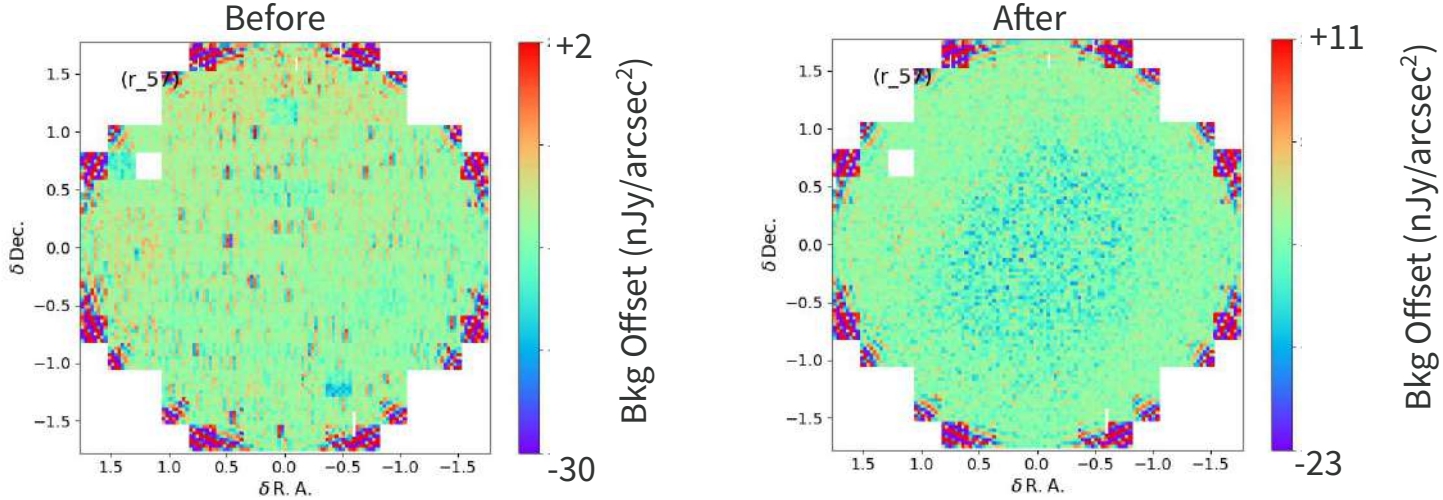
Instrumental Signature Removal Algorithmic Improvements in DP2

Speaker: Alex Broughton

- **Improved PTC** gain, noise, and saturation measurements and linearizer fitting for 3,024 science amplifiers (Alex)
- Updated crosstalk matrices from **in-dome data** (Chris)
- Updated gradient-aware defect finding (Agnes)
- **Entirely new brighter-fatter correction** using Astier+23 & Broughton+24 (Alex)
- Improved mask plane handling for outer PARTLY_VIGNETTED focal plane region (Eli)
- **Entirely new flat-field pipeline** to remove in-dome gradients and make use of all the **individual LEDs** (Eli)
- Masking of E2V edge bleeds (Agnes)
- Significant improvements to CTI measurement and application prior to gain fitting (Alex)
- Improved daily tracking of instrument performance (Chris)
- So many bug fixes …

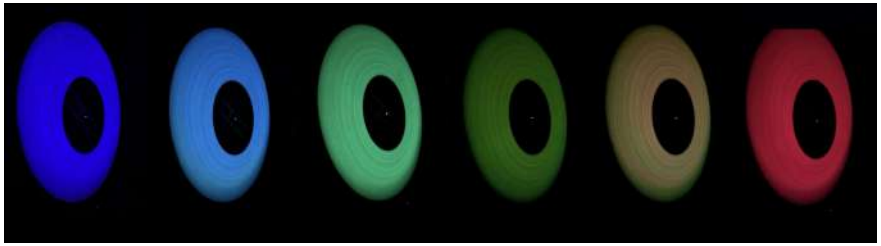
Highlight: Amp Offsets

- Through much of DP2 processing there were amp offsets which flowed through to photometry and PSF residuals
- In the end we decided to turn on the “amp offset correction” code, which isn’t most principled, but it does seem to work on average.

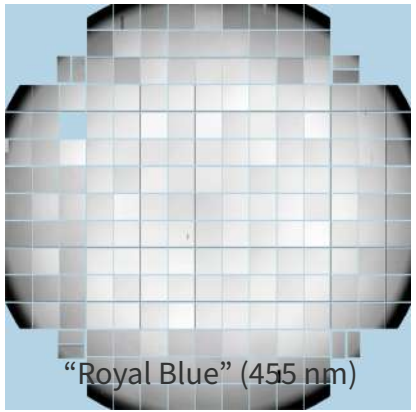


Highlight: Flat Fielding

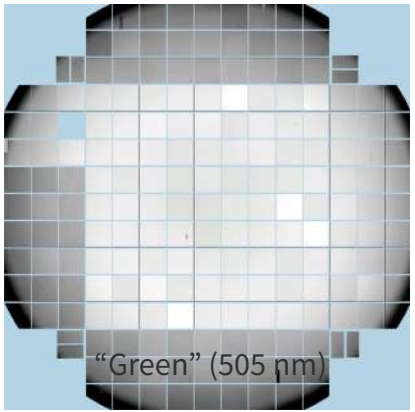
- **Flat-field acquisition improved across all three DP2 calibration epochs**
  - Tighter camera temperature, voltage (detectors 81, 82, 83), and gain control after Epoch 1
  - Fewer focal-plane restarts, which trigger new calibration sets
  - More uniform illumination across the focal plane
- Flat-field screen now illuminated by 10 LEDs (1 u; 2 each in griz; plus new 1050 nm LED shared between y/z)
- NEW combined “anaglyph" flats, with gradient and centering corrections



Collage from Yusra and Kris M



“Royal Blue” (455 nm)



“Green” (505 nm)

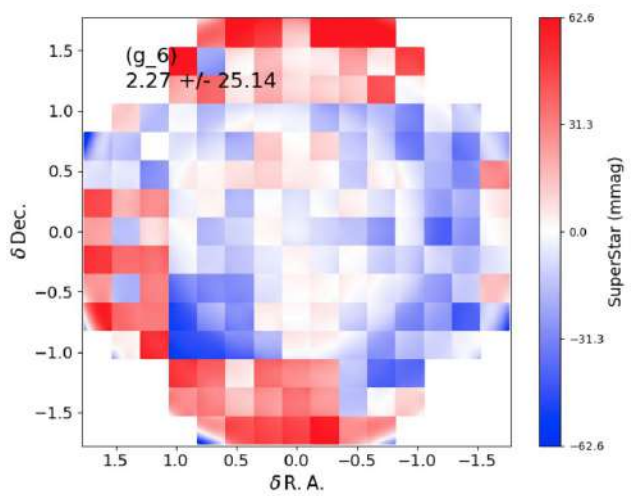
Photometric Calibration: Overview

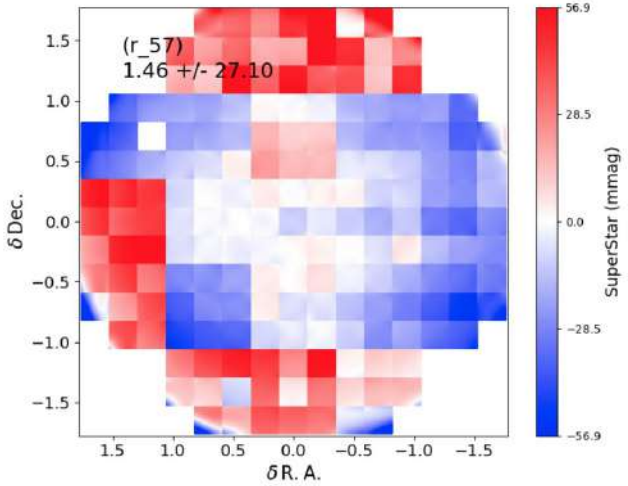
Speaker: Eli Rykoff
5 min

- Photometric calibration (still) based on the Forward Global Calibration Model (FGCM) (Burke, Rykoff, et al. 2018)
- Dense regions from DP2 were down-sampled for speed and not to overwhelm the statistics
- **Input was 29474 visits; 6 filters; 55 million stars with 380 million observations**
- Around 75% of the visits were determined to be “photometric” (variations can be explained with atmosphere/instrument, no gray clouds needed)
- Full band coverage was spotty, and the calibration in the regions with only a couple of bands is less well constrained because of chromaticity (these will not be in coadds)
- NB: Up to 1.2% of individual detectors have incorrect zeropoints due to significantly greater cloud structure in LSST data vs DES/HSC; this creates some offsets in the ragged edge of the DP2 footprint and some in the Galactic plane.

Photometric Calibration: Illumination Corrections

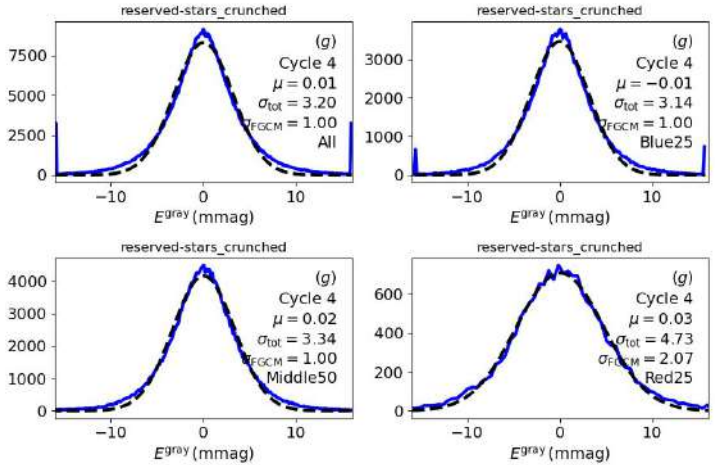
- Illumination corrections (from “background-flattened” to “photometrically-flattened” images) are handled as part of fgcm
  - DR1 will probably be split out separately
- Chromatic differences affect the illumination corrections (is this a problem?)

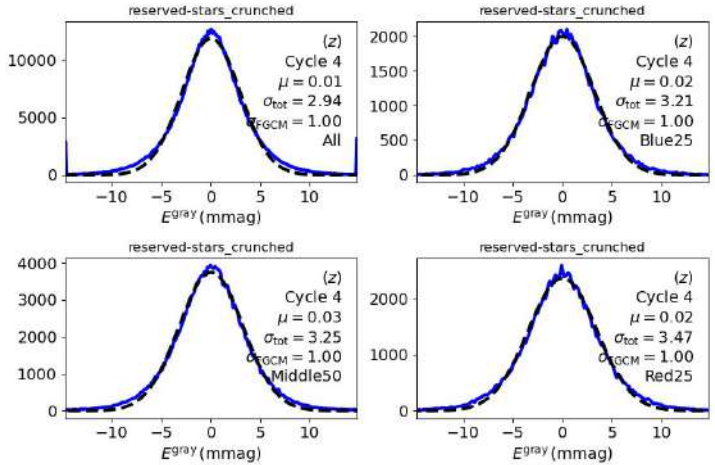




Photometric Calibration: Repeatability

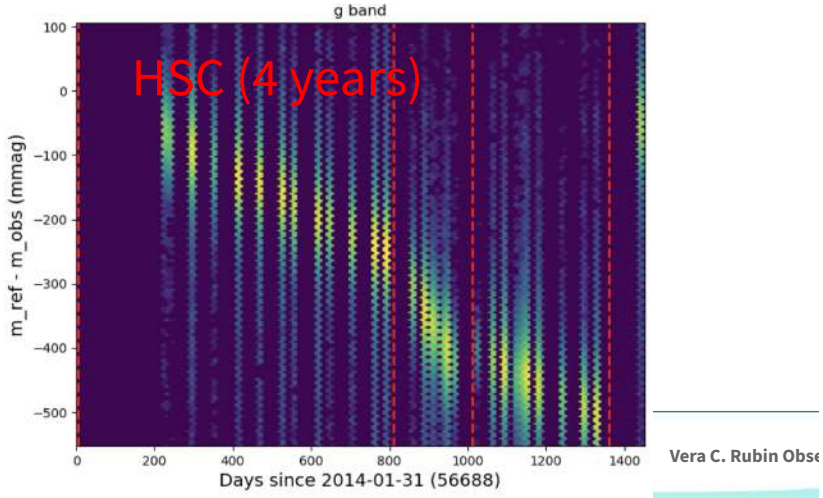
- Measure repeatability for s/n > 200 stars (10% reserved); all stars and 25% blue/50% middle/25% red for those with colors
- In most cases the width of the histogram is consistent with photometric error (intrinsic scatter is <1 mmag); u band is “worse” at 2-3 mmag.

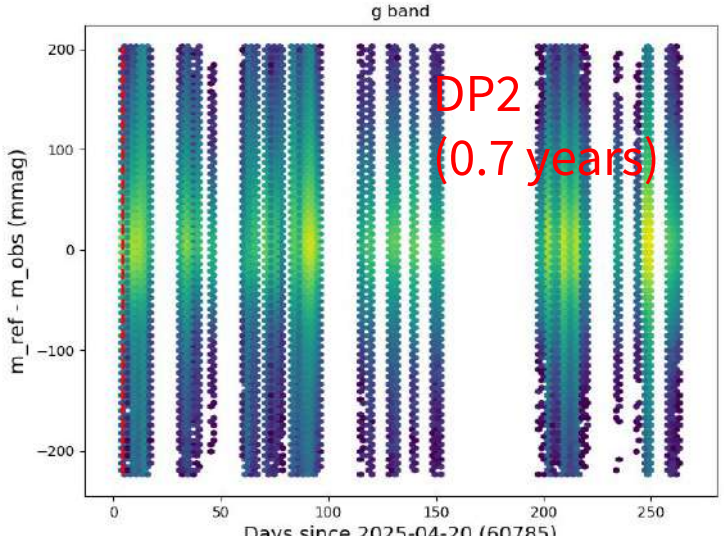




Photometric Calibration: Throughput Loss Over Time

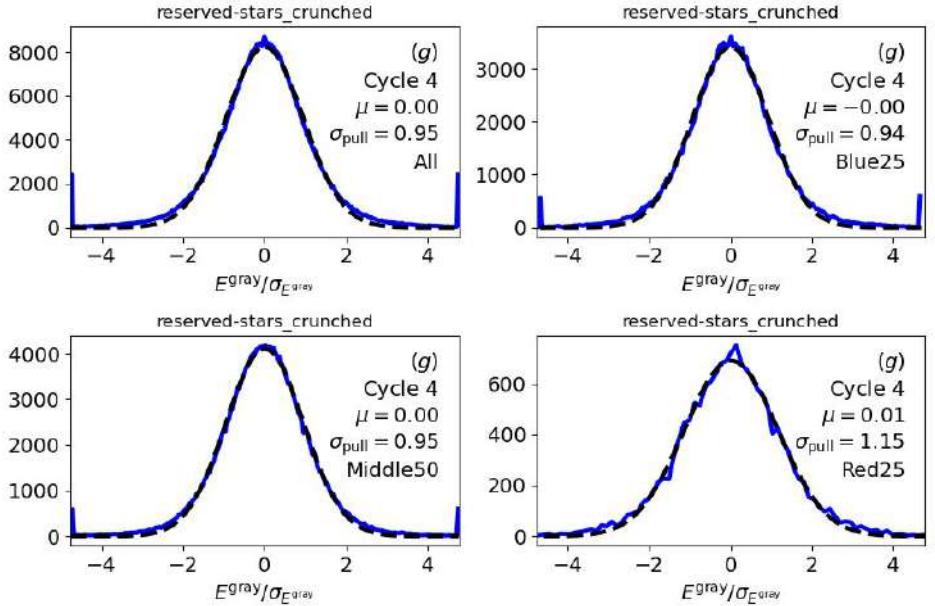
- **Over 250 days of DP2 there is no significant loss of throughput due to dust or mirror surface degradation**
- Looked at both individual stars (instrumental mags vs reference stars) or calibration zeropoints over time
- This is almost certainly due to the **excellent coating work of Tomislav Vucina** and our amazing triple-Ag mirrors





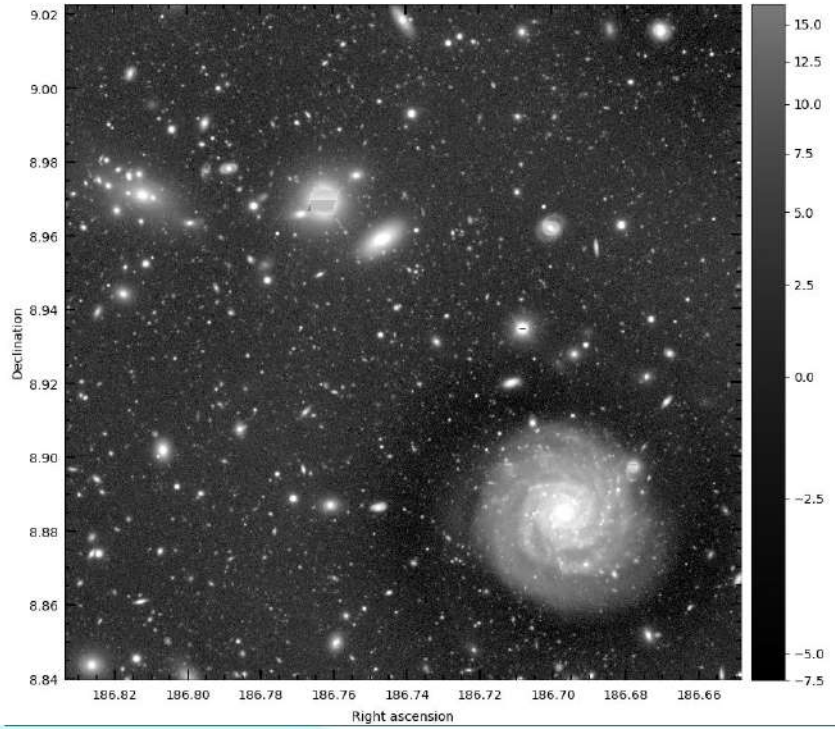
Photometric Calibration: Systematic Error Floor

- After addition of systematic error floor (in quadrature) the “pulls” (delta/sigma) are well behaved (except in u?)
  - After scaling for error, delta/sigma should have a width of 1.0
- For calibration fluxes fgcm adds 2.9 / 1.1 / 3.0 / 1.2 / 1.0 / 1.2 mmag in quadrature to ugrizy mags.
- These are not currently added to the object flux errors; this will be changed for DR1.



Cell-based coadds

tract=10804, patch=99, band="g"

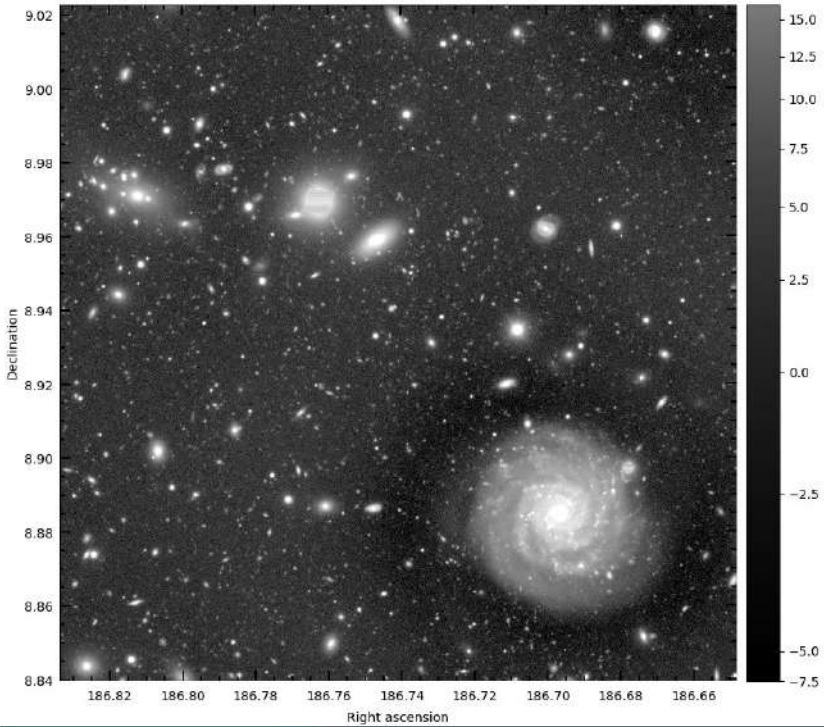


As has long been promised, we are now building cell-based coadds: whether a visit contributes to coadds is now *mostly* decided one 150x150 cell at a time.

- This makes the coadd's properties more consistent within a cell, but still *potentially* discontinuous across cells.
- Coadd PSFs are now much simpler to store and evaluate (one PSF image per cell).

Cell-based coadds

```
tract=10804, patch=99, band="g"
```

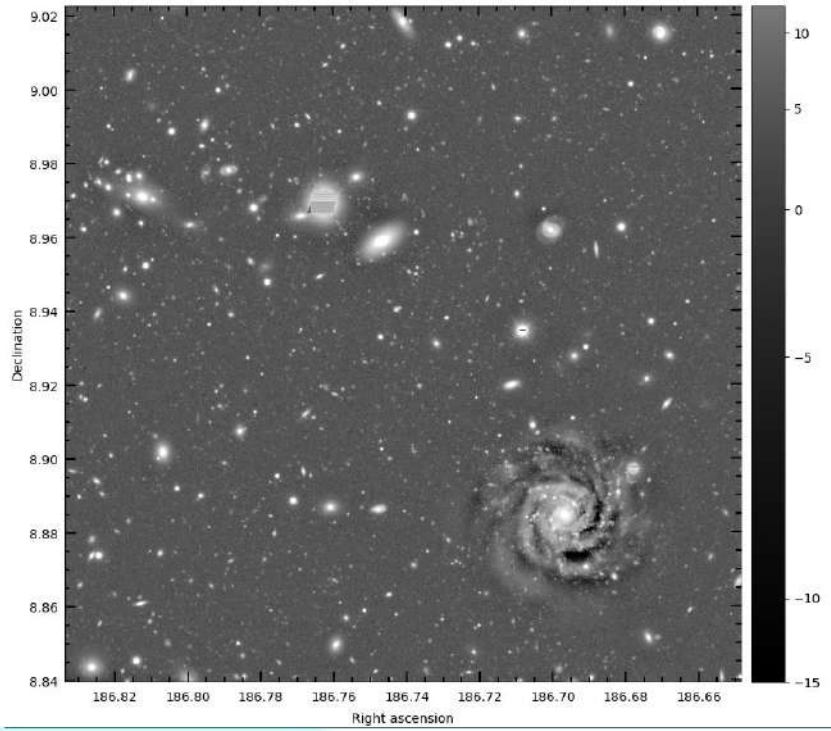


Measurement algorithms aren't fully leveraging this yet - they generally ignore the between-cell discontinuities right now. We don't know if this is net-positive or net-negative yet, but it's definitely better for the many-visits future of the survey.

About that *"mostly"*: if we dropped an entire visit-cell contribution any time any pixel in that visit-cell was masked, what would happen?

Coadd Backgrounds

```
tract=10804, patch=99, band="g"
```



The previous images were actually what we call "predetection" deep coadds, built directly from background-subtracted visit images.

Between detecting Objects and deblending them, we do another round of small-scale background subtraction. *This* is a deep_coadd:

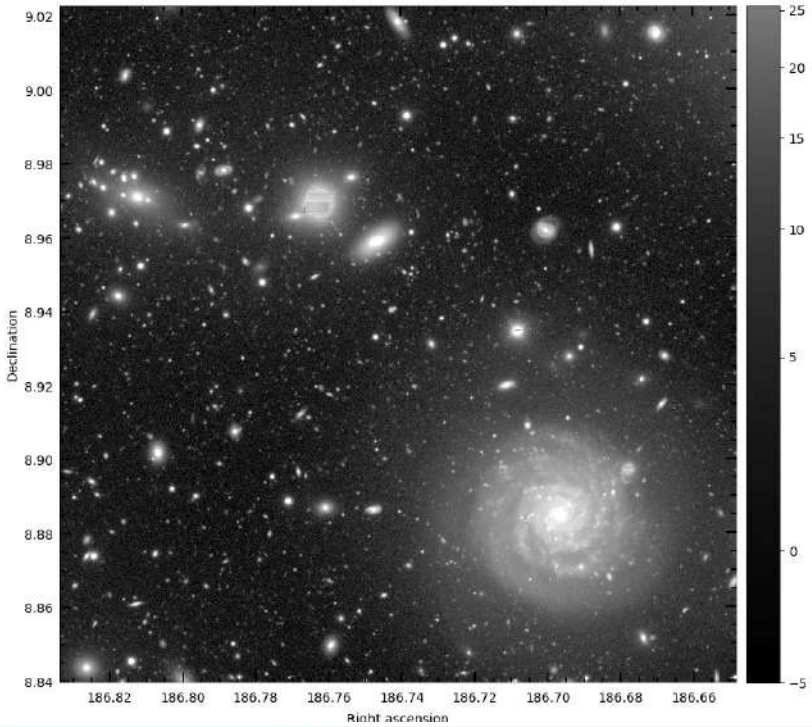
```
coadd = butler.get("deep_coadd", data_id)
```

You can recover the original with:

```
c.apply_background(None)
```

Coadd Backgrounds

`tract=10804, patch=99, band="g"`



Our "pretty" coadds are completely distinct: we built those using the SkyCorrection background subtraction only.

We didn't retain those coadds directly, but we did fit a spline to the difference between them and deep_coadd, so you can mostly recover them:

```
coadd.apply_background("pretty")
```

This is *very* experimental; trust this background only to the extent you can validate it independently.

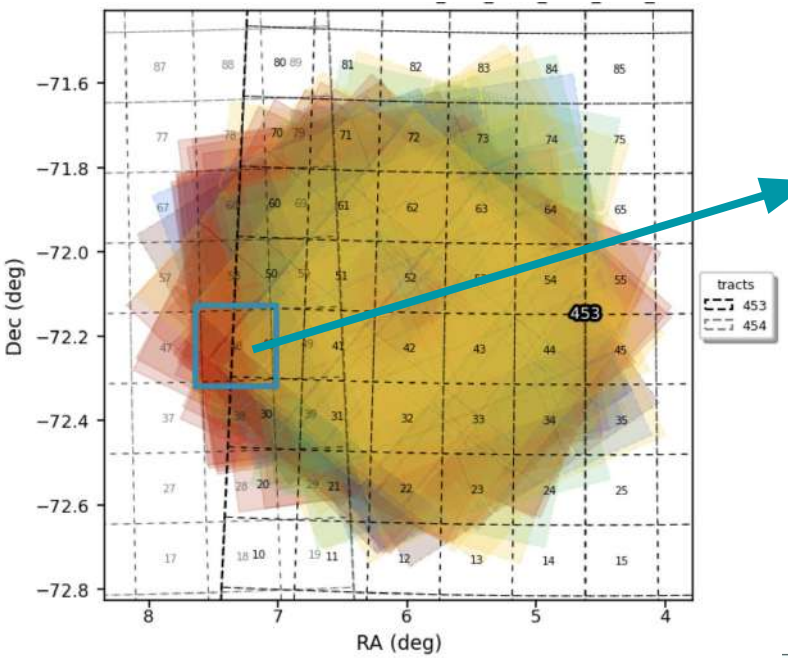
Coadd Deblending in DP1
Problem: Crowded fields too dense to deblend in DP1

Red = a single large blend that had to be skipped

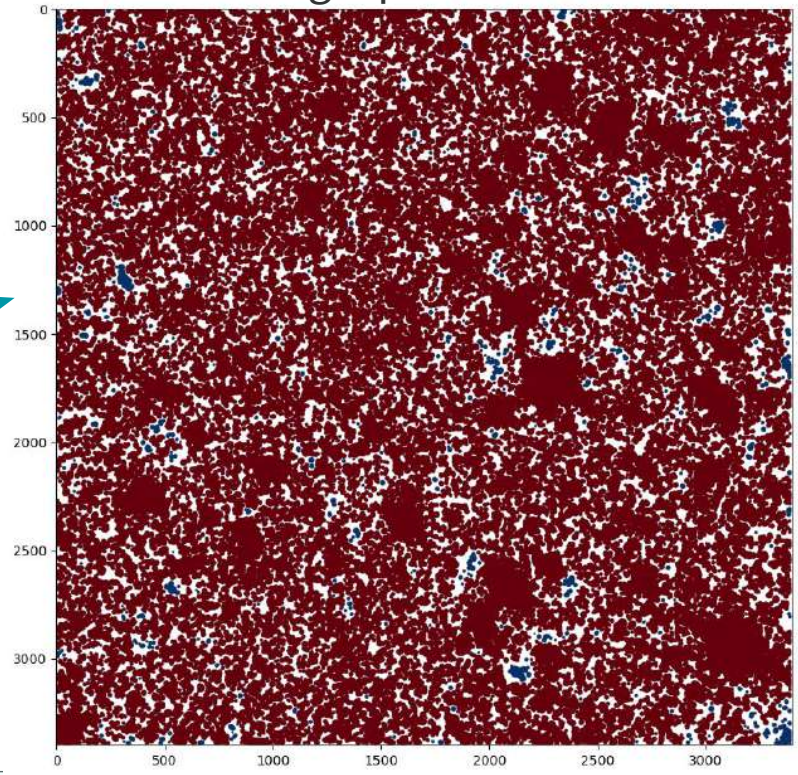
Blue = smaller blends that were deblended and measured



Fred Moolekamp

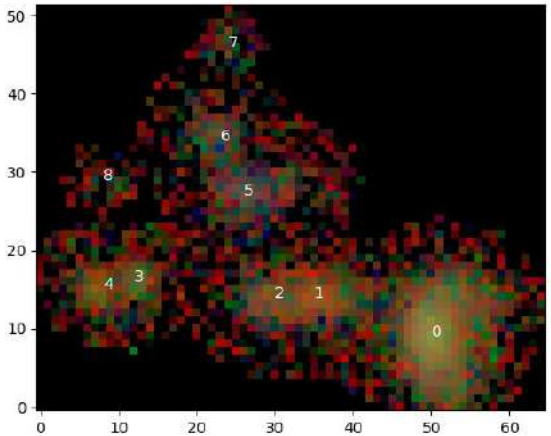


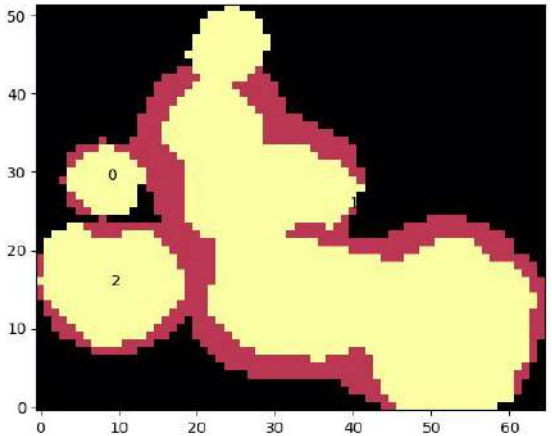
A single patch in 47 Tuc



Coadd Deblending in DP2
Solution: Deconvolution before deblending

- Partially deconvolve to a narrow, well-sampled Gaussian
- Allows us to subdivide a large percentage of blends
- Improves runtime by almost 10x
- Improves χ^2 (now published in the Object catalog)

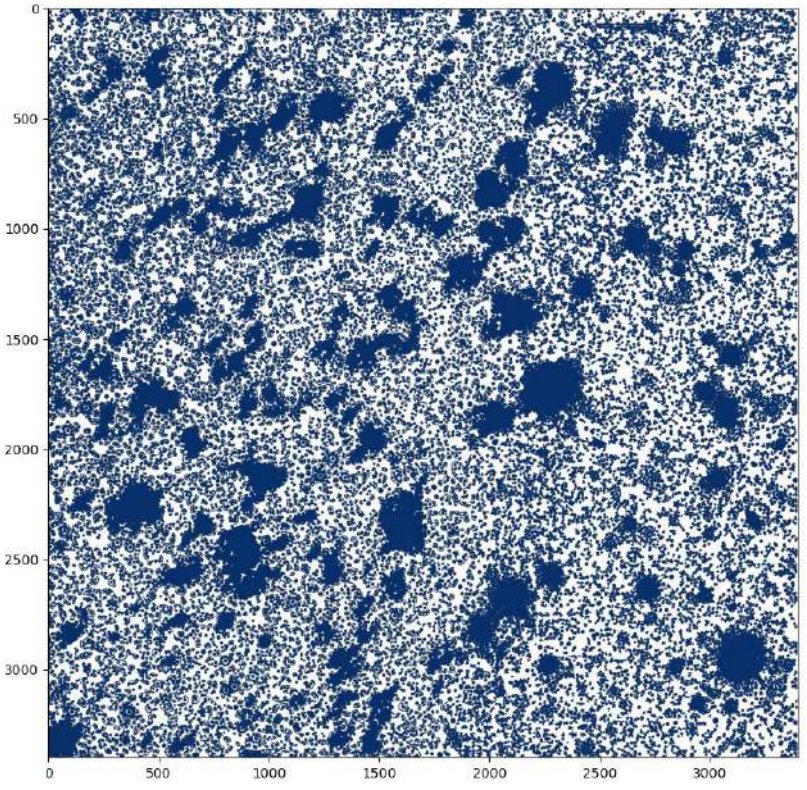




Coadd Deblending in DP2

- Improved our coverage in crowded fields
- Creates a hierarchical **object_scarlet_models** data product
- Also includes spans for isolated sources

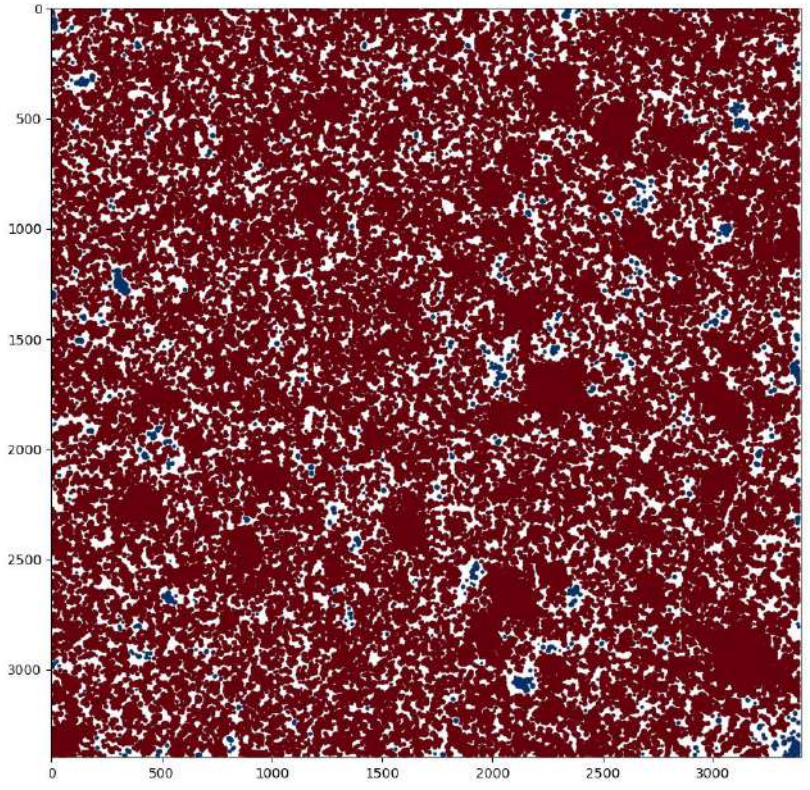
The same single patch after



Coadd Deblending in DP2

- Improved our coverage in crowded fields
- Creates a hierarchical **object_scarlet_models** data product
- Also includes spans for isolated sources

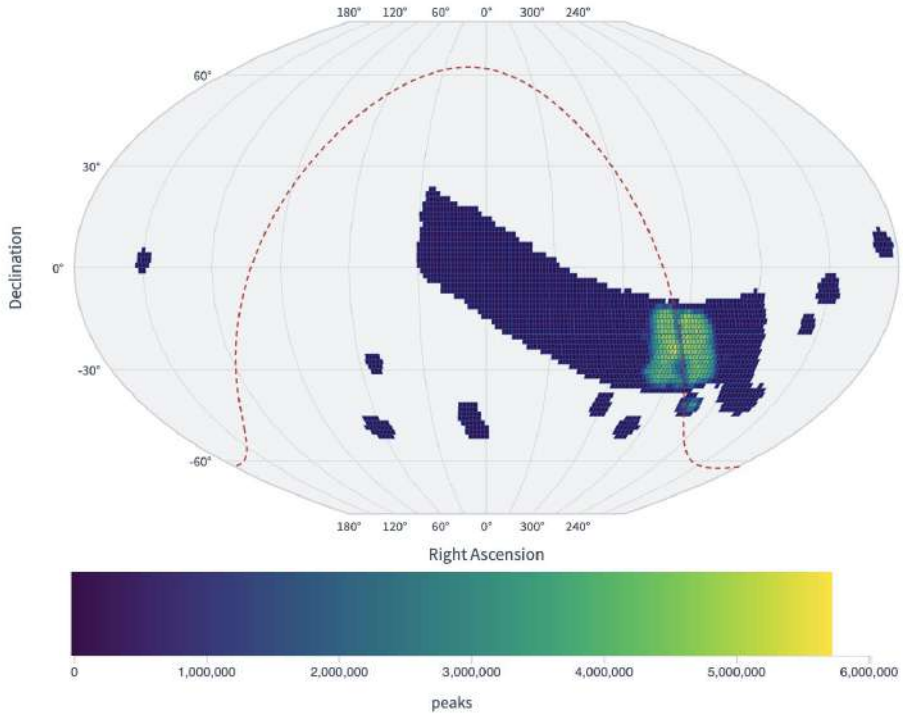
The same single patch before



Coadd Deblending in DP2

- Not perfect: small PSF sizes and poorly constrained wings in crowded fields cause rings, that connect many blends that should be subdivided

Skipped peaks in DP2



Deblending in DR1

- We have several ideas about how to process large blends/crowded fields
  - Subtract stars
  - Hierarchical deblending (subtract brightest first)
  - Implement a more robust scheme to subdivide with padding
- Test scarlet lite in visit deblending

Coadd Measurement

- We're now running a multiband exponential fit with MultiProFit and using it to initialize the Sersic model fit, reducing objects with overestimated sizes/fluxes.
- Errors on exponential/Sersic ellipse parameters are now available; axis ratio TBD.
- New model_extendedness classifier using exponential size & Sersic-to-PSF flux ratio.

More information in the (conflicting) Galaxies SC session; check out the recording.

the Metadetection ShearObject table

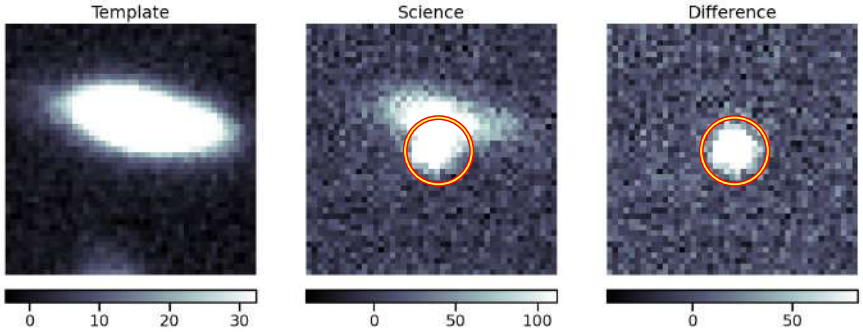
As for DP2, we're running a state-of-the-art self-calibrating weak-lensing shear estimation code, *metadetection* (Sheldon, Becker, Jarvis, & Armstrong, 2023), on our coadds. This produces the *ShearObject* table.

- It's rare to get shear working without biases on a new dataset on the first try.
  - We probably didn't, for reasons you all will discover.
  - Hopefully we're close enough for DP2 science.
- Do *not* match ShearObject to Object or anything else.
  - It contains independent detections to allow detection biases can be calibrated out as well, and matching wrecks that.
  - ShearObject has its own photometry measurements you can use instead.

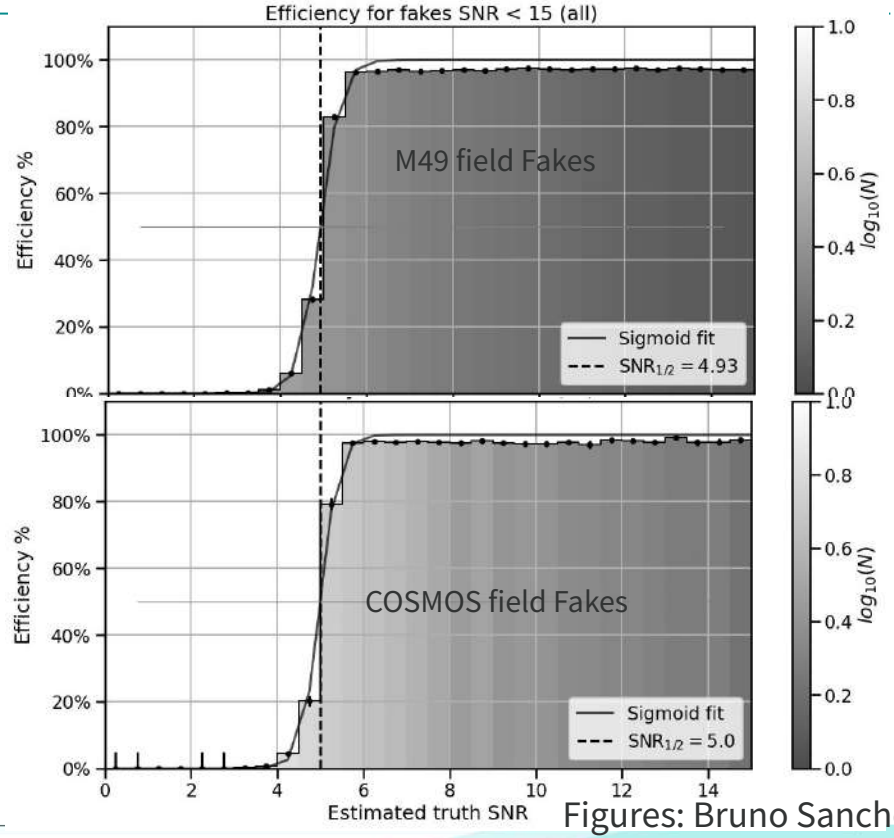
Difference imaging performance

Speaker: Ian Sullivan
5 min

Fake injection campaigns in M49 and COSMOS fields for 20260224 (first alerts)



98% completeness above SNR>6

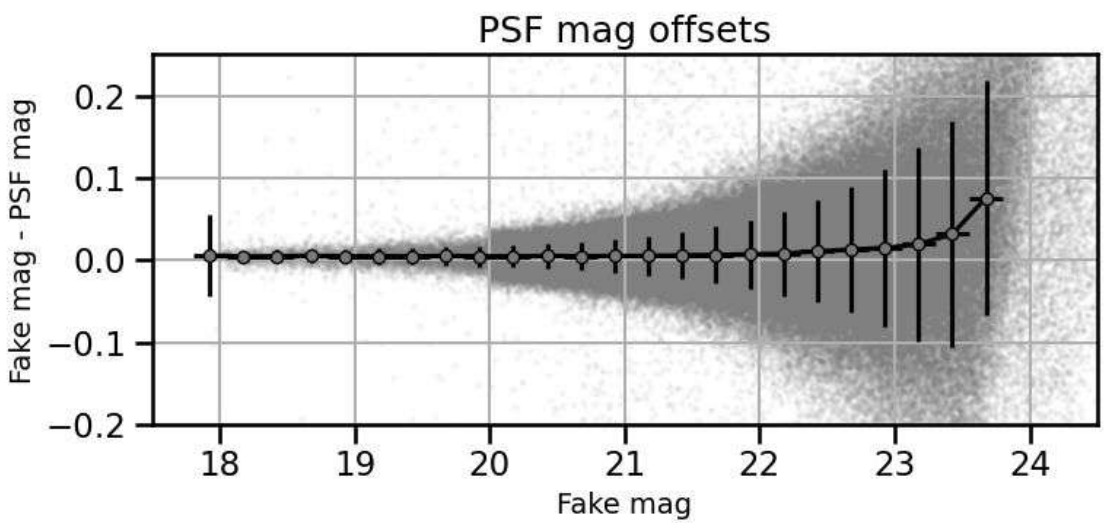


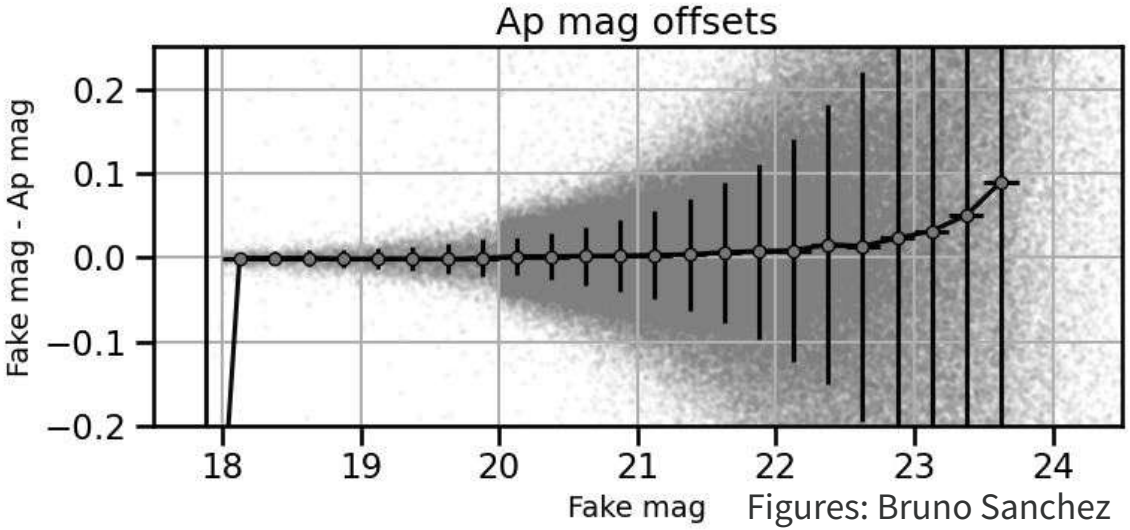
Figures: Bruno Sanchez

Difference imaging performance

Day_obs 20260224

- Over 400,000 synthetic sources injected into 1970 COSMOS images from first alerts data (2/24/2026)
- Injections into science and template as hosted and host-less transients, and variable stars
-



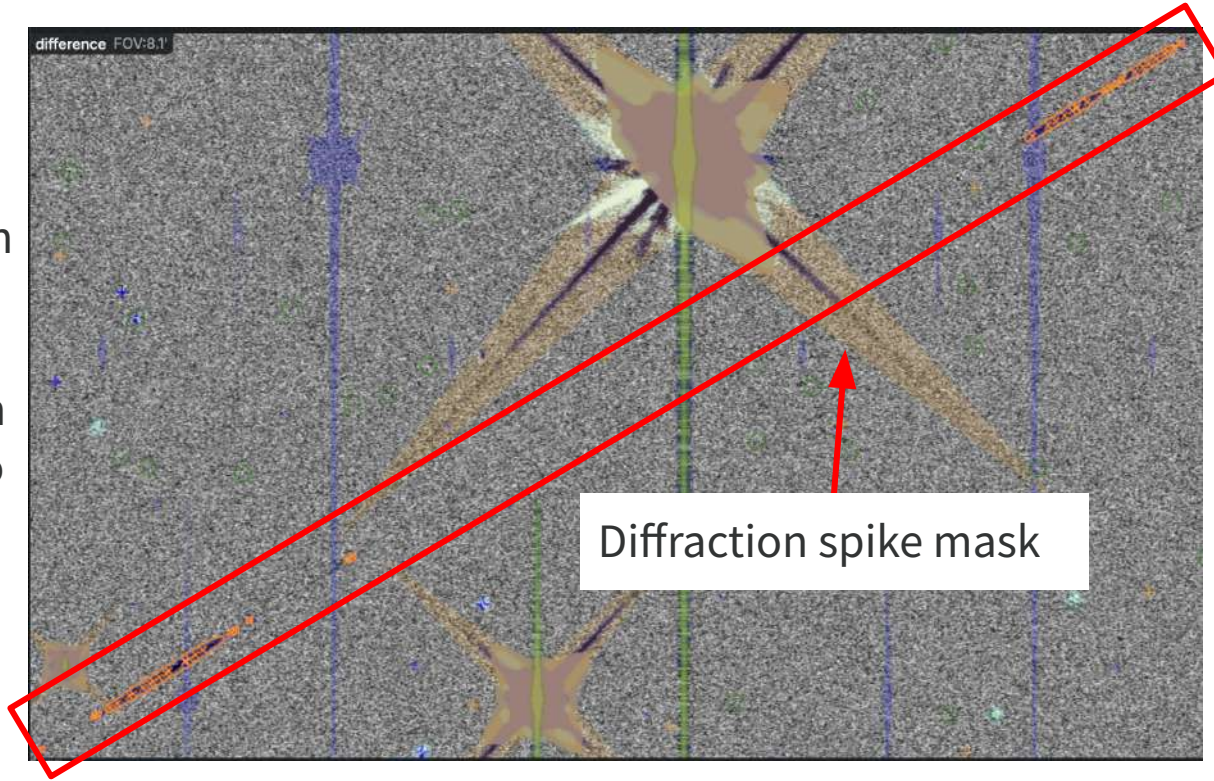


Figures: Bruno Sanchez

Image Differencing: Changes in DP2 since DP1: artifact rejection

Speaker: Ian Sullivan
5 min

- Flagging glint trails
  - Uses the solar system code to identify and flag linear features with **`glint_trail`**
- Cosmic rays
  - Re-run CR detection on the difference image to catch CRs that were missed in single frame due to blending with static sources
- Diffraction spike mask



Diffraction spike mask removes saturated star cores

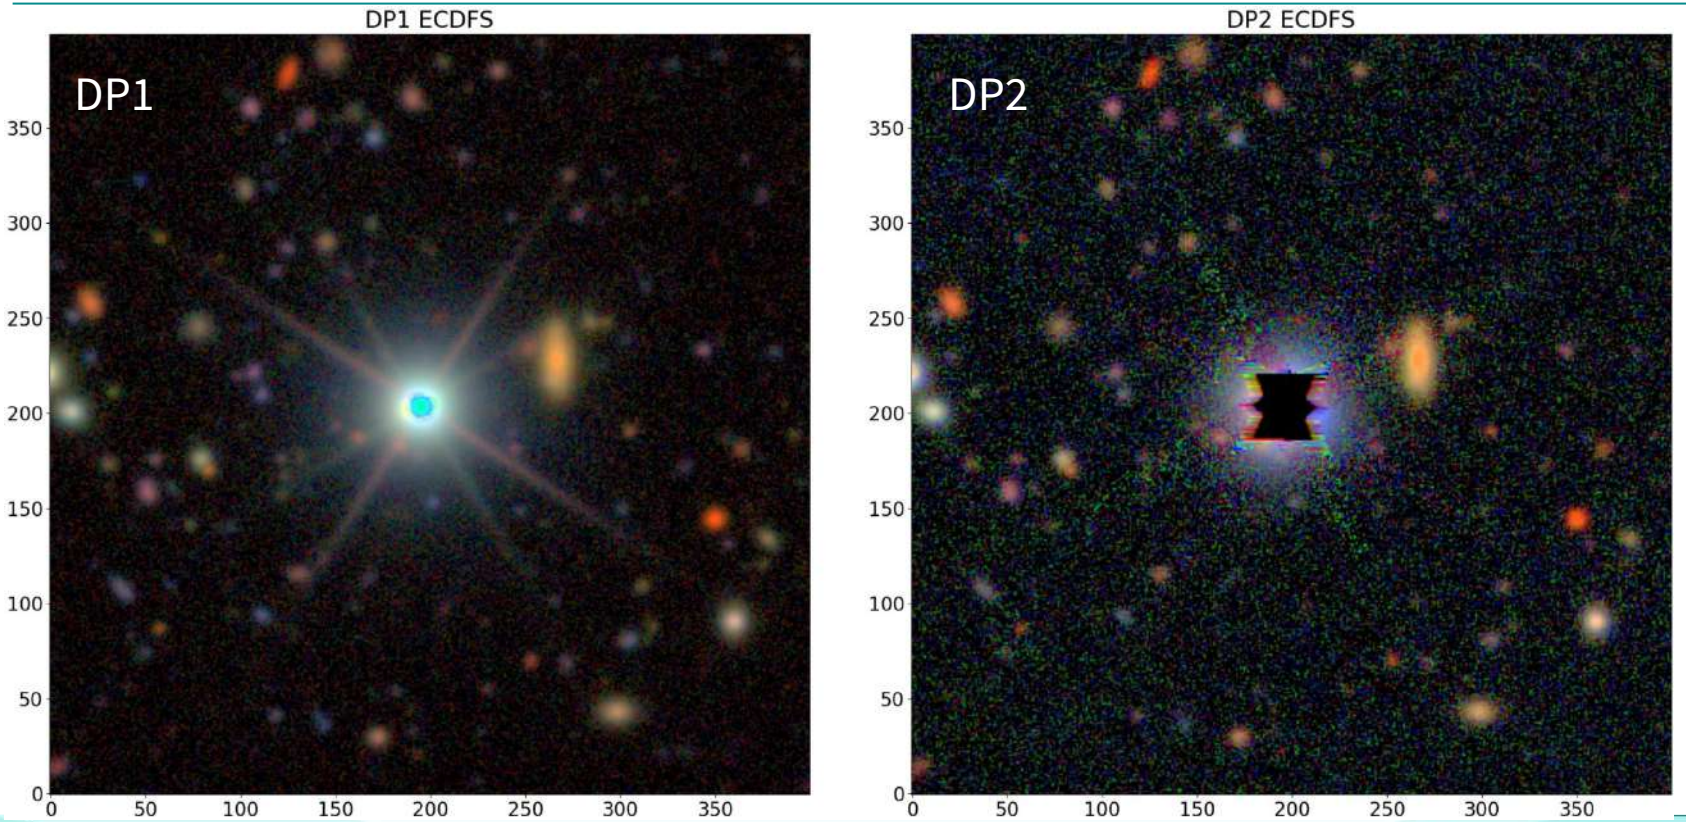
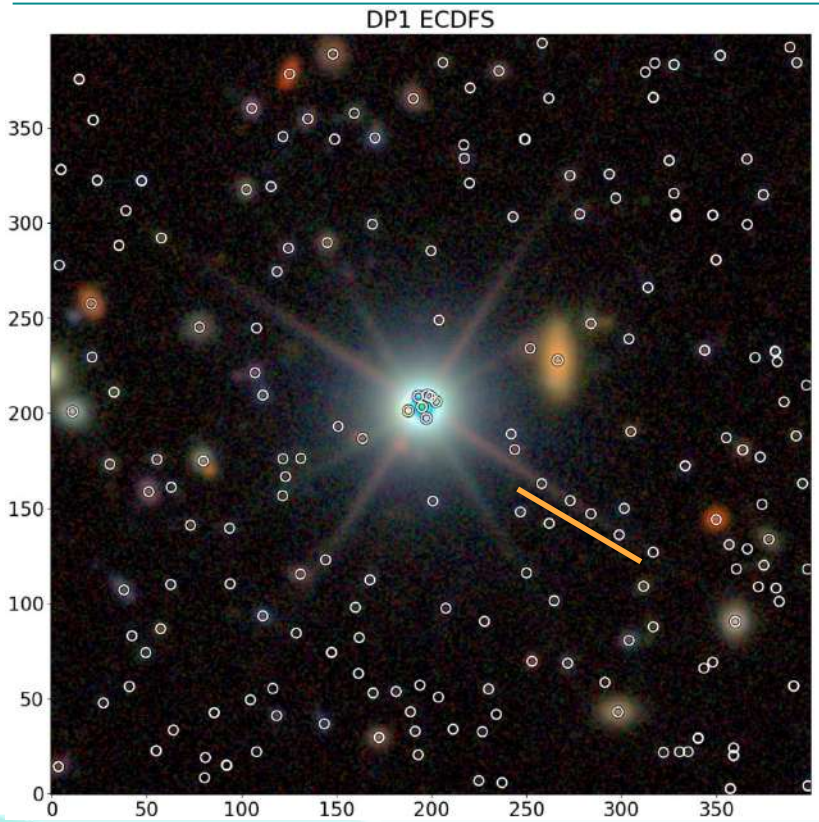


Image: Dan Taranu

Fewer false positives near spikes; some outside mask



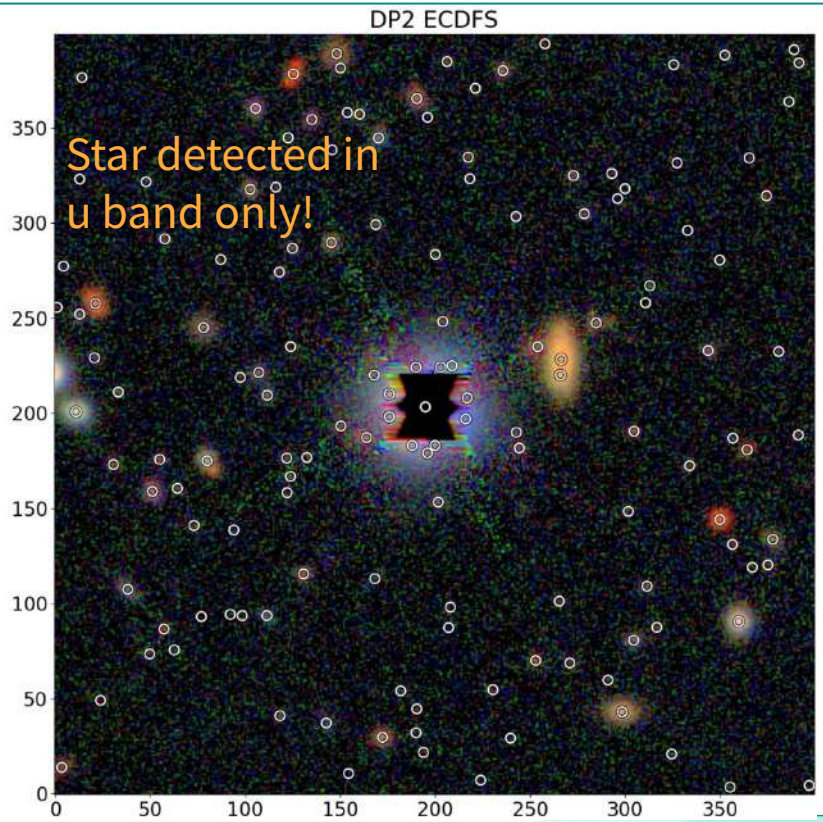


Image: Dan Taranu

Pretty coadds don’t mask diffraction spikes
HIPS maps also fill in saturated cores

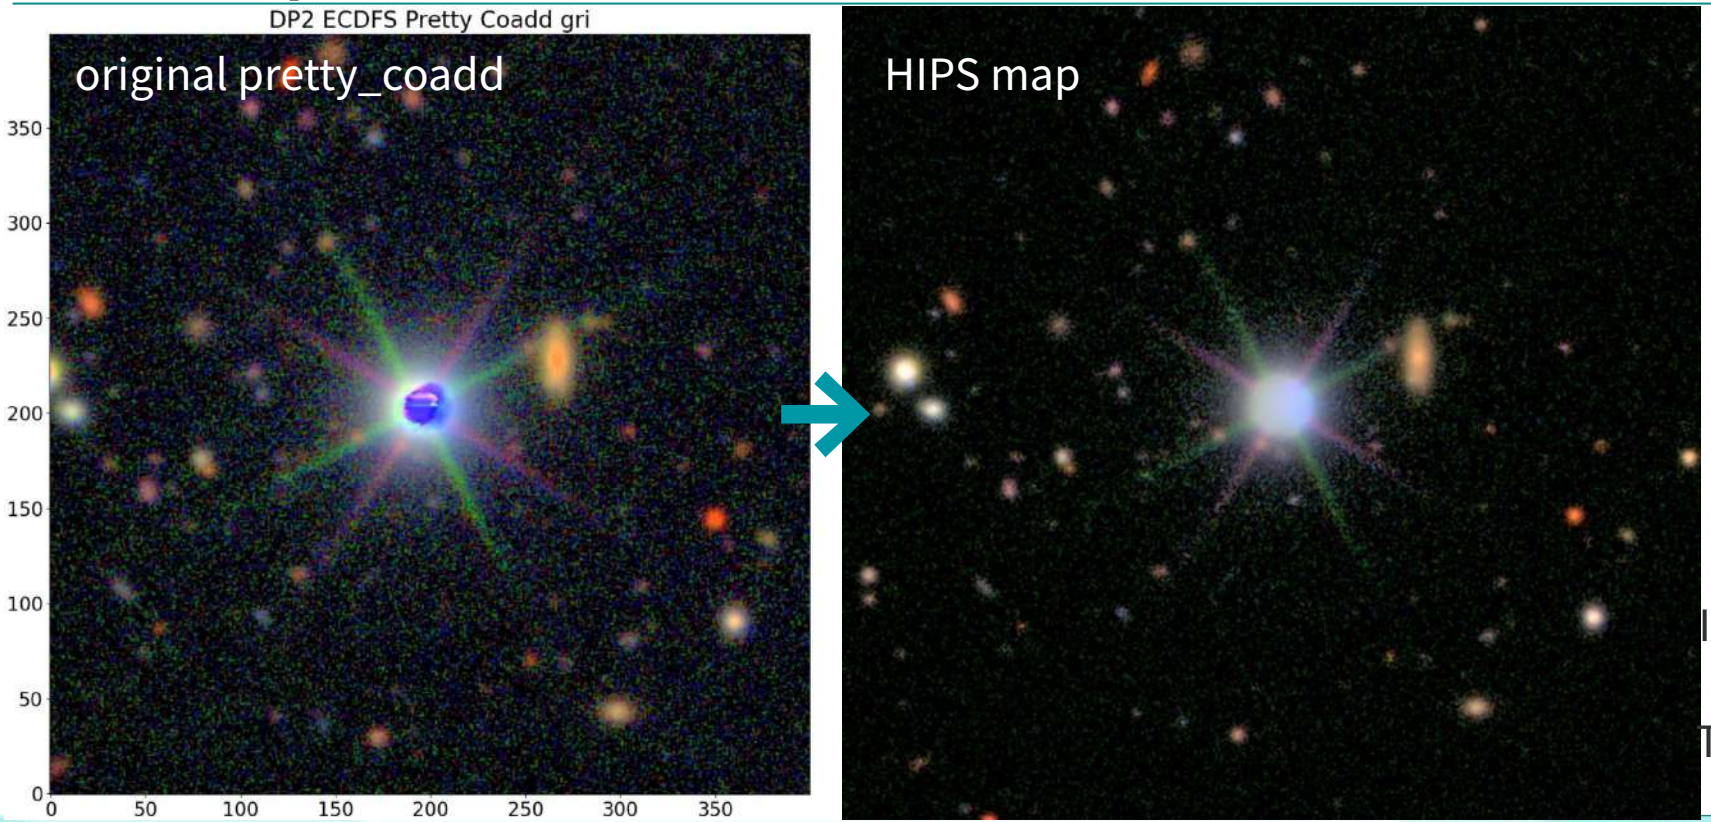


Image: Dan Taranu

Image Differencing: Changes in DP2 since DP1

Speaker: Ian Sullivan
5 min

- Image difference background subtraction
  - Additional round of background subtraction run before detection on the difference image suppresses artifacts around bright sources

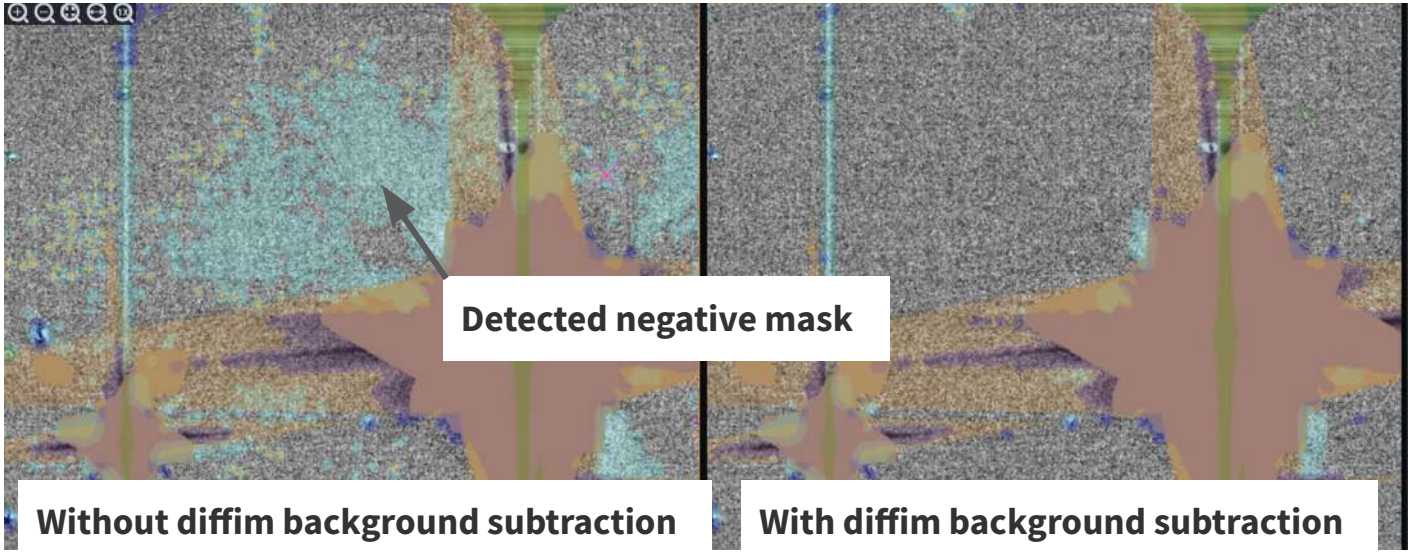
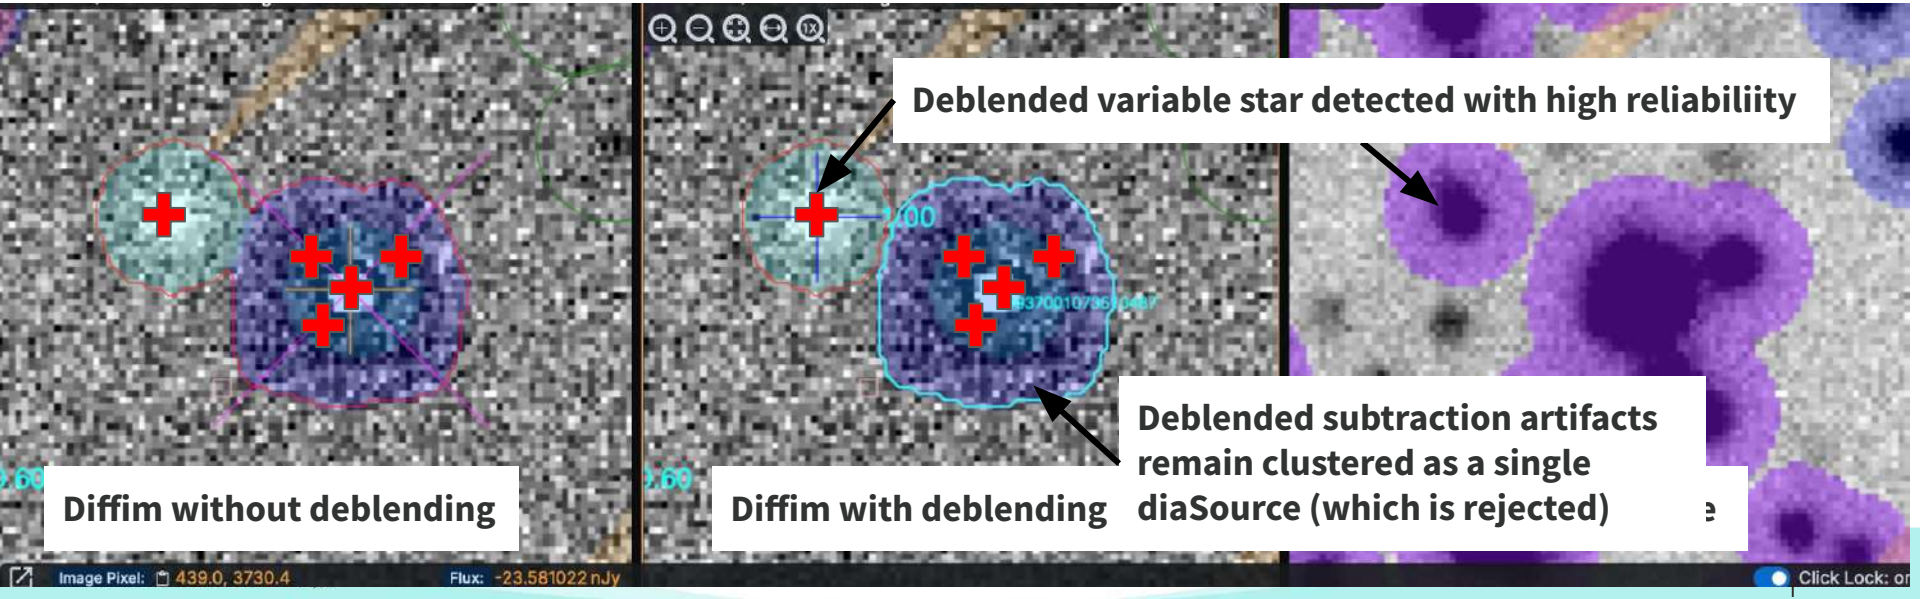


Image Differencing: Coming soon to the alert stream

Speaker: Ian Sullivan
5 min

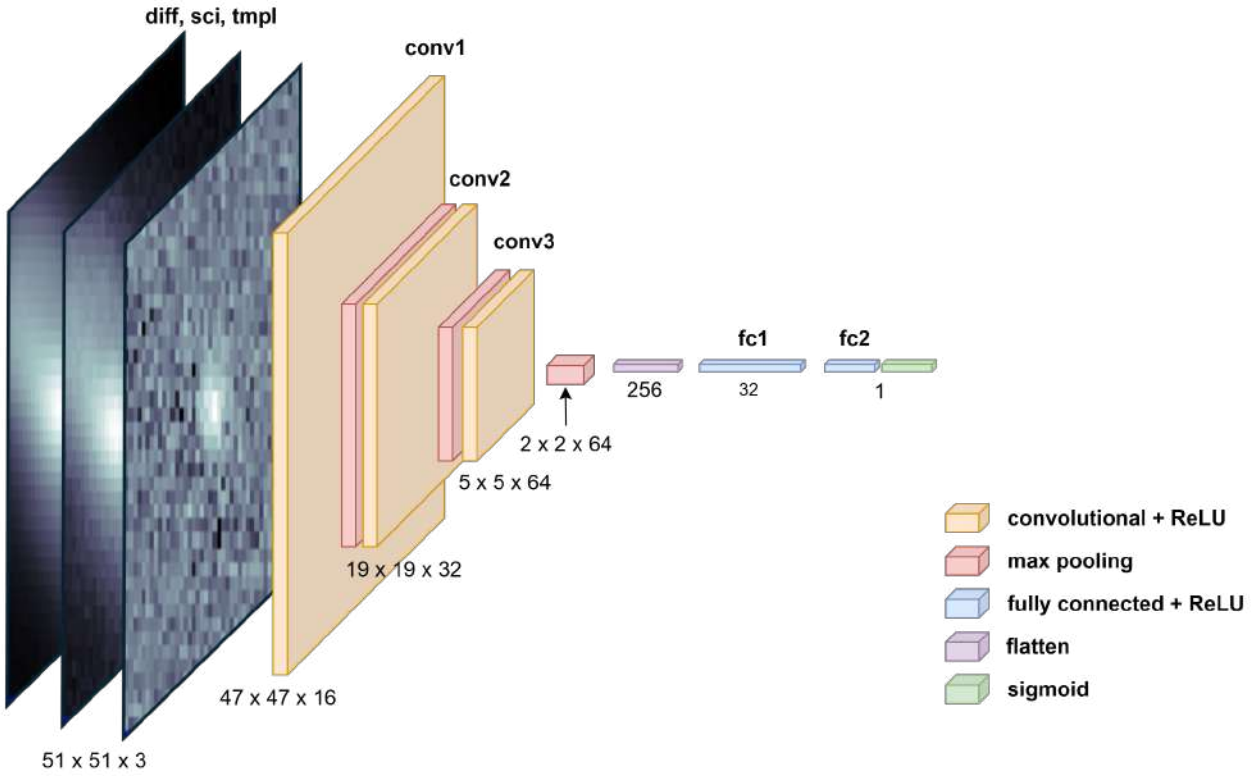
Deblending difference image footprints

- DiaSources with overlapping footprints are currently measured as a single source
- The new image difference deblender divides these footprints and allows measurement of these nearby diaSources



Real/Bogus: Reliability Score

Speaker: Tatiana Acero-Cuellar



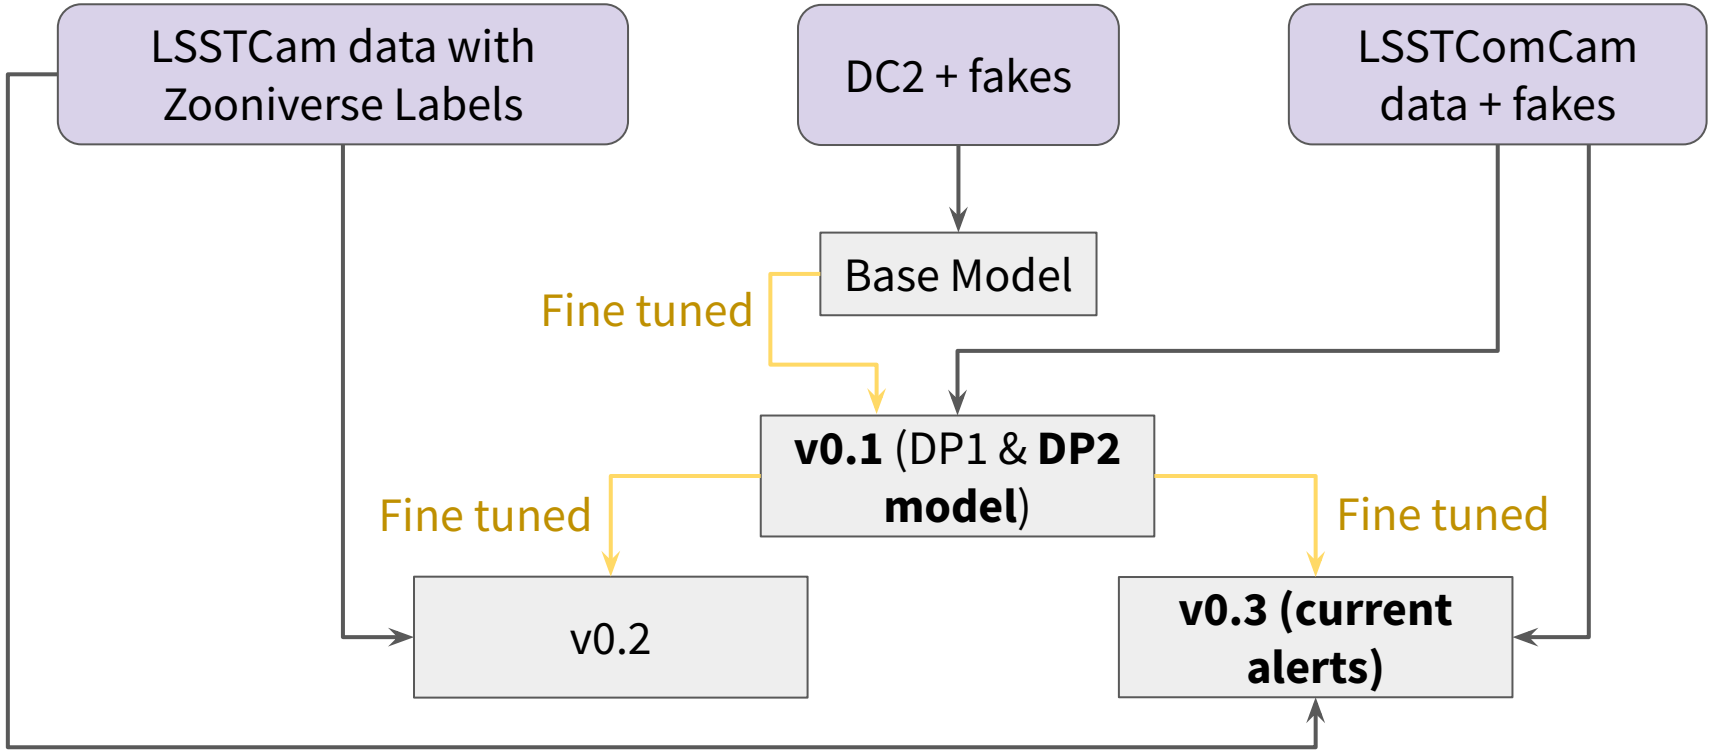
Real/Bogus: Reliability Score

Model version	AP deployment date	AP reliability cutoff	DRP usage
0.1	September 2025	0.1	DP1 & DP2 model
0.2	18 February 2026	0.1 (before Feb 24, 2026) 0.5 (after Feb 24, 2026)	
0.3	9 July 2026	0.5	

Starting with v0.3, the model version is now visible: “reliabilityVersion” field in DIASource table.

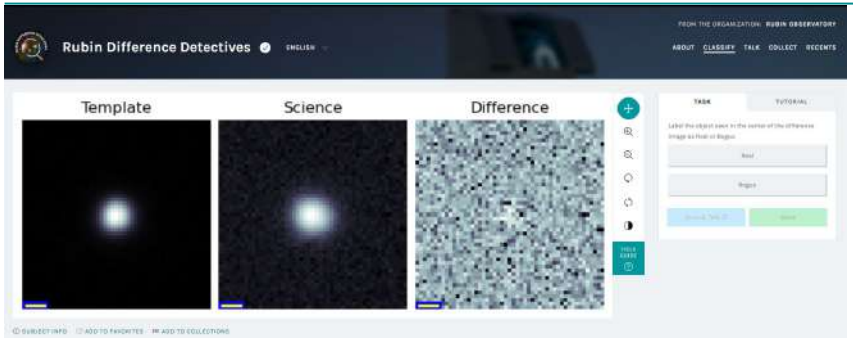
Real/Bogus model versioning and data usage

Speaker: Tatiana Acero-Cuellar



Rubin Difference Detectives - Zooniverse Project

Speaker: Tatiana Acero-Cuellar



2,075 Volunteers	1,339,732 Classifications
282,196 Subjects	230,236 Completed subjects

Transients

Solar system

Variables

TNS catalog match	4,054
Solar System catalog match	88,058
Gaia catalog match	89,410
0.5< reliability v0.1 model < 0.9	49,993
Sources w extendedness = 1	10,932

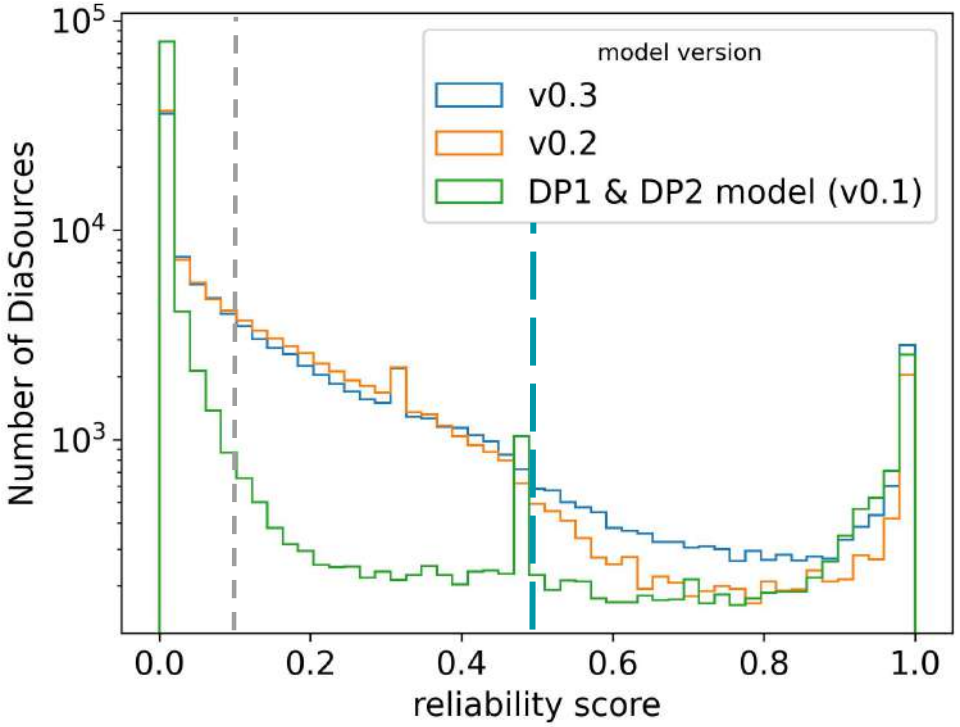
LSSTCam data with Zooniverse Labels

Reliability cutoff for Alerts

Speaker: Tatiana Acero-Cuellar

Cutoff applied to mitigate the impact on APDB when generating alerts.

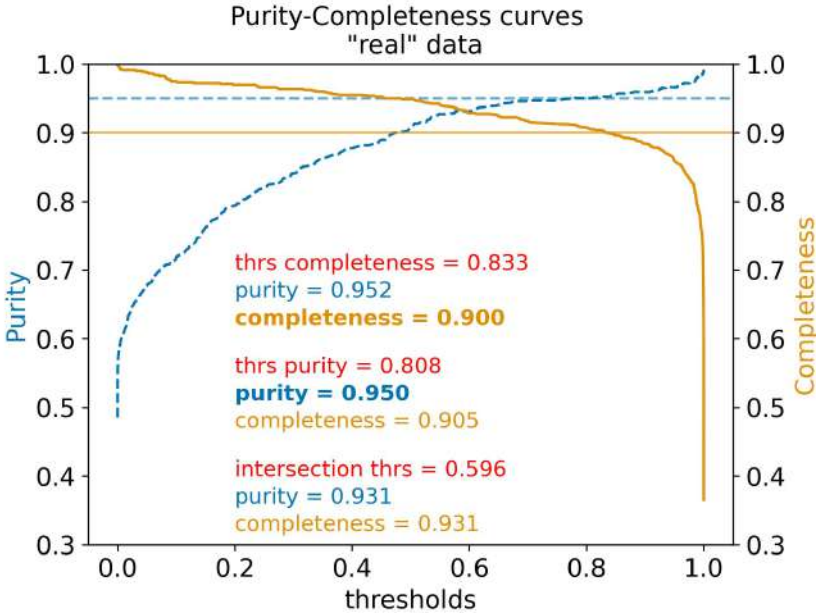
Model version	AP reliability cutoff
0.1	0.1
0.2	0.1 (before Feb 24, 2026) 0.5 (after Feb 24, 2026)
0.3	0.5



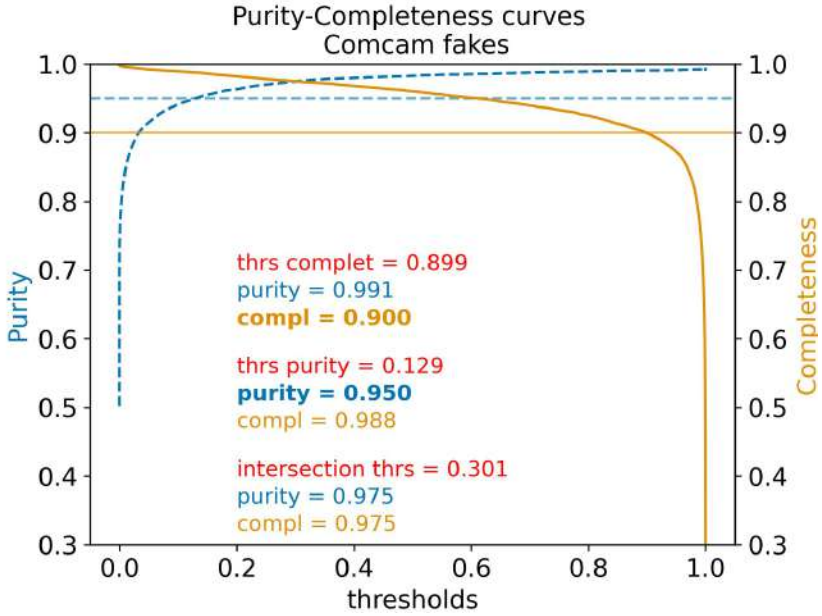
Speaker: Tatiana Acero-Cuellar

v0.3 model performance

‘real’ data with zooniverse labels

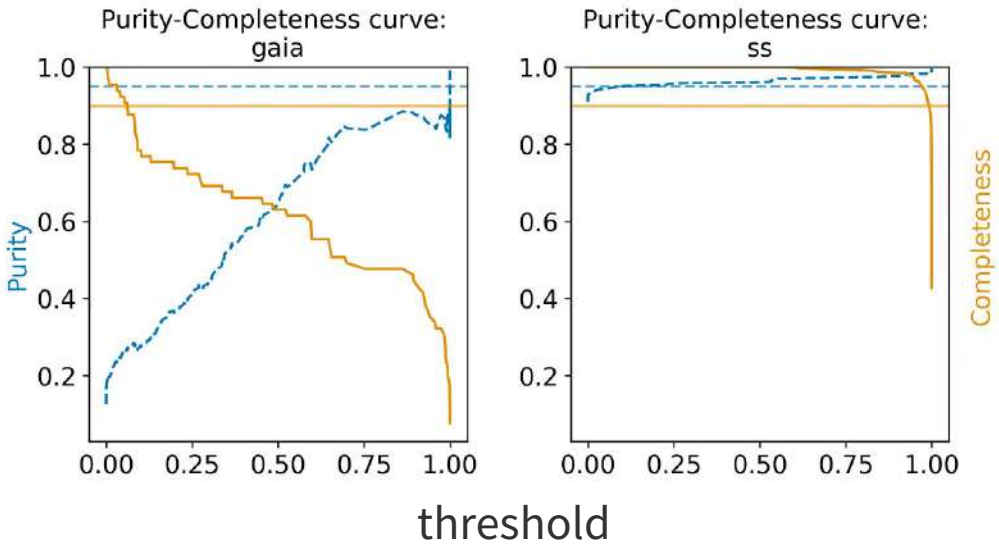


LSSTComcam Fakes



Note that performance is calculated on a test set and may not be fully representative of on-sky performance.

Current/Future work



The performance for variables is better with v0.3 but still room for improvement (see gaia fig).

Rubin Difference Detectives is on pause at the moment while we decide how to best proceed. Now we are collecting labels from recent data with the collaboration of group of students at UPenn.

Use LSSTCam fakes.

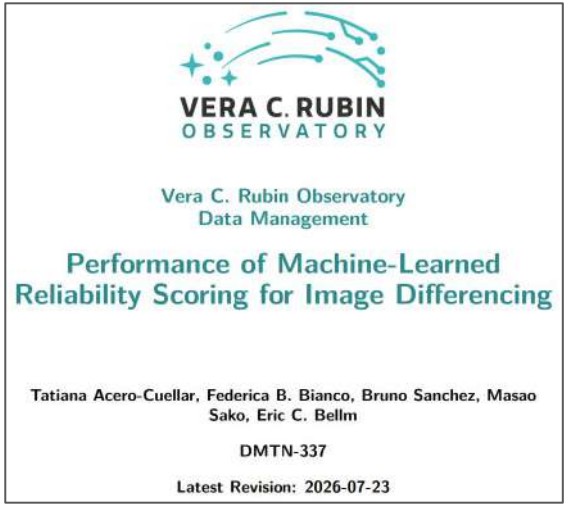
Real/Bogus: Reliability Score Documentation

Speaker: Tatiana Acero-Cuellar

More information about the most reliability score for alerts: https://prompt-products.lsst.io/processing/filtering/index.html#filtering-and-reliability-scoring

More detail information in DMTN-337





Solar System - Discoveries

Speaker: Jake Kurlander
4 min

11,000 discoveries pre-DP2 (no diffims)

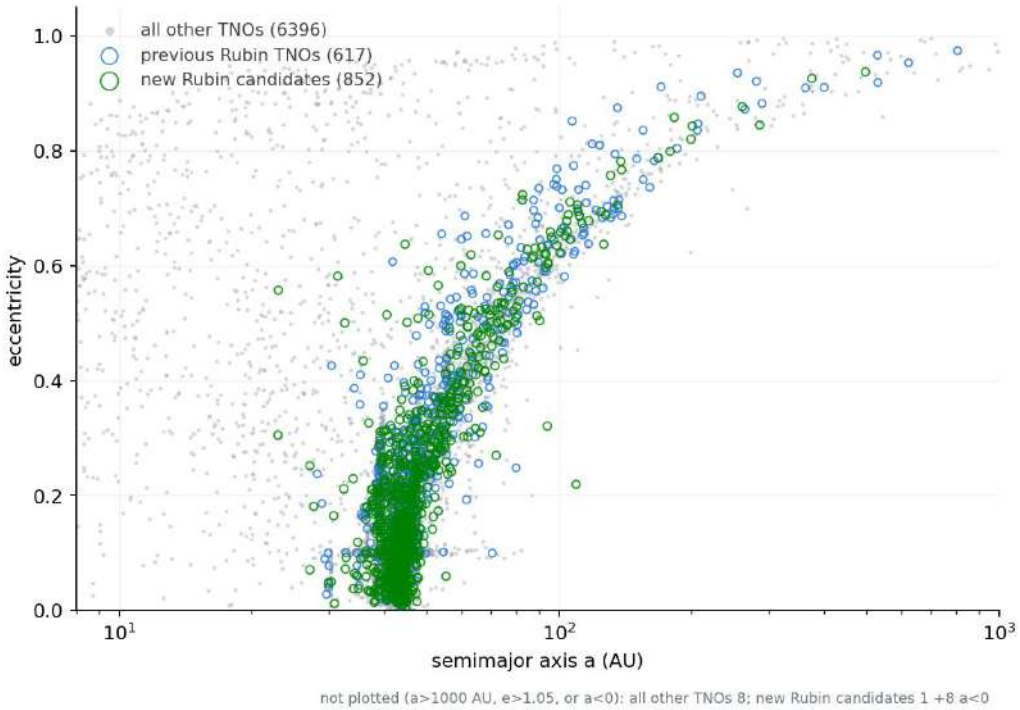
22,000 more with DP2 diffims!

…Plus ~1,400 TNOs

~ 2x the next-largest TNO survey!!!



Matt, Kevin & Ari

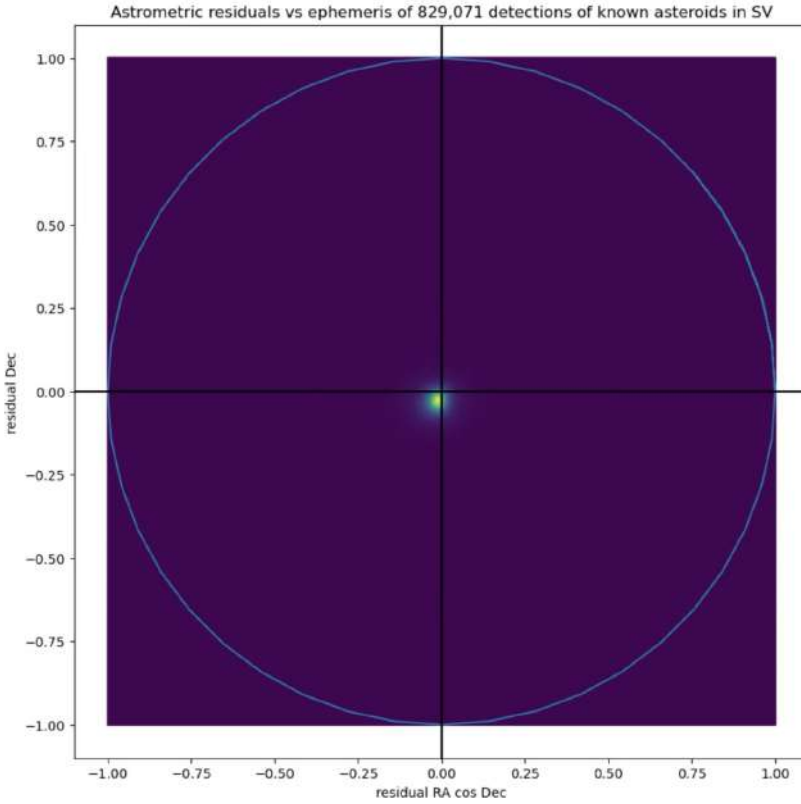


Solar System - DP2 SSSource + SSObject

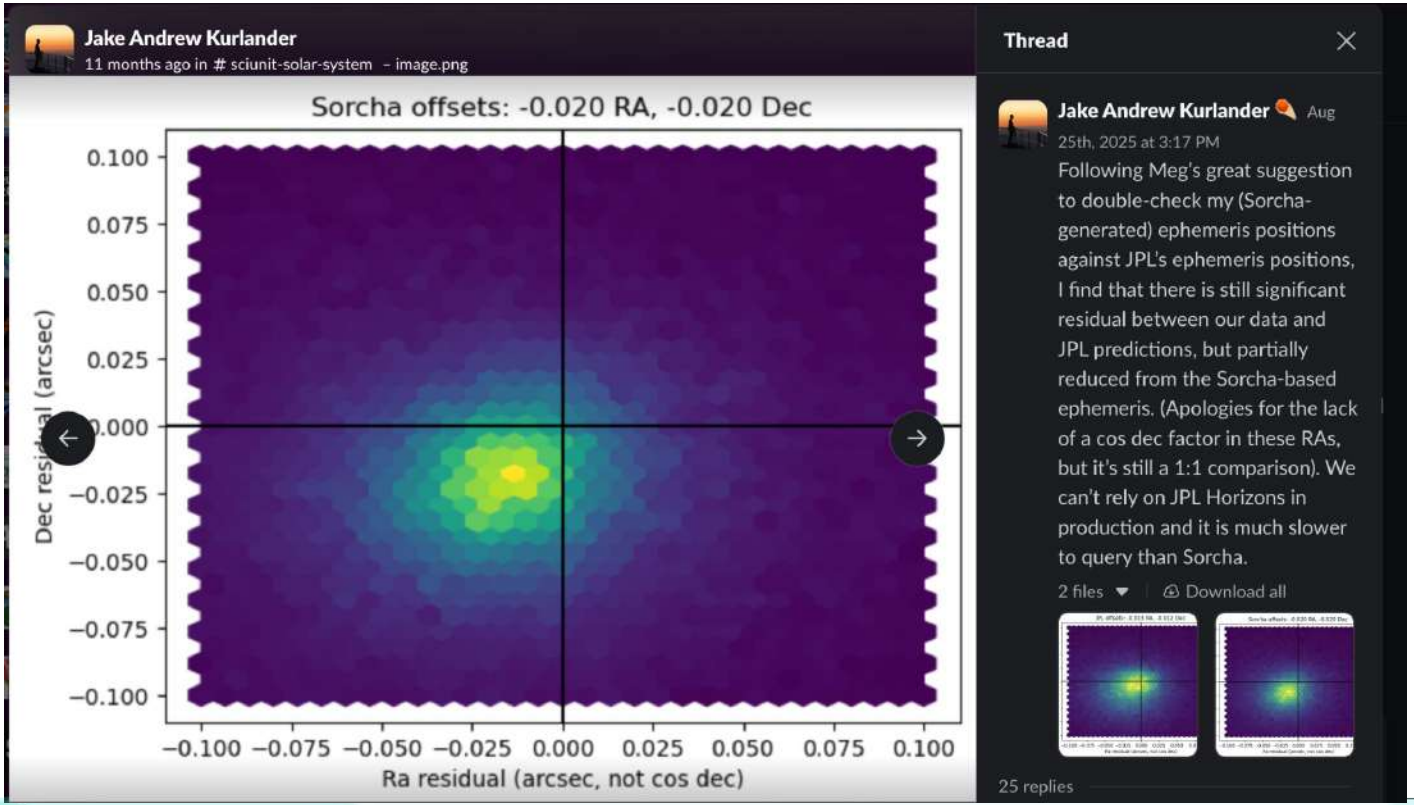
Speaker: Jake Kurlander
4 min

- Very good photometry, astrometry!
- 300k (20% of) known asteroids measured >1
- 8.1 million measurements (~2% of all, ever!)

- Sparse cadence poorly constrains most objects
- 1” association radius allows misassociations
- Sanity checks
  - High reliability
  - psfFlux > 0
  - Measured mag near ephVmag
  - looks like an asteroid, etc.



#sciunit-solar-system August 2025…



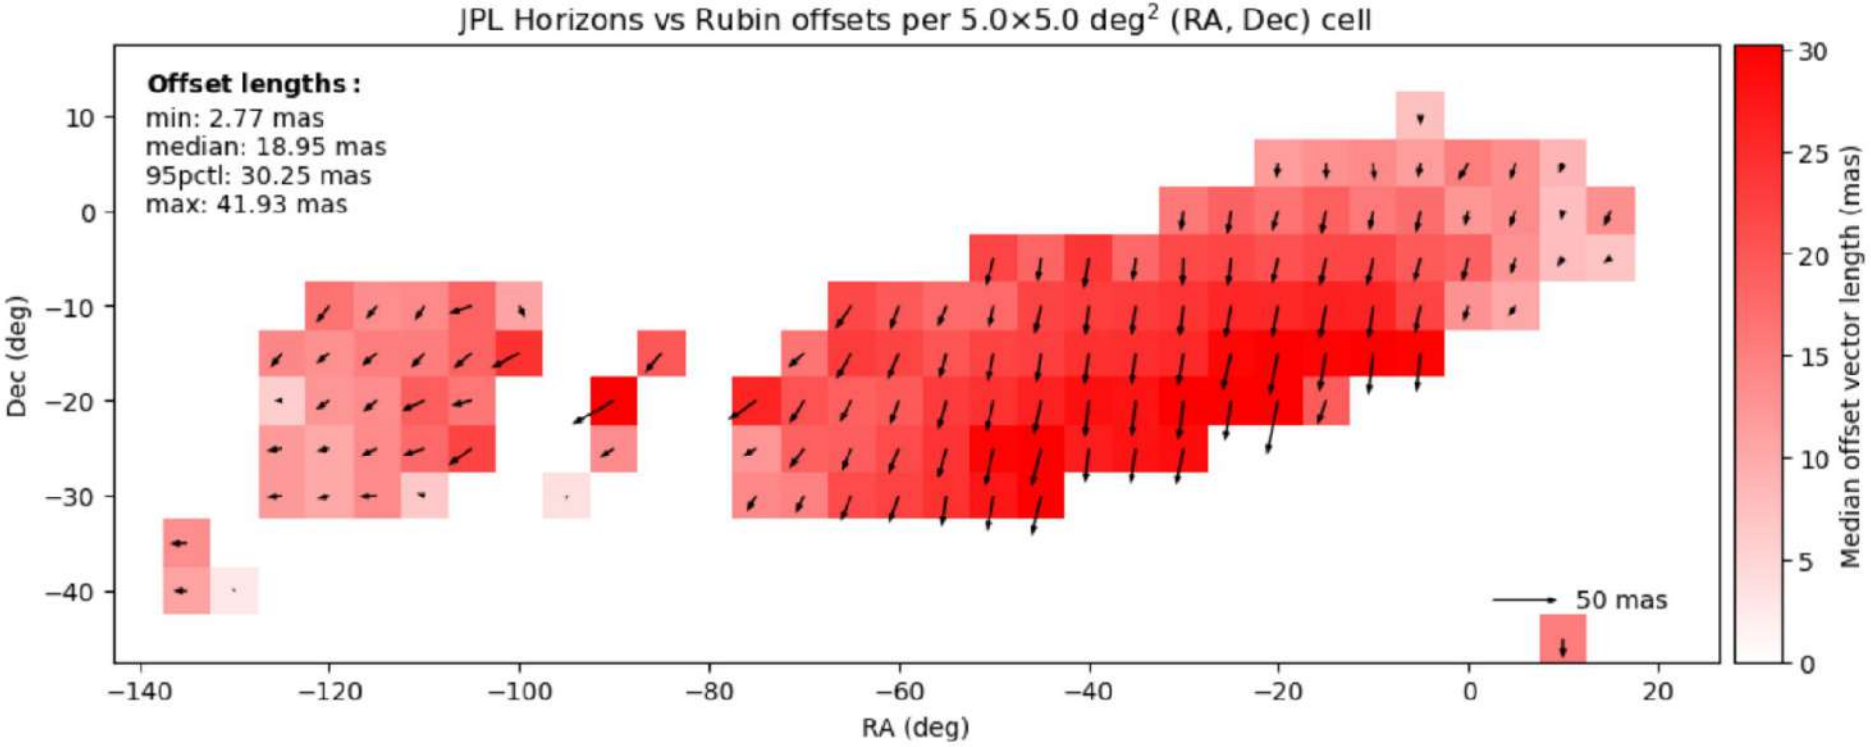


Figure: Pedro Bernardinelli (U. Sao Paolo, Brazil)
Rubin SolSys Science Unit

ATLAS data show a similar signature.

13 replies Last reply 11 months ago

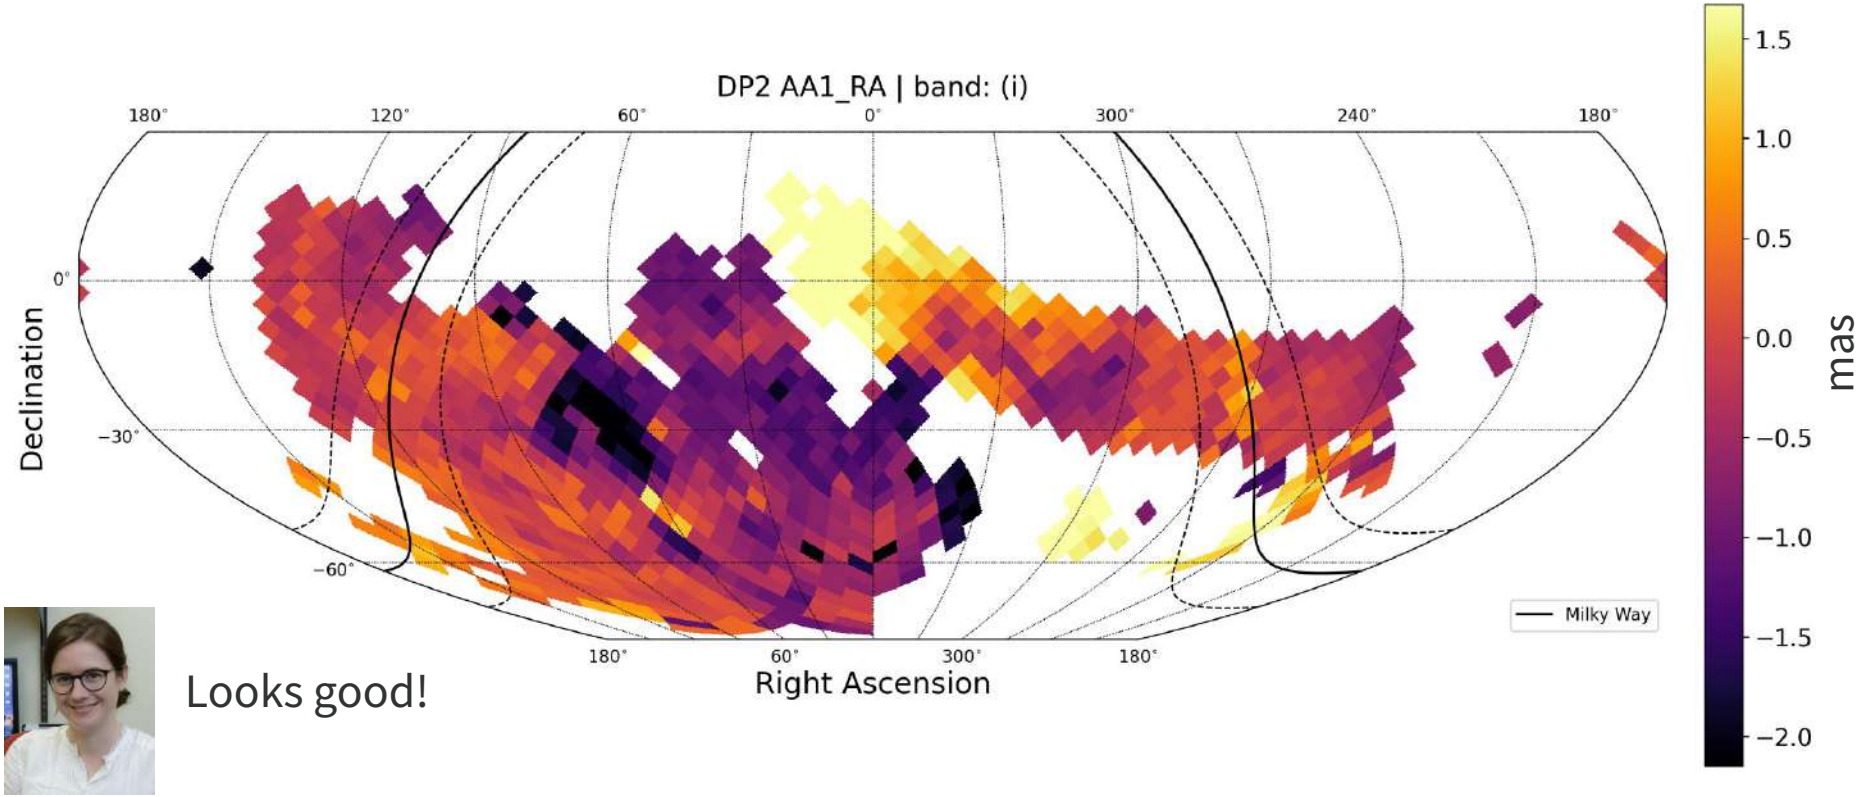
August 26th, 2025

Clare Saunders 1:19 PM

Looking at things from the opposite direction as @Jake Kurlander 's comment above, I've been investigating whether there is any evidence for an overall astrometric offset in our single-frame or fully calibrated catalogs. Since the 3l visits are not in a DRP yet, I looked at data from New Horizons from the collection `LSSTCam/runs/DRP/20250501_20250609/w_2025_30/DM-51933`. Here are some things I checked:

- Single-frame measurements vs Gaia stars: median offsets per visit are generally around 0.2 mas.
- Recalibrated (post-gbdes) measurements vs Gaia stars: median offsets per visit are about the same level (~0.2 mas).
- Offset with Gaia stars as a function of magnitude and as a function of parallax: consistent with zero across magnitude/parallax.
- Overfitting: training vs. validation stars are equally well-fit by gbdes. Checking this for SFM is a little harder, but I think there are enough sources in the tested visits that it's pretty unlikely.
- Offset with Gaia itself: I checked the position of Gaia stars shifted to the mean epoch of the New Horizons exposures accessed through the stack vs through astropy, and the results are equal down to machine precision.
- In the SFM sources, there can be coherent astrometric offsets up to ~15 mas, but these average out over a detector (see figure below) and, at least for the e2v chips, are not consistent visit-to-visit. These should not lead to

The mean offset with Gaia DR3 (AA1) is well within the design requirements



Looks good!

Rubin uncovered a systematic in the pre-Rubin/existing orbits

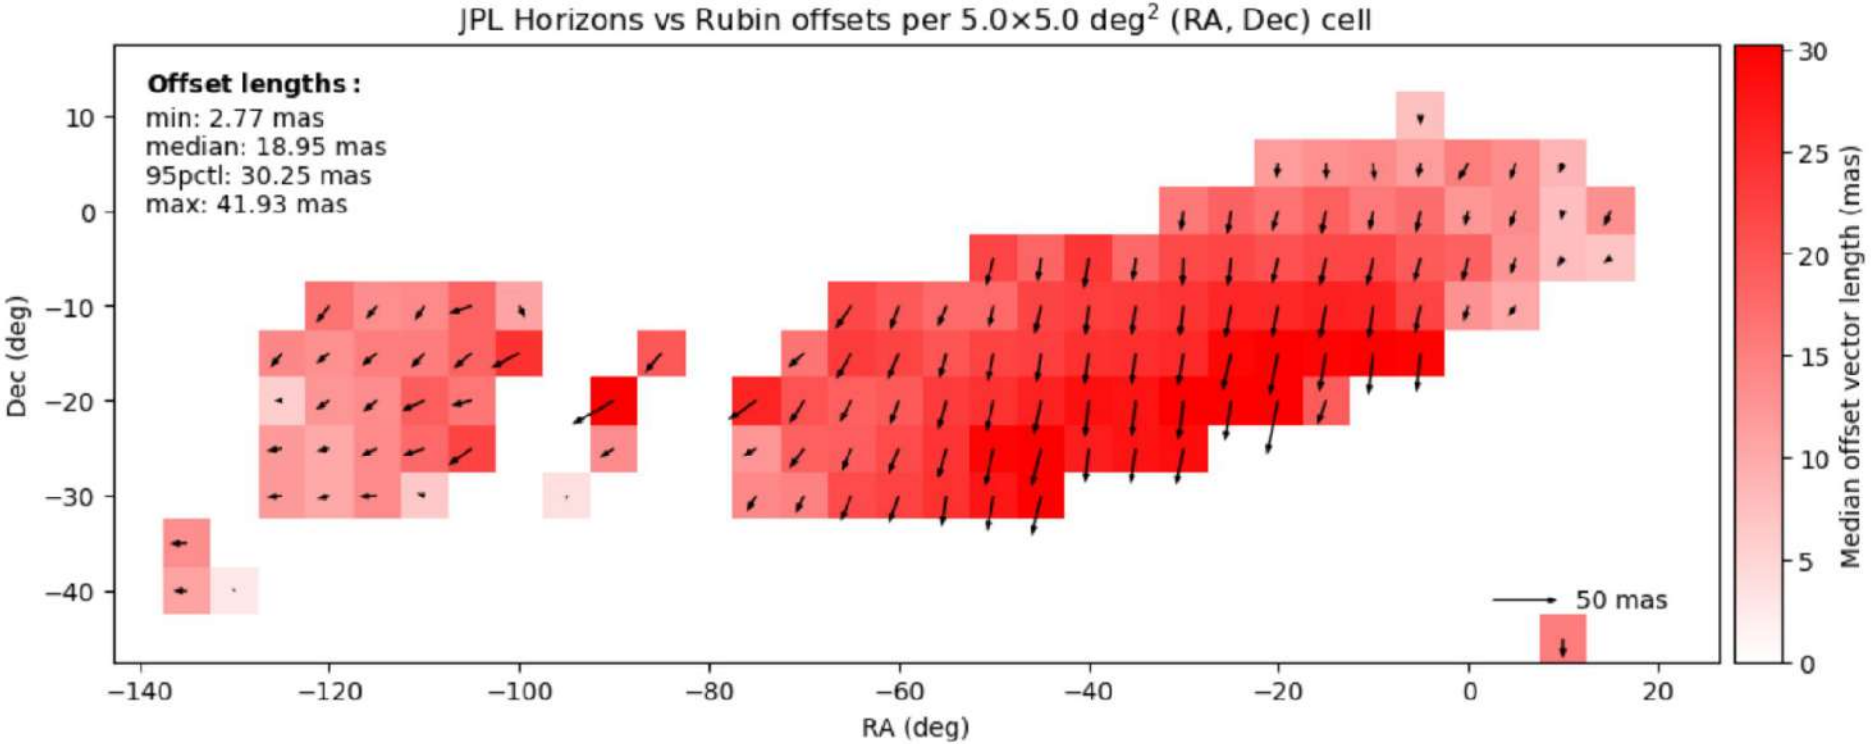


Figure: Pedro Bernardinelli (U. Sao Paolo, Brazil)
Rubin SolSys Science Unit

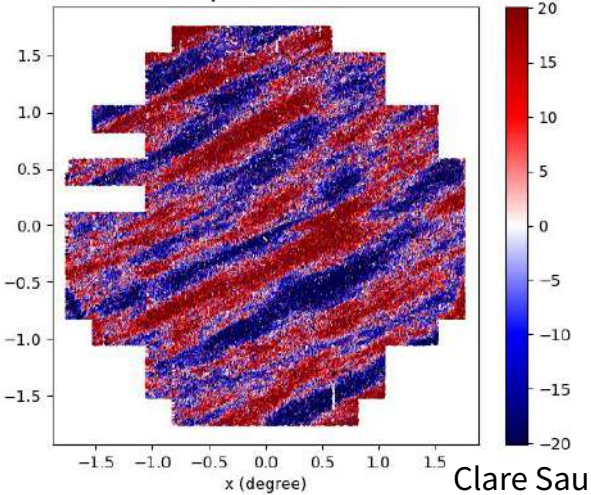
ATLAS data show a similar signature.

Astrometric Calibration
New for DP2:

- As of DP2, the final astrometric calibration is now a two-step process:
  - GBDES calculates the camera distortions and large-scale per-exposure shifts in the astrometric solution.
  - (New for DP2) A **Gaussian Processes model** is then applied on top of the gbdes solution to capture smaller scale variations due to **atmospheric turbulence**

- (New for DP2) **Proper motion and parallax** are fit for all isolated stars using this two-part astrometric solution and **served in the IsolatedStarStellarMotions table.**

Example of Astrometric Residuals due to atmospheric turbulence:

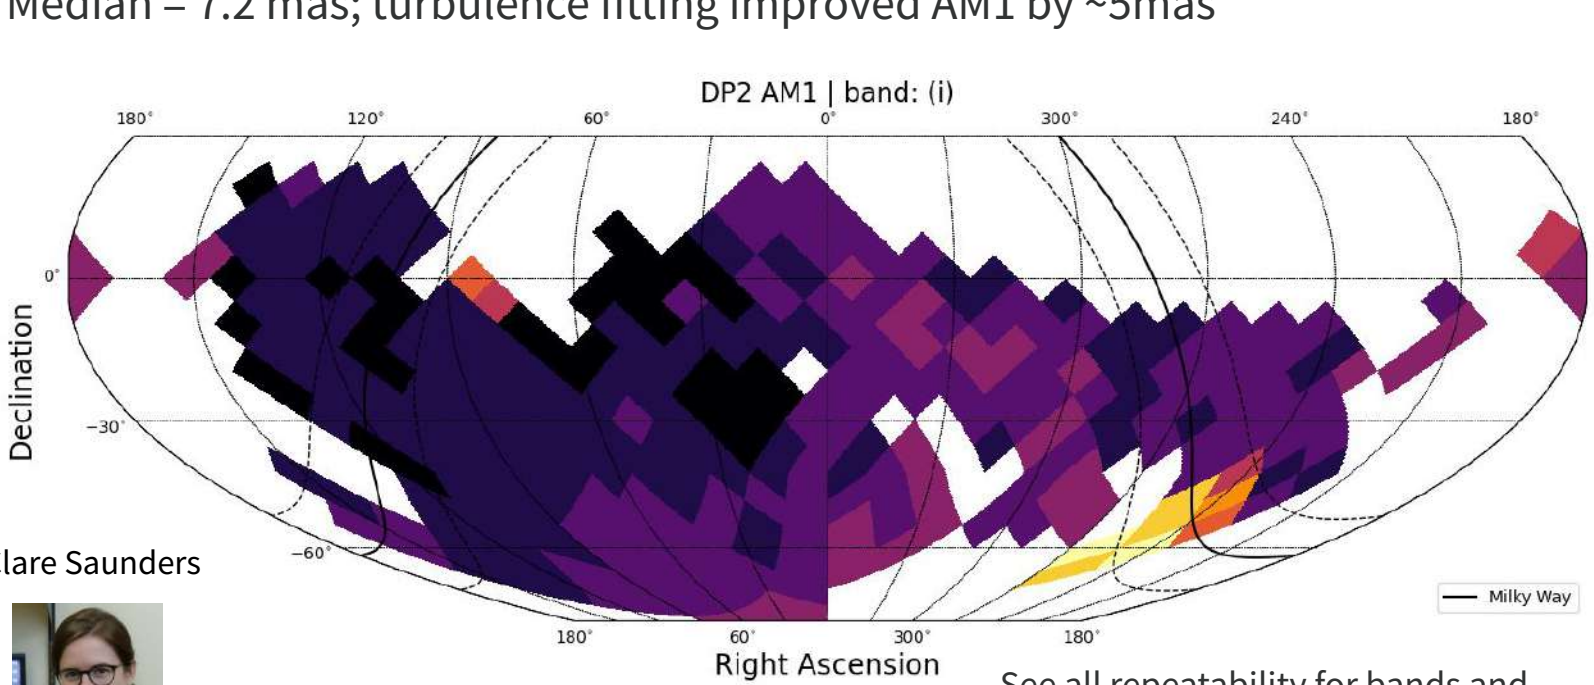


Clare Saunders



Current (DP2) Astrometric Performance: Repeatability of object separations at 5 arcmin (AM1).

Median = 7.2 mas; turbulence fitting improved AM1 by ~5mas



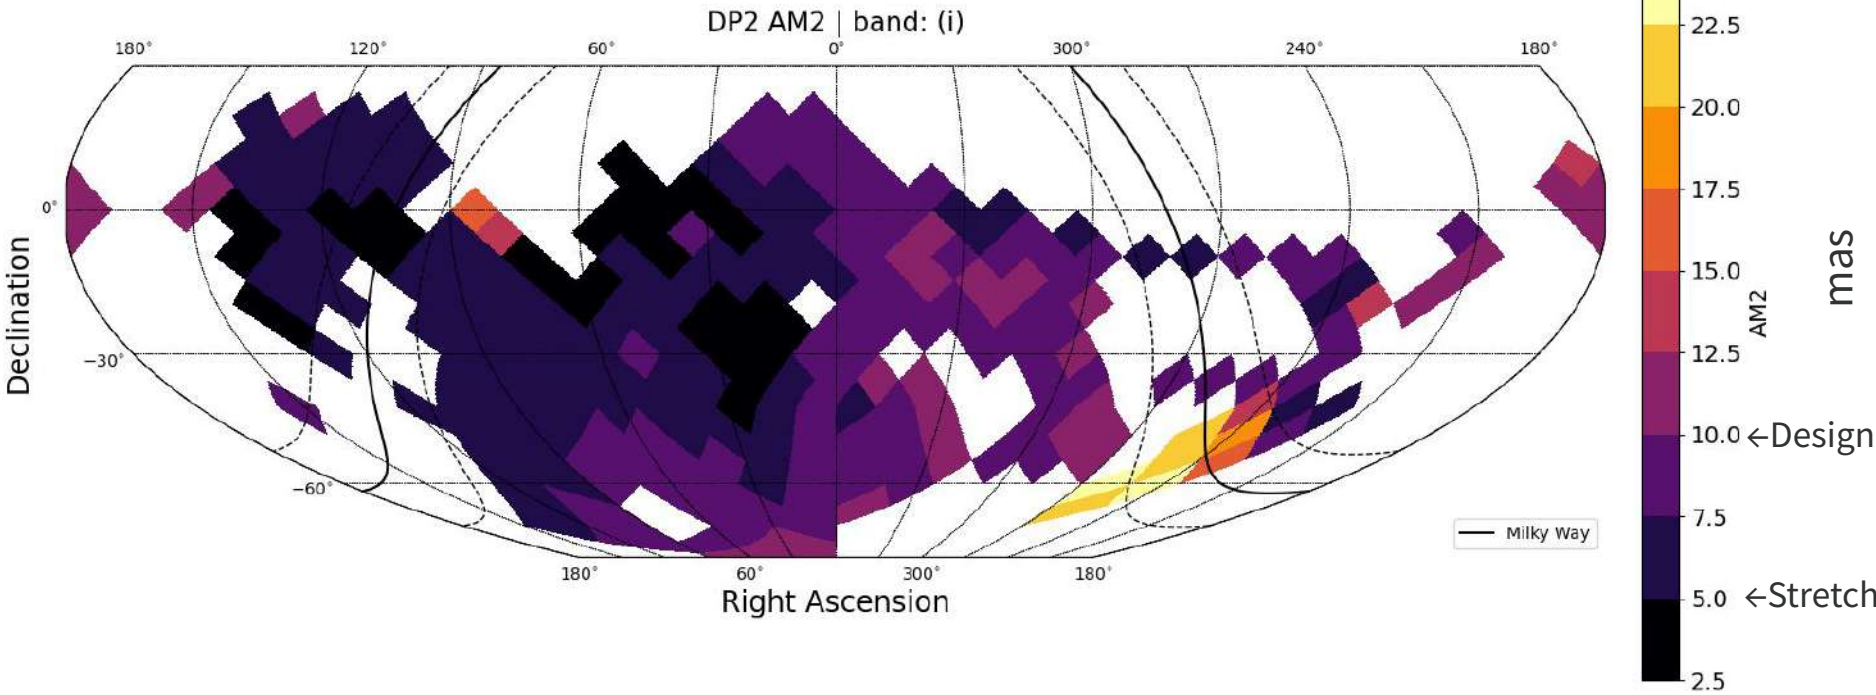
Clare Saunders



See all repeatability for bands and distance bins at [data quality viewer]

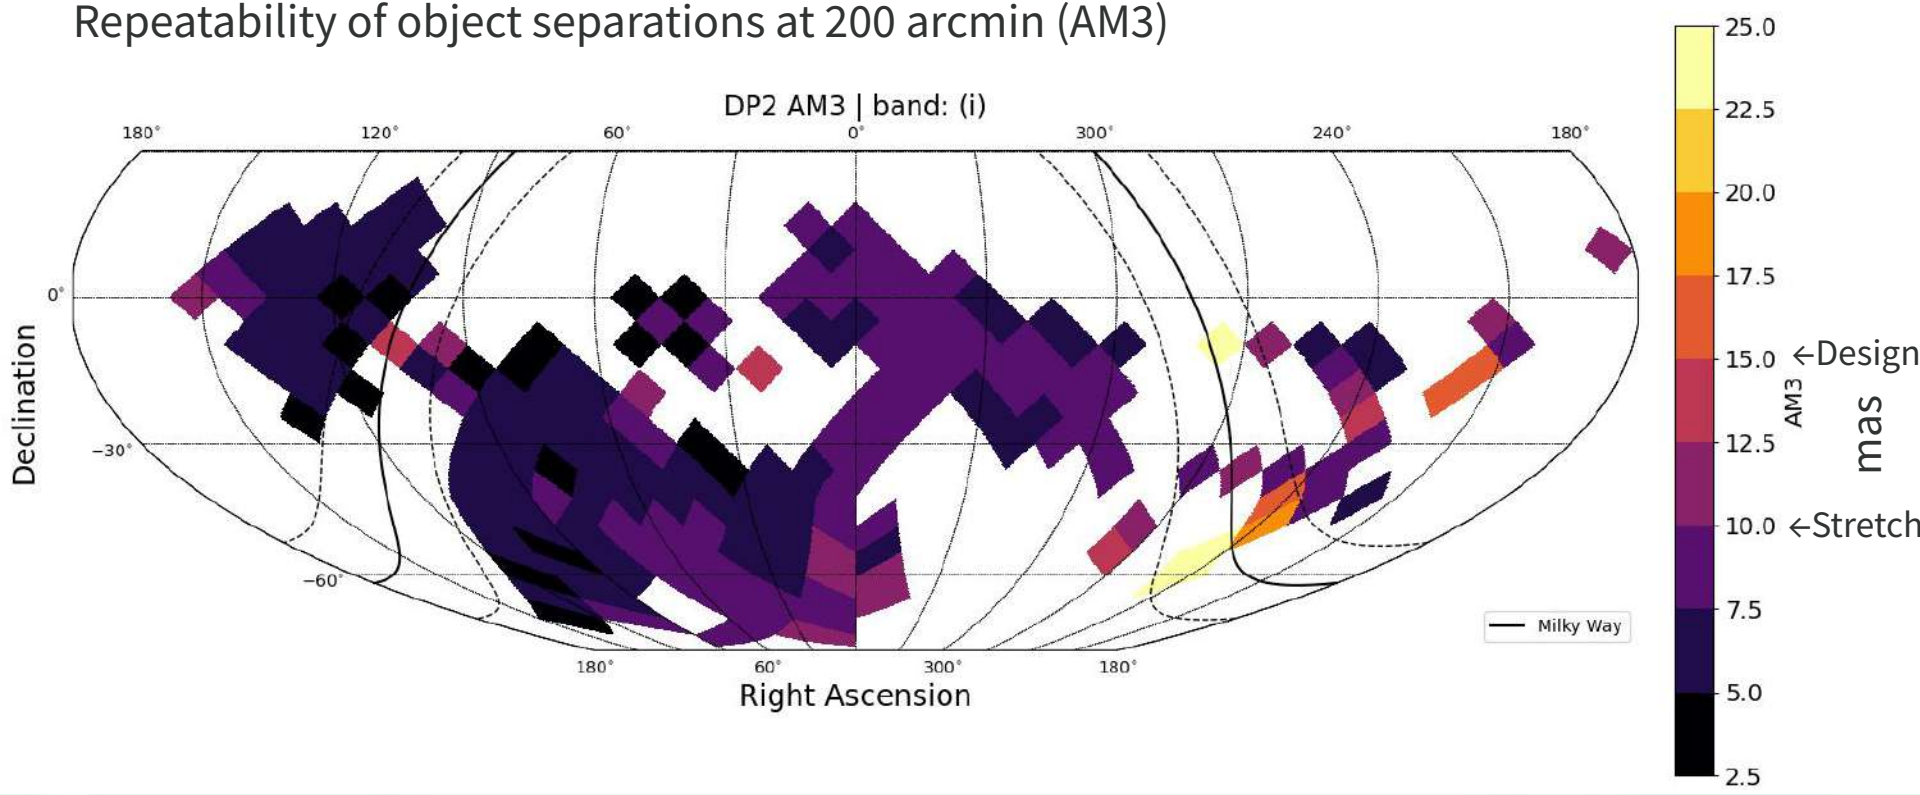
Astrometric Calibration - Current Performance

Repeatability of object separations at 20 arcmin (AM2)



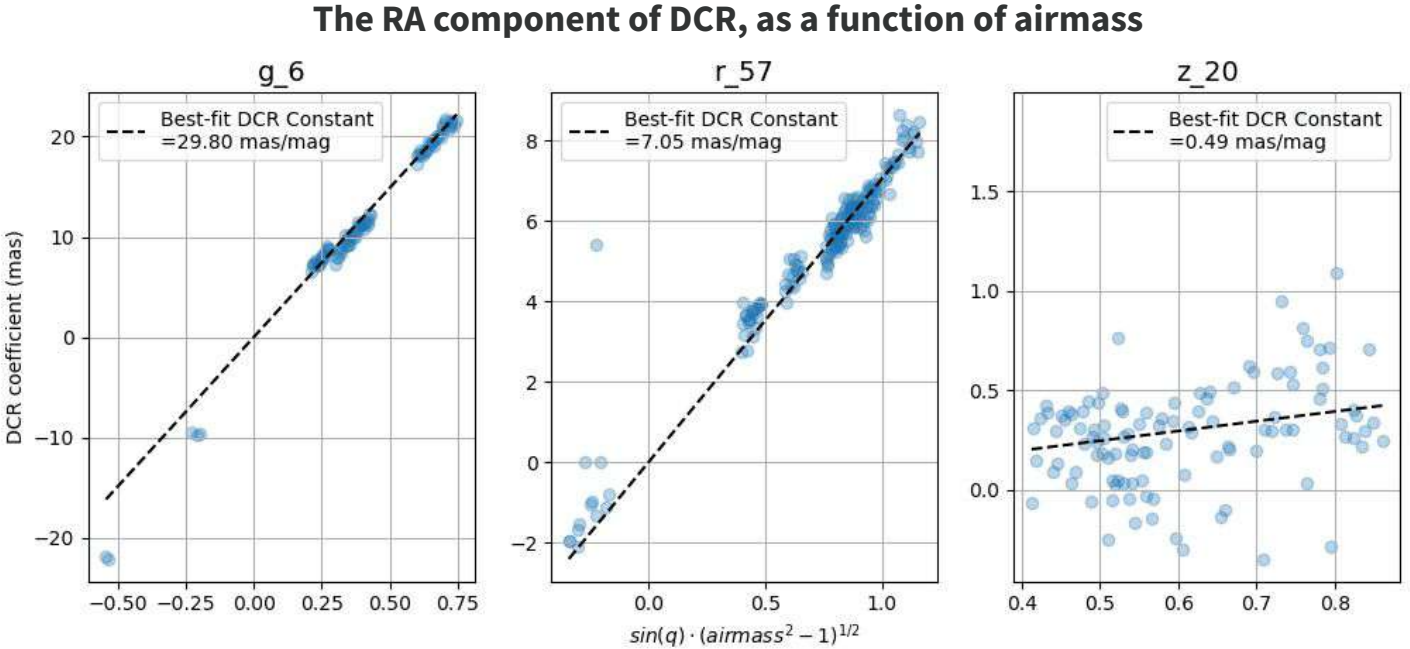
Astrometric Calibration - Current Performance

Repeatability of object separations at 200 arcmin (AM3)



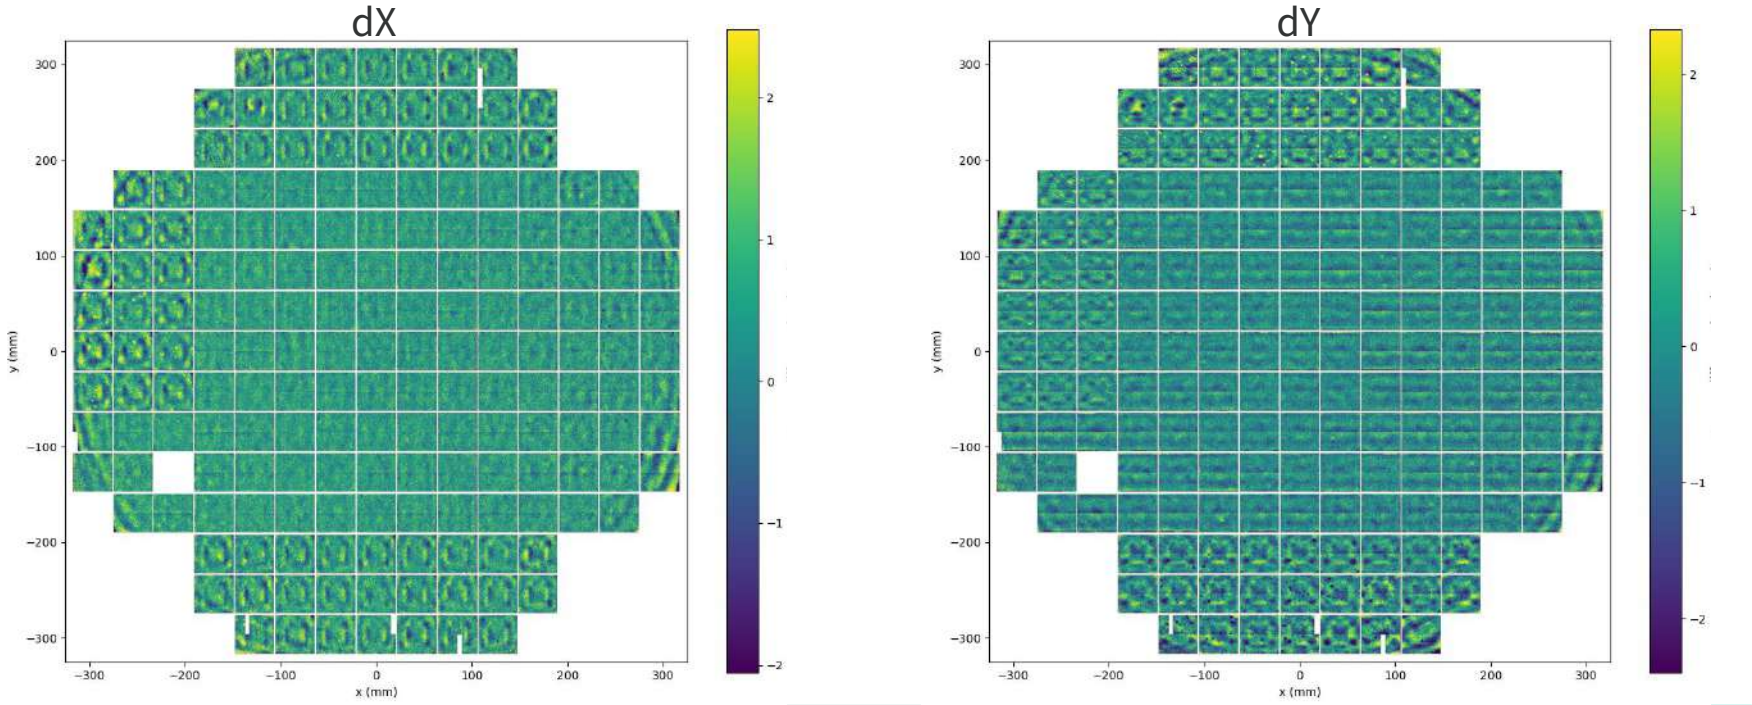
Astrometric Calibration - Upcoming Improvements

DCR correction will be turned on for DR1



Astrometric Calibration - Upcoming Improvements

Stacking residuals in focal plane coordinates shows unmodeled behaviour in the camera



Astrometric Calibration - Upcoming Improvements

Coming in DR1 or the near future:

- Including DCR in the astrometric model
- Computational improvements to speed up the fit and reduce memory usage
- Improvements to the Gaussian Processes fit of the atmospheric turbulence
- Improved modeling of chip distortions
- Proper motion and parallax fits for more categories of objects

Rubin uncovered a systematic in the pre-Rubin/existing orbits

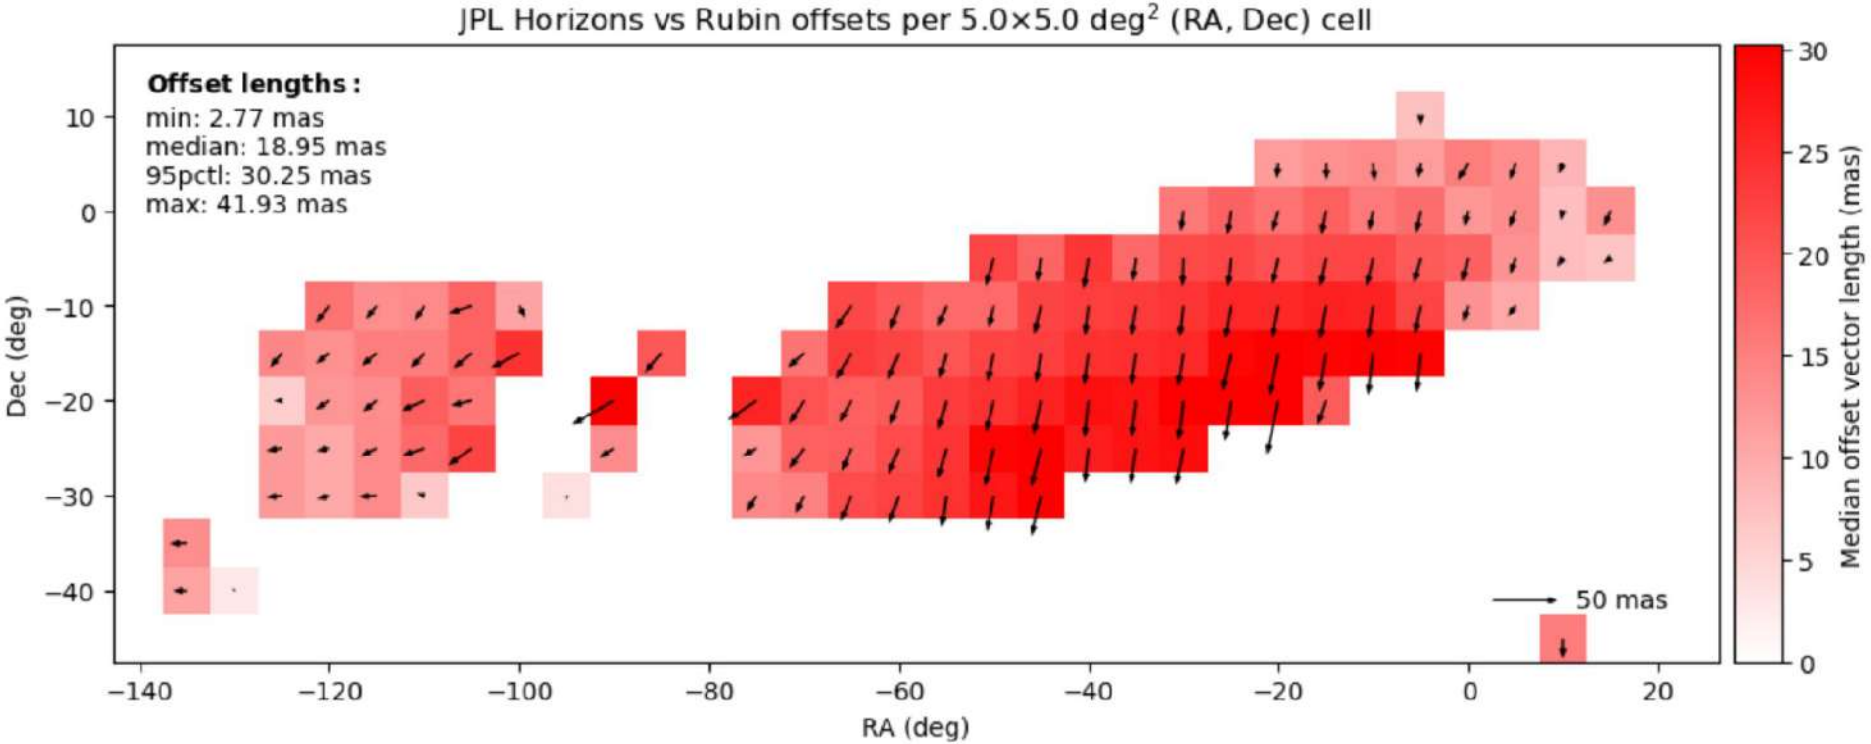


Figure: Pedro Bernardinelli (U. Sao Paolo, Brazil)
Rubin SolSys Science Unit

ATLAS data show a similar signature.