

October 19th, 2025
Astro-Colibri 4
Orsay, France (virtual)

“LSST is just ZTF 2.0” – Ž. Ivezić

ZTF: seeding and interoperating with LSST

Matthew J. Graham
mjg@caltech.edu

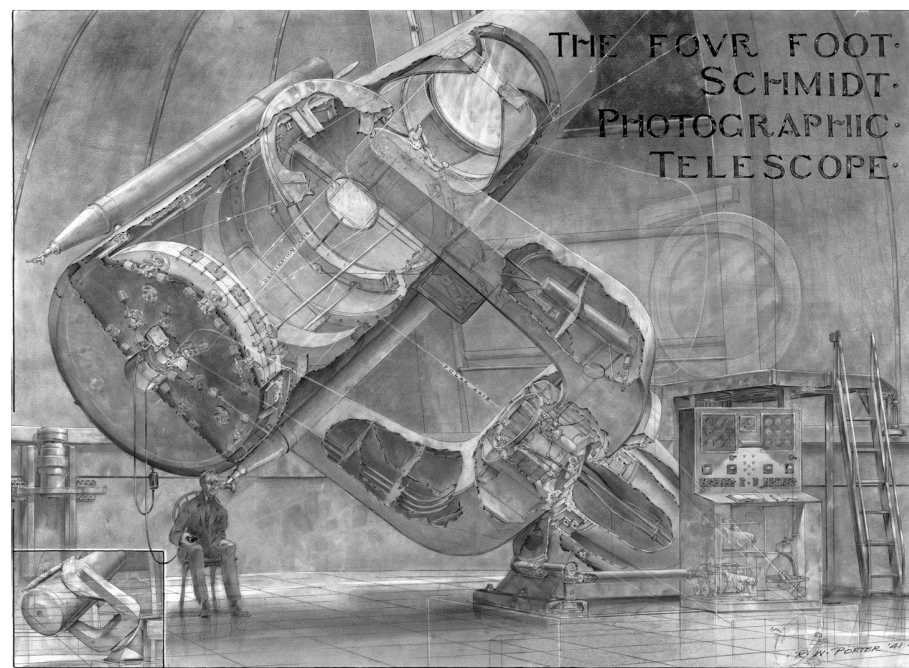
with contributions from
Theo du Laz, Michael Coughlin



ZTF: the first industrial astronomical stream



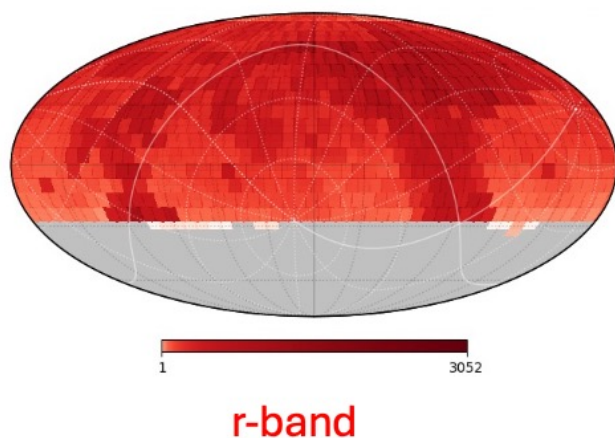
- The Palomar Oschin 48" telescope took its first image on Sept 30, 1948
- ZTF is its latest instrument: a 47 deg² field of view camera with > 660M pixels
- It can cover 3750 deg² / hr to 20.5-21 mag (30s exposures)
- It carries out a full northern sky every two nights in *g*, *r* (and *i*)
- First light was Oct 2017; survey started Mar 2018; first alerts issued Jun 2018
- Supports ToO programs for MMA
- ZTF is LSST 0.1



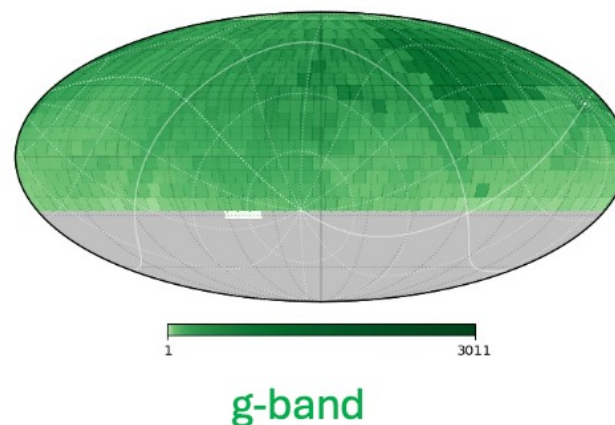
ZTF by the numbers



- Total sky area covered is > 50 million $\text{deg}^2 = \sim 2\times$ LSST 10-year total
- > 930 million photometric measurements for 4.9 billion sources
- ~ 1750 observations per source for ~ 2 billion sources
- Up to 1 million alerts per night
- Over 800 million alerts (> 65 TB) published



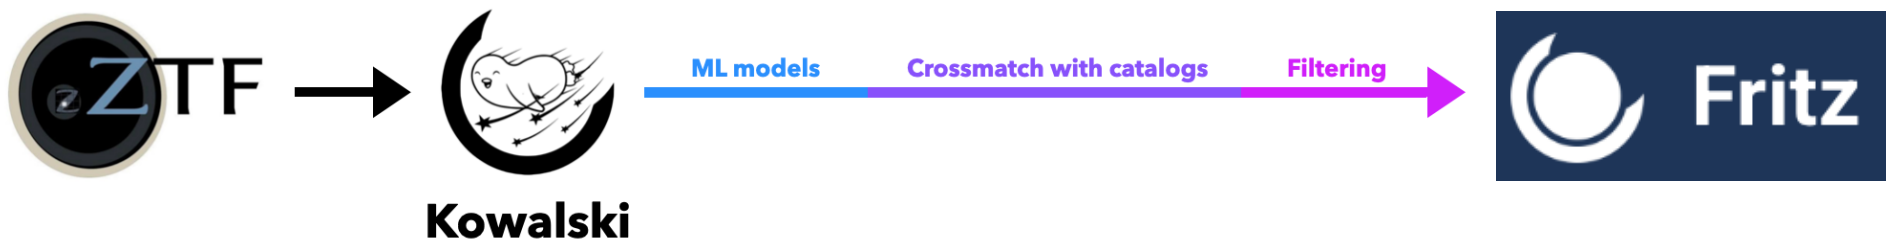
Coverage
 > -30 deg
declination



Transients (alerts)

- ZTF has generated >800 million alerts (5σ difference image detection)
- Switched to forced photometry light curves in ZTF alerts (summer 2024)
- ~500,000 of these are *interesting* (have passed a filter)

Pre-LSST:



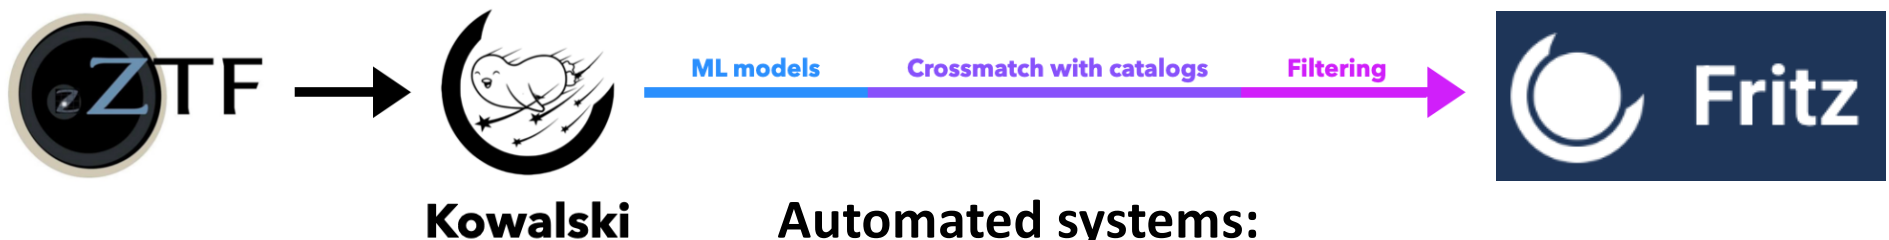
Kowalski:

- Receive alerts from ZTF
- Add spatial cross-matches with archival catalogs
- Update full light curve for each object
- Run ML models on the alerts, e.g. BTSBot
- Run user-defined filters and send filtered data to SkyPortal/Fritz

Transients (alerts)

- ZTF has generated >800 million alerts (5σ difference image detection)
- Switched to forced photometry light curves in ZTF alerts (summer 2024)
- ~500,000 of these are *interesting* (have passed a filter)

Pre-LSST:



Automated systems:

- **BTSBot**: 1 hr from detection to TNS reporting
- **BTSBot-nearby**: SWIFT observations **6 min** after detection

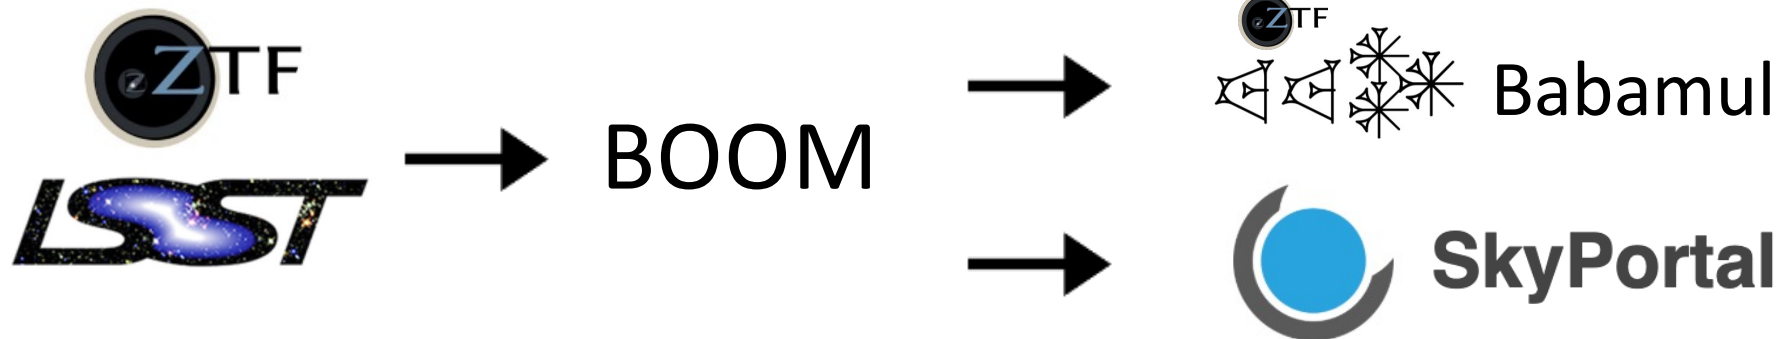
Kowalski:

- Receive alerts from ZTF
- Add spatial cross-matches with archival catalogs
- Update full light curve for each object
- Run ML models on the alerts, e.g. BTSBot
- Run user-defined filters and send filtered data to SkyPortal/Fritz

Transients (alerts)



With LSST:



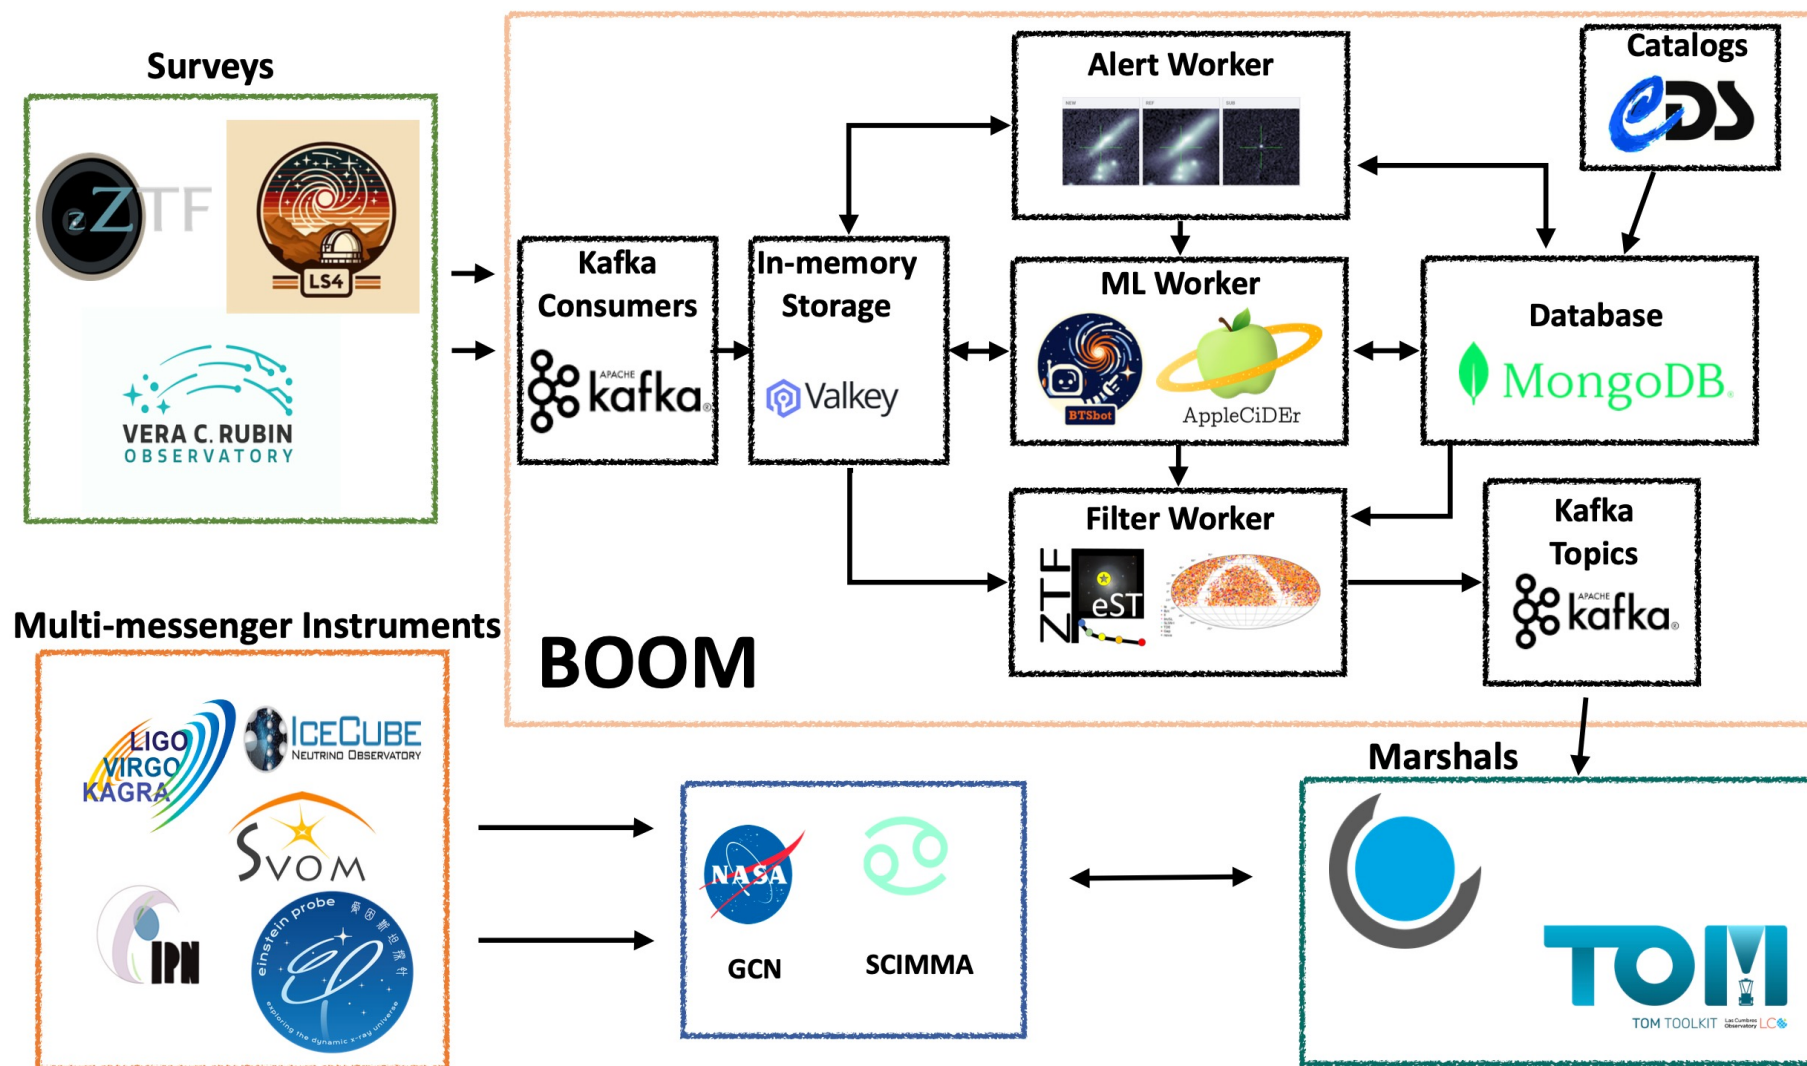
Why reimplement:

- Designed for ZTF volume (<1M alerts/night)
- Constant crossmatching
- Current filter building is not user friendly

Benefits:

- RUST-based code lowers memory and CPU usage,
- Greater parallelization, lower latency
- Modularity
- Distributed

Burst & Outburst Observations Manager (BOOM)



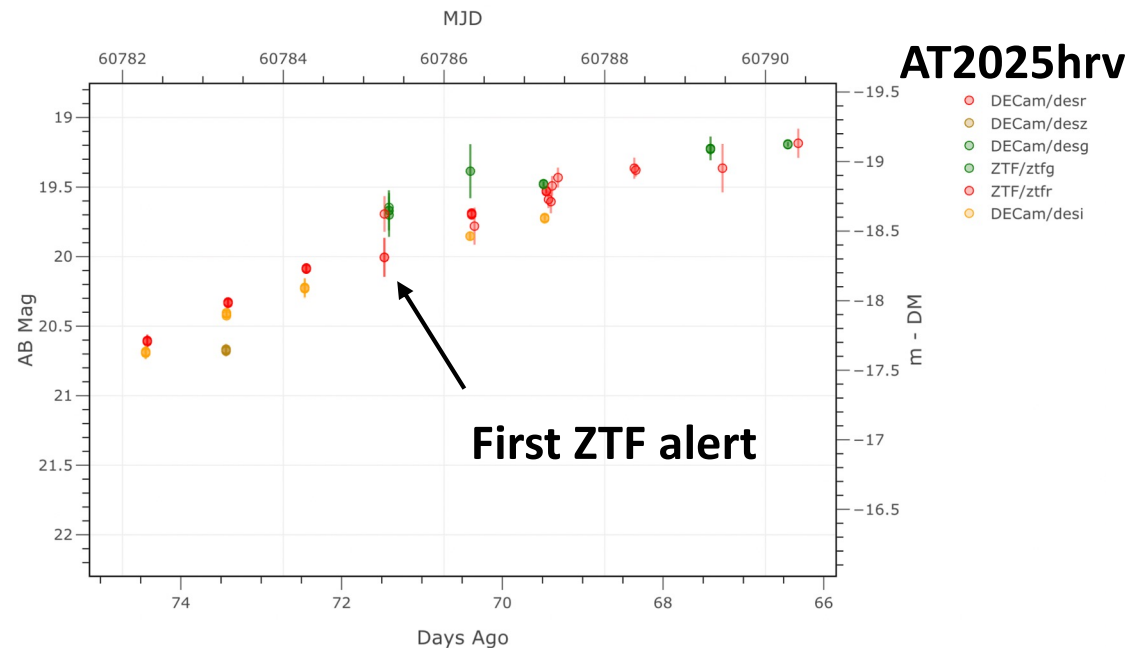
BOOM: Multistream alert architecture

Multisurvey support:

- We can successfully process the ZTF and Rubin alert stream
- Every new ZTF alert is crossmatched with the LSST alert-based object collection
- Future iterations will crossmatch by position, time, and brightness

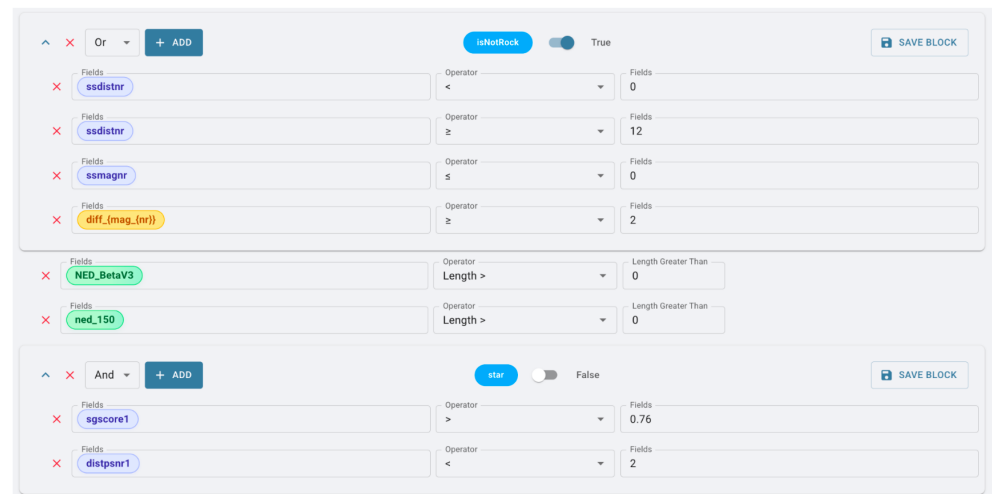
Preview with ZTF+DECAM:

- Sources with DECAM pre-detection



BOOM: Filtering

- Most brokers offer one of:
 - Static filters that users can read from for further downstream filtering
 - Limited filtering from a web app
 - Run user-provided Python filtering code
- Most users do not want to learn complex (non-Python) syntax
- Developing a UI/web-based no-code tool only requiring familiarity with the data product
 - Share blocks between users
 - Precompute common blocks
- **BABAMUL** will offer a number of phenomenological filters on alerts
 - Aim to serve a number of community-driven



The screenshot displays a web-based filtering interface with two main sections for creating filter blocks.

Top Section (Or filter):

- Header: **Or** (dropdown), **+ ADD** button, **listOfRock** toggle (True), **SAVE BLOCK** button.
- Filter 1: Fields: **ssdistnr**, Operator: **<**, Fields: **0**.
- Filter 2: Fields: **ssdistnr**, Operator: **≥**, Fields: **12**.
- Filter 3: Fields: **ssmagmr**, Operator: **≤**, Fields: **0**.
- Filter 4: Fields: **diff_(mag_inr)**, Operator: **≥**, Fields: **2**.

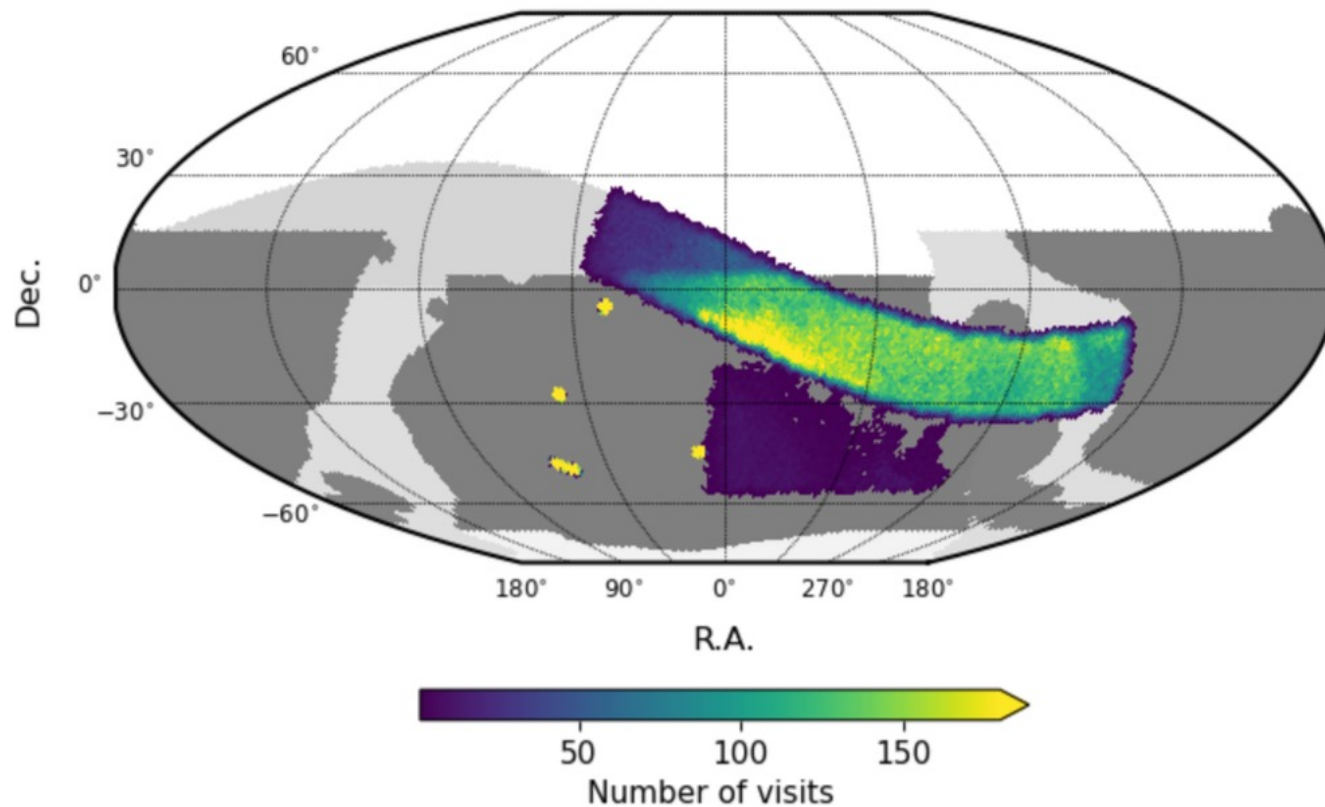
Bottom Section (And filter):

- Header: **And** (dropdown), **+ ADD** button, **star** toggle (False), **SAVE BLOCK** button.
- Filter 1: Fields: **sgscore1**, Operator: **>**, Fields: **0.76**.
- Filter 2: Fields: **distpsnr1**, Operator: **<**, Fields: **2**.

ZTF/LSST synergy already happening



- June – Sept 2025, ZTF shadowed the SV region every night in *gri*
- ZTF alerts released in real time
- Imaging data and light curves will be part of DR24 (Jan 2026)



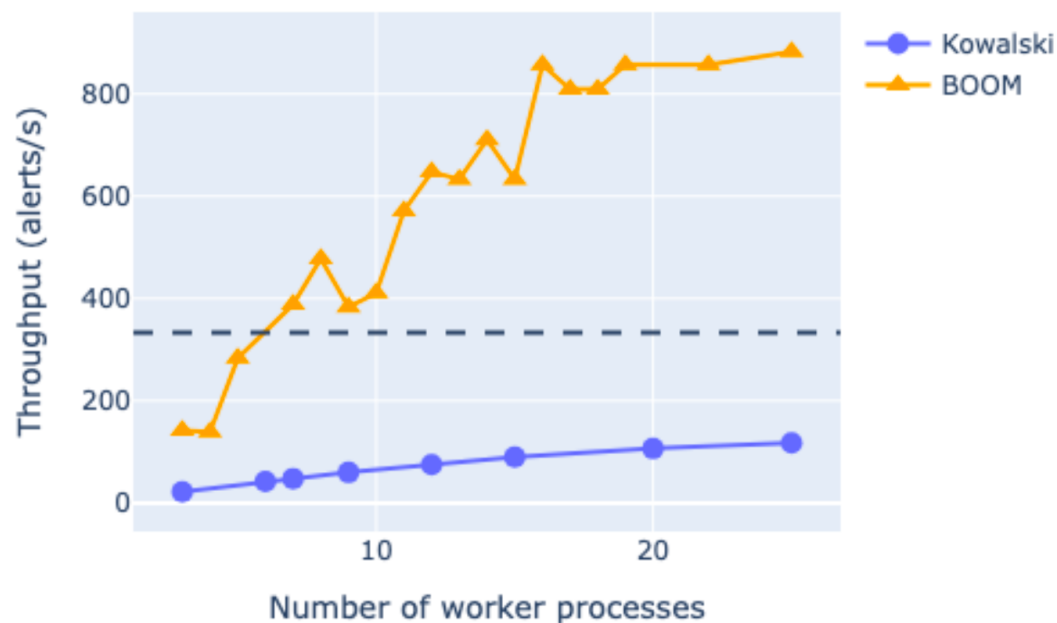
ZTF Plan for 2026



Summary: Nightly *gri* coverage of Rubin footprint

Details (TBC):

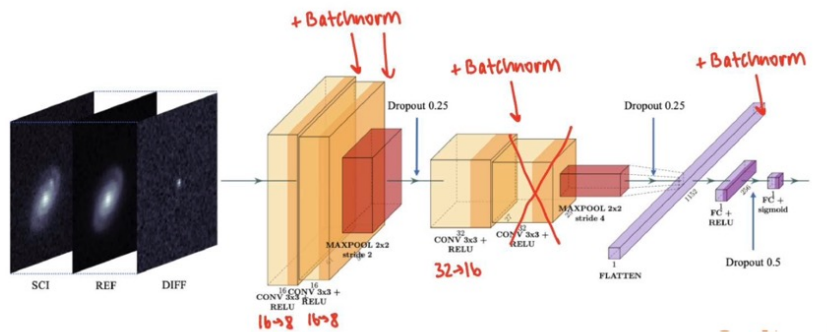
- Once LSST is fully operational, ZTF public time (50%) will do a nightly *gri* coverage of LSST footprint visible from Palomar (ZTF grid #1)
- Joint alert stream for alerts in ZTF and LSST
- A base set of filters characterizing events: blue, large amplitude, etc.



Fast inferencing

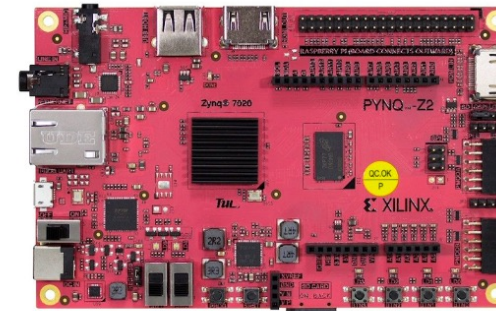
faster_braai (Deng+, in prep.)

- FPGA deployment
- Reduced CNN model using 6k parameters:



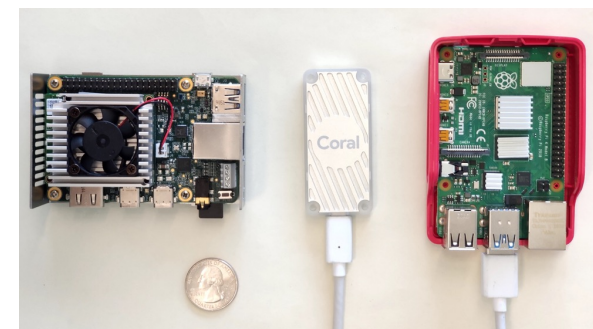
- 16-bit quantization-aware training:
 - quantized weights in forward pass
 - full-precision weights in backward pass and weight update

- Ported to PYNQ-Z2 FPGA using hls4ml:
 - optimized for latency and streaming data

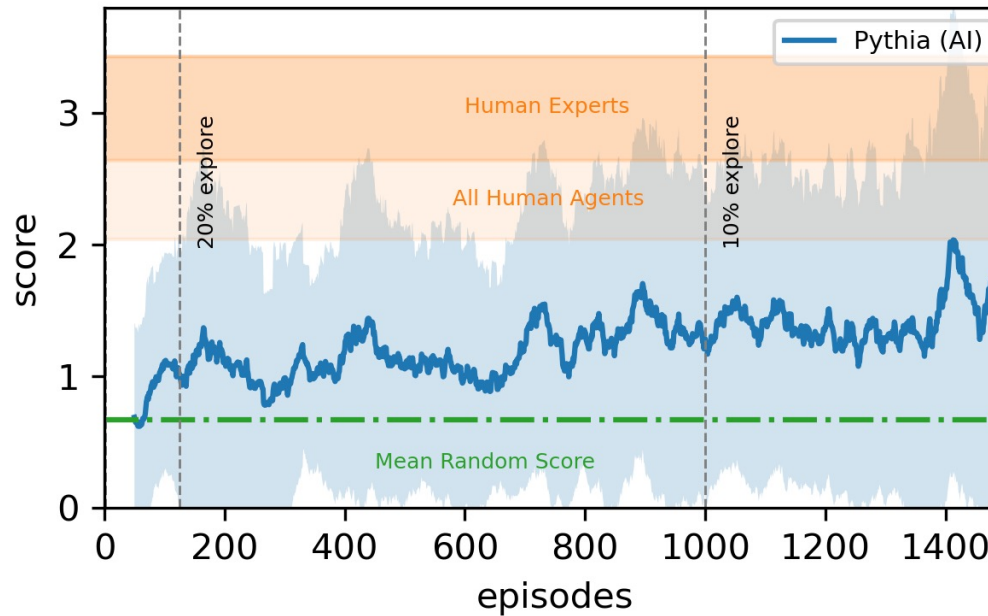


- Recent FPGA support for transformers, decision trees, and symbolic regression

Device	Inferences / sec
ZTF delivery rate	25
LSST delivery rate	300
USB Edge TPU (on RPi 3B)	333
PYNQ-Z2 FPGA	520
Coral Dev Board	1200



Pythia vs humans



Sravan+ 2023

agent	score	frac KN > 1 follow-up
Pythia	1.84	0.81
Non-expert 1	2.04	0.54
Non-expert 2	3.15	0.86
Expert 1	2.64	0.76
Expert 2	2.74	0.78
Expert 3	2.94	0.72
Expert 4	3.43	0.9

- Linear VFA hypothesis class not sufficiently rich representation of true Q function
 - Benefit is theoretical convergence guarantee. **Demonstrates problem learnable!**
 - Shifting to deep Q networks:
 - Will remove two-step learning, one for $x(s,a)$ in supervised/unsupervised learning and one for Q via Bellman updates in RL
 - Efficient evaluation of realistic large action space, can have vector instead of scalar output
 - SARSA is known to be the weaker of the simple policies, specifically for short horizons with low penalties: QL is expected to perform better
 - Use GNN instead of CNN for permutation invariance
 - Shifting to deep Q networks:
 - Will remove two-step learning: one for $x(s,a)$ in ((un)supervised learning and one for Q via Bellman update in RL
 - Efficient evaluation of realistic large action space with vector output instead of scalar
- => Sybil



ZVAR: ZTF Variability and Periodicity



Goal: Produce a *definitive* ZTF variability and periodicity data set open to the astronomical community

Phase 1 (Now) – 2027 release

1.5 billion sources (2/3 of sky)

$m < 22$ (essentially all Gaia sources)

Image difference forced photometry

Photometric errors $\sim 1 - 2\%$ ($m < 18$)

> 1750 observations per source (median)

Periodicity detection down to < 5 min

Phase 2 (“ZForce”) – 2028 release

3 billion sources

$m < 24$ (New ZTF Deep Catalog)

Science image forced photometry

Proper motion compensated

Detrended (co-trending analysis)

Photometric errors $\sim 0.5\%$ ($m < 18$)

