Machine learning in multi-messenger astrophysics: the present and the future

Michael W. Coughlin

May 20, 2024

The Time-Domain Astronomy Ecosystem

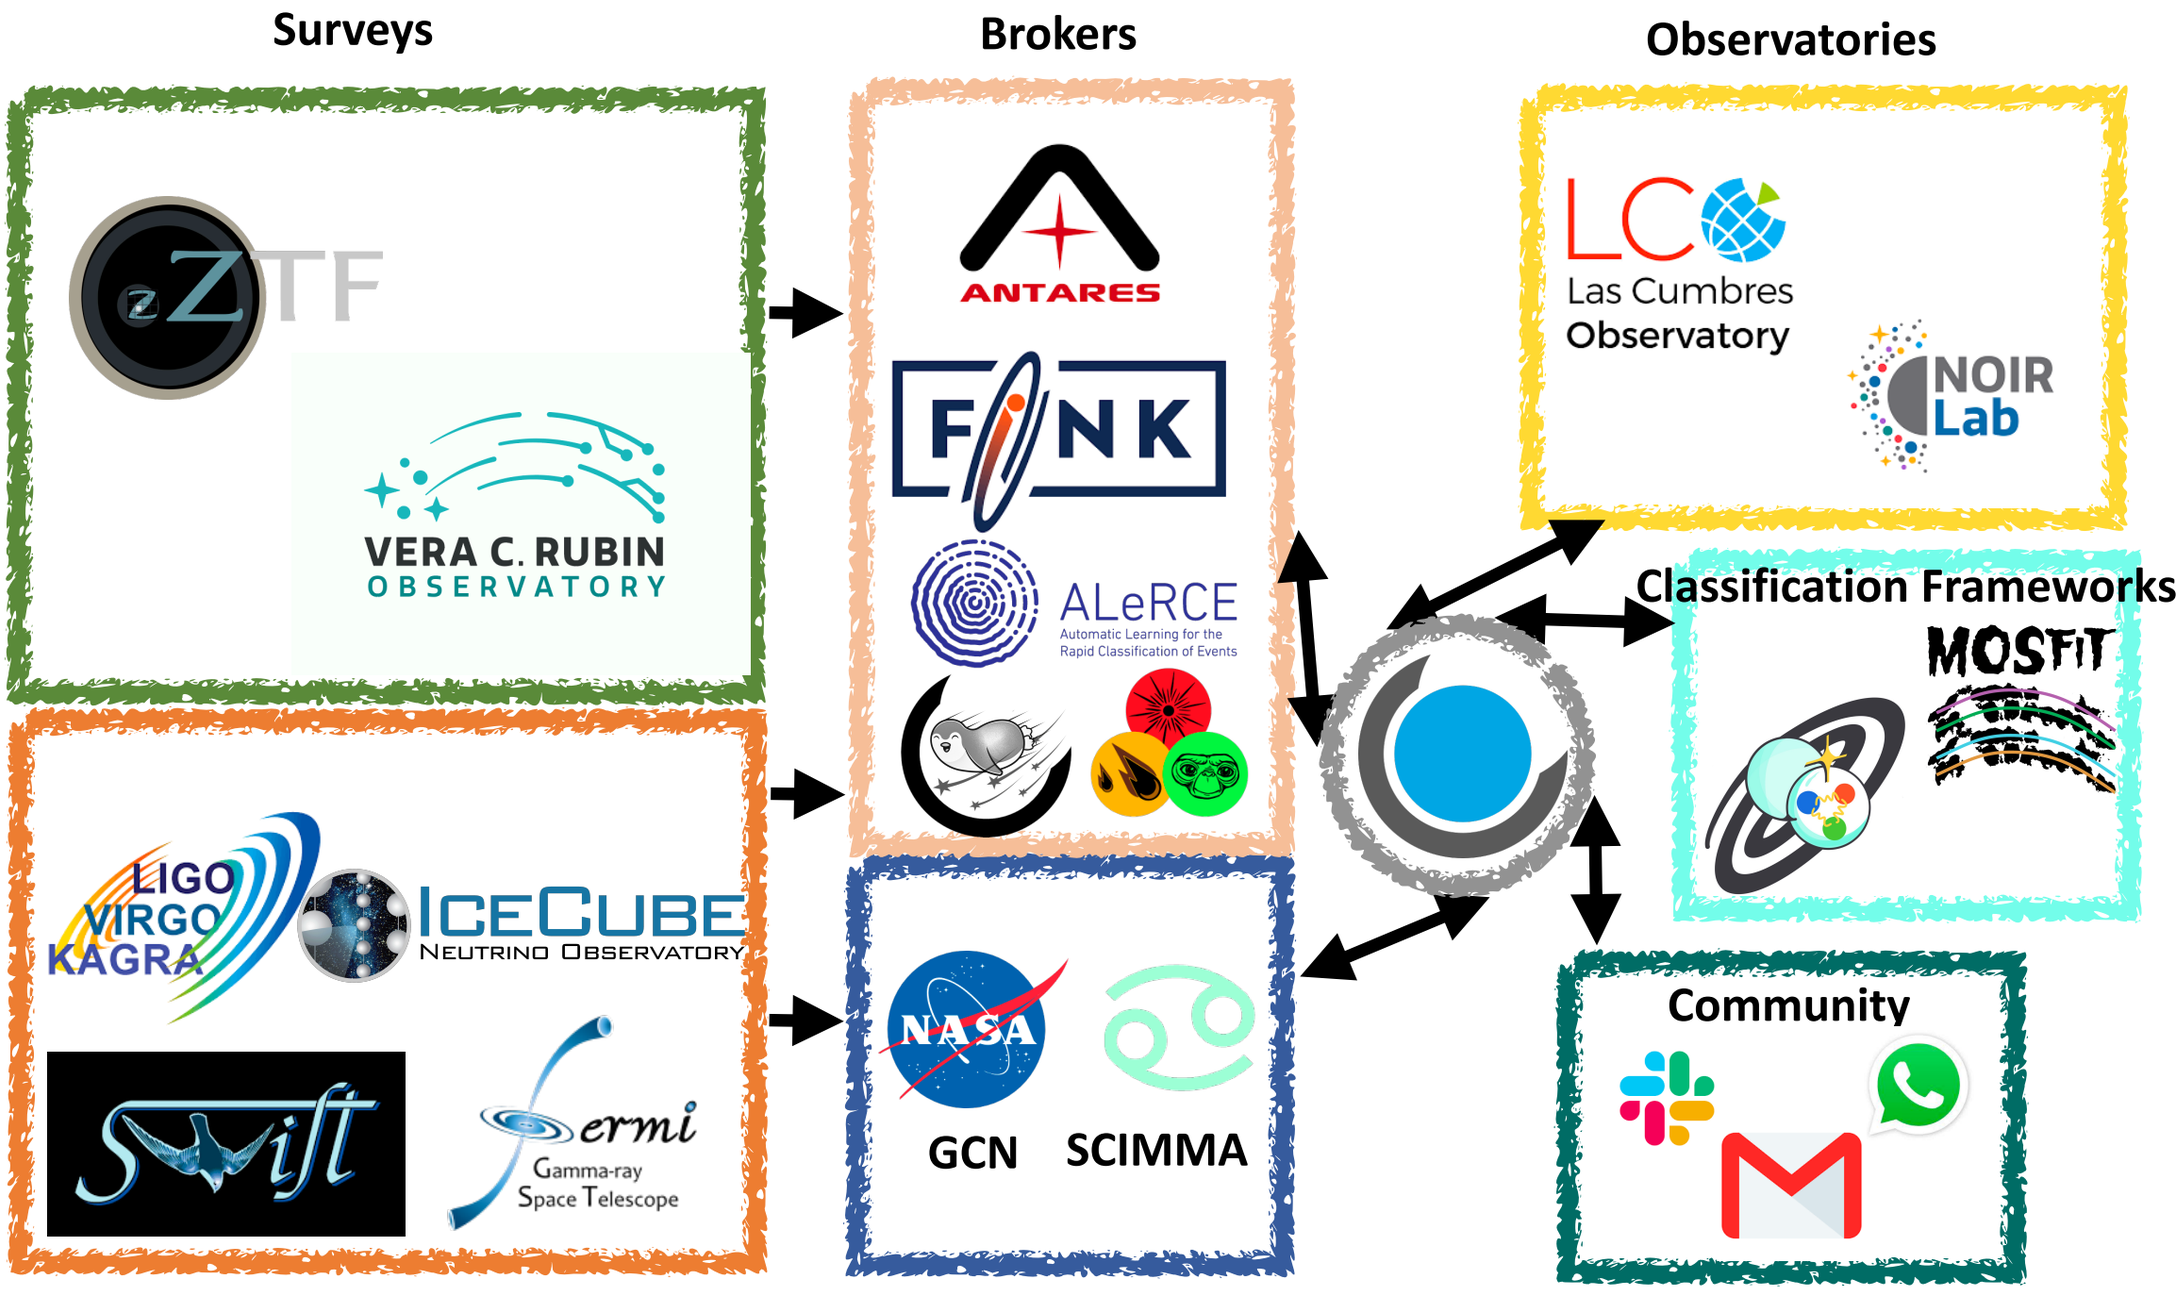

Multi-Messenger Astronomy: The (long) road from data to science

Kasliwal et al. (2020): 2008.00008

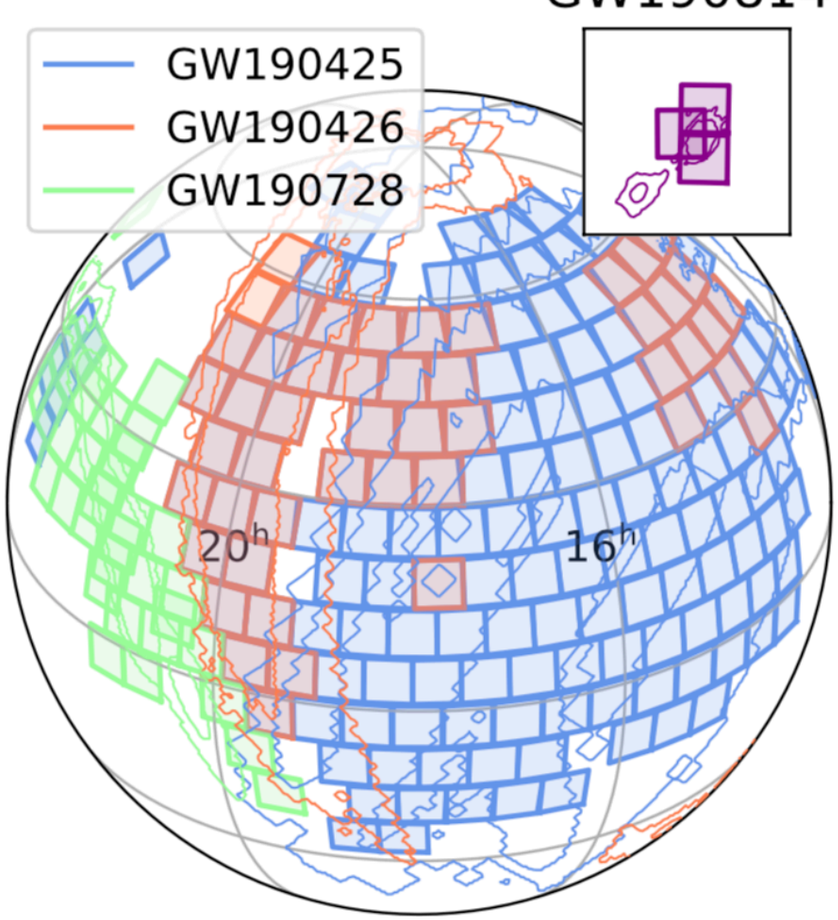

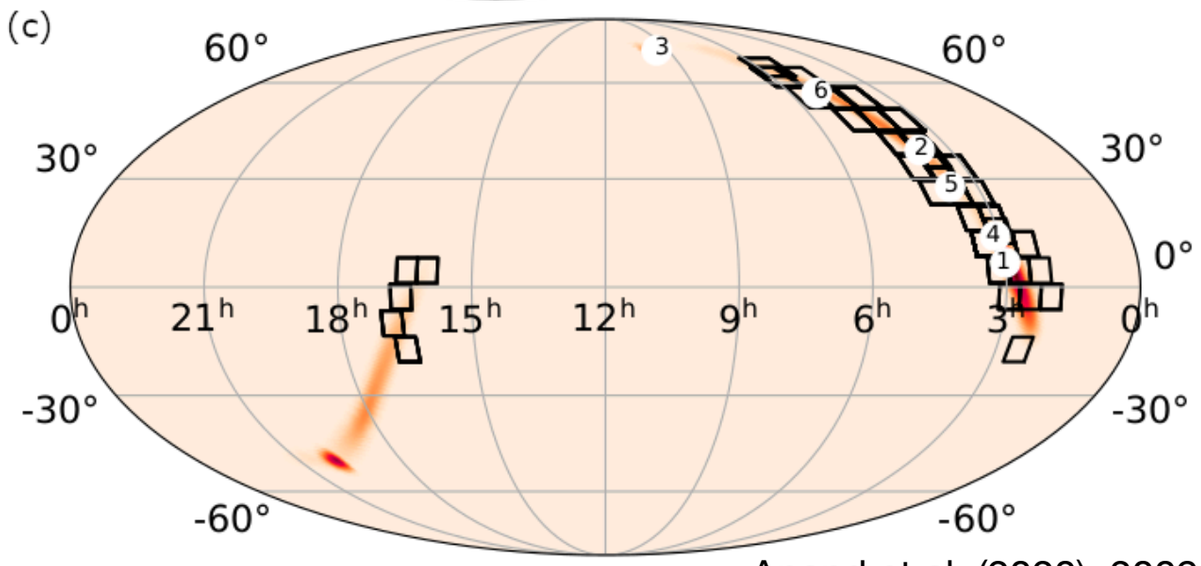

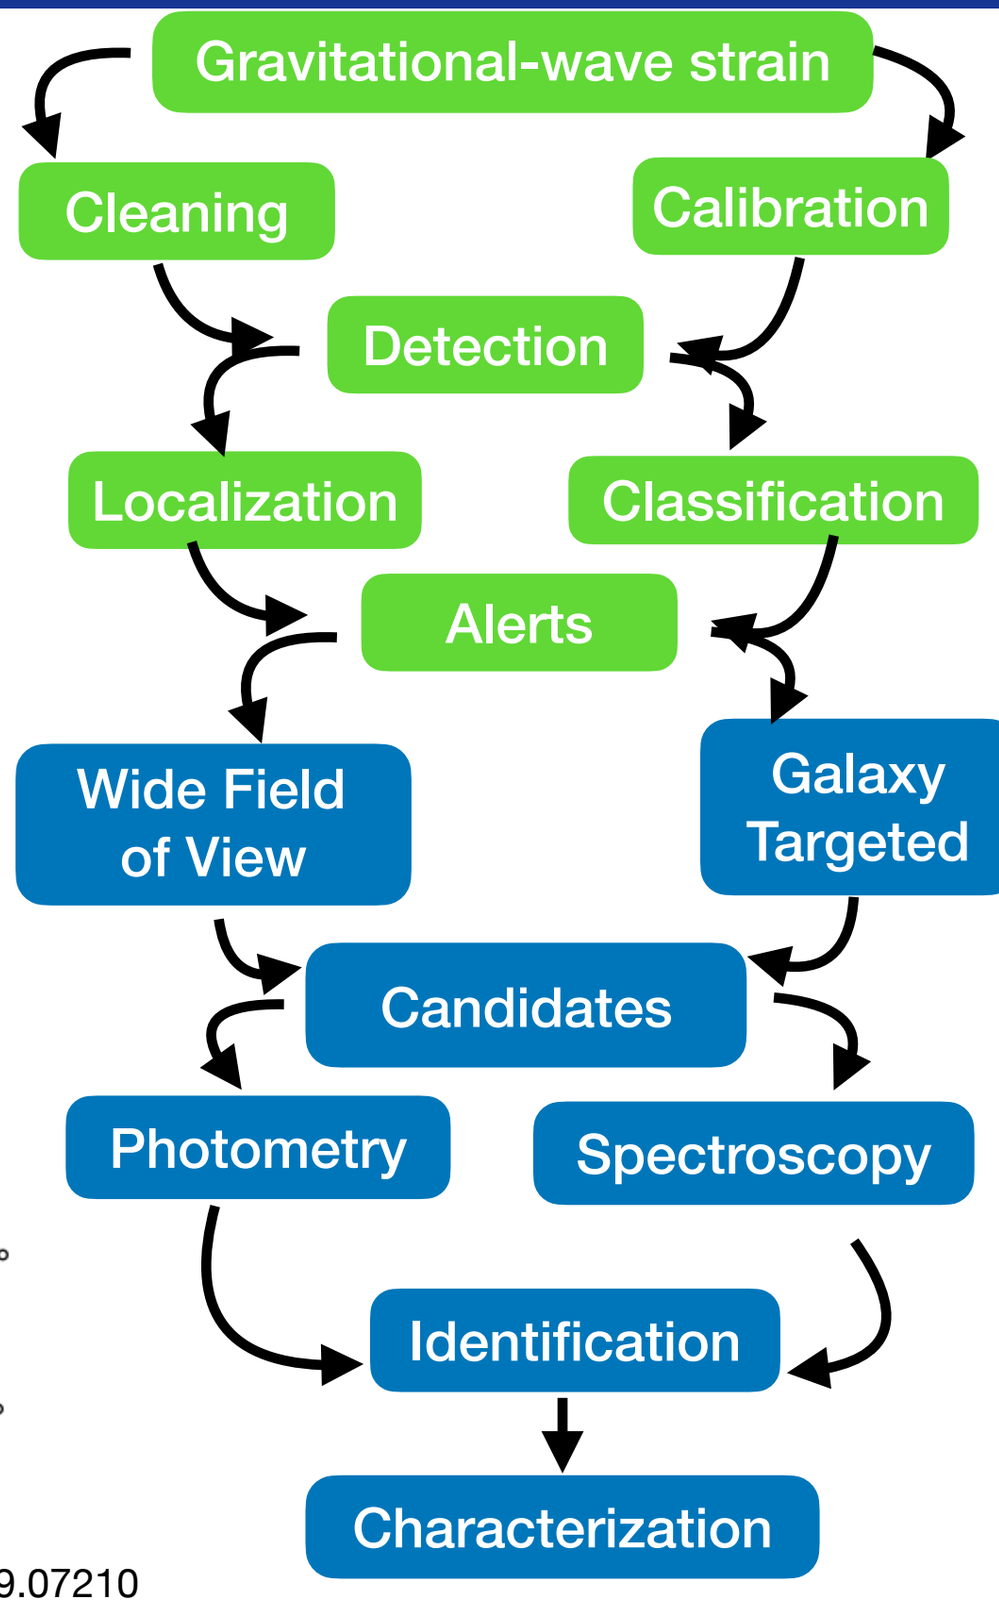

Anand et al. (2020): 2009.07210

Analyses

- **Online** (real-time; follow-up; alerts)
  - Modeled Compact Binary Coalescence
    - Pipelines: GstLAL, PyCBC, MBTA, SPIIR
    - Searches: AllSky, EarlyWarning, Sub-solar mass
  - Unmodeled searches
    - Pipelines: CWB, oLIB, MLy
    - Searches: AllSky, cwb-bbh
  - External coincidence search
    - GRB coincidences: Fermi, Swift, AGILE
    - Supernovae from SNEWS
    - Sub-threshold triggers via MoU
  - Real-time annotations:
    - Skymaps
    - Source-classification and properties
  - Online bayesian parameter estimation
- Offline (archival; catalog; rates and populations)

Alert Latencies

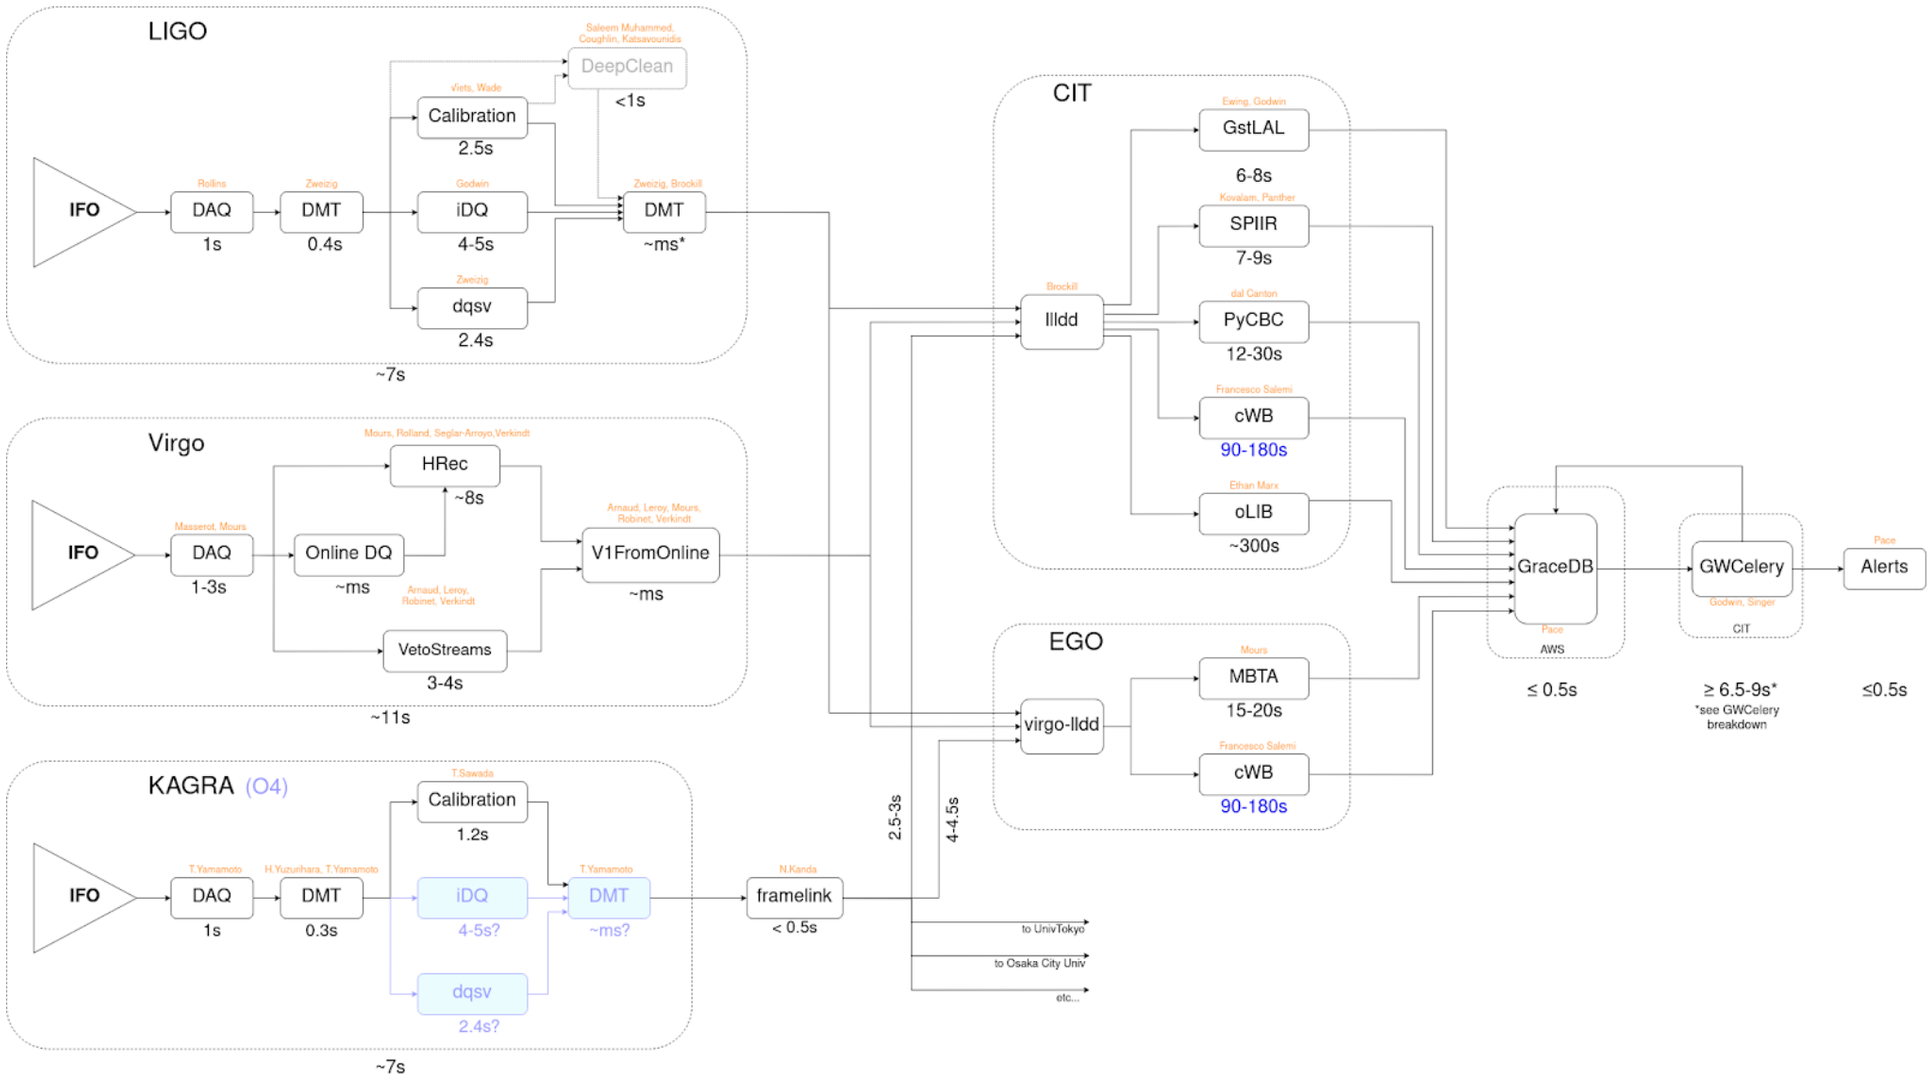

- Data Calibration ~3s
- Data transfer from sites (LLO/LHO) to CIT ~5s
- Pipeline analysis ~ 10 - 15s
- gwcelery + GraceDB alert ~ 8s (~2 seconds for skymaps)

Faster Searches

Latency Source	Latency (s)
Coalescence point exiting training kernel padding	0.25
Cropping corruption from whitening filter	0.50
Cropping corruption from resampling to 2048 Hz	1.0
Integrating network output	1.0
Reading data and transferring to GPU	$1.03^{+0.06}_{-0.05} \times 10^{-2}$
Estimating PSD and whitening	$8.77^{+1.35}_{-0.31} \times 10^{-4}$
Performing inference on whitened data	$9.63^{+0.38}_{-0.32} \times 10^{-3}$
Integrating and aggregating network output	$3.42^{+0.02}_{-0.01} \times 10^{-1}$
Identifying candidate events in integrated output	$1.40^{+0.62}_{-0.43} \times 10^{-4}$
Total	$3.114^{0.006}_{0.001}$

Marx, Benoit, et al. 2403.18661

NB: For the items listed in the upper section this table, the latency does not come from performing the computation, but rather from needing to wait for the data to exist before the action can occur…

Ethan Marx, MIT

Will Benoit, UMN

Discussions ongoing to put a GPU at each site (to bypass data transfer) and perform inference for ML detection and then use streaming ML parameter estimation models for these detections. **< 10 second latencies possible!**

Faster Parameter Estimation

AMPLFI

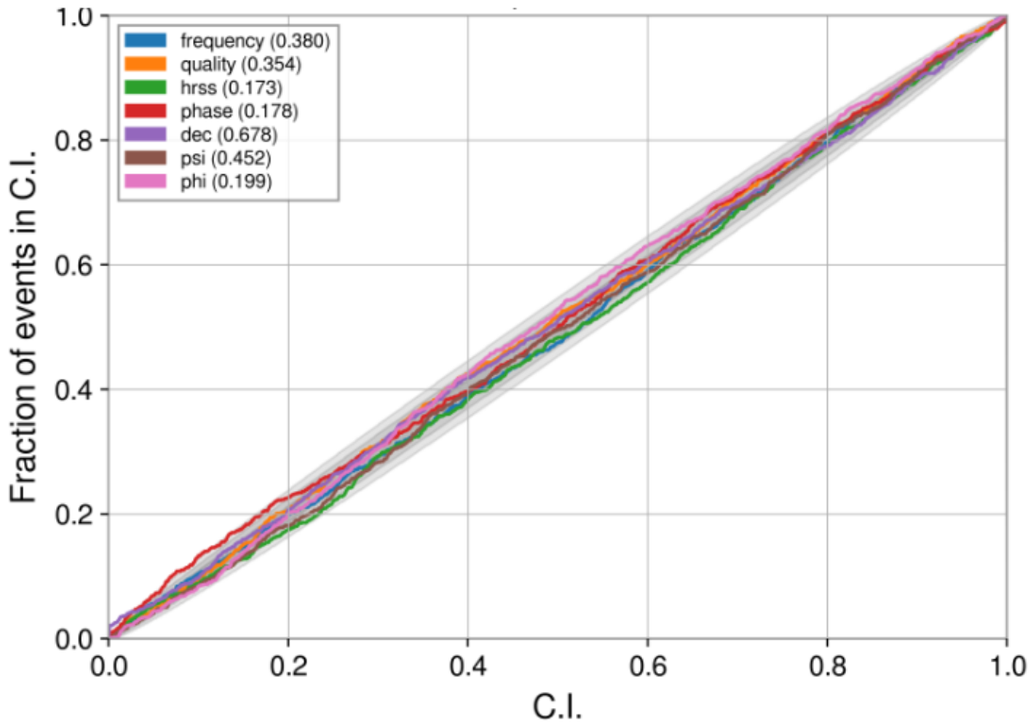

Chatterjee et al. NeurIPS 2023 ML and Physical Sciences workshop

Dingo

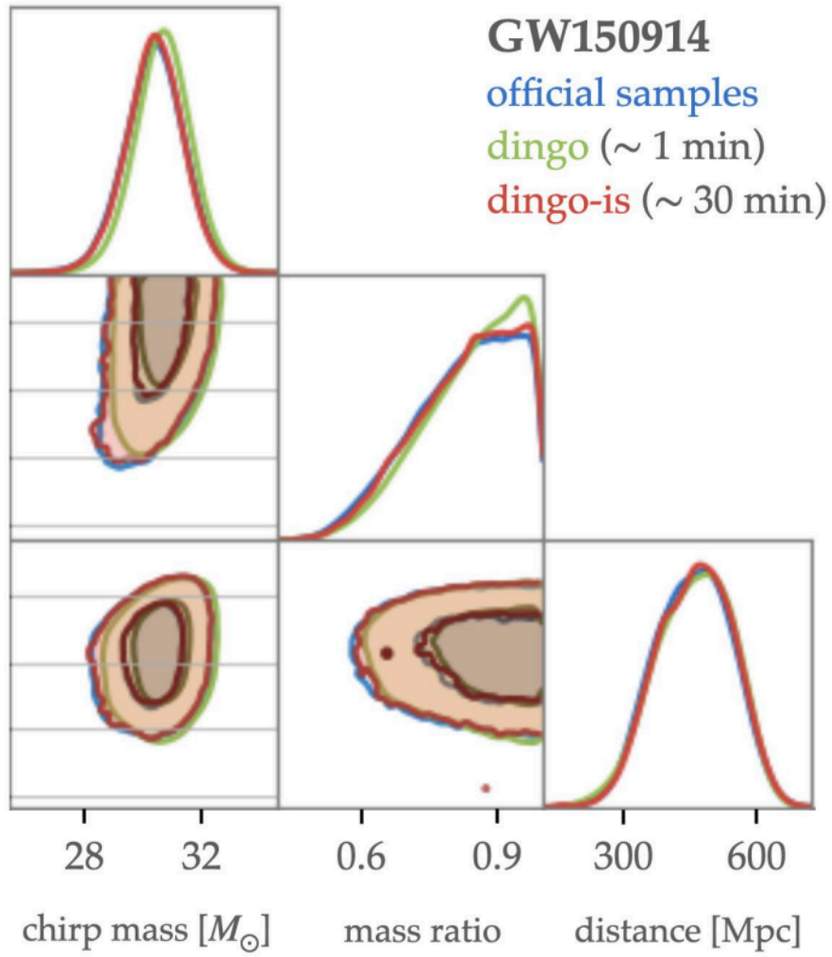

Dax et al. Phys.Rev.Lett. 127, 241103 (2021)

Promising initial results on BNS length waveforms, but needs work and review

More Unknown Unknowns

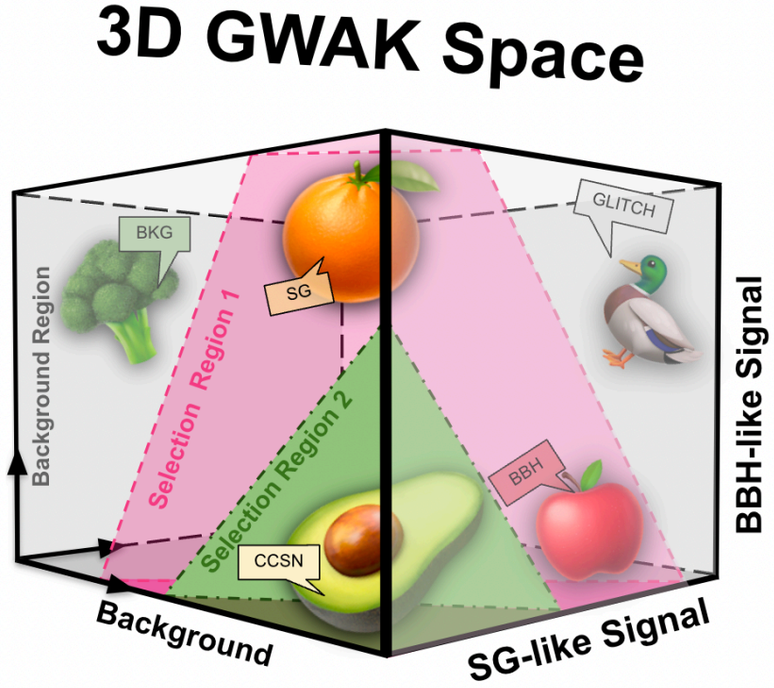

Raikman *et al* 2024 *Mach. Learn.: Sci. Technol.* **5** 025020

Semi-supervised approach to gravitational-wave anomaly detection

MLy PIPELINE

Skliris *et al*, *2009.14611*

CNN-based search running in real time

Need to encourage and support pipelines capable of detection of non-CBC signals!

The Time-Domain Astronomy Ecosystem

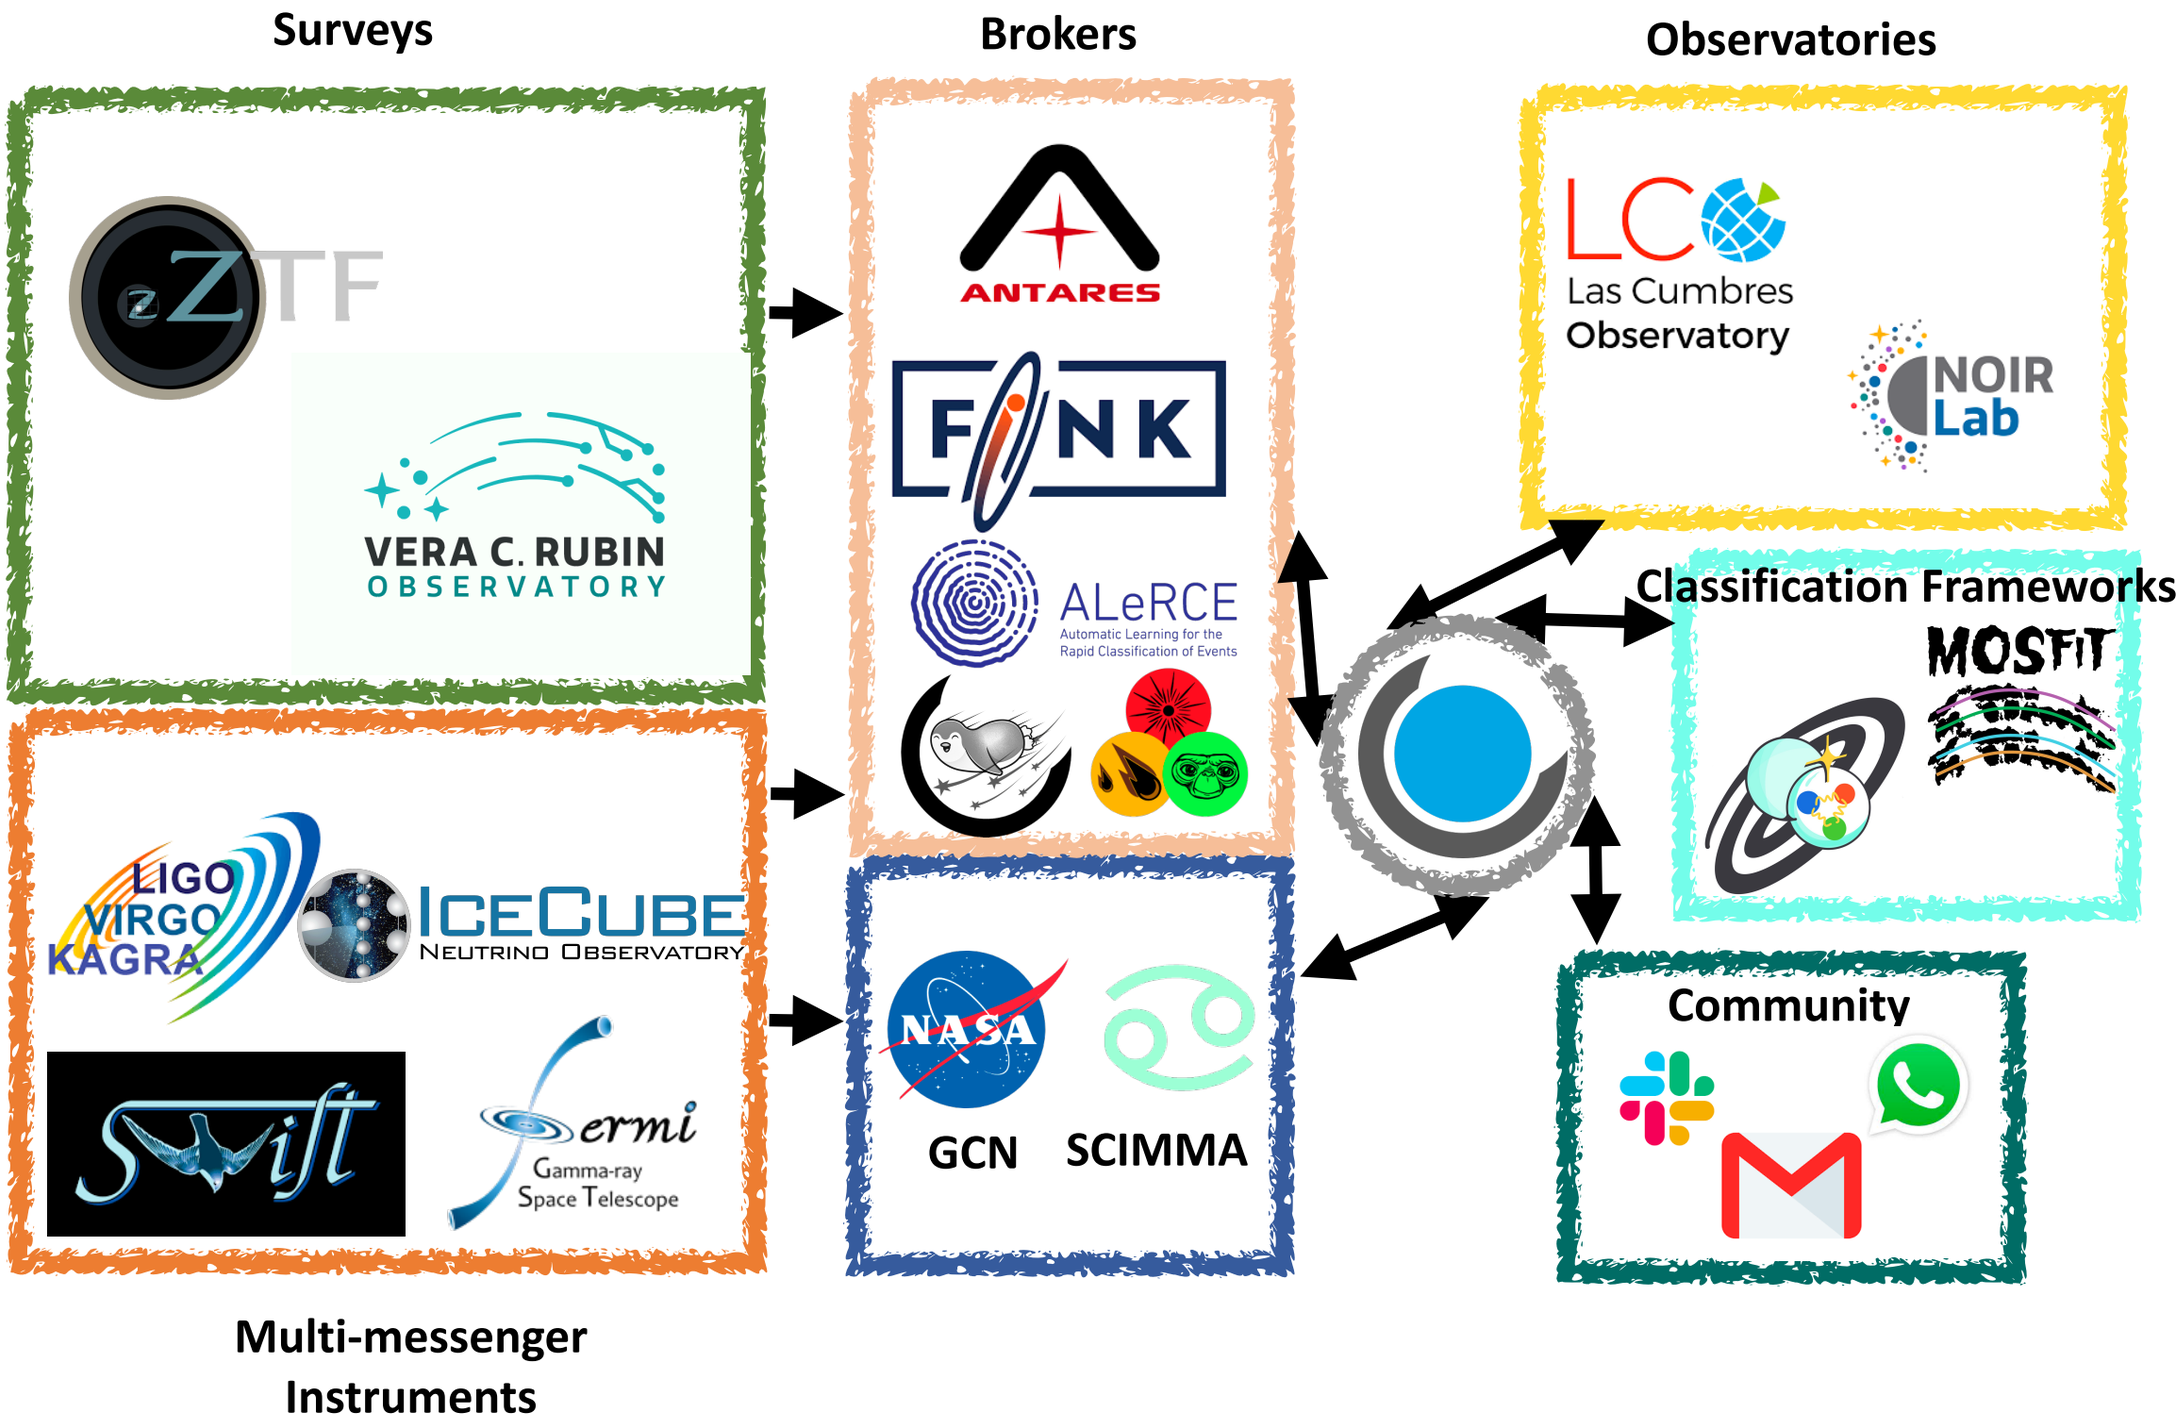

Let’s Discuss

What do you think are the minimum ingredients to make a good real time classification based on your experience?

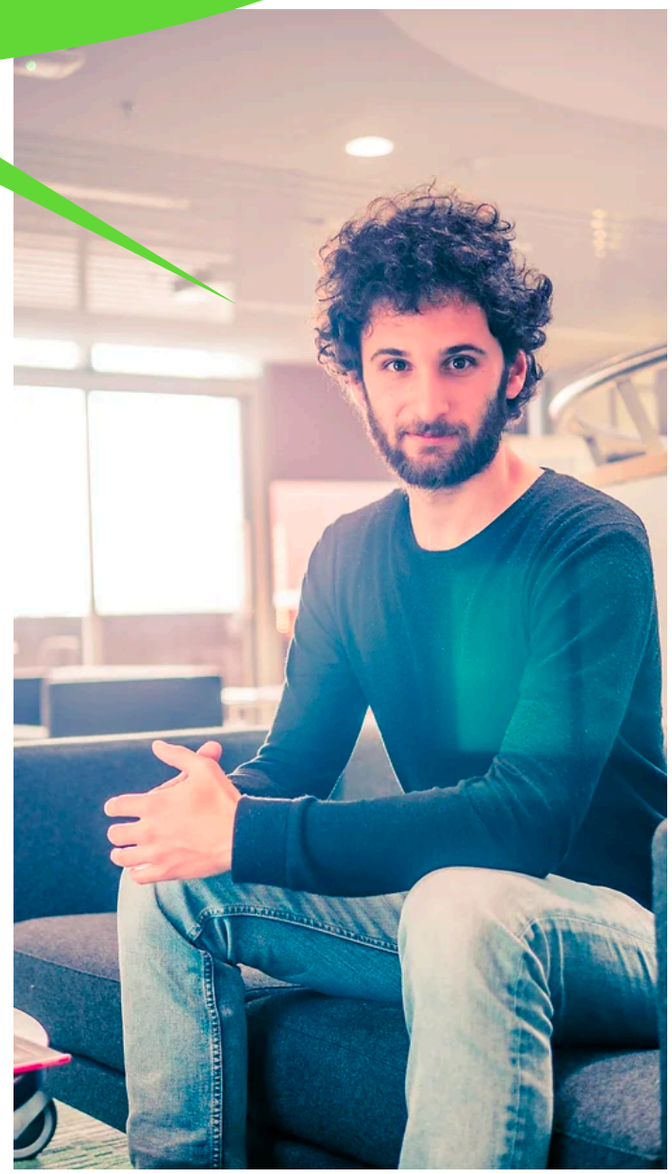

Let’s Discuss

What do you think are the minimum ingredients to make a good real time classification based on your experience?

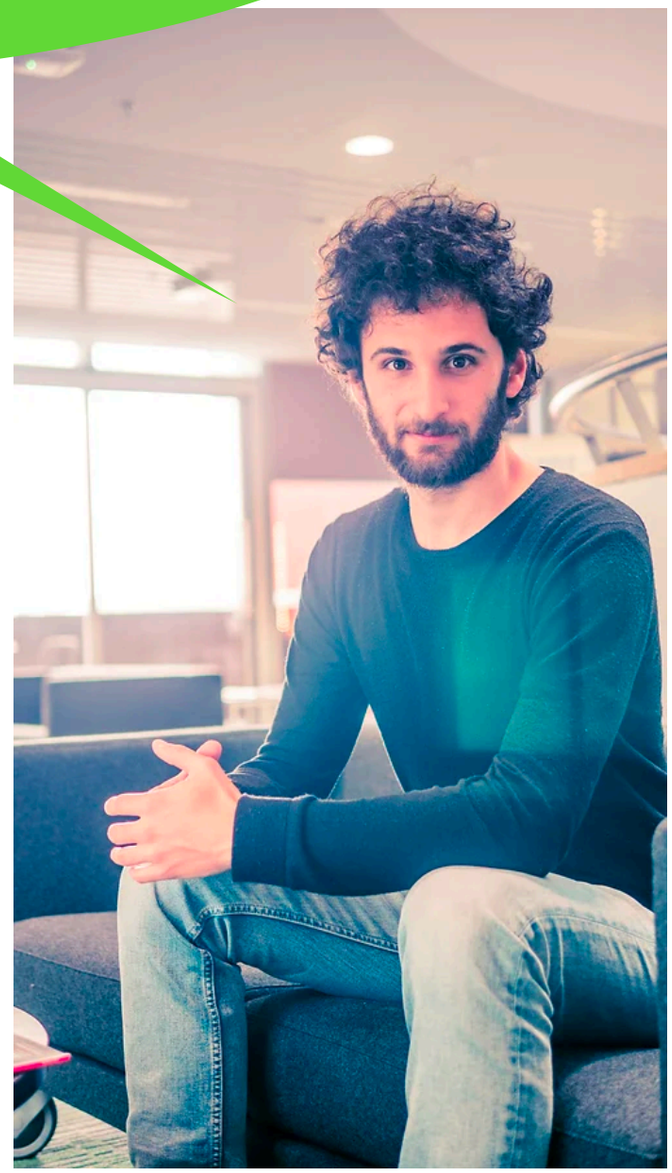

Let’s get on the same page:

- The answer is highly dependent on the science case
- The more information the better
  but it can be expensive (time, resources)

Let's Discuss

What do you think are the minimum ingredients to make a good real time classification based on your experience?

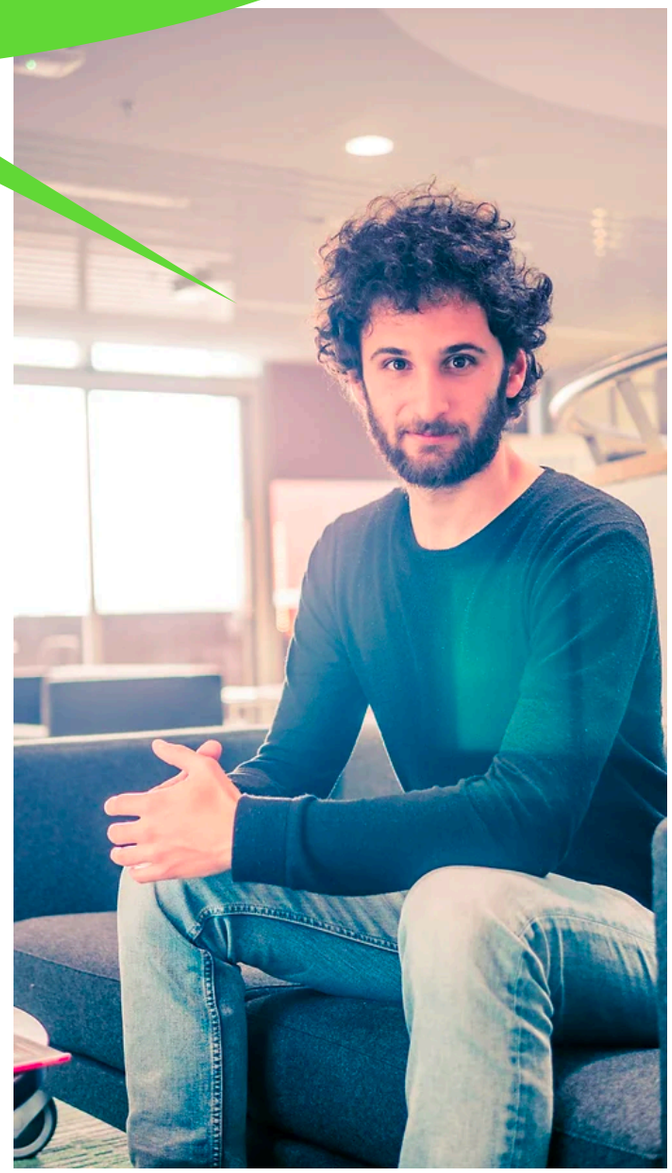

The Basics

- Transient being real
- Extragalactic: high b_{Gal} + deep archival images + star/galaxy separ.
- Fade rate ($\Delta m/\Delta t$) — importance of the cadence choice

Always welcome: same-night color, history, proximity to a galaxy, photoz

Kilonovae - Hard to find

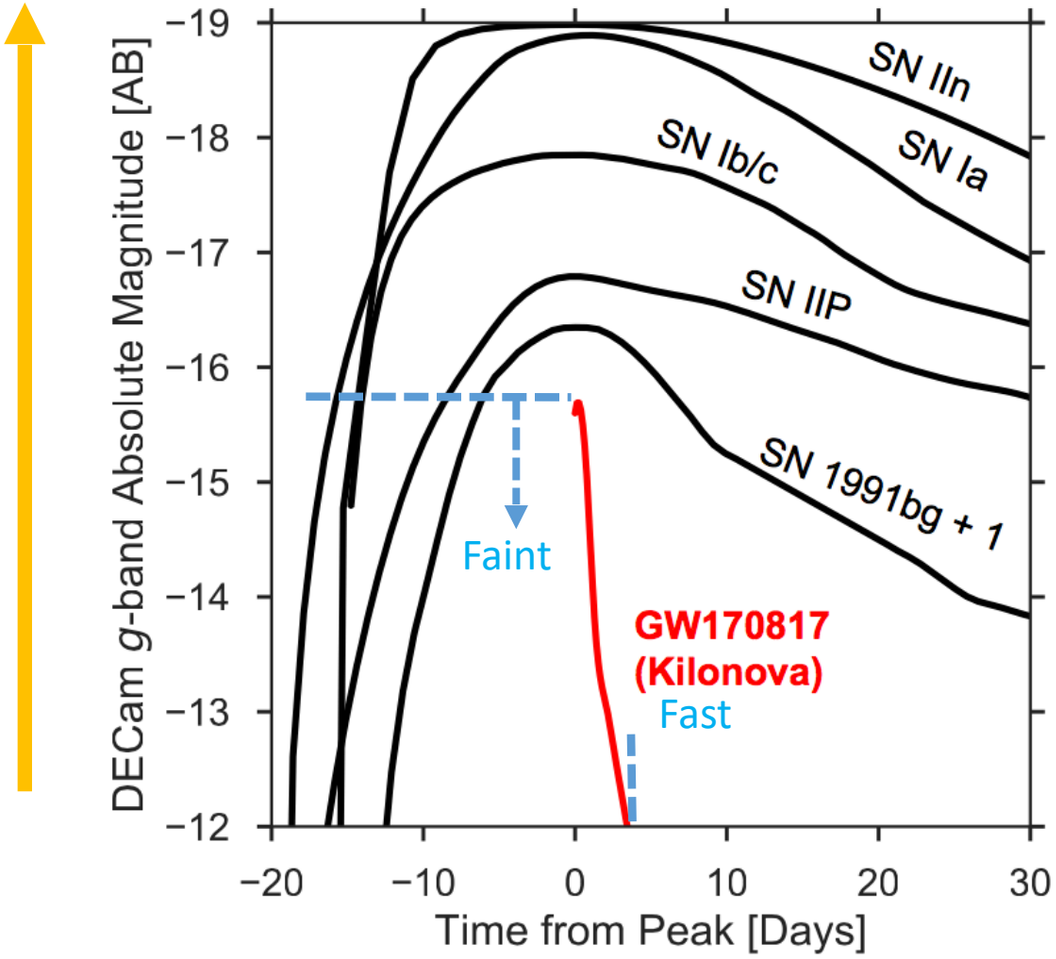

modified from
Andreoni+2018

Faint + Fast = hard to catch!

ZTFReST for fast transient discovery

growth-astro/ztfrest

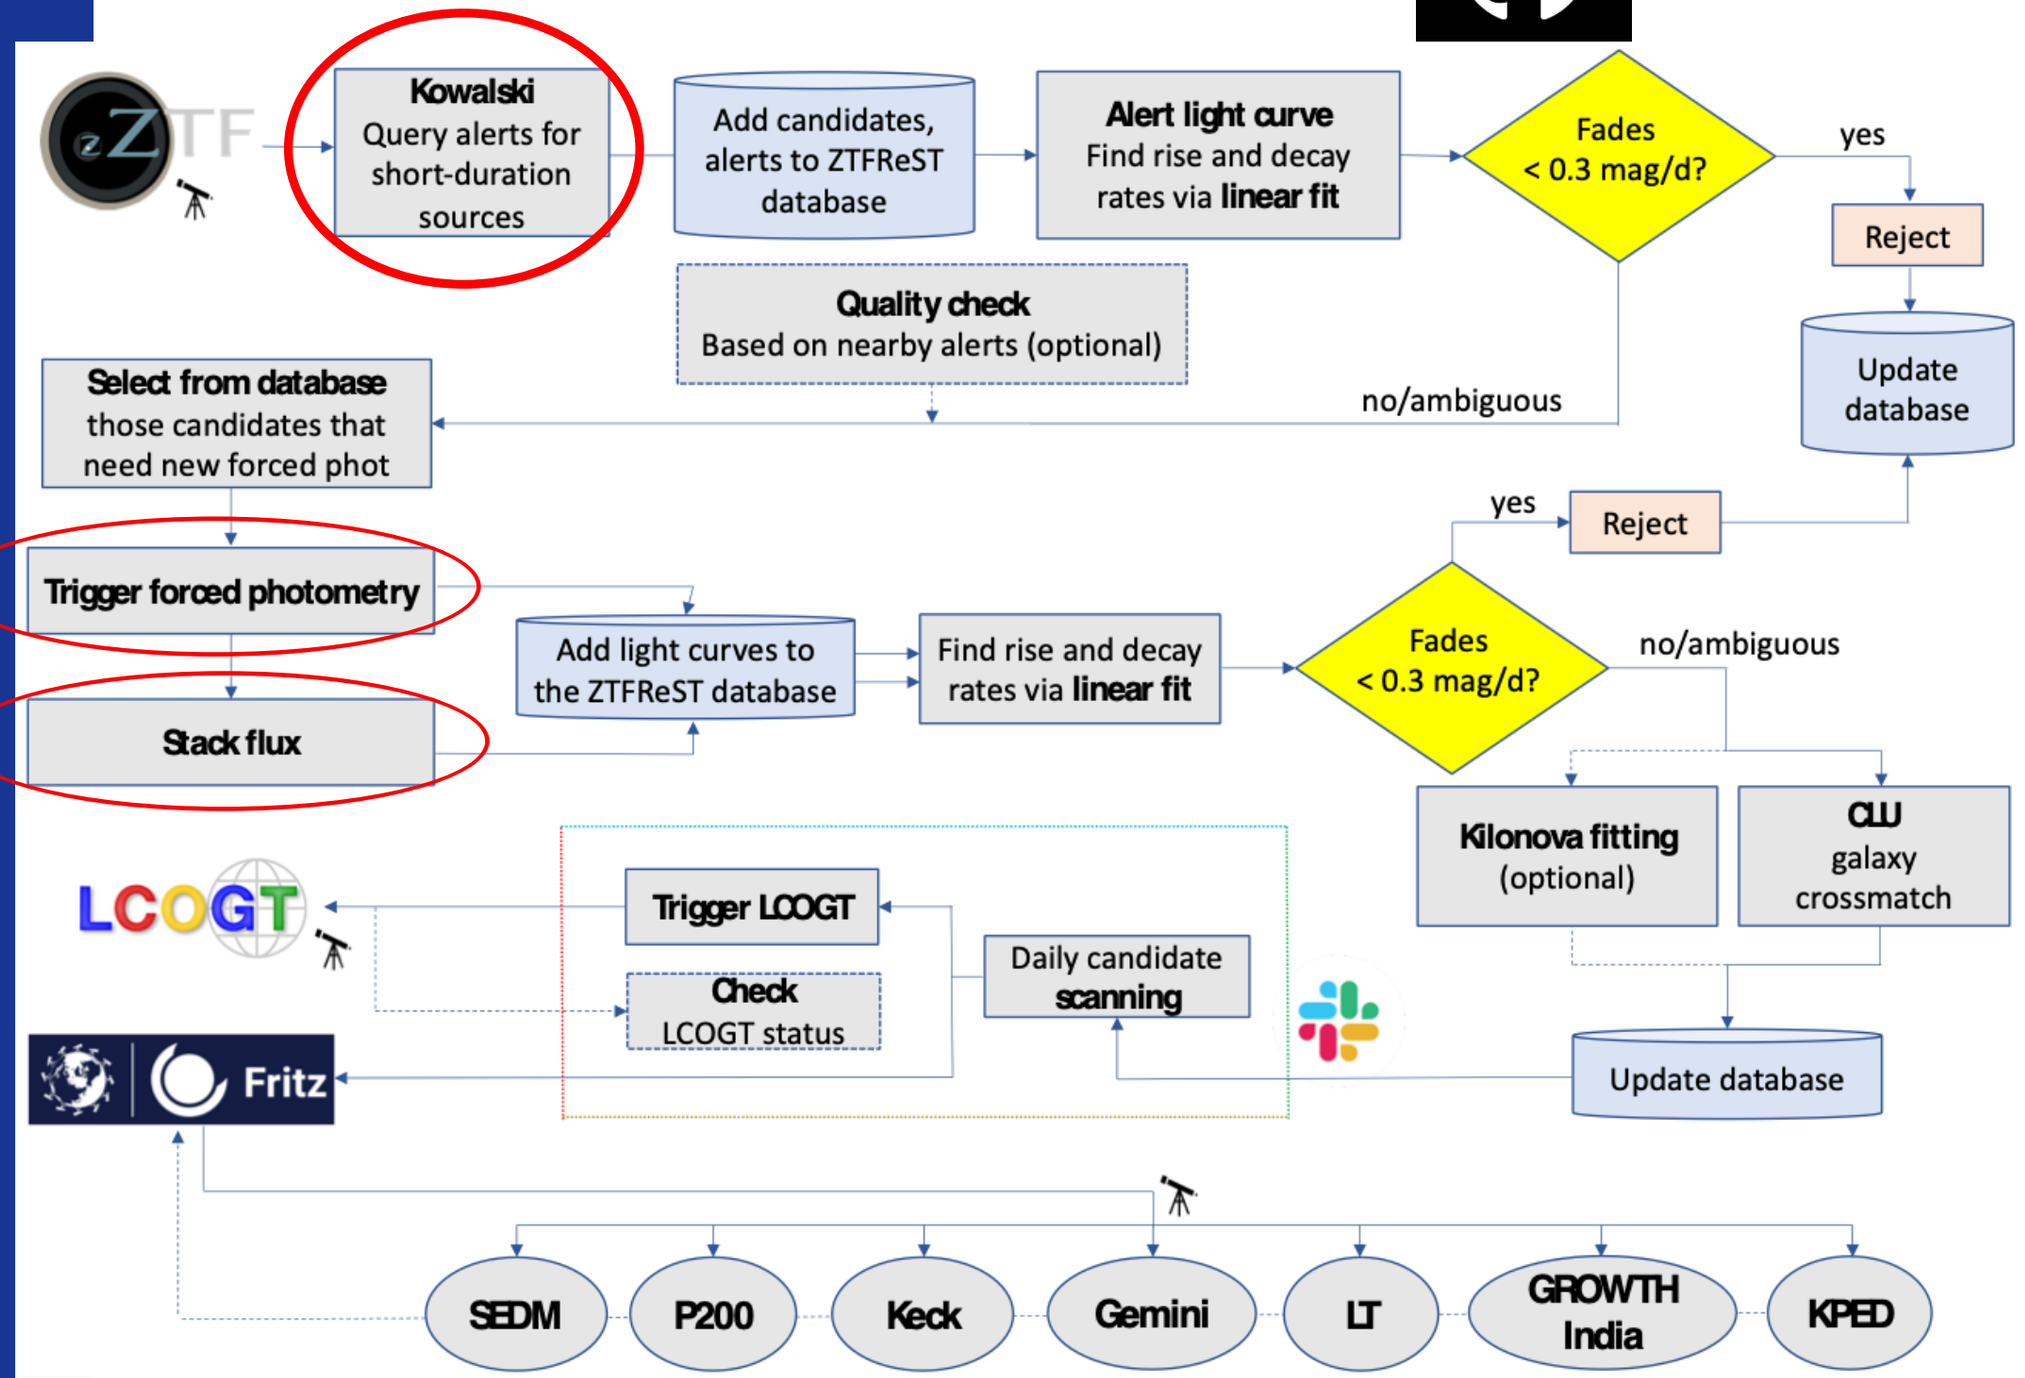

[Andreoni, Coughlin+2021, ApJ]

Real Time Discoveries

Igor Andreoni 11:05 AM
@channel Are we re-starting with a bang?!?! https://fritz.science/source/ZTF22aaajecp

LS photoz of the closest galaxy: z = 1.201284 +- 0.176194

Michael Coughlin 11:08 AM
WOW!

Igor Andreoni 11:08 AM
This looks a lot like an afterglow to me.. but caught on the rise would be crazy..

Daniel Perley 11:17 AM
I'll put it in for LT tonight

Igor Andreoni 11:18 AM
Thanks Dan

I am reporting to TNS and putting together a short astronote

Michael Coughlin 11:19 AM
@sganand Can you put in LCO?

Anna Ho 11:40 AM
Interesting!! I have a Gemini ToO program this semest[illegible] you would like me to trigger it.

↓ Latest messages

[Andreoni, Coughlin+2022, Nature]

Real Time Discoveries

Igor Andreoni 11:05 AM
@channel Are we re-starting with a bang?!?! https://fritz.science/source/ZTF22aaajecp

LS photoz of the closest galaxy: z = 1.201284 +- 0.176194

Michael Coughlin 11:08 AM
WOW!

Igor Andreoni 11:08 AM
This looks a lot like an afterglow to me.. but caught on the rise would be crazy..

Daniel Perley 11:17 AM
I'll put it in for LT tonight

Igor Andreoni 11:18 AM
Thanks Dan

I am reporting to TNS and putting together a short astronote

Michael Coughlin 11:19 AM
@sganand Can you put in LCO?

Anna Ho 11:40 AM
Interesting!! I have a Gemini ToO program this semester, so let me know if you would like me to trigger it.

↓ Latest messages

My favorite part

wrong galaxy, correct photoz!

Flagged as special

Immediate public reporting

Rapid follow-up

[Andreoni, Coughlin+2022, Nature]

We have found many transients…

Andreoni & Coughlin et al. (2021), ApJ, 918, 2, 63

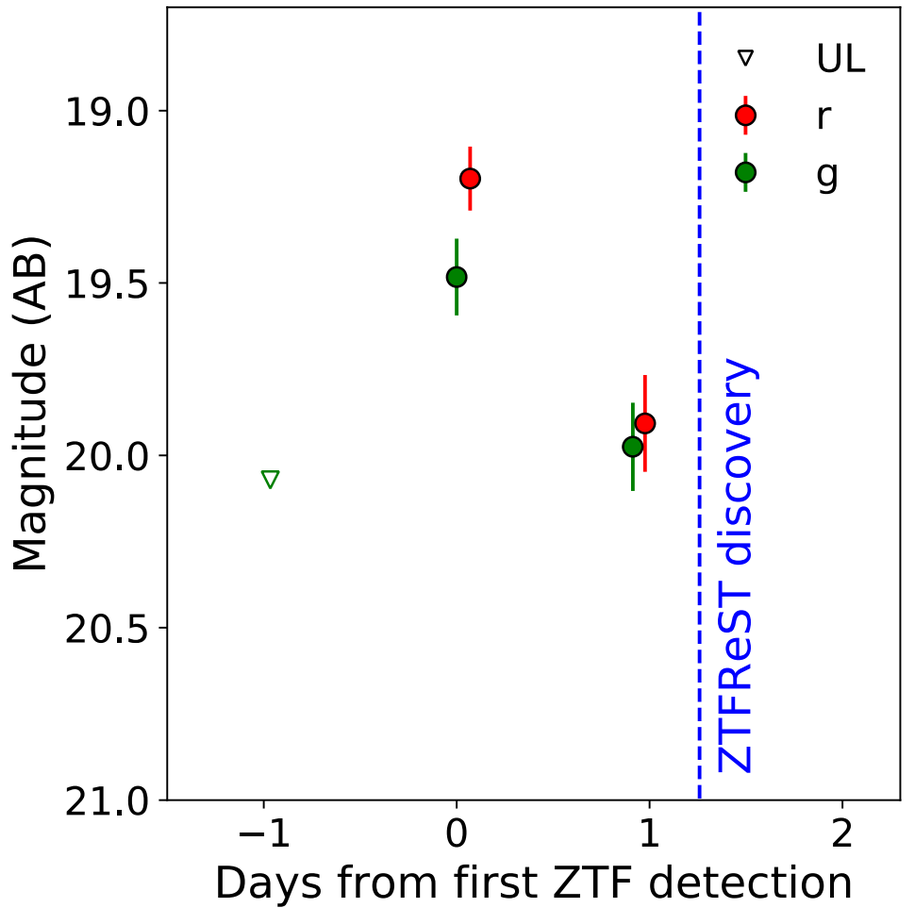

Near real-time implementation of the search methods used in Andreoni et al. (2020d)

Supernova shock cooling
~ a dozen

Serendipitous GRB afterglows
- 5 with GRB association (long)
- four confirmed, un-triggered afterglows w/o GRB association
see also Ho+2020,2021,2022; Andreoni+2020d

Kilonovae
Still waiting…

AT2022cmc
a jetted TDE!

Including some we were not expecting…

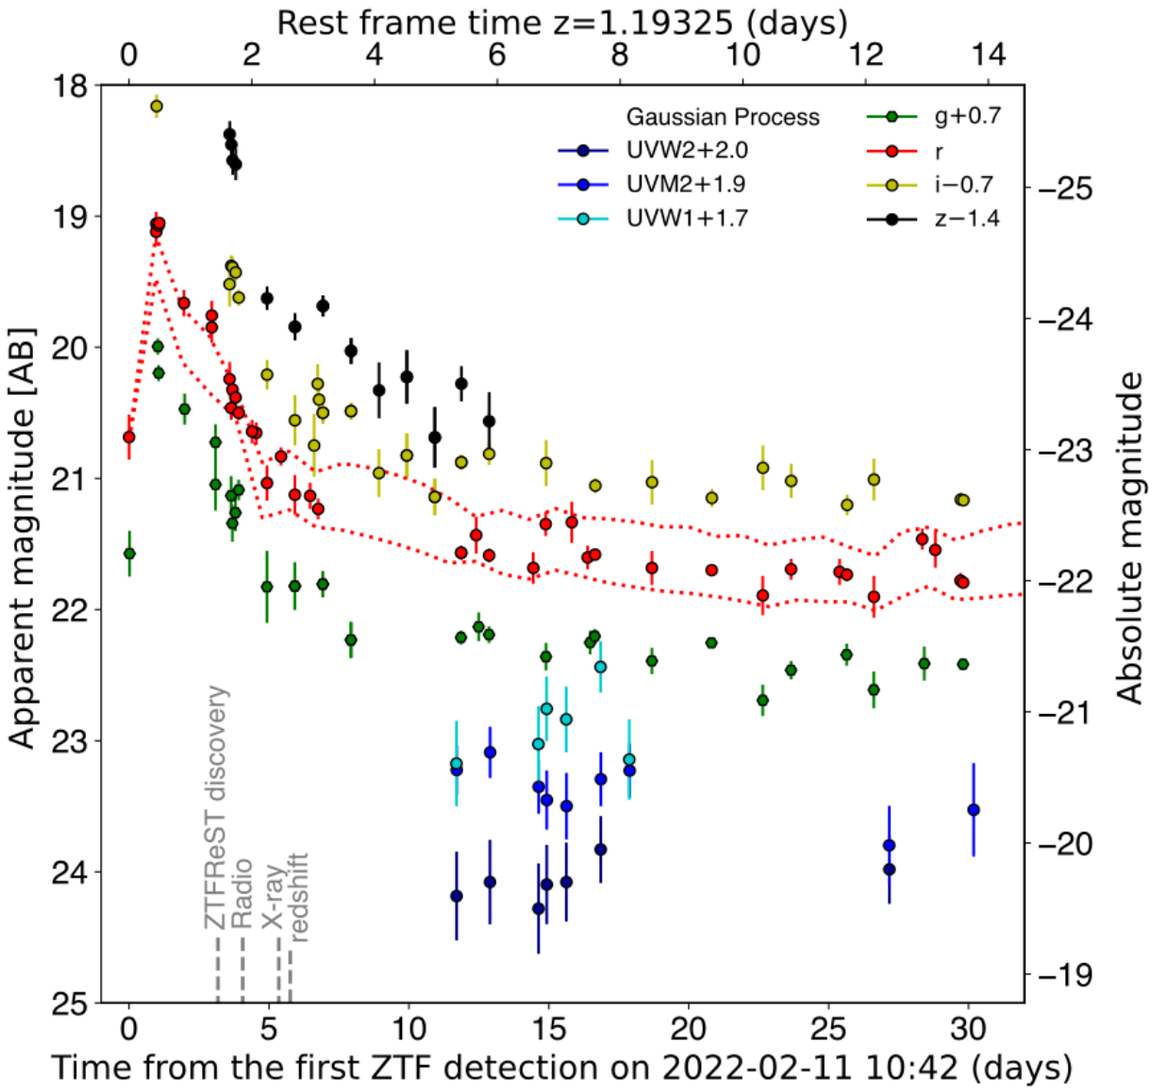

A couple of fun facts:

- First relativistic TDE identified in near real time by an optical survey
- Furthest TDE ever found

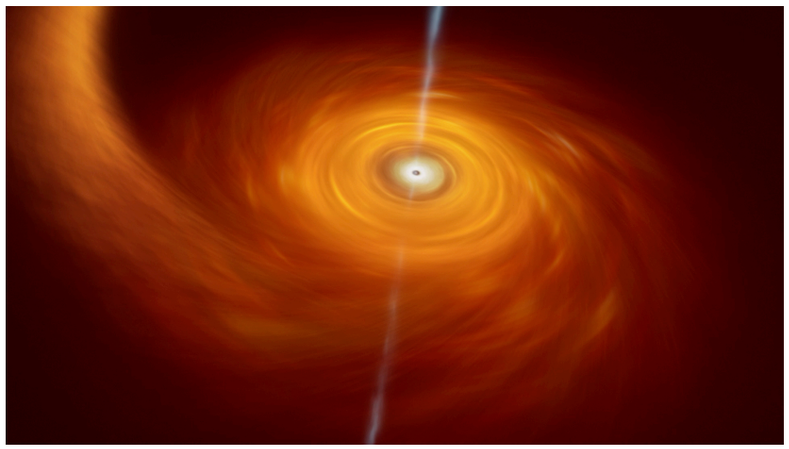

Andreoni & Coughlin et al. (2022), Nature, 612, 7940

And strongly constrained the rates of those we were…

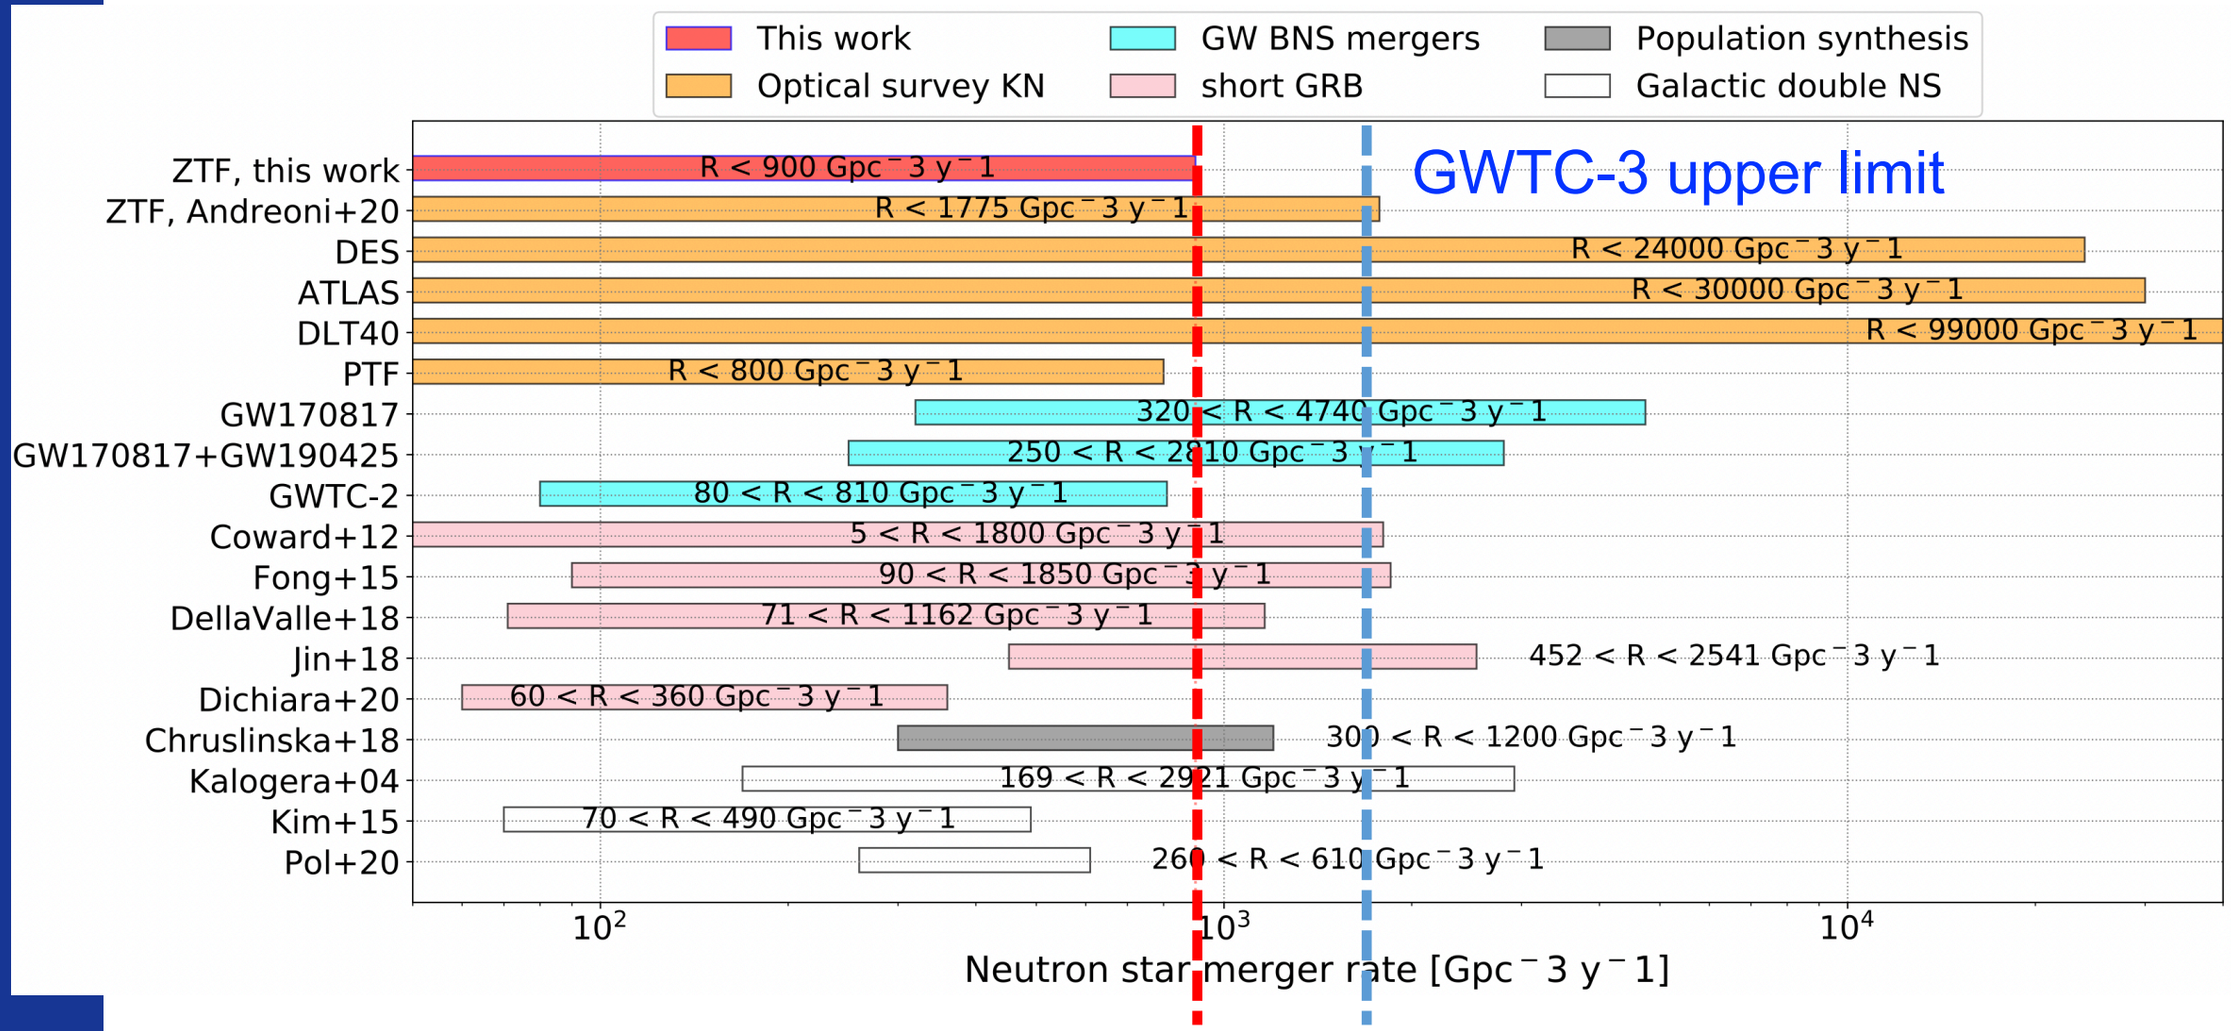

Andreoni & Coughlin et al. (2021), ApJ, 918, 2, 63
Model grid in Andreoni et al. (2020d), ApJ, 904, 2, 155

ZTF constrained the rate of GW170817-like kilonovae to be $R < 900\ \mathrm{Gpc^{-3}\ y^{-1}}$

[Andreoni, Coughlin+2021, ApJ]

And what about Rubin?

Assuming a rate of 800 $Gpc^{-3}y^{-1}$ and a uniform luminosity function,
~350 GW170817-like kilonovae are expected to be present in Rubin data,
but only 3 – 32 might be identified using ZTFReST-like algorithms

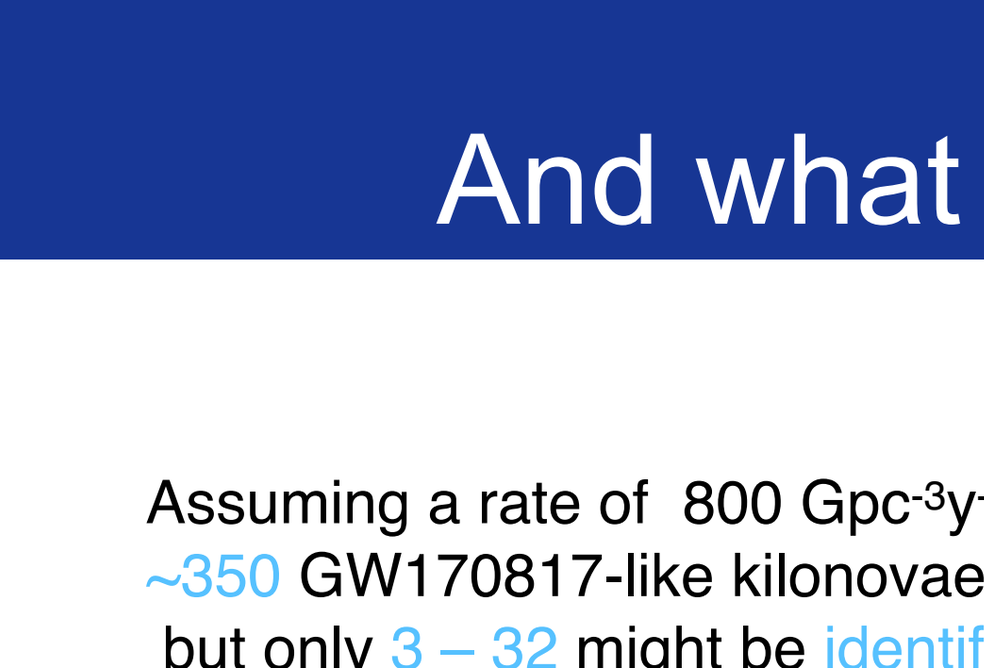

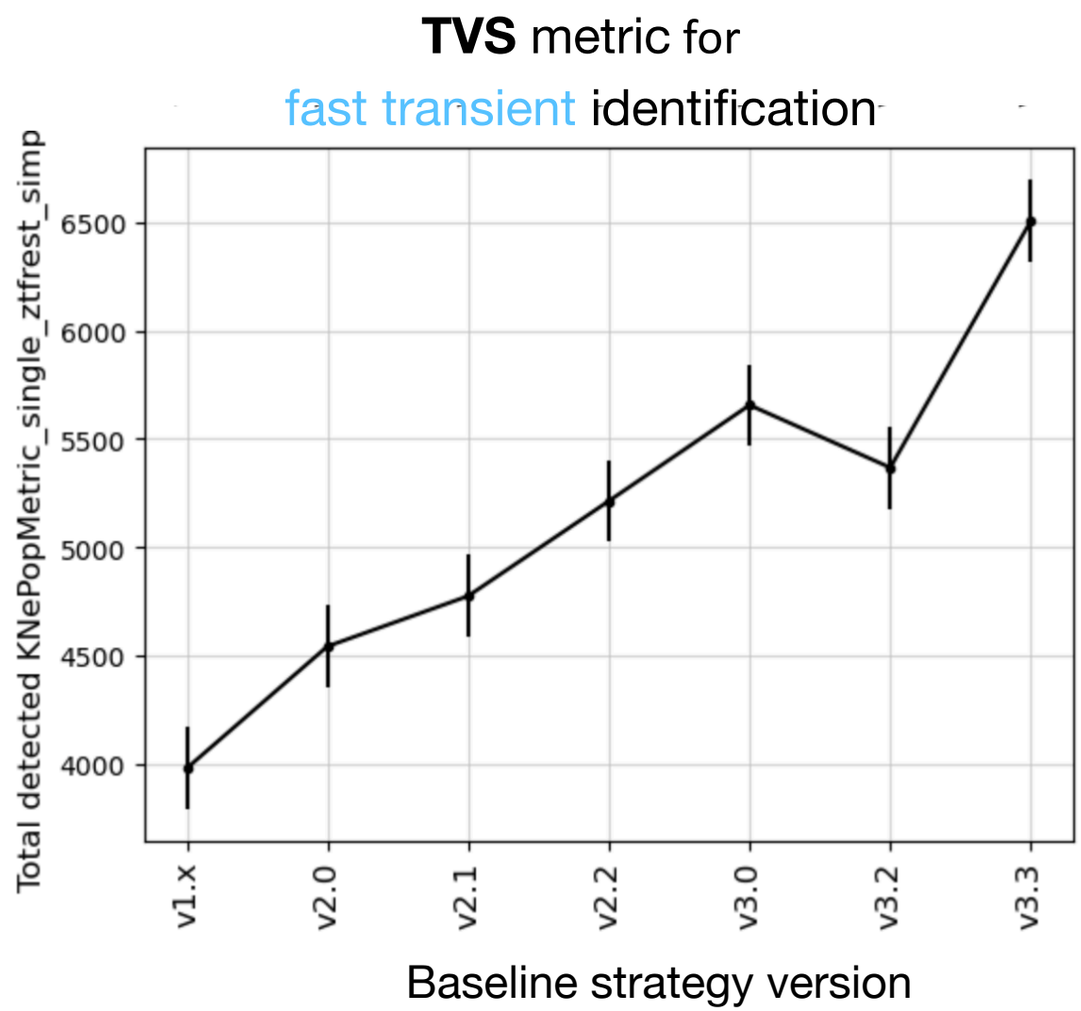

[Andreoni et al 2022, ApJS]

And what about Rubin ToOs?

"We expect to discover ~10 counterparts to NS–NS mergers and probe the existence of EM counterparts to ~15 NS–BH mergers per year during O5" [Andreoni et al 2022b, ApJS]

Update: *"6 BNS* and up to *2 NS-BH* events per year with $\Omega < 100$ square degrees".

https://lssttooworkshop.github.io/images/Rubin_2024_ToO_workshop_final_report.pdf

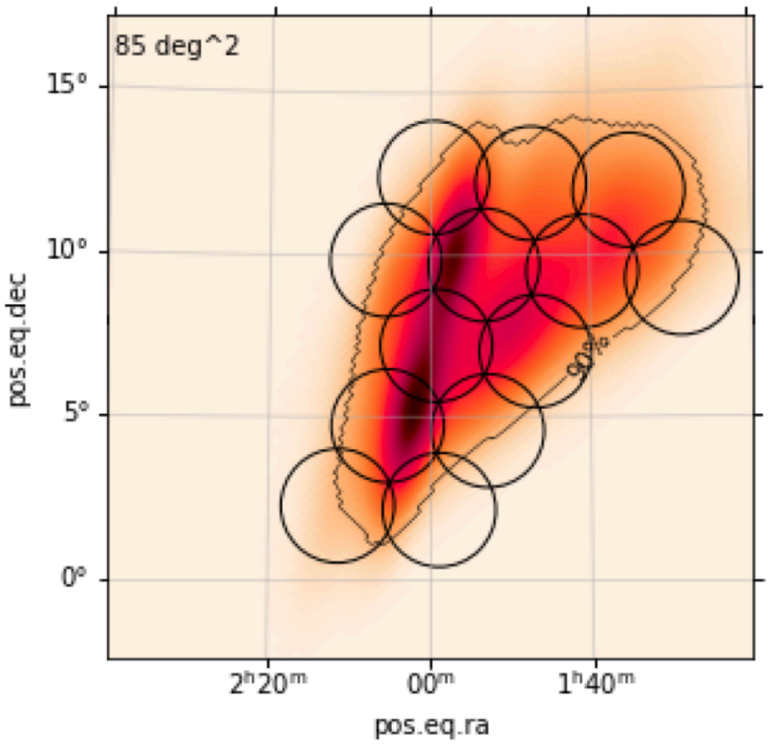

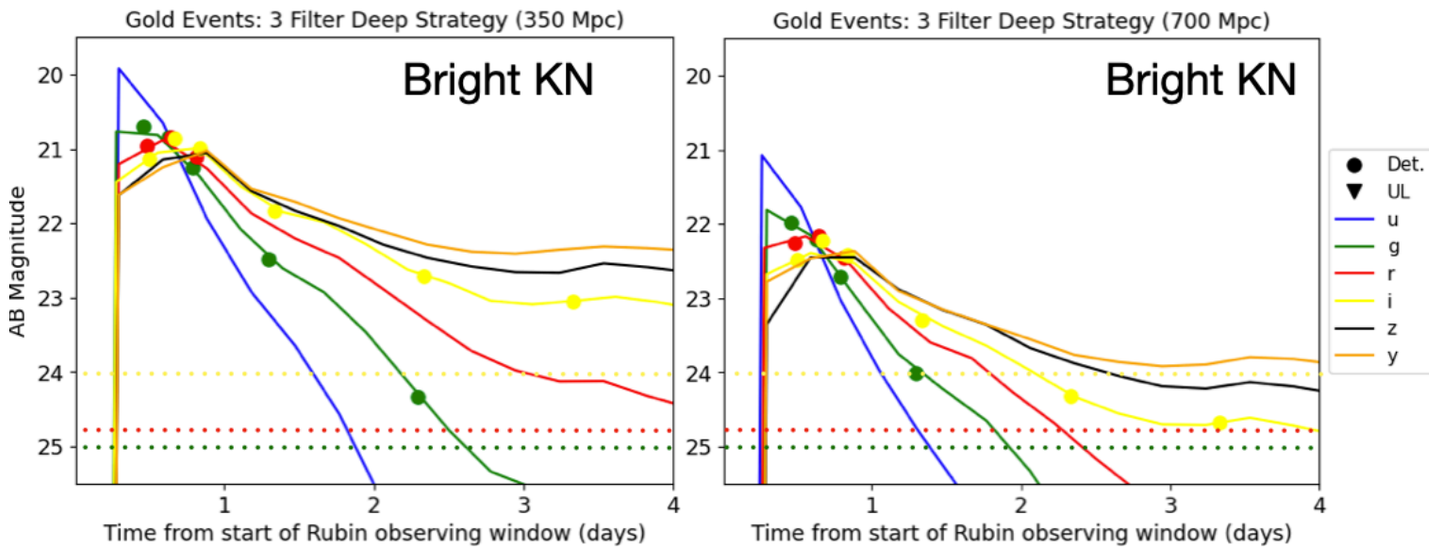

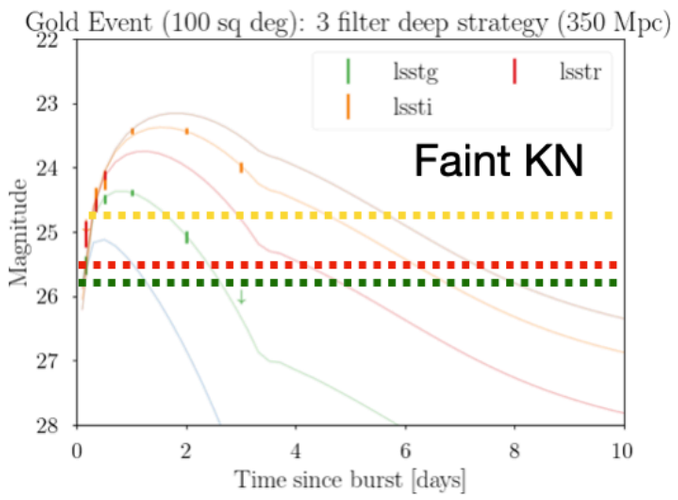

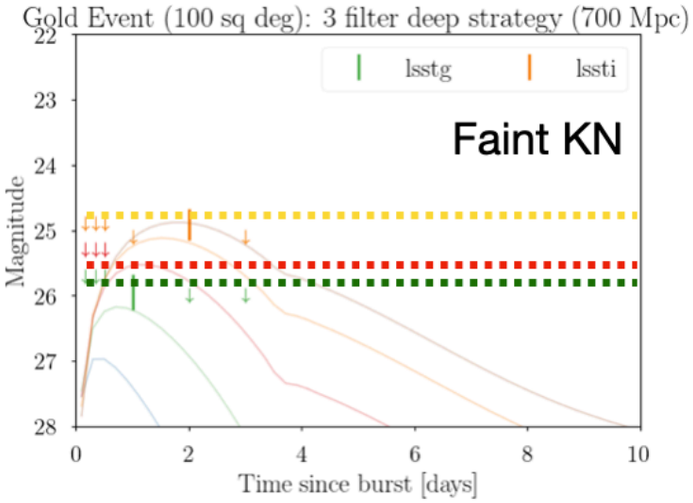

A nuclear physics and multi-messenger framework (NMMA)

github.com/nuclear-multimessenger-astronomy

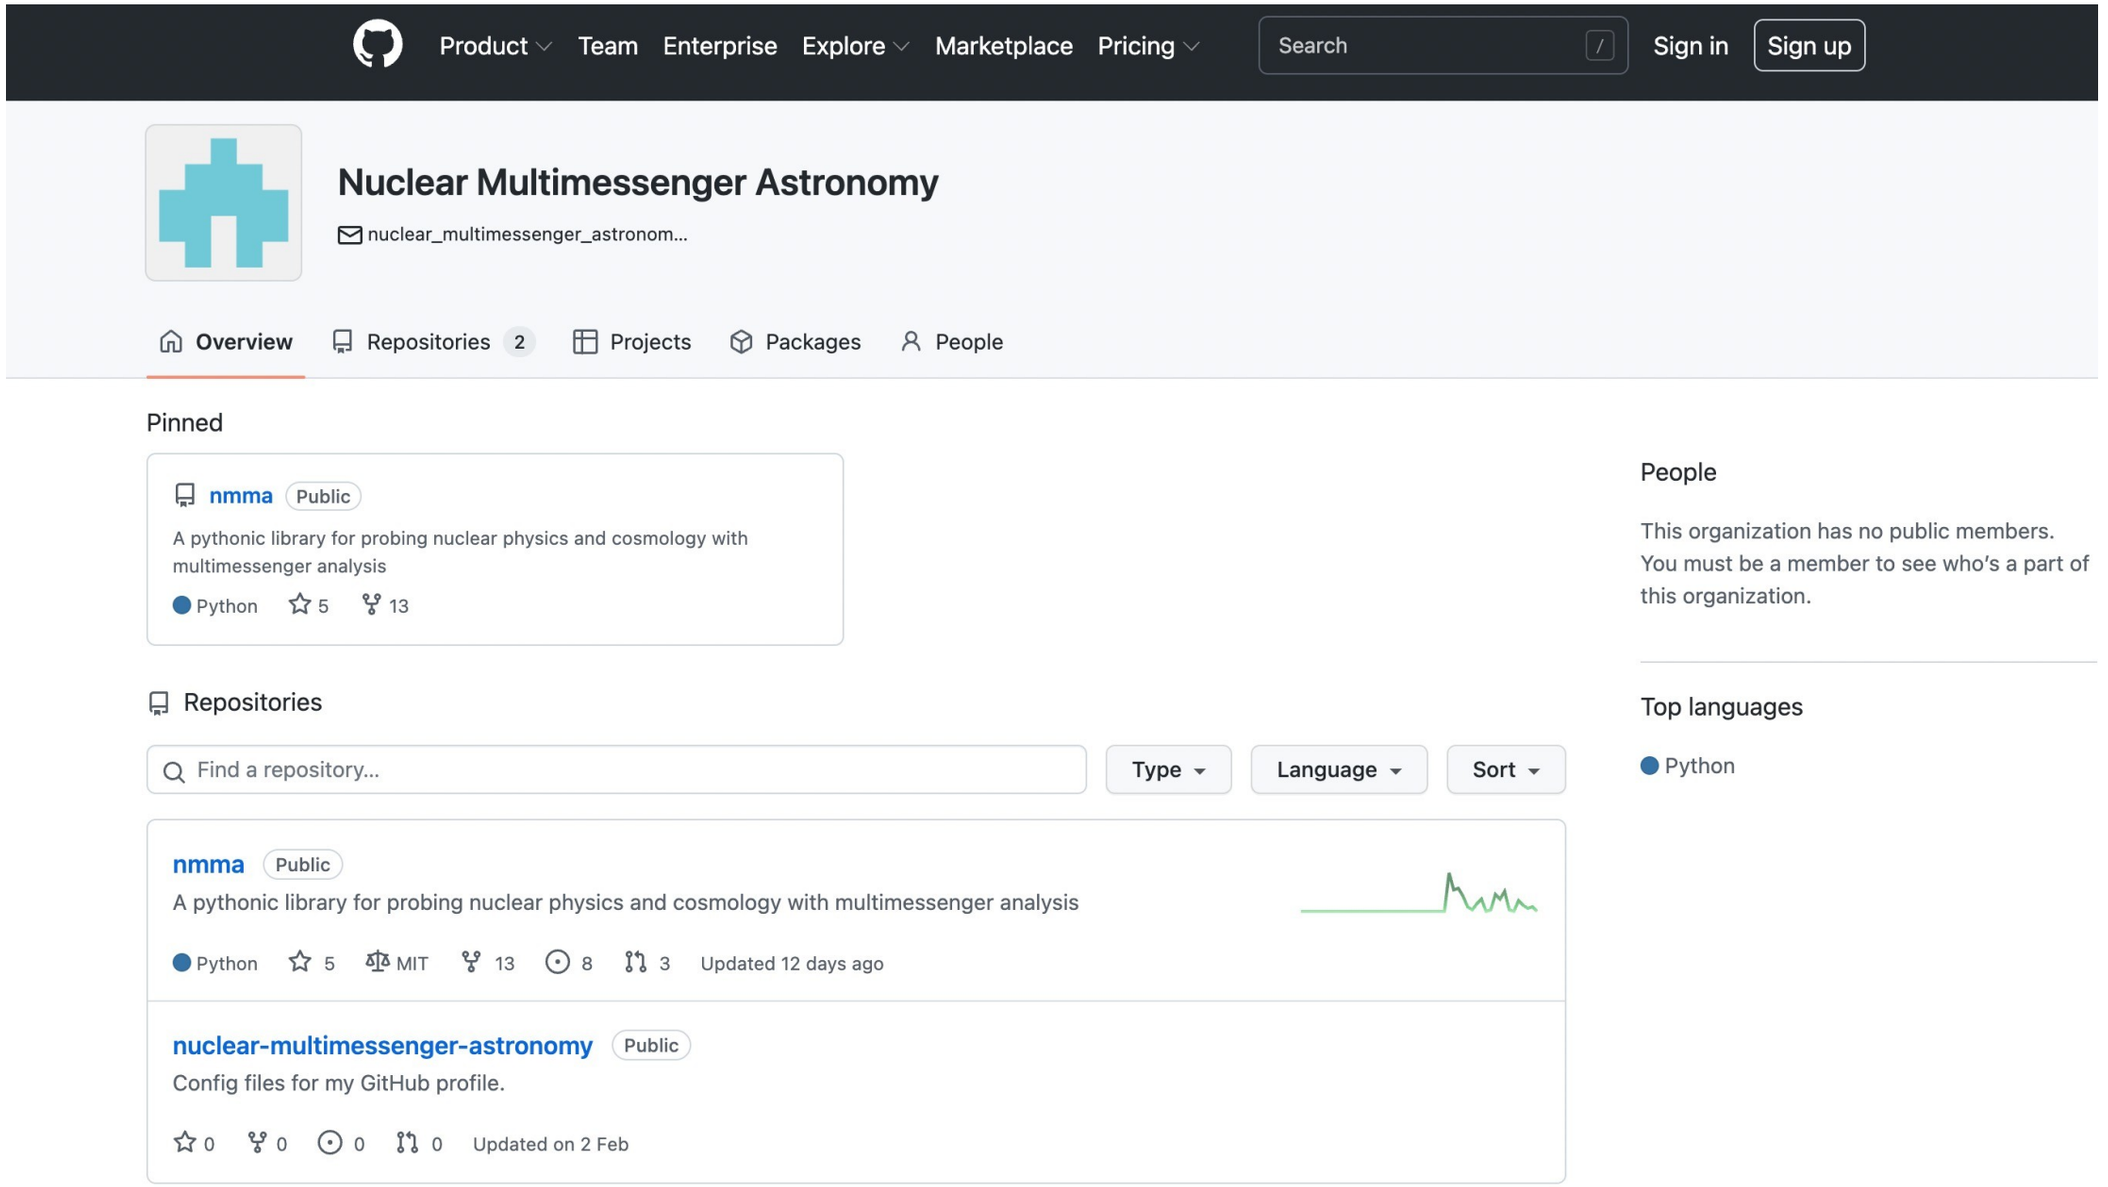

NMMA - Consider using it!

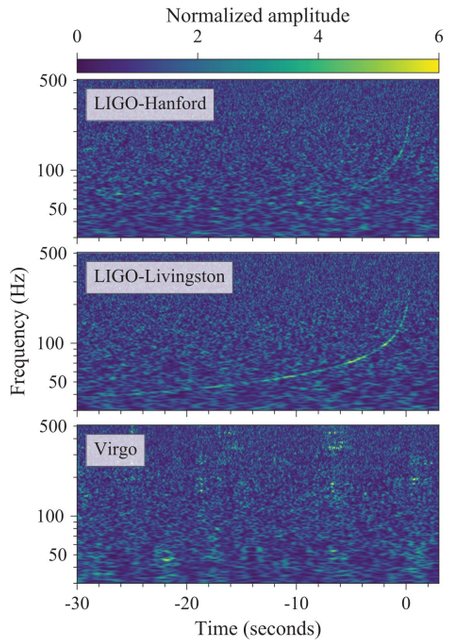

+

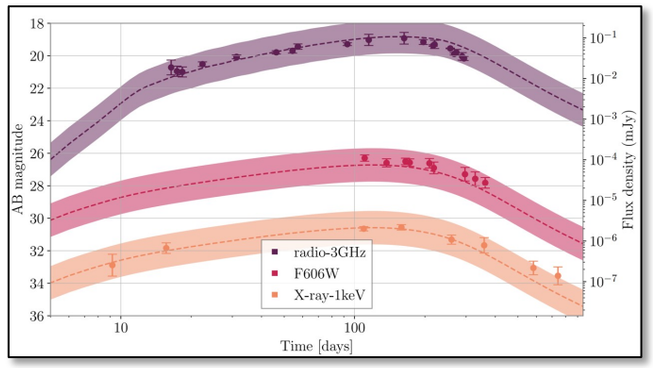

+

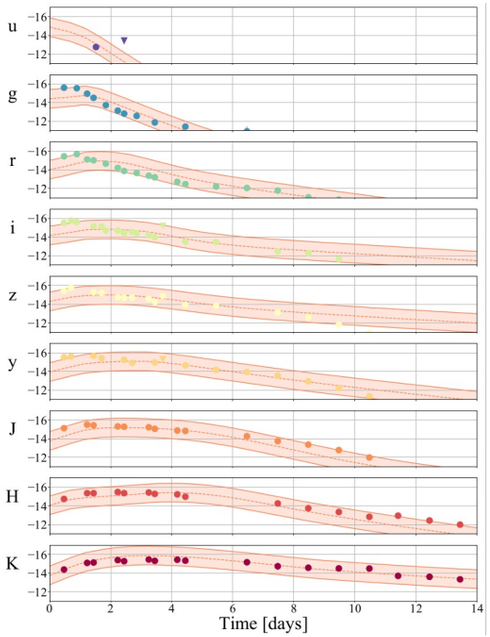

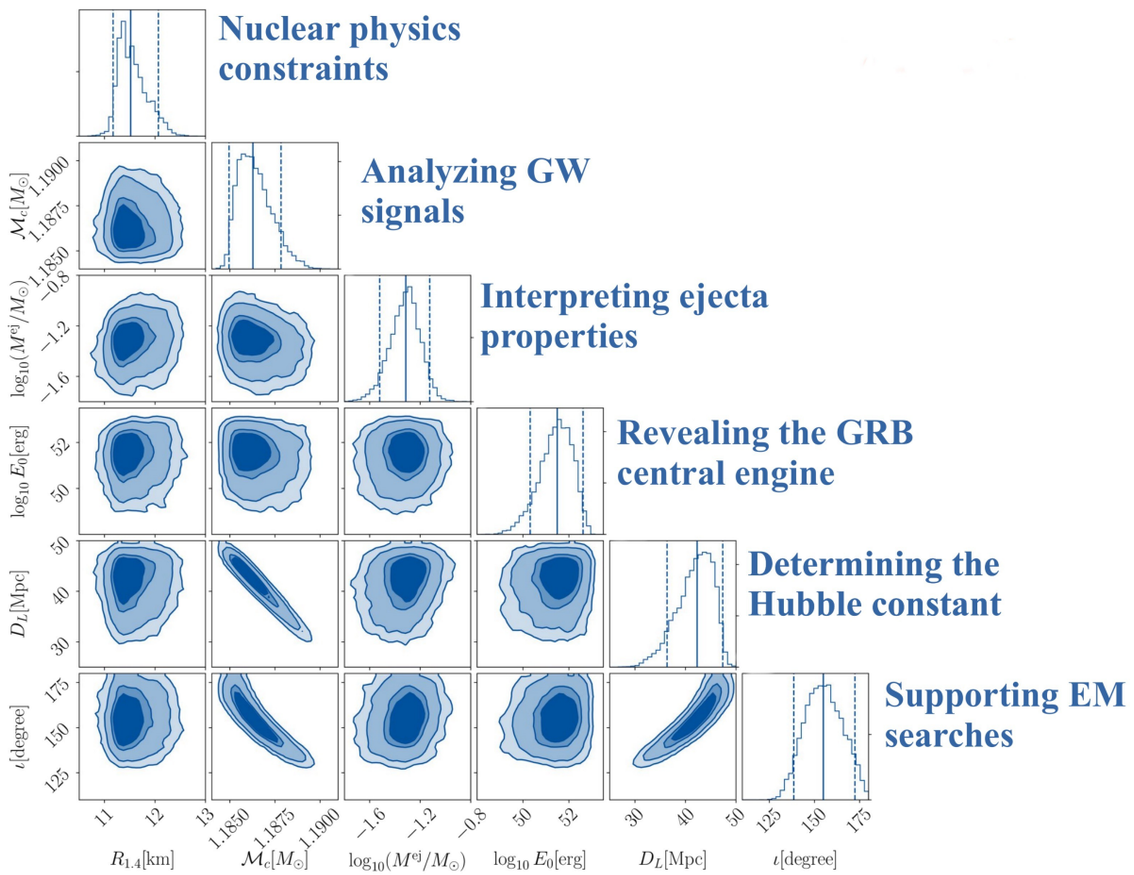

- incorporation of nuclear-physics information
- <u>simultaneous</u> analysis of GW, kilonova, and GRB afterglow
- (some) unit testing and (growing) documentation!
- Used in online kilonova searches
- HPC facilities needed

Similarity Imbeddings in Likelihood Free Inference

Proof of concept demonstration of incorporating symmetries and optimizing likelihood free inference

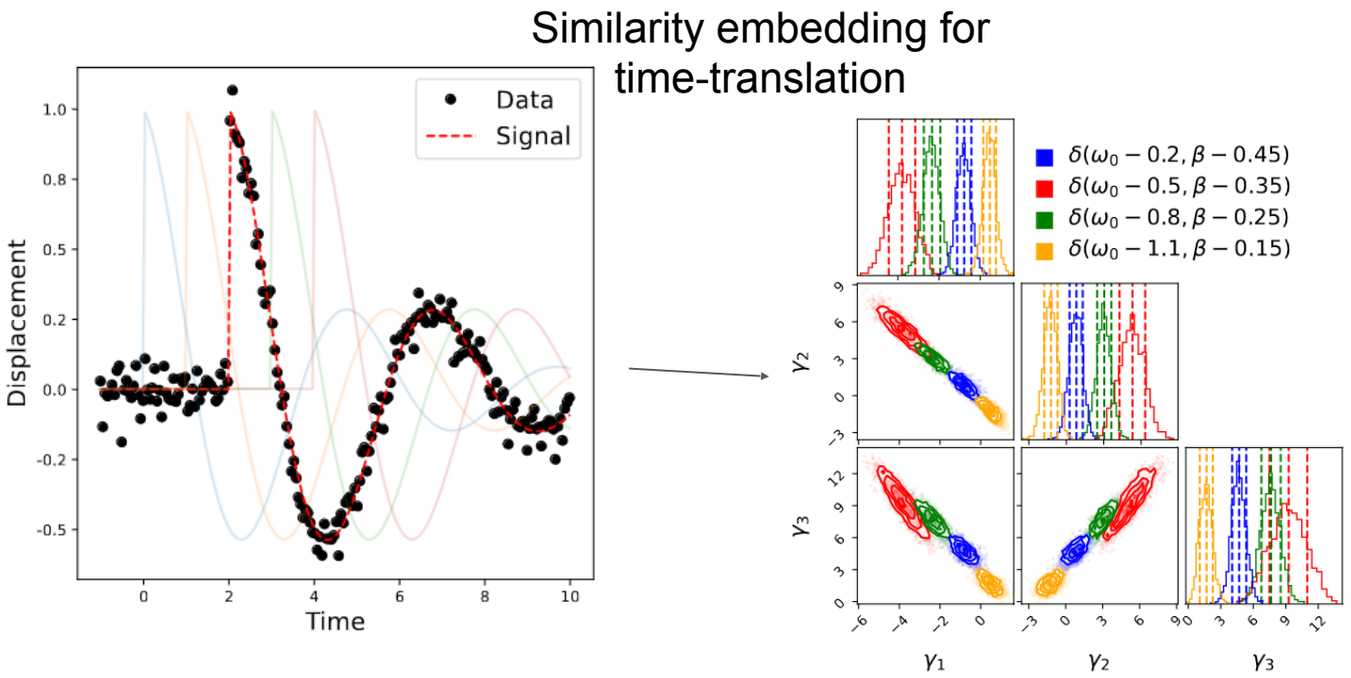

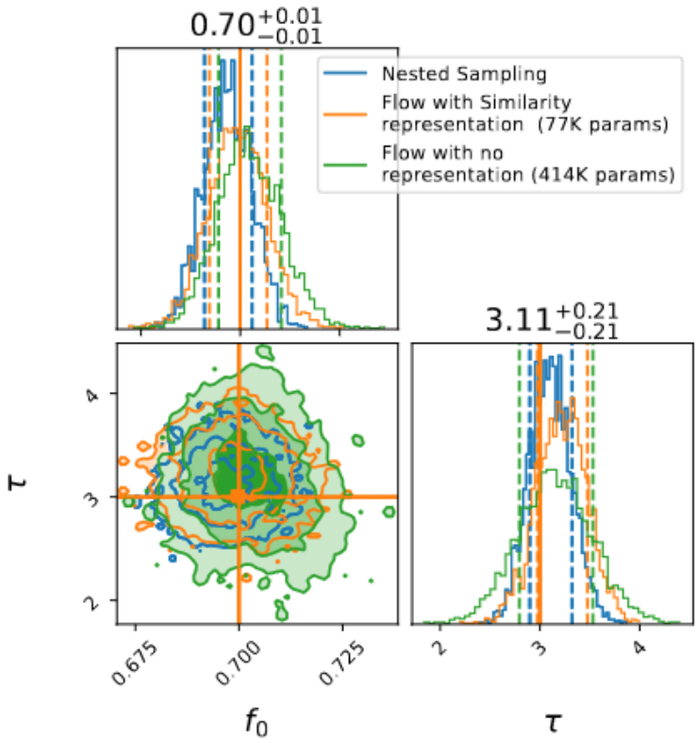

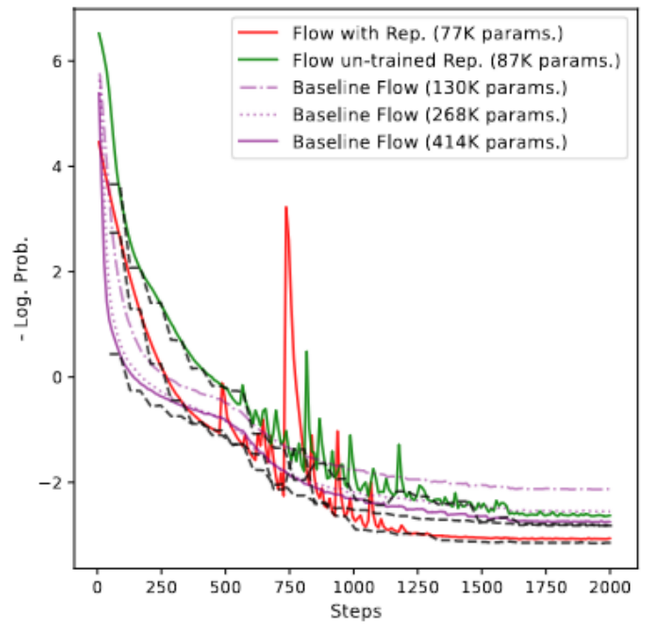

Similar performance at lower number of parameters. Currently exploring avenues to marginalize nuisance parameters in GW and lightcurve parameter estimation.

The Time-Domain Astronomy Ecosystem

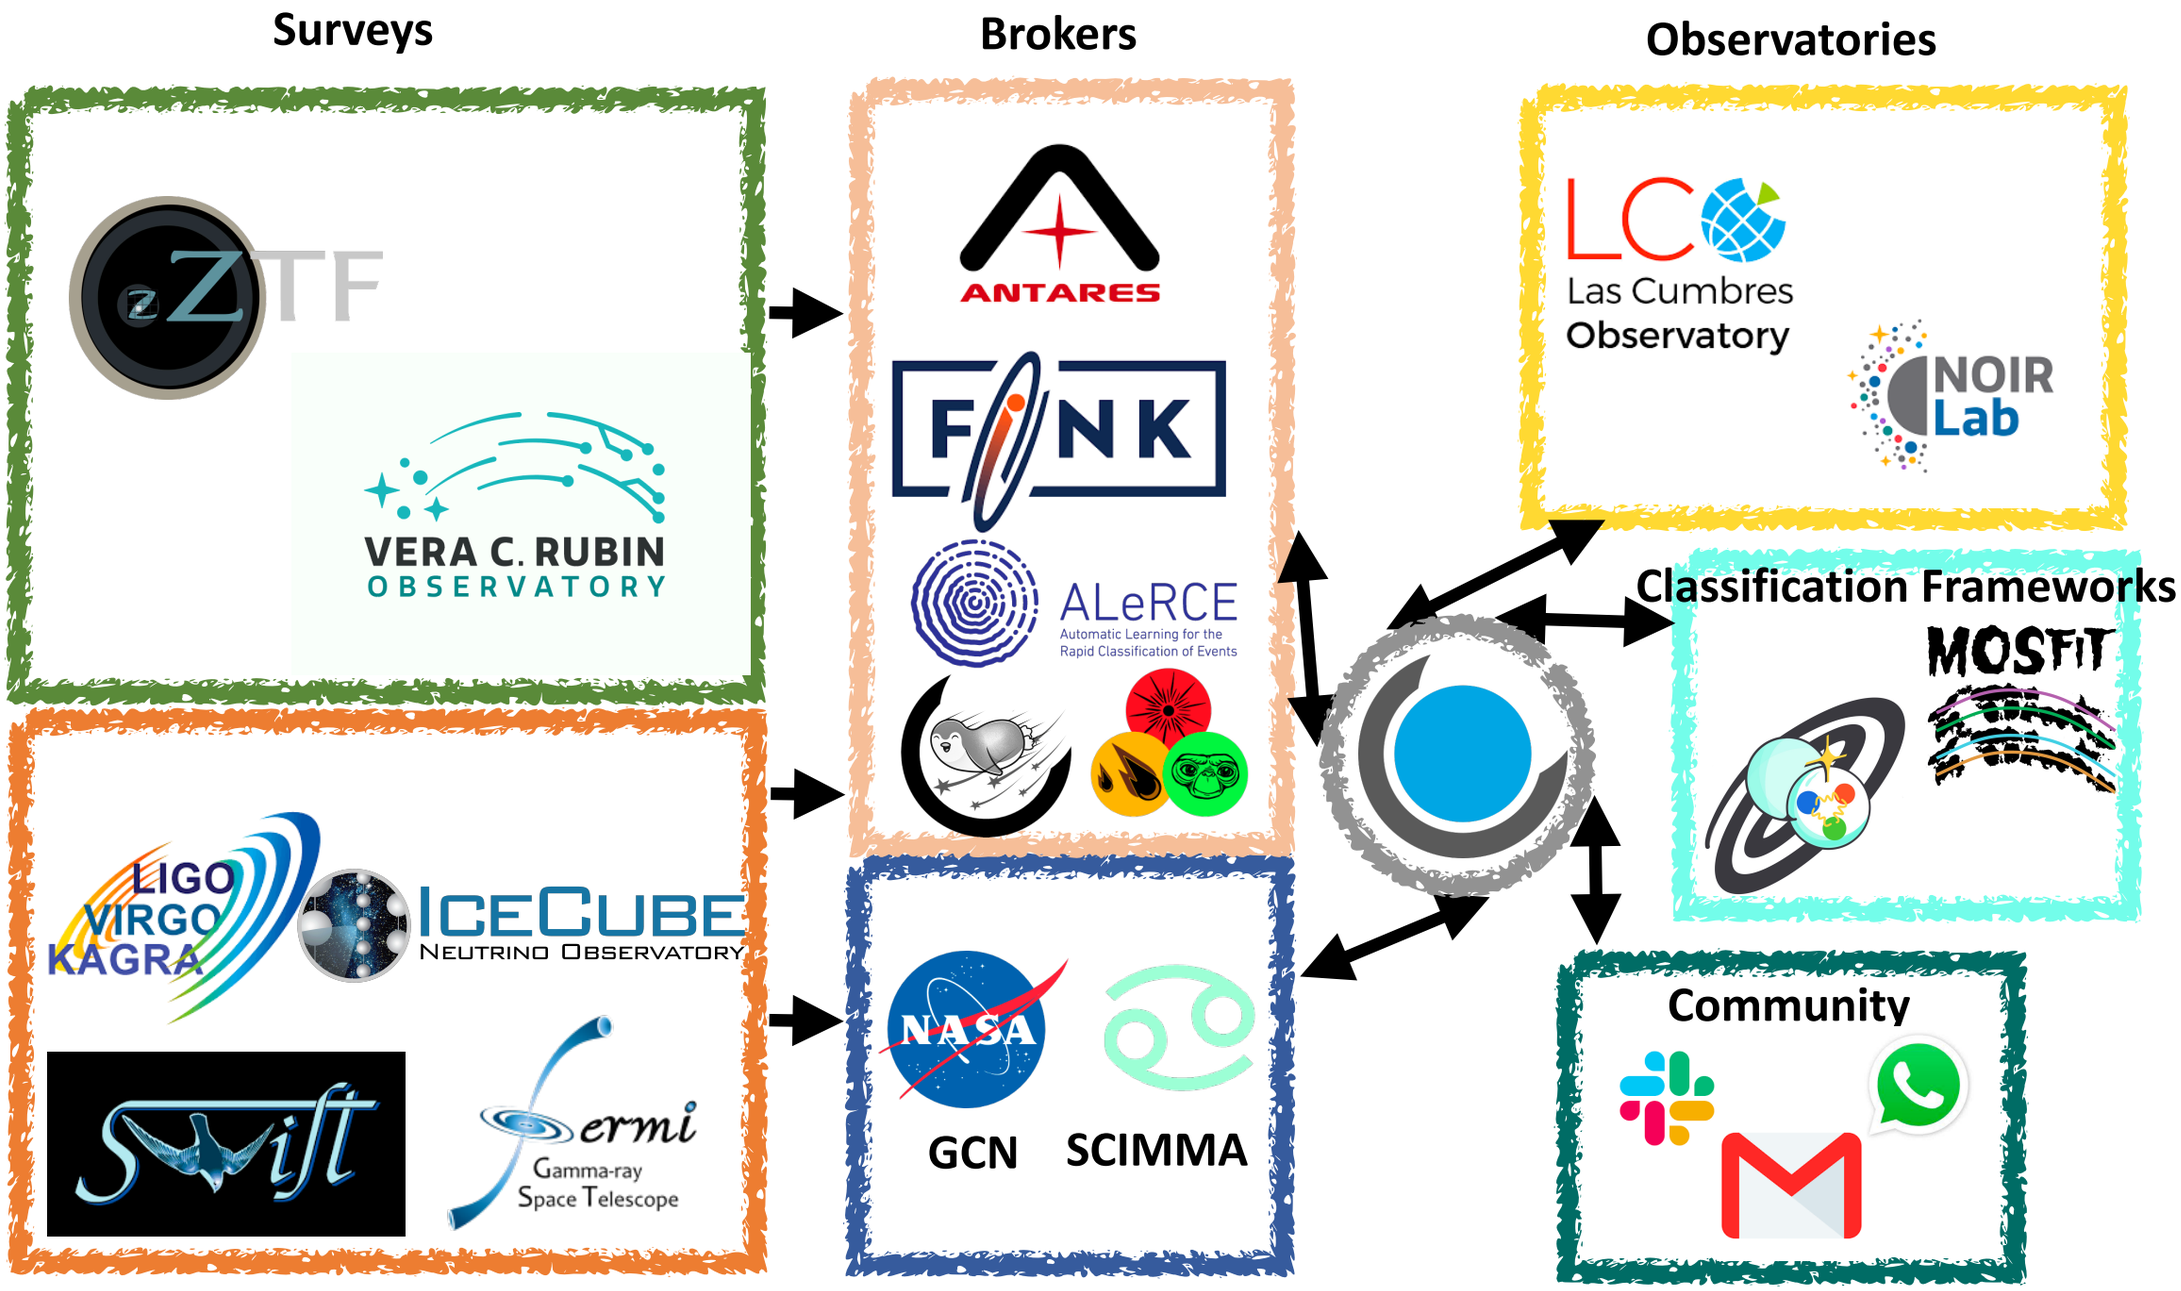

SkyPortal: Overview

A **portal** utilizing secure modern web technologies, **scaling** effectively, and is highly **highly customizable** and extensible to various astronomy workflows related to ZTF, LSST, LS4, and other surveys. A single-source-of-truth **marshal** for transient, variable, and Solar system science cases. Facilitates **follow-up observation management**: robotic and classical facilities

- Open source (free to use, modify, and distribute)
- API-first system: rich APIs for machine usage
- Extensible & scalable design
- Fine-grained access control, Authentication via Social/OAuth
- Real-time Slack-like messaging, notifications
- Rich visualization capabilities
- MMA planning, telescope triggering, follow-up management
- 3rd Party Source Analysis integration
- Distributed computation via Dask
- Docker compose or Kubernetes deployment
- Well-tested, extensive docs, CI/CD

SkyPortal: In Context

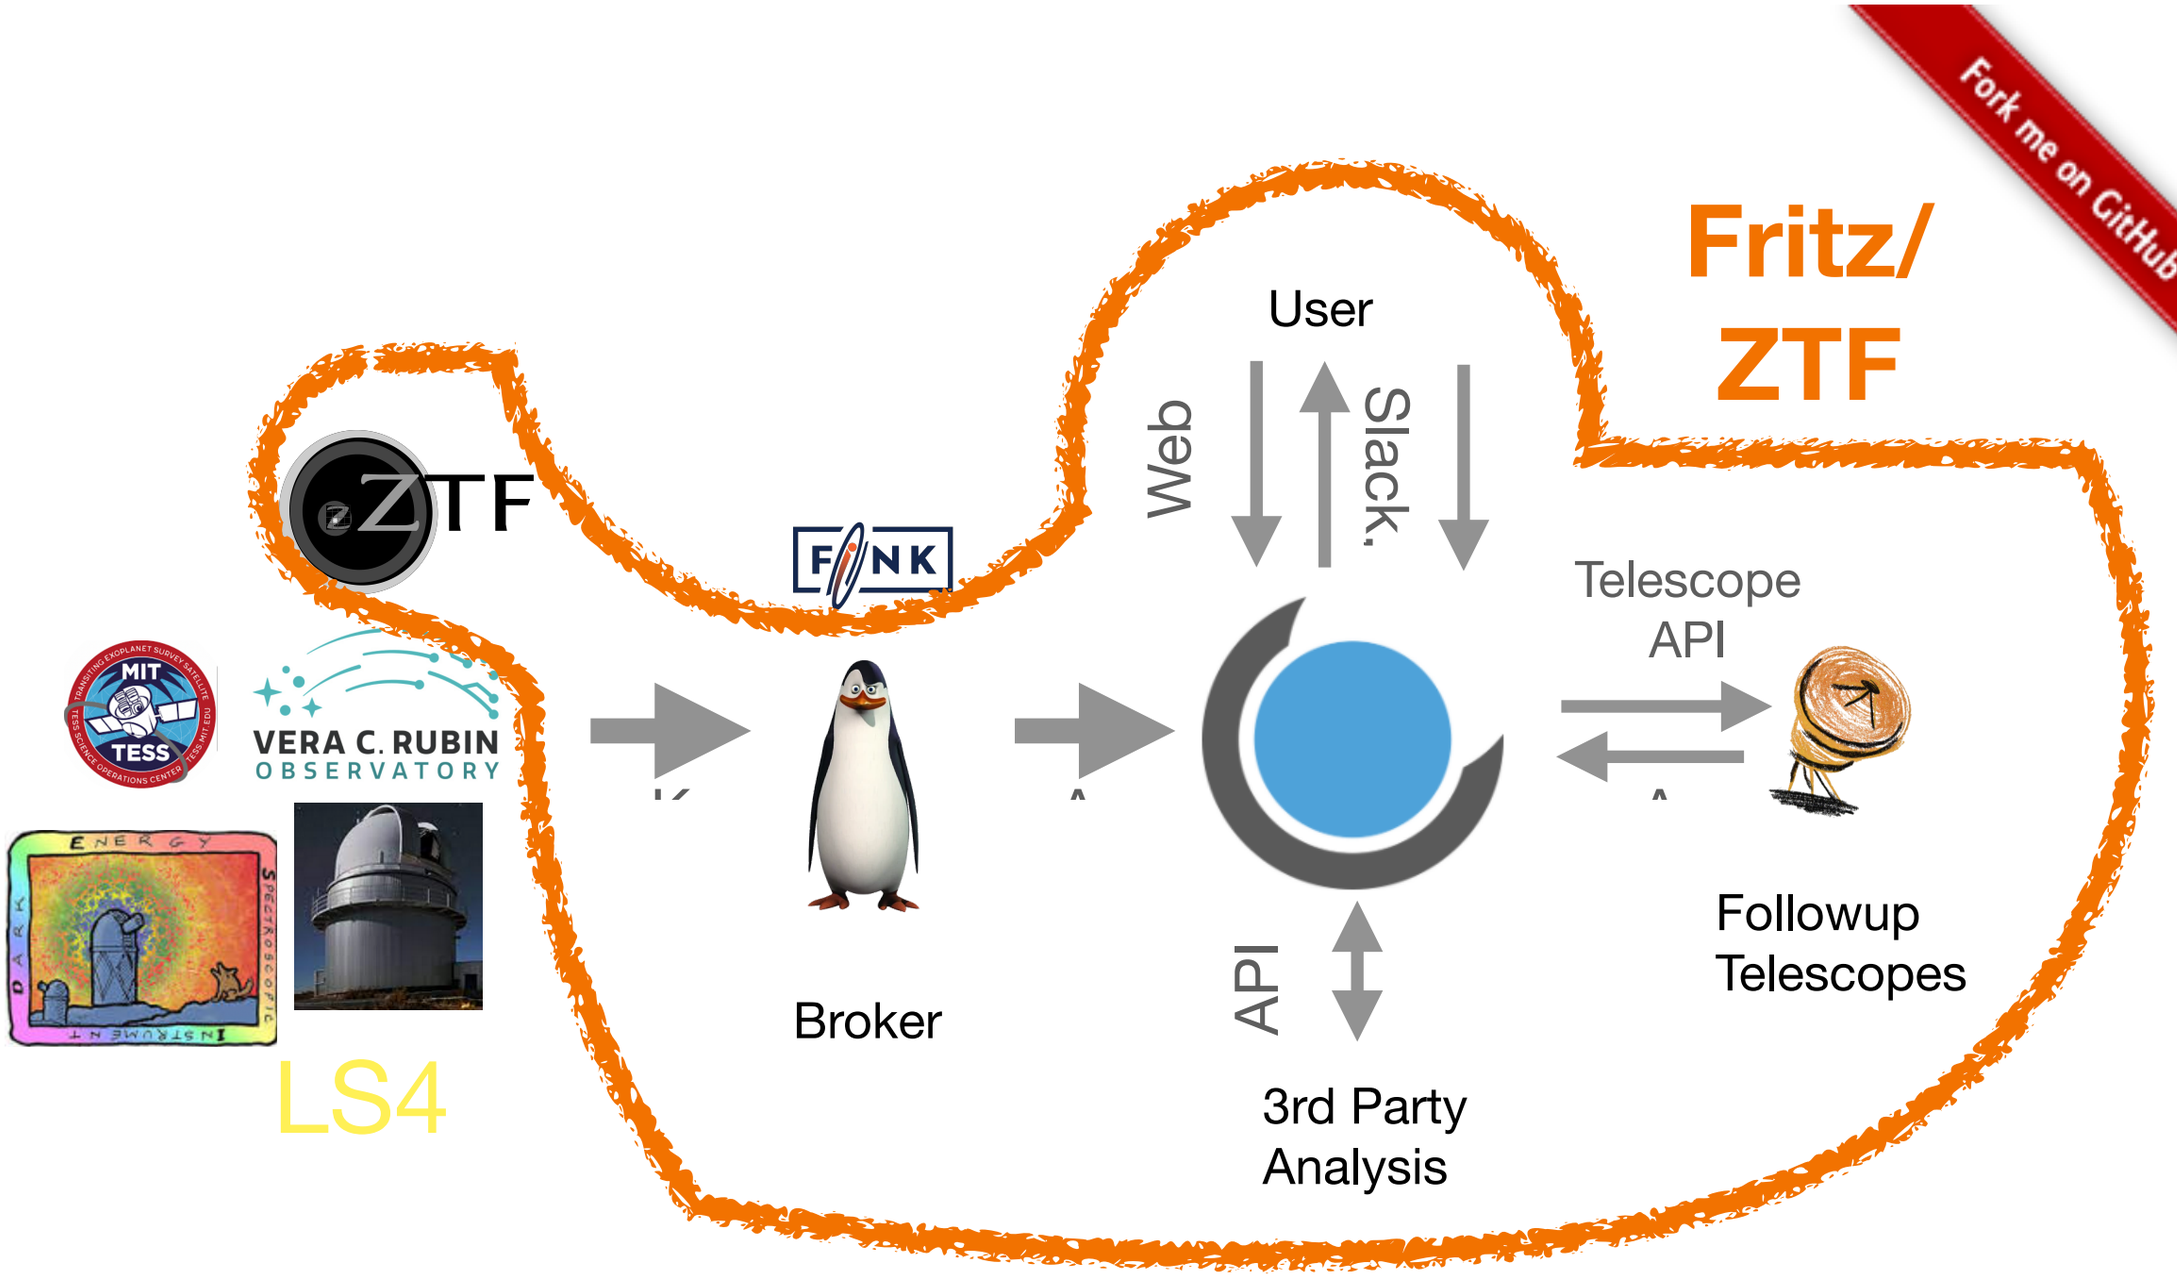

SkyPortal: O4 workflow

Ingestion by `gcn-kafka`

Network-level **scheduling** using `gwemopt`. Validation using `simsurvey`.

API interactions with follow-up telescopes (SEDM, LCO, Swift, NICER, etc.) and photometry services (ATLAS, ZTF, etc.)

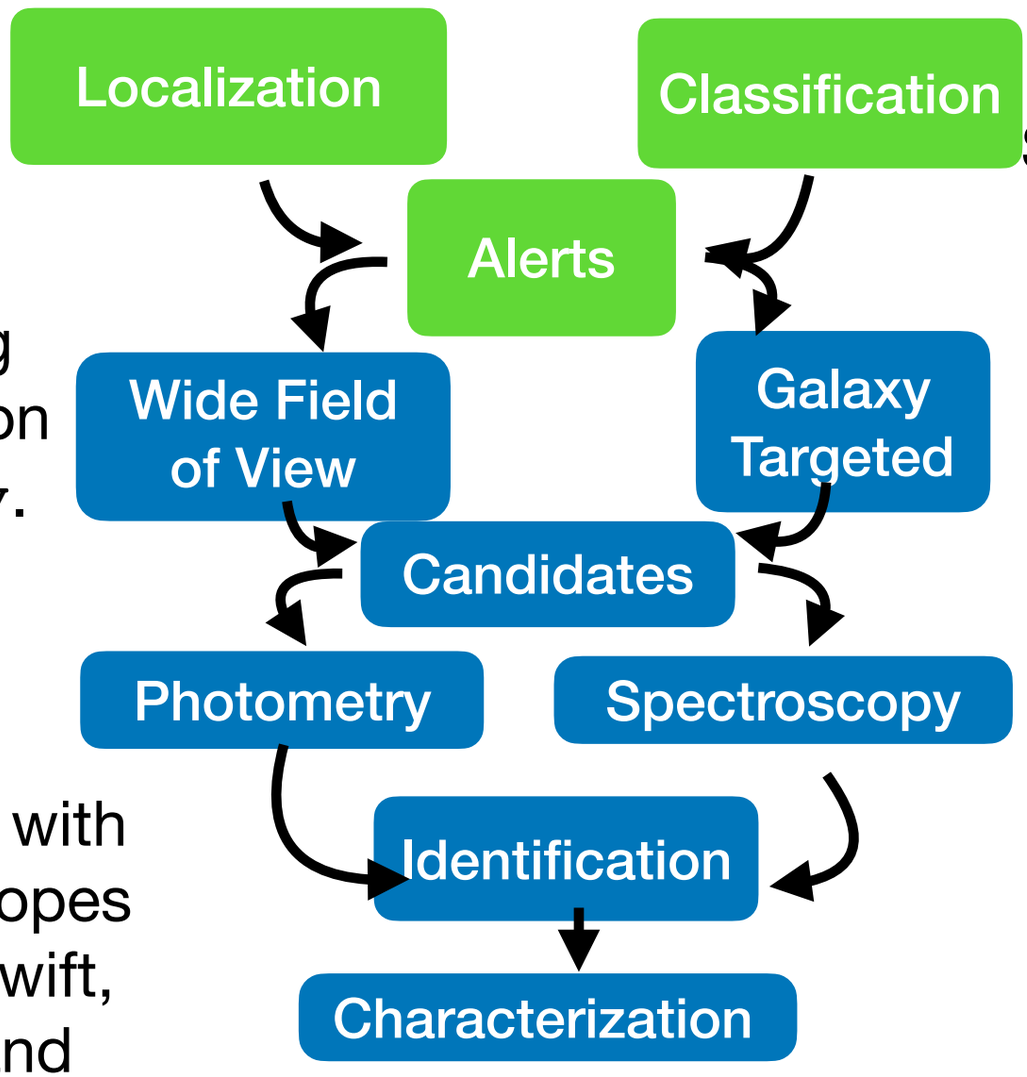

Notification framework by email, SMS, slack, phone etc. Time-zone cognizant shifts

Cross matching of filter streams with skymaps. **Automated summaries** for distribution through GCN.

Analysis framework to interact with light curve / spectroscopy fitting frameworks

ZTF Source Classification Project

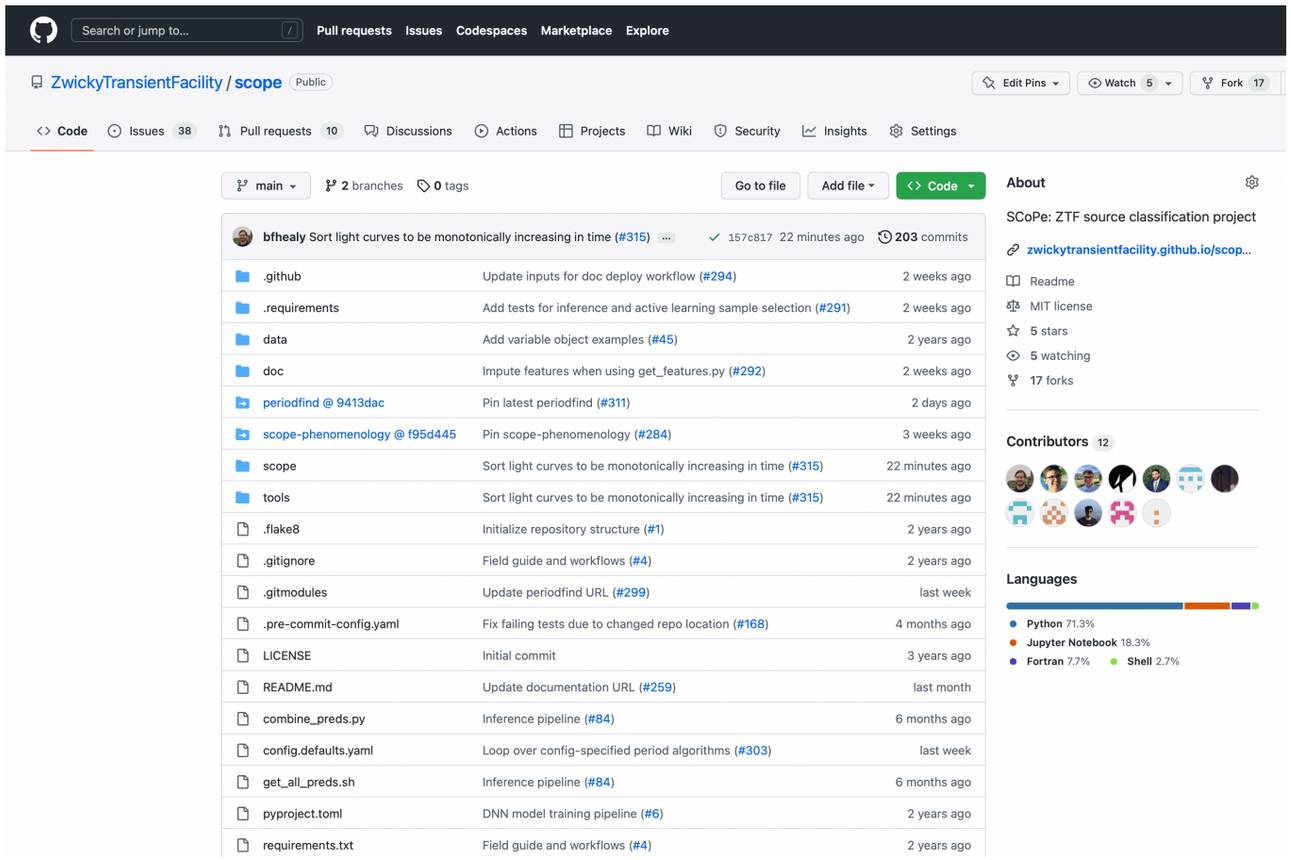

- Open-source
- Python-based
- CI/CD pipeline
- Regularly updated docs

Supervised, active learning: training set built up over time (w/human input)

Two taxonomies: ontological (intrinsic), phenomenological (light curve shape)

Data release on Zenodo: https://zenodo.org/records/11127912

(van Roestel et al. 2021, Coughlin et al. 2021, Healy et al. 2024)

SkyPortal: In Production

6.8M
Source Views

3.6M
Sources

540M
Photometry Points

6.4k
GCN Events

403
Users

223k
Comments

4.3M
Annotations

18.3M
Thumbnails

11.1M
Candidates

The Bright Transient Survey

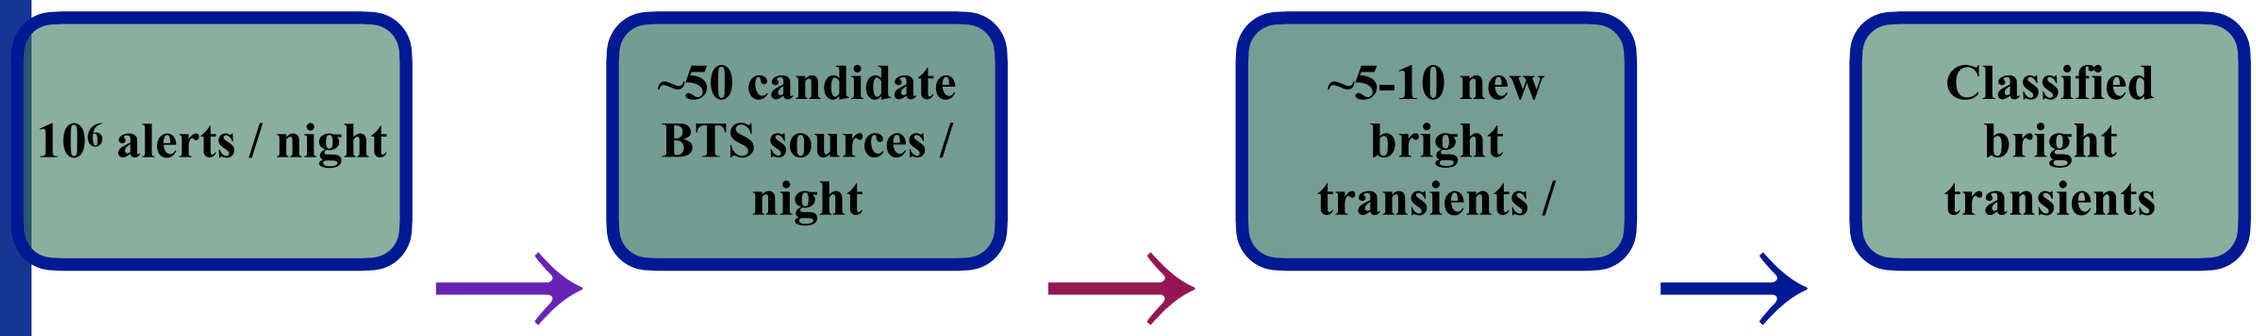

BTSBot

10^6 alerts / night

→ `sgscore`
Tachibana & Miller 18
`braai`
Duev+19

~50 candidate BTS sources / night

→ `BTSbot`
Rehemtulla+24

`SkyPortal`
Coughlin+23

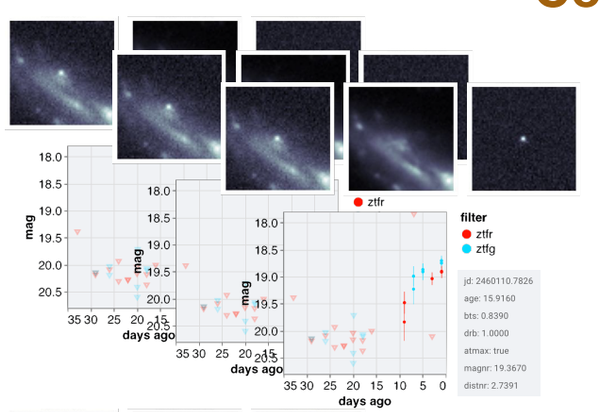

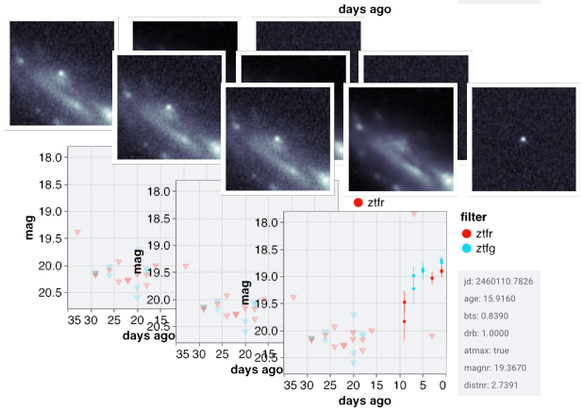

~5-10 new bright transients /

→ `pySEDM`
Rigault+19

Classified bright transients

`SNIascore`
Fremling+21

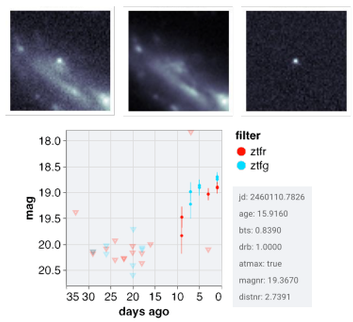

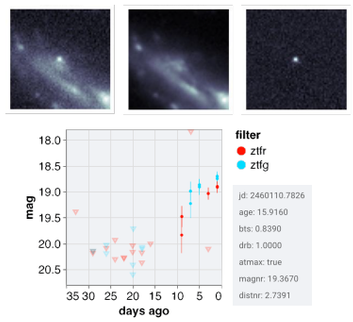

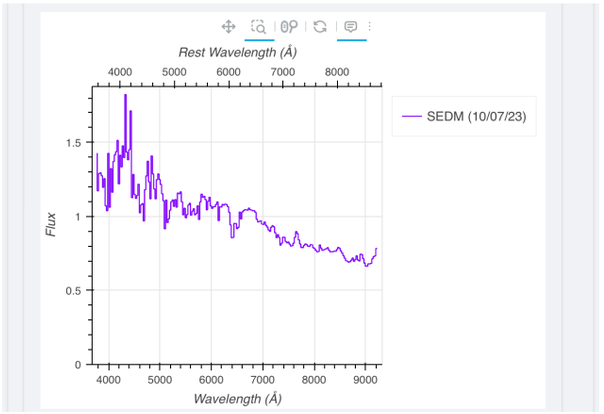

SN 2023tyz

World's First: Fully automatic from discovery to TNS

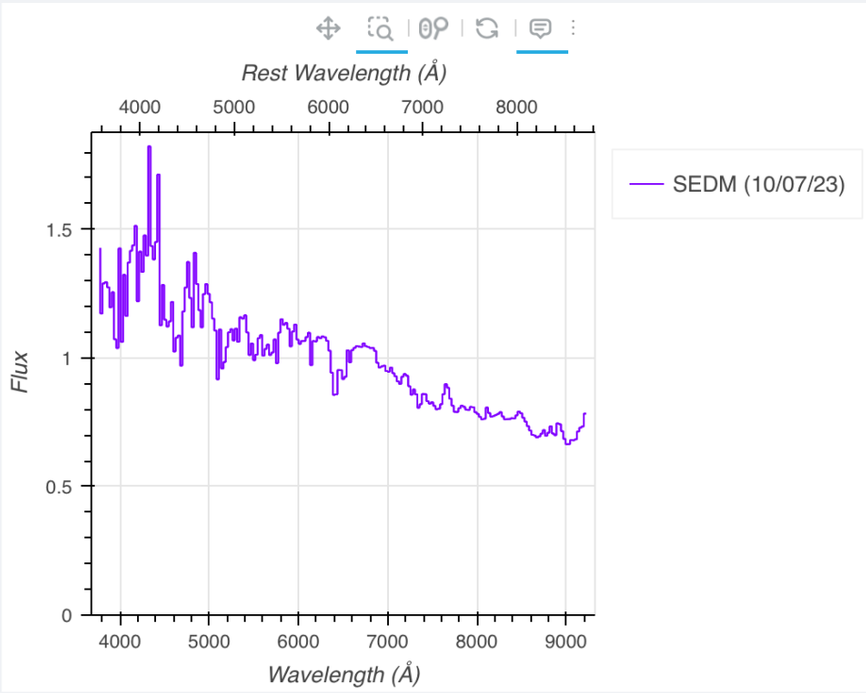

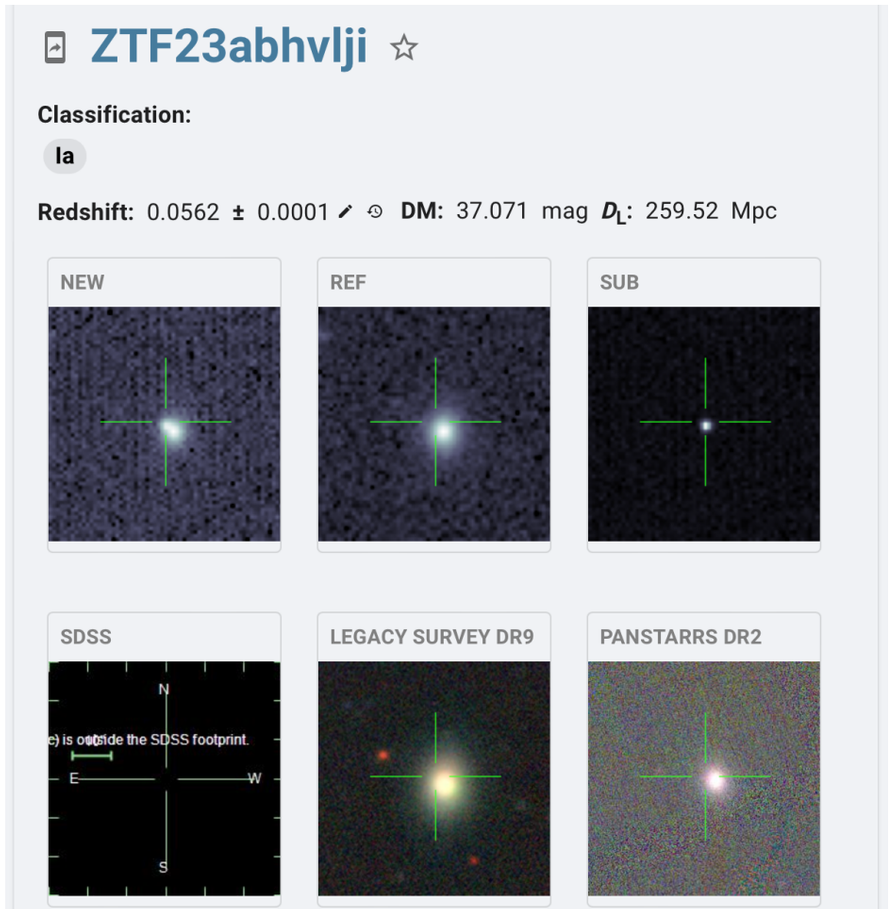

Nabeel Rehemtulla
Northwestern

Theophile Jegou du Laz
Caltech

SN 2023tyz

World's First: Fully automatic from discovery to TNS

Up to ~70 fully automated classifications, ~96% purity for BTSBot, "very very complete" - Nabeel

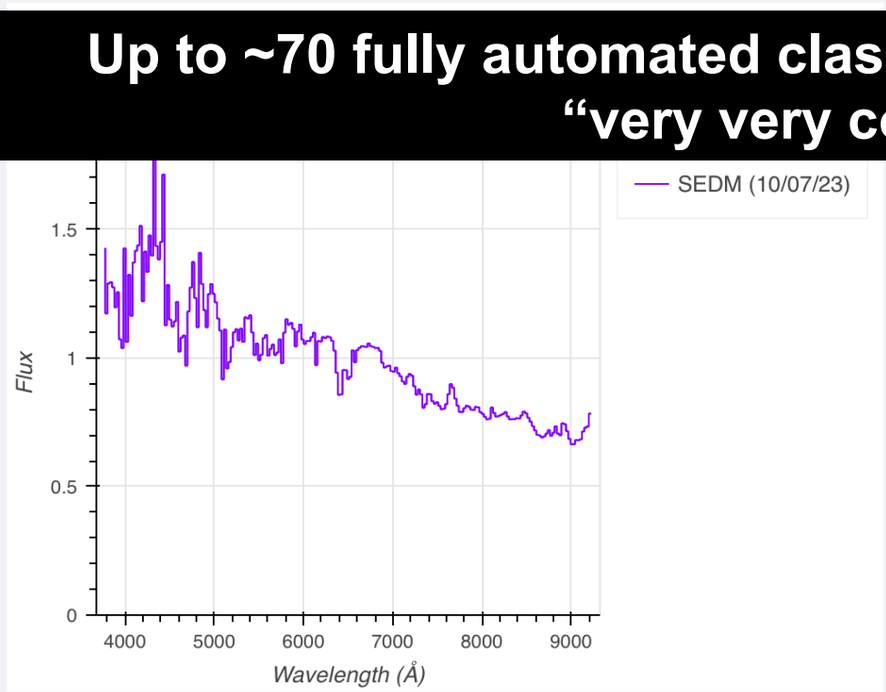

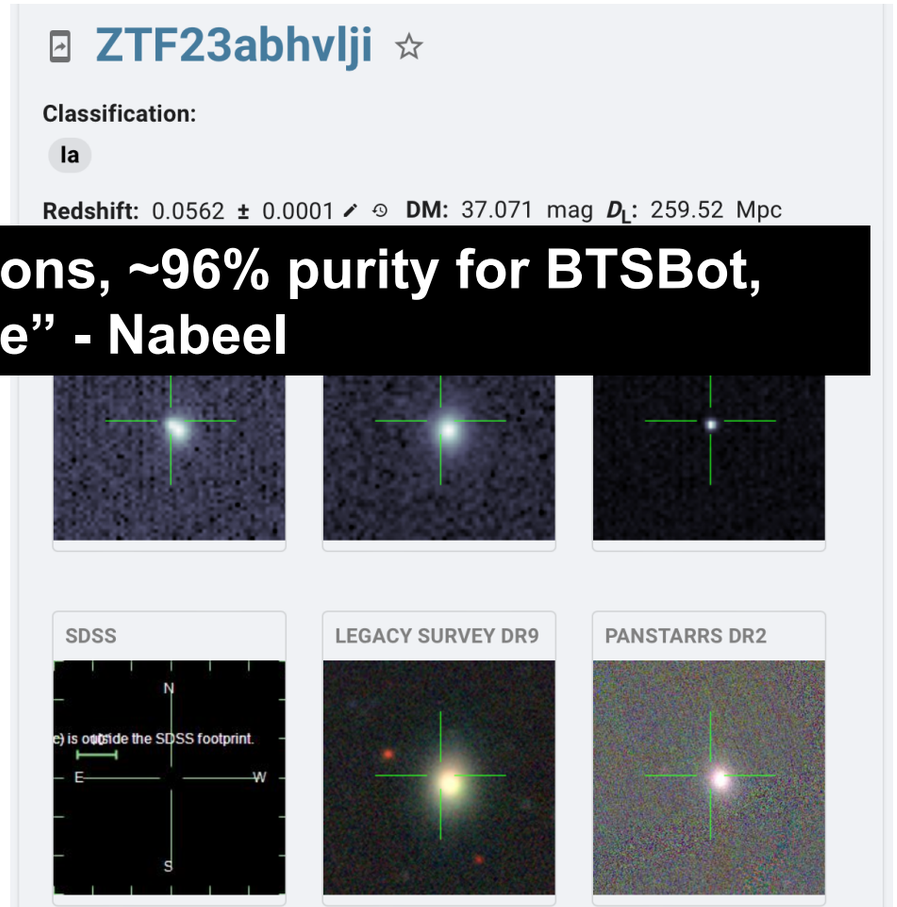

Nabeel Rehemtulla
Northwestern

Theophile Jegou du Laz
Caltech

SED Machine - Kitt Peak: First Science Observations

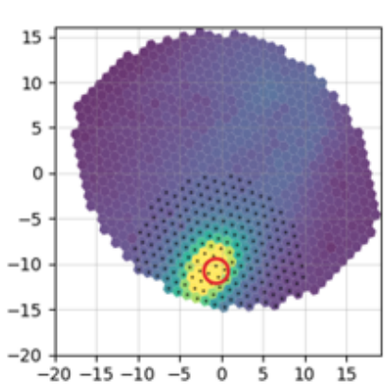

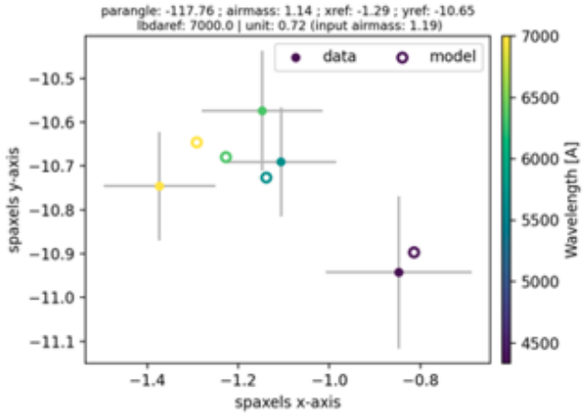

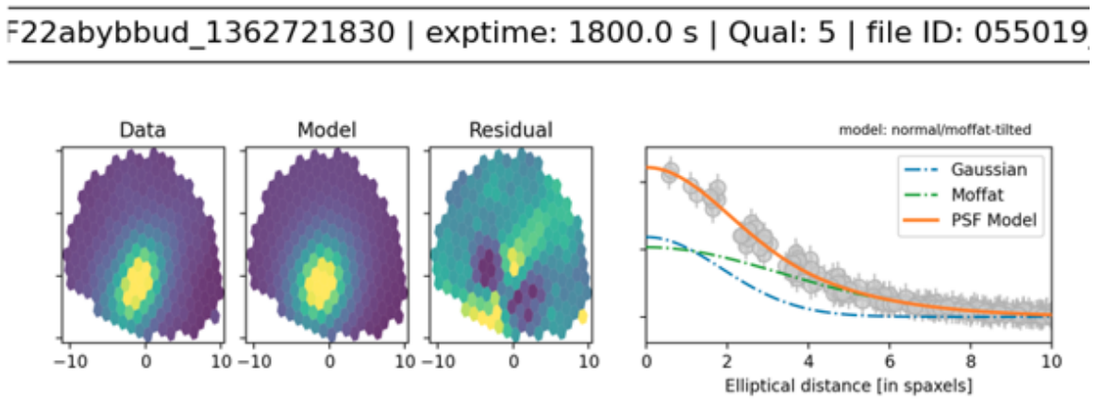

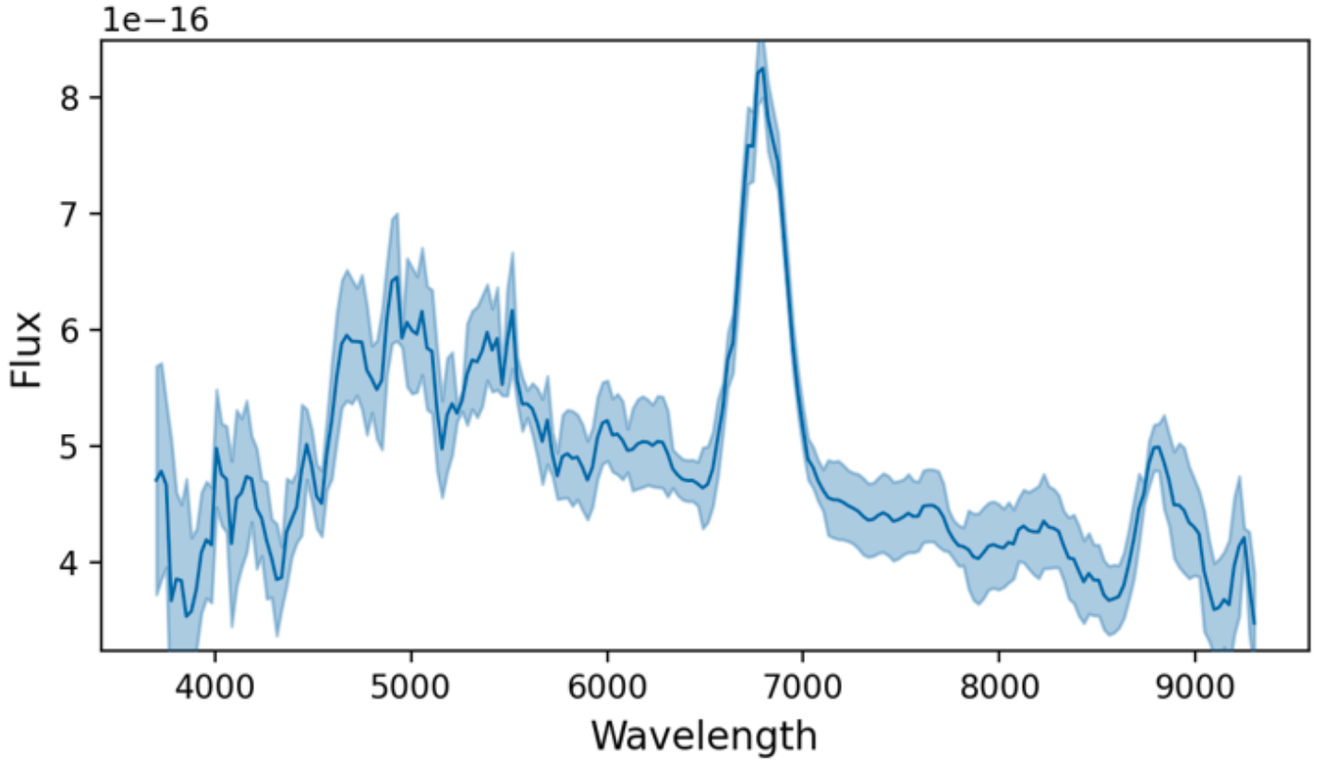

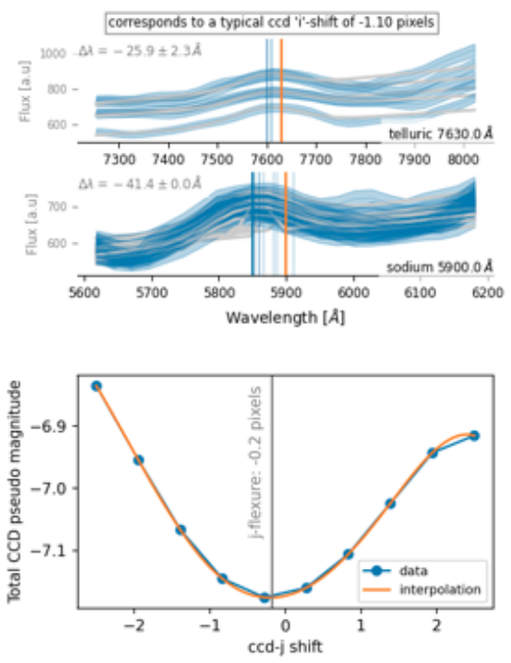

pysedm version 0.30.0 | made the 2023-04-16 at 18:33:00

SED Machine - Kitt Peak

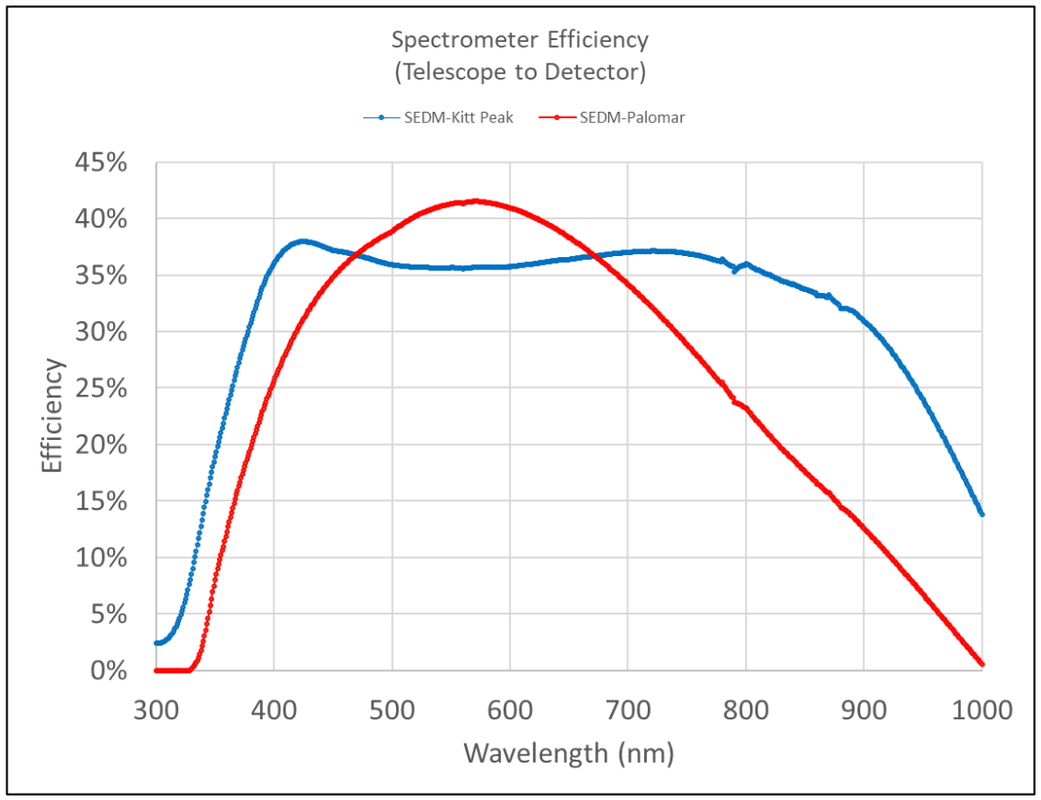

Kitt Peak 2.1m: Facility Specs

- Primary: 2.1m (84in)
- 2x P60 area = +0.75mag
- Secondary: f/7.6
- Automated for KPED

Kitt Peak 2.1m: Facility History

- 3yrs with RoboAO
- 2yr with KPED

Instrument improvements over v1

- Optimize IFU wavelength coverage and throughput
- Optimize imager FOV
- Reduce number of optics
- Improved QE response in imager
- Use filter wheel for imager instead of fixed quadrant design
- Use fold mirror with central hole instead of pickoff mirror

SED Machine - Kitt Peak

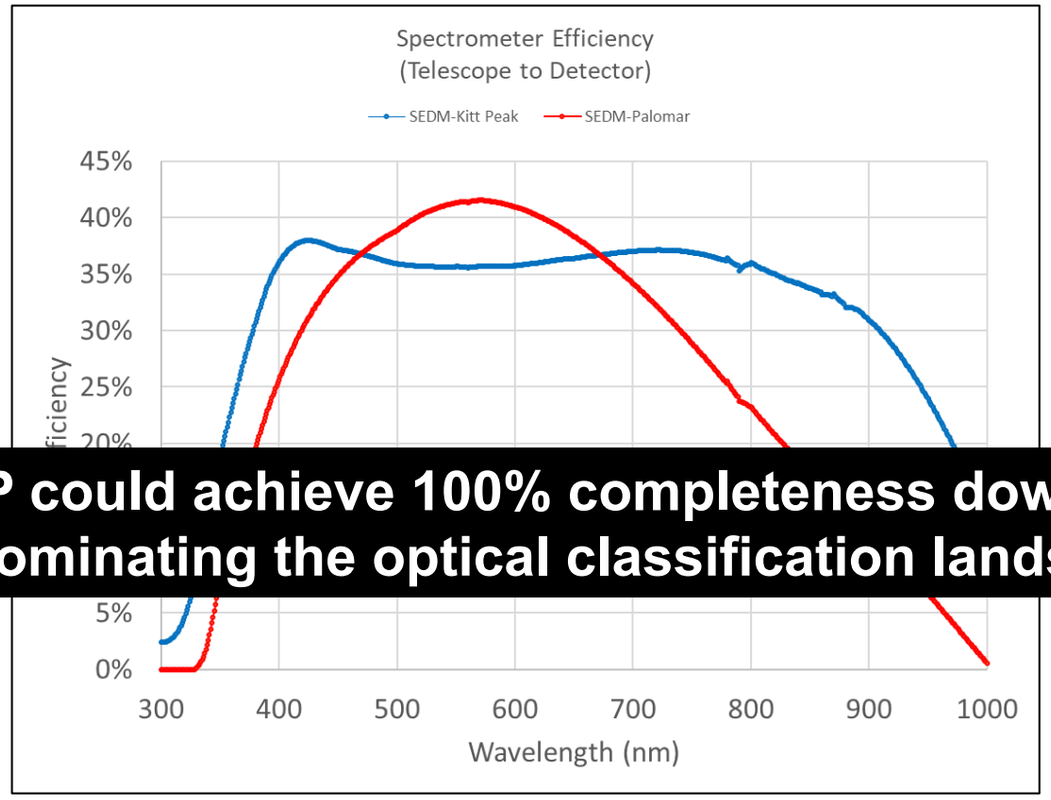

SEDM-KP could achieve 100% completeness down to 19+ mag (dominating the optical classification landscape)

Kitt Peak 2.1m: Facility Specs

- Primary: 2.1m (84in)
- 2x P60 area = +0.75mag
- Secondary: f/7.6
- Automated for KPED

Kitt Peak 2.1m: Facility History

- 3yrs with RoboAO
- 2yr with KPED

Instrument improvements over v1

- Optimize IFU wavelength coverage and throughput
- Optimize imager FOV
- Reduce number of optics
- Improved QE response in imager
- Use filter wheel for imager instead of fixed quadrant design
- Use fold mirror with central hole instead of pickoff mirror

NSF HDR Institute: Accelerated Artificial Intelligence Algorithms for Data Driven Discovery

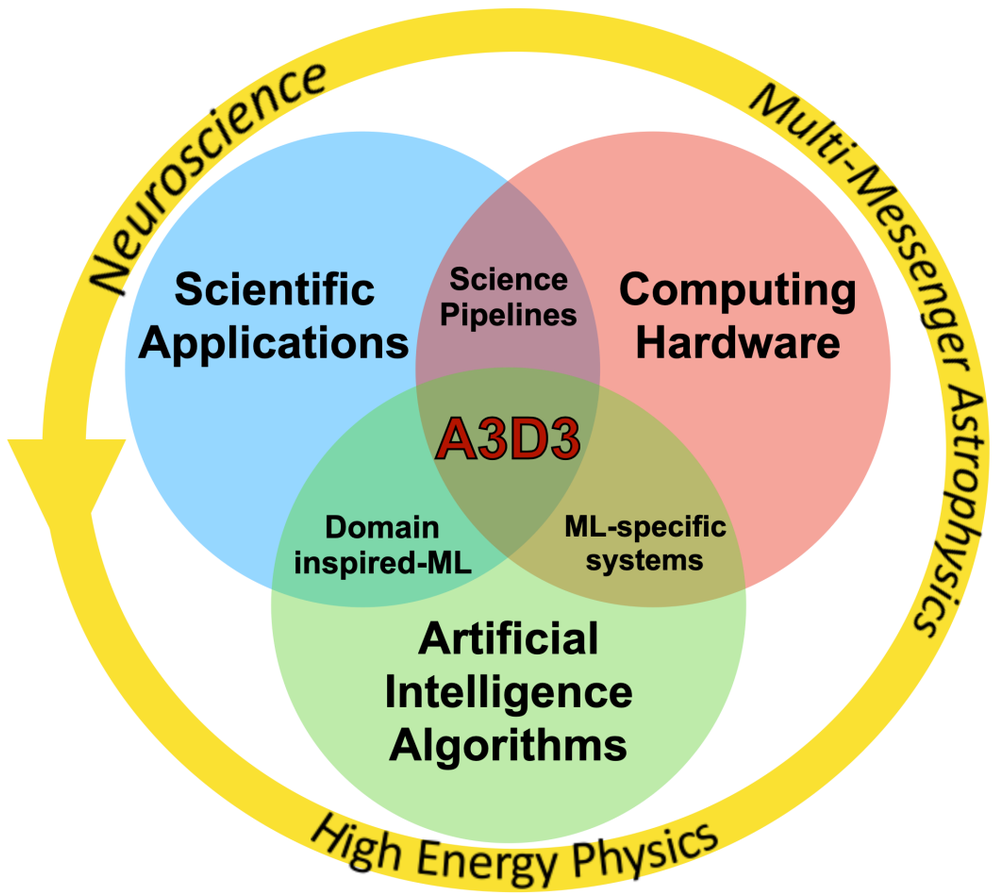

- Integration of three core components: **domain scientists**, **computer scientists** and **hardware engineers**

- Leading a paradigm shift in the application of **real-time AI at scale** to advance science and engineering discoveries.

NSF HDR Institute: Accelerated Artificial Intelligence Algorithms for Data Driven Discovery

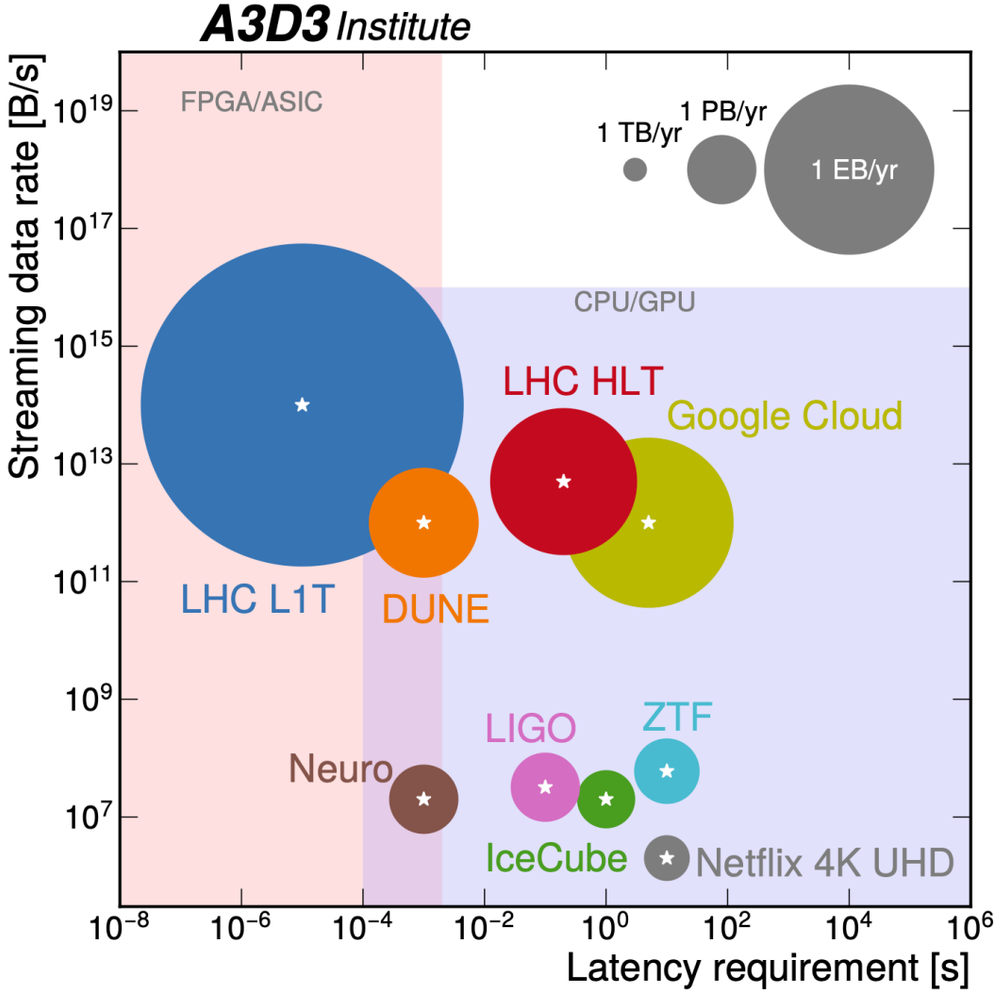

- **High Energy Physics at the LHC driving technology frontier**
  - Both data size and streaming rates of large-scale experiments exceed those handled by industry leaders.
- **New opportunities for applications by accelerating ML/AI algorithms with co-processors**

A3D3: Post-Bac Program

- Unique program within our domains
- 1-year program
- Aim to increase opportunities for women and underrepresented groups in STEM (African American/Black, Chicano/ Latino, Native American/ Alaska Native, Native Hawaiin/Pacific Islander, and Filipino)
- Act as a “bridge” to graduate school or industry

Year 1 Cohort

Abby Gray

Host Institute : University of Minnesota

Research supervisor(s): Dr. Michael Coughlin

Short Intro: Hi! I’m Abby. I’m from New Jersey and graduated in 2021 with a degree in Astrophysics from Colorado College (CC). I spent the last academic year working for the CC physics department and am excited to continue my physics journey this year! Aside from physics, I have a deep love for the outdoors, the ocean, flying, and solo traveling.

Research domain:
Astrophysics

3

Lauren Peterson

Host Institute : University of Washington

Research supervisor(s): Dr. Amy Orsborn

Short Intro: I graduated this June from the UW with a degree in electrical engineering. I have experience with human-subject research, neuroscience, and machine learning and am excited to work with everyone. I also love hiking and kayaking.

Research domain:
Neuroscience

Van Tha Bik Lian

Host Institute : Duke University

Research supervisor(s): Dr. Kate Scholberg

Short Intro: Hello! My name is Van Tha Bik, I graduated from UW with an Applied physics degree in June. I enjoy the outdoors and playing video games. If anyone would like to join me some time in the not so close future, I’d like to try paragliding. Maybe even wingsuit jumping!

Research domain:
Neutrino Physics

5

Andrew J. Skivington (he/him)

Host Institute : University of California - San Diego

Research supervisor(s): Dr. Javier M. Duarte

Short Intro: Greetings all! My name is Andrew, but my preferred name is Andy. I recently graduated this past May 2022 with my B.Sc. in physics from the University of Florida. My long term goal is to enter a physics PhD program where I hope to research machine learning to solve modern physics problems via data driven methods. Hence, I feel I am one of the four luckiest people to be given this opportunity to be a part of the inaugural cohort of A3D3 post-bacc fellows, and I am eager and excited to work with you all this coming year.

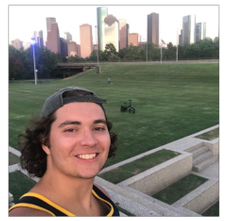

Research domain:
Particle physics and ML

6

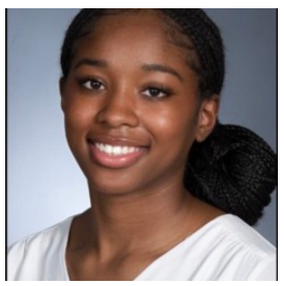

Year 2 Cohort

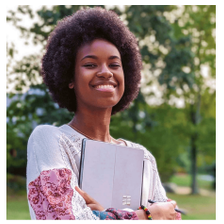

- **Lucie Afko** (W, URM) - Kate Scholberg - Duke (particle physics/astrophysics)
- **Kira Nolan** (W) - Matthew Graham - Caltech (astrophysics)
- **Jada Marshall** (W, URM) - Maria Makin - Purdue (neuroscience)
- **Malina Desai** (W) - Phil Harris & Erik Katsavounidis - MIT (particle physics/astrophysics)

The Time-Domain Astronomy Ecosystem

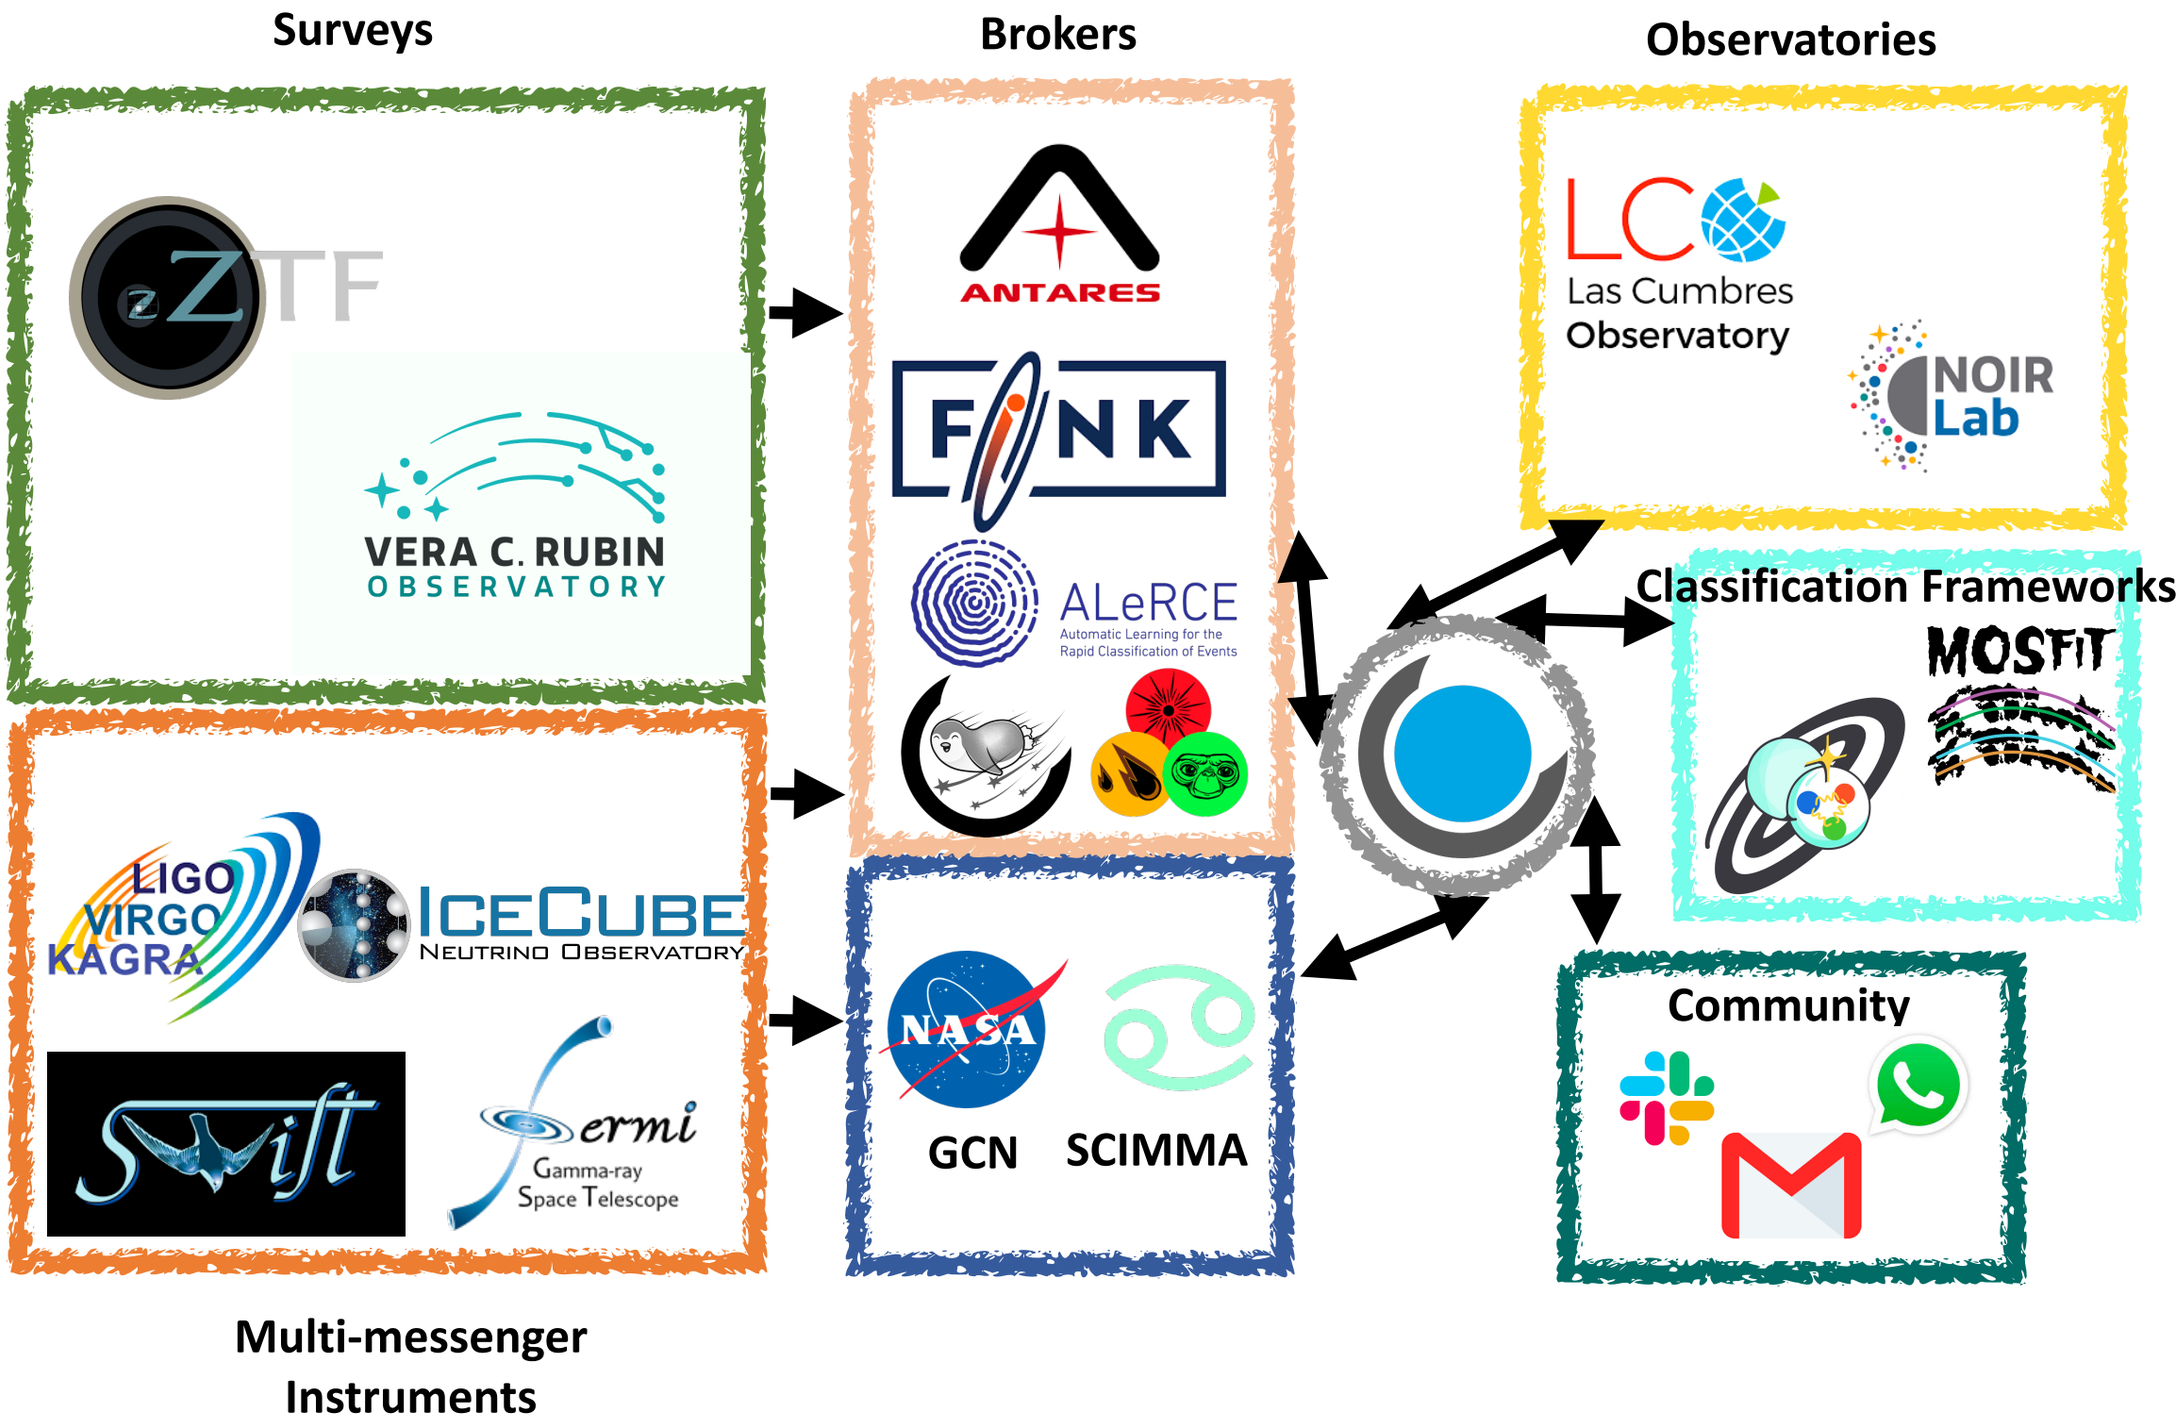

Thank you!