A technical ecosystem to enable multi-messenger astrophysics

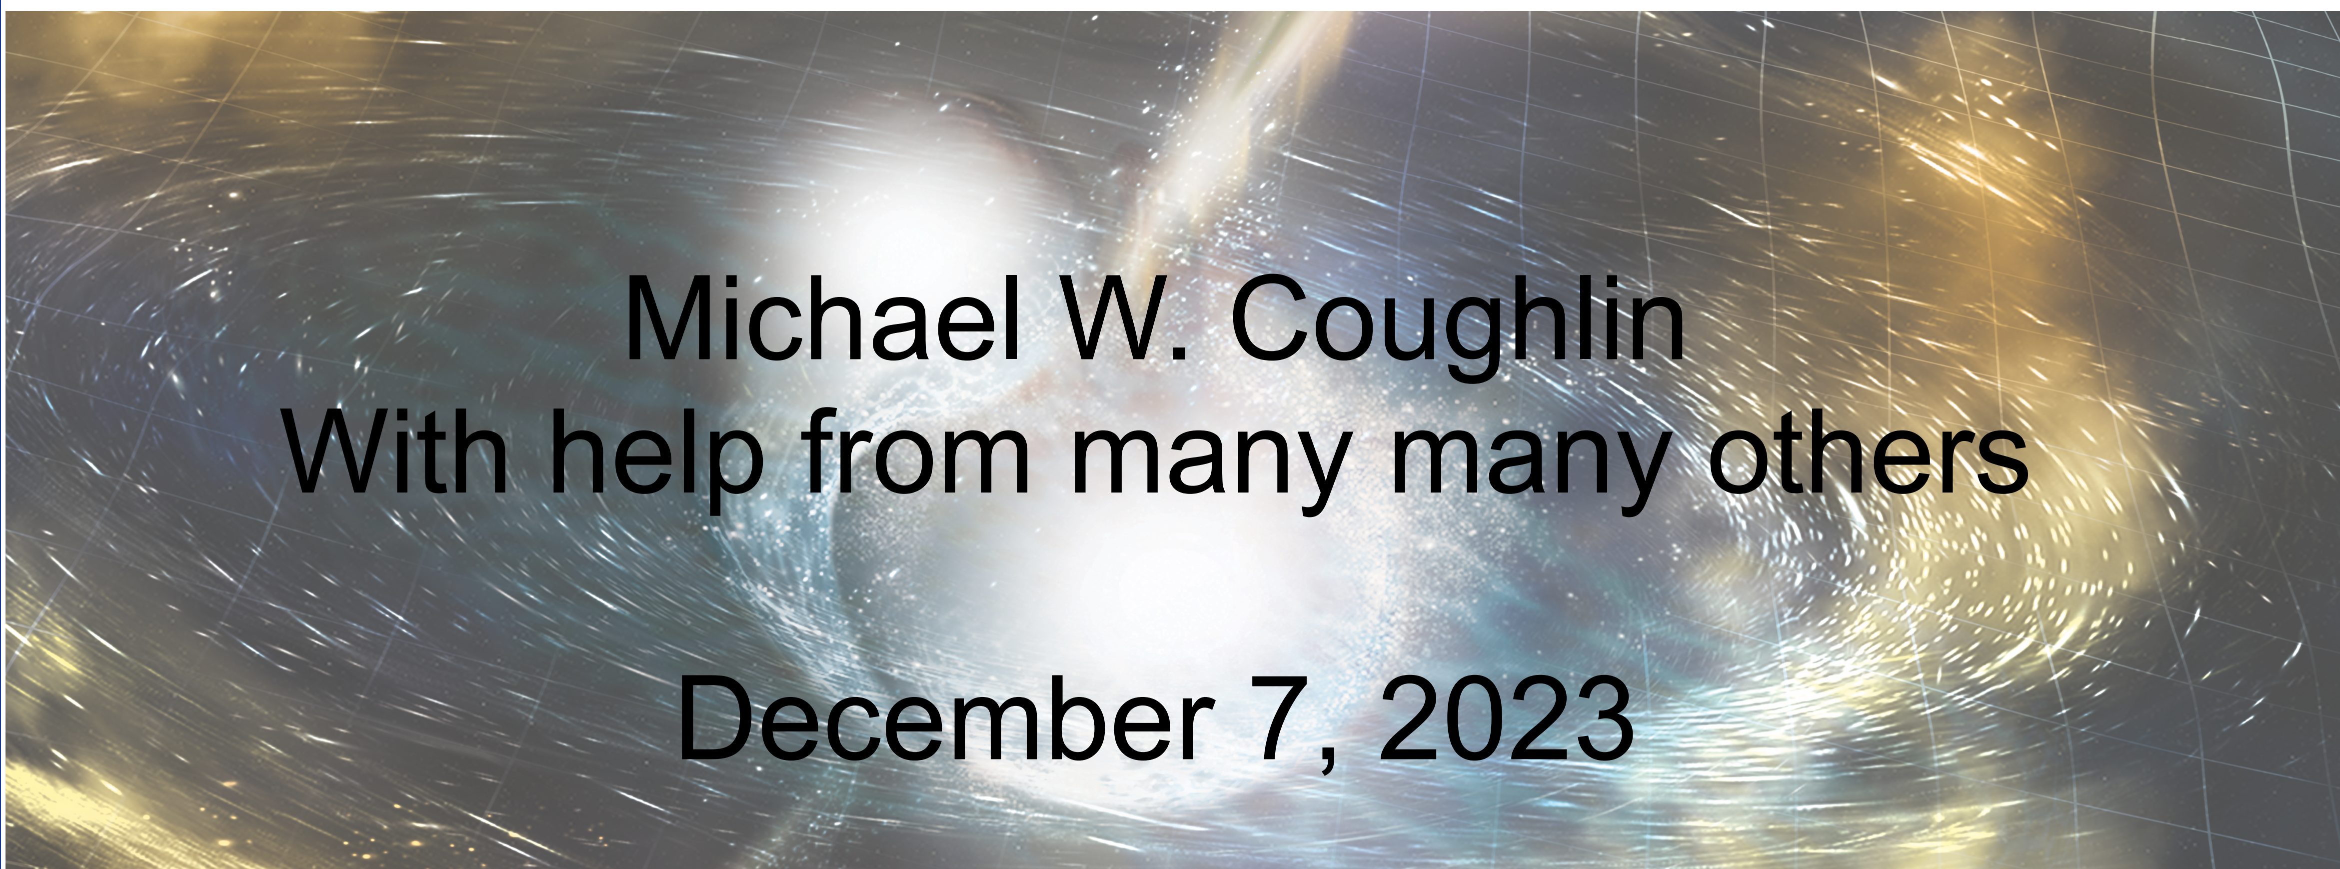

Michael W. Coughlin
With help from many many others

December 7, 2023

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

Kasliwal et al. (2020): 2008.00008

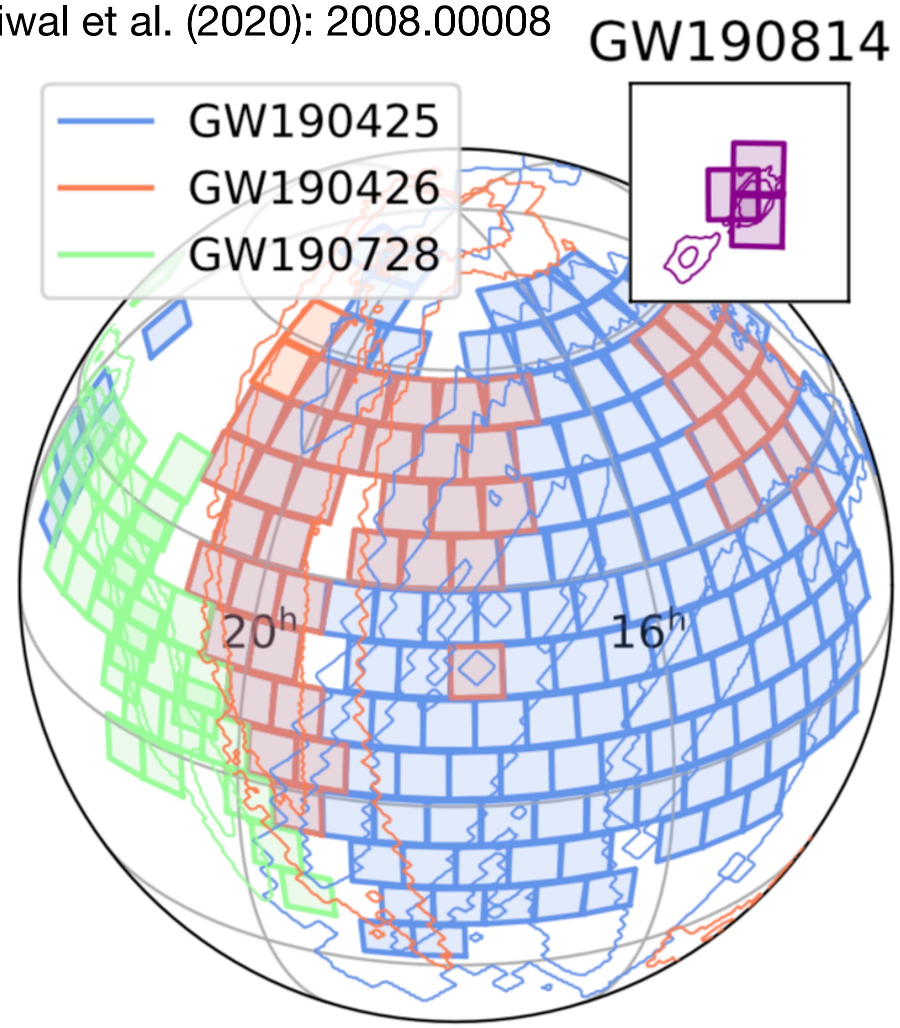

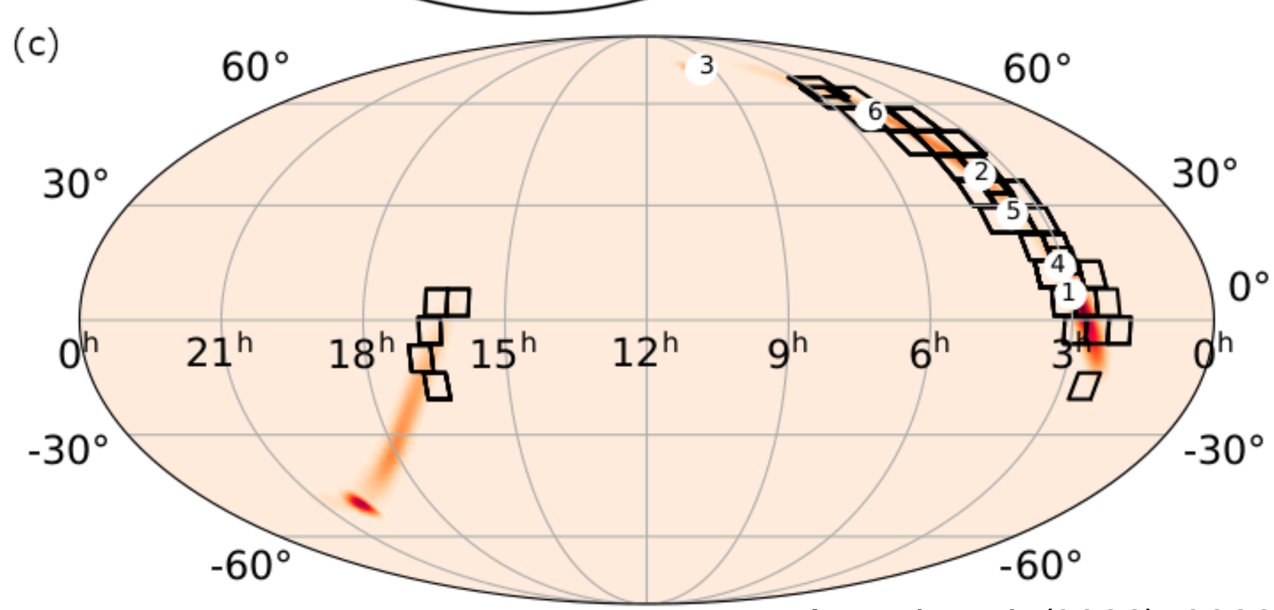

Anand et al. (2020): 2009.07210

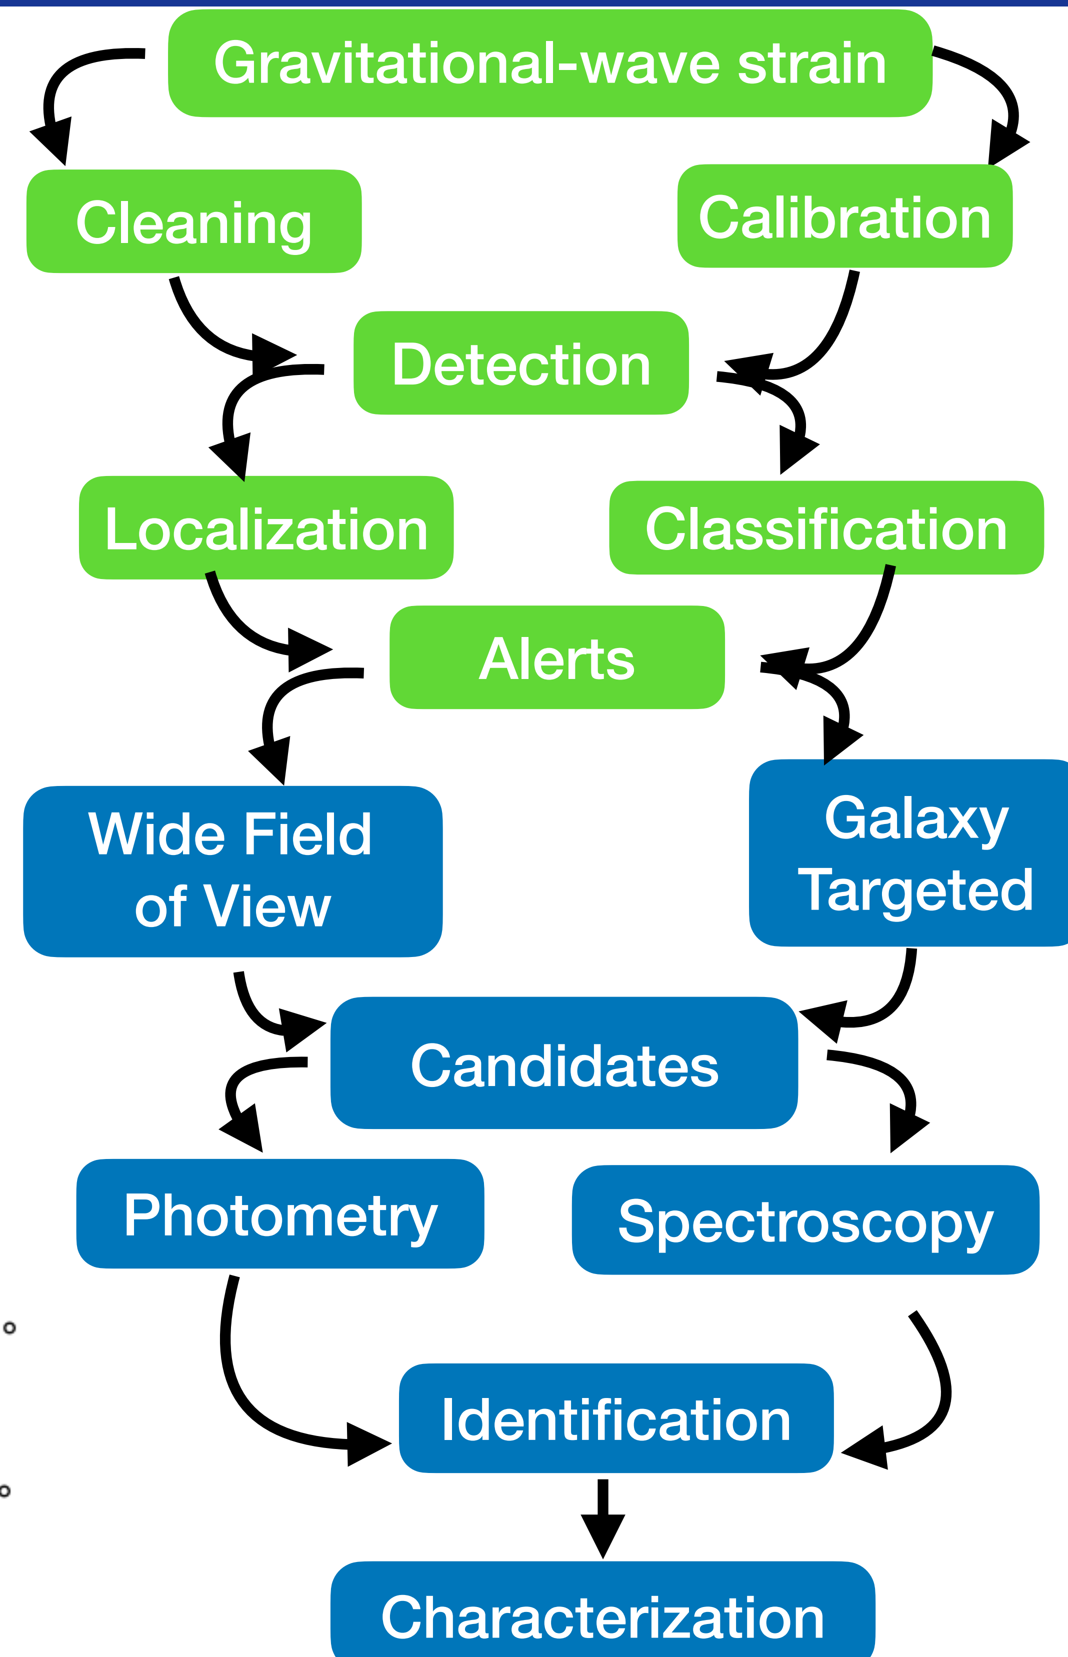

What did we learn from O3?

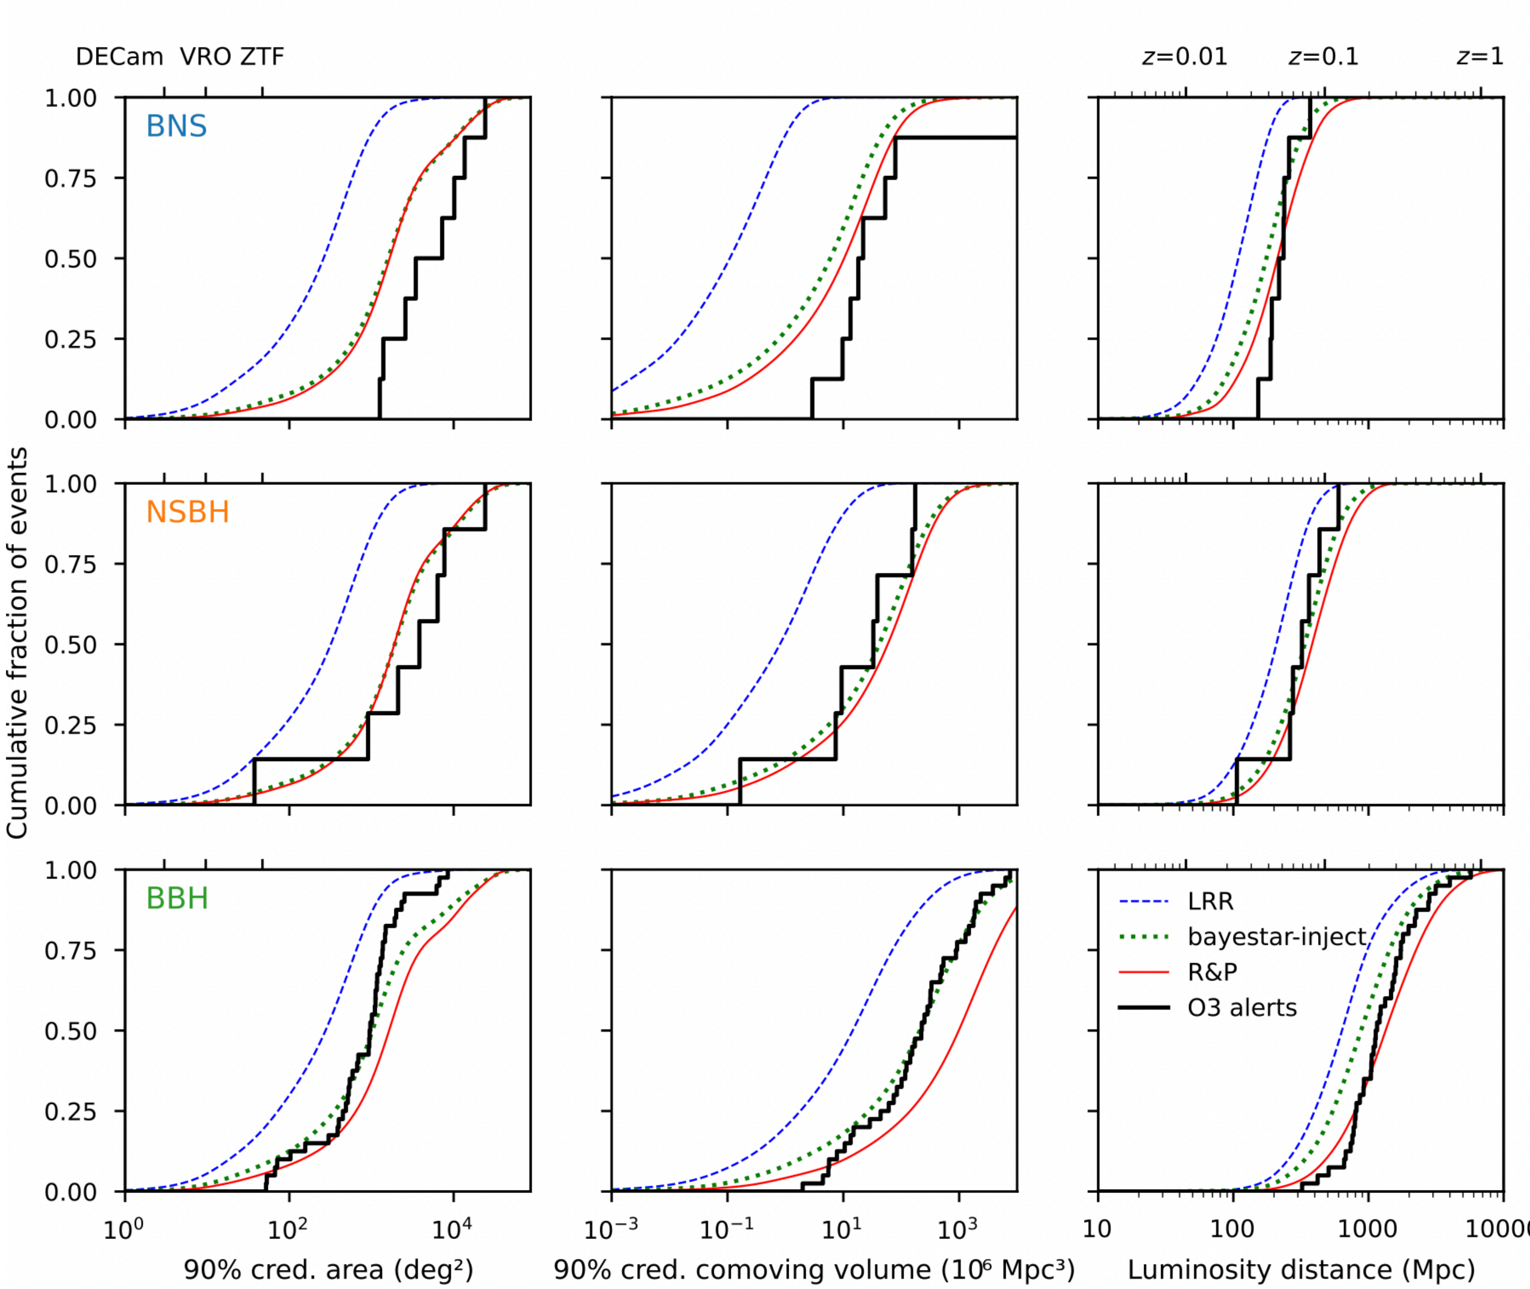

Weizmann Kiendrébéogo
Observatoire de la Côte d'Azure

[Petrov+2021, ApJ; Kiendrebeogo+, ApJ]

What are we learning from O4?

SNR	Network	BNS	NSBH	BBH
	Annual number of detections			
10	H measured PSD	1^{+4}_{-1}	0^{+0}_{-0}	21^{+29}_{-14}
	HL measured PSD	5^{+10}_{-5}	0^{+0}_{-0}	60^{+78}_{-36}
8	HL measured PSD	12^{+17}_{-9}	1^{+4}_{-2}	115^{+147}_{-67}
	HL ideal PSD	11^{+17}_{-8}	2^{+4}_{-2}	123^{+157}_{-71}
	HLVK	36^{+49}_{-22}	6^{+11}_{-5}	260^{+330}_{-150}

Weizmann Kiendrébéogo
Observatoire de la Côte d'Azure

[Petrov+2021, ApJ; Kiendrebeogo+, ApJ]

A Multi-Messenger Ecosystem

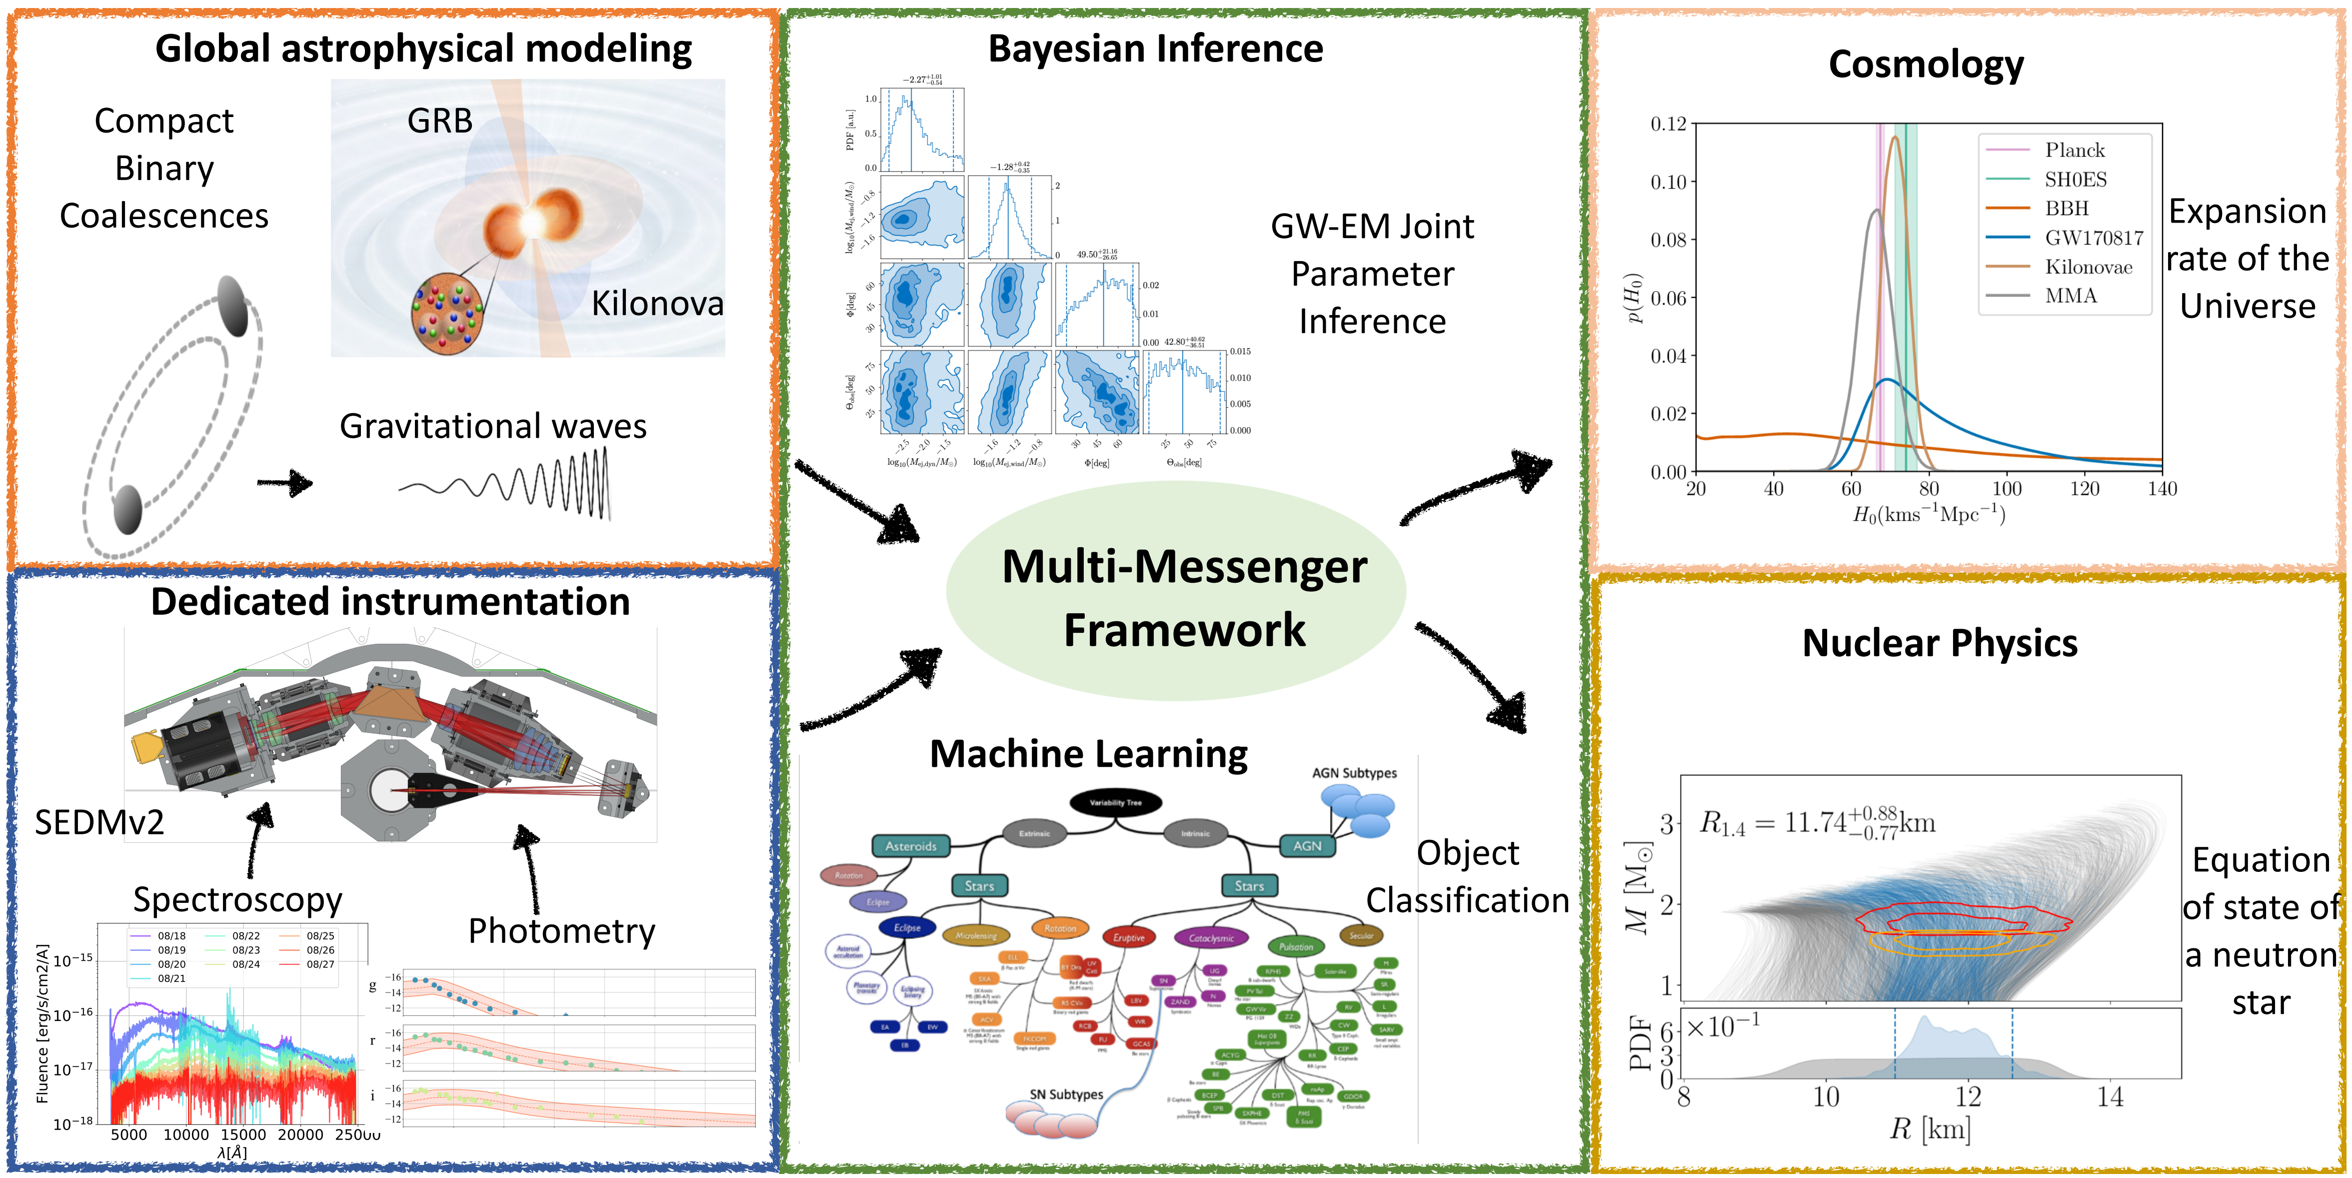

Gravitational Wave Data

Continuous **time series** (1Hz, 128Hz … 16kHz)

Gravitational Wave channel:
~20GB/day (per instrument)

Physical Environment Monitors (seismometers, accelerometers, magnetometers, microphones etc)

Internal Engineering Monitors (sensing, housekeeping, status etc)

Together with various intermediate data products >2TB/day (per instrument)

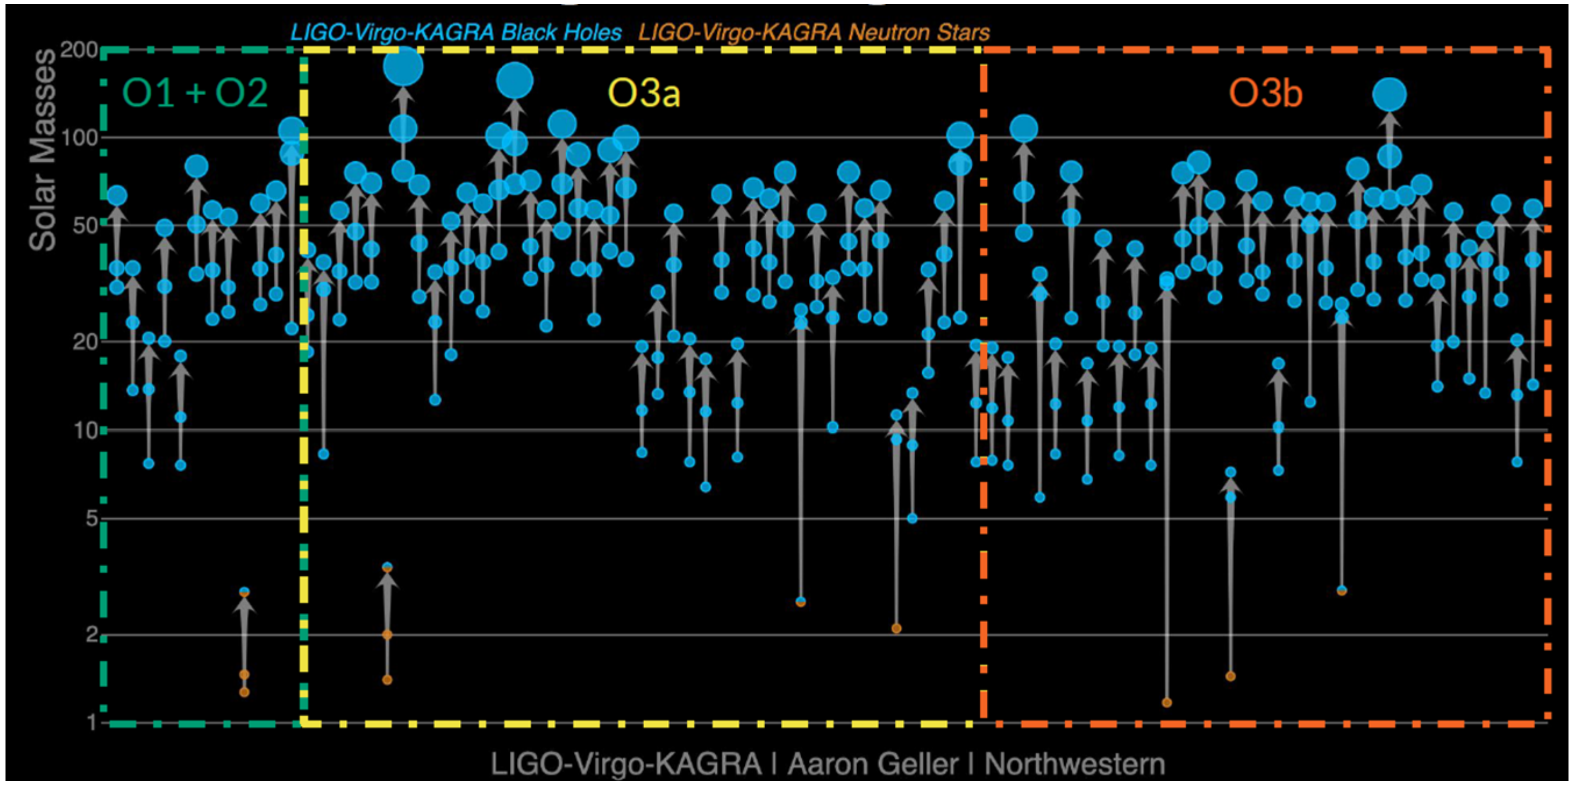

11 events from O1+O2

44 events in O3a, 55 total
1041 “subthreshold” events in O1,O2,O3a

35 events in O3b, 90 total (catalogs are cumulative)

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

Kasliwal et al. (2020): 2008.00008

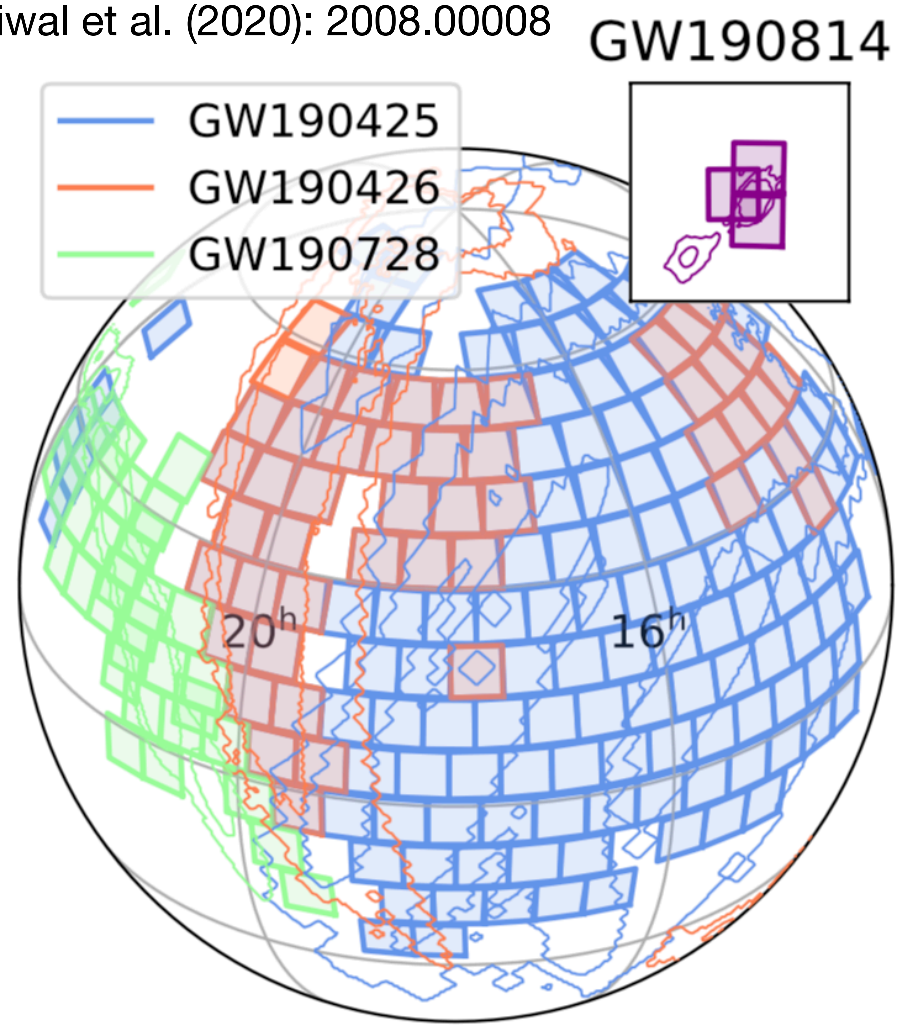

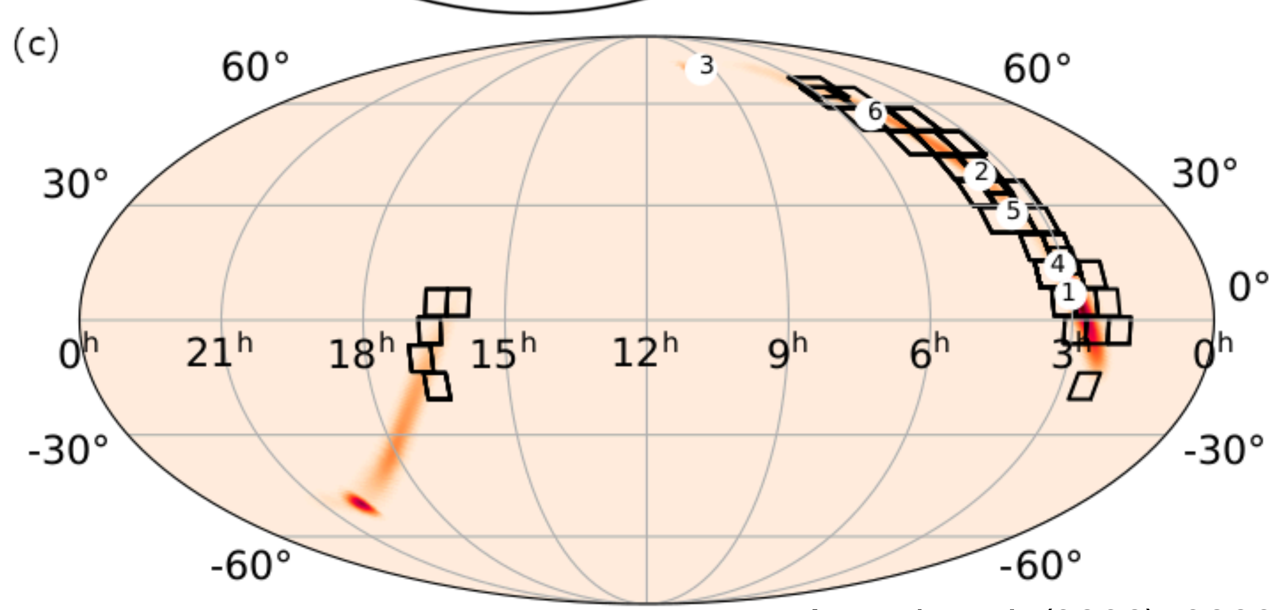

Anand et al. (2020): 2009.07210

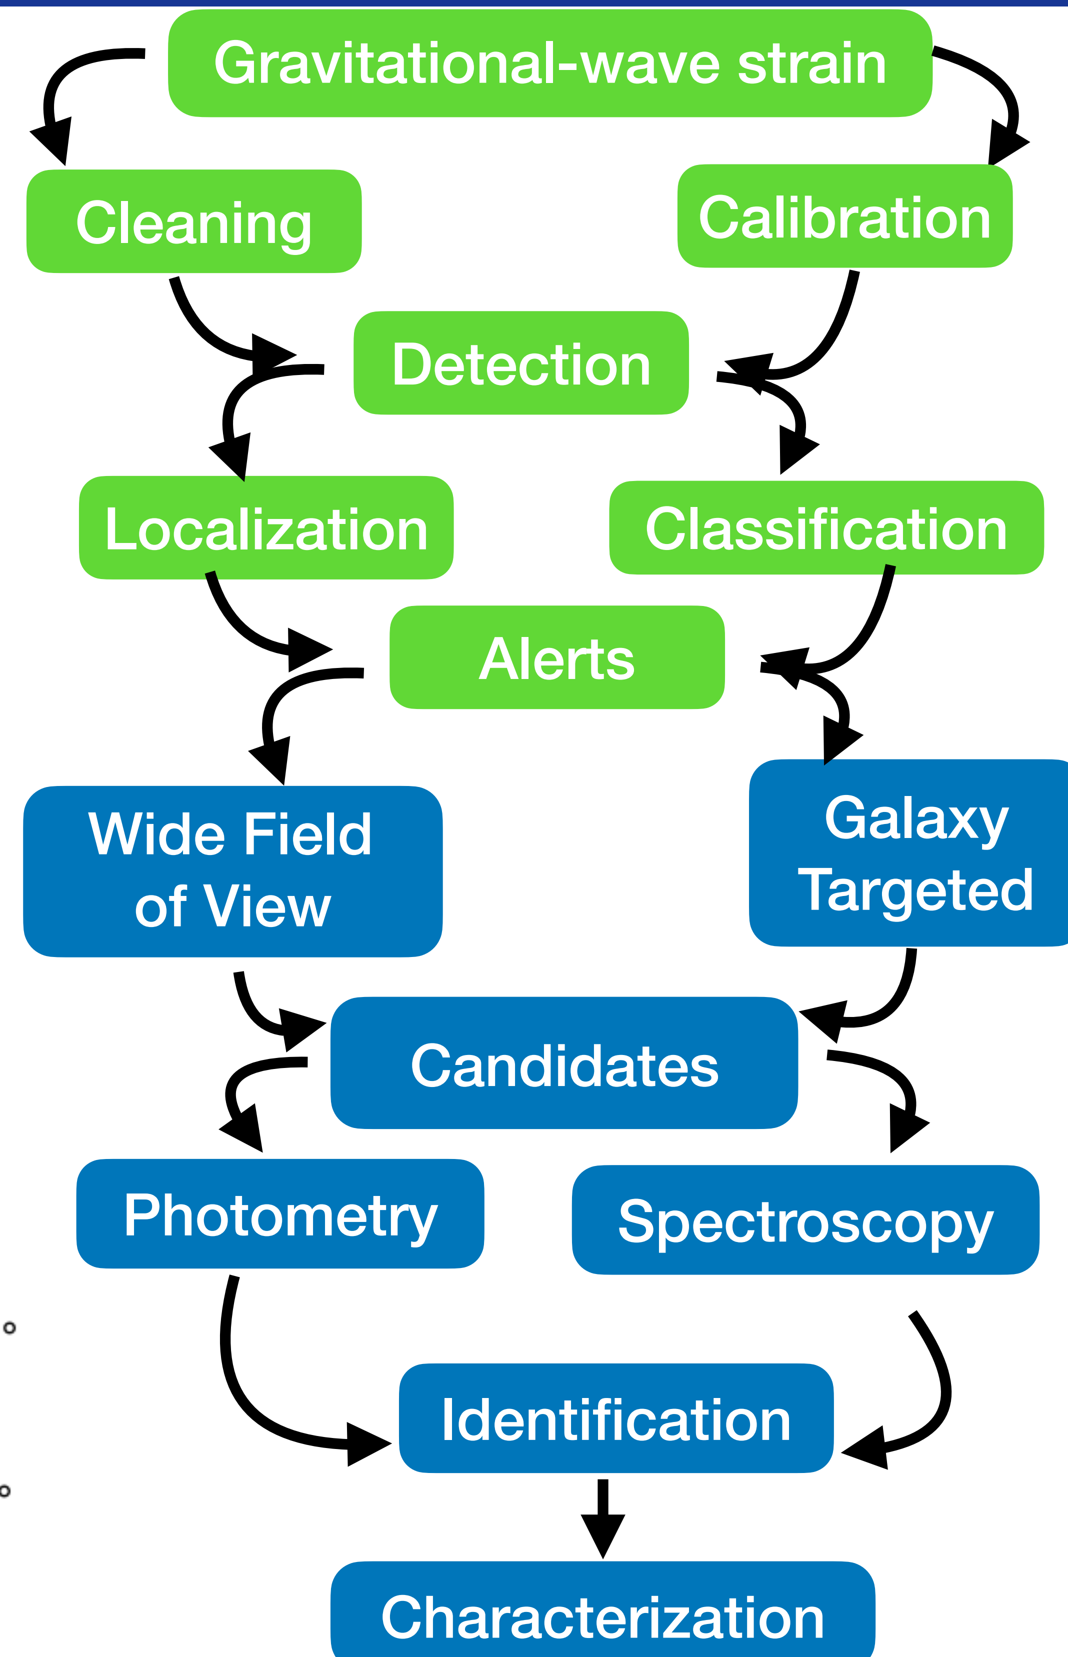

Gravitational Wave Alerts

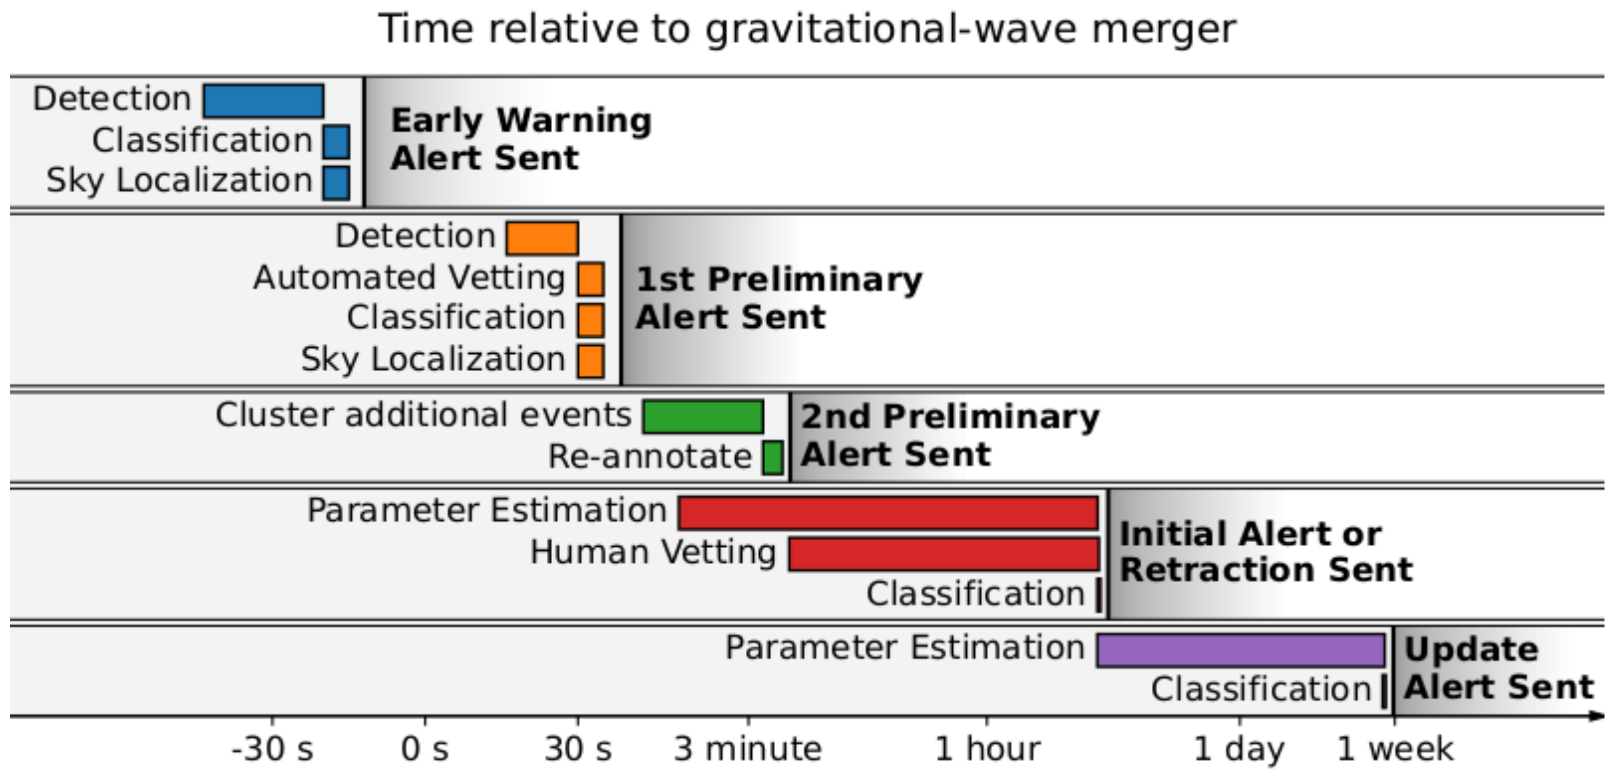

Tested the end-to-end alert infrastructure using a MDC (Sharma-Chaudhary & Toivonen et. al. 2308.04545)

Latencies In O4 so far

Median: 28s Max: 38s

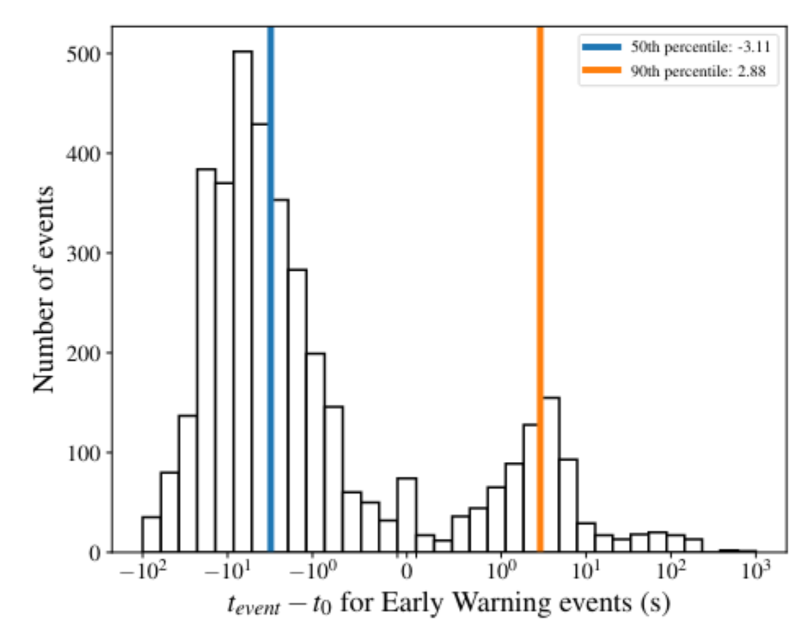

Machine Learning in Gravitational Waves

Detector Characterization

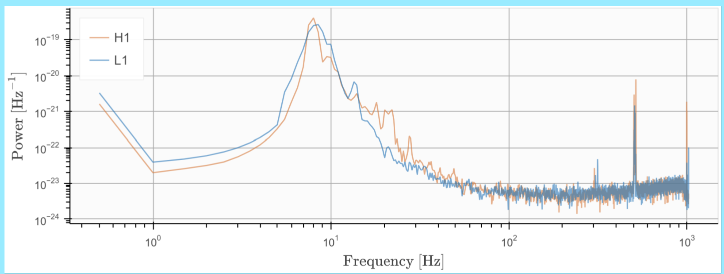

Feedback/controls reinforcement learning

Noise subtraction nonlinear regression

Event Detection

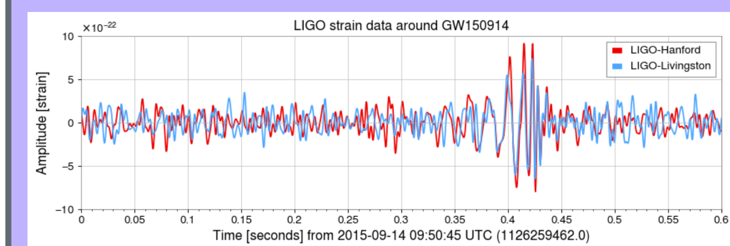

Modelled supervised

Unmodelled semi-supervised

Vetoes Glitch identification

Event Characterization

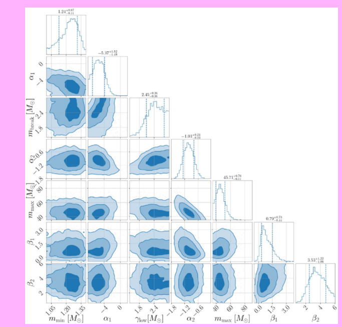

Parameter estimation Normalizing flows

Population modelling Normalizing flows (reprise)

Fast inference times, small compute footprint

Machine Learning in Gravitational Waves

Papers

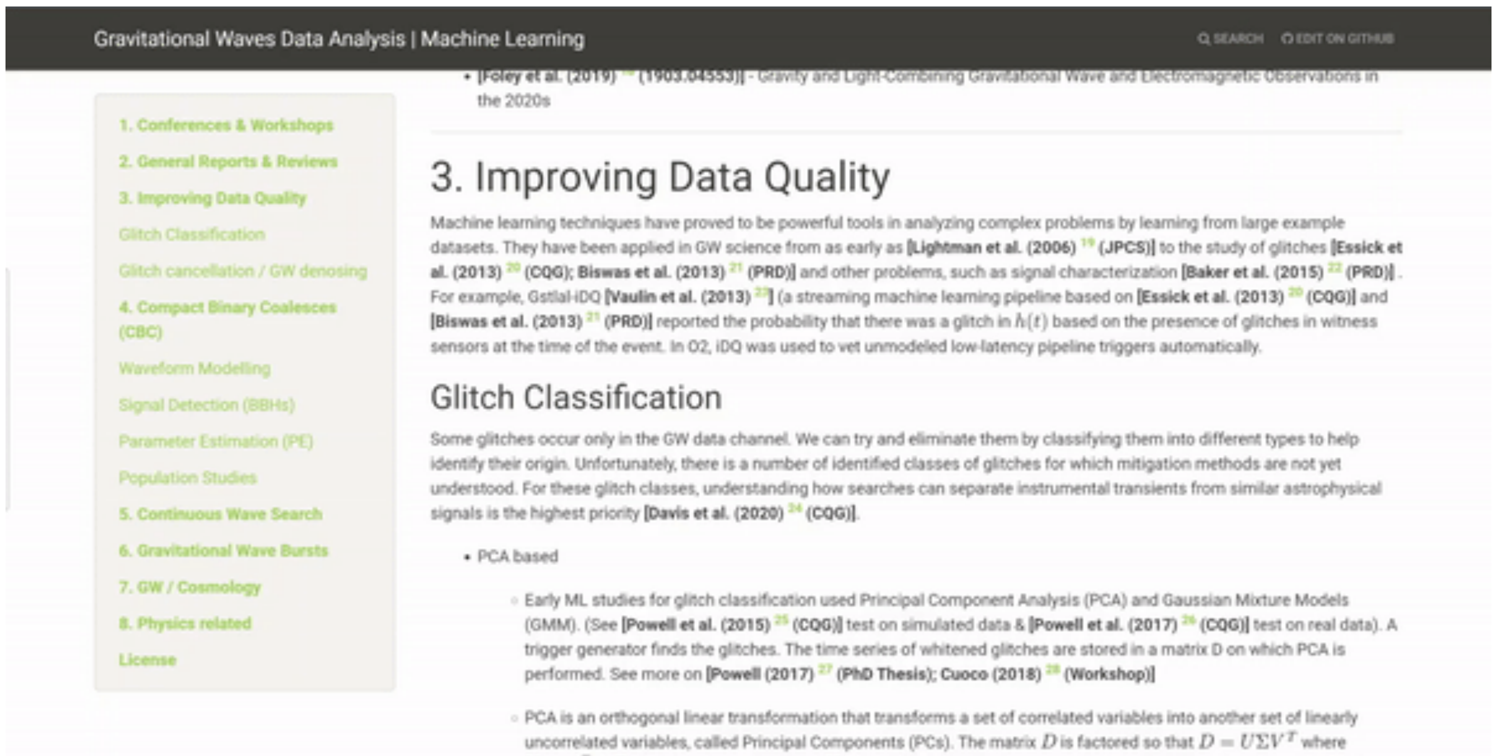
Gravitational Waves Data Analysis | Machine Learning

SEARCH EDIT ON GITHUB

- 1. Conferences & Workshops
- 2. General Reports & Reviews
- 3. Improving Data Quality
- Glitch Classification
- Glitch cancellation / GW denosing
- 4. Compact Binary Coalesces (CBC)
- Waveform Modelling
- Signal Detection (BBHs)
- Parameter Estimation (PE)
- Population Studies
- 5. Continuous Wave Search
- 6. Gravitational Wave Bursts
- 7. GW / Cosmology
- 8. Physics related
- License

- [Foley et al. (2019) [18] (1903.04553)] - Gravity and Light-Combining Gravitational Wave and Electromagnetic Observations in the 2020s

3. Improving Data Quality

Machine learning techniques have proved to be powerful tools in analyzing complex problems by learning from large example datasets. They have been applied in GW science from as early as **[Lightman et al. (2006)** [19] **(JPCS)]** to the study of glitches **[Essick et al. (2013)** [20] **(CQG); Biswas et al. (2013)** [21] **(PRD)]** and other problems, such as signal characterization **[Baker et al. (2015)** [22] **(PRD)]**. For example, GstLal-iDQ **[Vaulin et al. (2013)** [23]**]** (a streaming machine learning pipeline based on **[Essick et al. (2013)** [20] **(CQG)]** and **[Biswas et al. (2013)** [21] **(PRD)]** reported the probability that there was a glitch in $h(t)$ based on the presence of glitches in witness sensors at the time of the event. In O2, iDQ was used to vet unmodeled low-latency pipeline triggers automatically.

Glitch Classification

Some glitches occur only in the GW data channel. We can try and eliminate them by classifying them into different types to help identify their origin. Unfortunately, there is a number of identified classes of glitches for which mitigation methods are not yet understood. For these glitch classes, understanding how searches can separate instrumental transients from similar astrophysical signals is the highest priority **[Davis et al. (2020)** [24] **(CQG)]**.

- PCA based
  - Early ML studies for glitch classification used Principal Component Analysis (PCA) and Gaussian Mixture Models (GMM). (See **[Powell et al. (2015)** [25] **(CQG)]** test on simulated data & **[Powell et al. (2017)** [26] **(CQG)]** test on real data). A trigger generator finds the glitches. The time series of whitened glitches are stored in a matrix D on which PCA is performed. See more on **[Powell (2017)** [27] **(PhD Thesis); Cuoco (2018)** [28] **(Workshop)]**
  - PCA is an orthogonal linear transformation that transforms a set of correlated variables into another set of linearly uncorrelated variables, called Principal Components (PCs). The matrix D is factored so that $D = U\Sigma V^T$ where

Approaching production

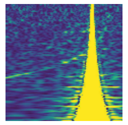

[GravitySpy]()

MLY PIPELINE [MLy]()

What else?

Machine Learning for Machine Learning Sake

ALGORITHMIA

2020 state of enterprise machine learning

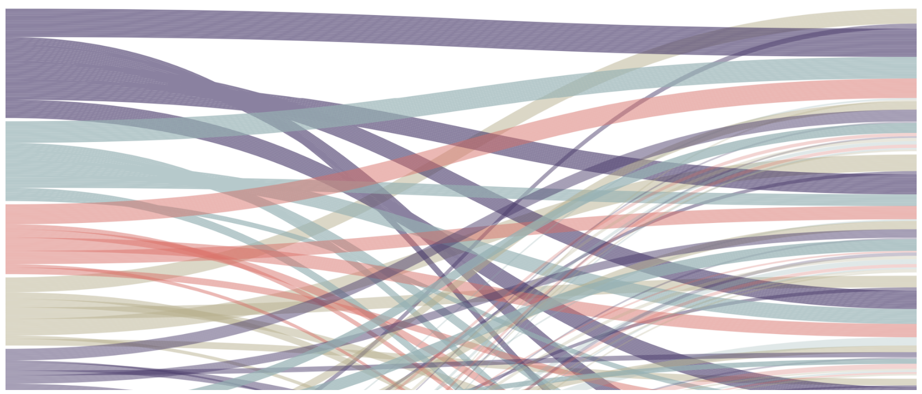

Key finding 3: Overcrowding at early maturity levels and AI for AI's sake

55% of companies surveyed have not deployed a machine learning model

9% more companies have gotten models into production since 2018

Model deployment timeline

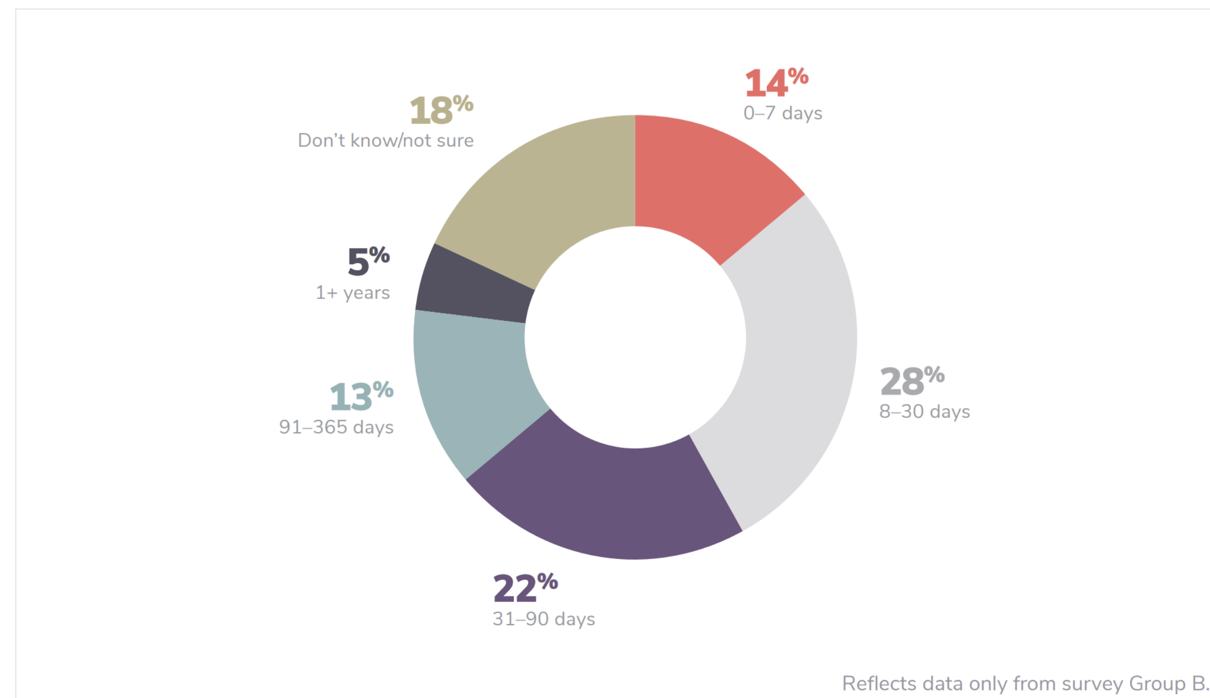

https://info.algorithmia.com/hubfs/2019/Whitepapers/The-State-of-Enterprise-ML-2020/Algorithmia_2020_State_of_Enterprise_ML.pdf

Are we having fun yet?

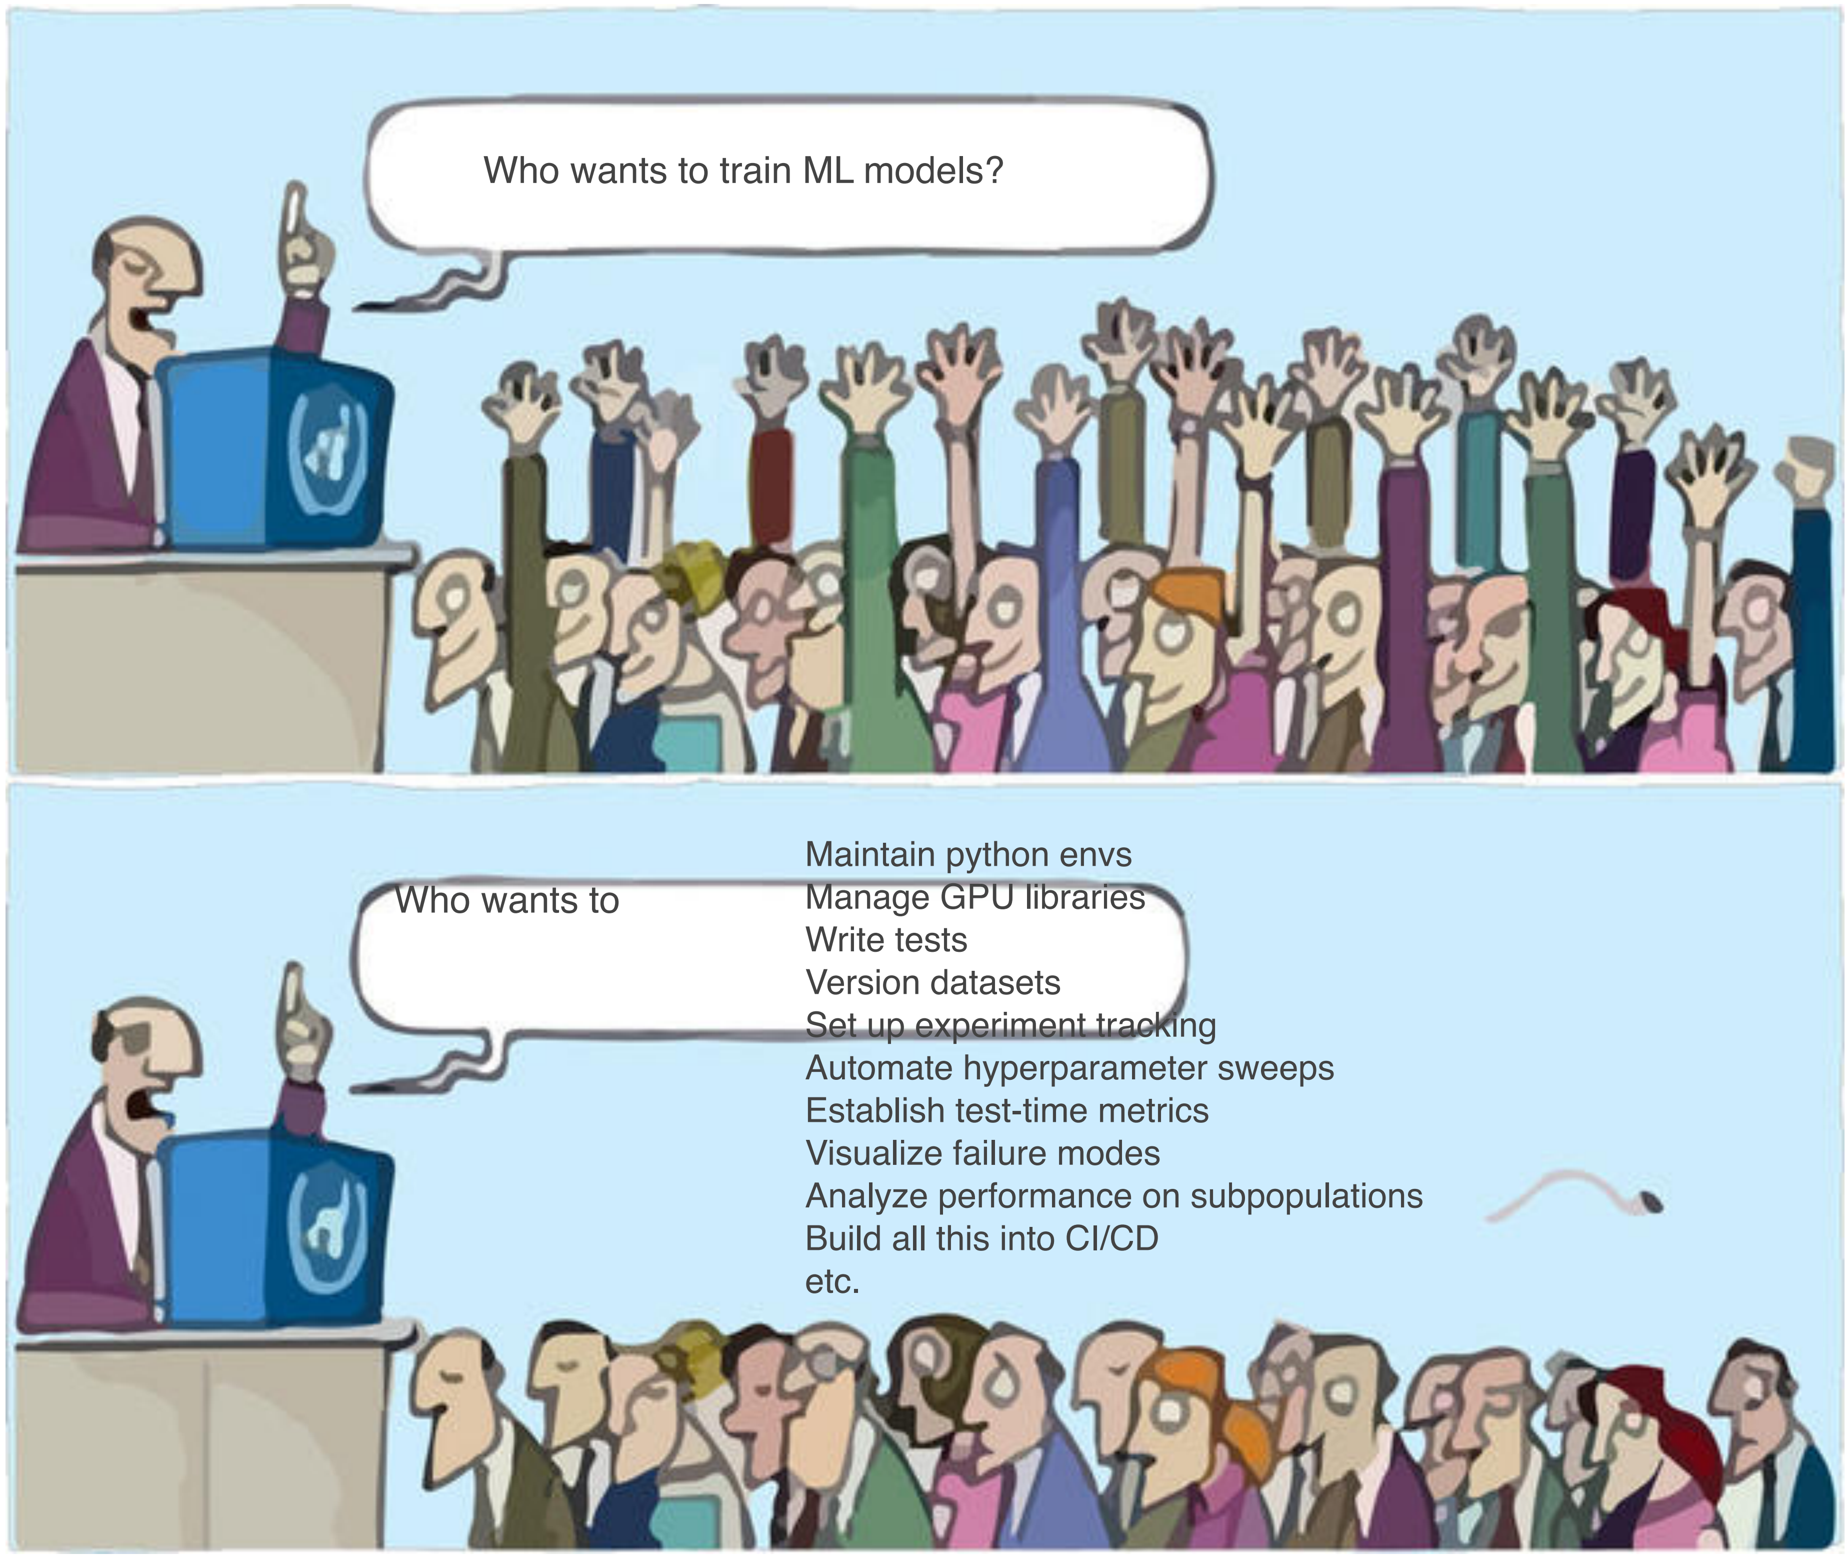

MLOps: Chicken and the Egg

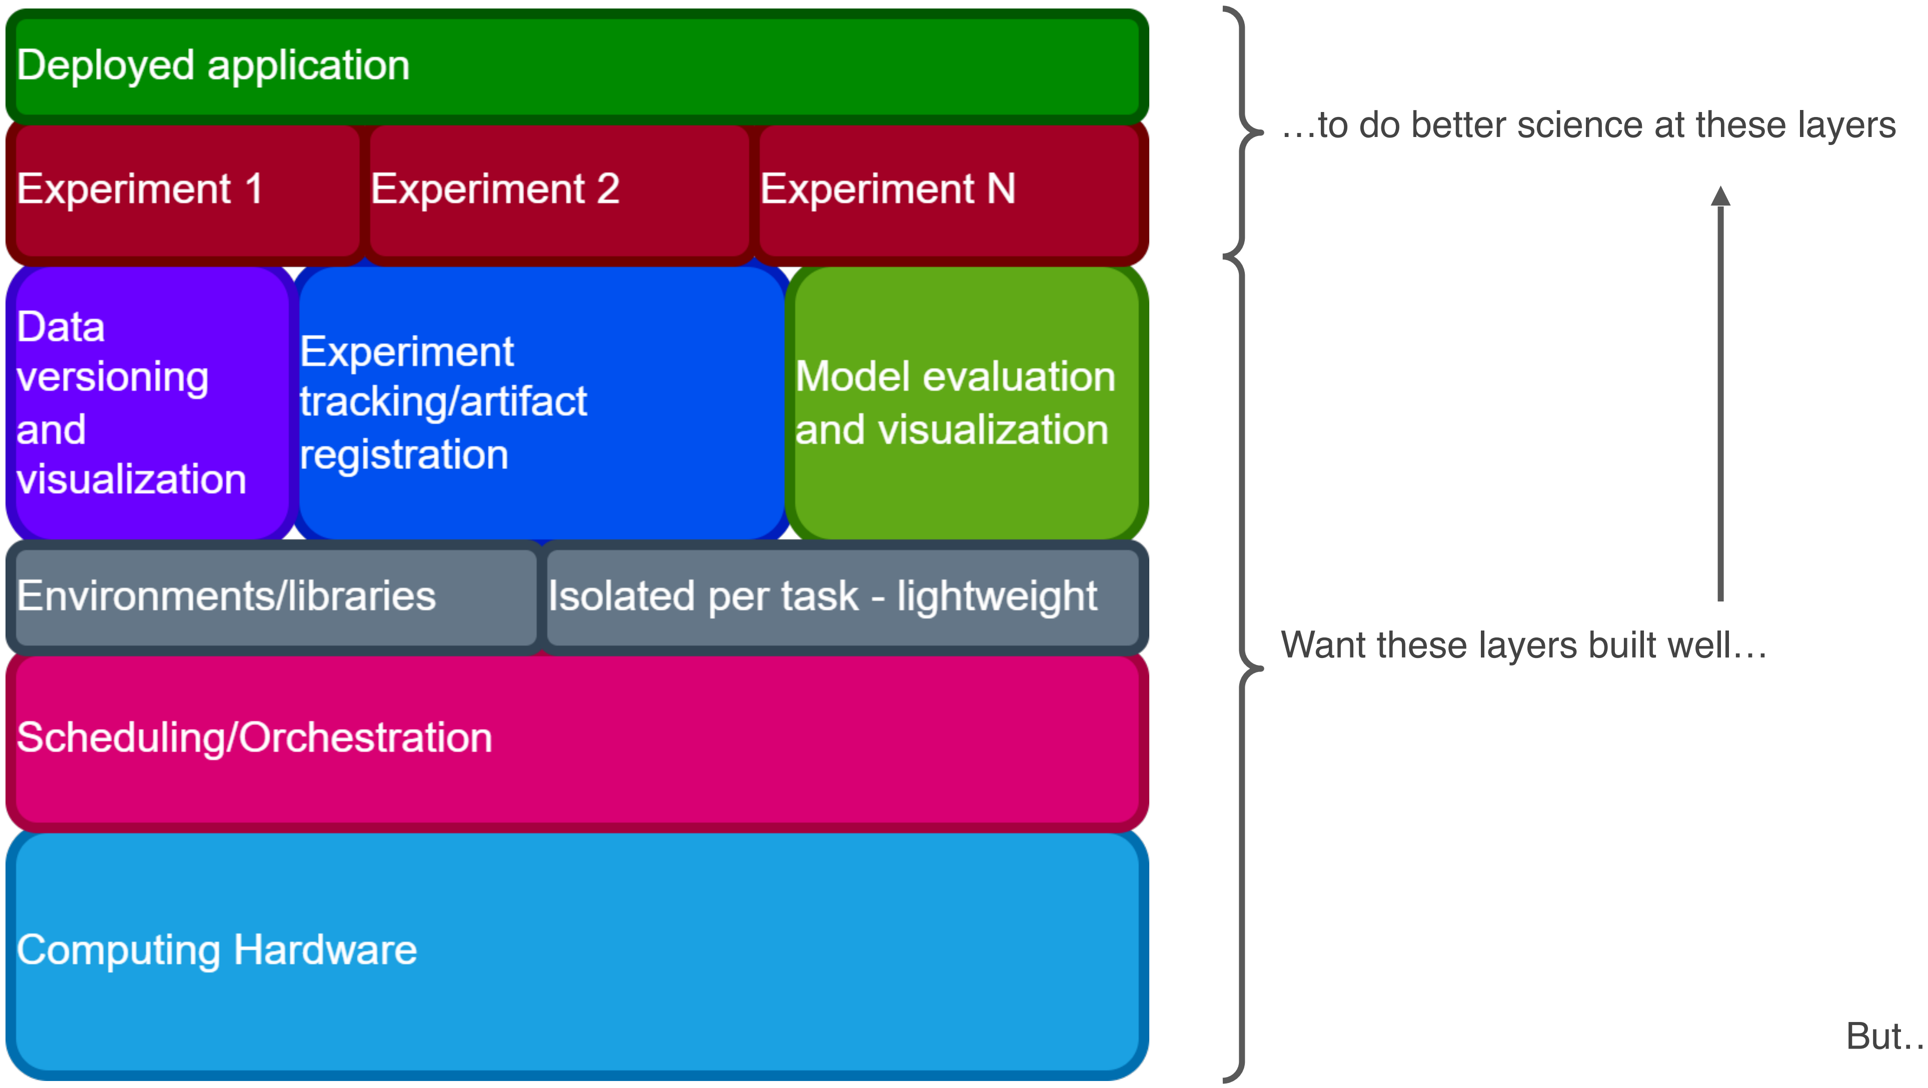

MLOps: Chicken and the Egg

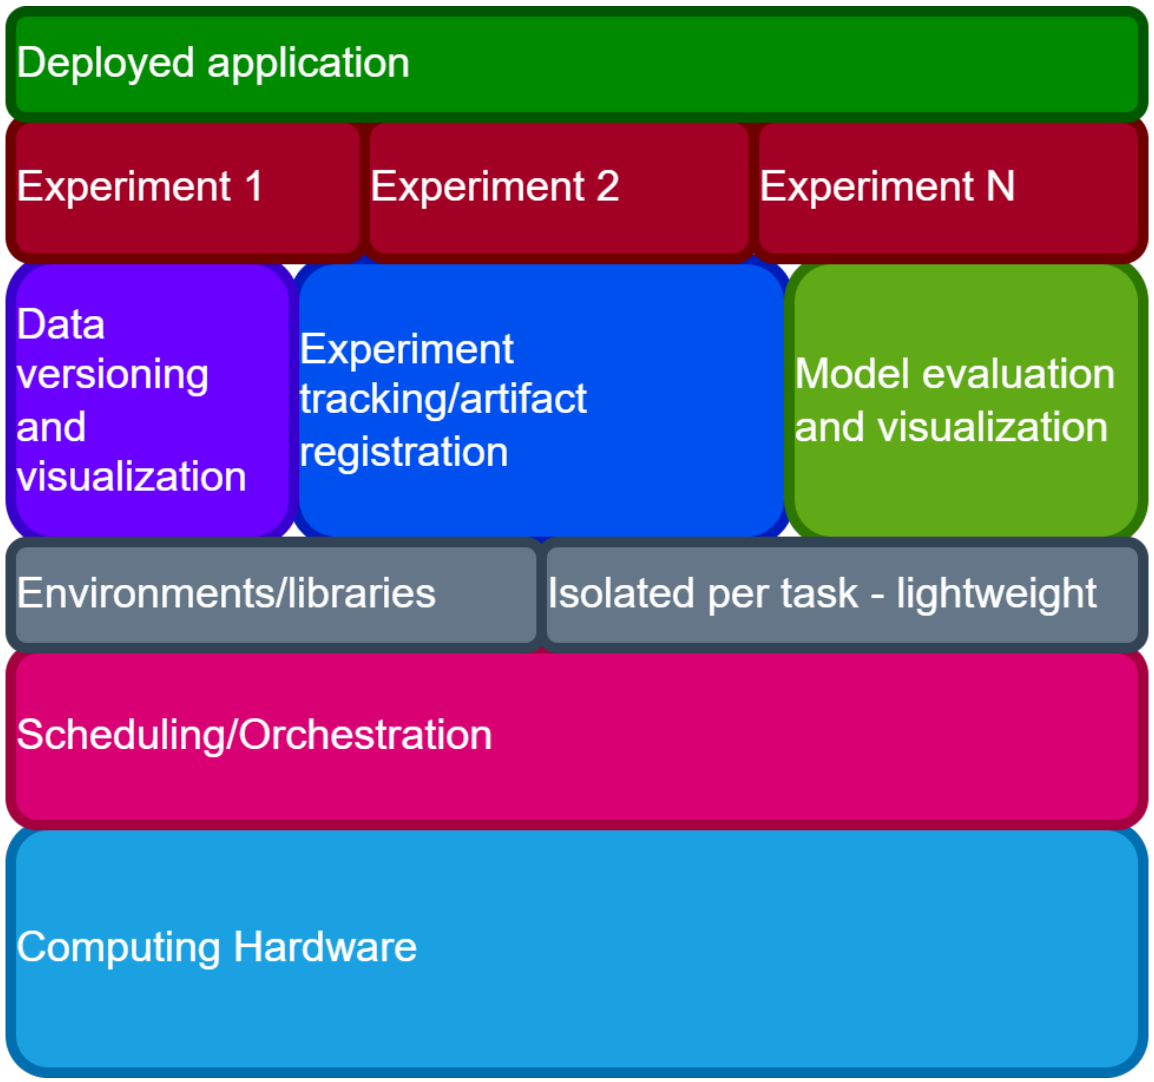

Needs here driven by science goals:

- What models (e.g. TD vs. FD)?
- What kinds of data, how much?
- Runtime latency/throughput
- Performance requirements

Drives needs at these layers

- Only discovered through experimentation
- Defined by domain experts
- Built by?

Frictionless Reproducibility

Ben Recht:

Our machine learning algorithm development is what Stephen Boyd calls "graduate student descent." Given the industrial interest, I think these days it's better designated "GitHub descent." Find a model on the internet, tweak a parameter or two, see if it gets better test error. If it does, that's a paper. We're most definitely optimizing, as we really care about these competitions on dataset benchmarks. But our algorithm is some sort of massively parallel genetic algorithm, not a clean, rigorous, and beautiful convex optimization method.

https://argmin.substack.com/p/rigor-vs-github-descent

https://arxiv.org/abs/2310.00865

Data Science at the Singularity
Version 1.00

David Donoho *

October 3, 2023

[**FR-1: Data**] datafication of everything, with a culture of research data sharing. One can now find datasets publicly available online on a bewildering variety of topics, from chest x-rays to cosmic microwave background measurements to uber routes to geospatial crop identifications.

[**FR-2: Re-execution**] research code sharing including the ability to exactly re-execute the same complete workflow by different researchers.

[**FR-3: Challenges**] adopting challenge problems as a new paradigm powering scientific research. The paradigm includes: a shared public dataset, a prescribed and quantified task performance metric, a set of enrolled competitors seeking to outperform each other on the task, and a public leaderboard. Thousands of such challenges with millions of entries have now taken place, across many fields.

The maturation of [FR-1]+[FR-2]+[FR-3] and emergence of FRX did not spring out of a vaccum. Nor out of slidedecks presented to Silicon Valley VC's who funded Github, Kaggle, and Hugging Face. Nor out of today's hegemon research labs. Rather, they developed organically from efforts by data scientists and technologists across at least 4 decades, witnessed by my own eyes.

Hacking Graduate Student Descent with Better Tools

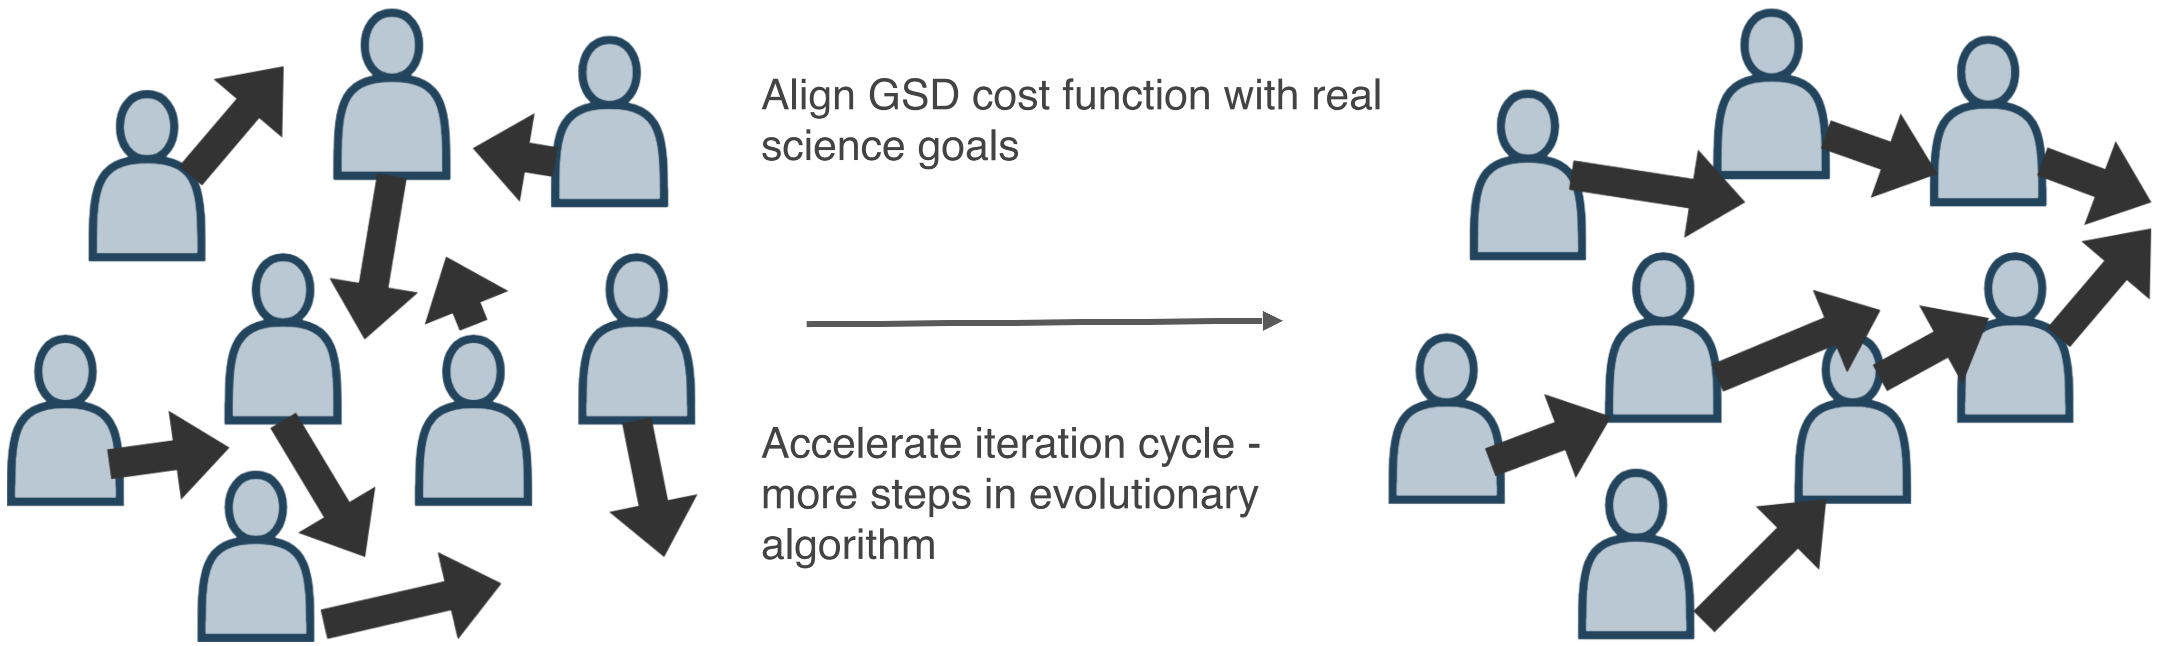

High-level requirements

- Lower barriers to entry
- Allow users to focus on physics
- Produce real *knowledge* - is this method better or not?

More specifically

- Modular, composable
- Successive layers of abstraction
- Fast, efficient usage of heterogeneous computing resources

An Example Use Case

- Deployed application
- Experiment 1 | Experiment 2 | Experiment N
- Data versioning and visualization | Experiment tracking/artifact registration | Model evaluation and visualization
- Environments/libraries | Isolated per task - lightweight
- Scheduling/Orchestration
- Computing Hardware

aframe - Binary black hole merger detection

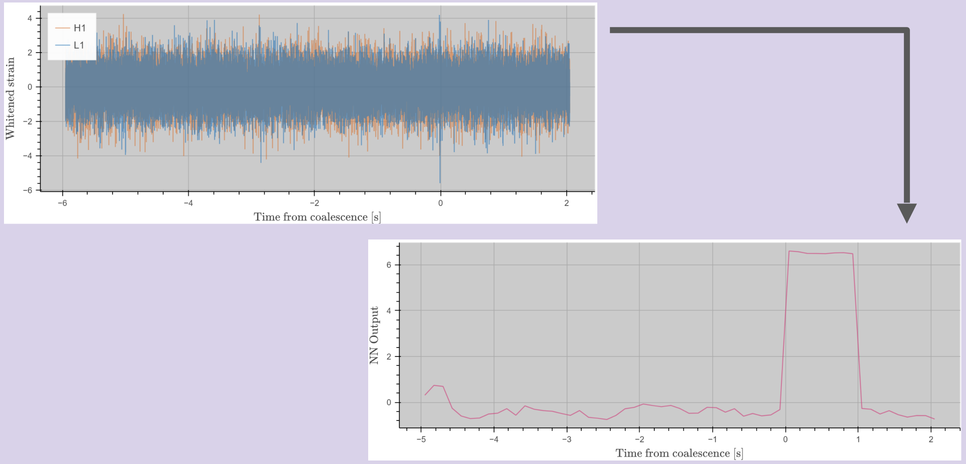

Simple supervised binary classification

99% of work has been physics/engineering

Finally nearing position to do good ML research at scale - what's next?

Paper with SOA sensitivity in the works

Machine Learning: ML4GW

Tools and Infrastructure

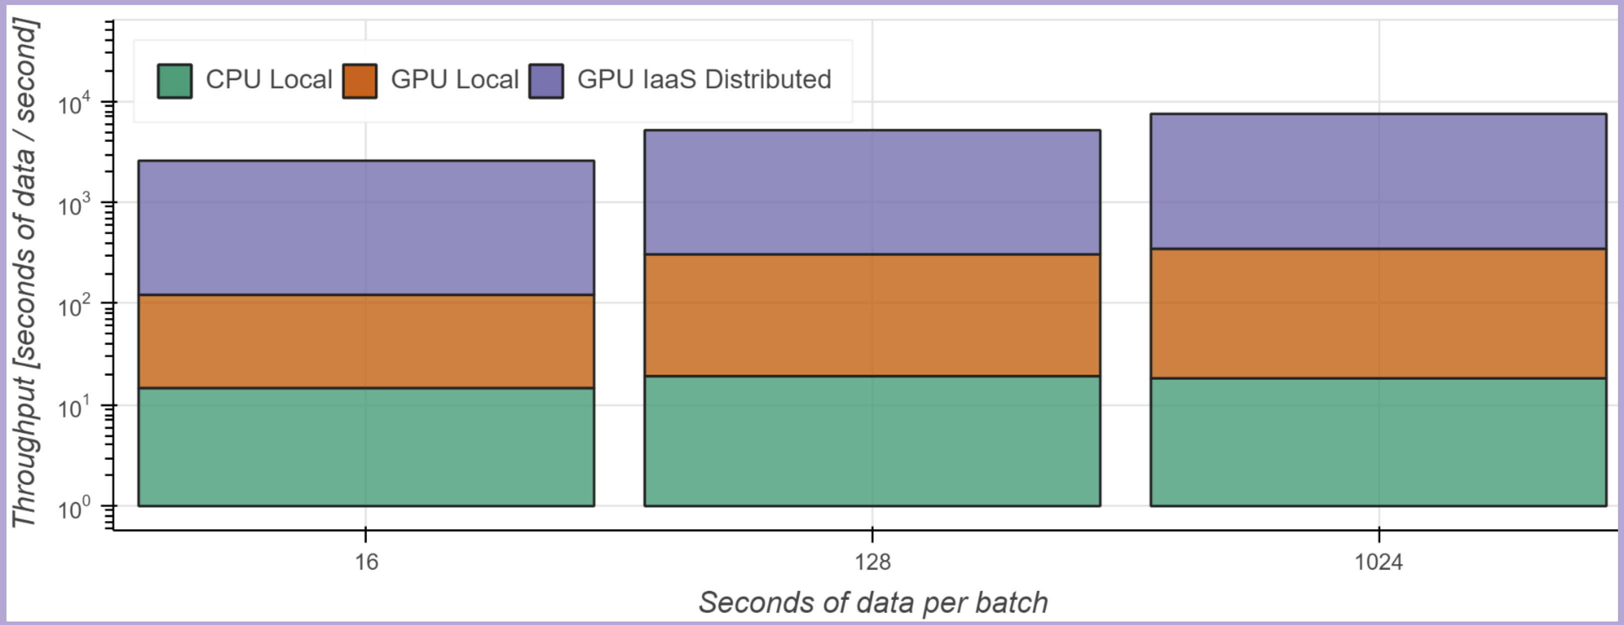

hermes: extending inference-as-a-service prototype to scale up model inference/validation

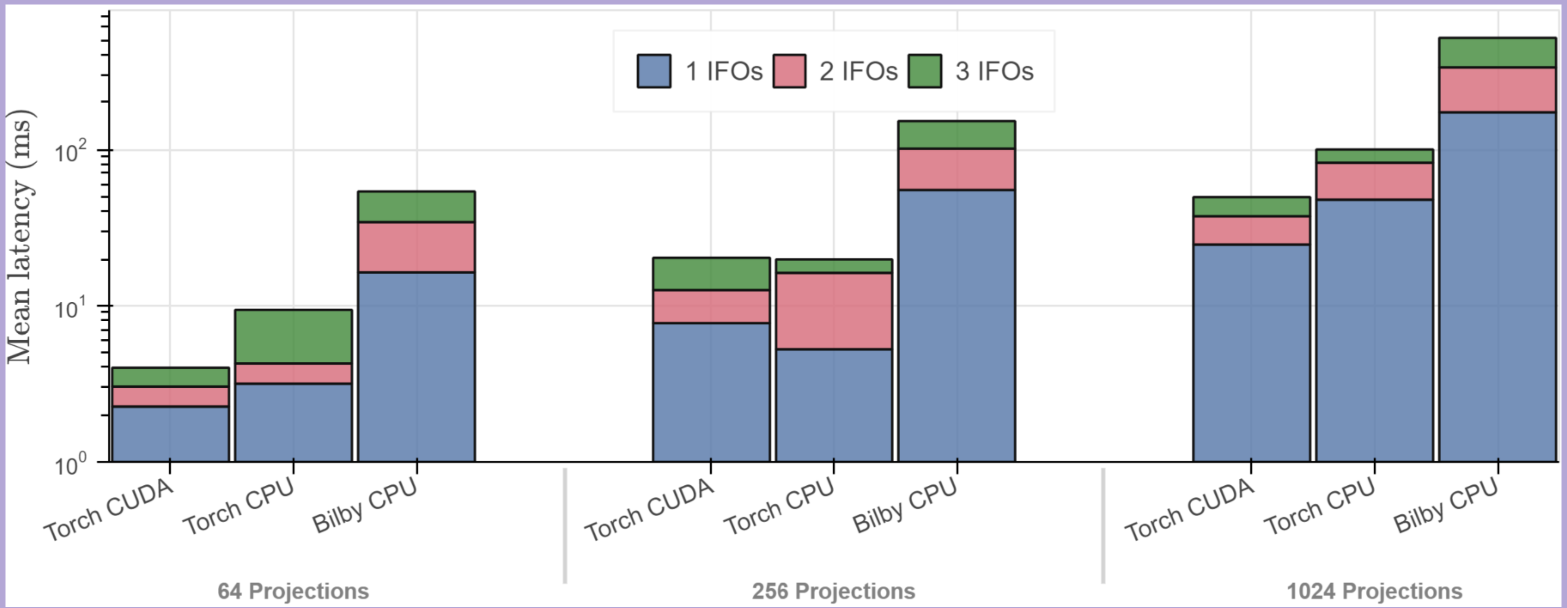

ml4gw: GPU-efficient Pytorch implementations of common GW ops, e.g. waveform → IFO response

Top-level github work area: ML4GW

Clean the Data: DeepClean (CNN)

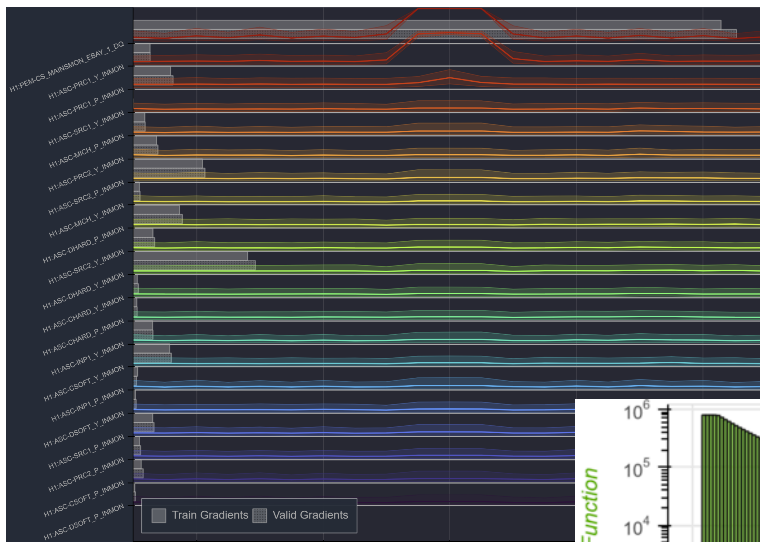

Detect the GWs: aframe (CNN)/GWAK (autoencoders)

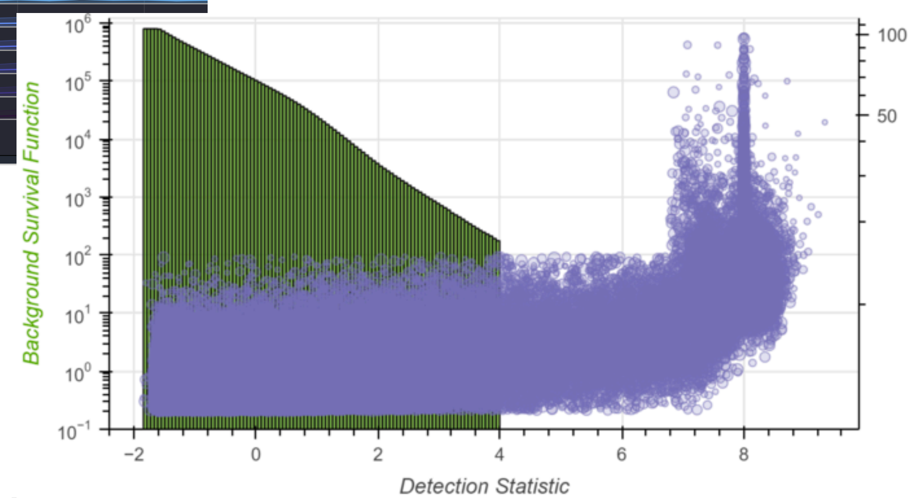

Characterize the GWs: (MAF*)

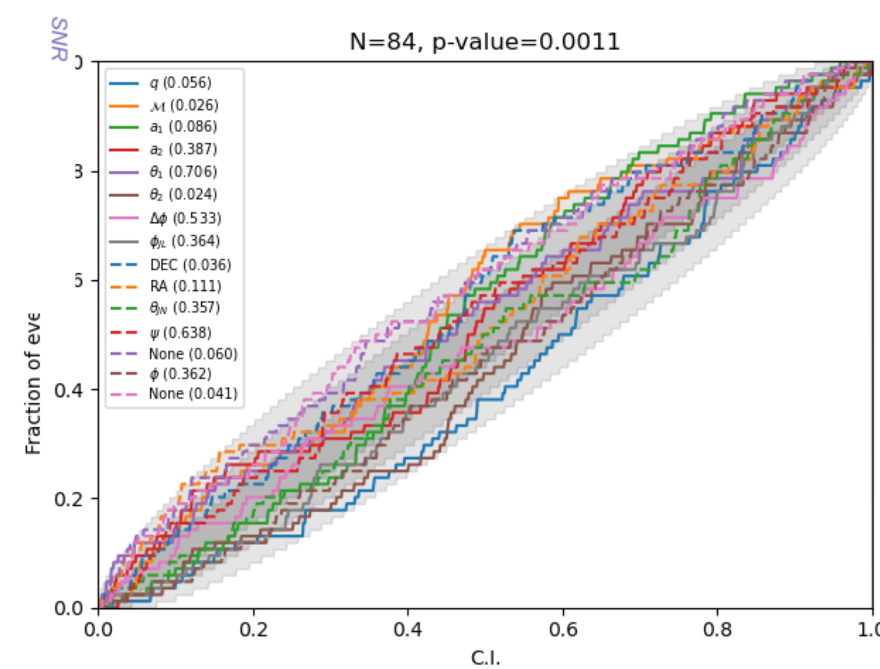

MAF: Masked Autoregressive FLow

Main focus: All algorithms use our inference-as-a-service (IaaS) prototype to implement a real-time noise subtraction pipeline (DeepClean), detection (aframe/GWAK), and parameter estimation for use during the fourth observing run (O4) of LIGO-Virgo-KAGRA on dedicated hardware at the detector sites.

The Time-Domain Astronomy Ecosystem

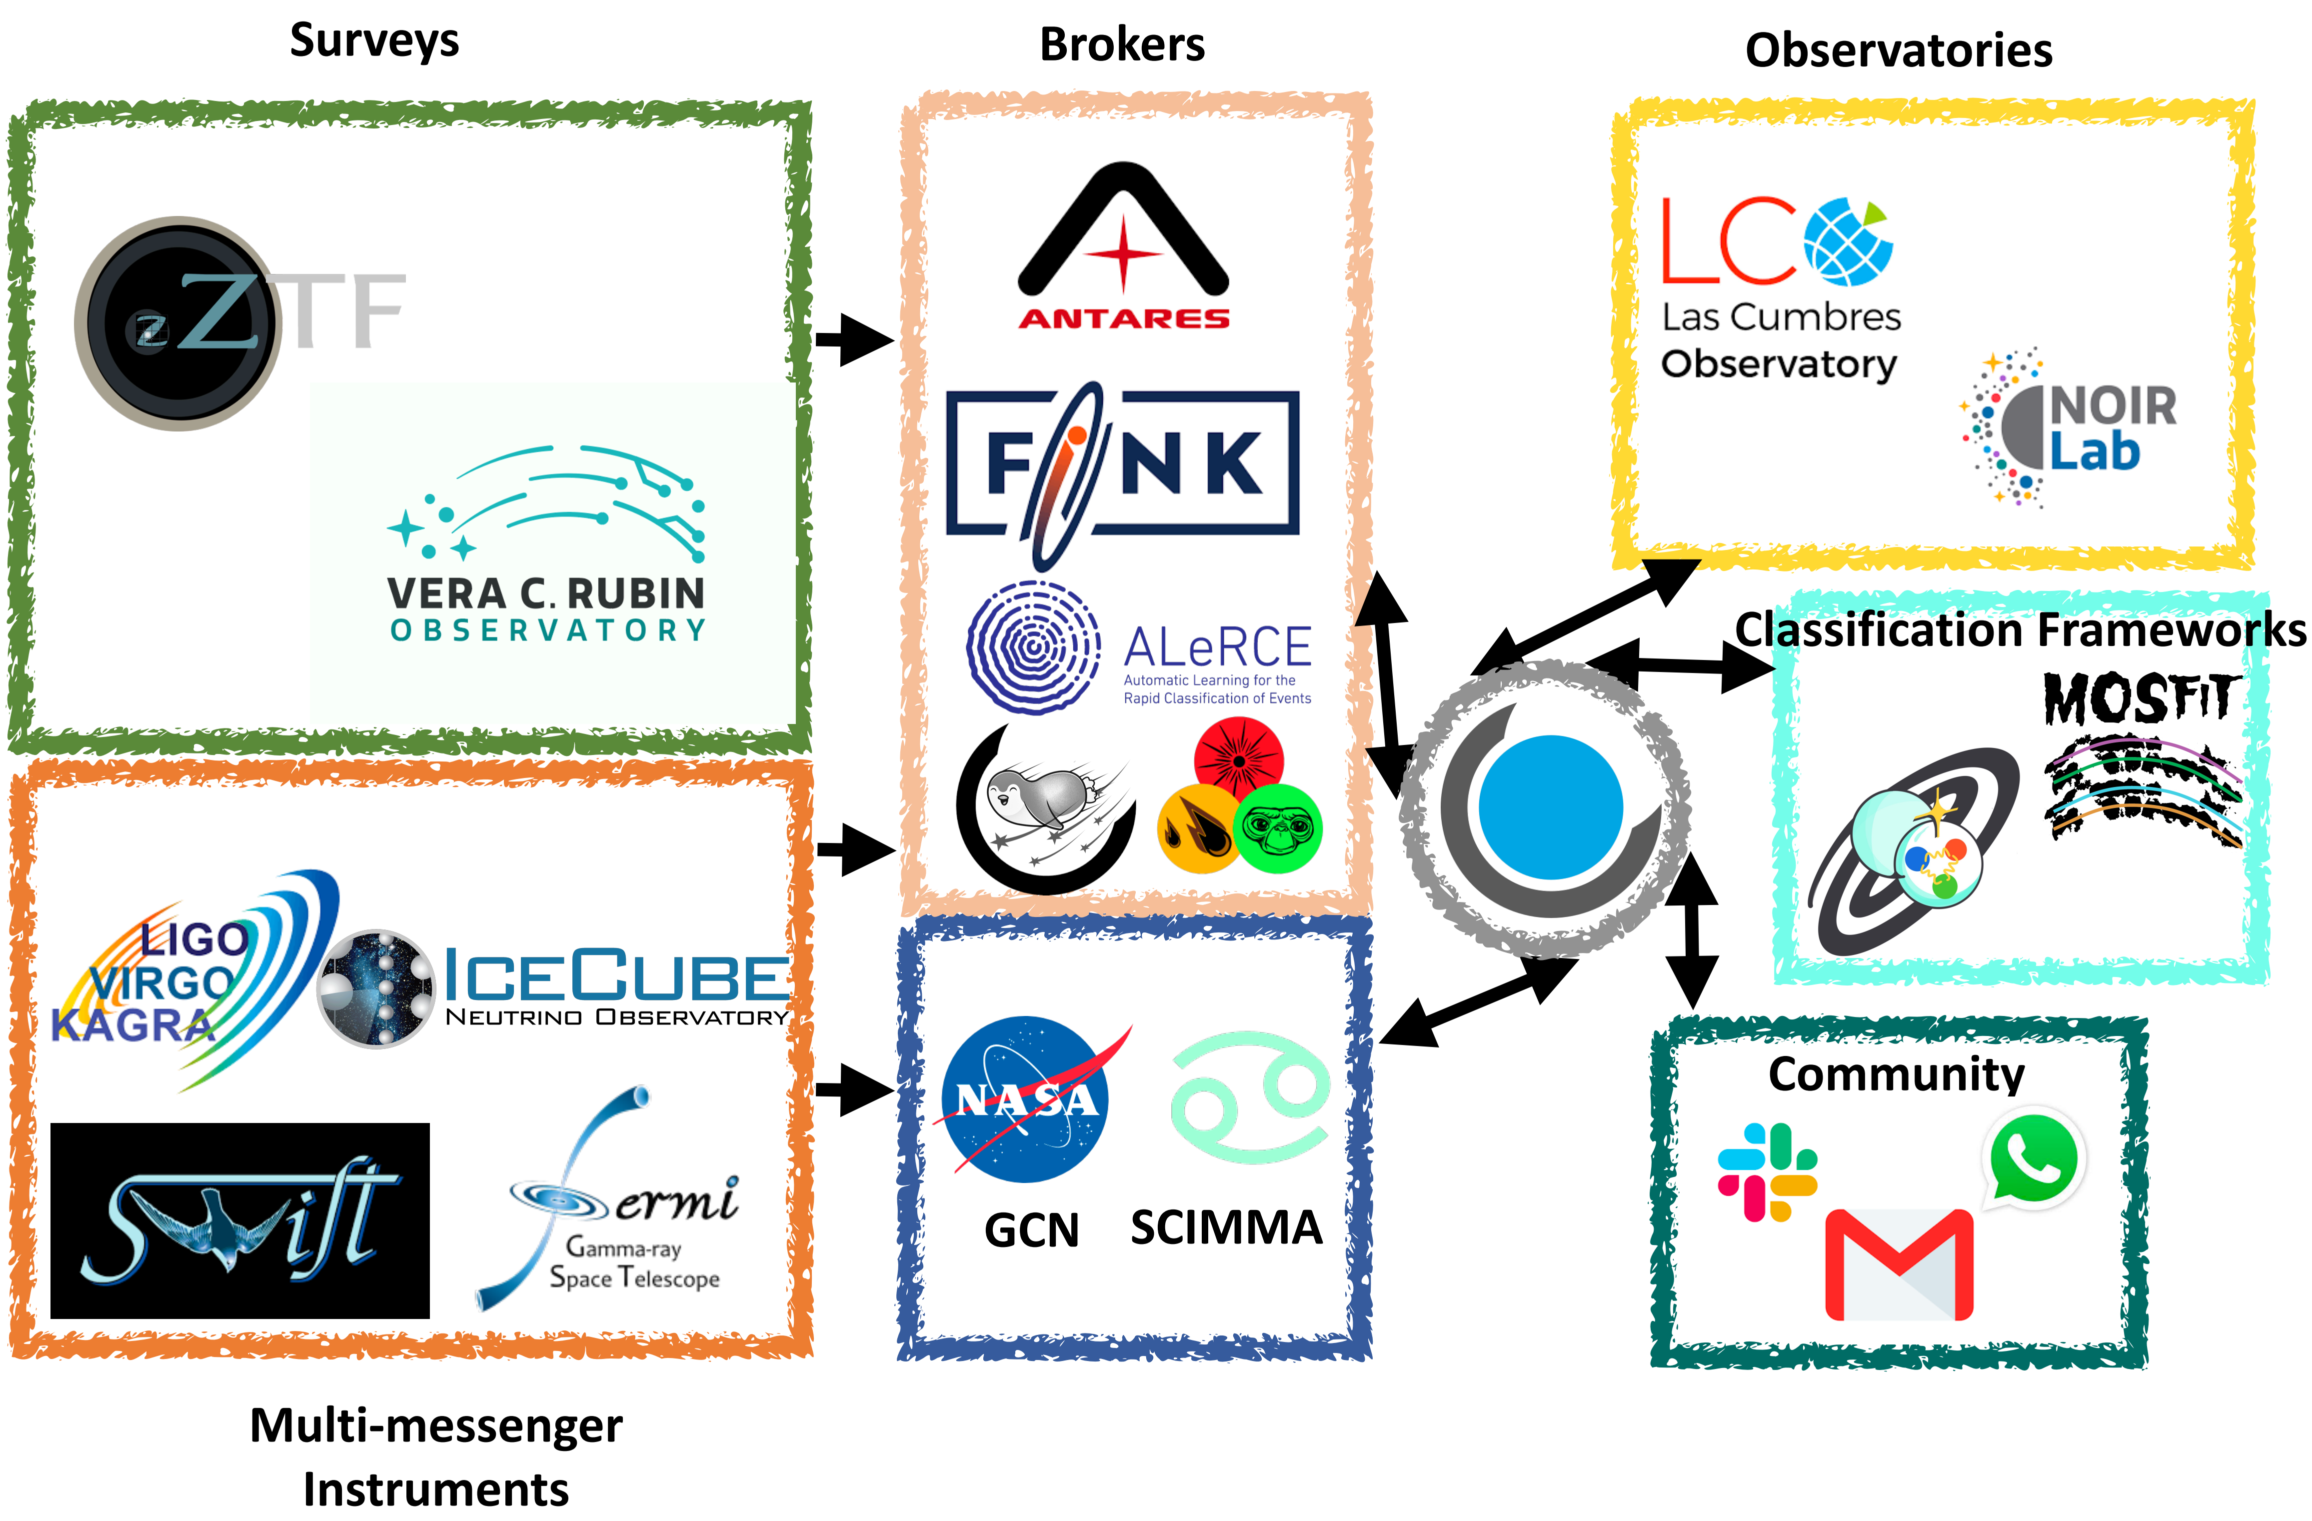

SkyPortal: Overview

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

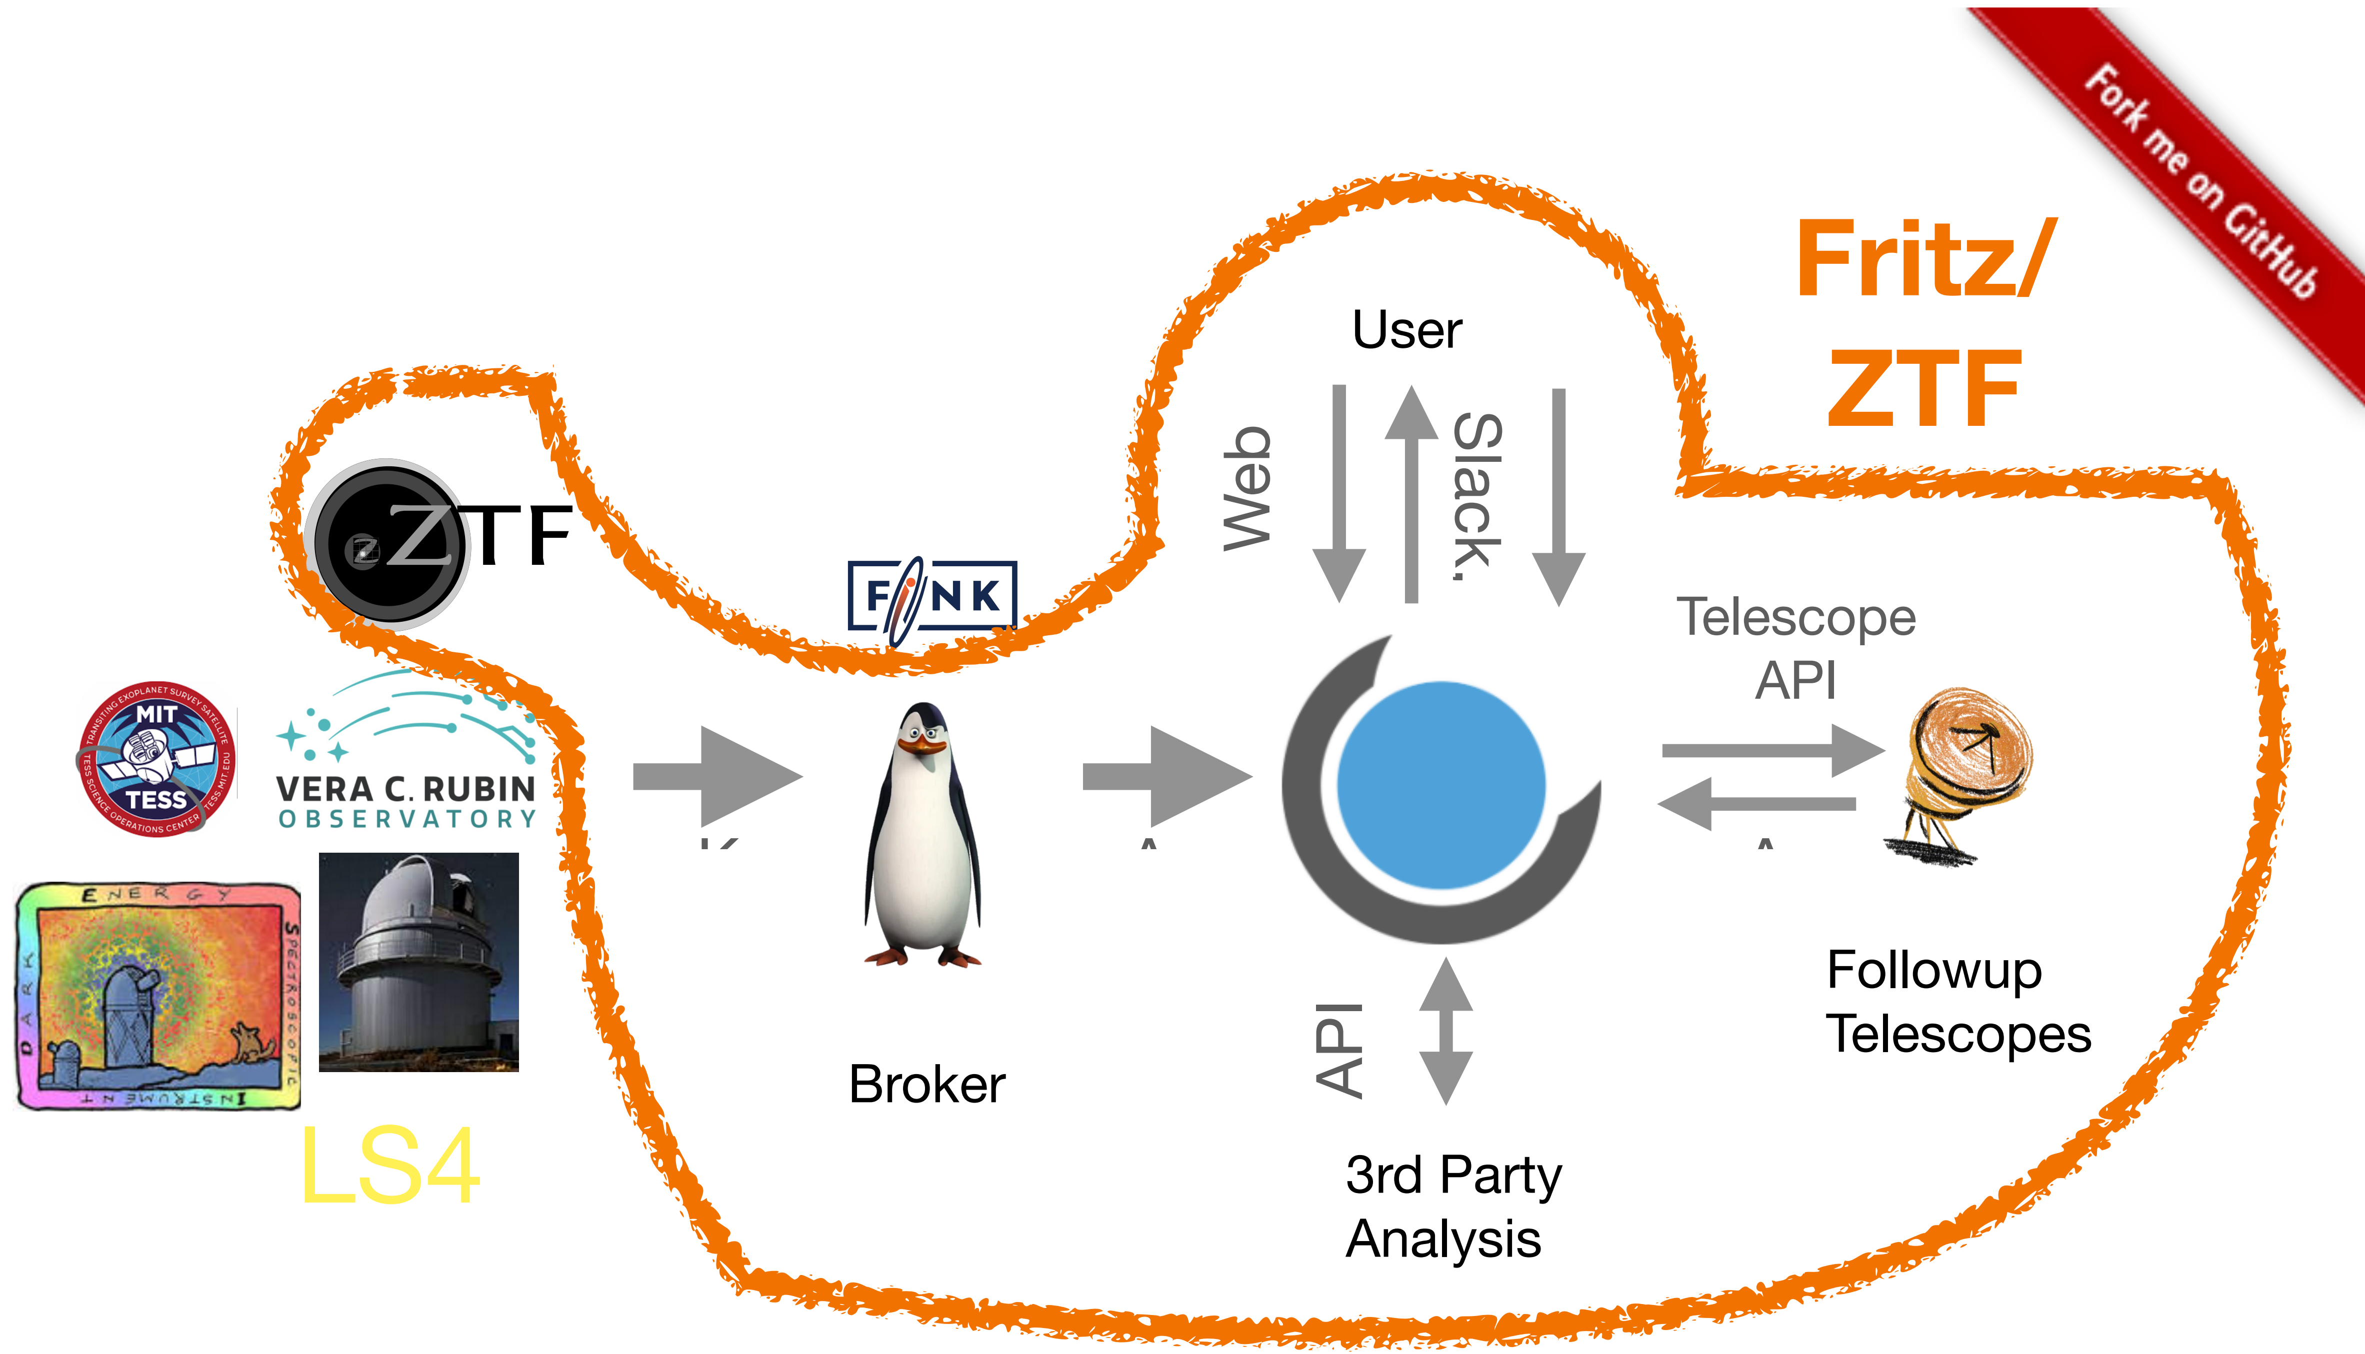

SkyPortal: O4 workflow

Ingestion by `gcn-kafka`

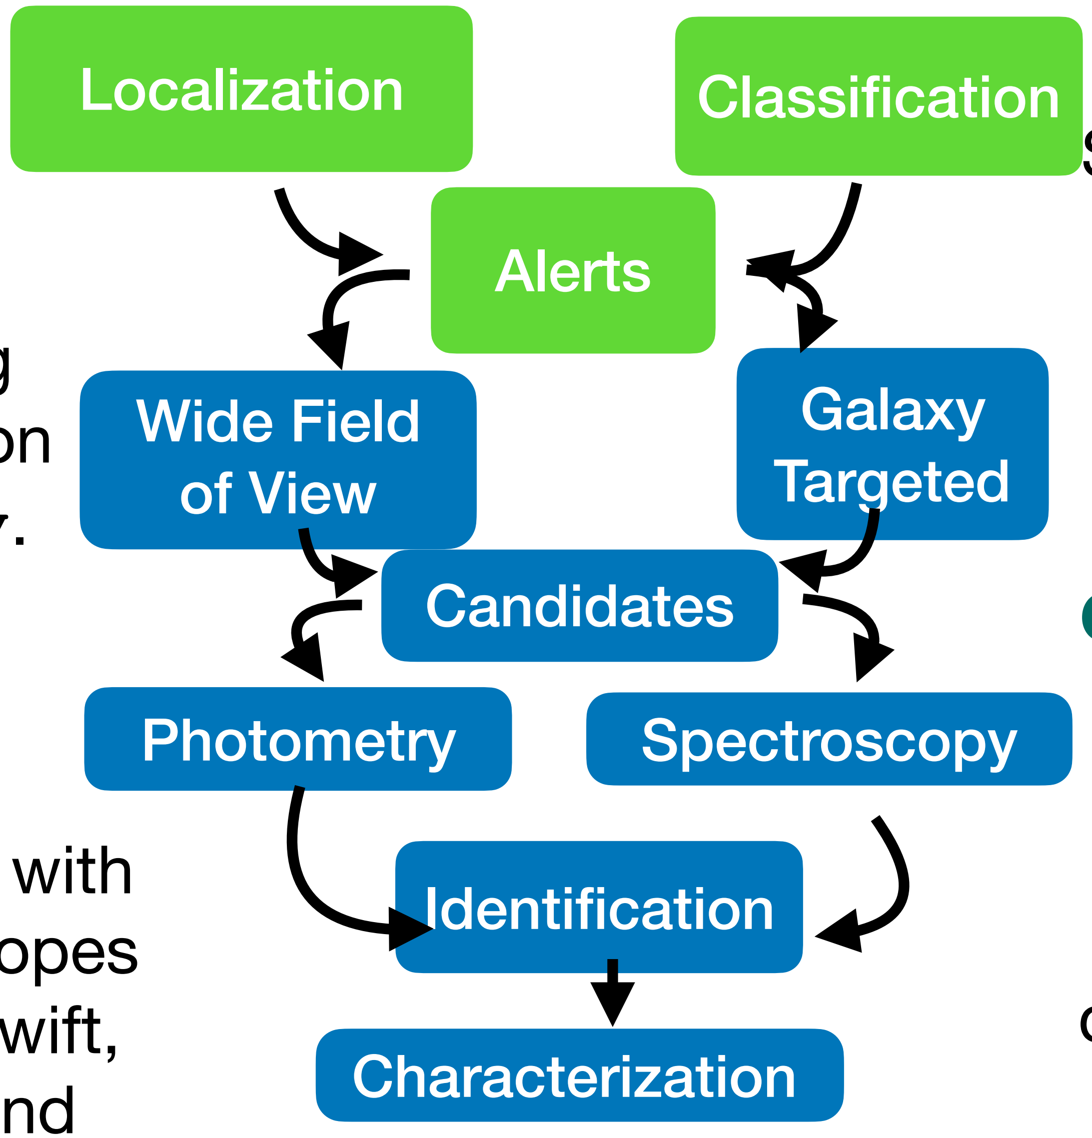

Network-level **scheduling** using `gwemopt`. Validation using `simsurvey`.

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

SkyPortal: In Production

5.9M
Source Views

3.3M
Sources

489M
Photometry Points

5.1k
GCN Events

198k
Comments

366
Users

3.3M
Annotations

9.4M
Candidates

14.7M
Thumbnails

The Bright Transient Survey

10^6 alerts / night → ~50 candidate BTS sources / night → ~5-10 new bright transients / → Classified bright transients

Fritz alert filter

Manual Scanning

BTSBot

10⁶ alerts / night → ~50 candidate BTS sources / night → ~5-10 new bright transients / → Classified bright transients

sgscore
Tachibana & Miller 18
braai
Duev+19

BTSbot
Rehemtulla+23

SkyPortal
Coughlin+23

pySEDM
Rigault+19

SNIascore
Fremling+21

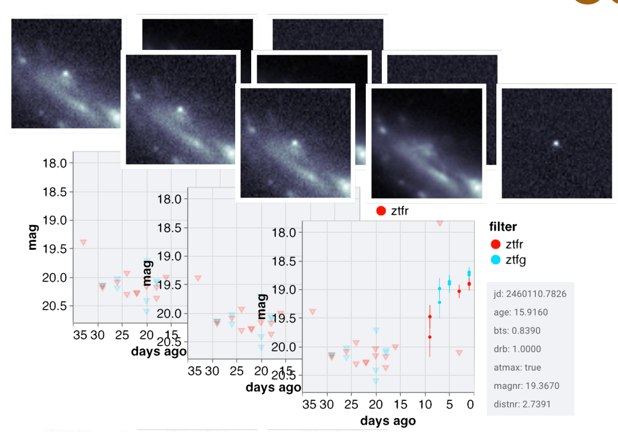

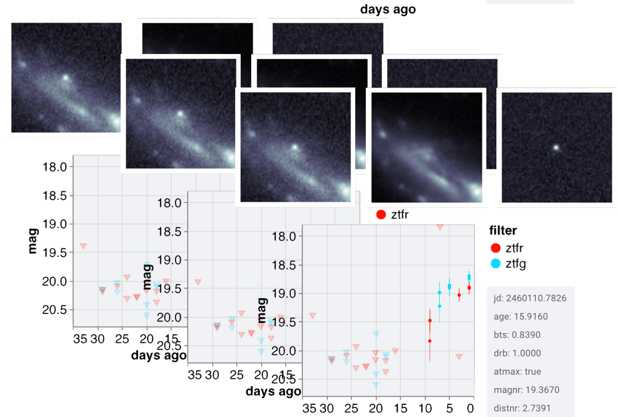

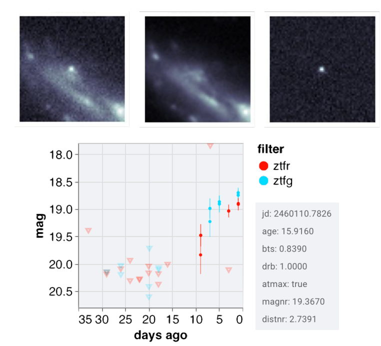

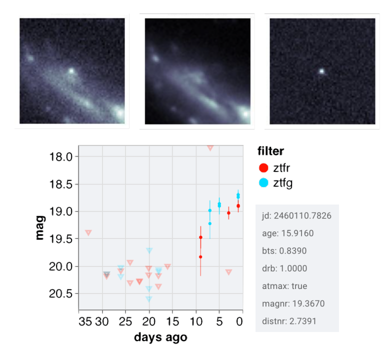

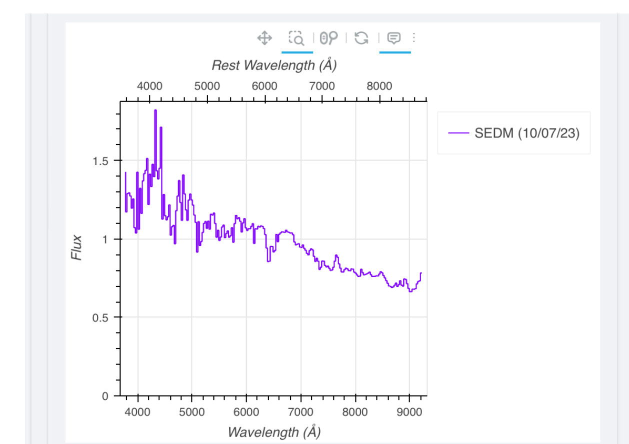

SN 2023tyz

World's First:
Fully automatic from discovery to TNS

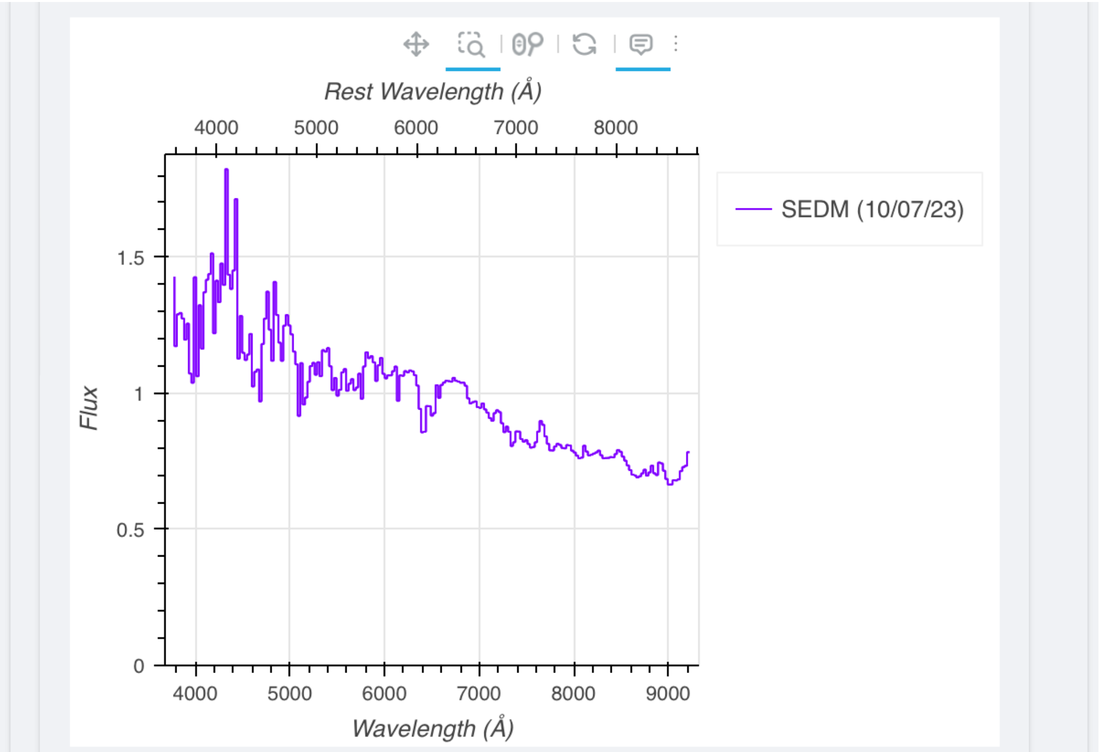

Nabeel Rehemtulla
Northwestern

Theophile Jegou du Laz
Caltech

From gravitational waves to photons

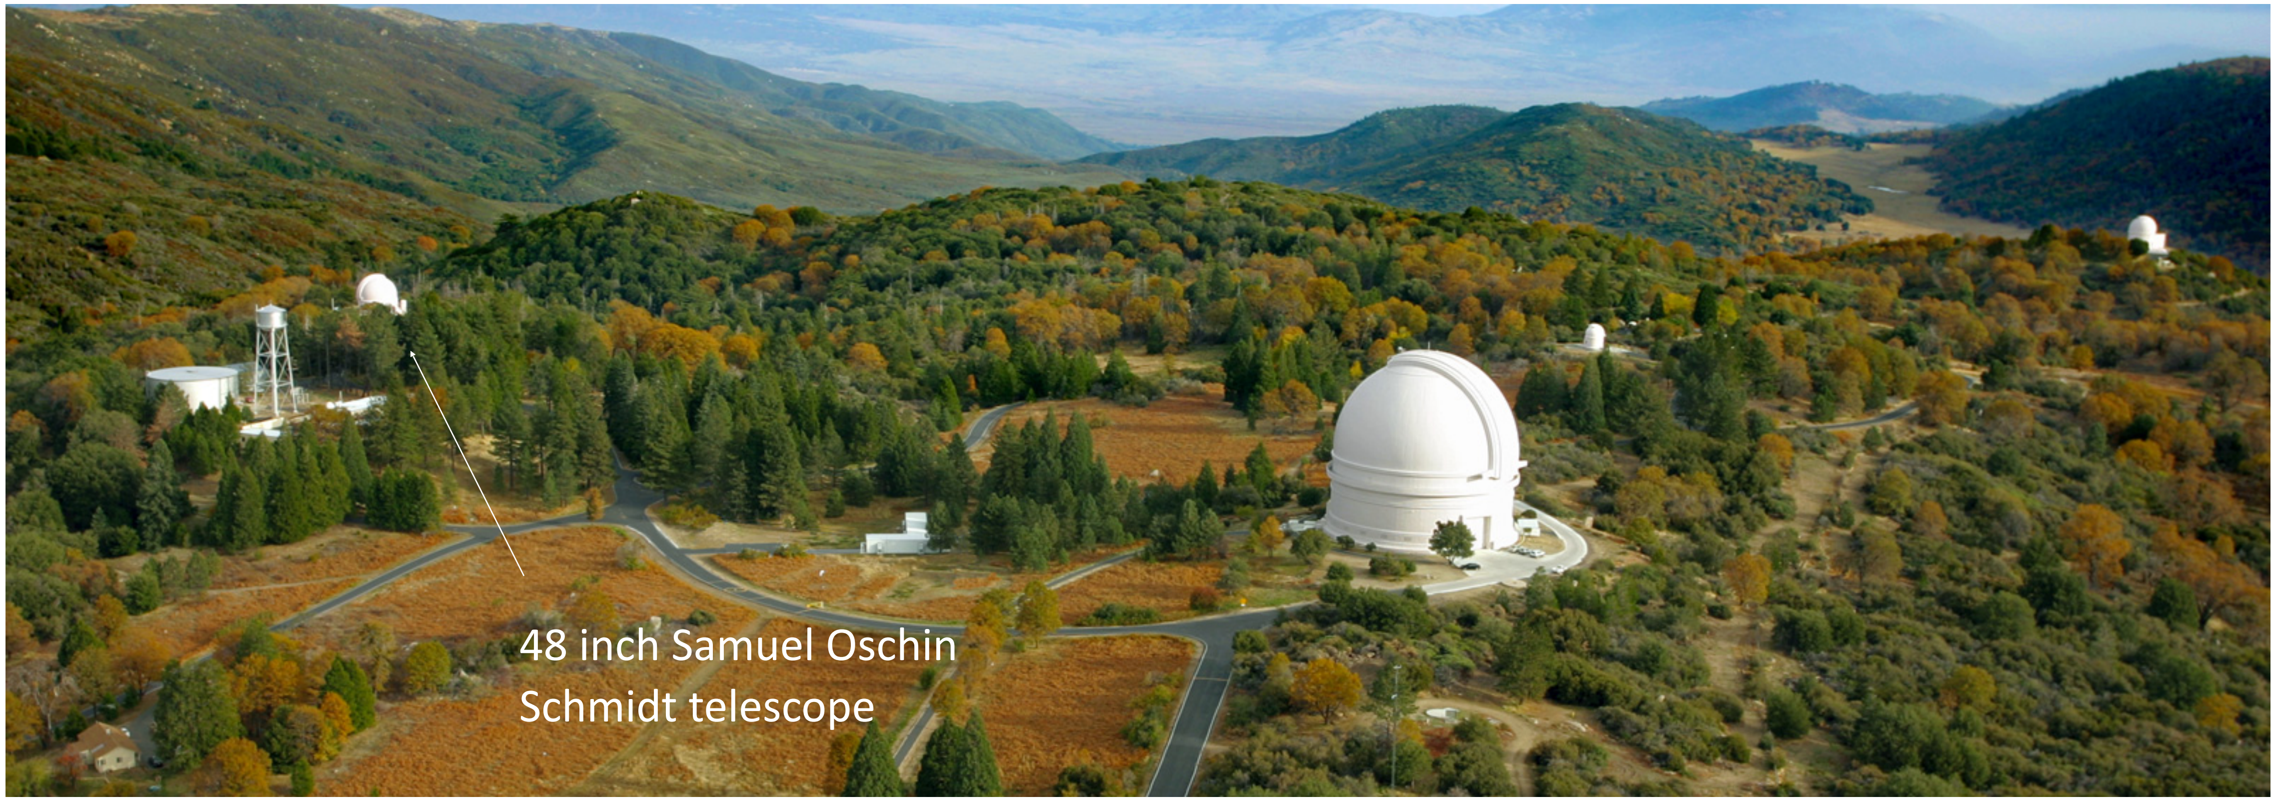

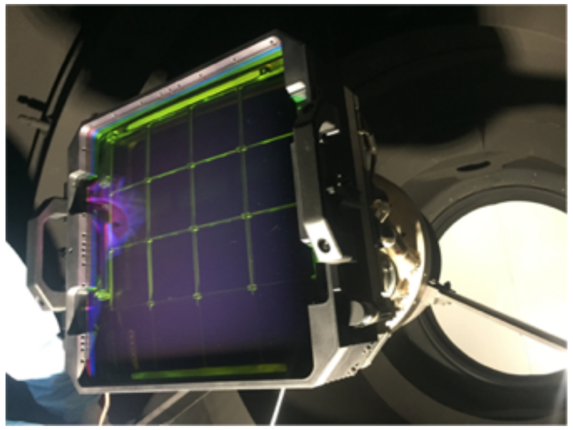

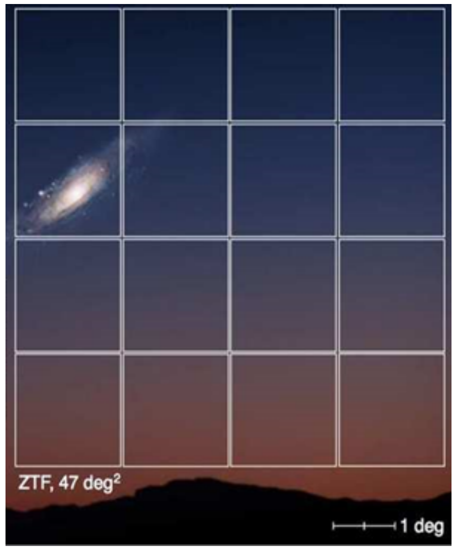

The crucial element: ZTF has a large field of view, and accumulates many images quickly

Optical Data

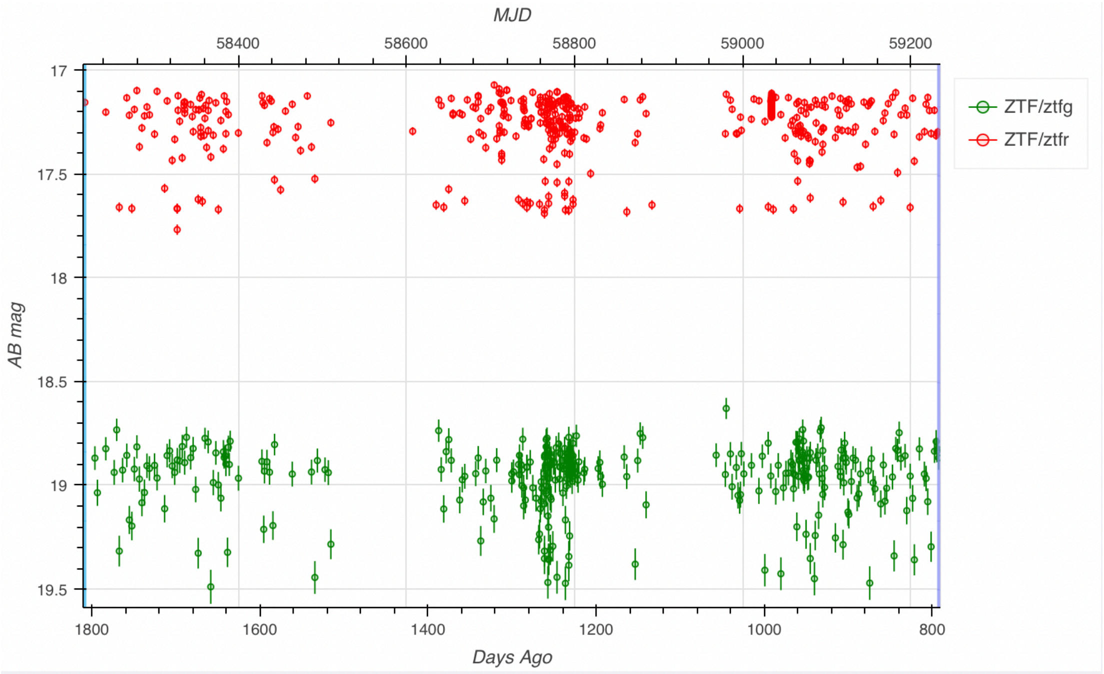

Beta Lyrae

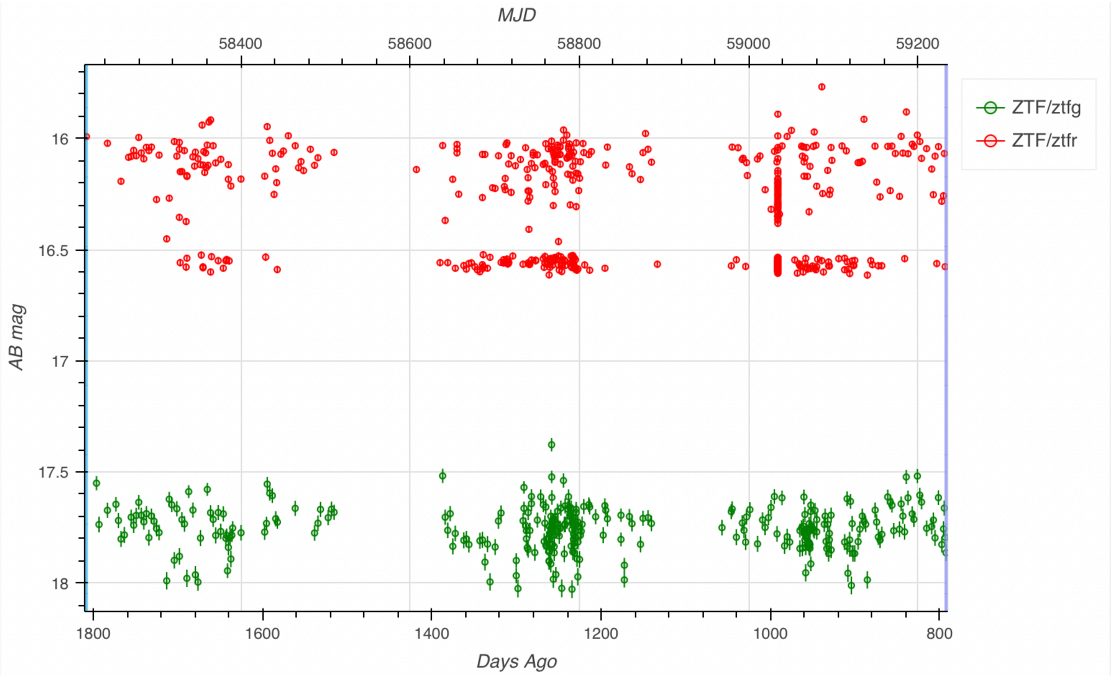

Blend

And with the help of some other kinds of hardware…

Optical Data

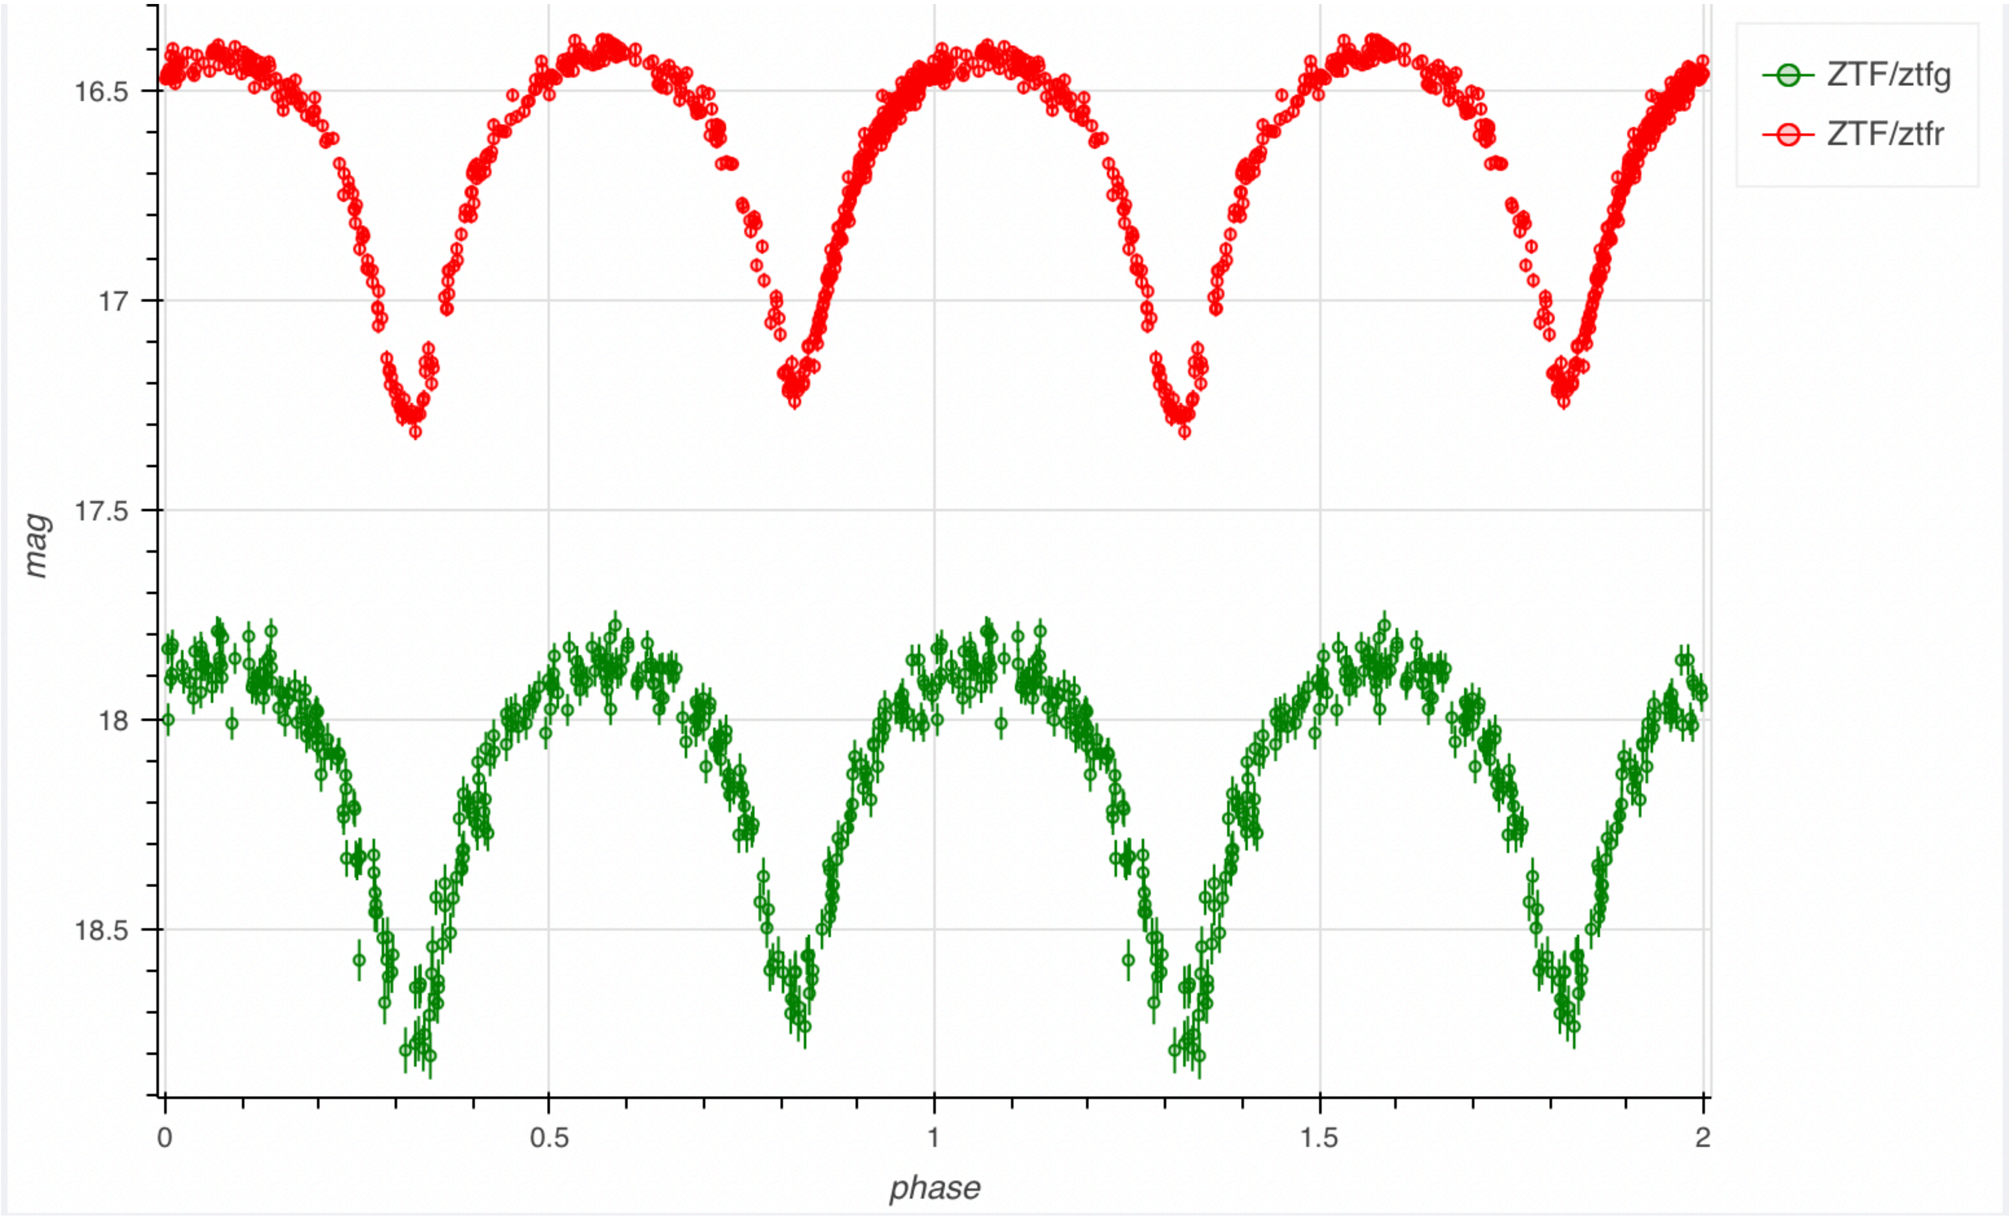

W Ursae Maj
(contact binary)
P = 0.35017 d

ZTF Source Classification Project

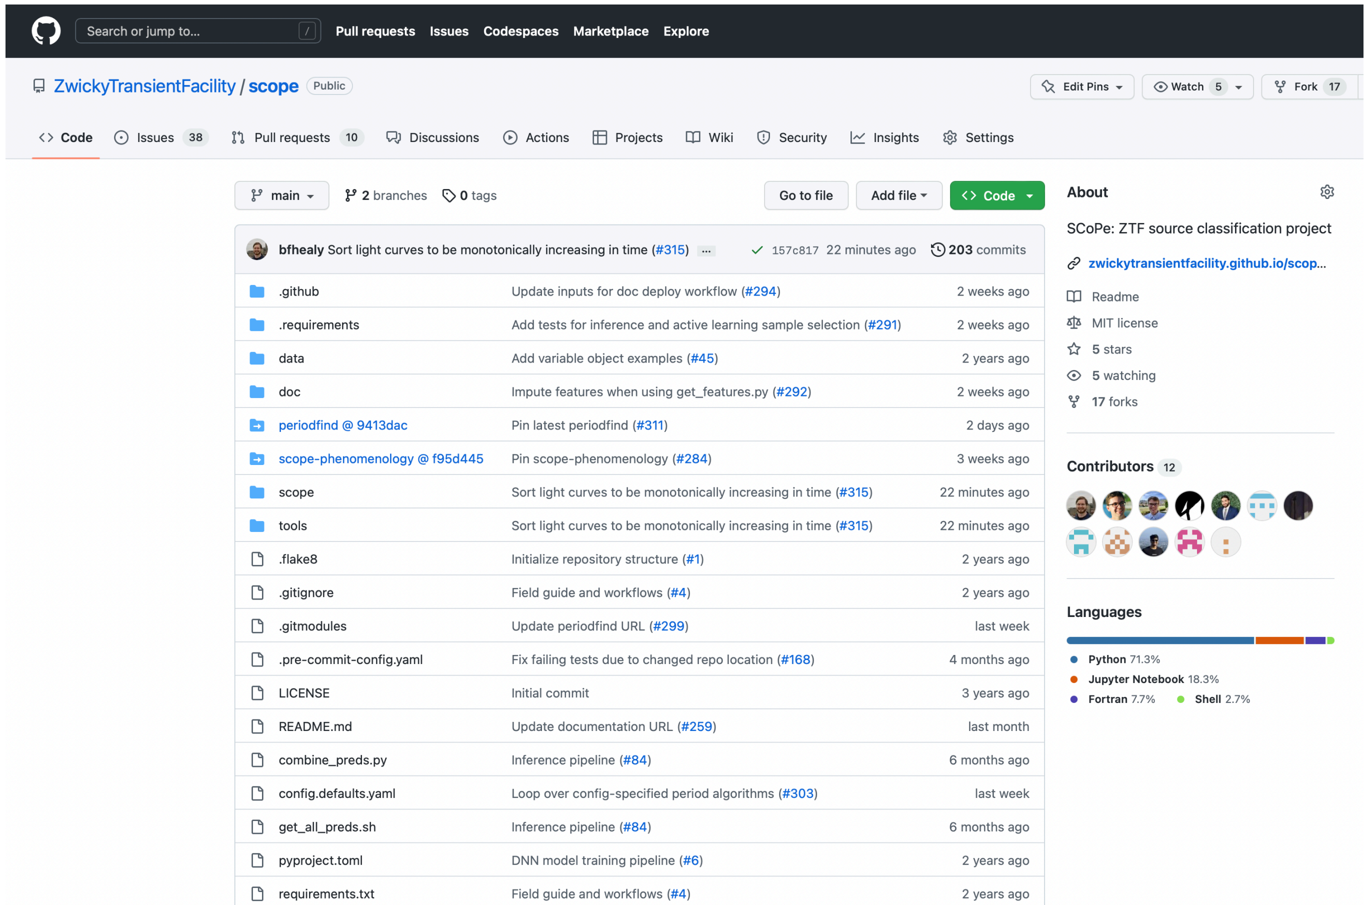

- Open-source
- Python-based
- CI/CD pipeline
- Regularly updated docs

- **Supervised, active learning:** training set built up over time (w/human input)
- **Two taxonomies:** ontological (intrinsic), phenomenological (light curve shape)
  - Provides useful information for anomalous sources
  - Avoids complications of overlapping classes

(van Roestel et al. 2021, Coughlin et al. 2021, Healy et al. 2023)

ZTFReST for fast transient discovery

growth-astro/ztfrest

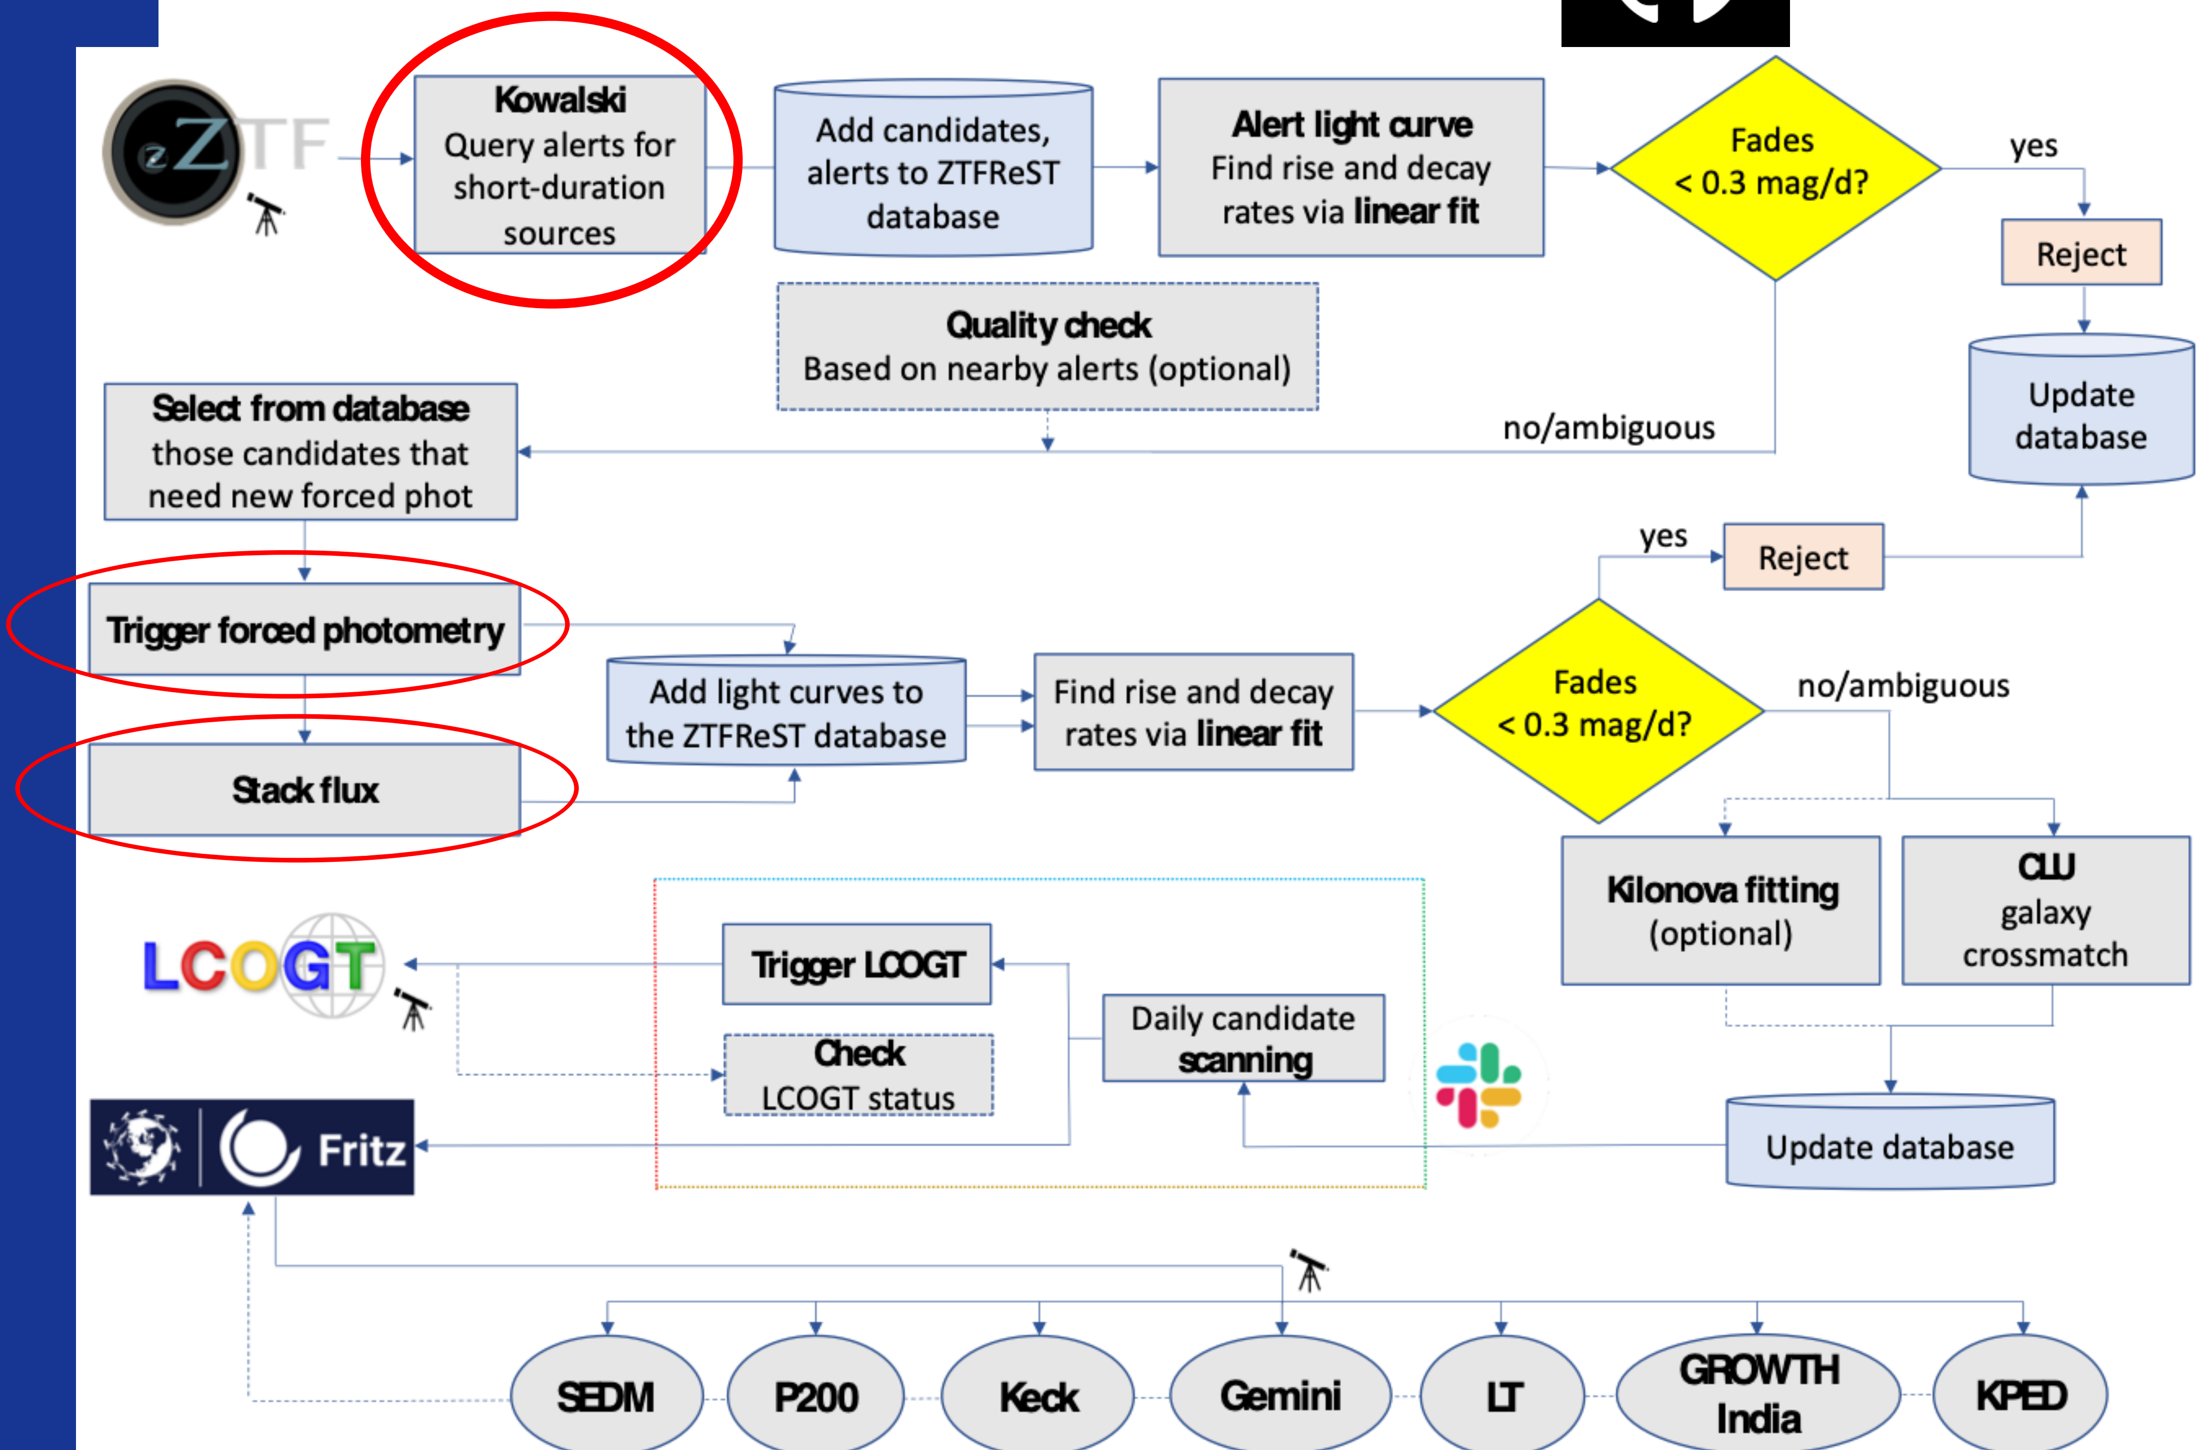

[Andreoni, Coughlin+2021, ApJ]

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 … 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 astron…

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… 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]

Constraints on Kilonova and neutron star merger rates

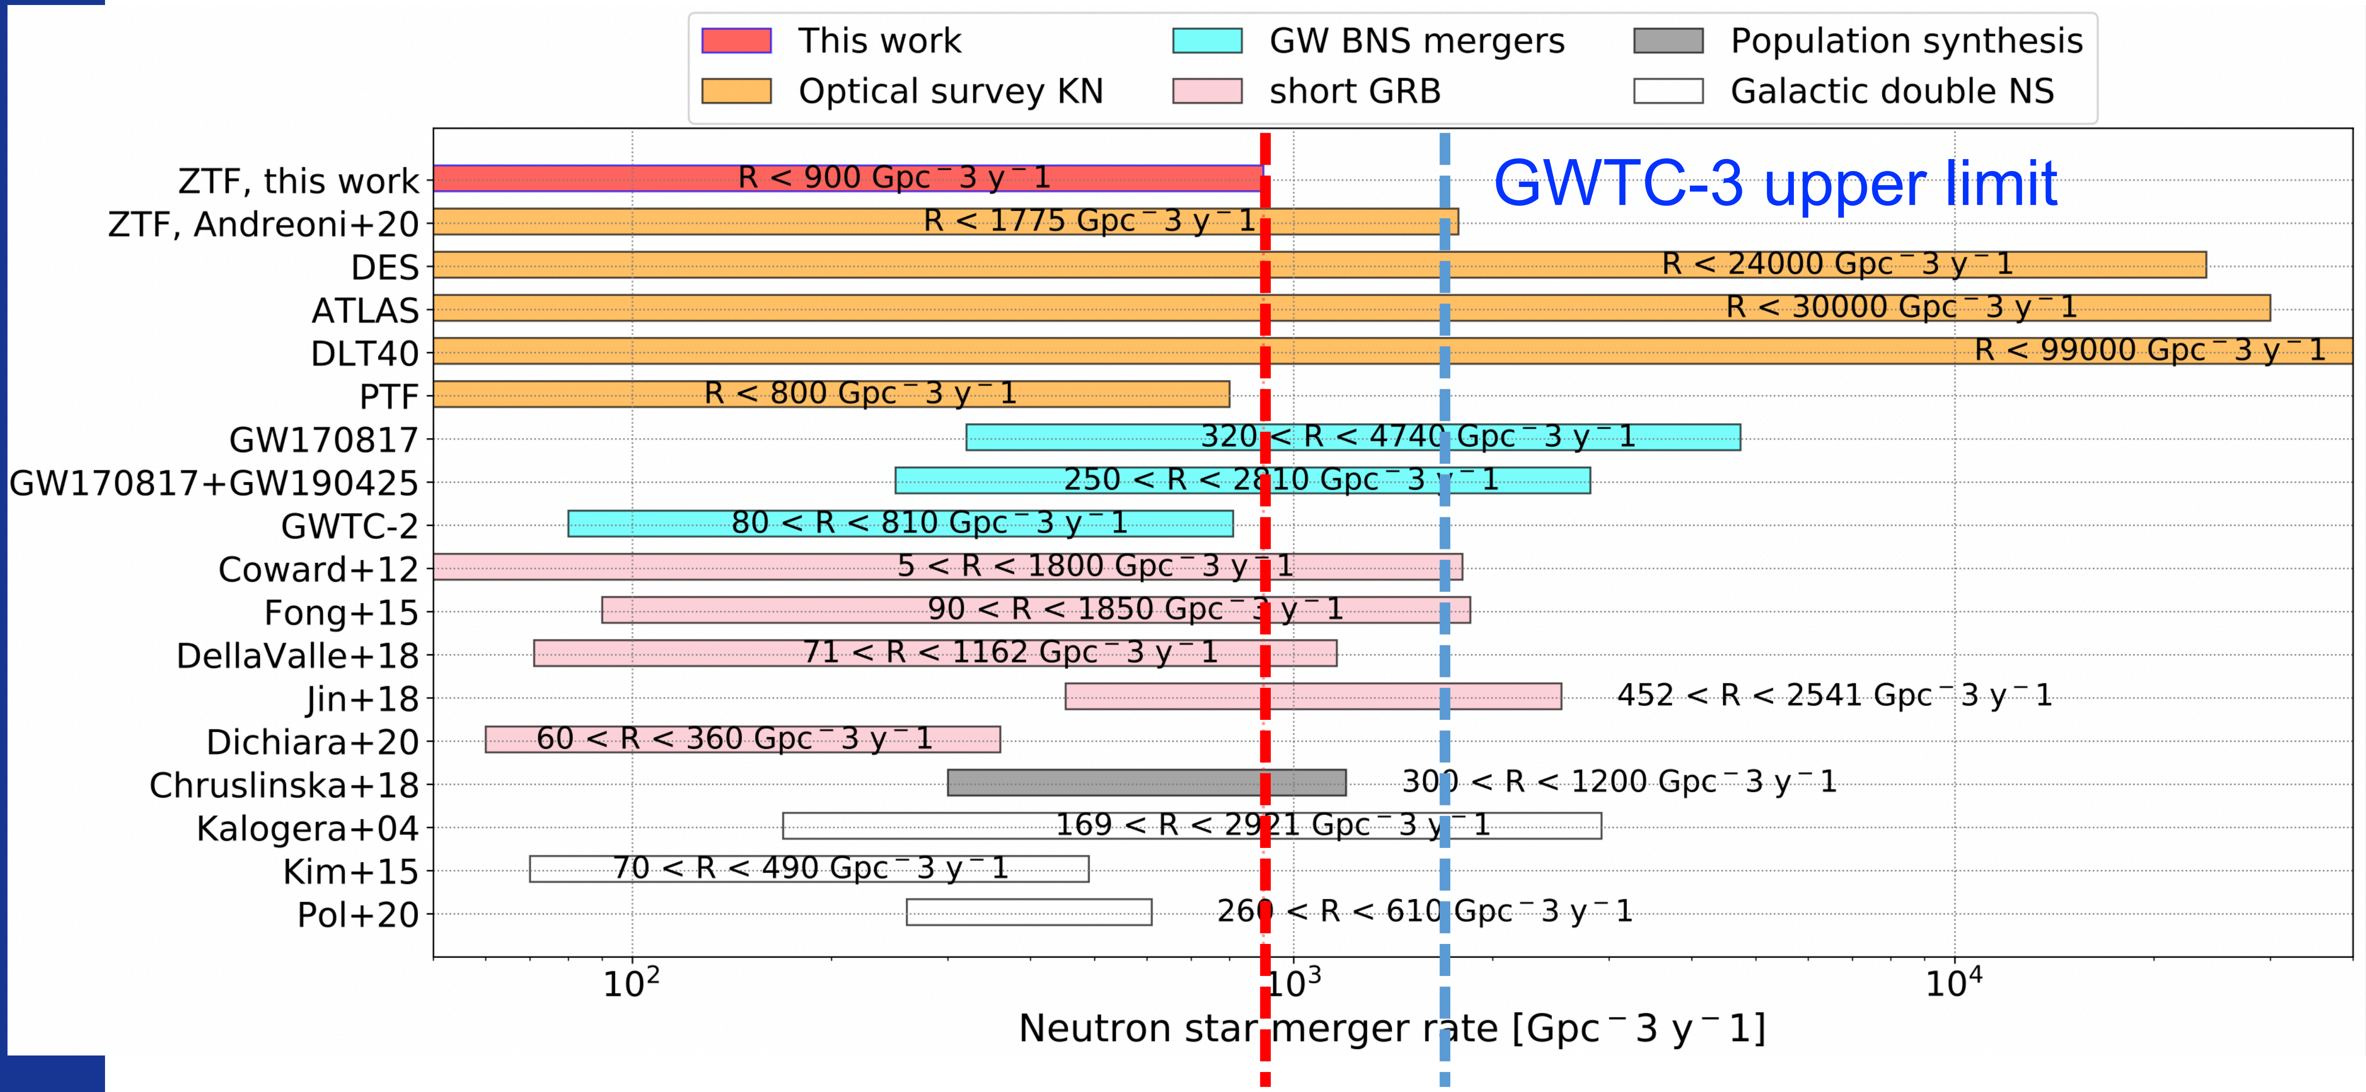

Andreoni & Coughlin et al. (2021), ApJ, 918, 2, 63
Model grid in Andreoni et al. (2020d), ApJ, 904, 2, 155

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

[Andreoni, Coughlin+2021, ApJ]

NS Equation of State

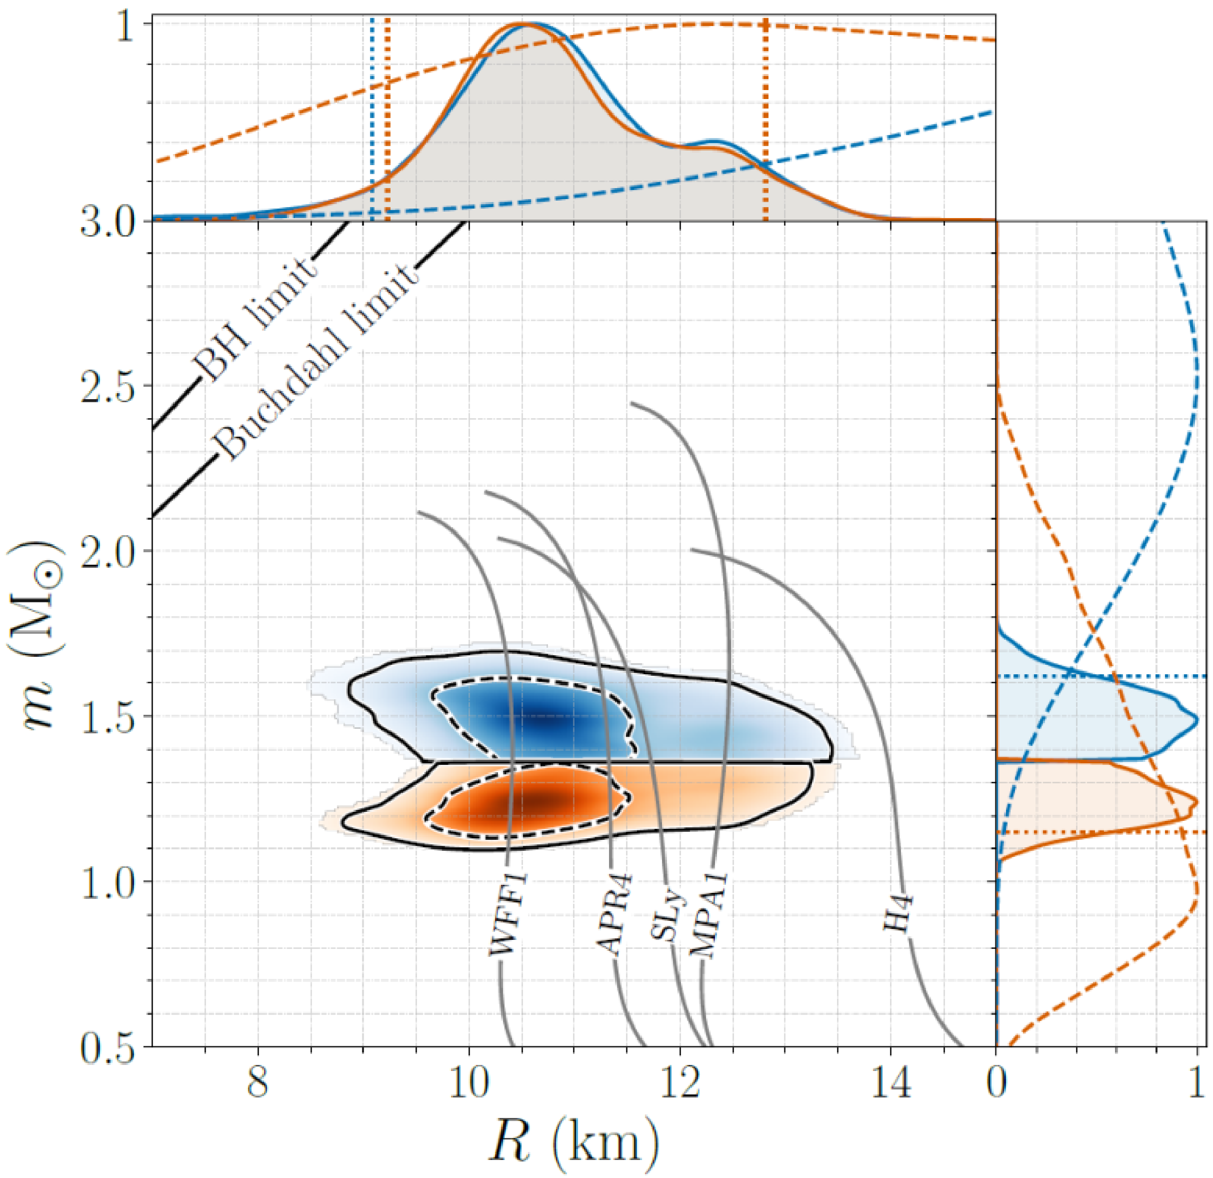

Phys. Rev. Lett. **121**, 161101 (2018)

Described by the Equation of State:

- Pressure-density or equivalently mass-radius.
- Numerous models proposed.

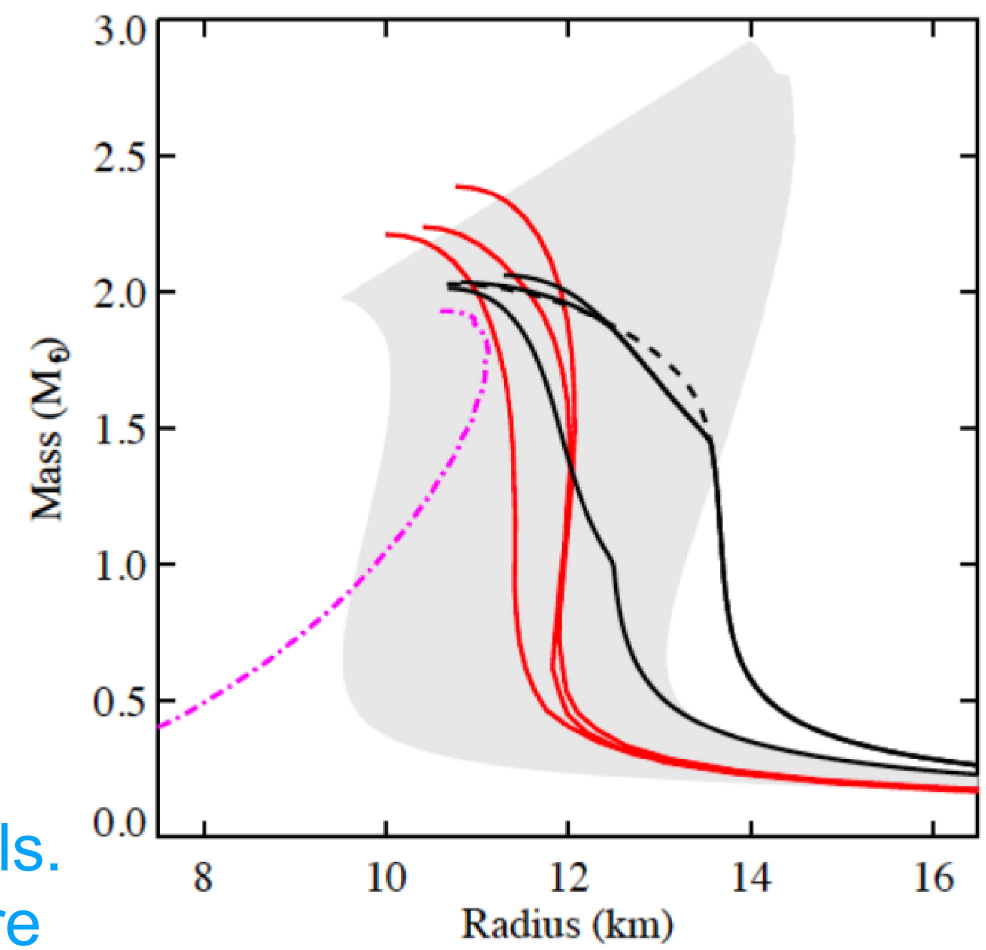

Watts et al, Rev. Mod. Phys. 88. 021001 (2016)

- Stiffness of NS determines how much it deforms under gravity (tidal effects).
- Estimate NS deformations from GW signals.
- Some EoS models are ruled out, others are consistent with observations.

MMA Equation of State Constraints

Nuclear Theory

NICER - Pulsars

Optical Counterpart

GWs

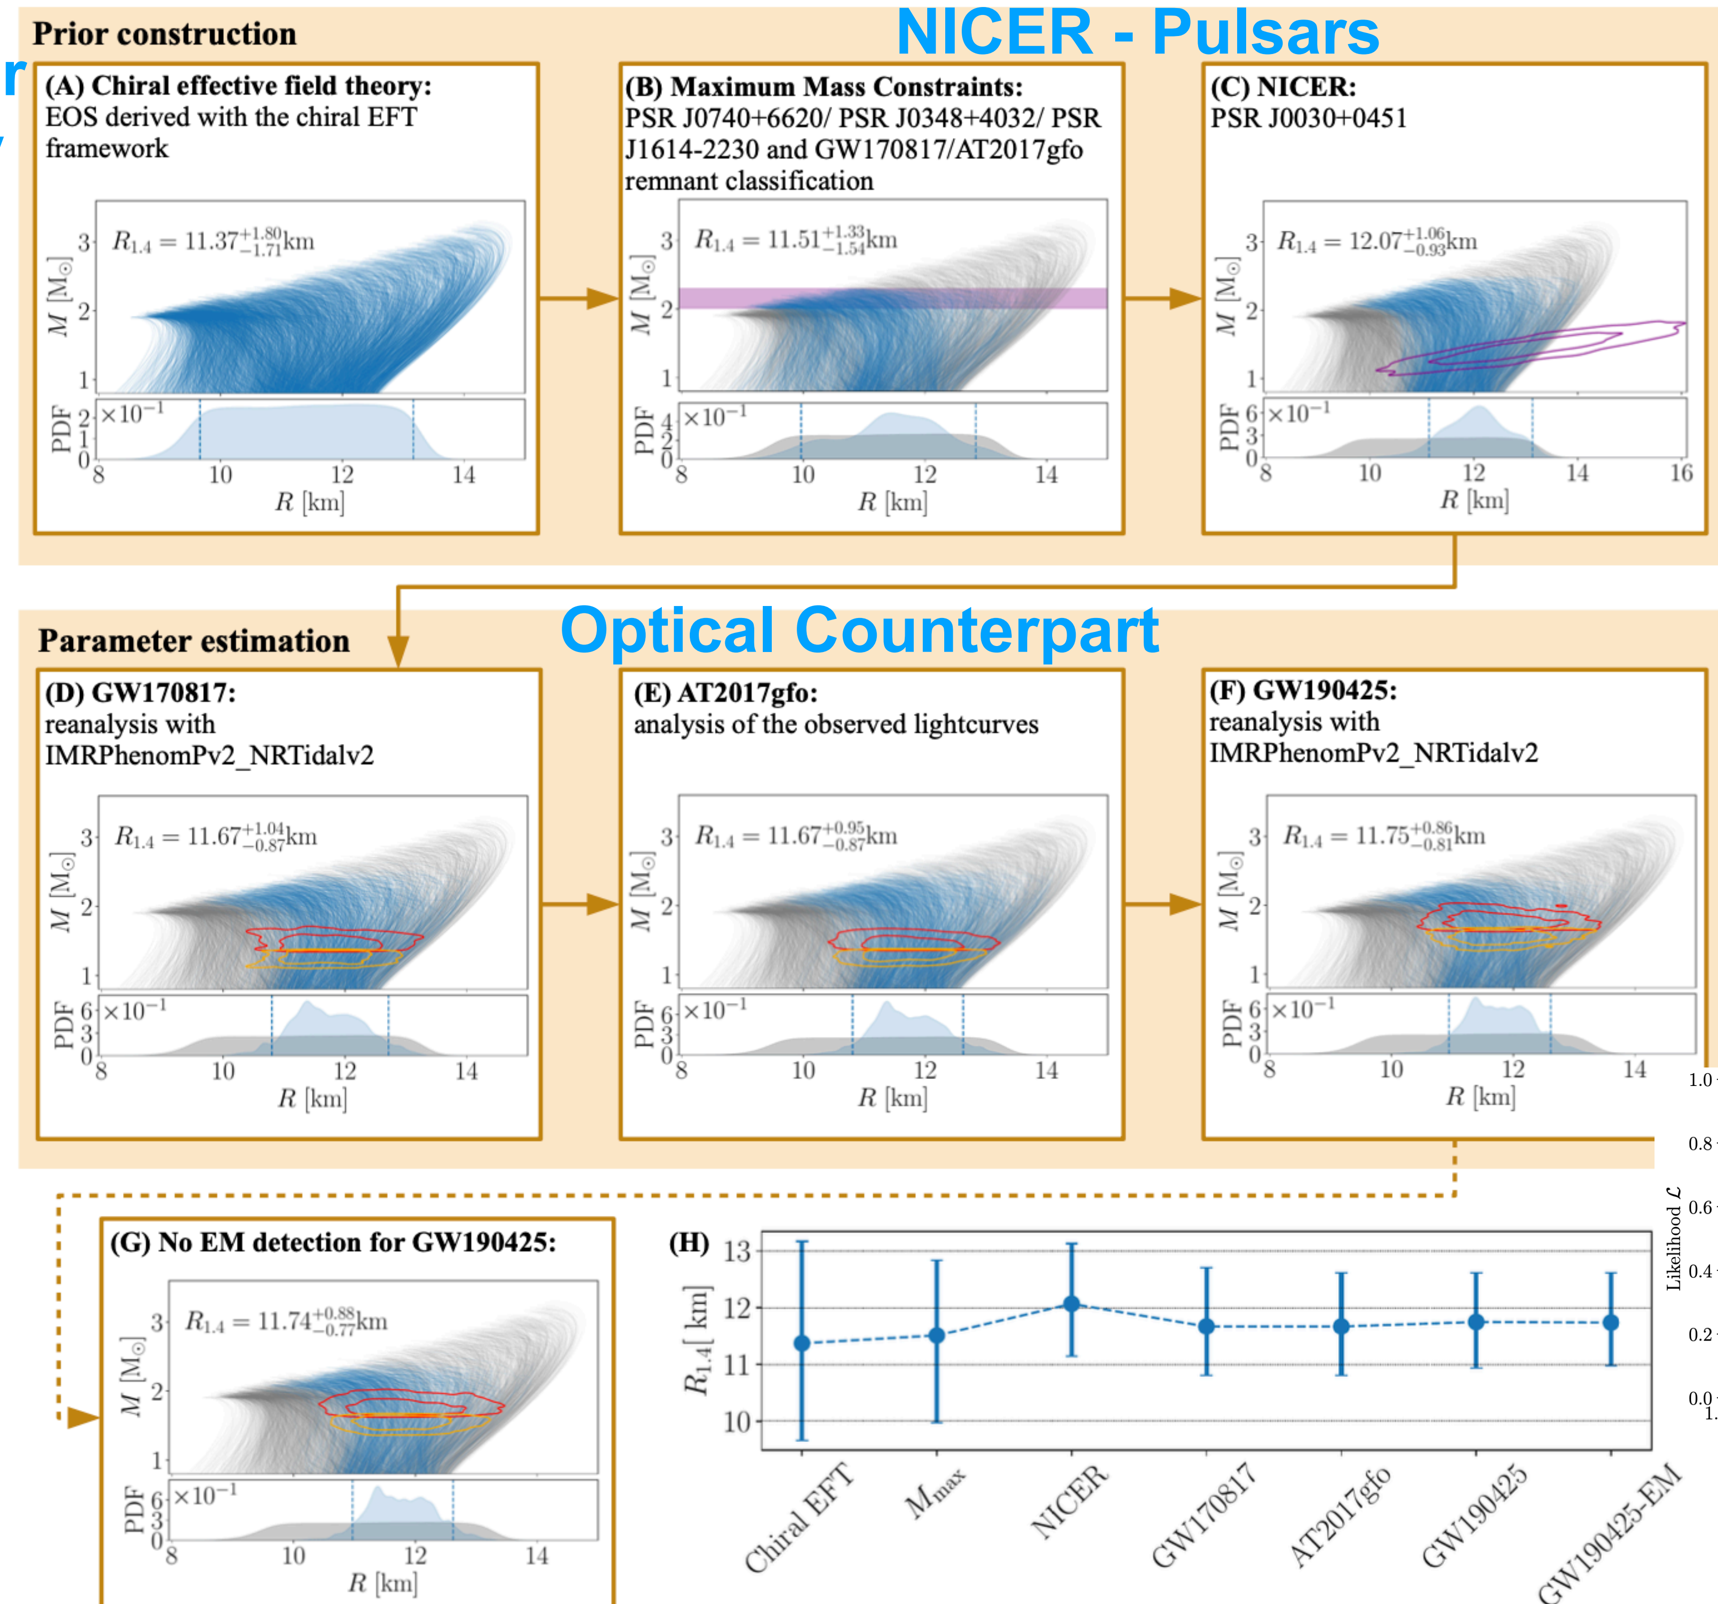

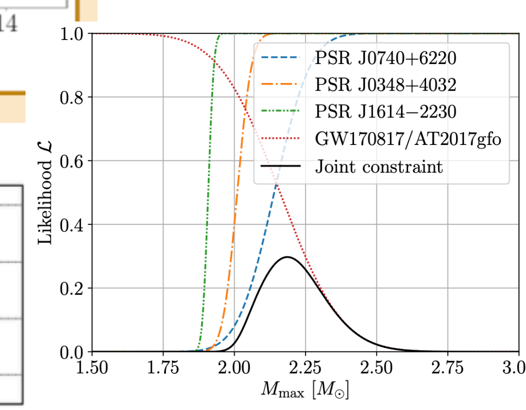

[Dietrich, Coughlin, Pang+2020, Science]

MMA Equation of State Constraints

Good agreement between macroscopic and microscopic collisions

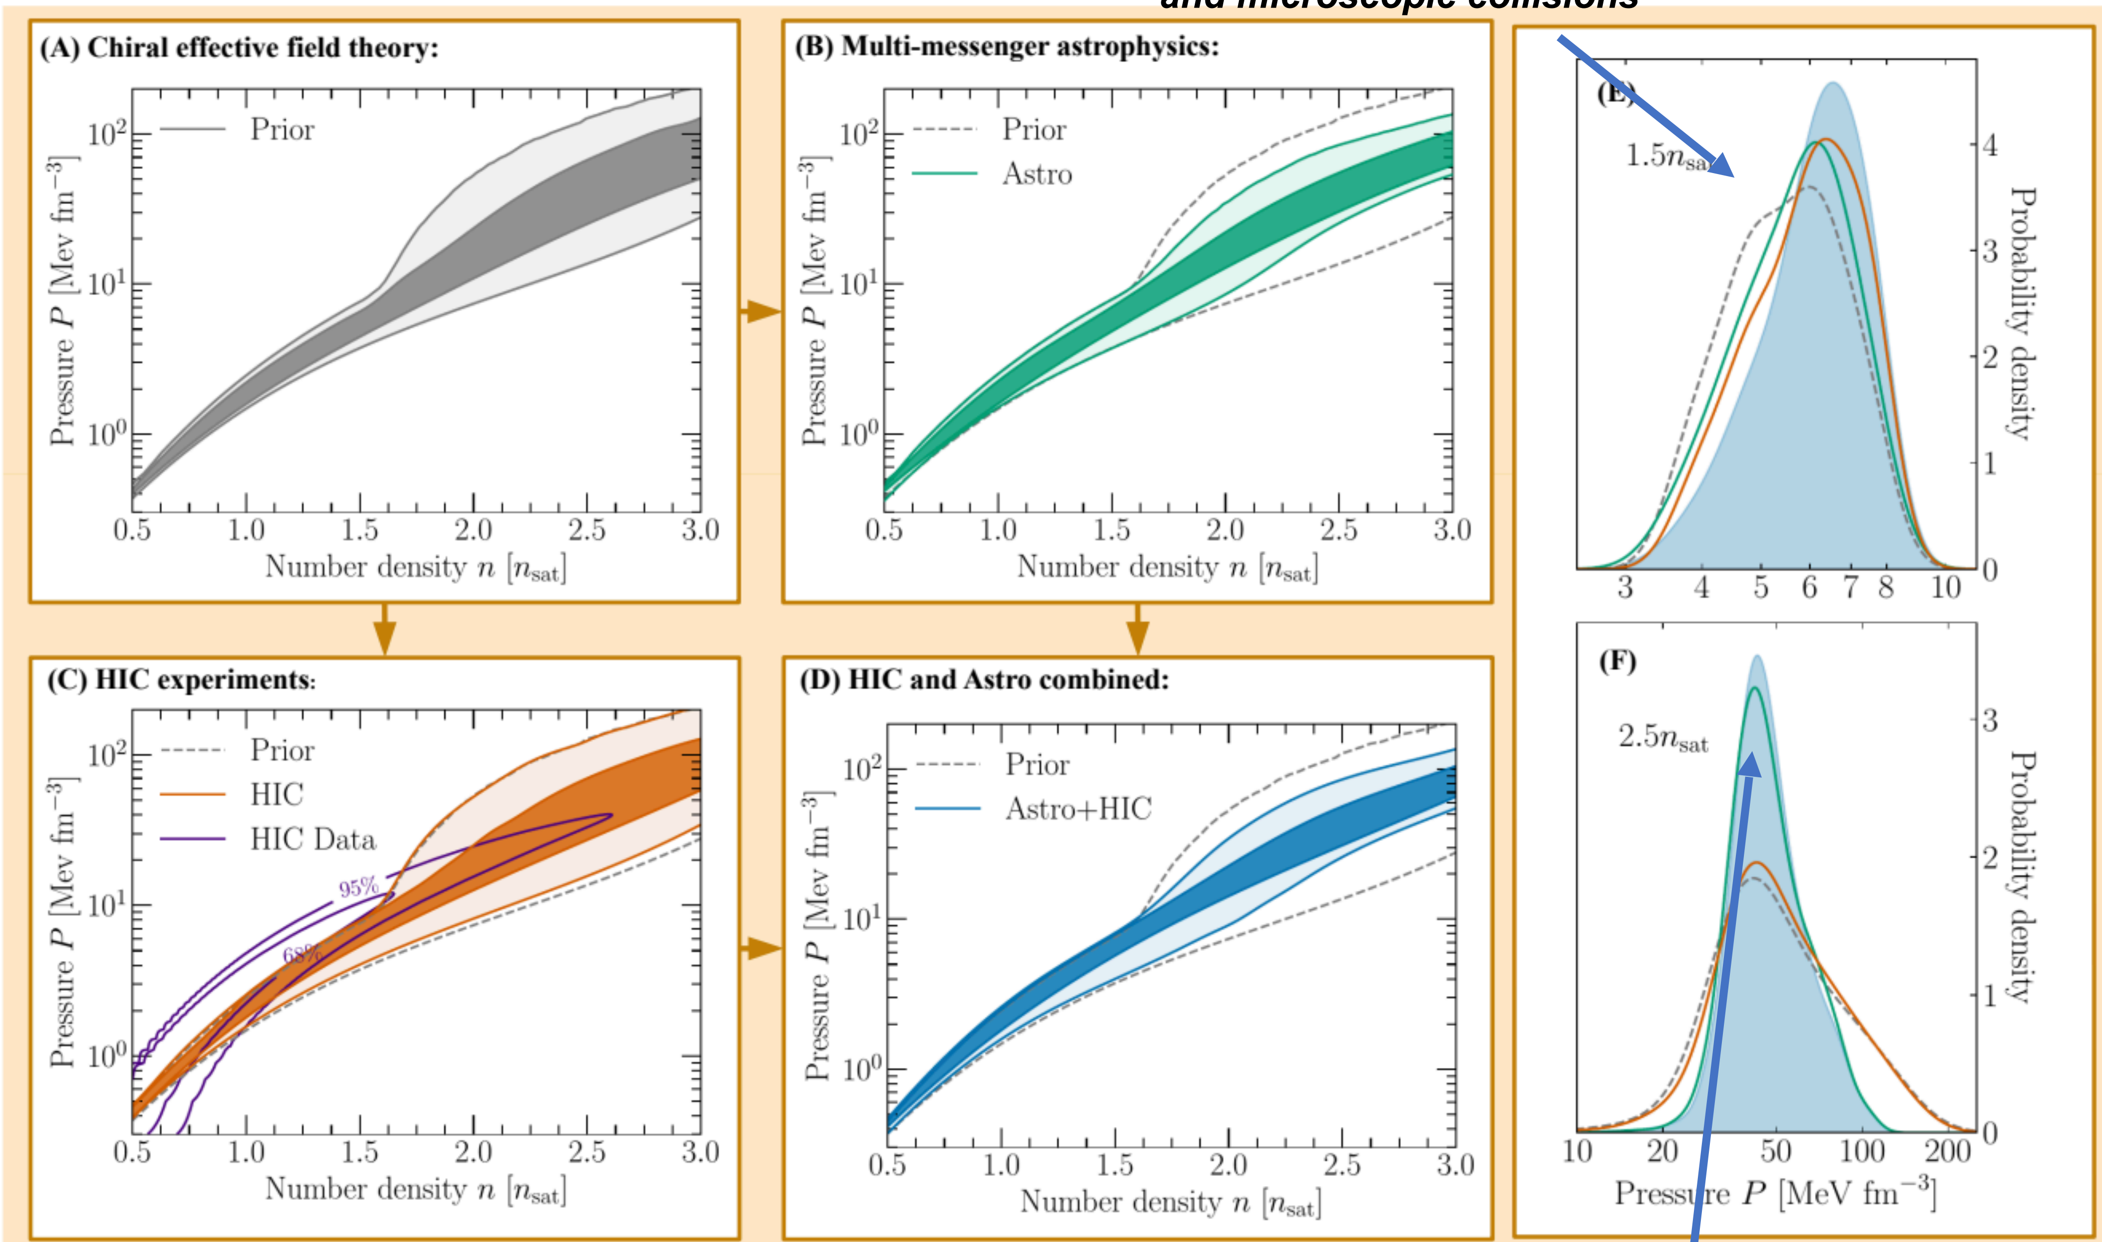

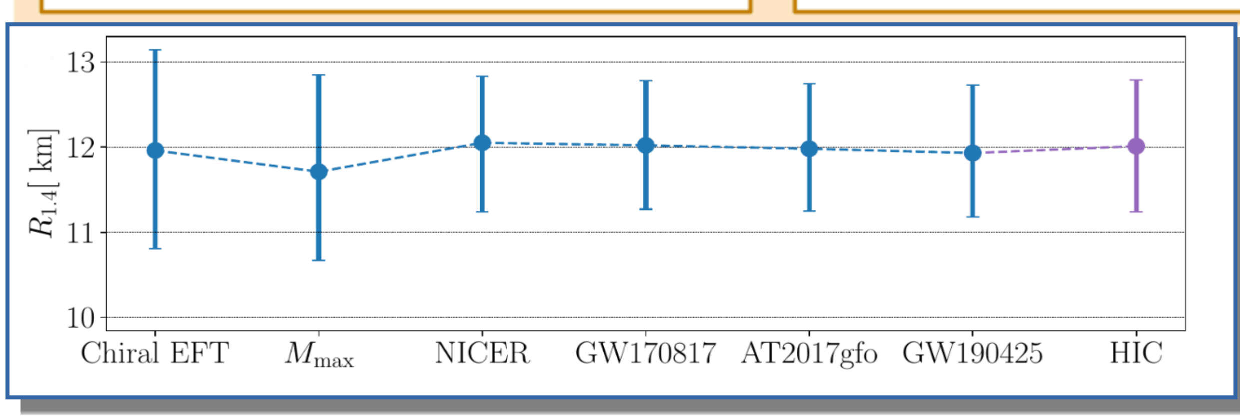

high-density information from astrophysical studies

[Huth+2022, Nature]

A nuclear physics and multi-messenger framework (NMMA)

github.com/nuclear-multimessenger-astronomy

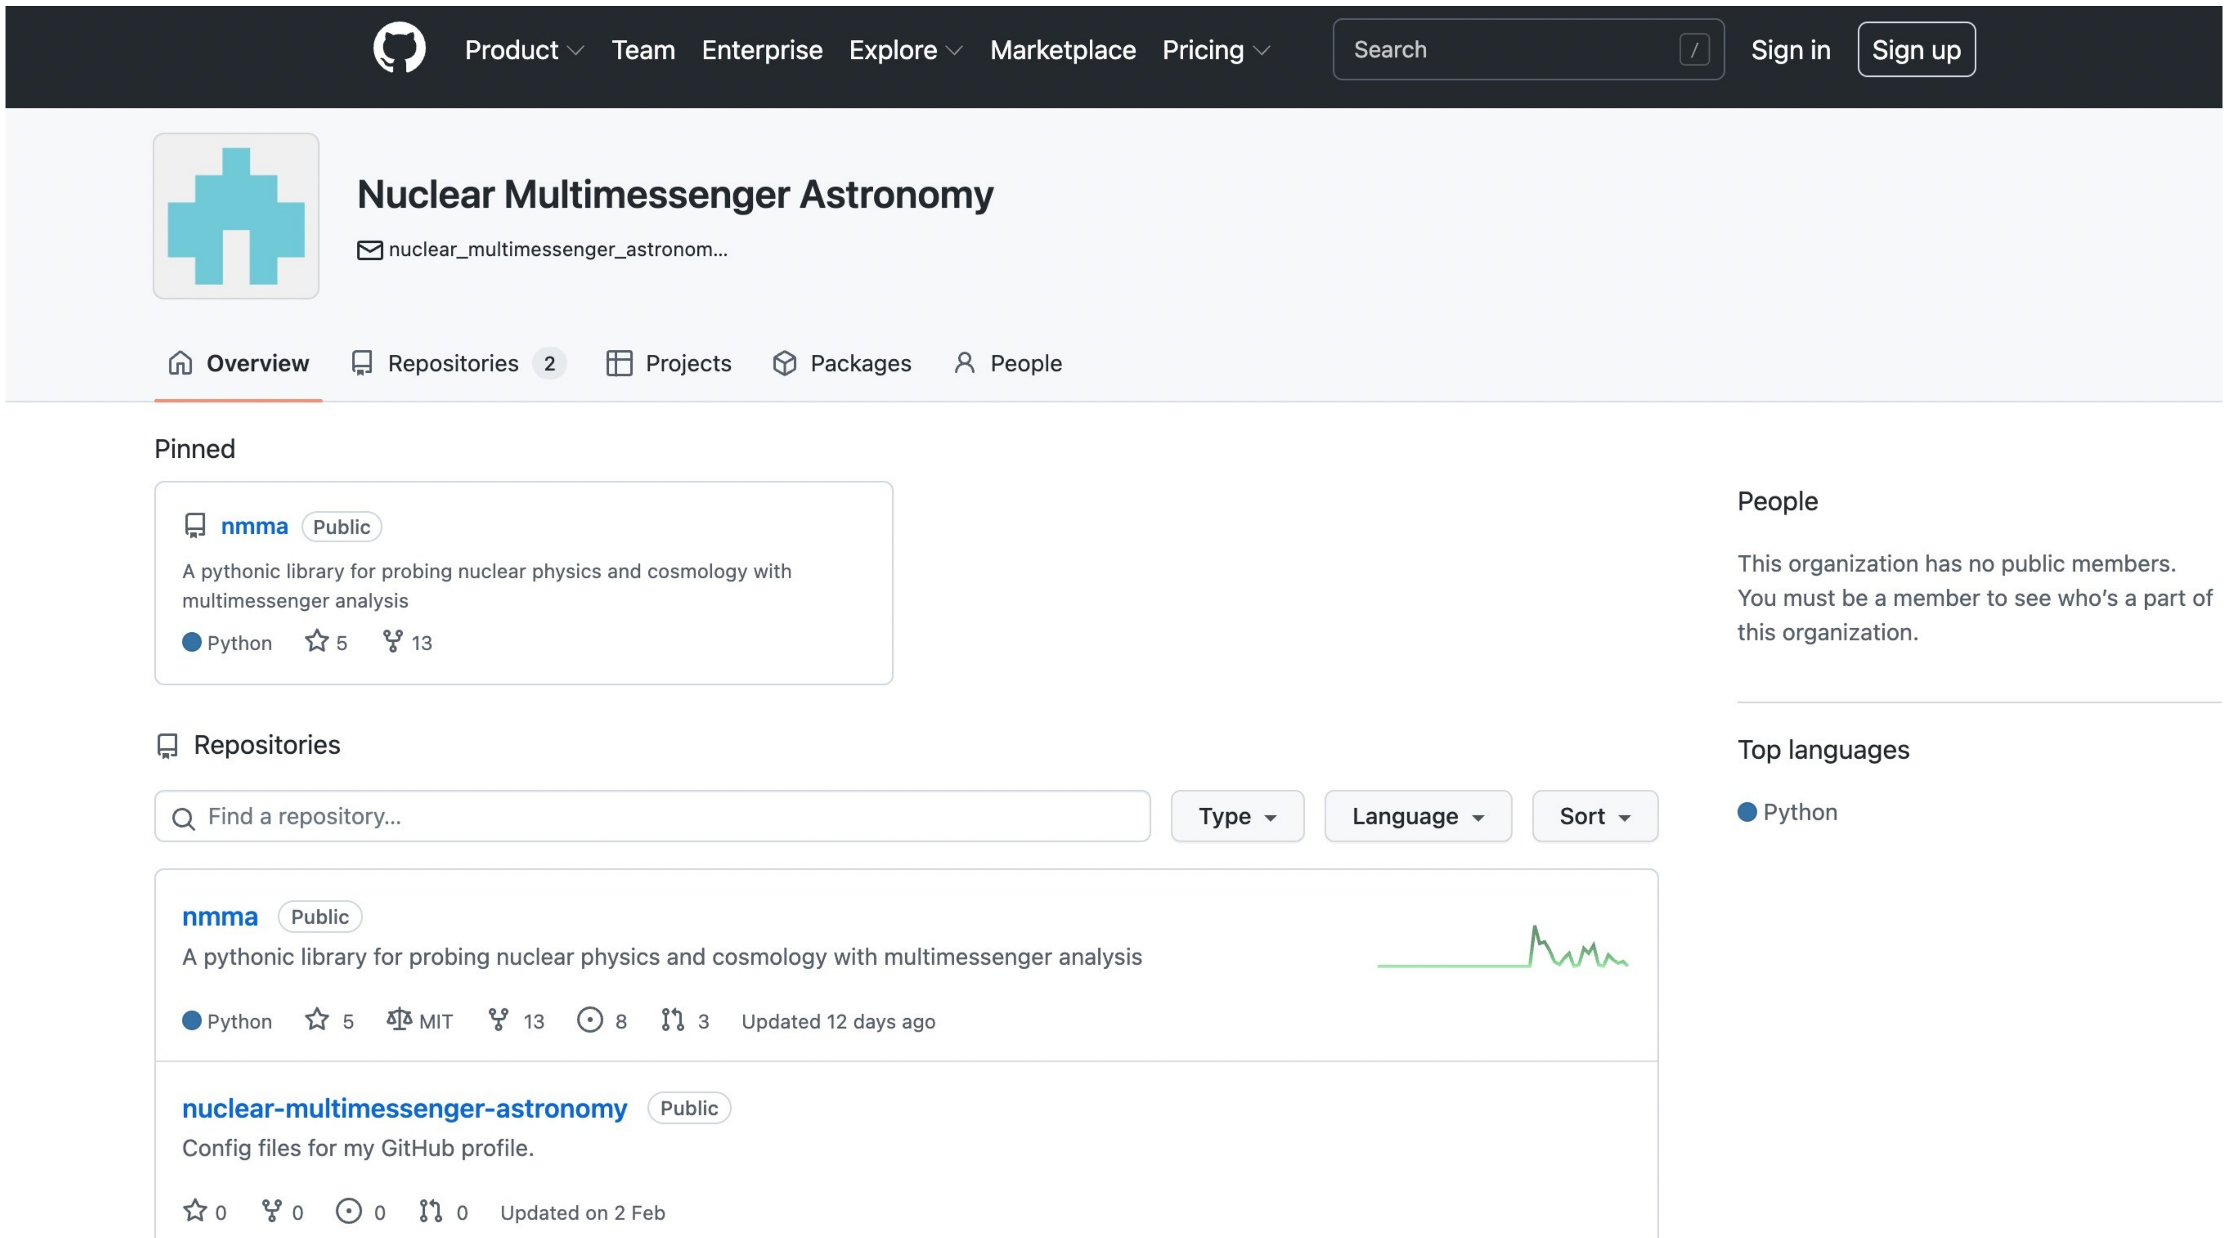

NMMA - Consider using it!

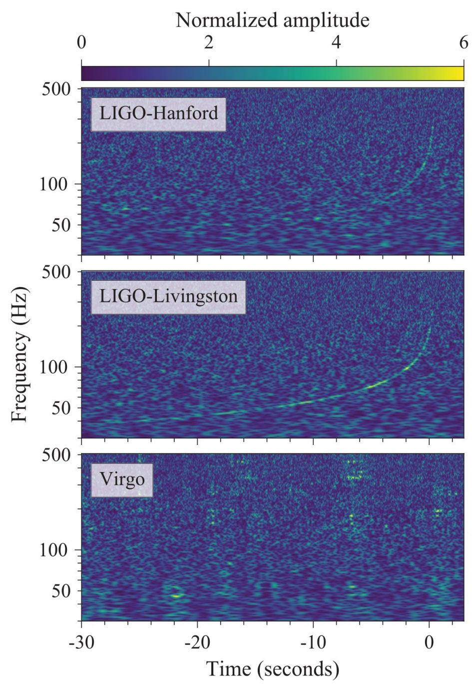

+

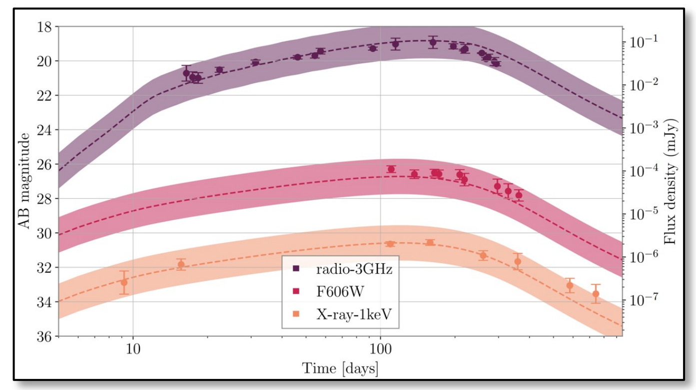

+

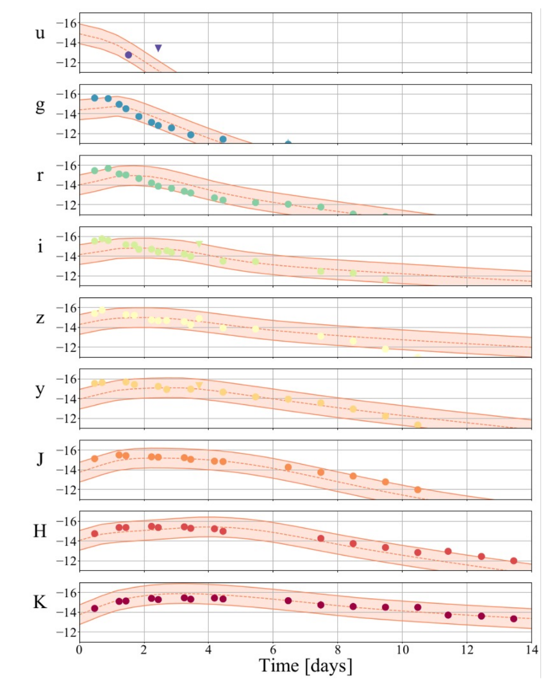

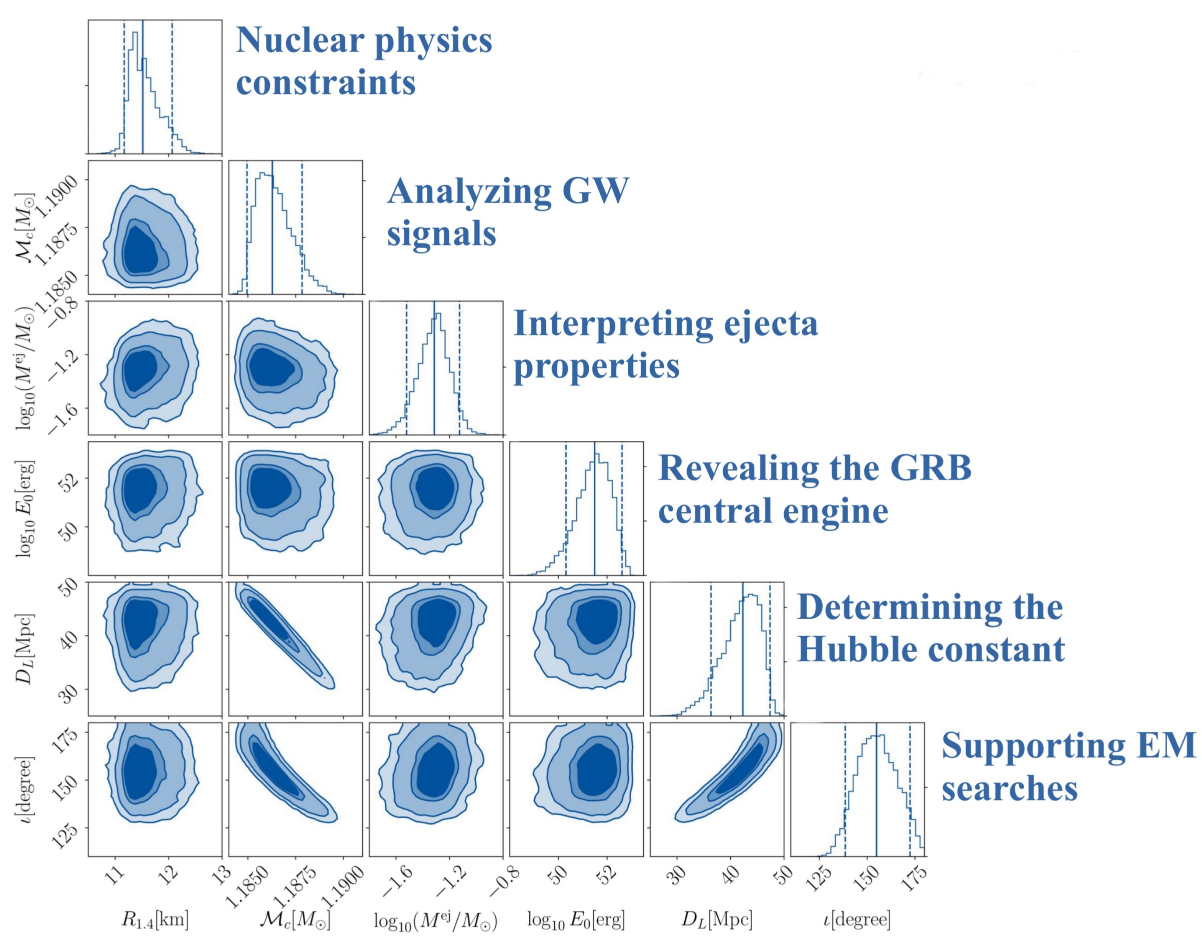

- incorporation of nuclear-physics information
- <u>simultaneous</u> analysis of GW, kilonova, and GRB afterglow
- (some) unit testing and (growing) documentation!
- Used in online kilonova searches
- HPC facilities needed

A Multi-Messenger Ecosystem

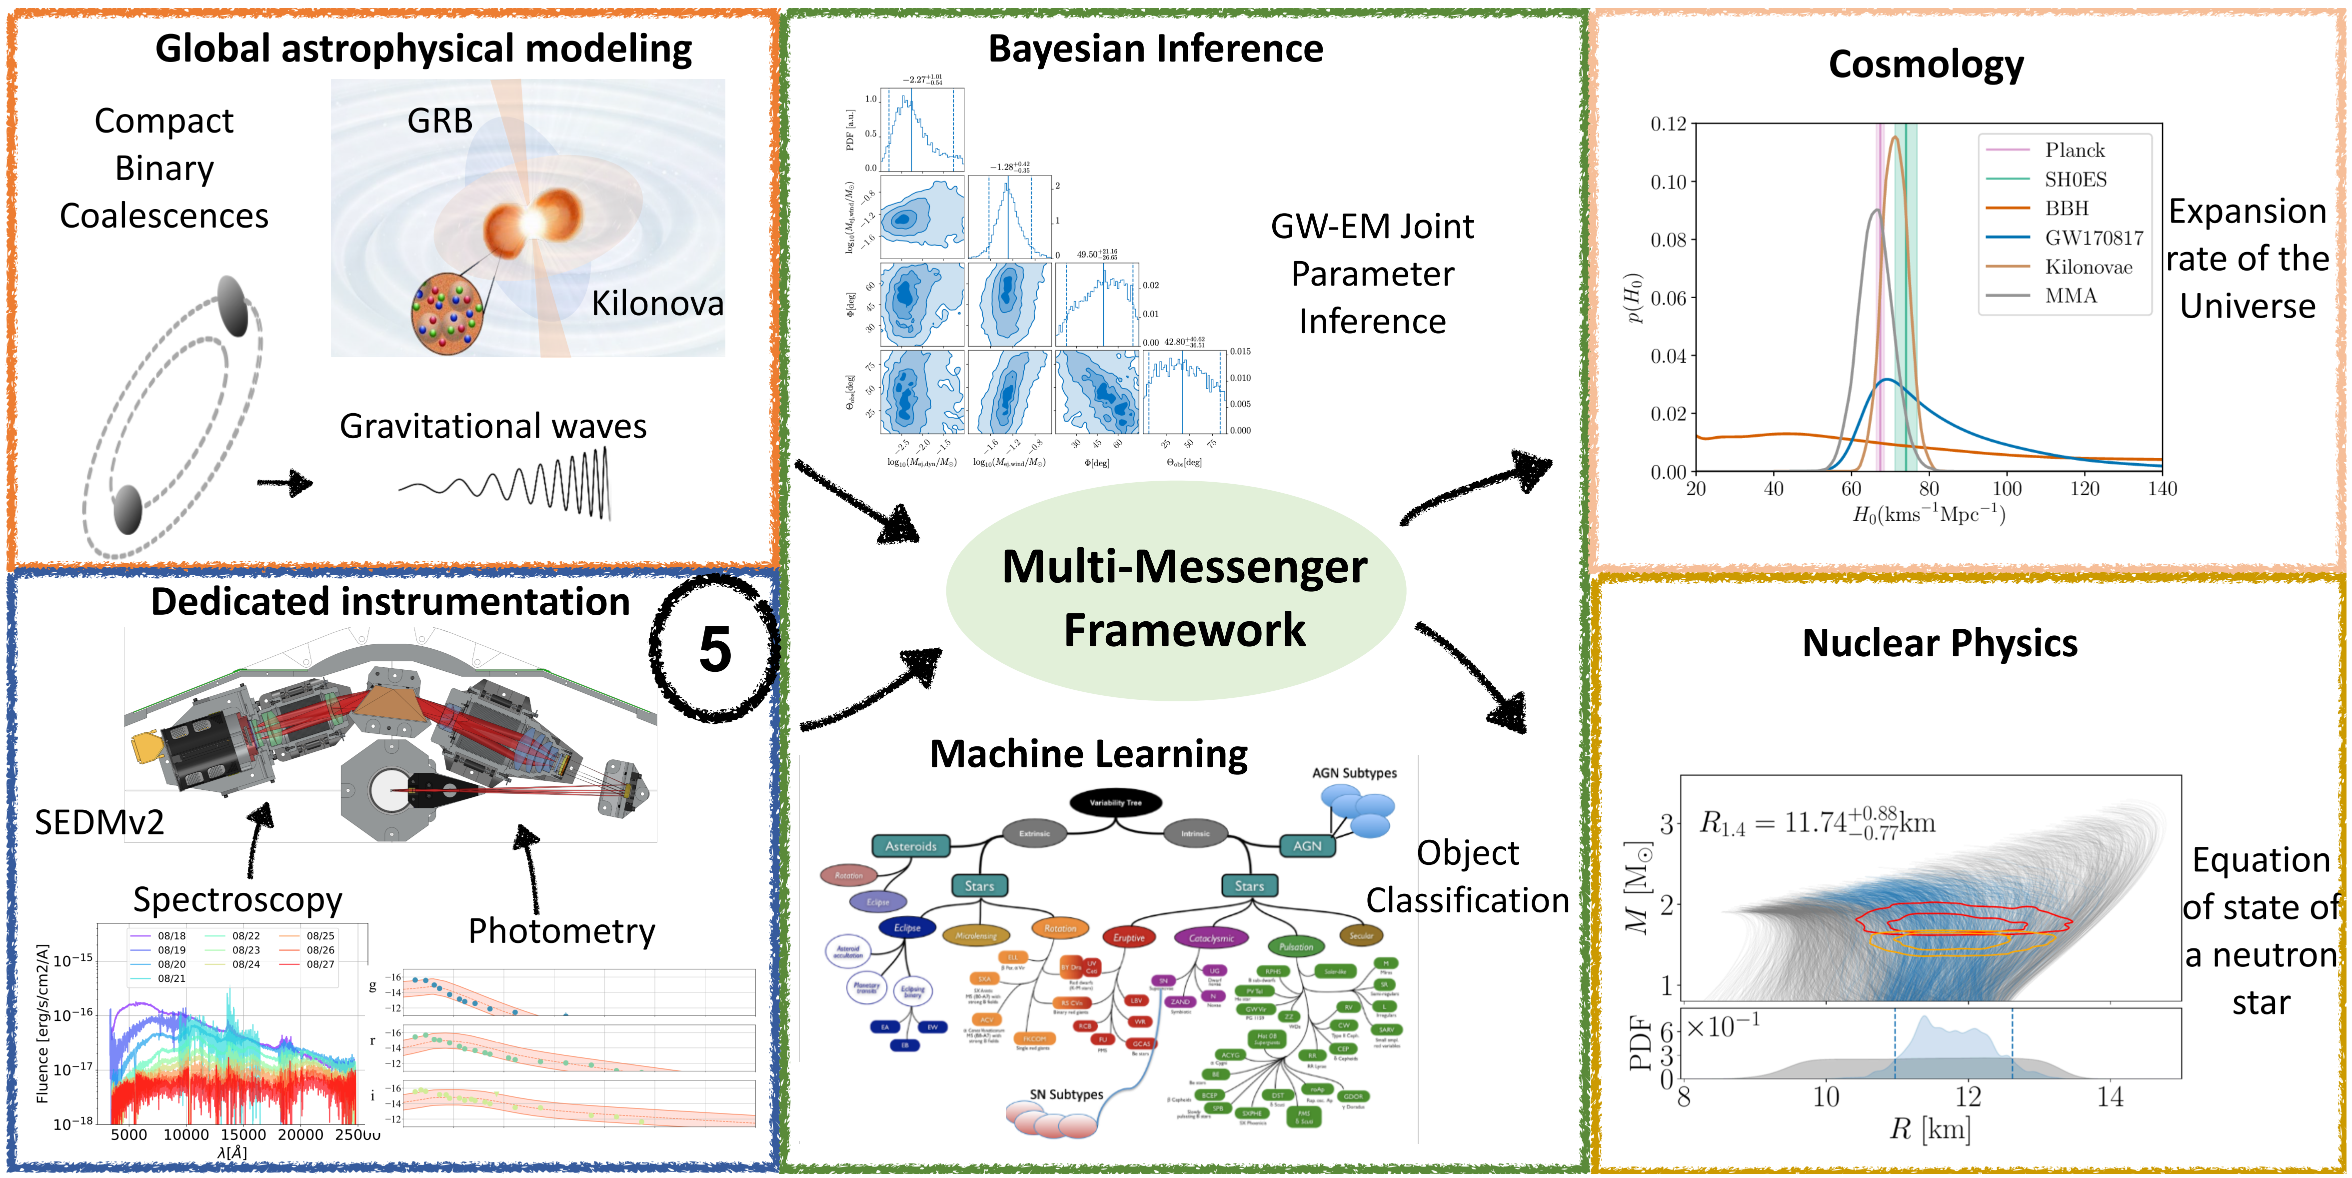

Kilonovae - Hard to find

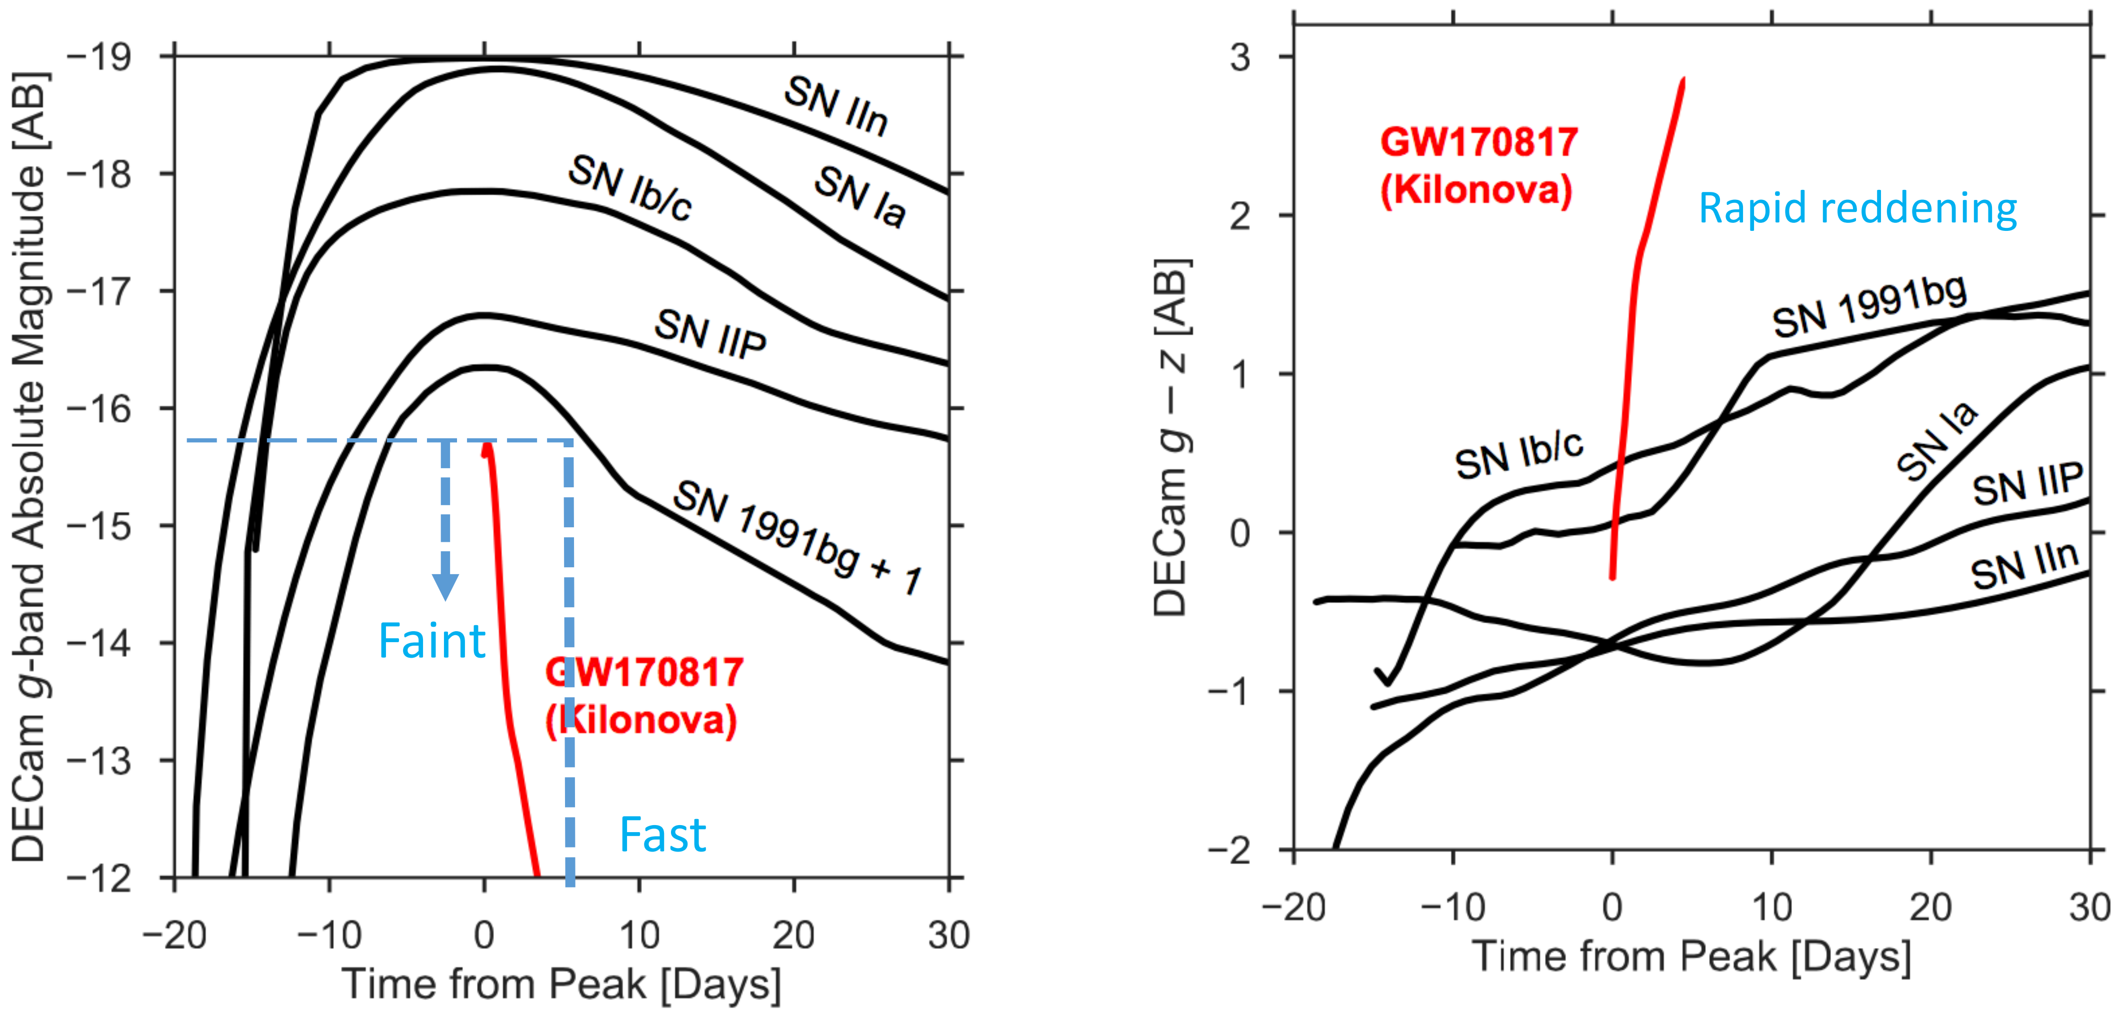

modified from Andreoni+2018, LSST White Paper

The Observational Landscape

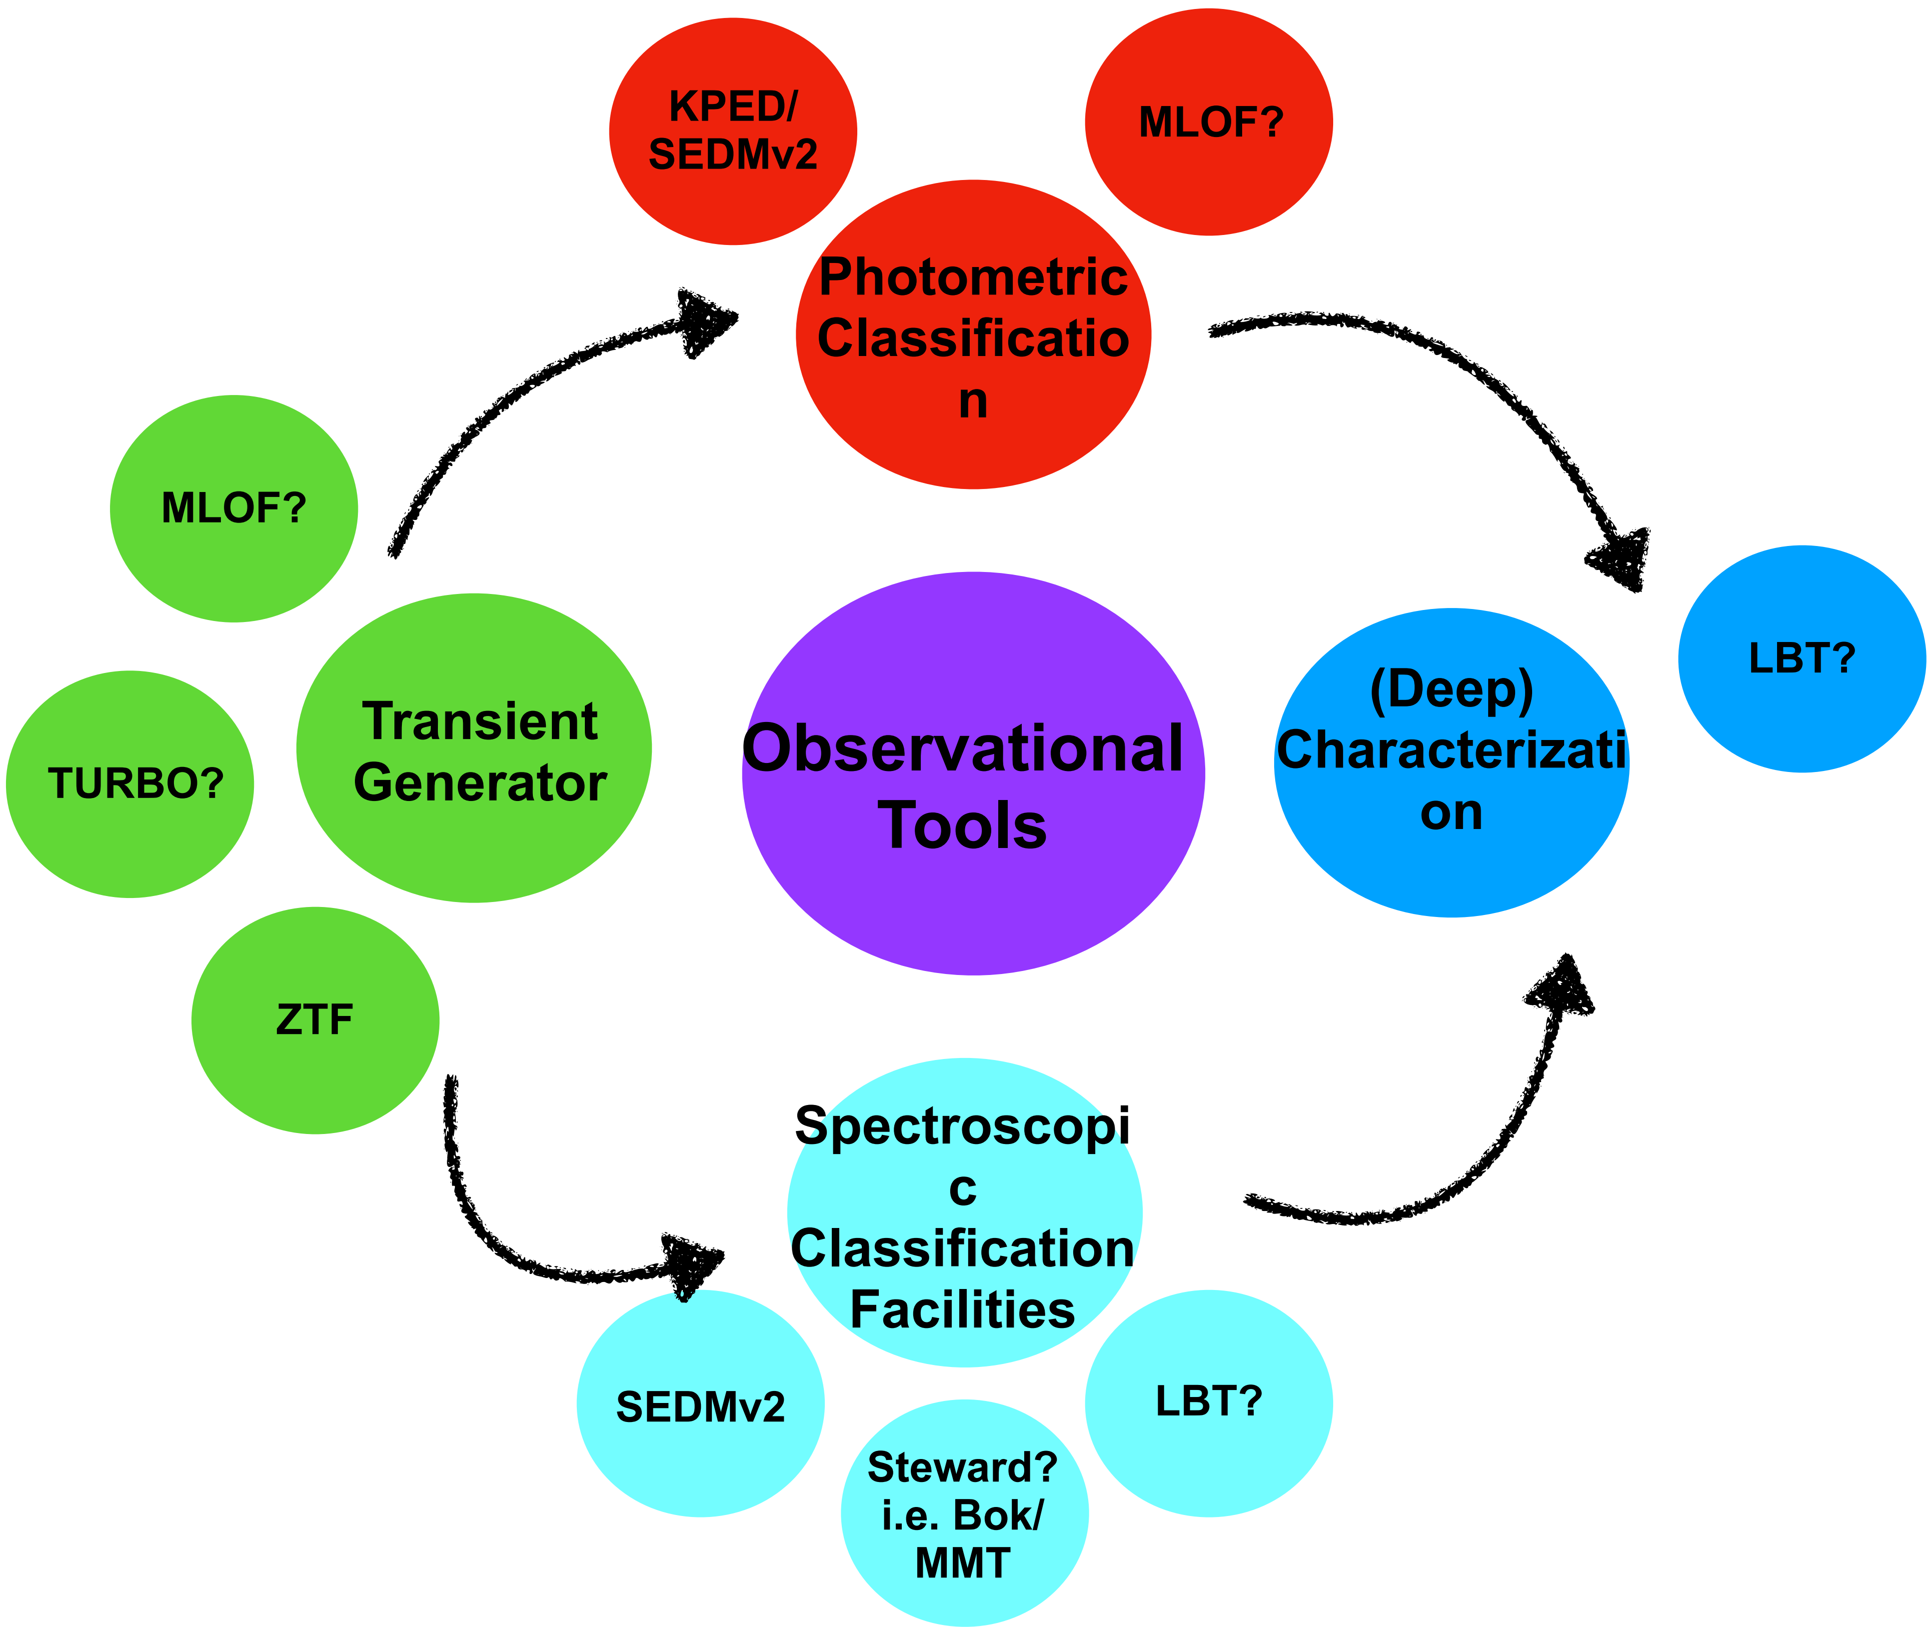

The Observational Landscape

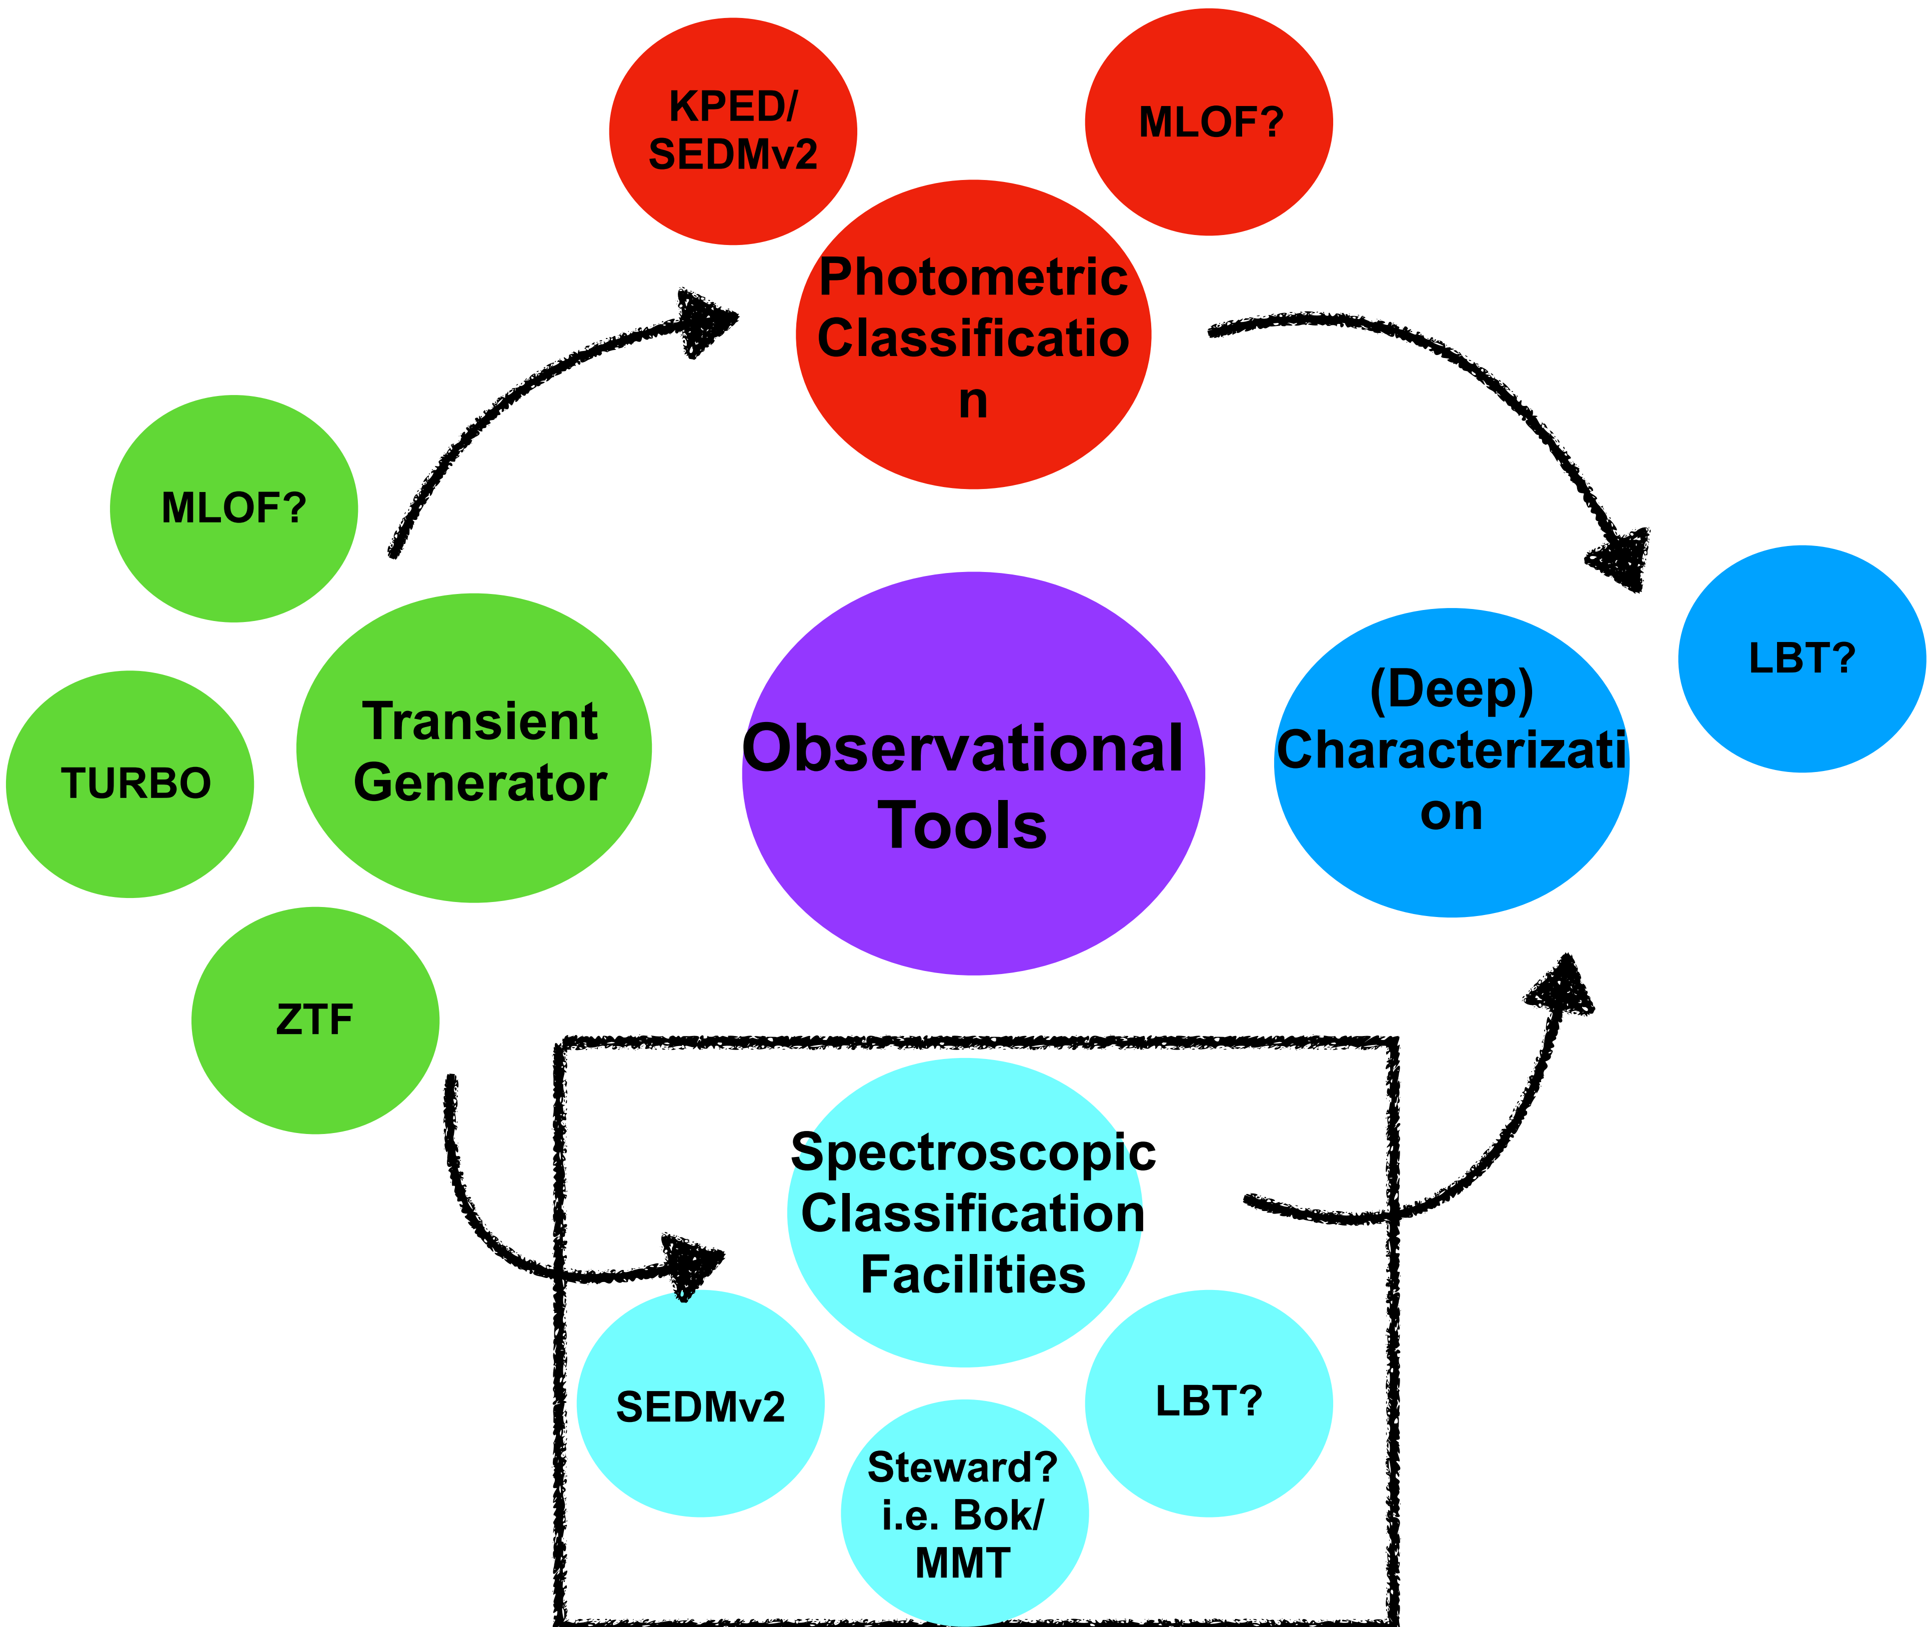

Classifying Astronomical Transients

In 2019 alone, 18,296 ATs were identified by various surveys

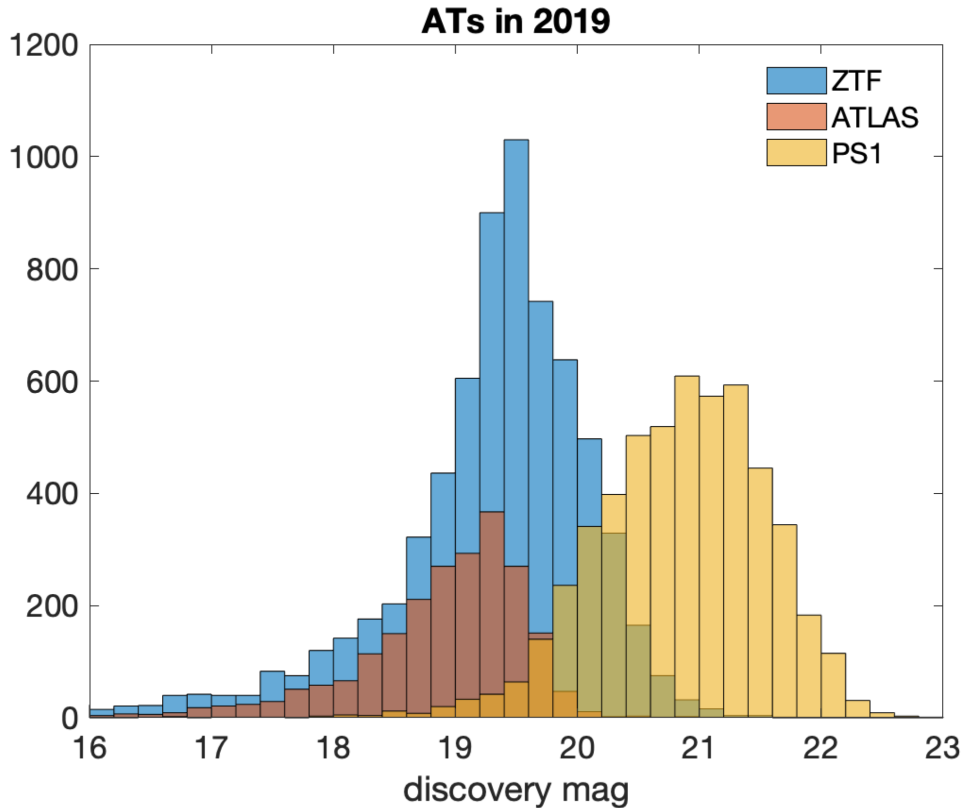

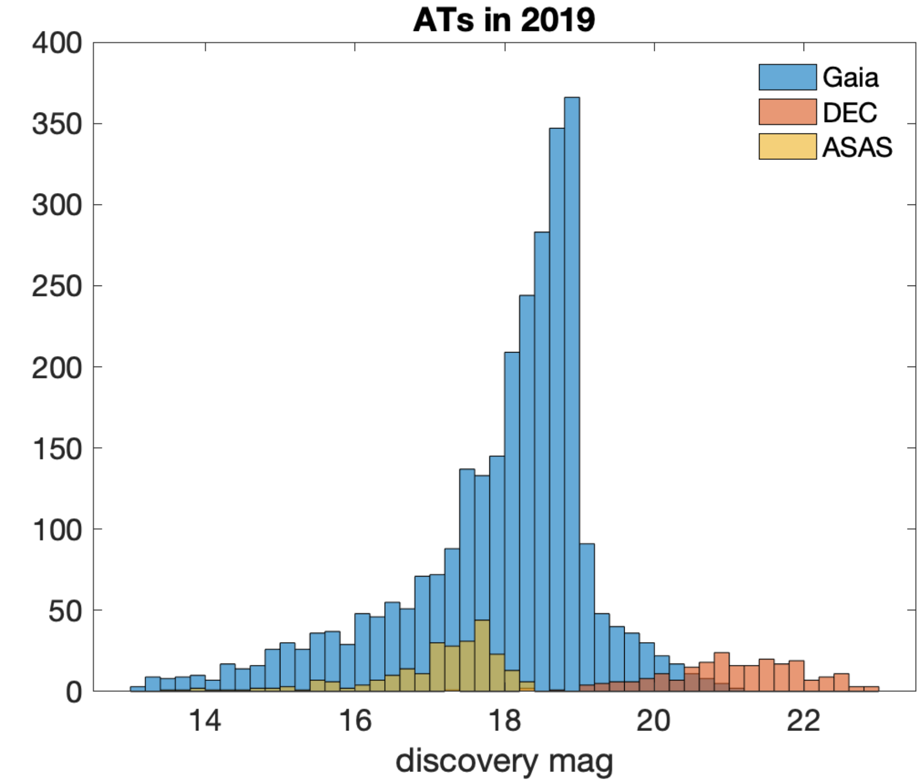

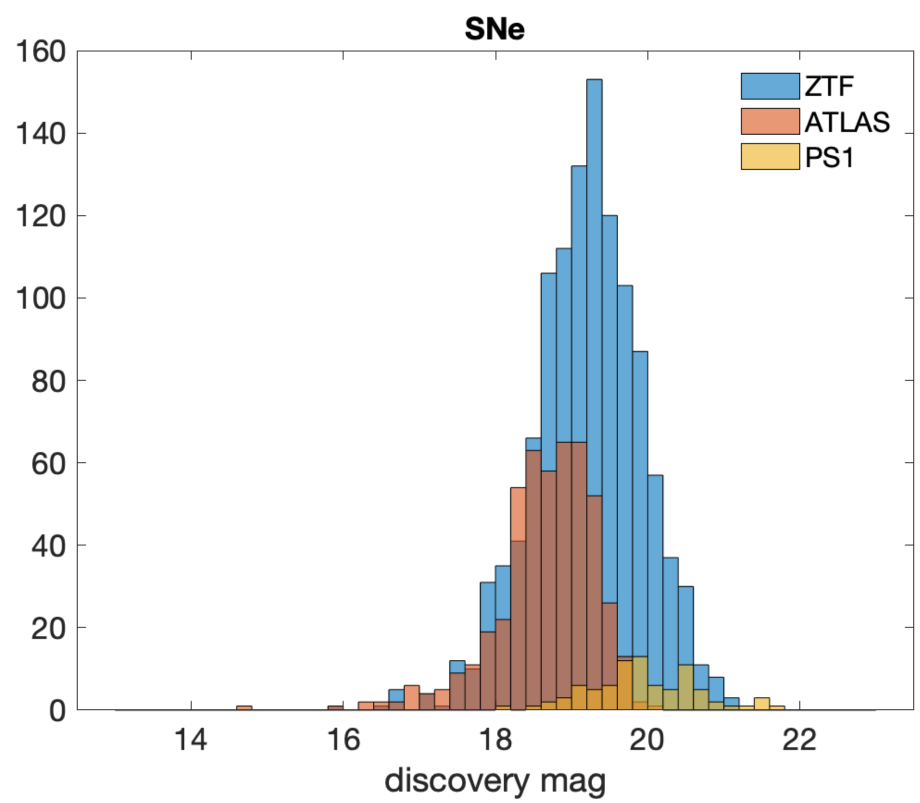

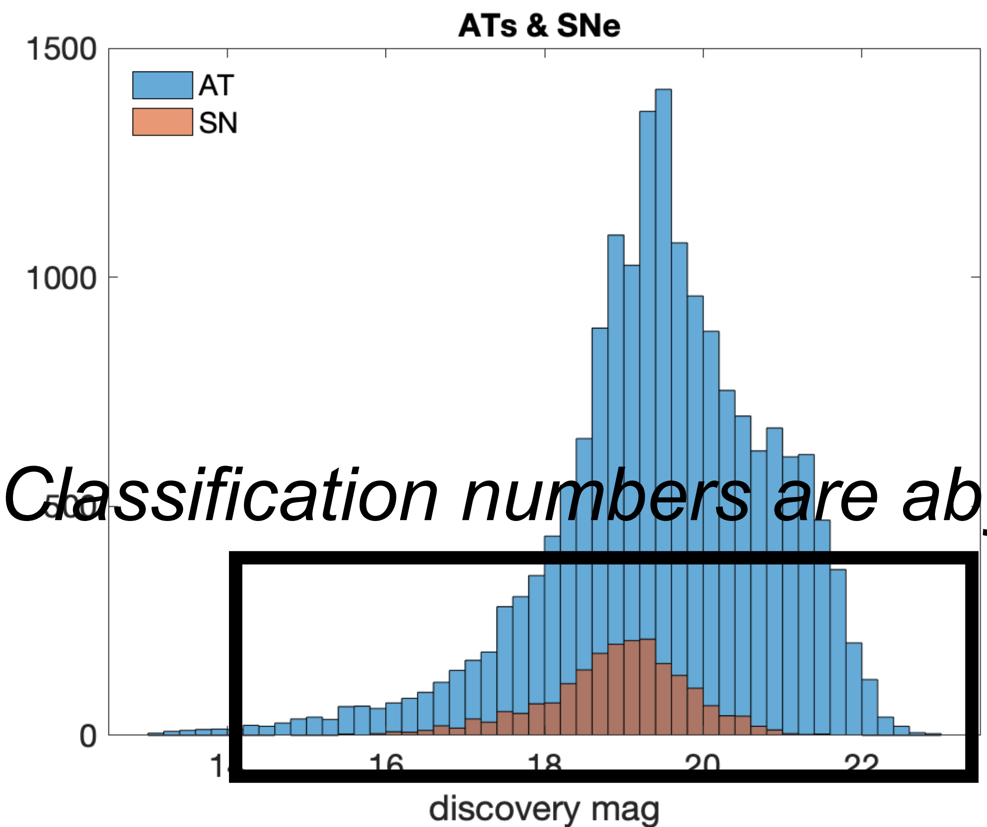

Classification numbers are abysmal

The SED Machine

Hyperspectral imaging spectrograph

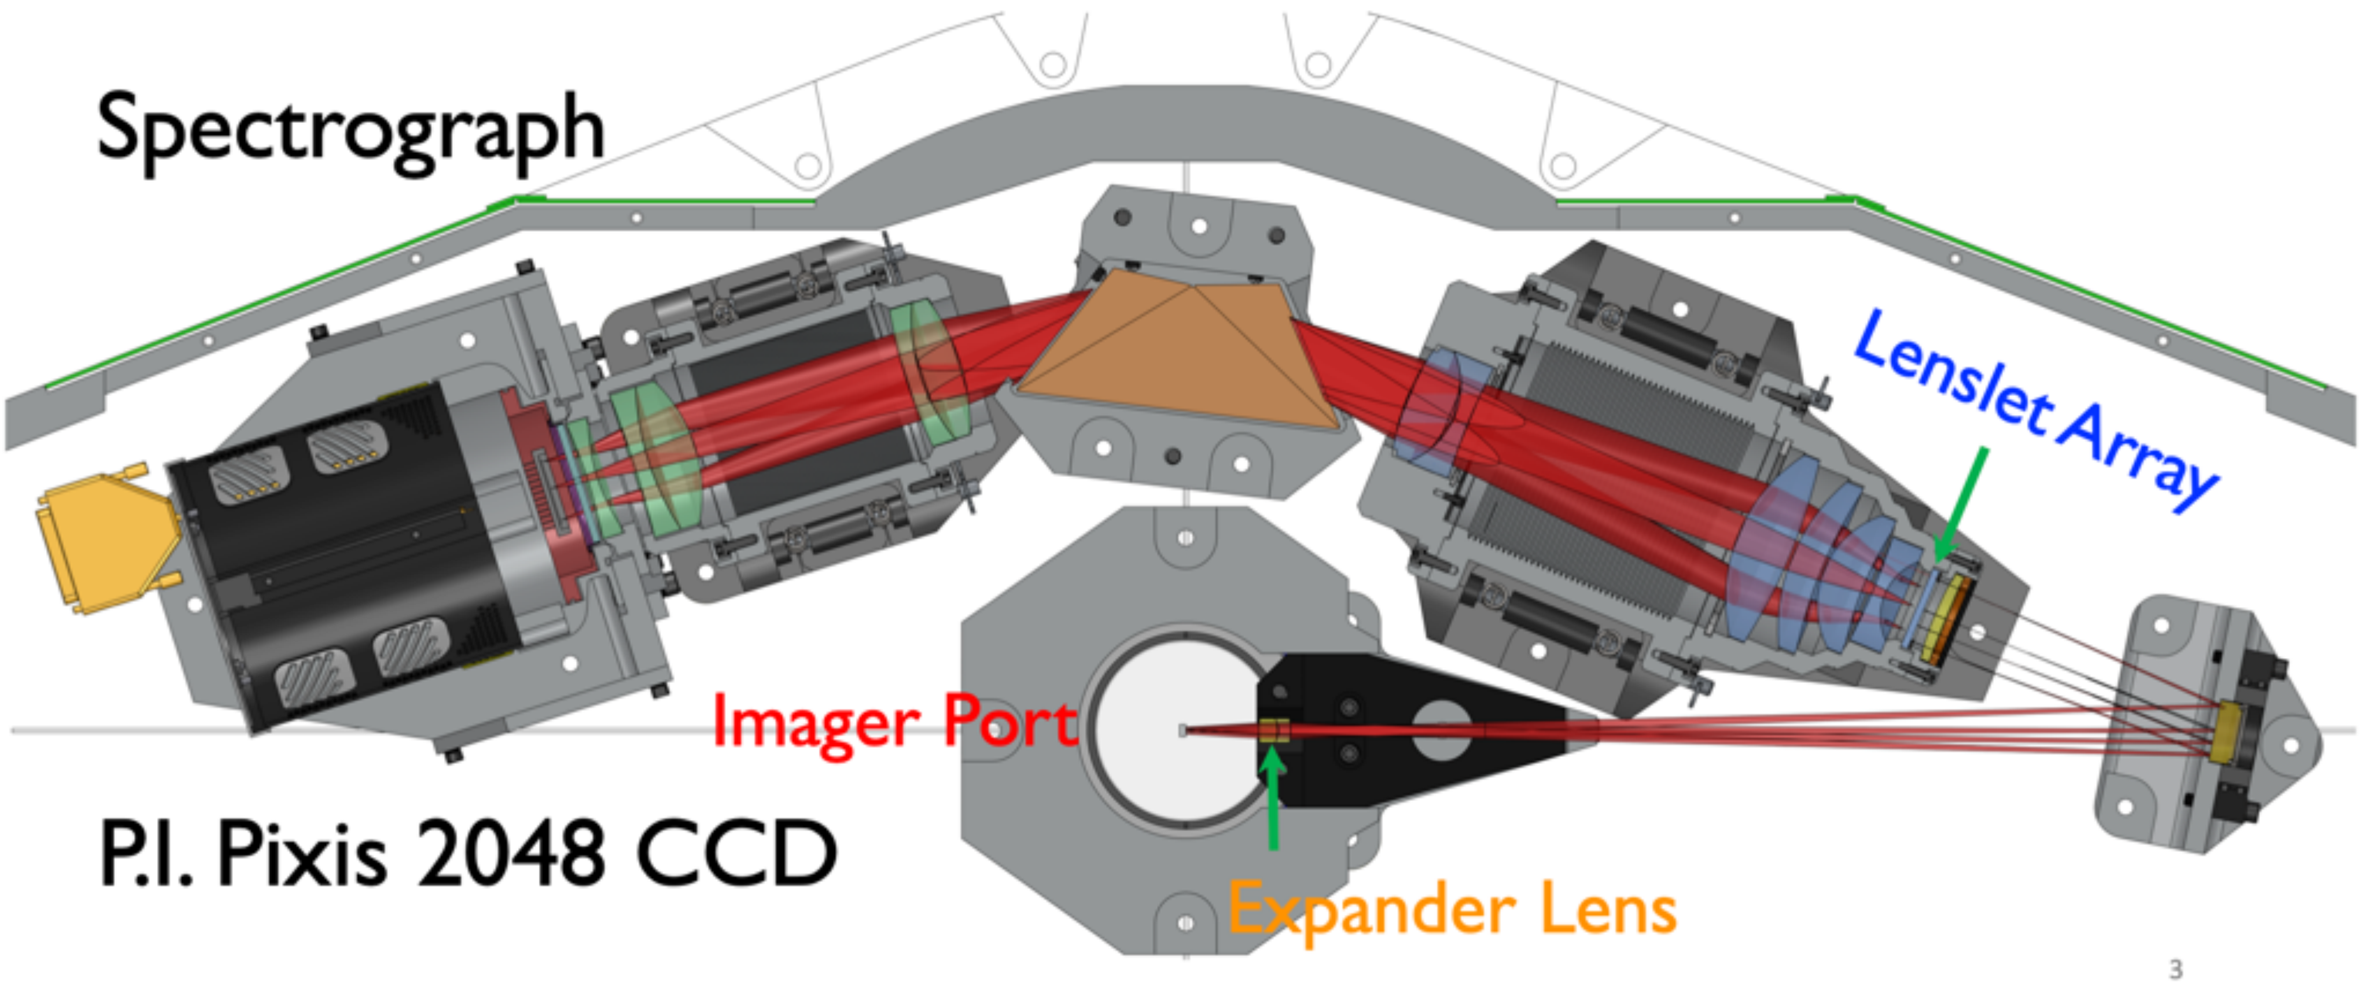

The SED Machine

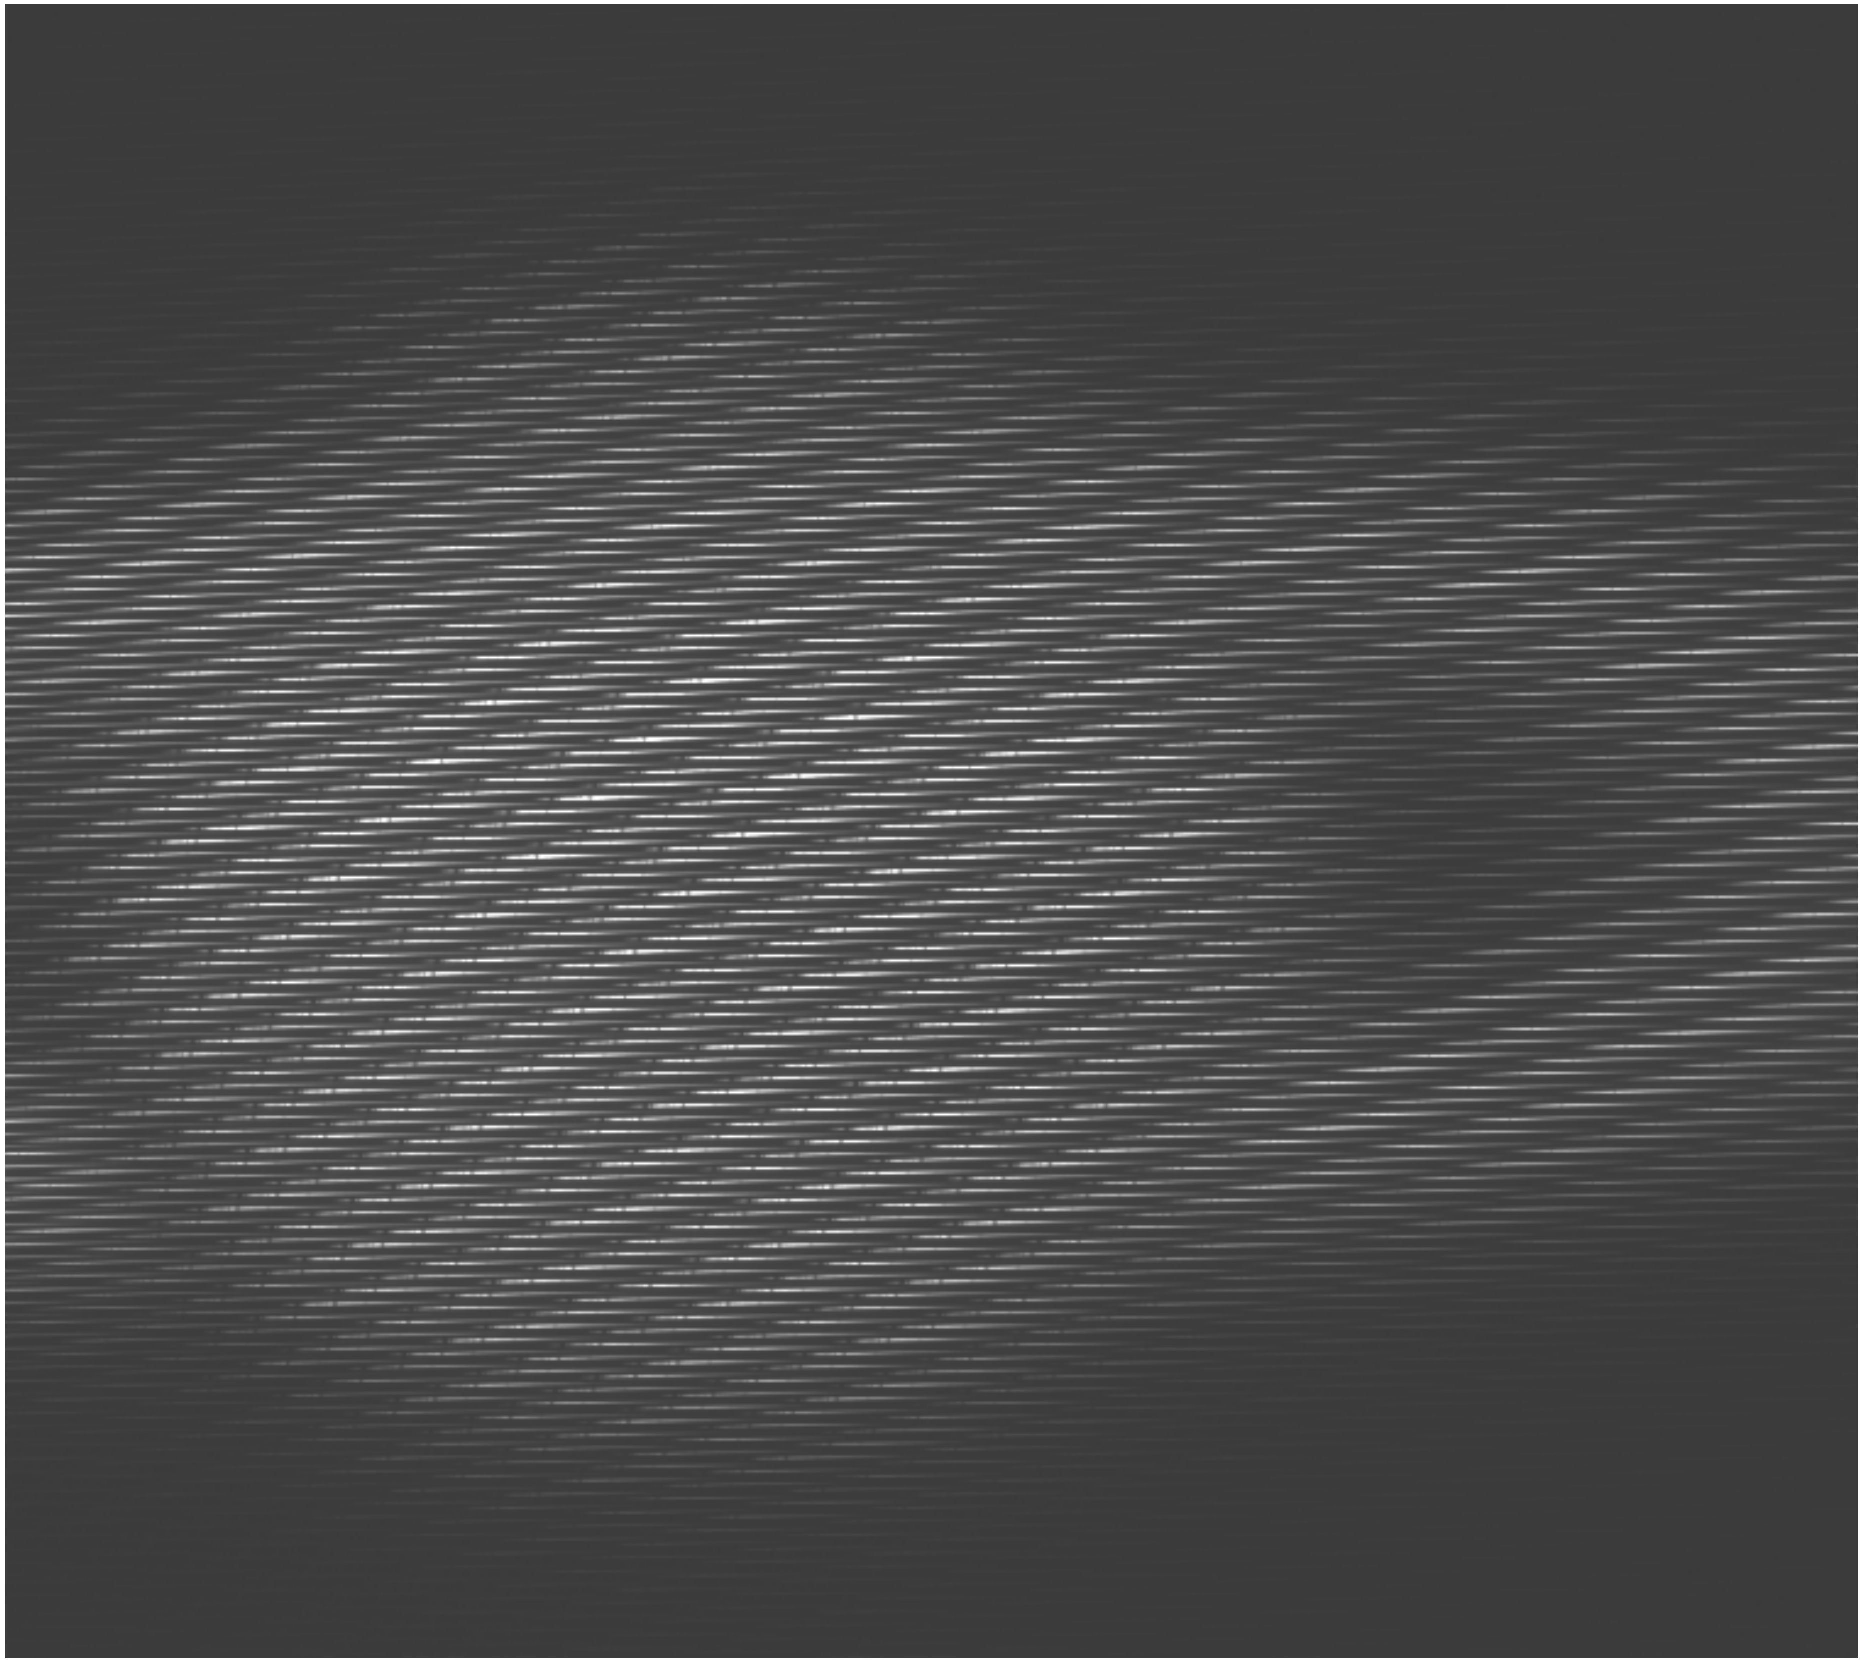

The SED Machine (zoomed)

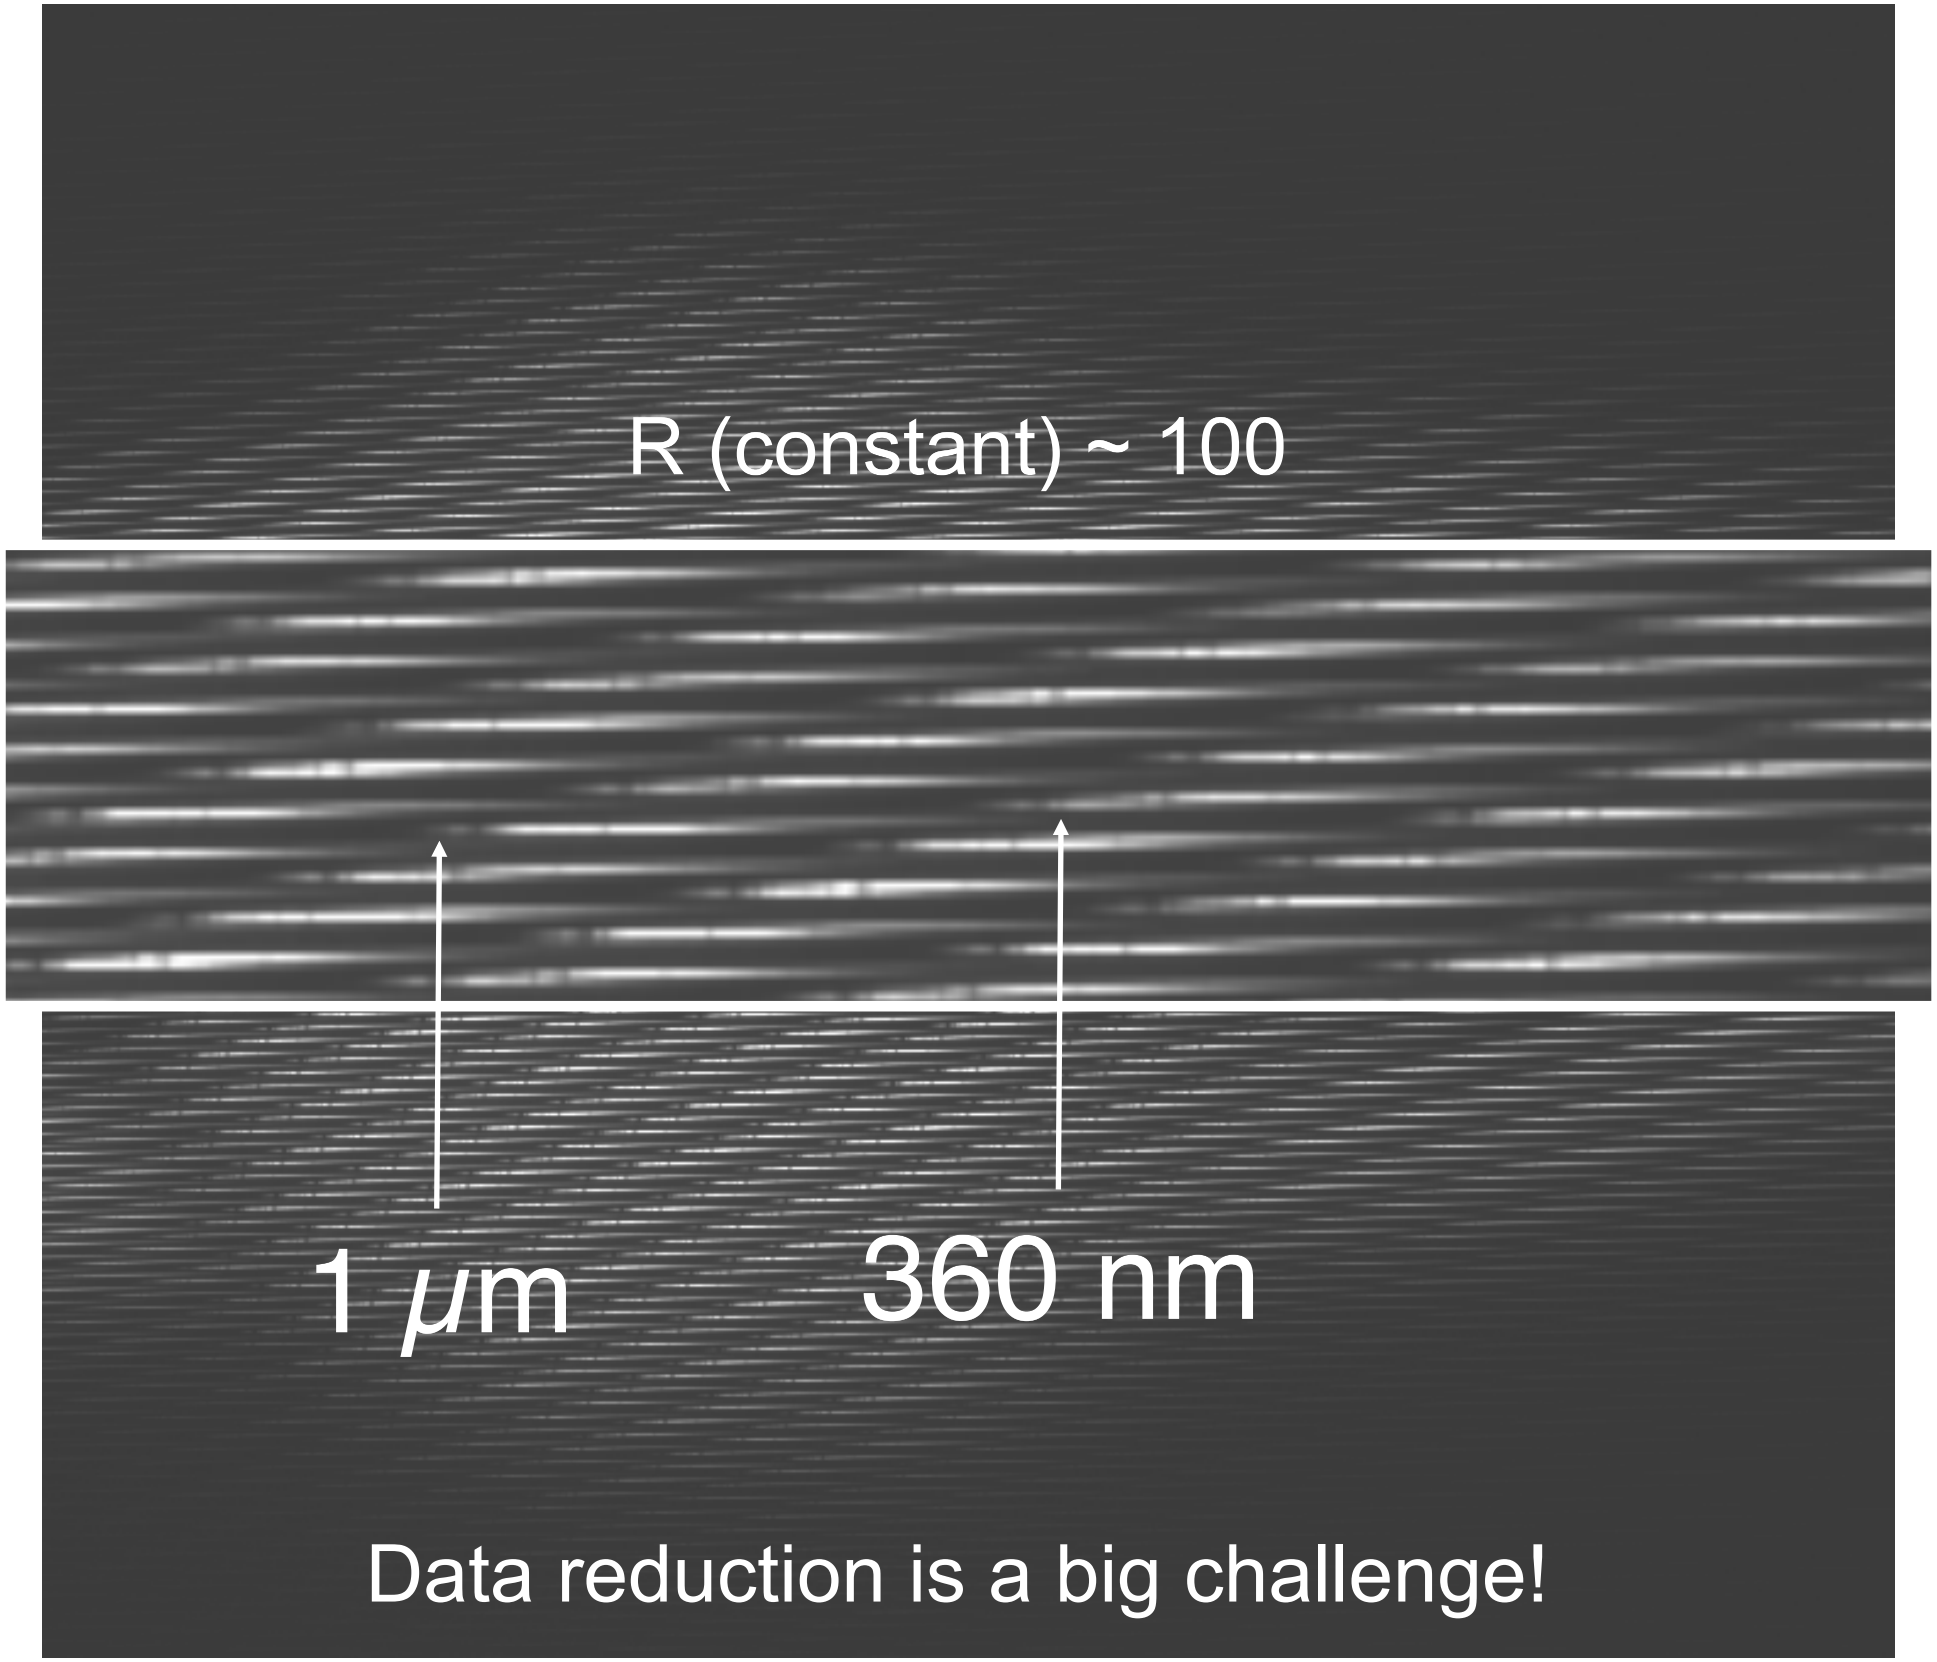

Comparison to higher resolution instruments

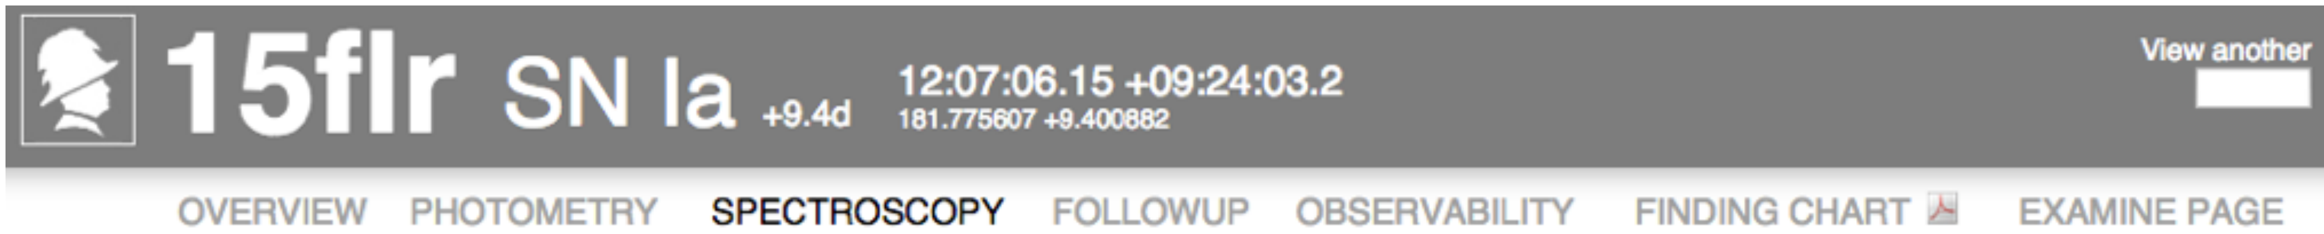

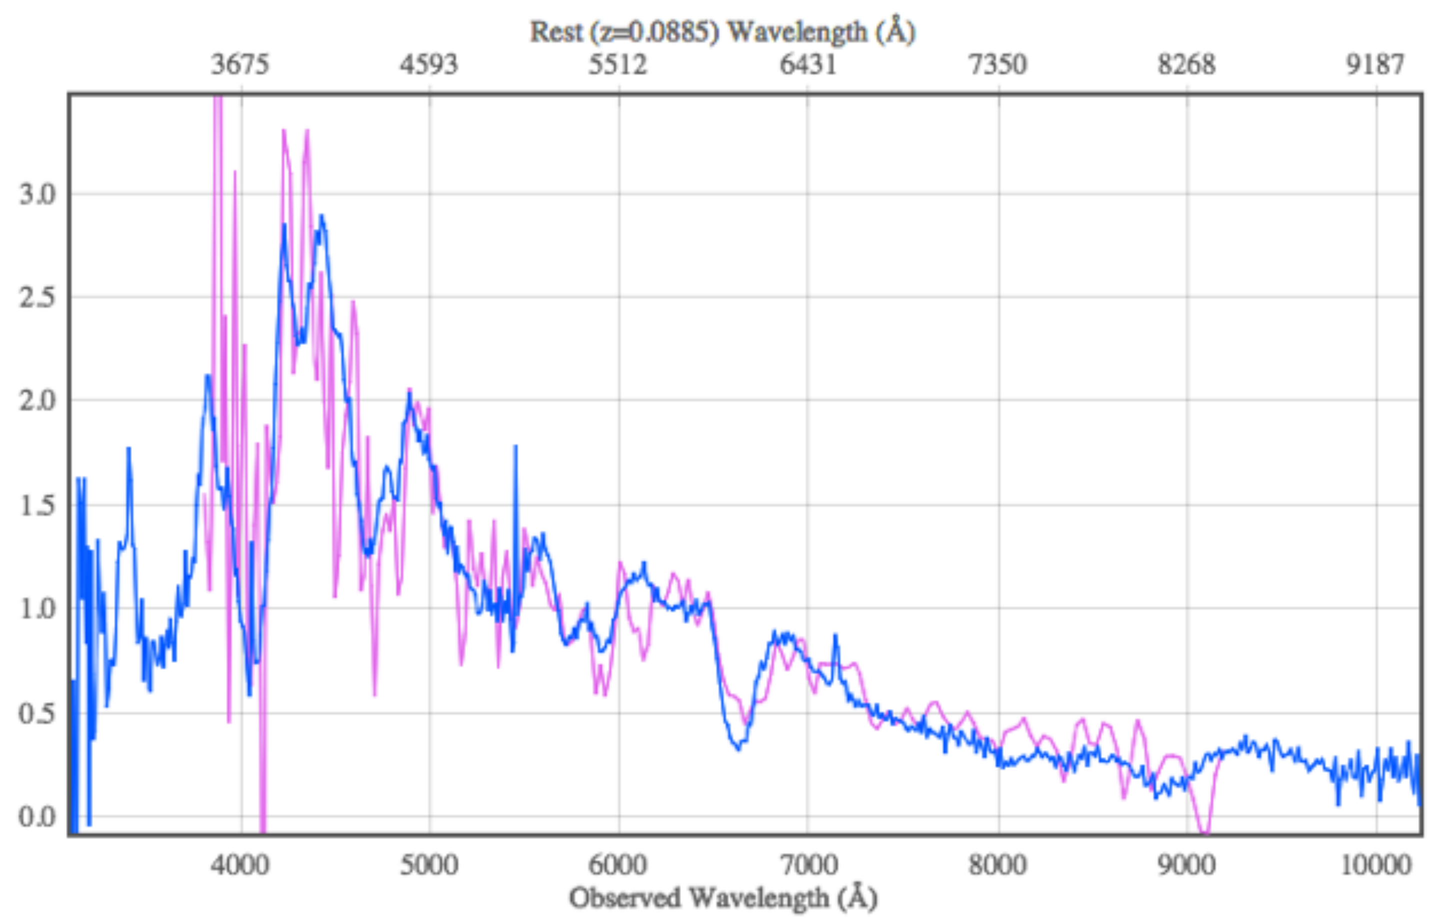

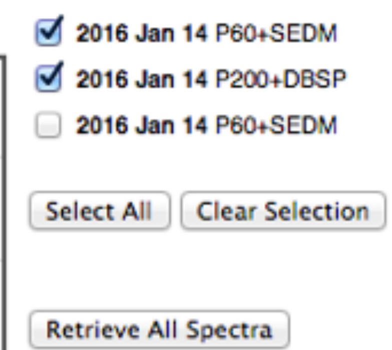

SEDM: Pink, 1800s, 60-inch
DBSP: Blue, 500s, 200-inch

- 50% of all SN classifications on TNS website
- Integral part of ZTF impact
- 10-20+ Targets Per Night (Depending on time of year)

SED Machine - Kitt Peak

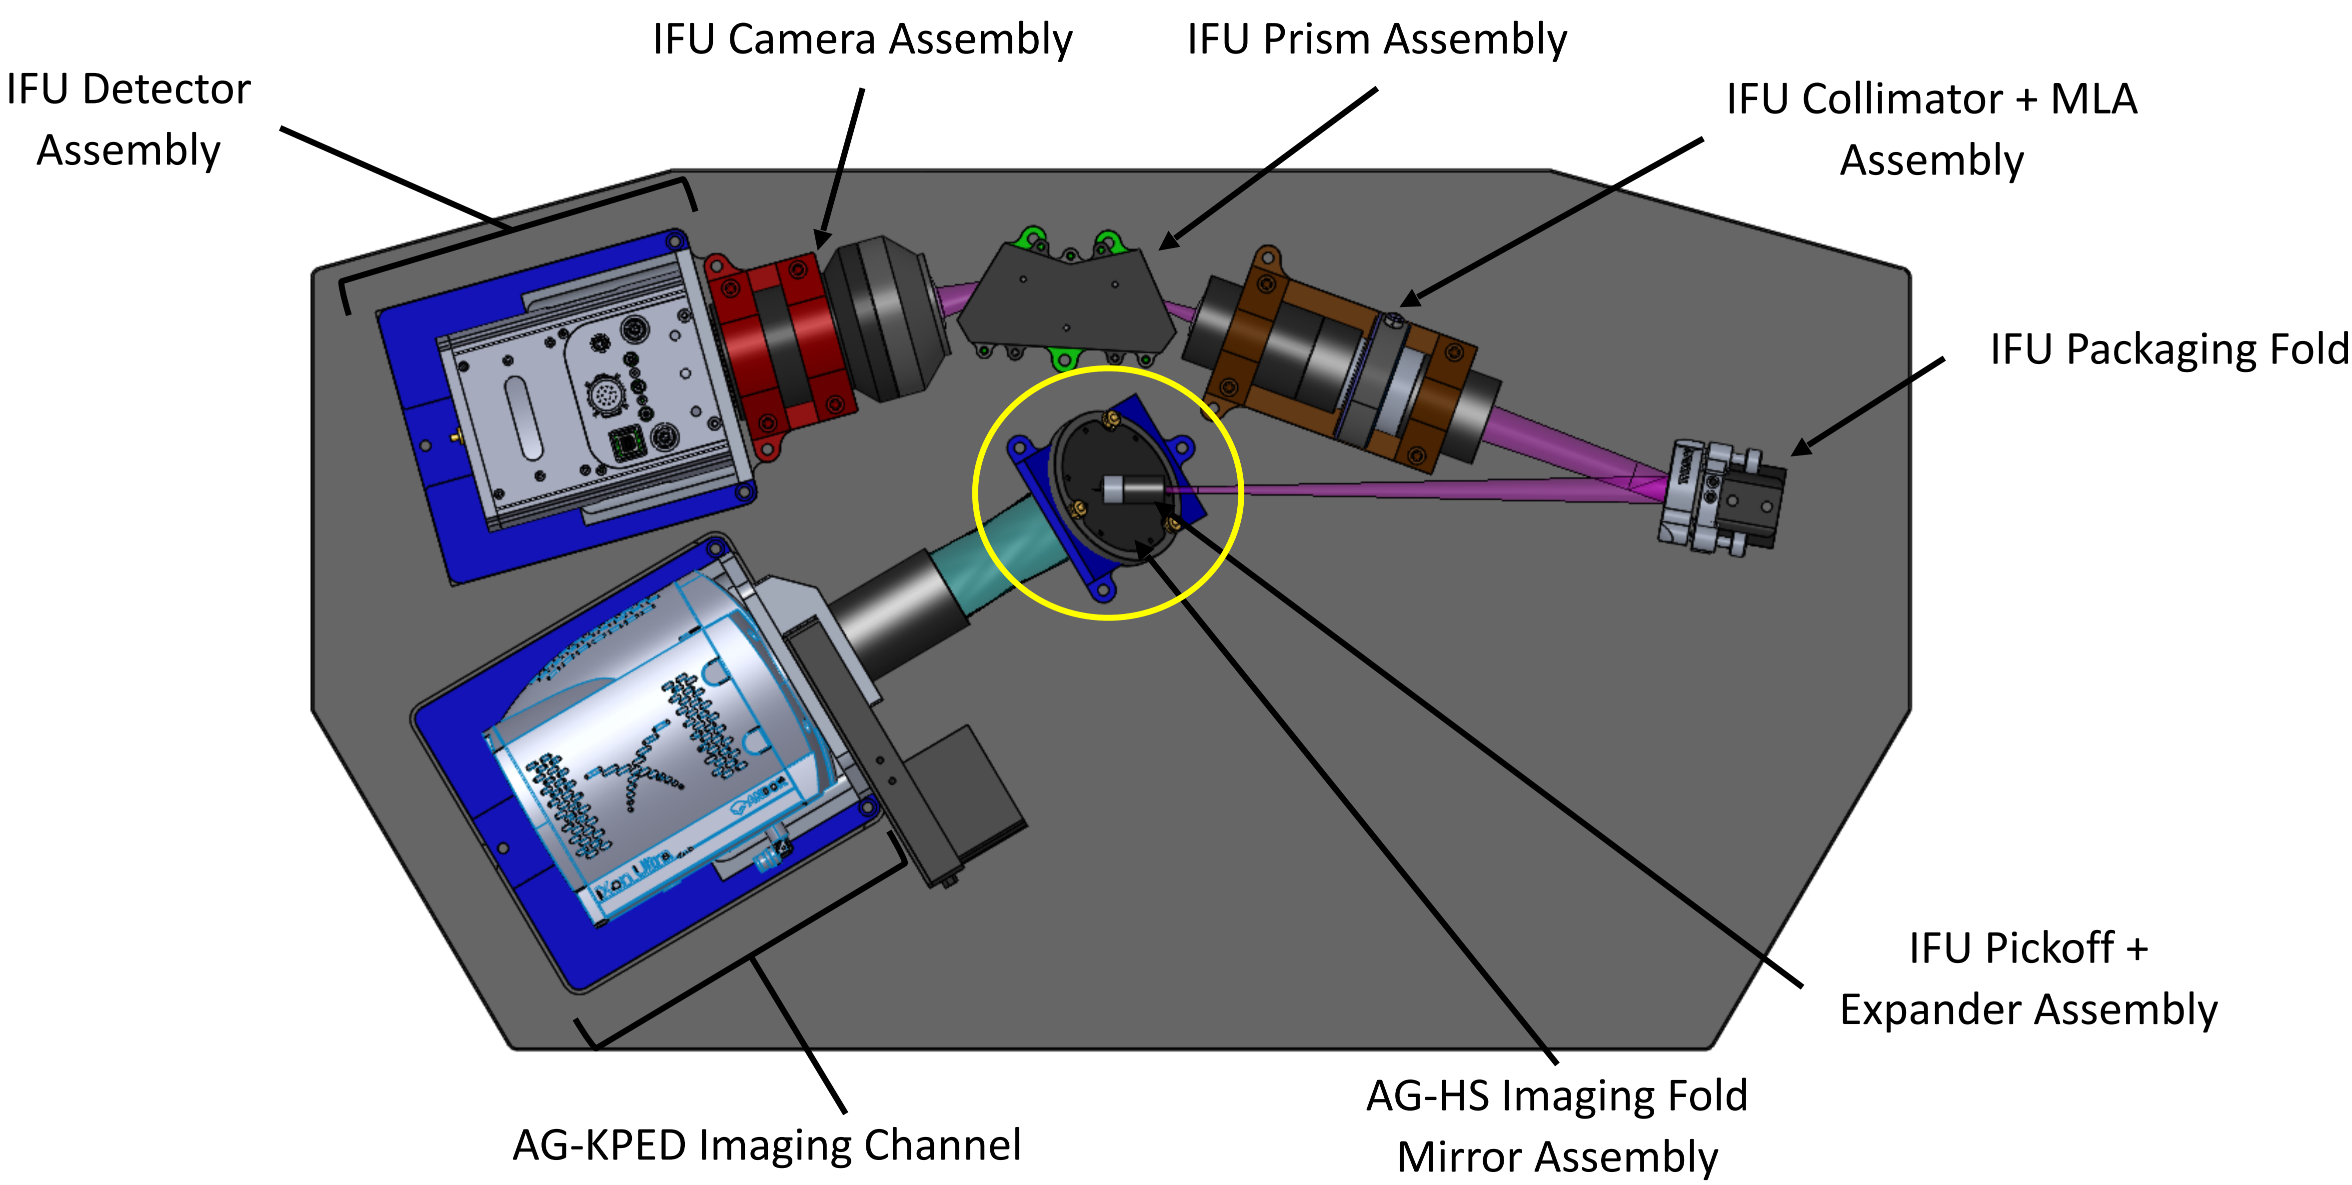

SED Machine - Kitt Peak

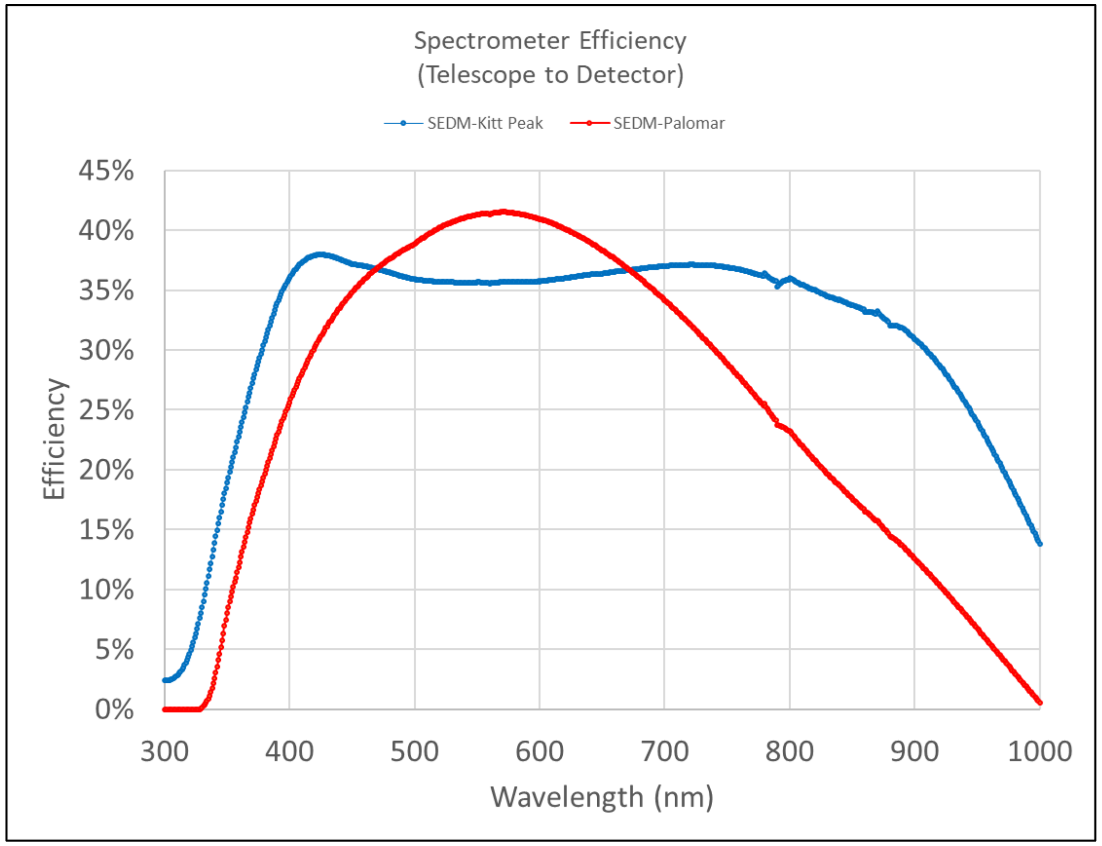

Kitt Peak 2.1m: Facility Specs

- 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

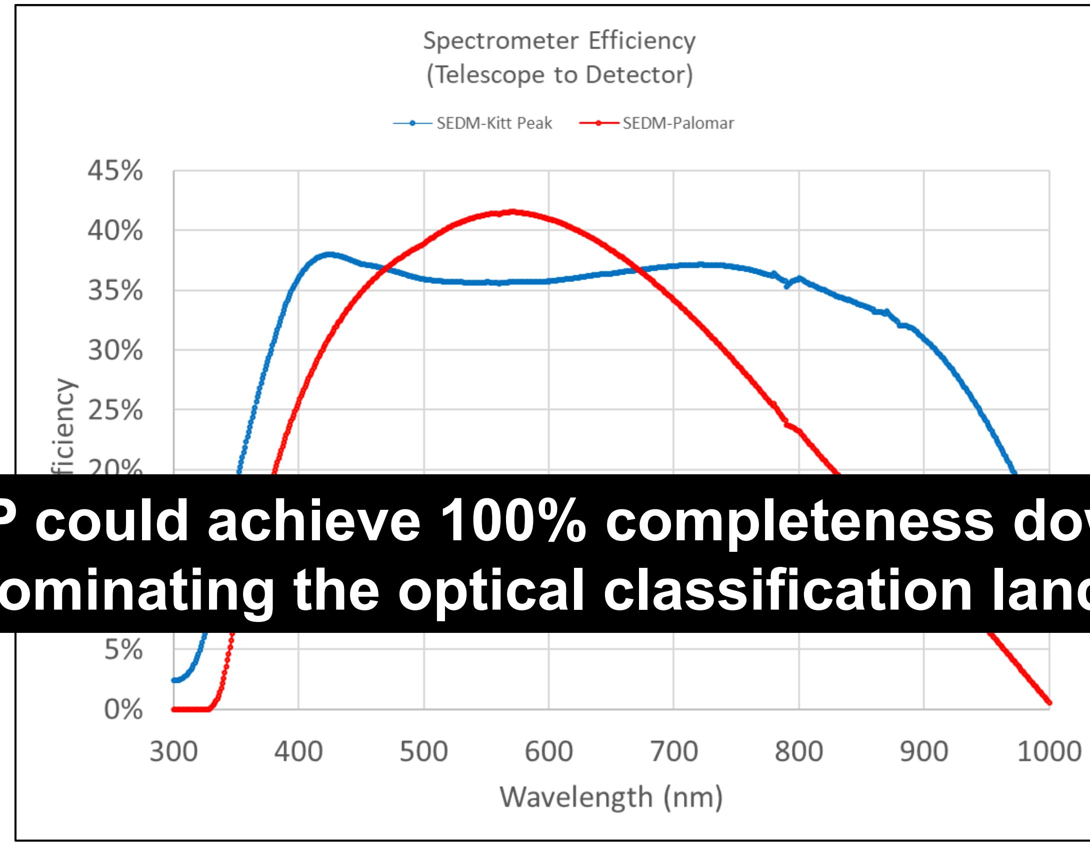

SEDM-KP could achieve 100% completeness down to 19+ mag (dominating the optical classification landscape)

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

SED Machine - Kitt Peak: First Science Observations

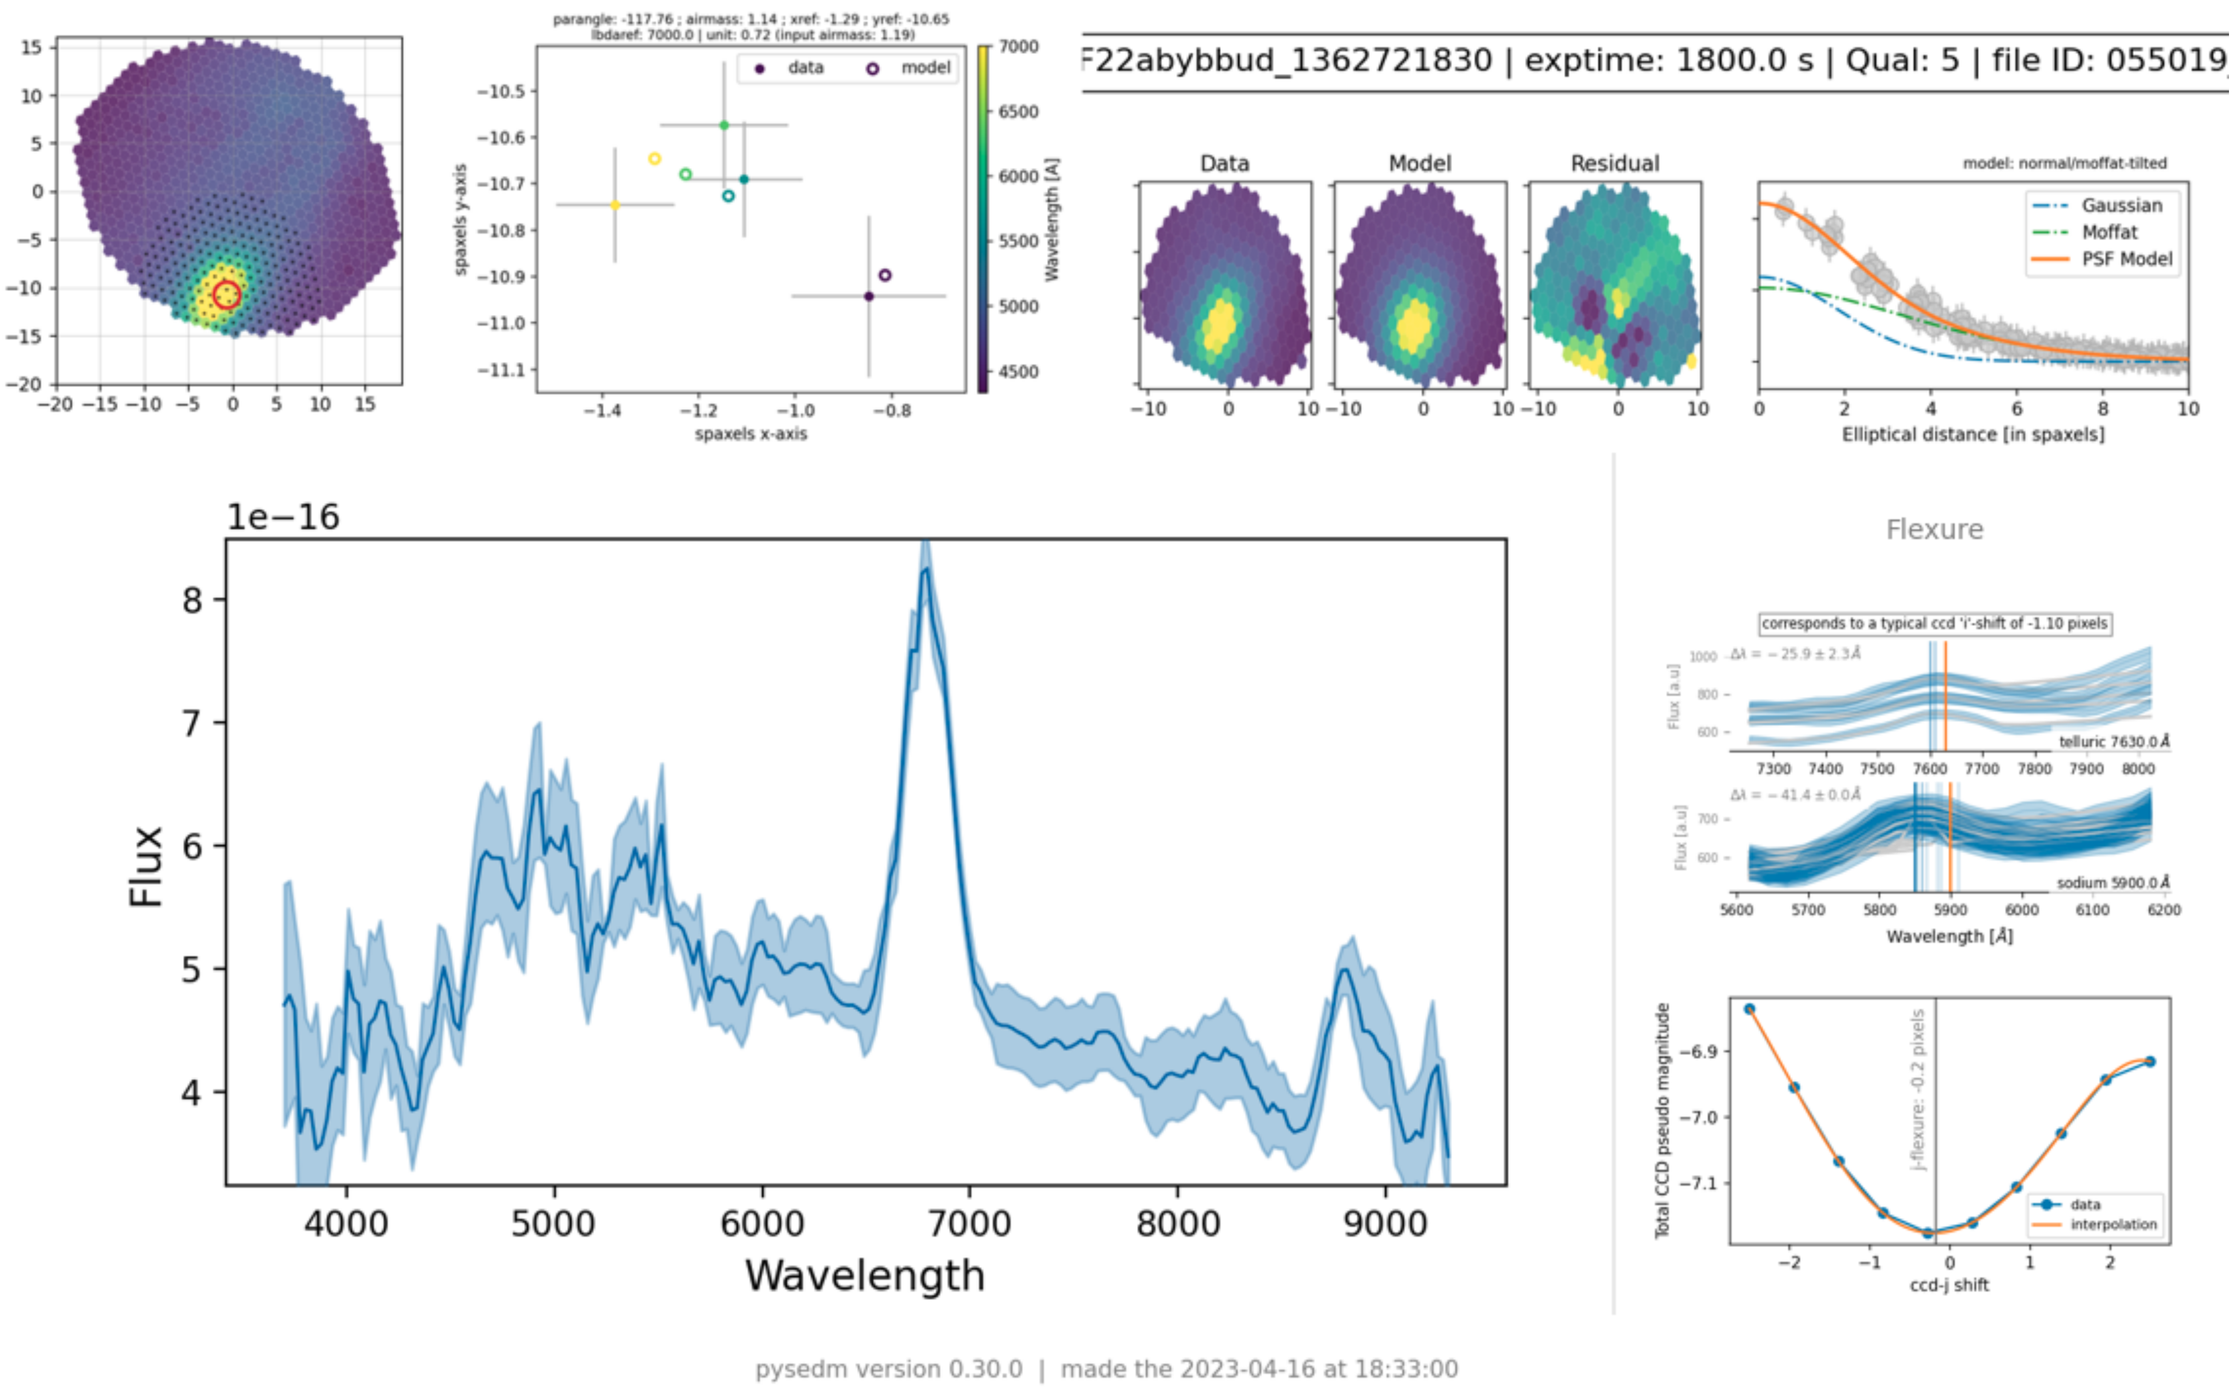

SED Machine - Kitt Peak: Timeline

June 5 - June 8
First commissioning trip
-Removed KPED
-Put SEDMv2 on with stop-gap prism

June 15
Contreras wildfire - KP evacuated

Sep 19 - Sep 30
Second commissioning trip(s)
-No damage to instrument
-Replaced stop-gap prism with tri-prism and realigned
-Put SEDMv2 back
-Operations paused until access to stable power supply

Oct 18
Line power back
Dec 25
Internet restored
BUT
KP84 UPS failure!

Nov 12 - April 2023
Multiple trips to KP84
-Robotic operations tested
-Taking commissioning data
-Onsite data analysis

To Do
-Fix dome drive
-Mirror recoating
-Fix minor software bugs as we go!

June 8
FIRST LIGHT!
System not yet fully robotic

Feb 7
-New UPS installed
-Improved KP84 drive performance
-Guiding implemented

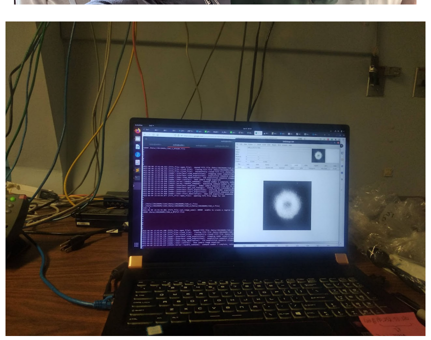

-Extensive damage to KP power poles, internet cables
-All scientific buildings saved
-Extensive damage to summit road because of landslides during monsoon

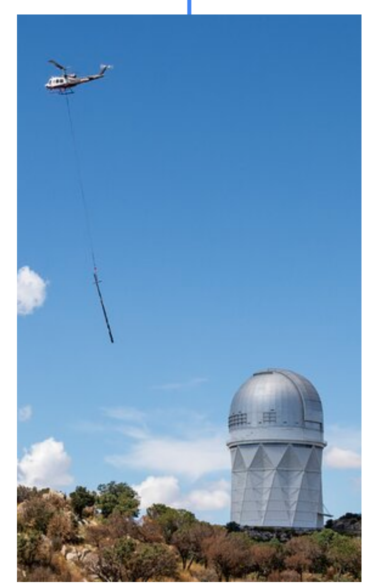

SED Machine Version - Kitt Peak

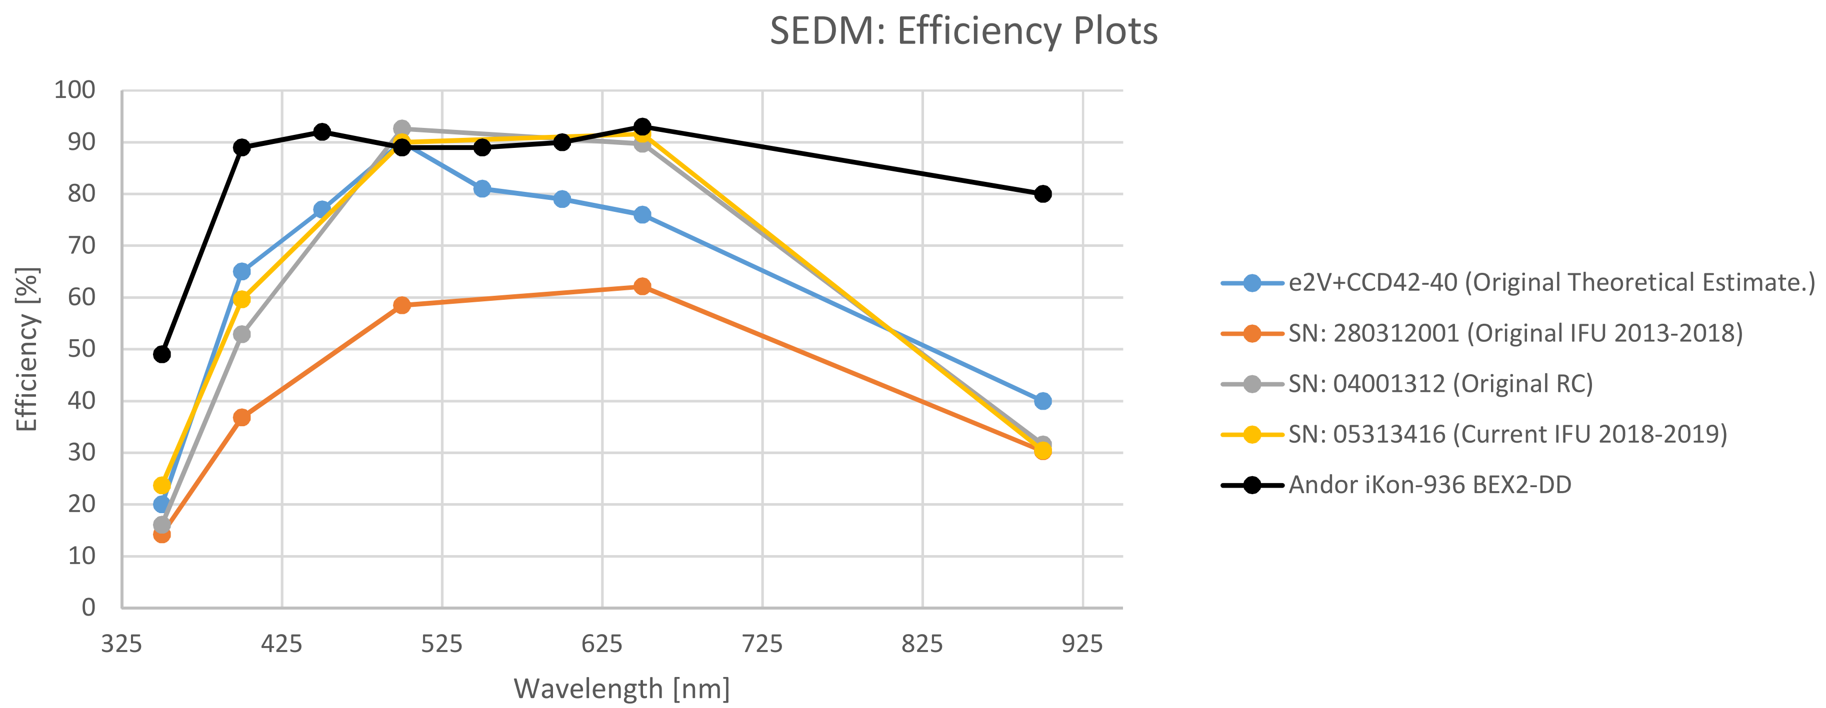

Kitt Peak 2.1m: Facility Specs

- 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 Version - Kitt Peak

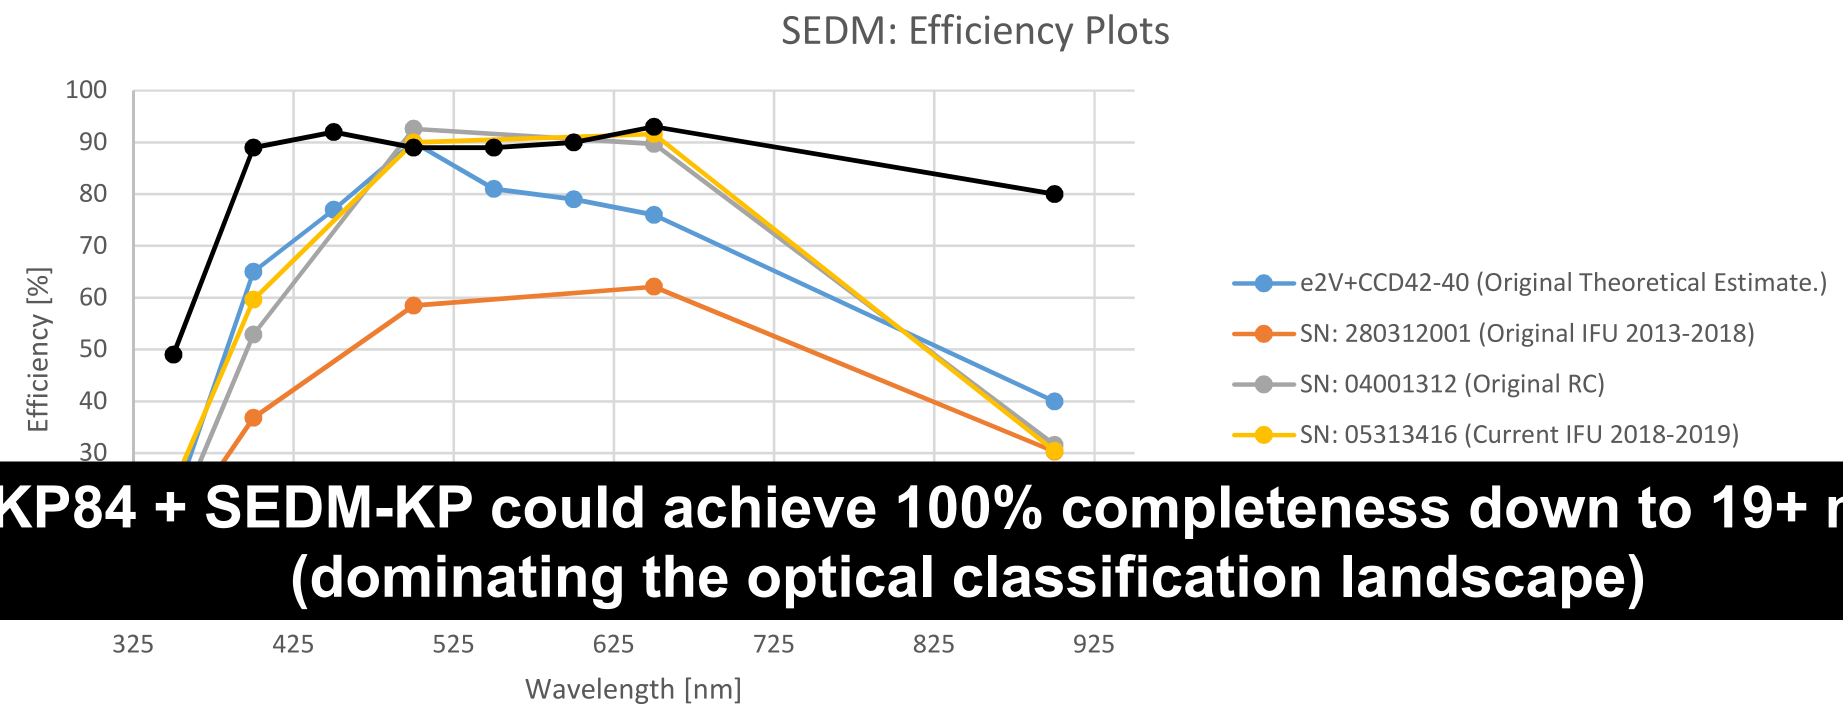

KP84 + SEDM-KP could achieve 100% completeness down to 19+ mag (dominating the optical classification landscape)

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

A Multi-Messenger Ecosystem

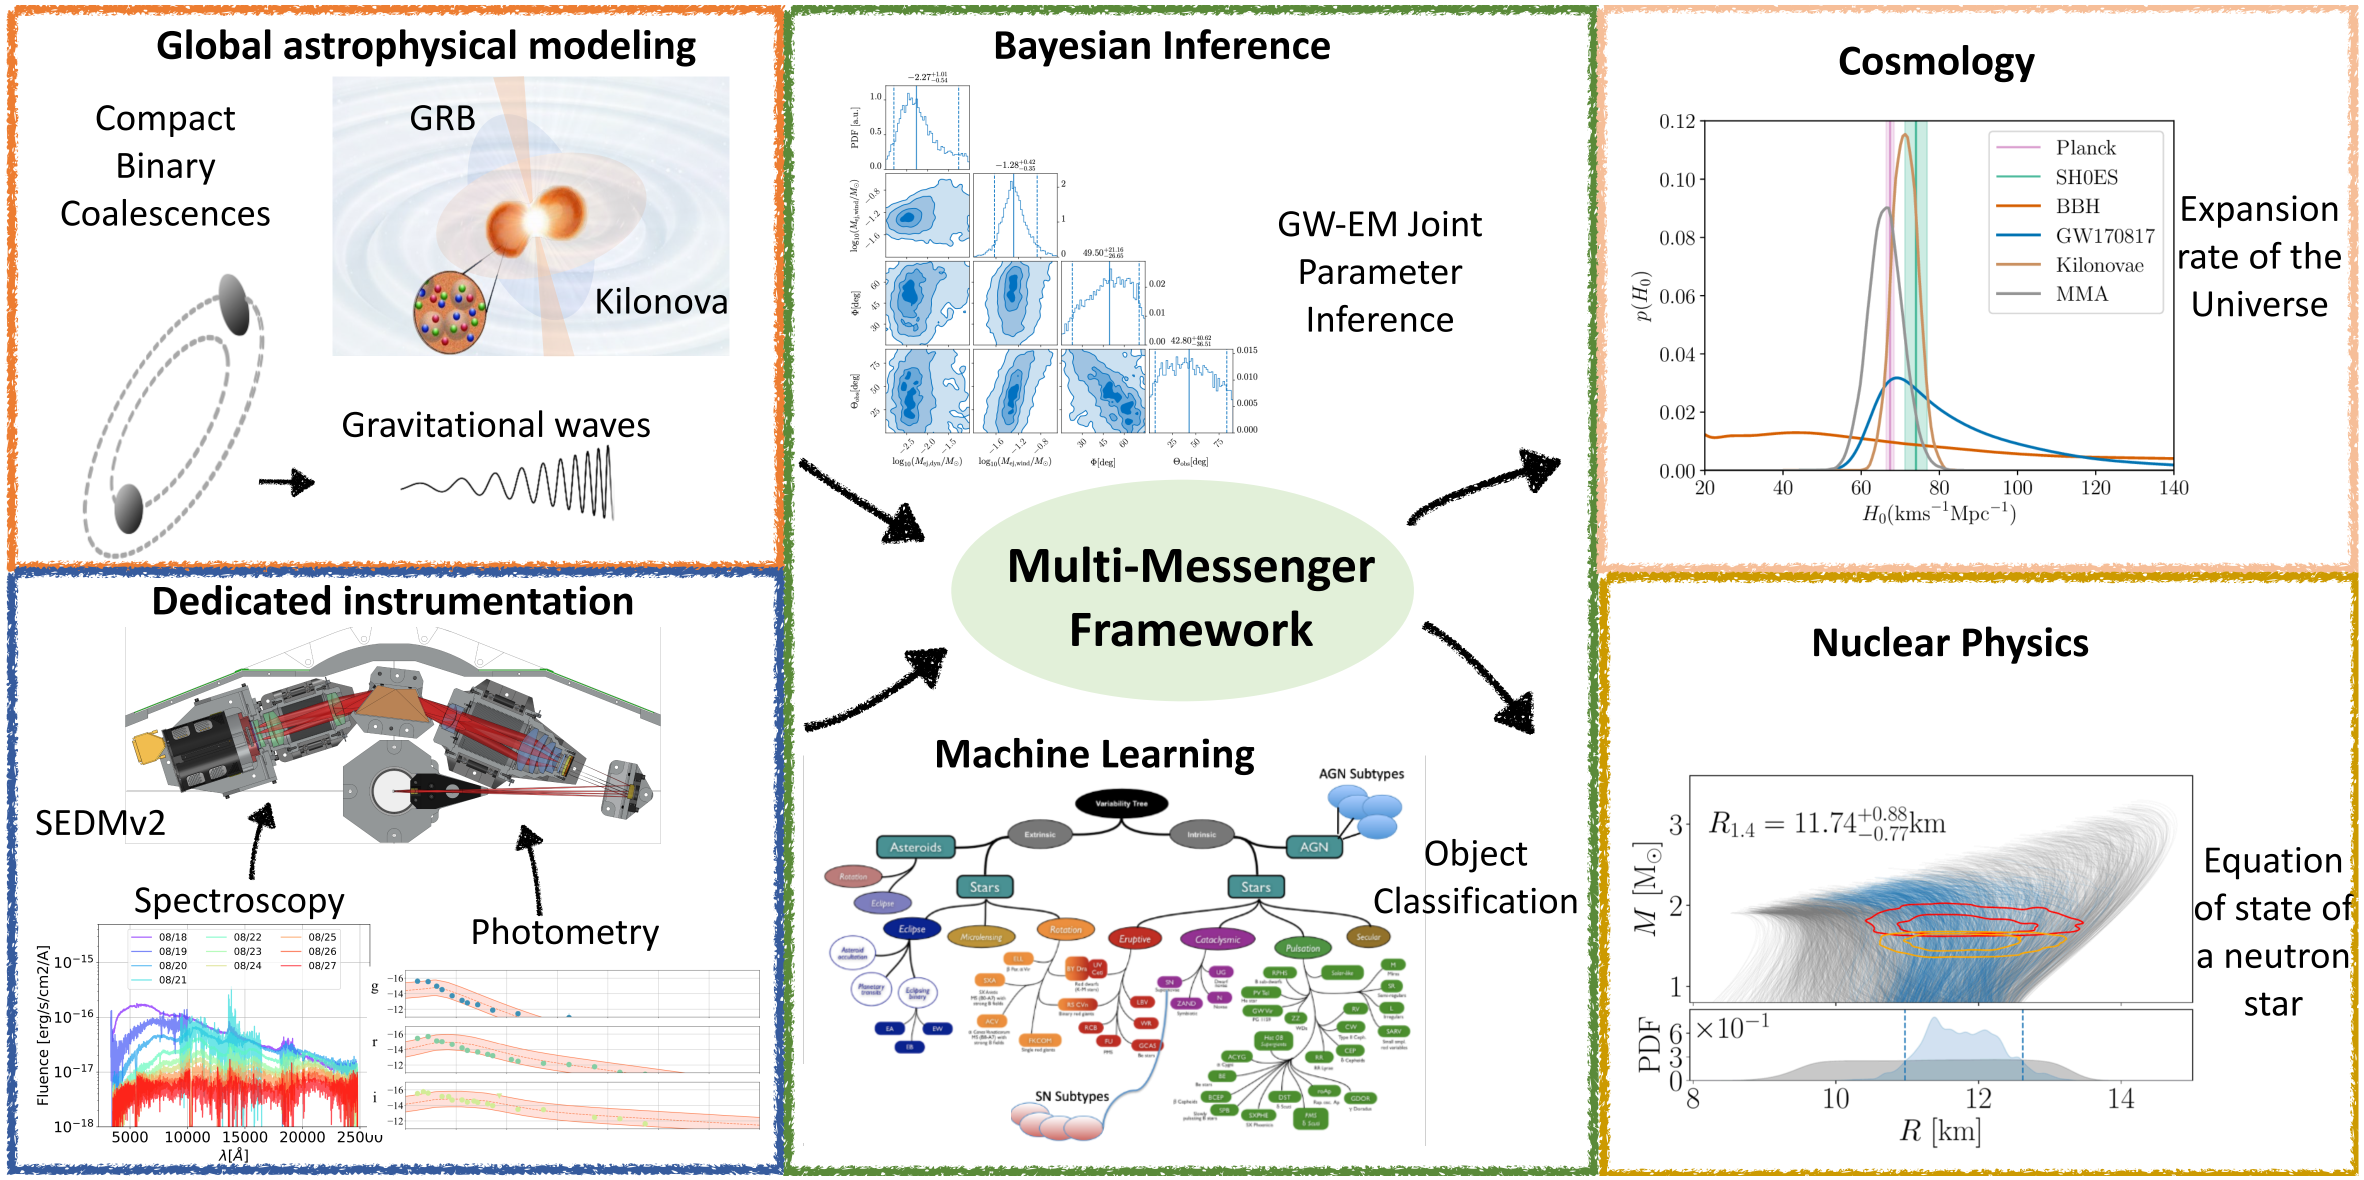

Thank you!