

# Transient-optimised source classification with Bayesian convolutional neural networks

Thomas Killestein (he/him)

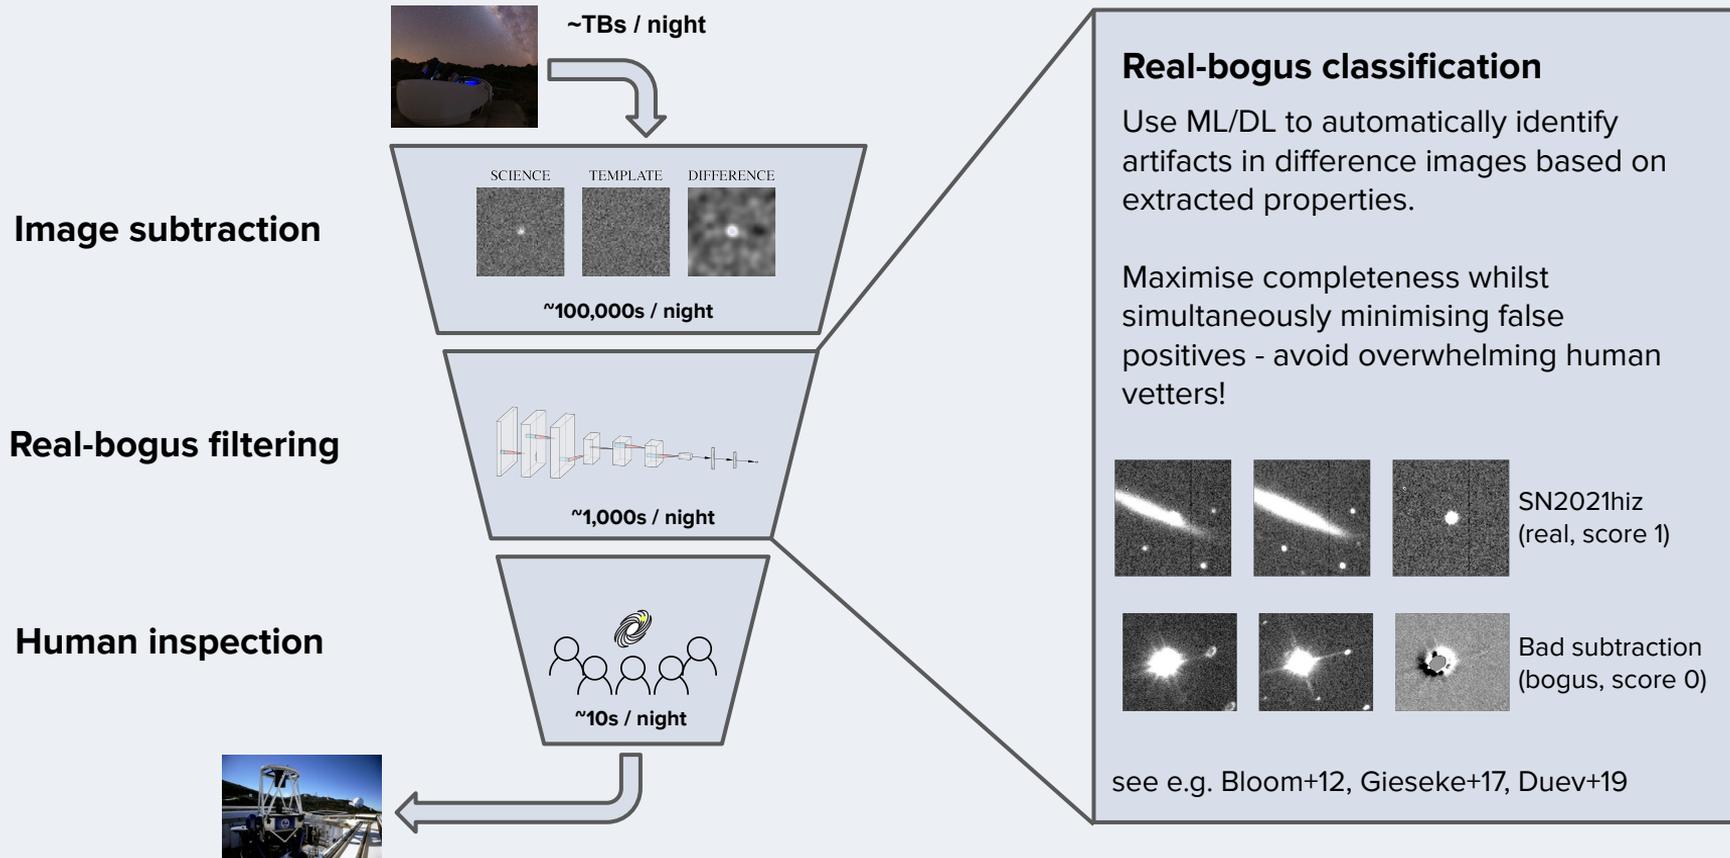
*University of Warwick*



*Presented at NAM 2021 AstroML session*

# Time-domain astronomy: the real-bogus problem

How can we efficiently sift the deluge of data from current/upcoming transient surveys to maximise science yield?

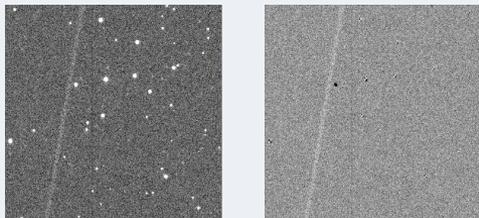


# The Gravitational-wave Optical Transient Observer

## Each telescope/node:

8x40cm astrographs, combined FoV of **40 square degrees** - reaches **L ~ 20.5 in 180s** enabling rapid and deep surveys of GW localisation regions in real-time.

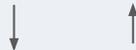
## Difference imaging



**More info:** talks in Transients Diversity by Kendall Ackley and Ben Gompertz tomorrow!

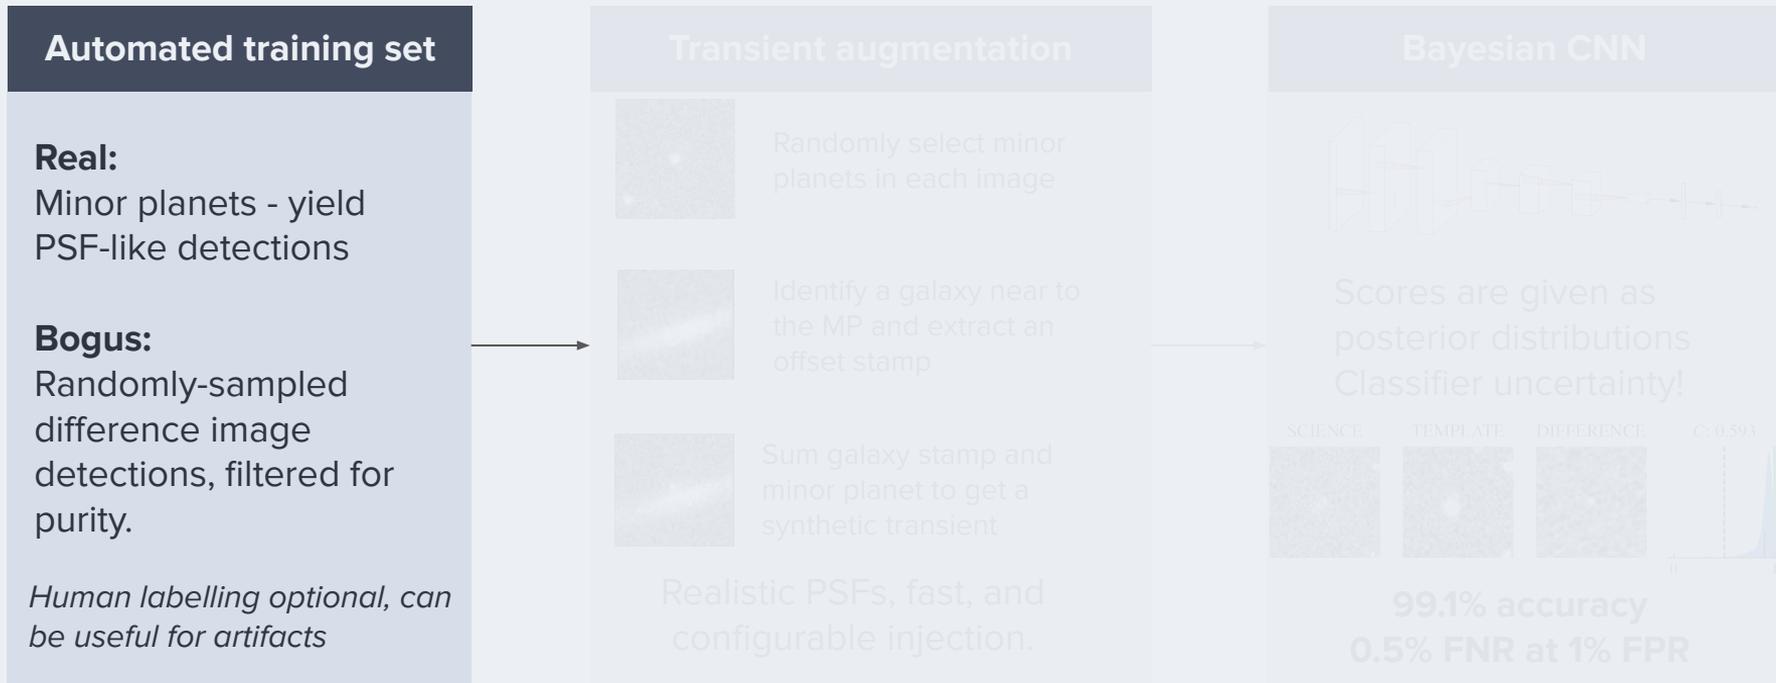


Rapidly fading (~hours) GW counterparts necessitates **real-time, accurate** source classification - only possible with high-performance ML/DL.



The significant data volume generated by GOTO (and other surveys) provides an ideal **data-rich proving ground** for novel ML/DL algorithms applied to a diverse range of data modalities.

# Training set generation



Highly configurable, image-level parallelism allows generating a **400k example dataset in under 24h!**

For full implementation details: **Killestein et al., (2021)** [arXiv: 2102.09892, in MNRAS]

All project code is open source via GitHub - *GOTO-OBS/gotorb*

# Transient augmentation

## Automated training set

### Real:

Minor planets - yield  
PSF-like detections

### Bogus:

Randomly-sampled  
difference image  
detections, filtered for  
purity.

*Human labelling optional, can  
be useful for artifacts*

## Transient augmentation



Randomly select minor  
planets in each image



Identify a galaxy near to  
the MP and extract an  
offset stamp



Sum galaxy stamp and  
minor planet to get a  
synthetic transient

Realistic PSFs, fast, and  
configurable injection.

## Bayesian CNN



Scores are given as  
posterior distributions  
Classifier uncertainty!



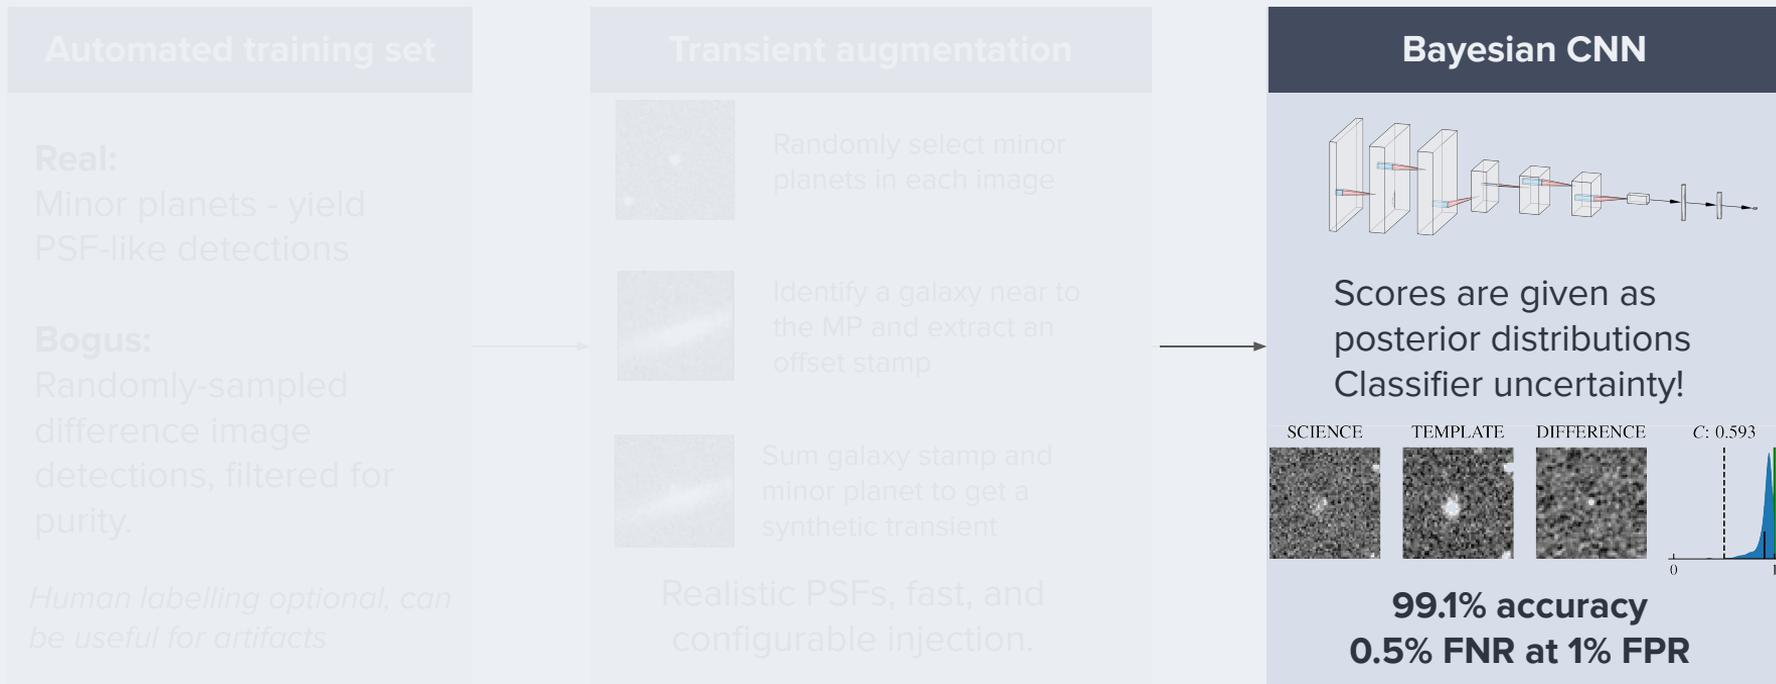
**99.1% accuracy**  
**0.5% FNR at 1% FPR**

Highly configurable, image-level parallelism allows generating a **400k example dataset in under 24h!**

For full implementation details: **Killestein et al., (2021)** [arXiv: 2102.09892, in MNRAS]

All project code is open source via GitHub - *GOTO-OBS/gotorb*

# Model architecture



Highly configurable, image-level parallelism allows generating a **400k example dataset in under 24h!**

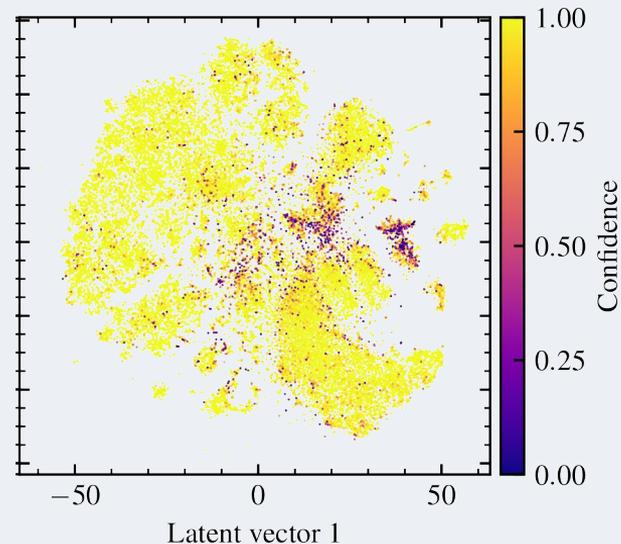
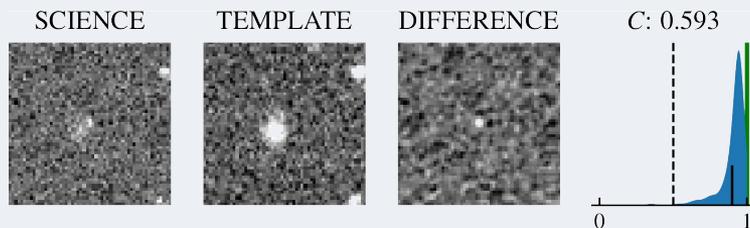
For full implementation details: **Killestein et al., (2021)** [arXiv: 2102.09892, in MNRAS]

All project code is open source via GitHub - *GOTO-OBS/gotorb*

# Bayesian neural networks: uncertainty-aware predictions

**Dropout:** at each train step, **disconnect** some fraction of neurons randomly for regularisation and generalisation.

**Monte Carlo Dropout:** sampling different realisations of dropped-out models is  $\sim$ equivalent to sampling from the model output posterior (via an underlying Gaussian process). This allows quantification of **epistemic uncertainty** arising from our choice of model weights, and marginalisation.



t-SNE gives idea of coverage of training space, underrepresented classes, intrinsically difficult source types.

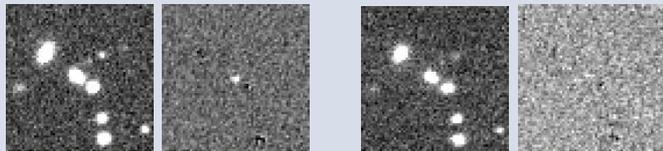
For in-depth theoretical discussion, see **Gal and Ghahramani (2014a,b and 2016)**

# Dataset-induced biases and alignment

Think carefully about the inductive biases your dataset imposes!

## Dataset-induced biases

Although introducing synthetic transients *alone* markedly improves the recovery of supernovae, this introduces an unintended side-effect: because every single one of our synthetic transients has an underlying galaxy, the network learns that all stamps containing a galaxy are **real detections...**



Confirmed SN (real)

Residual (false pos.)

**Fix:** add galaxies into the negative class as balance

Neural networks only care about **minimising the loss!**  
It's up to you to make sure your loss function and dataset reflect what you want it to learn.

## Combining labelled datasets

Human labelling is a valuable tool for augmenting the performance of a classifier trained on synthetic data - but need to be aware of (hidden) differences in labelling conventions between them.



Minor planets graded bogus by humans (contaminant in GW searches) [FN]



Low-amplitude stellar variability (ambiguous so graded bogus) [FN]

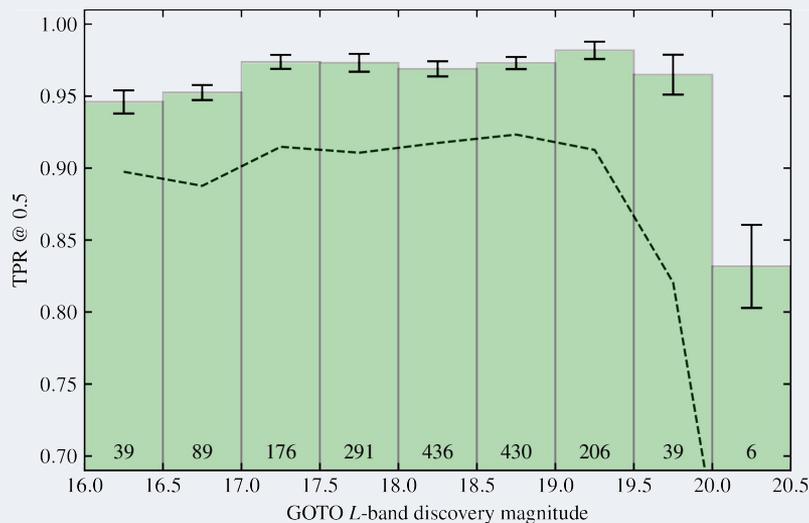


Genuine human errors (misclicks, vetting gets boring!) [FP]

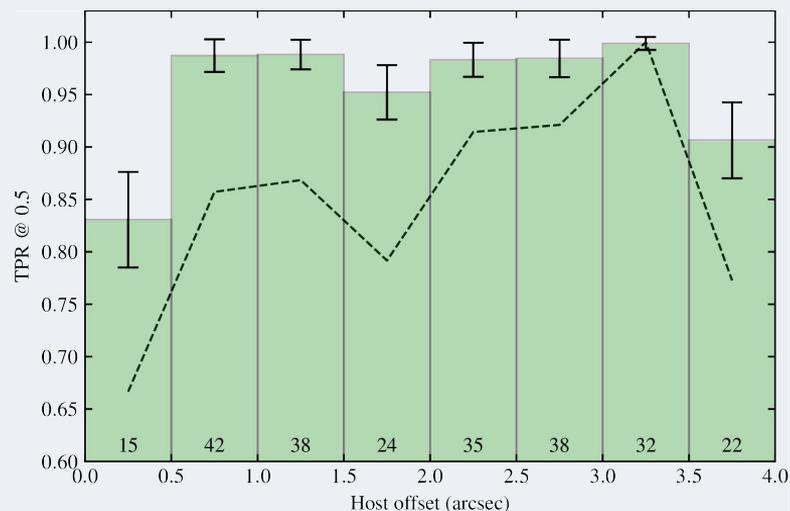
Can exploit the scarcity of mislabelled examples + BNN posterior confidences to semi-automatically filter these out - see Killestein+21 for more details

# Transient optimisation: marked performance improvements

Test the classifier on >900 spectroscopically-confirmed transients recovered in the GOTO prototype phase — performance likely to be even better with in-progress instrumentation and pipeline upgrades.



**Significant increase in recovery of faint transients** - excellent prospects for kilonova recoveries in local Universe.



Uniform coverage of host-offset space leads to **boosted recovery of nuclear transients**

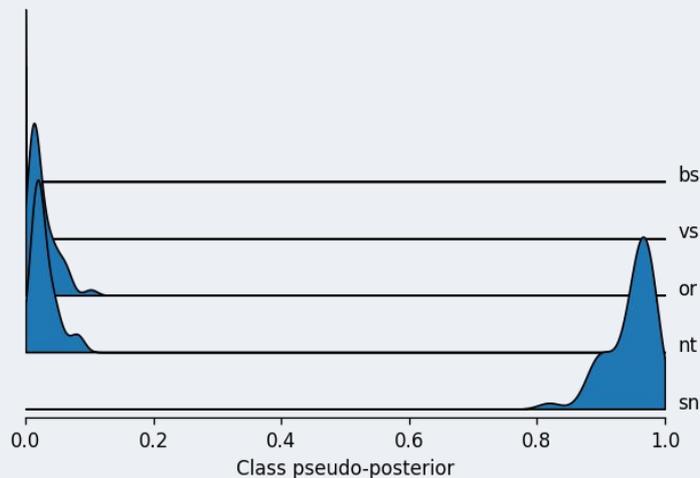
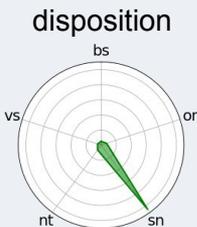
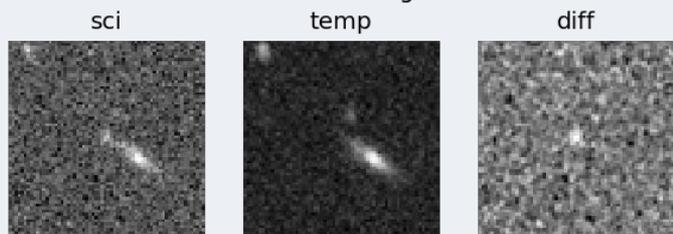
# Next steps: context-aware, granular classification

Image-level data can only go so far: need to fold in complex multi-modal data to give the full picture.

bs: 0.000 vs: 0.000 nt: 0.028 sn: 0.947 or: 0.025

$\text{argmax}(\text{pred})$  class: sn

reconstructed real-bogus: 1.000



Early work  
(image only)

Use methods developed to build a prompt stamp classifier (similar to ALeRCe, Carrasco-Davis+ 2019) -- **principled, Bayesian, high-performance multi-label classification** for low-latency alerts of interesting candidates

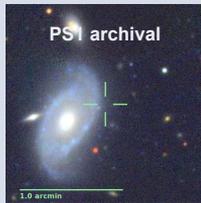
**Maximise available information**  
If a human uses it in their vetting the classifier should too!

- + contextual awareness
- + hierarchical meta-classification
- + adaptive stream learning

# Fusing image-level data and contextual catalogs

Overarching aim: make available all human-interpretable information to any ML/DL classifier

## Image-only:



- + Always available (discovery image)
- Limited by image quality/resolution
- Apparent type != actual type

## Catalog-only:



```
ra, dec, gmag, ...  
32.042, 12.321, 19.042, ...  
31.987, 12.021, 18.428, ...  
32.321, 12.124, 21.421, ...
```

Salient information already extracted

- + Far deeper than survey images
- Incomplete (galaxy catalogs)
- Incorrect (misclassifications)
- Inhomogeneous (variable coverage)

## Image-level information can compensate for catalog issues!

Ensemble catalogs into VS/AGN/GLX sources, and compute basic nearest-neighbour distances - combine this with image-only CNN scores in a random forest

Train on censored data - randomly exclude contextual information to **simulate incompleteness** and enable **robust classifications** with all information available.

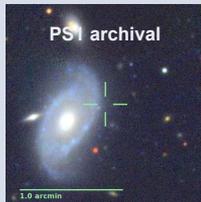


96% recovery of supernovae  
98% test set accuracy

# Fusing image-level data and contextual catalogs

Overarching aim: make available all human-interpretable information to any ML/DL classifier

## Image-only:



- + Always available (discovery image)
- Limited by image quality/resolution
- Apparent type != actual type

## Catalog-only:



```
ra, dec, gmag, ...  
32.042, 12.321, 19.042, ...  
31.987, 12.021, 18.428, ...  
32.321, 12.124, 21.421, ...
```

Salient information already extracted

- + Far deeper than survey images
- Incomplete (galaxy catalogs)
- Incorrect (misclassifications)
- Inhomogeneous (variable coverage)

## Image-level information can compensate for catalog issues!

Ensemble catalogs into VS/AGN/GLX sources, and compute basic nearest-neighbour distances - combine this with image-only CNN scores in a random forest

Train on censored data - randomly exclude contextual information to **simulate incompleteness** and enable **robust classifications** with all information available.

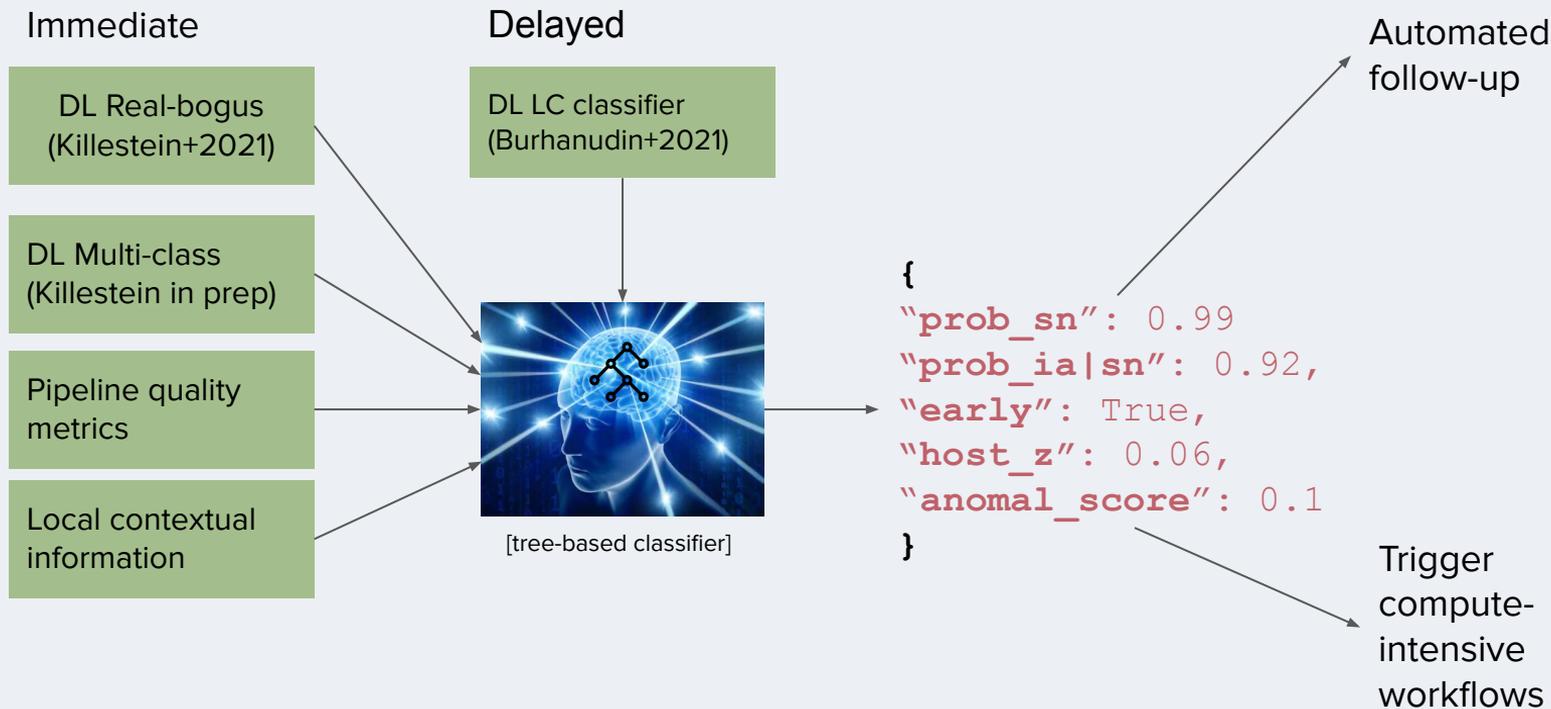
work in progress

True label \ Predicted label	bs	vs	nt	sn	or
bs	97.9% (3263)	0.0% (0)	0.0% (0)	0.1% (2)	2.1% (69)
vs	0.4% (16)	98.8% (4028)	0.6% (24)	0.0% (0)	0.2% (9)
nt	0.0% (1)	0.9% (26)	98.6% (2776)	0.0% (0)	0.4% (11)
sn	0.3% (8)	0.0% (1)	0.0% (0)	96.1% (2660)	3.6% (100)
or	0.3% (19)	0.0% (0)	0.0% (0)	0.1% (4)	99.6% (6050)

**96% recovery of supernovae**  
**98% test set accuracy**

# End goal: hierarchical meta-classification

Strong tree-based **ensembles of expert classifiers** - interpretable, efficient, and performant.  
Hierarchical approaches **mimic human approaches** to transient classification.



We combine **Bayesian convolutional neural networks** with a **novel training set generation and augmentation procedure** to achieve state of the art results on real-bogus classification with **minimal human labelling**.

In active development: **context-aware deep learning ensembles** to bridge the gap between image-level classification and catalog-based contextual classification in a principled, robust manner for even better source filtering!



[t.killestein@warwick.ac.uk](mailto:t.killestein@warwick.ac.uk)



[tkillestein](https://github.com/tkillestein), [GOTO-OBS/gotorb](https://github.com/GOTO-OBS/gotorb)



**GOTO**

GRAVITATIONAL-WAVE OPTICAL TRANSIENT OBSERVER



[goto-observatory.org](https://goto-observatory.org)