AUTONOMOUS REAL-TIME DECISION-MAKING IN THE ERA OF MULTI-MESSENGER ASTRONOMY

NIHARIKA "ARI" SRAVAN

DREXEL UNIVERSITY

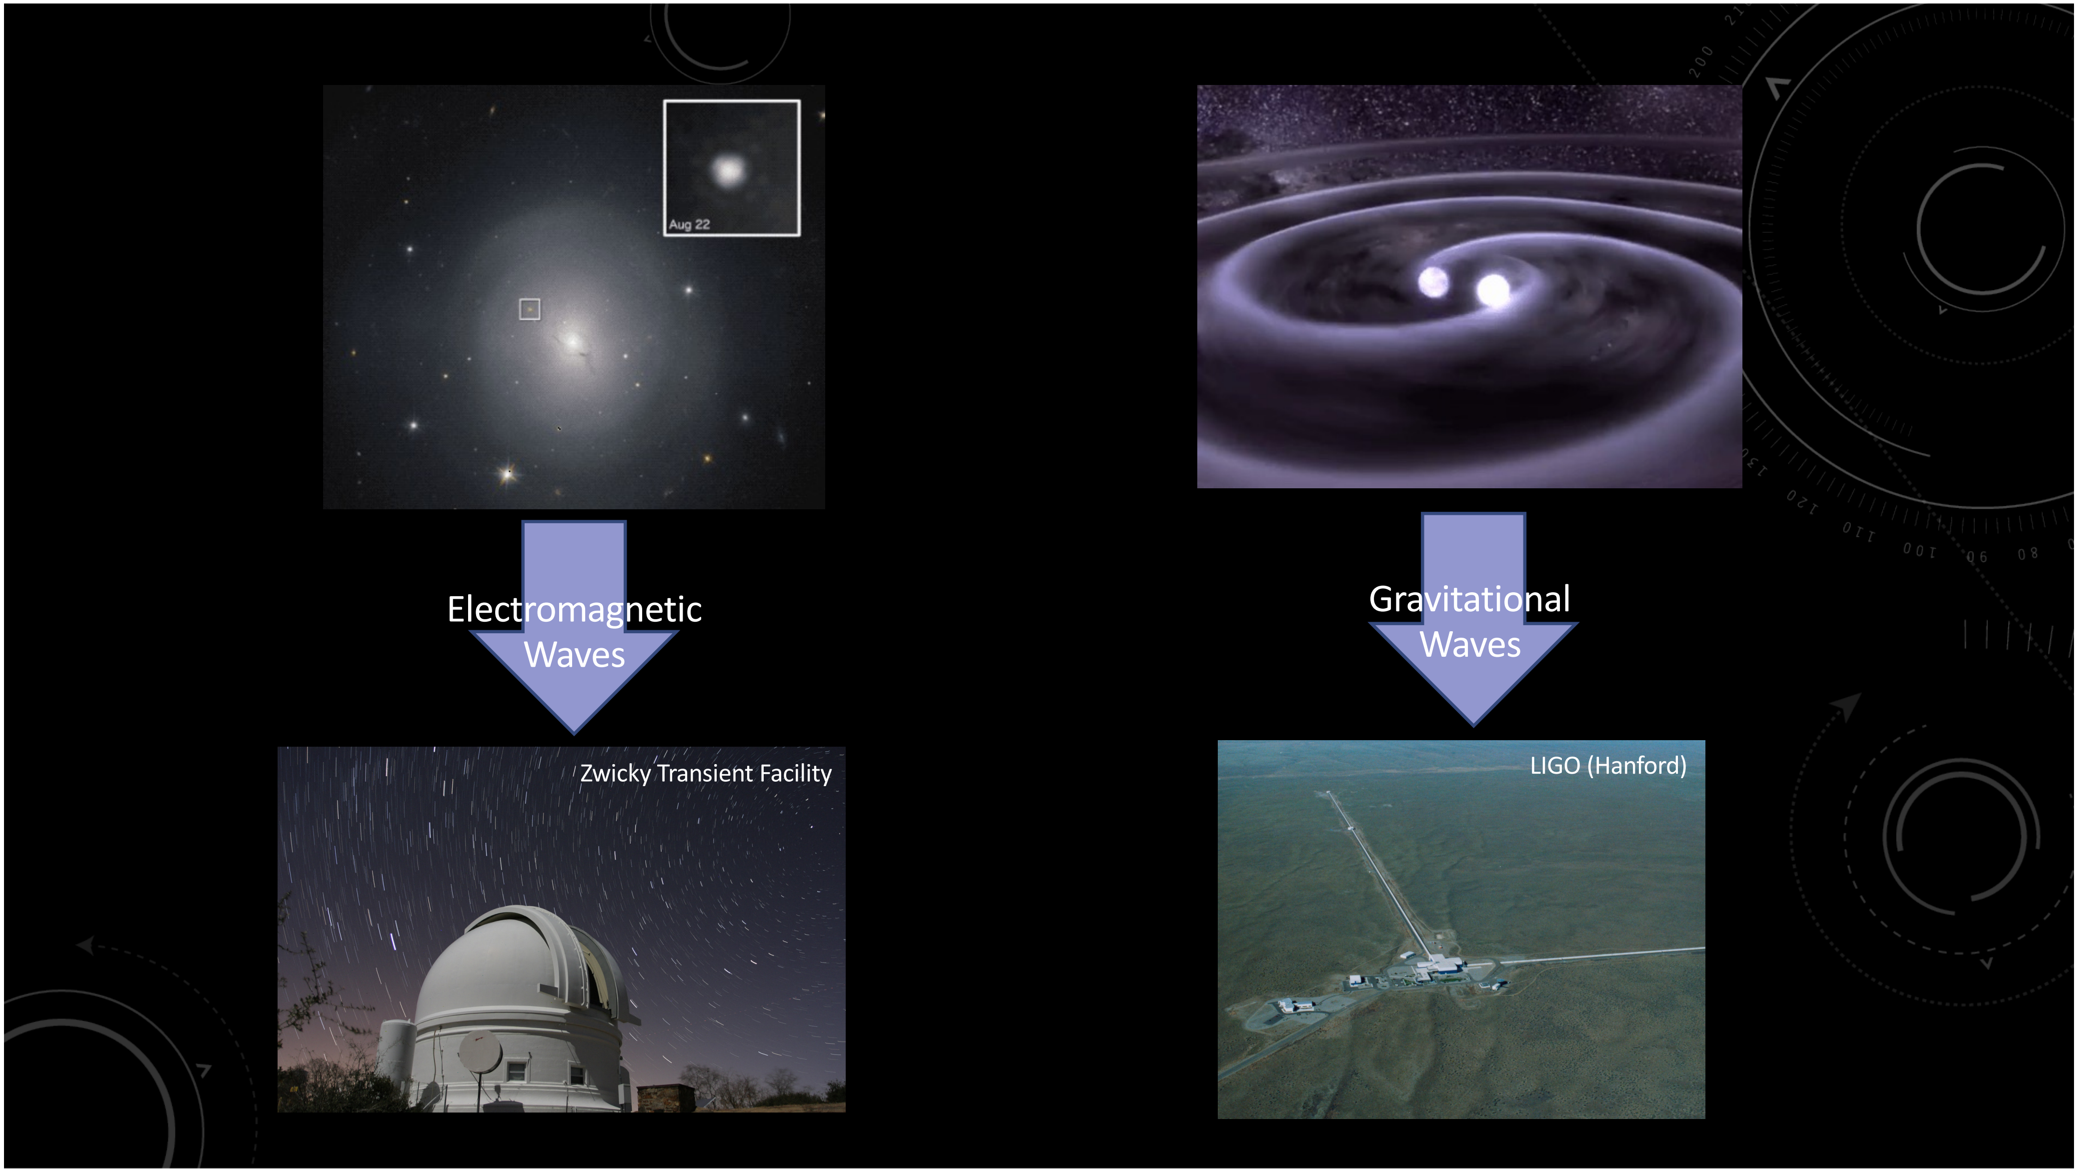
Aug 22
Electromagnetic Waves
Zwicky Transient Facility
Gravitational Waves
LIGO (Hanford)

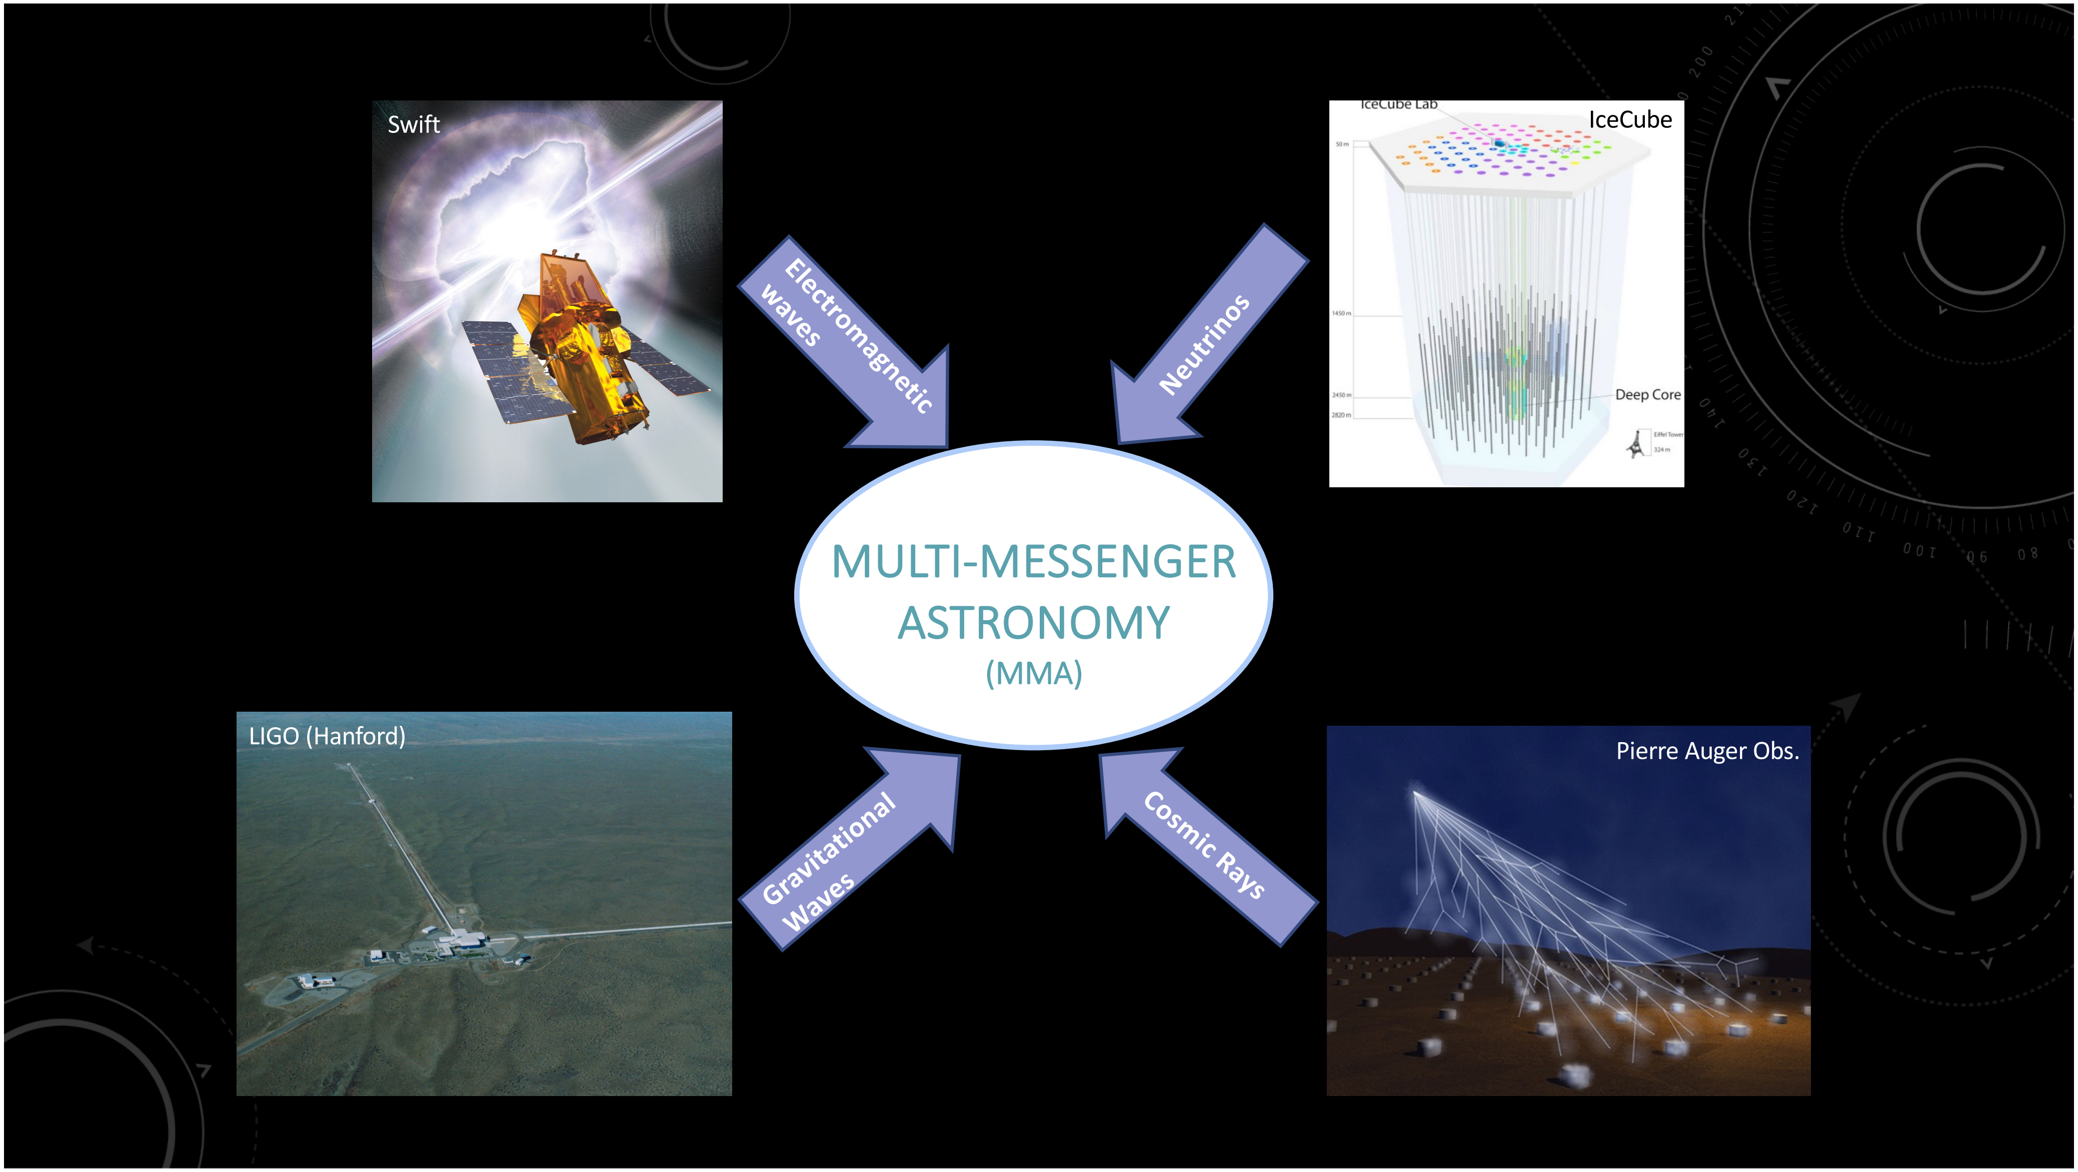
Swift
IceCube
Deep Core
Electromagnetic waves
Neutrinos
MULTI-MESSENGER ASTRONOMY
(MMA)
LIGO (Hanford)
Gravitational Waves
Cosmic Rays
Pierre Auger Obs.

Supernovae

Kilonovae

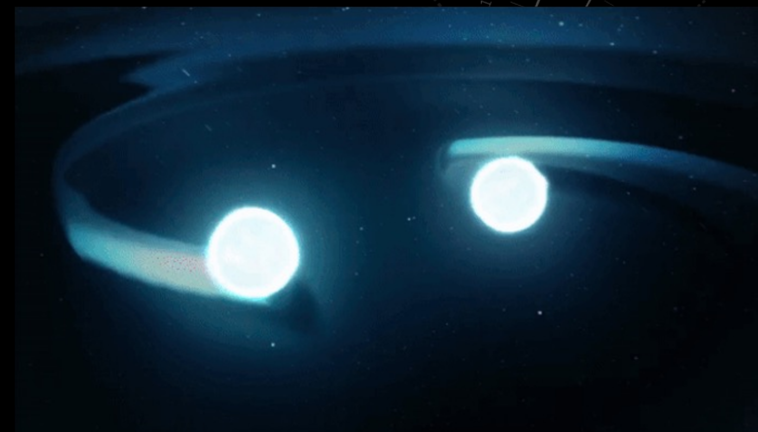

Active Galactic Nuclei

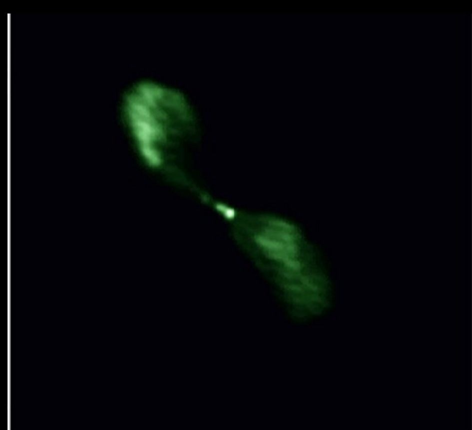

Tidal Disruption Events

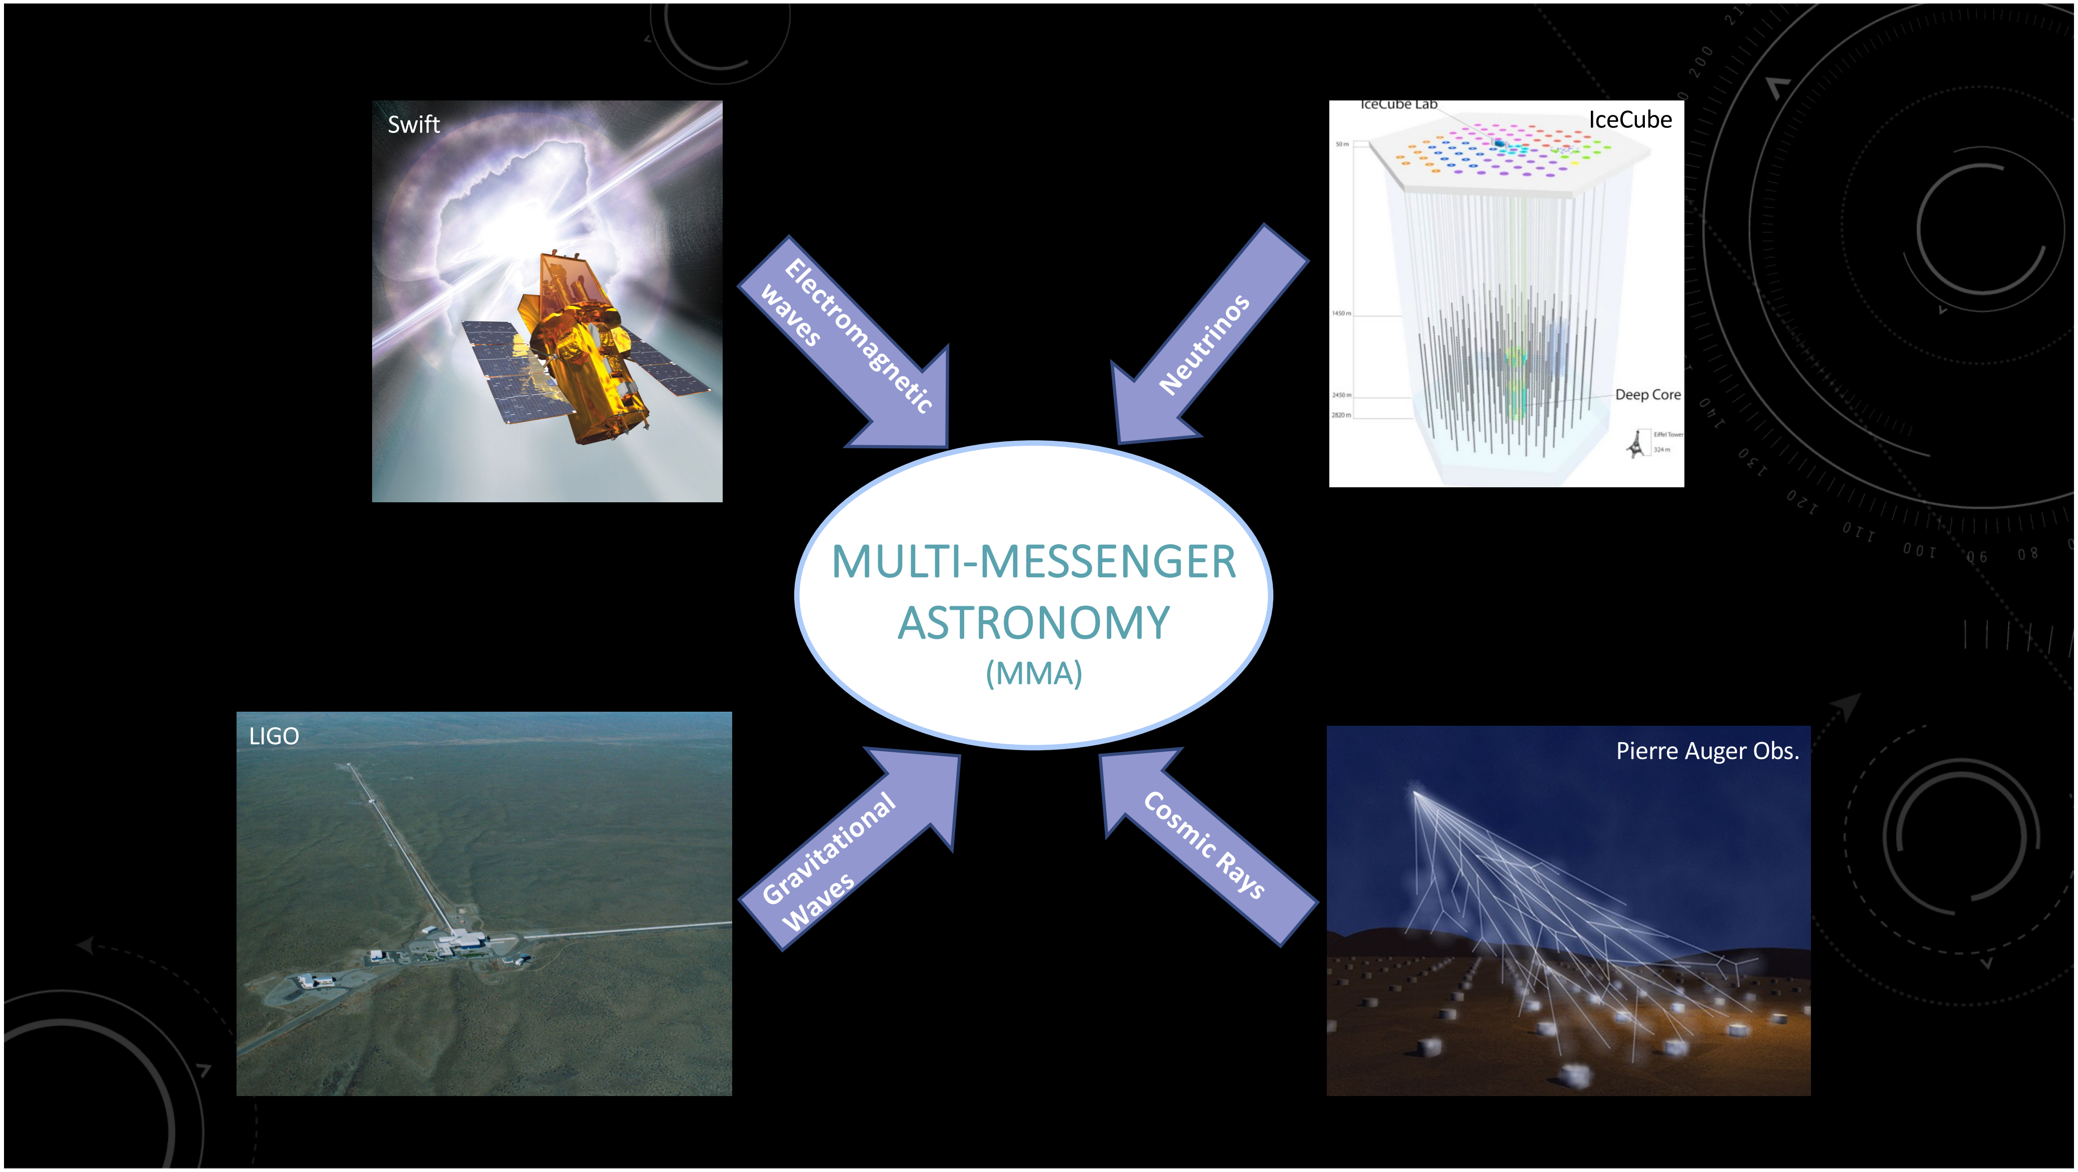

Swift
Electromagnetic waves
IceCube Lab
IceCube
Deep Core
Neutrinos
MULTI-MESSENGER ASTRONOMY
(MMA)
LIGO
Gravitational Waves
Cosmic Rays
Pierre Auger Obs.

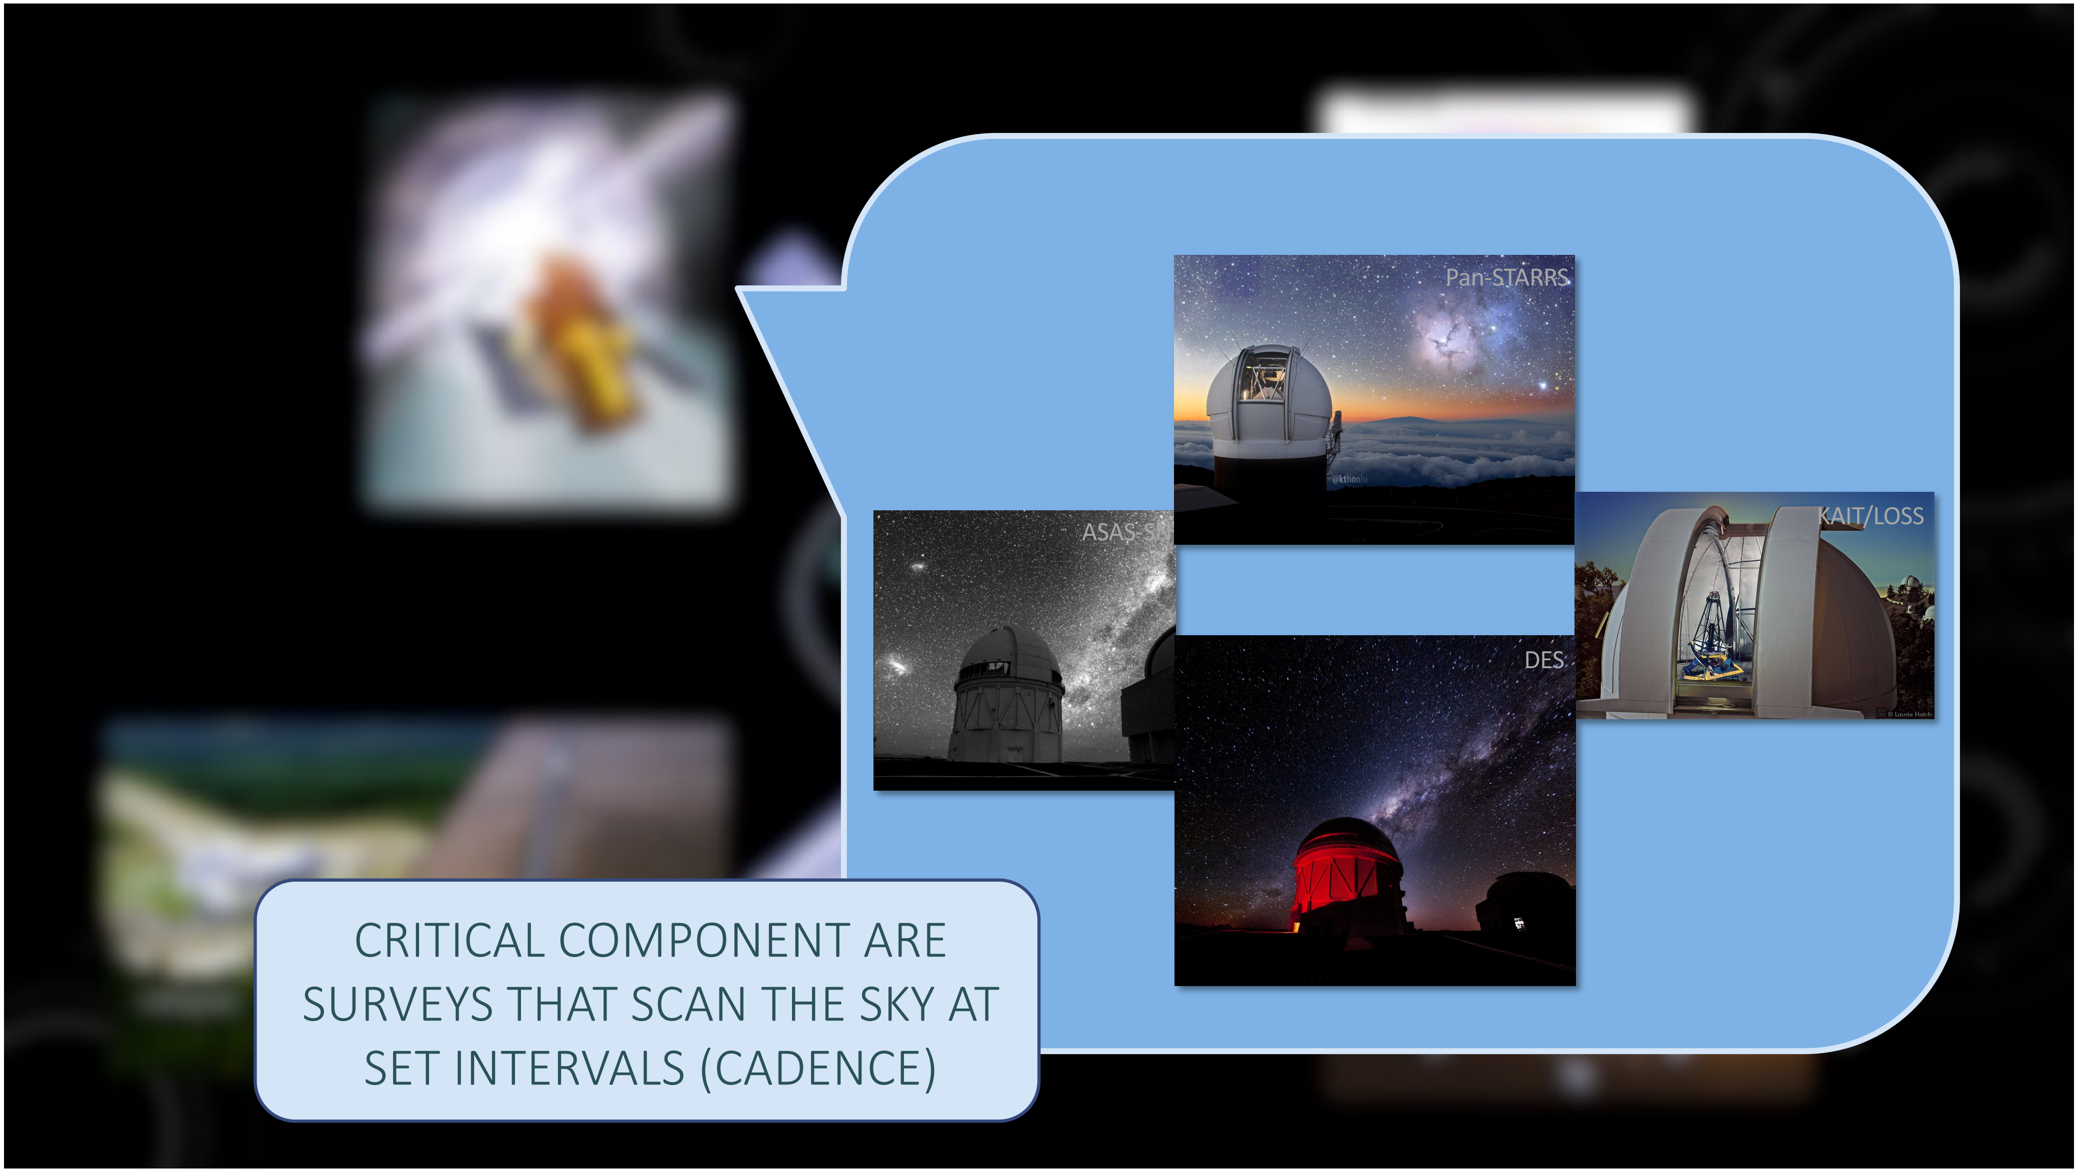
Pan-STARRS
ASAS-S
KAIT/LOSS
DES
CRITICAL COMPONENT ARE SURVEYS THAT SCAN THE SKY AT SET INTERVALS (CADENCE)

MULTI-COLOR LIGHT CURVES

(MULTI-CHANNEL TIME SERIES)

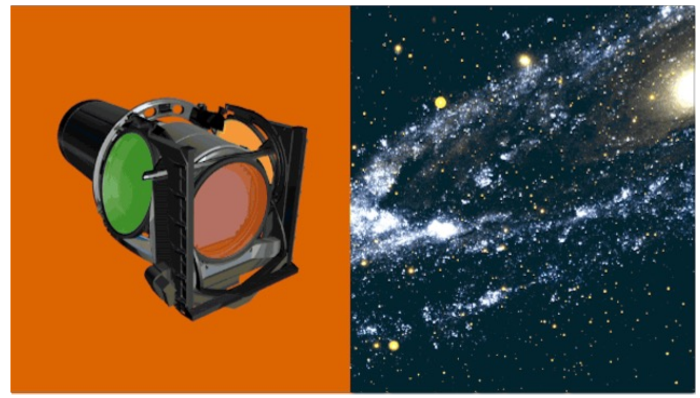

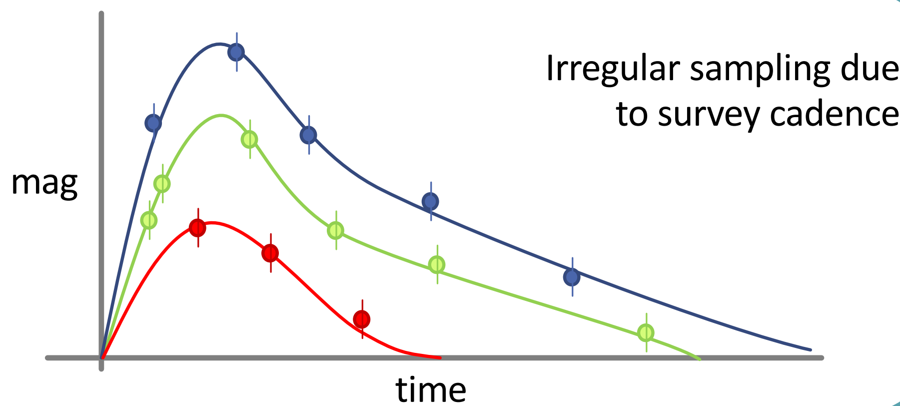

RUBIN OBSERVATORY

- #1 flagship of US community, recommended by National Research Council since 2010, jointly funded by NSF and DOE
- Ten year survey (LSST) starting 2026
- ~ 10 million alerts/night or 20TB per night
- Broadcast worldwide within 60s
- 37 billion light curves
  - 6 filters (u,g,r,i,z,y or 320–1050 nm, 3 day cadence)

25TB per day

RUBIN OBSERVATORY

- #1 flagship of US community, recommended by National Research Council since 2010, jointly funded by NSF and DOE
- Ten year survey (LSST) starting 2026
- ~ 10 million alerts/night or 20TB per night
- Broadcast worldwide within 60s
- 37 billion light curves
  - 6 filters (u,g,r,i,z,y or 320–1050 nm, 3 day cadence)

ZWICKY TRANSIENT FACILITY (ZTF)

- Pathfinder to Rubin especially for transients
  - Operating at 10% Rubin scale
- Alerts broadcast within ~20 mins
  - Only survey with public real-time alerts
- 2 billion light curves and counting…
  - public survey in 2 filters (g,r, 2 day cadence)

  private-partnership in i (+ other programs)

MULTI-MESSENGER SCIENCE

CURRENT PROTOCOLS AND LIMITATIONS

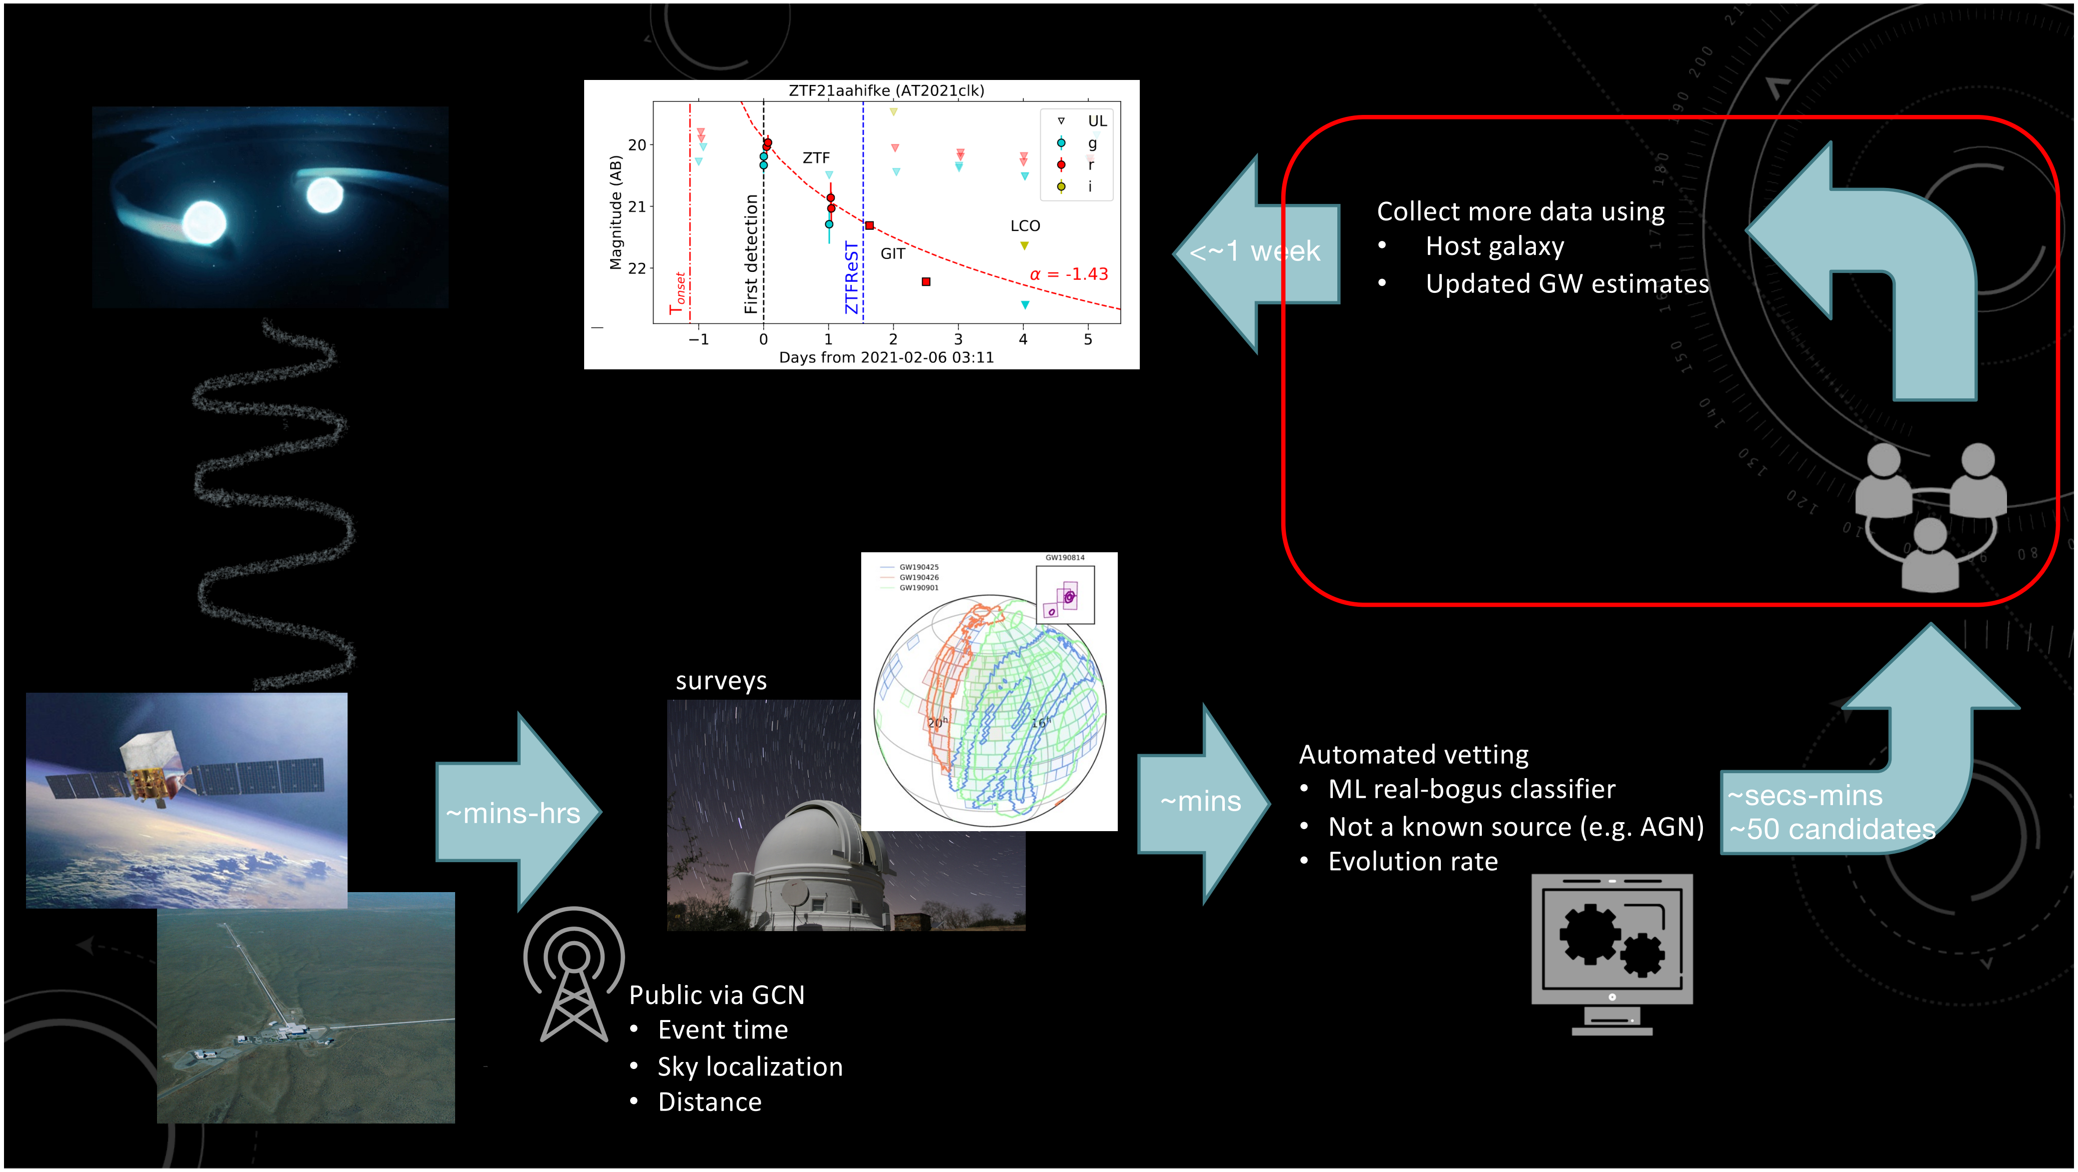

ZTF21aahifke (AT2021clk)
Magnitude (AB)
Days from 2021-02-06 03:11
First detection
ZTFReST
ZTF
GIT
LCO
UL
g
r
i
α = -1.43
<~1 week
Collect more data using
• Host galaxy
• Updated GW estimates
surveys
~mins-hrs
~mins
Automated vetting
• ML real-bogus classifier
• Not a known source (e.g. AGN)
• Evolution rate
~secs-mins
~50 candidates
Public via GCN
• Event time
• Sky localization
• Distance

~50 candidates

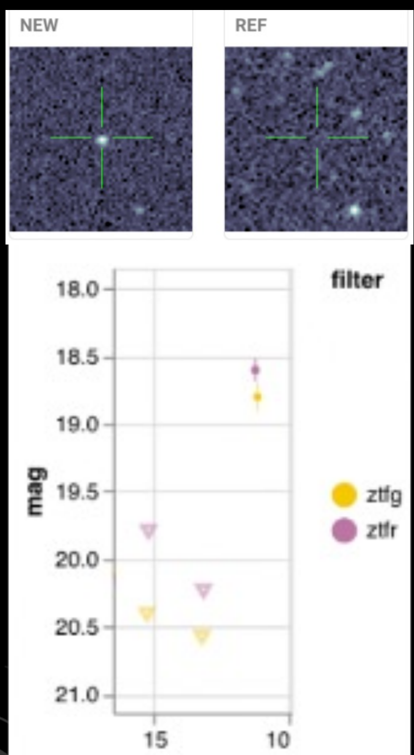

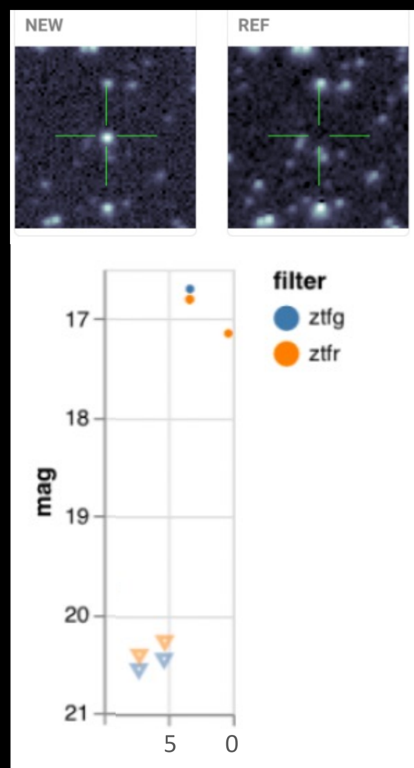

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
    - True event prob <0.05
  - Maximize constraints on interesting light curve physics

- Additional follow-up is critical!
- Classifiers don’t answer what to do next and how to adapt

Process needs to be:

✓ Free from fatigue/bias
✓ Low-latency
✓ Scalable

BY MID 2023

LVK will operate at twice the sensitivity

- 50-250 detections a year compared to 20 last time
- Localizations will not improve by much

BY LATE 2025

Rubin will come online and produce 10x as many candidates for human experts to analyze

Follow-up resources will not increase at nearly the same rate

Current protocol not sustainable or suitable to get at statistics

AUTONOMOUS REAL-TIME DECISION-MAKING

DECISION MAKING UNDER UNCERTAINTY

- Classically: Optimal experiment design
- Contemporary ML: Reinforcement learning, optimal sensing

Kirstine Smith (1878-1939)

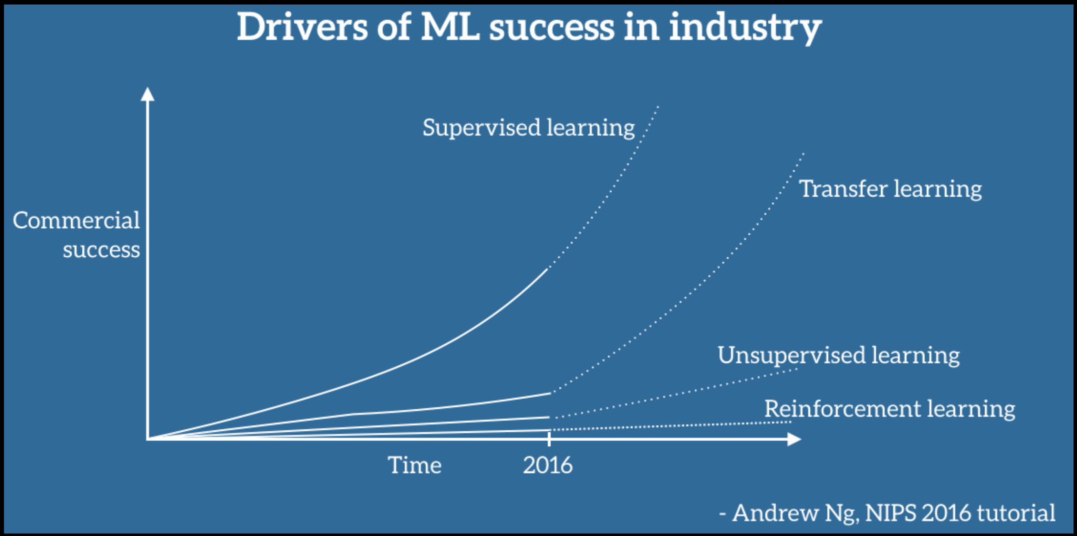

~50 candidates

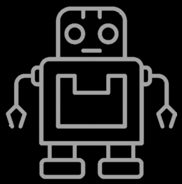

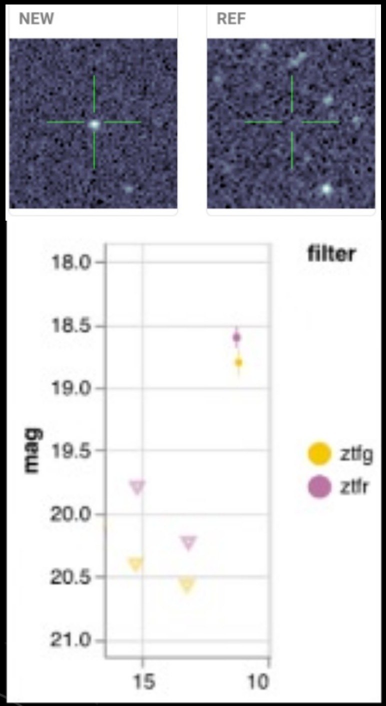

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

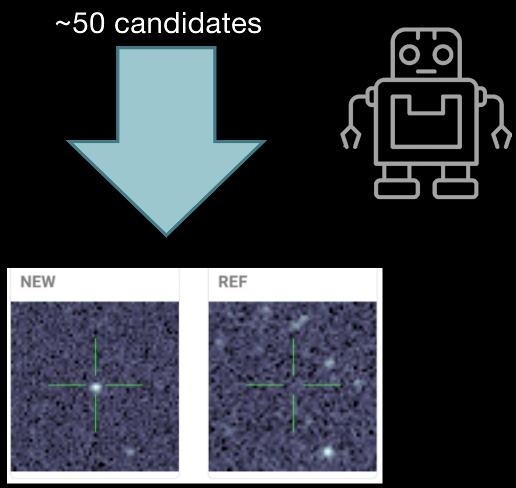

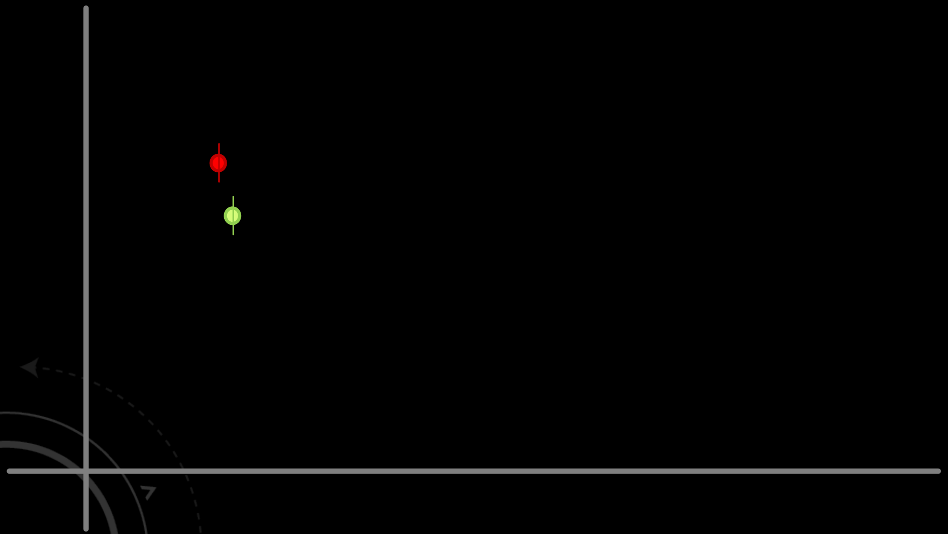

Goal:
- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

Step 1:

Define *state* space
- e.g. observed multi-channel light curves

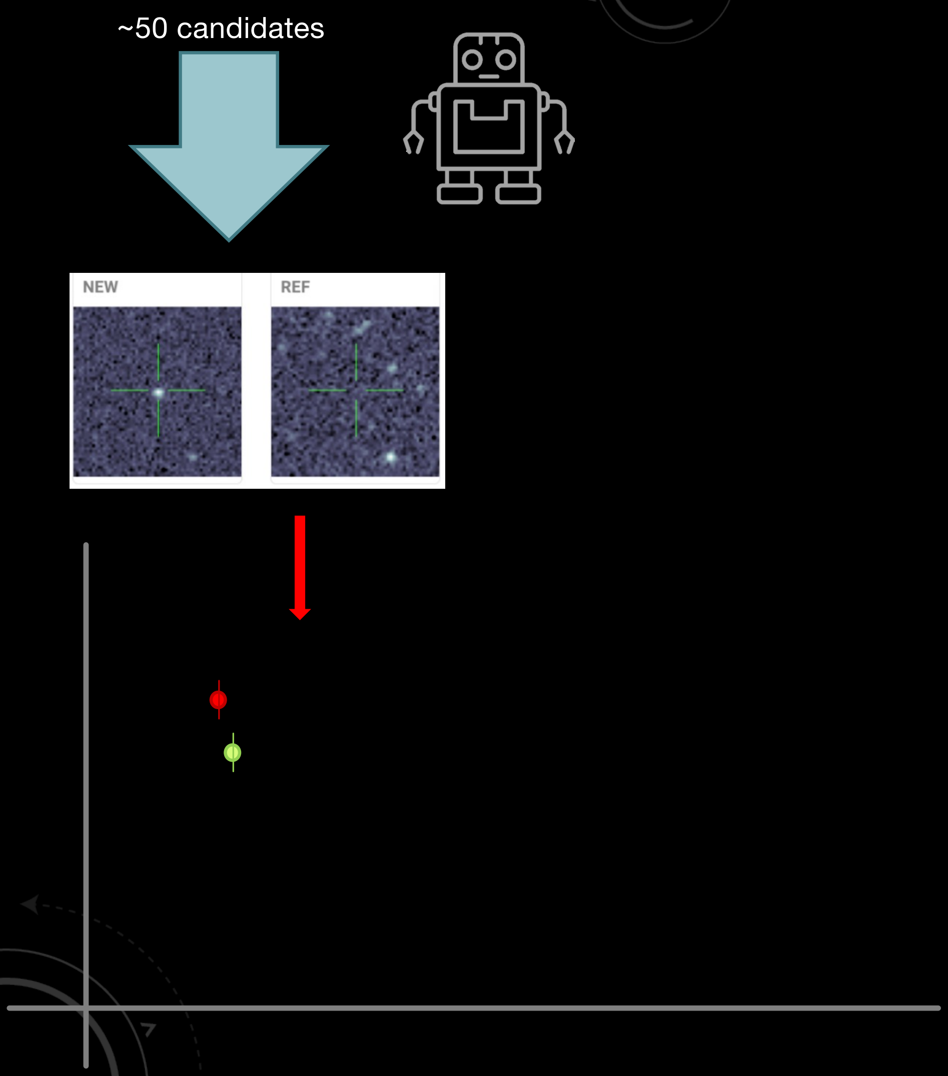

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

Step 2:

Define *action* space

- e.g. add data in {g, r, g+r, do nothing}

In general, actions can be:

Continuous

Stochastic

With variable cost and/or subject to budget

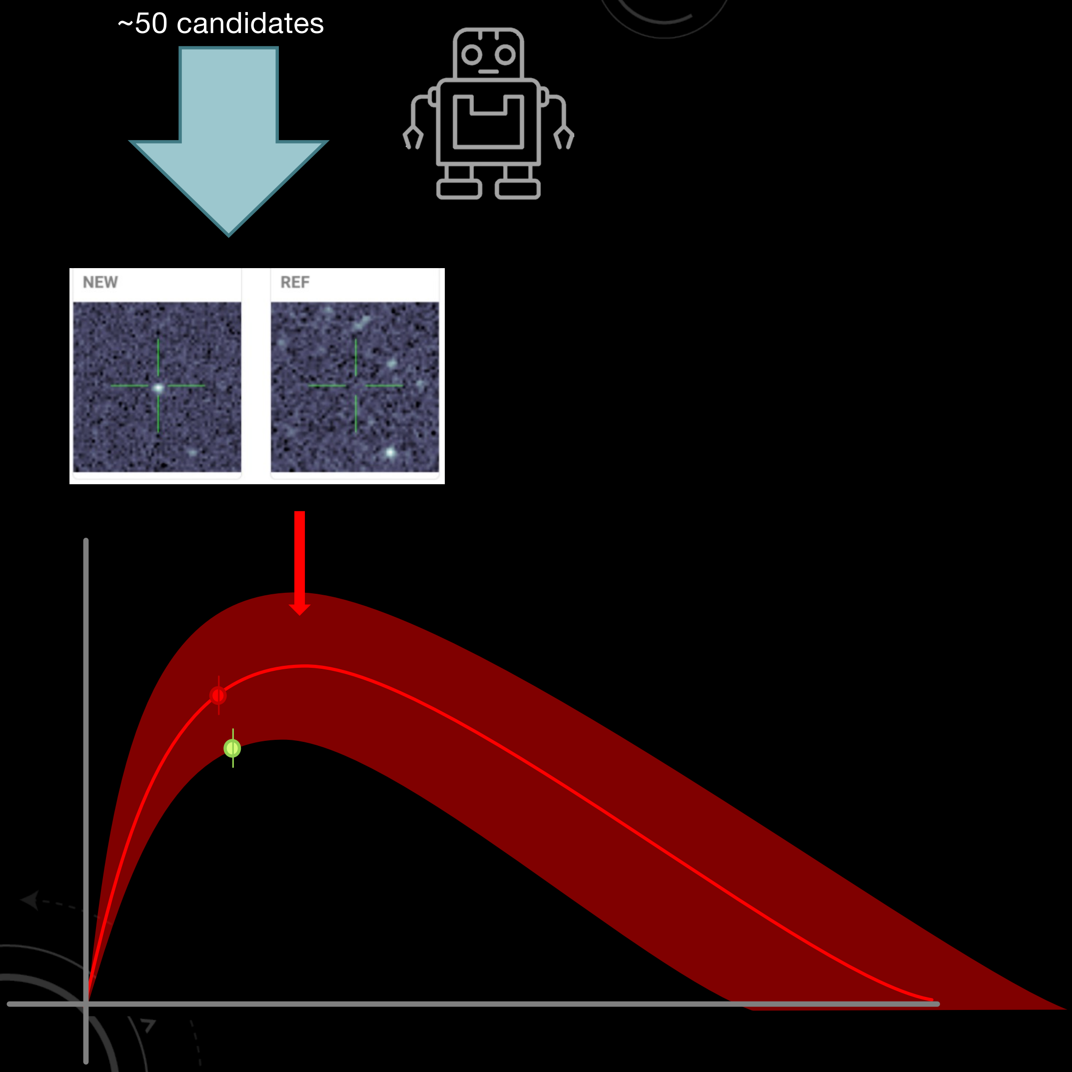

Goal:

- Use limited resources to acquire more information to:
    - Identify the event
    - Maximize constraints on interesting light curve physics

Step 3:

Estimate *outcome states* given actions

Dynamics/transition:

In general can be unknown and/or stochastic

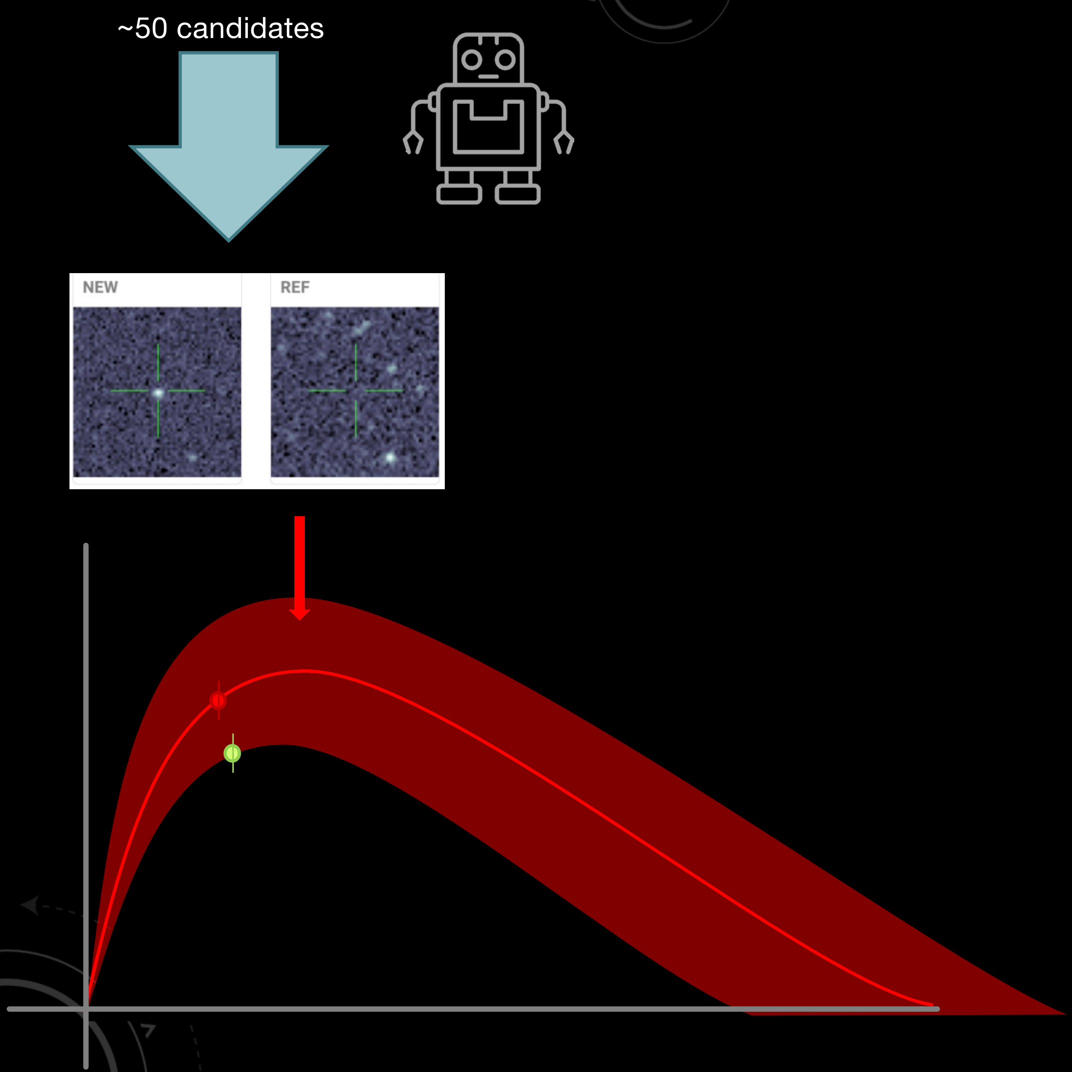

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

Step 4:

Estimate *utility* of outcome states:

- classification accuracy, TPR, F1-score, etc
- improvement in physics model parameters
- all of the above

In general can be stochastic

Your science!

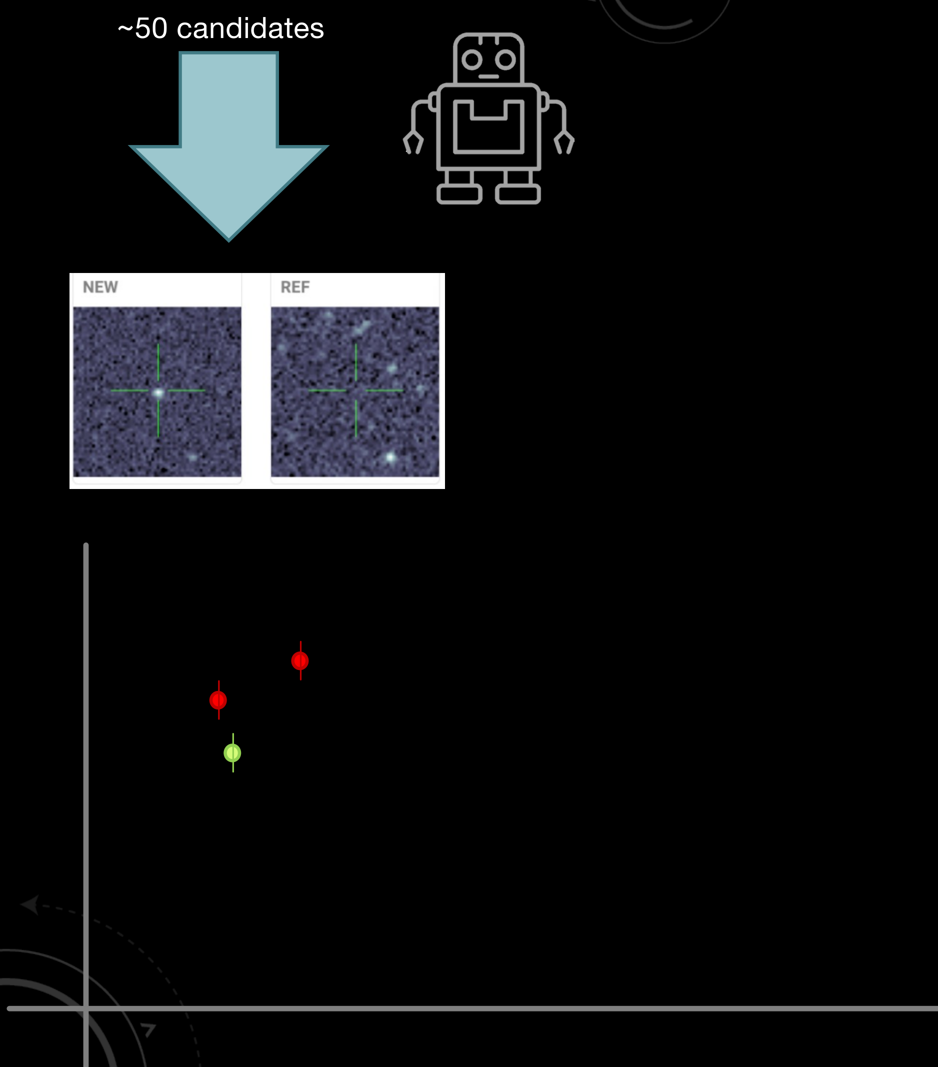

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

Step 5:

Take action according to *policy*

- e.g. greedy policy: take action with maximum reward; does not guarantee optimal series of actions
- Commonly $\arg\max_a Q(s,a)$

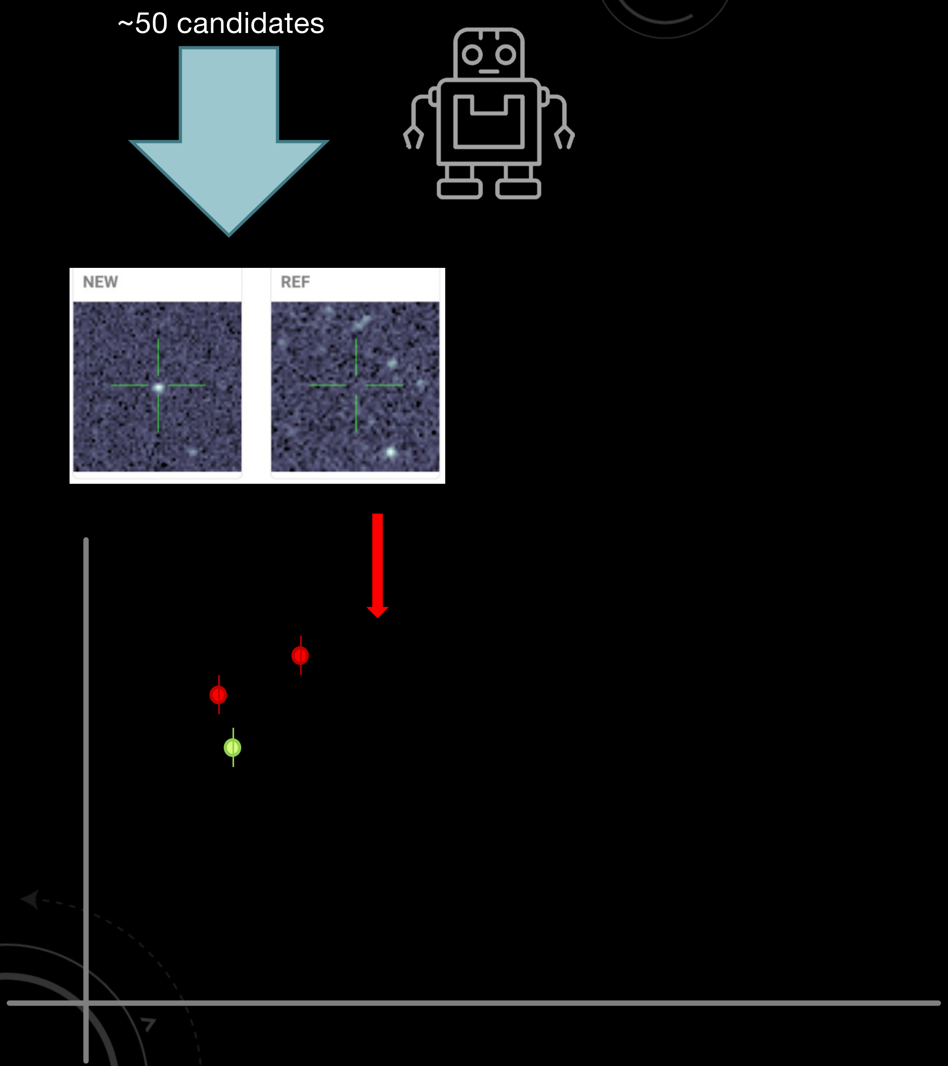

Goal:

- Use limited resources to acquire more information to:
  - Identify the event
  - Maximize constraints on interesting light curve physics

Step 6:

Adapt to new information
(inc. acquisition failure/latency, survey data)

Finish when episode ends or exhausted budget or repeat forever

EXAMPLE 1:

MEASURING EXPANSION OF THE UNIVERSE

Type Ia Supernovae

Constrain cosmological parameters:
Hubble constant, Dark Energy equation of state, …

Low redshift samples constrain local large-scale structure properties:
growth rate, velocity flows

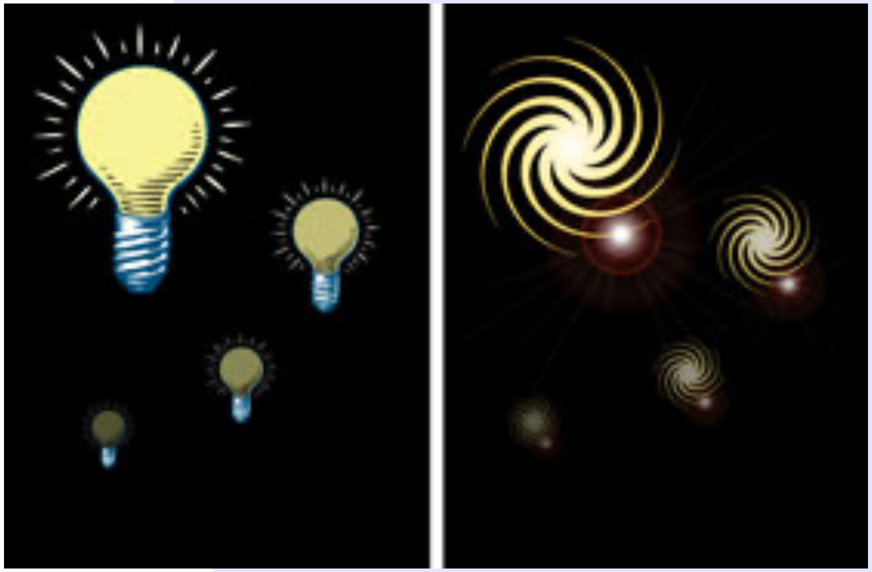

Real-time SN Ia LC augmentation to maximize cosmology

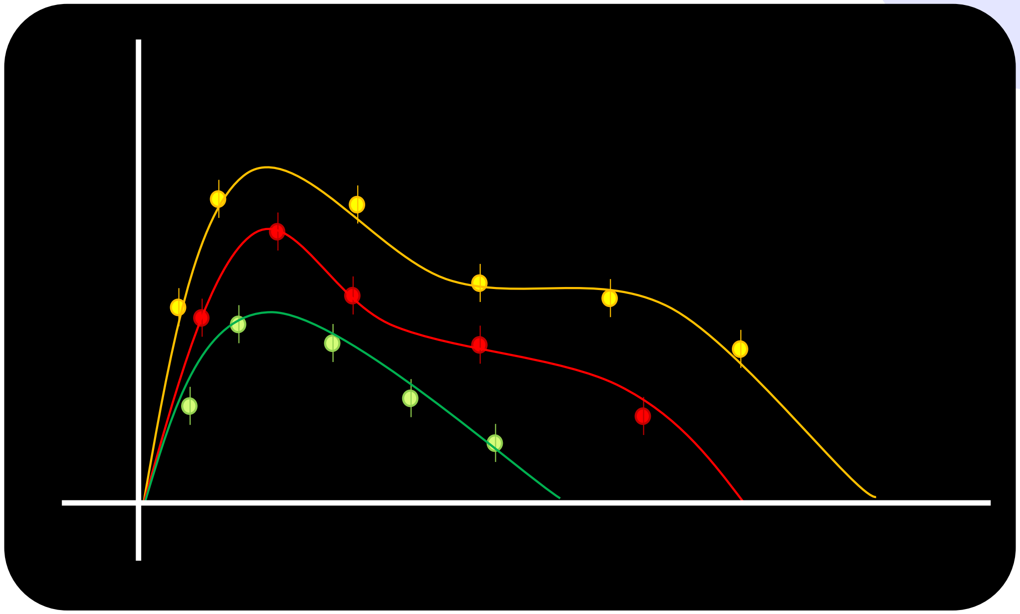

x_0, x_1, c are light curve fit parameters (z is spectroscopic)

Cosmology is $f(x_0, x_1, c, z)$ for a sample of SNe

Real-time SN Ia LC augmentation to maximize cosmology

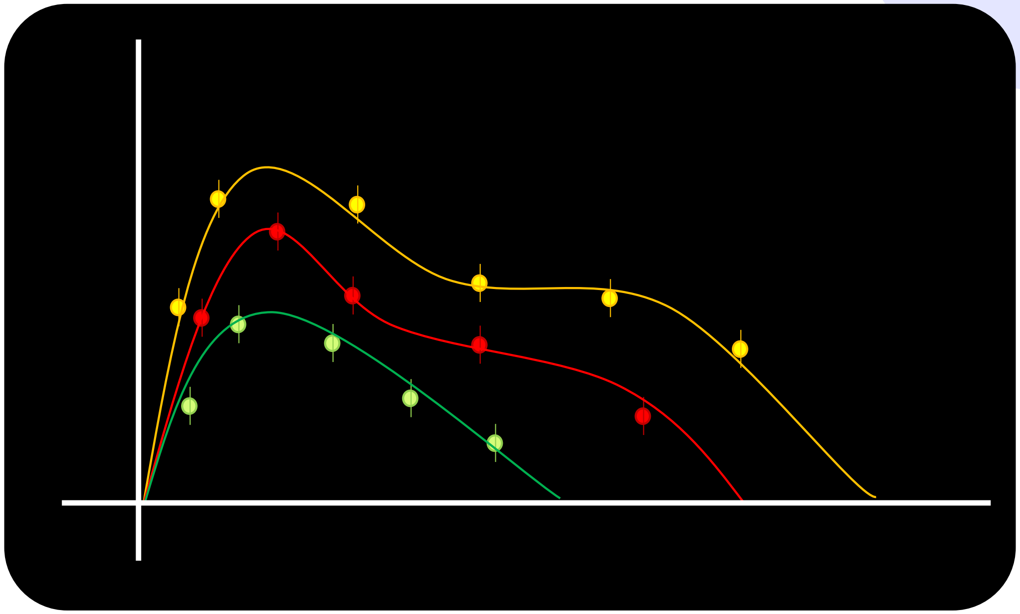

Minimize uncertainty on light curve fit parameters (in quadrature)

Minimize uncertainty on cosmology

Problem statement:

Augment photometry to branch-normal SN Ia light curves from ZTF-I public survey (g and r) in g,r, and i to minimize net uncertainty on SALT2 parameters

Problem statement:

Augment photometry to branch-normal SN Ia light curves from ZTF-I public survey (g and r) in g,r, and i to minimize net uncertainty on SALT2 parameters

i-band important for precisely estimating H_0 (Burns+ 2018)

Second peak could help probe SN Ia explosion mechanisms (Folatelli+ 2010)

Data in UV or IR can help better calibrate models (Milne+ 2015)

Algorithm

Non-stationary MDP with finite horizon (60 day episodes)

- State space: Observed photometry and expected data from survey (stochastic, 10x monte carlo). Remaining budget allocated randomly*
- Action space: {no action, g, r, i, gr, ri, ig, gri}
- Deterministic reward model: SALT2 + photo z A-optimality**

Deterministic dynamics model

- Upper limit using 2-D Gaussian Process regression over full LC
- For real time estimated with encoder-decoder LSTM trained on 10^5 simulated ZTF SNe Ia (slightly lower performance)
- New state simulated using 2-D Gaussian Process fit to full LC and fed back the next day

Deterministic on-policy, fixed budget and unit cost

- Take modal action with maximum reward and least cost across all rollouts. Threshold ϵ over no action (hyperparameter)

* substitutes expected **optimal** actions for expected **naïve** actions

** Distance error in quadrature. Note: min $\chi 2 \neq$ max liklihood

Gap filling

Resolves phase with high variability (first and second peaks+valleys)

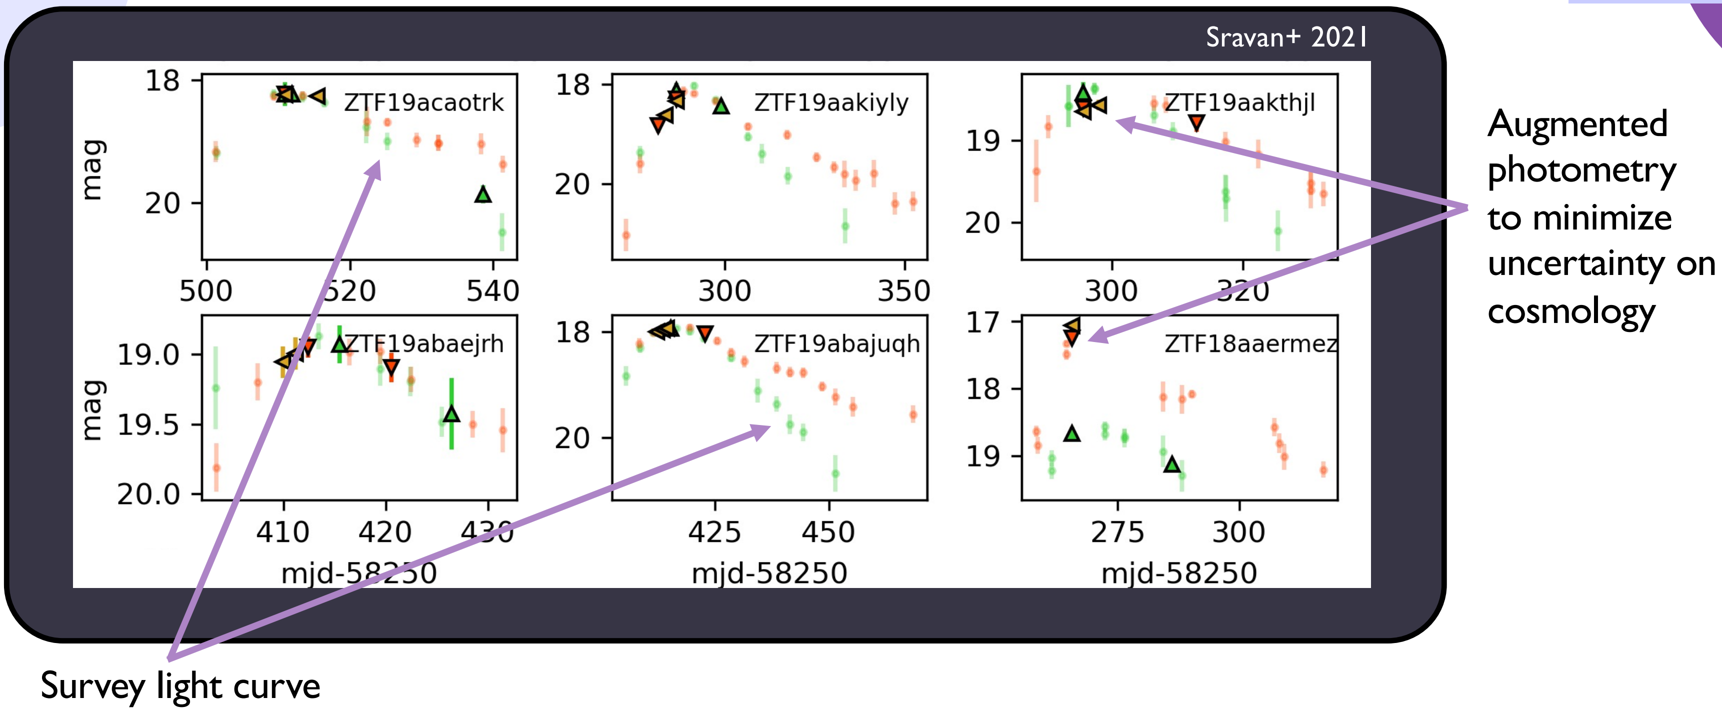

Improvement in SALT2 parameters over naïve** strategy

2-5% more improvement for faint SNe Ia (peak>18.5 mag)

Due to gap filling, strong prospects for Rubin

Budget	Usage	$\delta(\sigma_{x0})$	$\delta(\sigma_{x1})$	$\delta(\sigma_c)$	$\delta(\sigma_z)$
3	2	2%	3%	5%	6%
6	5	3%	5%	4%	6%
9	7	5%	6%	5%	11%

Sravan+ 2021

**Adding data itself can lead to improvement!

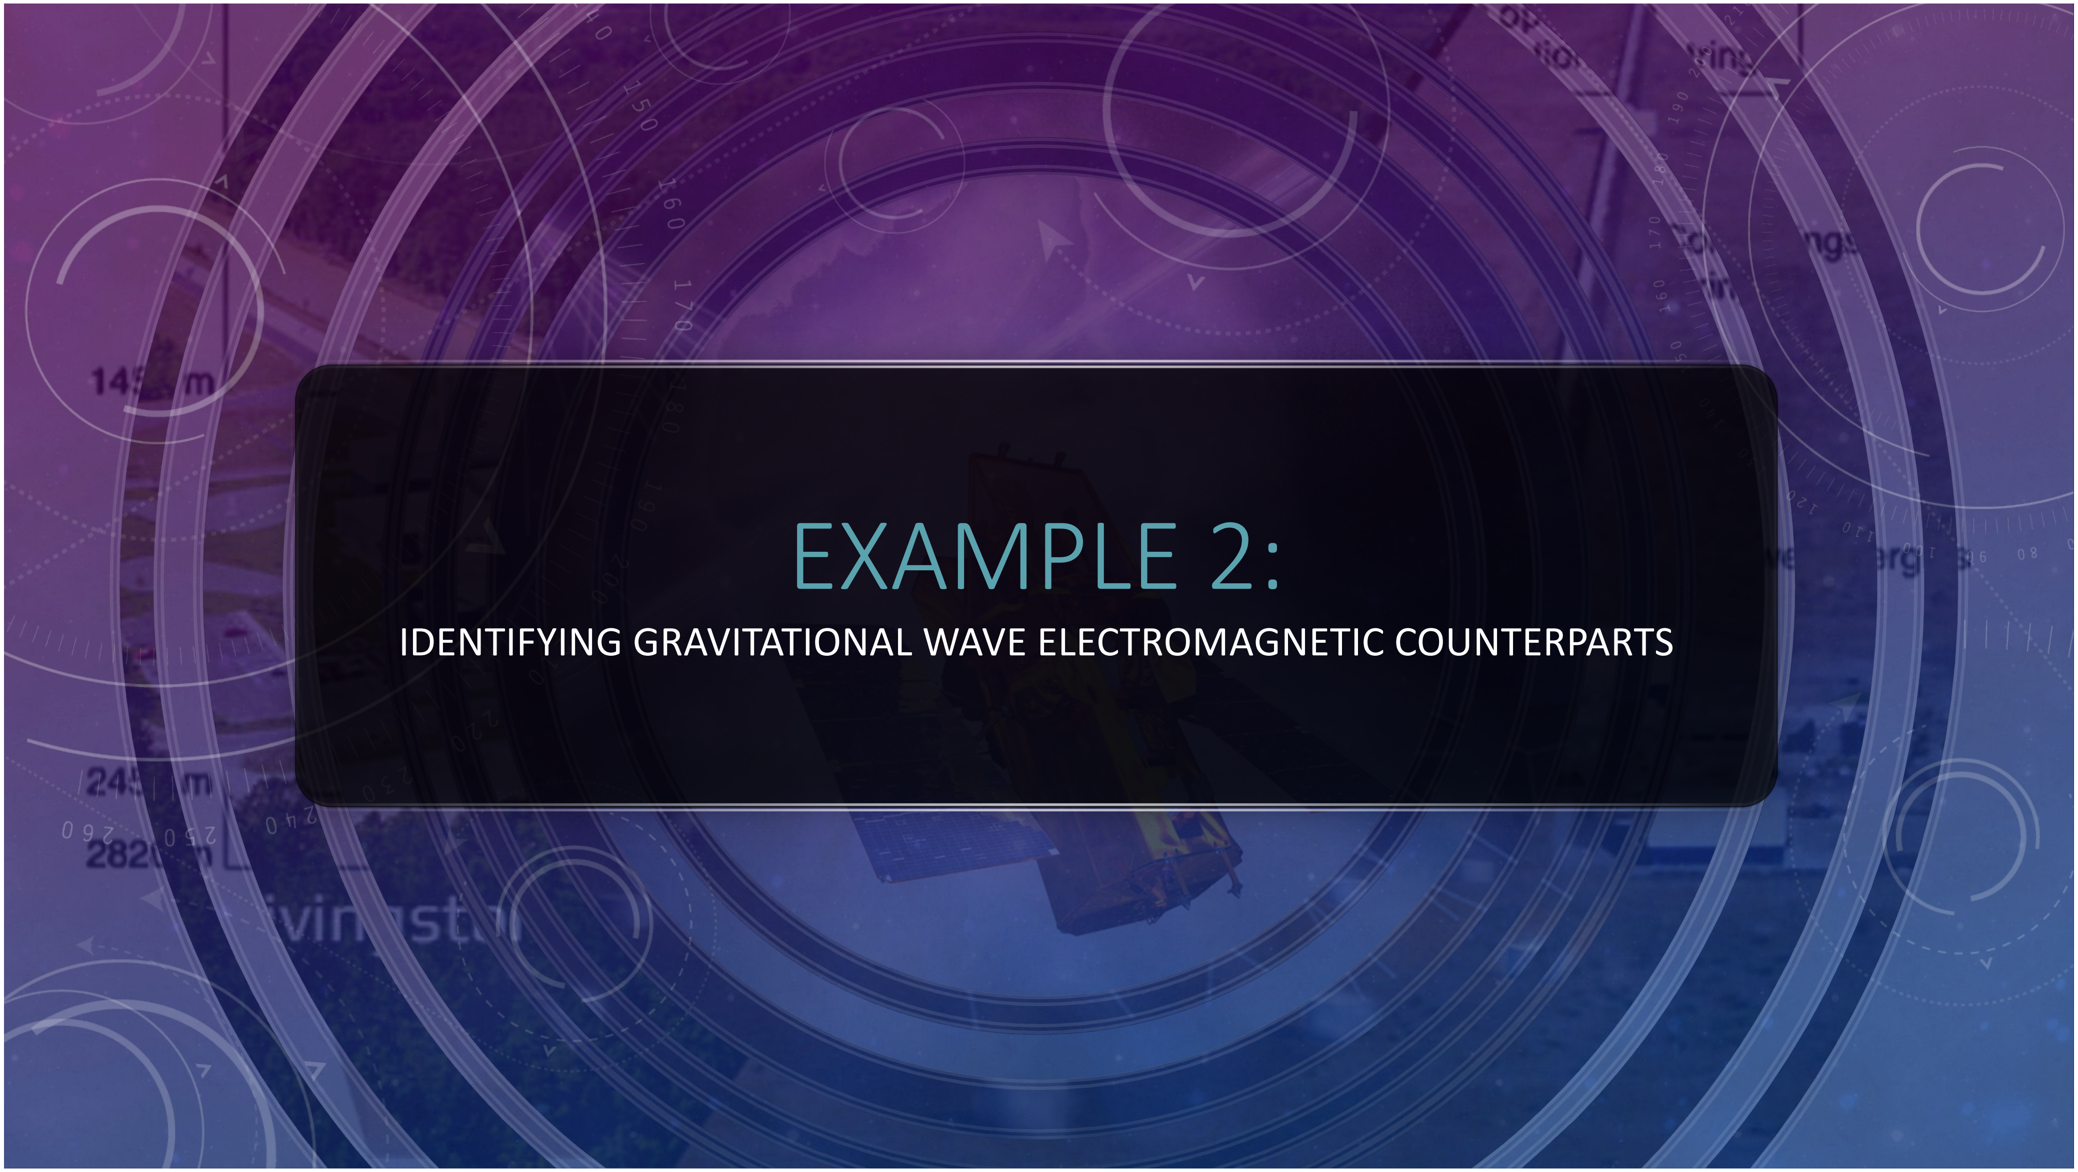
EXAMPLE 2:
IDENTIFYING GRAVITATIONAL WAVE ELECTROMAGNETIC COUNTERPARTS

Kilonovae

- UVOIR transients
- Probe nucleosynthesis in ejecta due to merger and associated power sources and the NS EoS
- Robust counterparts to most BNS and some NSBH mergers
- Short lived (<~1 week) and faint

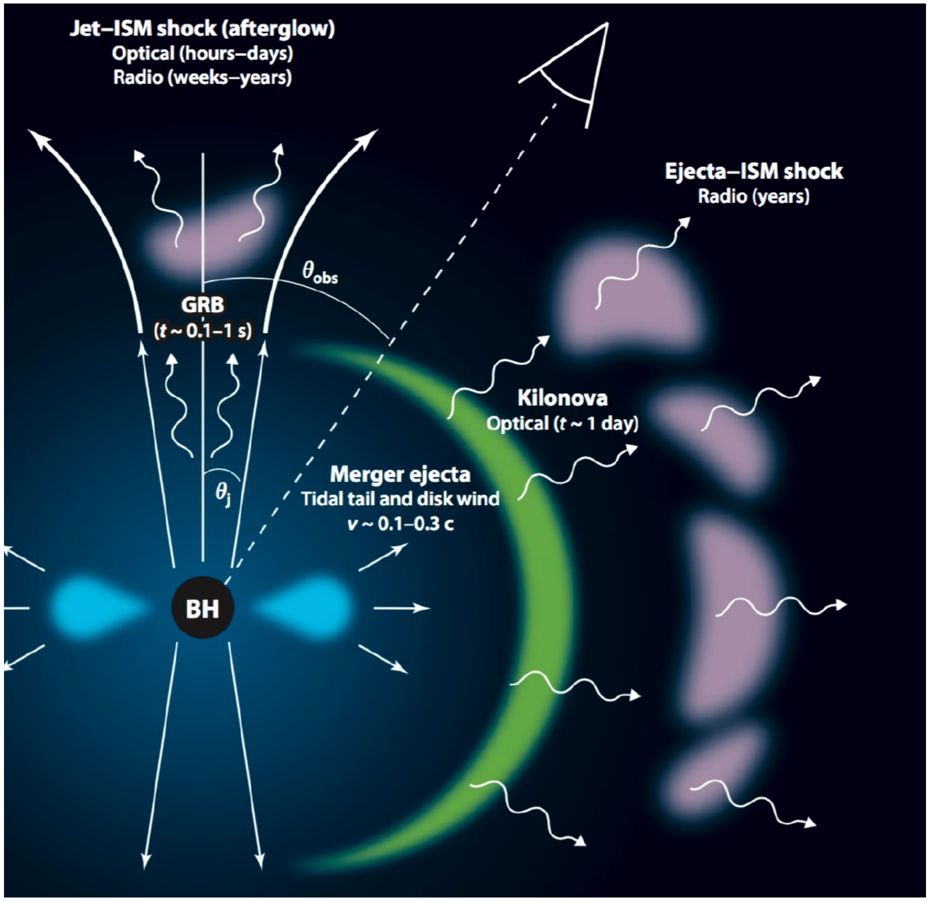

Berger 2014

Pythia

Reinforcement learning agent that strategizes follow-up to identify kilonovae

- Learns to evaluate the explore/exploit tradeoff
- Solves the credit assignment problem from any delayed consequences
- Adapts to new information, from its own actions or other sources

Toy sequential decision making under uncertainty problem:

- 9 transients, one of which (always) is the true kilonovae (min photometry = 1)
  - Contaminants are SNe, unassociated GRB afterglows, shock breakout (do not include observational significance)
- Follow-up in ZTF g, r, or i (300s exposure) per day
  - Finite horizon – 6 days (no action on day 1)
- Reward 1 if agents adds data to the kilonova else 0
  - Maximize the number of follow-up to the true kilonova (non-model specific objective with the expectation that more data ~ better constraints)

Pythia

- Learns online (collecting new experiences) in simulated environment
- Linear VFA (state-action value Q = x(s,a)^T ω
- x(s,a) is an CNN-autoencoder (for order invariance) representing the light curves with forecasted outcomes per action
- Learns ω via stochastic gradient decent and Adam optimizer

Algorithm SARSA and TD(0) target

Initialize w to small random weights
Set $\epsilon_0 = 1$
for k = 1, M **do** ▷ For each episode
  $\epsilon \leftarrow \epsilon_0 / k^n$
  Initialize s_1
  for t = 1, horizon **do**
    With probability ϵ select random action a_t
    otherwise select $a_t = \max_a \hat{Q}(s_t, a_t; \hat{w})$
    Execute action and observe reward r_t and next state s_{t+1} from environment
    With probability ϵ select random action a_{t+1}
    otherwise select $a_{t+1} = \max_a \hat{Q}(s_{t+1}, a_{t+1}; \hat{w})$
    Set $\Delta\hat{w} \leftarrow [r_t + \gamma\hat{Q}(s_{t+1}, a_{t+1}; \hat{w}) - \hat{Q}(s_t, a_t; \hat{w})]\nabla_w \hat{Q}(s_t, a_t; \hat{w})$
    ▷ : Loss is MSE between TD(0) target[3]substitute for Q^*) and current Q
    Update $\hat{w} \leftarrow \hat{w} + \alpha\Delta\hat{w}$ ▷ α is using Adam
  end for
end for

AI v humans

agent	score	frac
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**

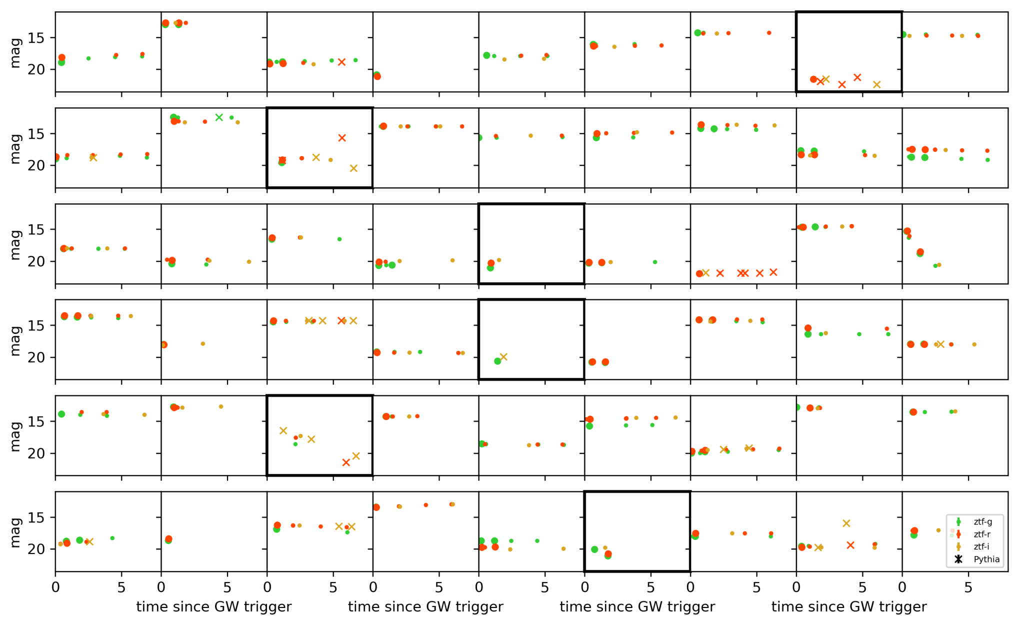

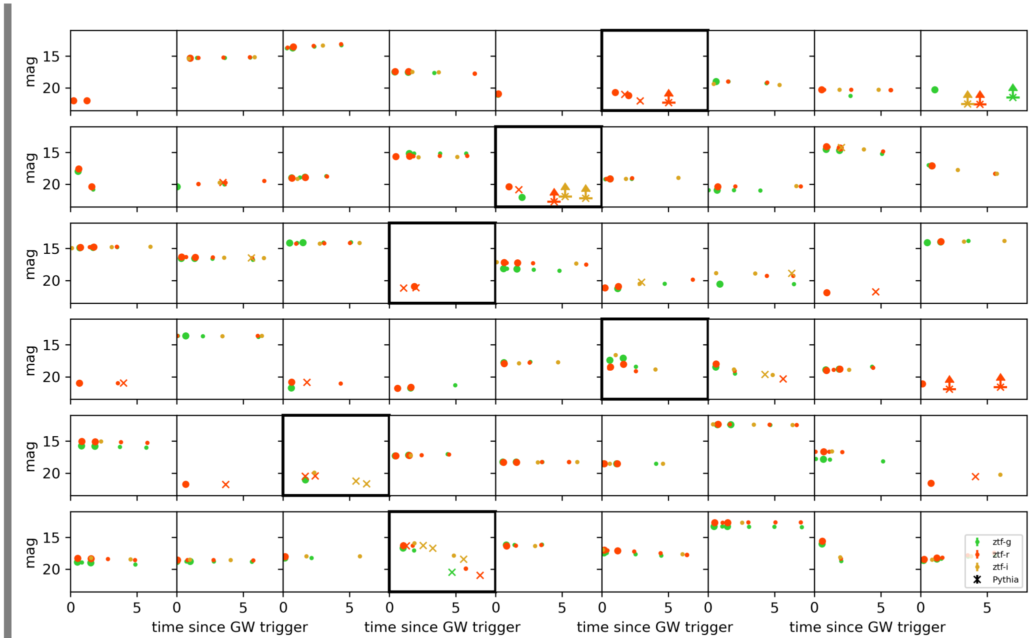

Sravan et al, in prep

PYTHIA

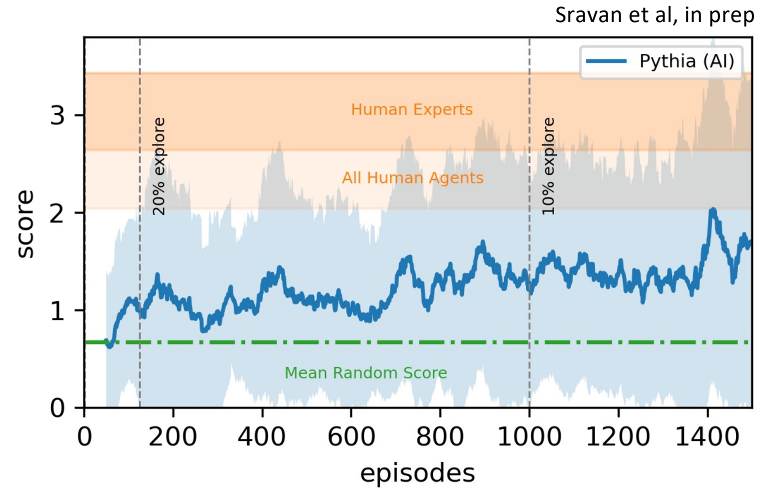

Linear VFA hypothesis class not sufficiently rich representation of true Q function

- 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

Carbon Footprint

Estimated emissions: 1210 kg of CO_2eq. assuming carbon efficiency of 0.432 kgCO_2eq/kWh

Approximately equal to:

- One round trip LAX-JFK (1180 Kg CO_2)
- 4900 km driven in an average combustion engine car

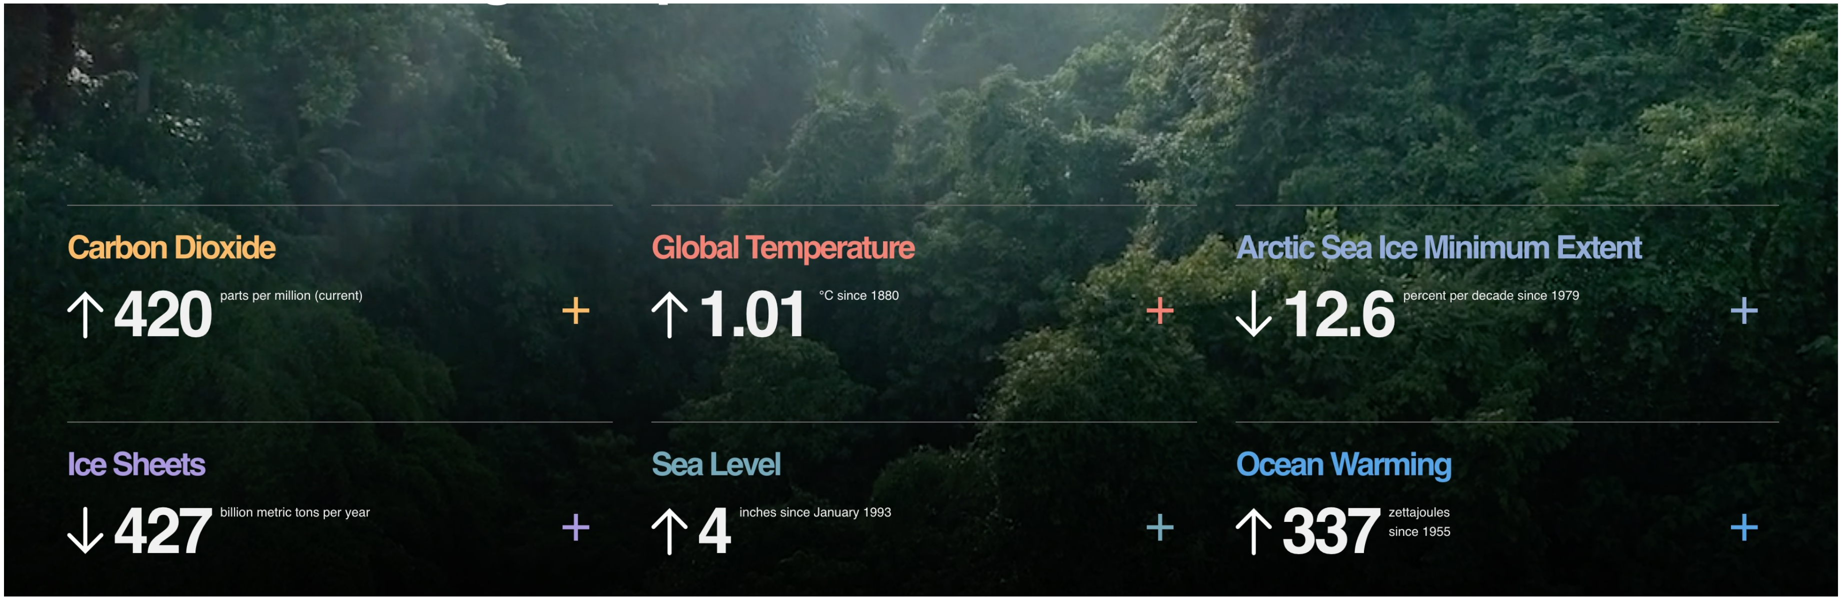

climate.nasa.gov

OUTLOOK

- Target maximizing constraints on the NS EoS and e.g. place constraints on maximum non-rotating NS mass
- Kilonova diversity with large samples
- Prepare for Rubin
  - Deeper and high SNR events
  - Motivates effective low-latency use of expensive space-based follow-up resources
- Other messengers!

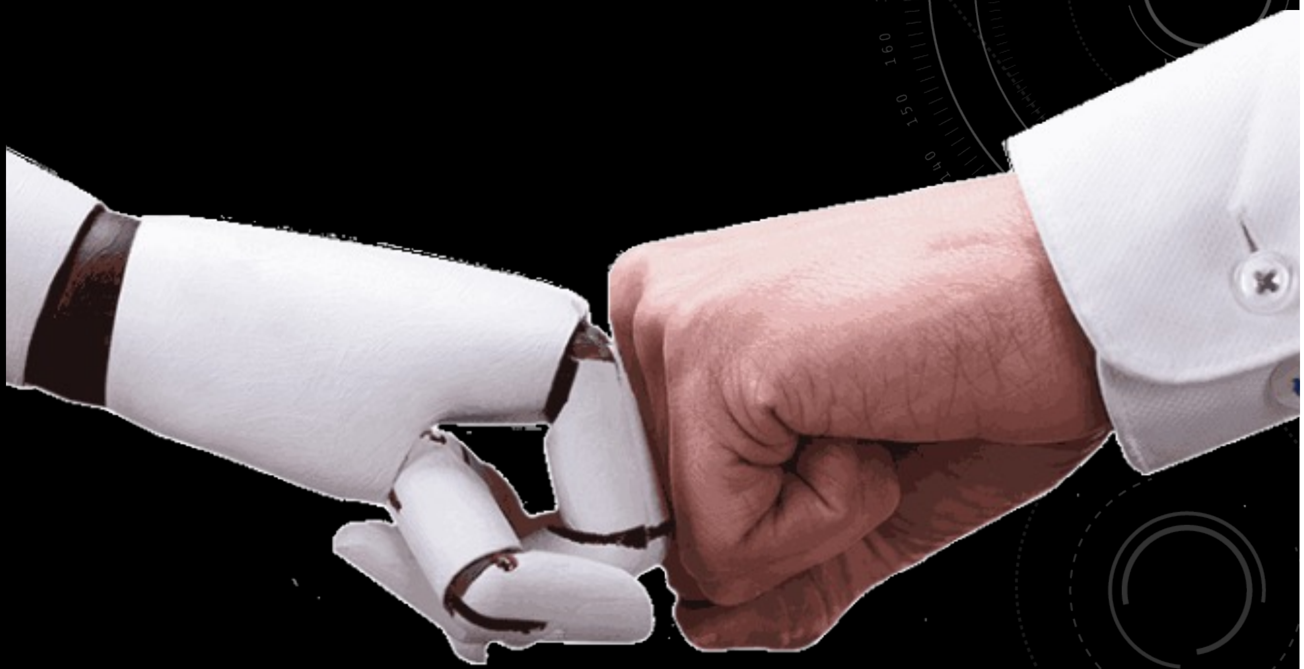

OUTLOOK

- Flexible to address any situation where real-time decisions need with resource limitations
- Approaches such as these are the ultimate human-machine symbiosis
  - Reduce burden of tedious work (especially for well-defined science cases)
  - Leave innovation and discovery to humans (?)

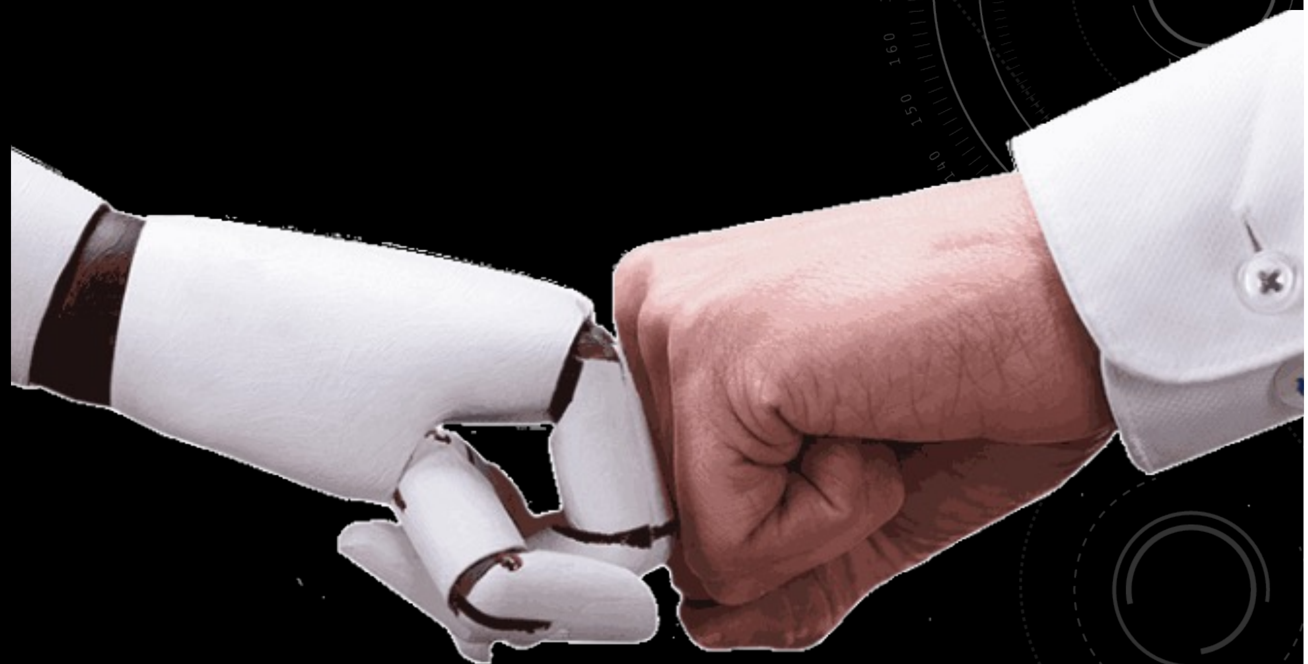

BACKUP CONTENT

REFITT FOR ZTF

- Training dataset:
  - “Classical” SNe (Ia, II, IIn, IIb, Ib, Ic, Ic-BL) simulated from ZTF BTS, PS1, and all historic classified SNe with sdssg, sdssr photometry
  - 5k per type spanning flat z space 0-0.8 for SN Ia, 0-0.1 for CC SNe
- ML and forecasting:
  - Multi-D Gaussian Process LC fit
  - use k similar training LCs using Xception penultimate vector of modal training class (implemented as balltree)
  - align with cross-correlation
- Daily run at 0900PT (10 mins on 24 cores)
  - Ingested via Antares: within 60 d of trigger, <21 d since last photo, at least 3 photo with at least two in the same band >5 hours apart
- Recommendations for:
  - Photometry for events approaching peak and not in ZTF’s observing plan
  - Classifications within |1wk| of forecasted peak
  - Anomalies (poor forecast)

WHY ONLY GW170817?

- For O3:
  - Median skymap size ~4000 sq deg
  - Median distance:
    - BNS ~ 240Mpc
    - NSBH ~ 320 Mpc

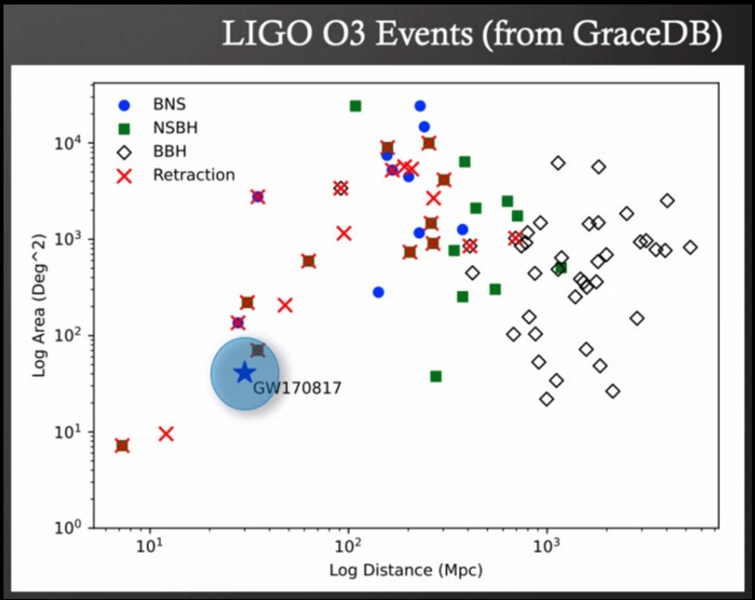

Petrov+ 2021 (adapted)

Run	BNS	NSBH	BBH
	Median 90% credible area (deg^2) [a]		
O3	1672^{+94}_{-110}	1970^{+110}_{-110}	1069^{+43}_{-41}
O4	1820^{+190}_{-170}	1840^{+150}_{-150}	335^{+28}_{-17}
O5	1250^{+120}_{-120}	1076^{+65}_{-75}	$230.3^{+7.8}_{-6.4}$
	Median luminosity distance (Mpc) [a]		
O3	$176.1^{+6.2}_{-5.7}$	$337.6^{+10.9}_{-9.6}$	871^{+31}_{-28}
O4	$352.8^{+10.3}_{-9.8}$	621^{+16}_{-14}	1493^{+25}_{-33}
O5	620^{+16}_{-17}	1132^{+19}_{-23}	2748^{+30}_{-34}
	Annual number of detections [cd]		
O3	5^{+14}_{-5}	13^{+15}_{-9}	24^{+18}_{-12}
O4	34^{+78}_{-25}	72^{+75}_{-38}	106^{+65}_{-42}
O5	190^{+410}_{-130}	360^{+360}_{-180}	480^{+280}_{-180}