Unveiling AGN through Time-Domain Variability and Machine Learning

In collaboration with Demetra De Cicco, Stefano Cavuoti, Ylenia Marruccia and Giuseppe Riccio

Funded by the Horizon 2020 Framework Programme of the European Union Grant Agreement No. 101168906

Natale De Bonis

My background

I took both my bachelor's and master's degree in Naples at Università degli studi di Napoli Federico II

Research interests

- Galaxies evolution
- AGN physics
- Machine learning

ACTIVE GALACTIC NUCLEI (AGN)

- Bolometric luminosities up to $10^{15} L_{\odot}$
- Emission over the whole electromagnetic spectrum
- Highly variable
- Extremely compact
- Polarized emission

AGN UNIFIED MODEL

- Super massive black hole $M > 10^6 M_{\odot}$
- Accretion disk
- Corona
- Broad and Narrow line region
- Torus
- Bipolar Jets

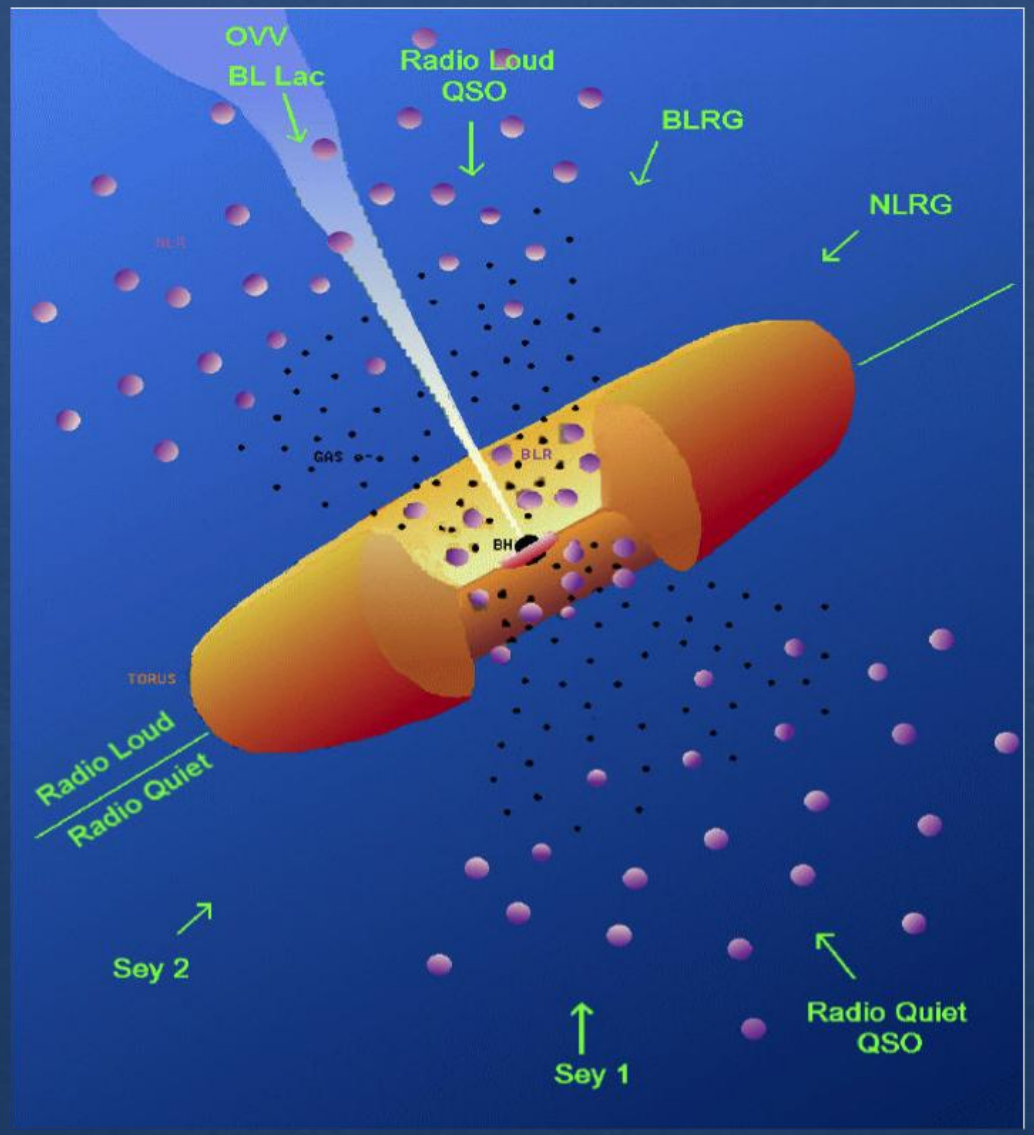

As their spectral energy density (SED) covers the whole electromagnetic spectrum, we can study AGN in multiple ways:

- Photometrically
- Spectroscopically
- Polymetrically

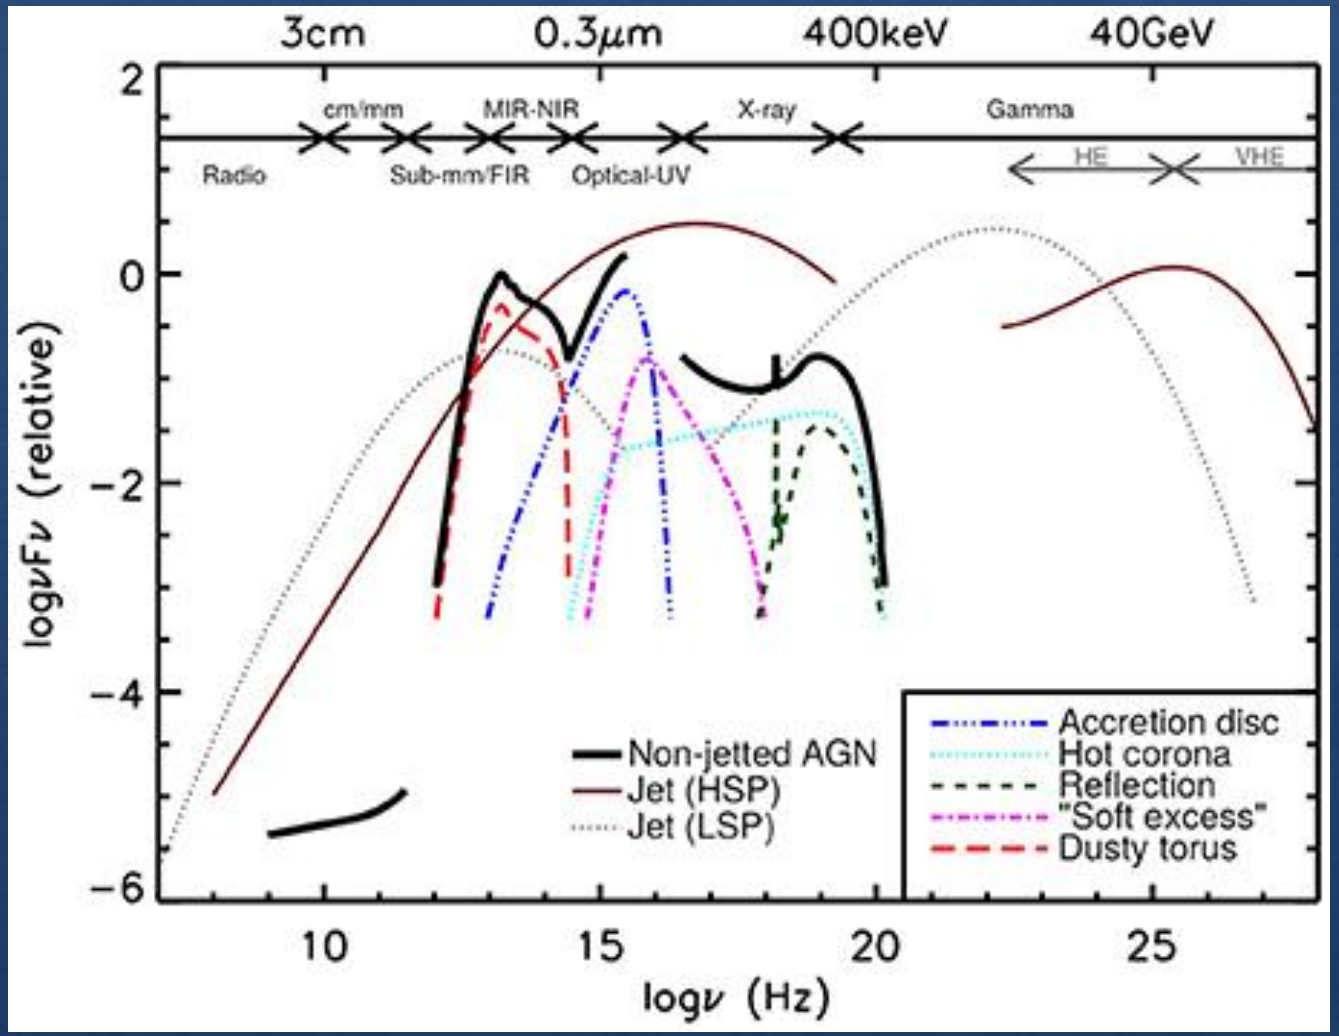

Variability:

- AGN show brightness variation at all wavelengths and they are more variable at higher frequencies (shorter wavelengths)
- They show very rapid variation, even on time scales of days or shorter
- Variability can help us constrain the geometry and dimension of the various AGN components: if they vary this fast, they should be this small

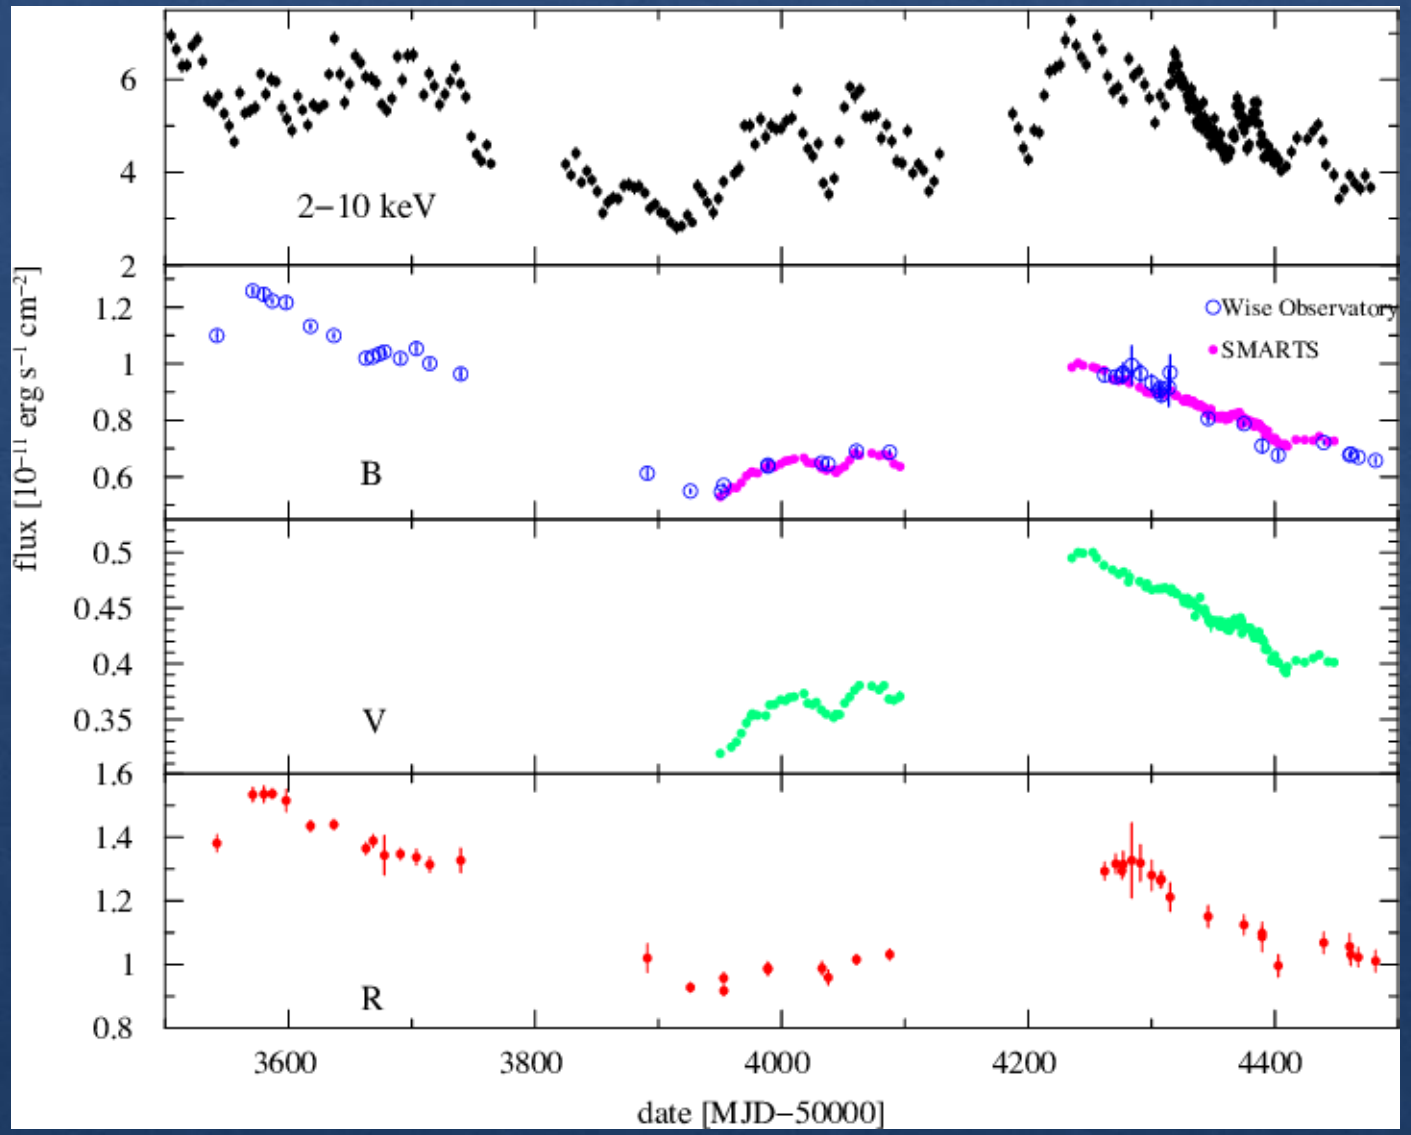

Reverberation mapping

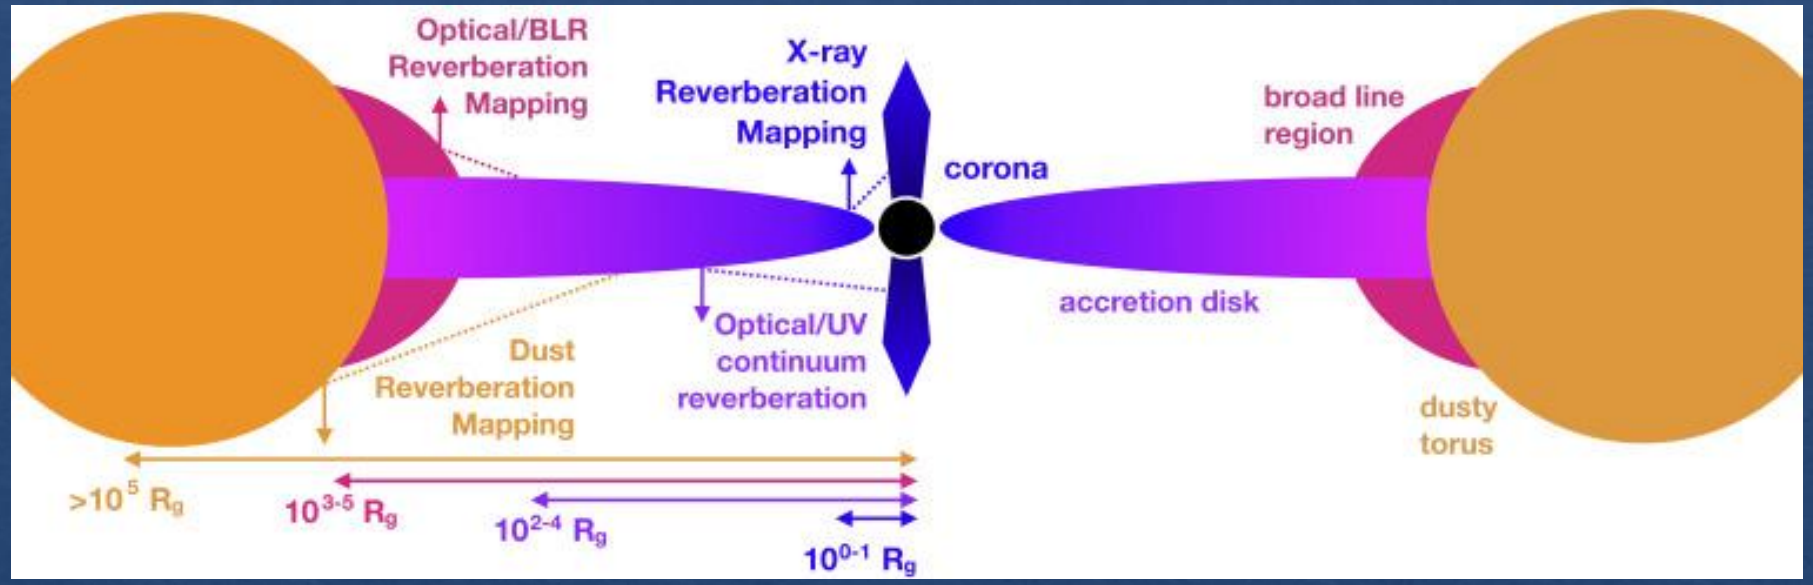

- A technique that exploits the intrinsic variability of AGN.
- Measures the time delay between a variation in the accretion disk emission and the response of other emission regions.
- This delay allows us to estimate the distance of the region from the supermassive black hole (SMBH): Distance=c×Delay

Anomaly detection in AGN light curves

Anomaly detection can help us uncover:

- Systematic problems (instrumental artifacts, calibration errors, or data processing issues)
- Rare or unexpected phenomena (unusual variability patterns, changing-look AGN)
- Transient phenomena (flares, dips, or short-lived events)

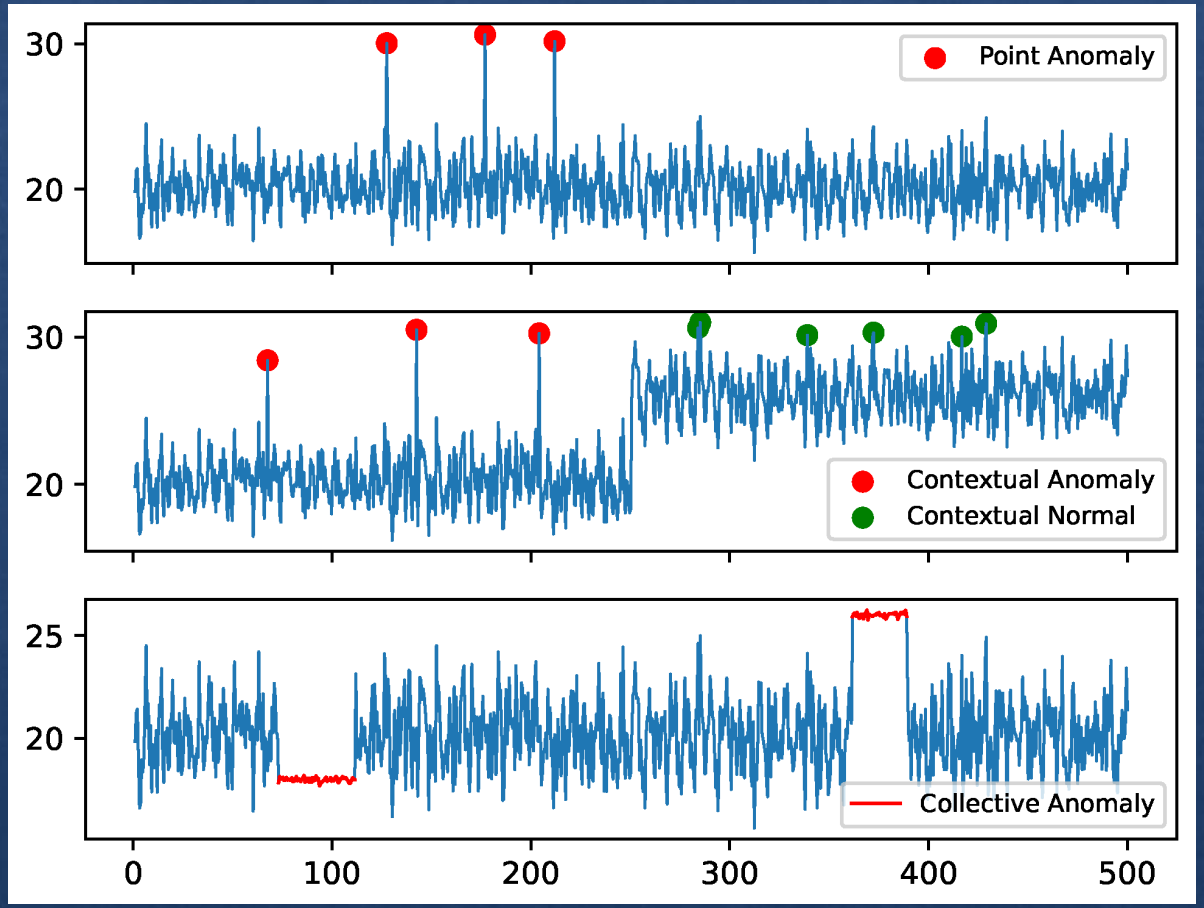

Anomaly detection regimes

Light curve-based

Pro:

- **Detect anomalies directly in the curve (**
- **Not to ally expensiv**

Cons:

- **AG ar ary in leng king d modeling or comparison more challenging**

Machine learning

Feature-based

Pro:

- **We can easily work with light curves of different lengths**
- **We can have a more 'physical' interpretation of the results**

Cons:

- **Computationally expensive (we have to compute the features)**

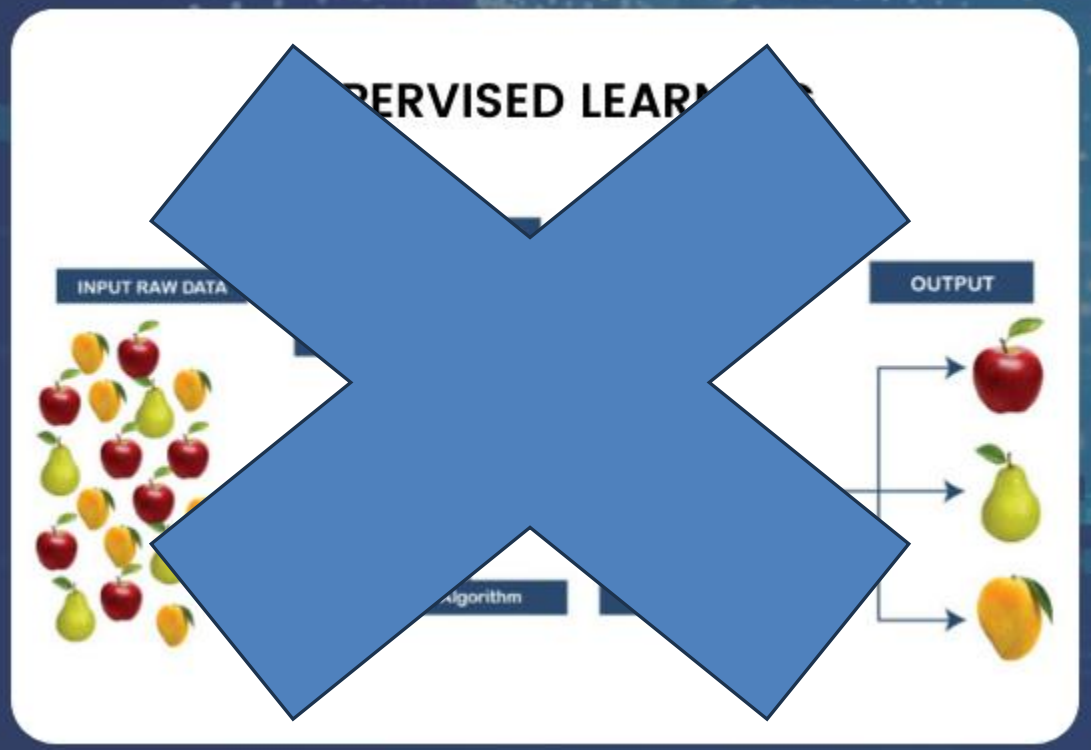

Supervised learning

- Requires **prior knowledge** in the form of labeled data
- Can perform both **regression** (predicting continuous values) and **classification** (assigning categories)
- Learns a mapping from inputs to known outputs

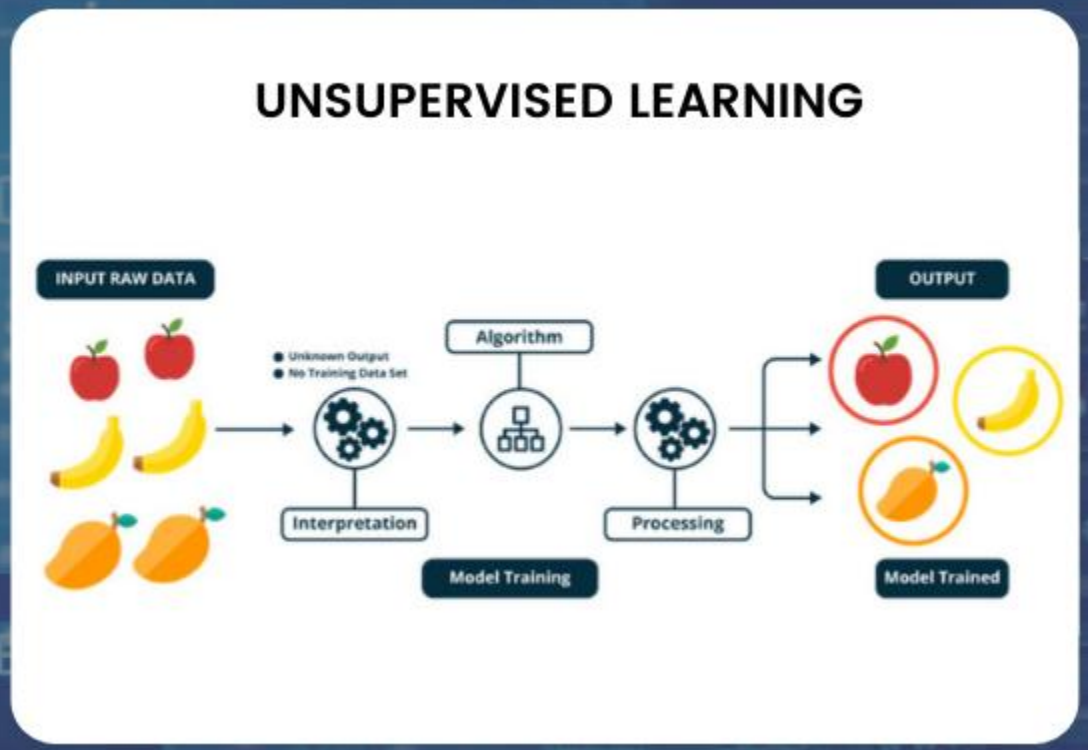

Unsupervised learning

- Works **without labeled data**, so no prior knowledge is needed
- Focuses on discovering hidden patterns or structures in the data
- Common techniques include **clustering** and **dimensionality reduction**

AutoEncoder (AE)

Autoencoders are a special class of unsupervised learning algorithms.

They have two main components:

- Encoder, which compresses the input data in a latent space
- Decoder, which reconstructs the input data from its compressed representation

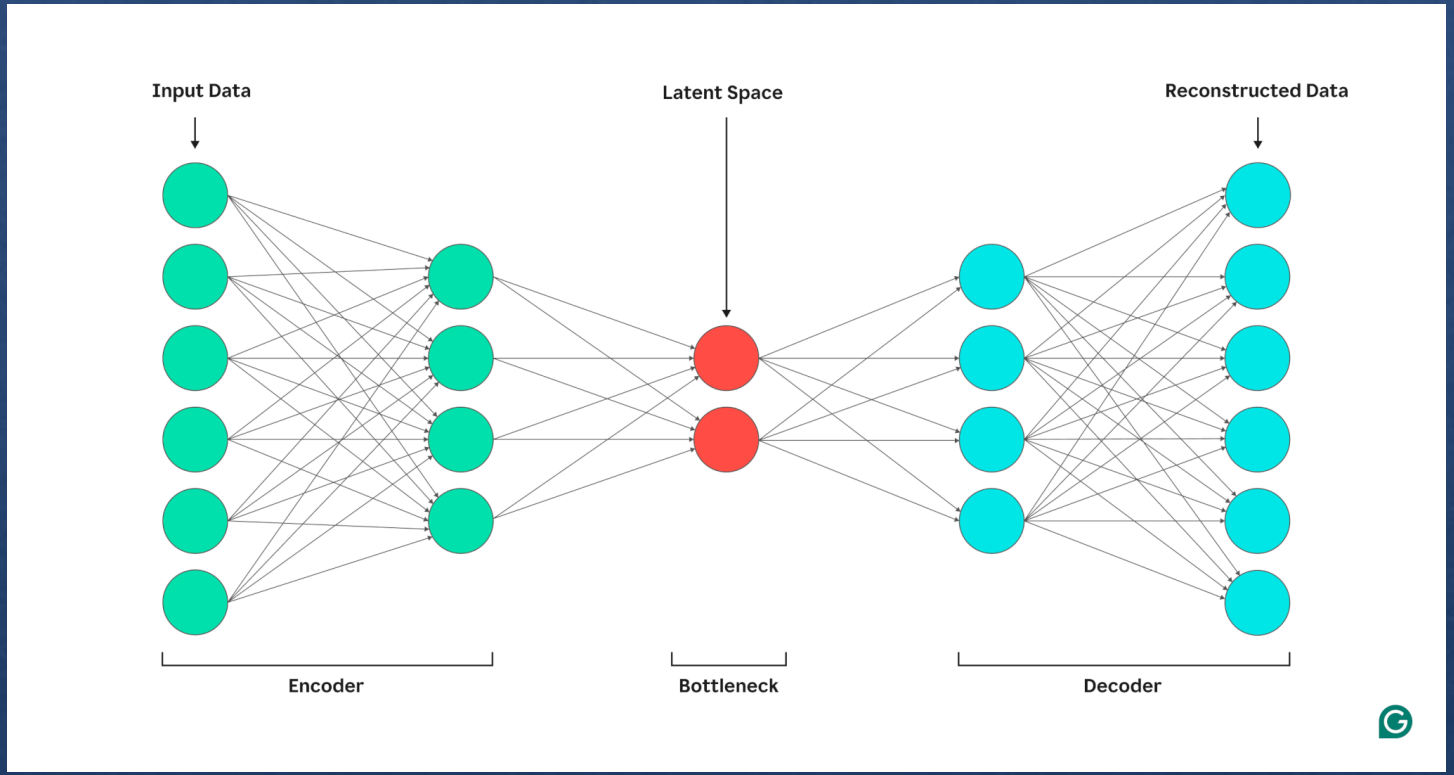

Case study: Anomaly detection for obscured AGN using an autoencoder

Why use an autoencoder:

- **They can learn non-linear representations of complex data distributions**
- **They offer a flexible architecture, so we can easily tailor the neural networks to our specific needs**

Why study obscured AGN:

- **They are notoriously difficult to distinguish from normal galaxies based solely on their optical variability properties**
- **Finding an anomalous obscured AGN could lead to a correction in its classification or rare/unexpected behavior**

Dataset description

The dataset used consists of *r*-band observations from the three observing seasons of the COSMOS field by the VLT survey telescope.

Only light curves with at least 27 points were considered, with an avarege mag $r \leq 23.5$ within an 1” arcsec aperture

We have two main samples: the unlabeled set (US) and the labeled set (LS)

The US consists of 20647 sources

The LS is composed of 2414 sources for which we have a classification in star, galaxy or AGN

For 413 of these sources we have a subclassification which gives us information about the AGN type and the selection criteria (X-ray, medium infrared (MIR) or variability)

Feature extraction

We trained our algorithms on a set of 37 features.

Some of these features were computed using the Python library Feature Analysis for Time Series (FATS), and the others are either photometric features (colors), morphological features (class star) or Damped Random Walk parameters.

We will divide them in three classes: Variability features, photometric features and morphological features.

Feature	Description	Reference
	Variability features	
A_{SF}	RMS magnitude difference of the Structure Function (SF), computed over a 1 yr timescale	Schmidt et al. (2010)
γ_{SF}	Logarithmic gradient of the mean change in magnitude	Schmidt et al. (2010)
GP_DRW_τ	Relaxation time (i.e., time necessary for the time series to become uncorrelated), from a Damped Random Walk (DRW) model	Graham et al. (2017)
GP_DRW_σ	Variability of the time series at short timescales ($t \ll \tau$), from a DRW model	Graham et al. (2017)
ExcessVar	Measure of the intrinsic variability amplitude	Allevato et al. (2013)
P_{var}	Probability that the source is intrinsically variable	McLaughlin et al. (1996)
IAR_ϕ	Level of autocorrelation using a discrete-time representation of a DRW model	Eyheramendy et al. (2018)
Mean	Mean magnitude	Nun et al. (2015)
Con	Index introduced for the selection of variable stars from the OGLE database	Kim et al. (2011)
Amplitude	Half of the difference between the median of the maximum 5% and of the minimum 5% magnitudes	Richards et al. (2011)
AndersonDarling	Test of whether a sample of data comes from a population with a specific distribution	Nun et al. (2015)
Autocor_length	Lag value where the autocorrelation function becomes smaller than η^e	Kim et al. (2011)
Beyond1Std	Percentage of points with photometric mag that lie beyond 1σ from the mean	Richards et al. (2011)
η^e	Ratio of the mean of the squares of successive mag differences to the variance of the light curve	Kim et al. (2014)
Gskew	Median-based measure of the skew	—
LinearTrend	Slope of a linear fit to the light curve	Richards et al. (2011)
MaxSlope	Maximum absolute magnitude slope between two consecutive observations	Richards et al. (2011)
Meanvariance	Ratio of the standard deviation to the mean magnitude	Nun et al. (2015)
MedianAbsDev	Median discrepancy of the data from the median	Richards et al. (2011)
MedianBRP	Fraction of photometric points within amplitude/10 of the median magnitude	Richards et al. (2011)
MHAOV Period	Period obtained using the P4J Python package (`https://github.com/phuijse/P4J`)	Huijse et al. (2018)
PairSlopeTrend	Fraction of increasing first differences minus decreasing ones over the last 30 time-sorted mag measures	Richards et al. (2011)
PercentAmplitude	Largest percentage difference between either max or min mag and median mag	Richards et al. (2011)
Q31	Difference between the third and the first quartile of the light curve	Kim et al. (2014)
Period_fit	False-alarm probability of the largest periodogram value obtained with LS	Kim et al. (2011)
Ψ_{CS}	Range of a cumulative sum applied to the phase-folded light curve	Kim et al. (2011)
Ψ_η	η^e index calculated from the folded light curve	Kim et al. (2014)
R_{CS}	Range of a cumulative sum	Kim et al. (2011)
Skew	Skewness measure	Kim et al. (2011)
Std	Standard deviation of the light curve	Nun et al. (2015)
StetsonK	Robust kurtosis measure	Kim et al. (2011)
	Morphological features	
class_star	HST stellarity index	Koekemoer et al. (2007); Scoville et al. (2007)
	Photometric features	
$u-B$	CFHT u magnitude – Subaru B magnitude	Laigle et al. (2016)
$B-r$	Subaru Suprime-Cam B mag – Subaru Suprime-Cam $r+$ mag	Laigle et al. (2016)
$r-i$	Subaru Suprime-Cam $r+$ mag – Subaru Suprime-Cam $i+$ mag	Laigle et al. (2016)
$i-z$	Subaru Suprime-Cam $i+$ mag – Subaru Suprime-Cam $z++$ mag	Laigle et al. (2016)
$z-y$	Subaru Suprime-Cam $z++$ mag – Subaru Hyper-Suprime-Cam y mag	Laigle et al. (2016)
ch21	*Spitzer* 4.5μm (channel2) – 3.6μm (channel1)	Laigle et al. (2016)

Table 1. List of the features used in this study, grouped into three categories: variability features, morphological features, and photometric features.

Autoencoder architecture and training process

The logic behind the use of the autoencoder for anomaly detection lies in its ability to reconstruct the input data accurately when the data resembles those seen during training. As a consequence, if a light curve shows a big reconstruction error (RE),it will be flagged as anomalous.

We trained the neural networks on the US

We use dropout layers to avoid overfitting, a common problem that can occur in machine learning, i.e., when the algorithms fail to generalize new inputs. We also applied batch normalization layers, which standardize the inputs to each layer during training, helping to stabilize and accelerate the learning process

Encoder	**Decoder**
Input (37)	Latent (10)
Dense (1024, ReLU)	Dense (128, LeakyReLU)
Dropout (0.3) + BatchNorm	Dropout (0.5) + BatchNorm
Dense (768, ReLU)	Dense (256, LeakyReLU)
Dropout (0.3) + BatchNorm	Dropout (0.5) + BatchNorm
Dense (512, ReLU)	Dense (512, LeakyReLU)
Dropout (0.3) + BatchNorm	BatchNorm
Dense (256, ReLU)	Output (37, Linear)
Dropout (0.3) + BatchNorm	
Dense (128, ReLU)	
Dense (10, Tanh)	

We computed the reconstruction error by means of the mean squared error (MSE):

$$RE_i = \frac{1}{N}\sum_{j=1}^{N}(X_{ij} - \hat{X}_{ij})^2$$

Where N is the dimension of the starting space (37) X is the the original data and $\hat{X}$ is the reconstructed data.

The threshold is defined as the 95° percentile in the Re distribution of the US set.

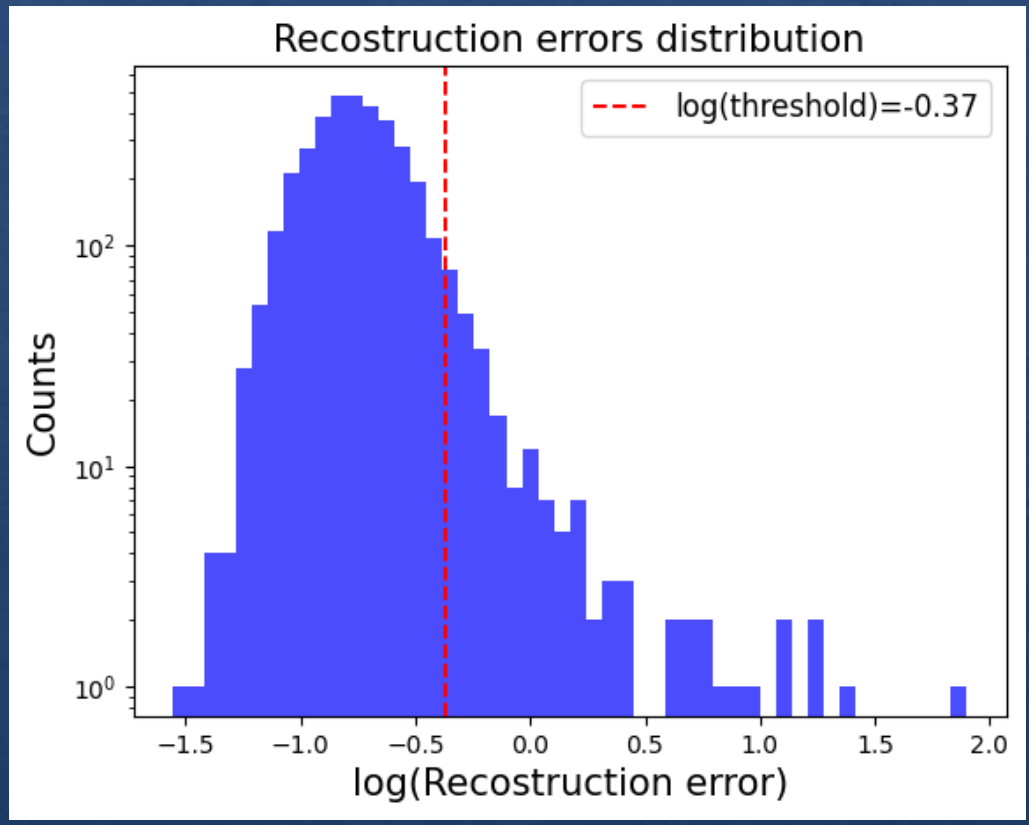

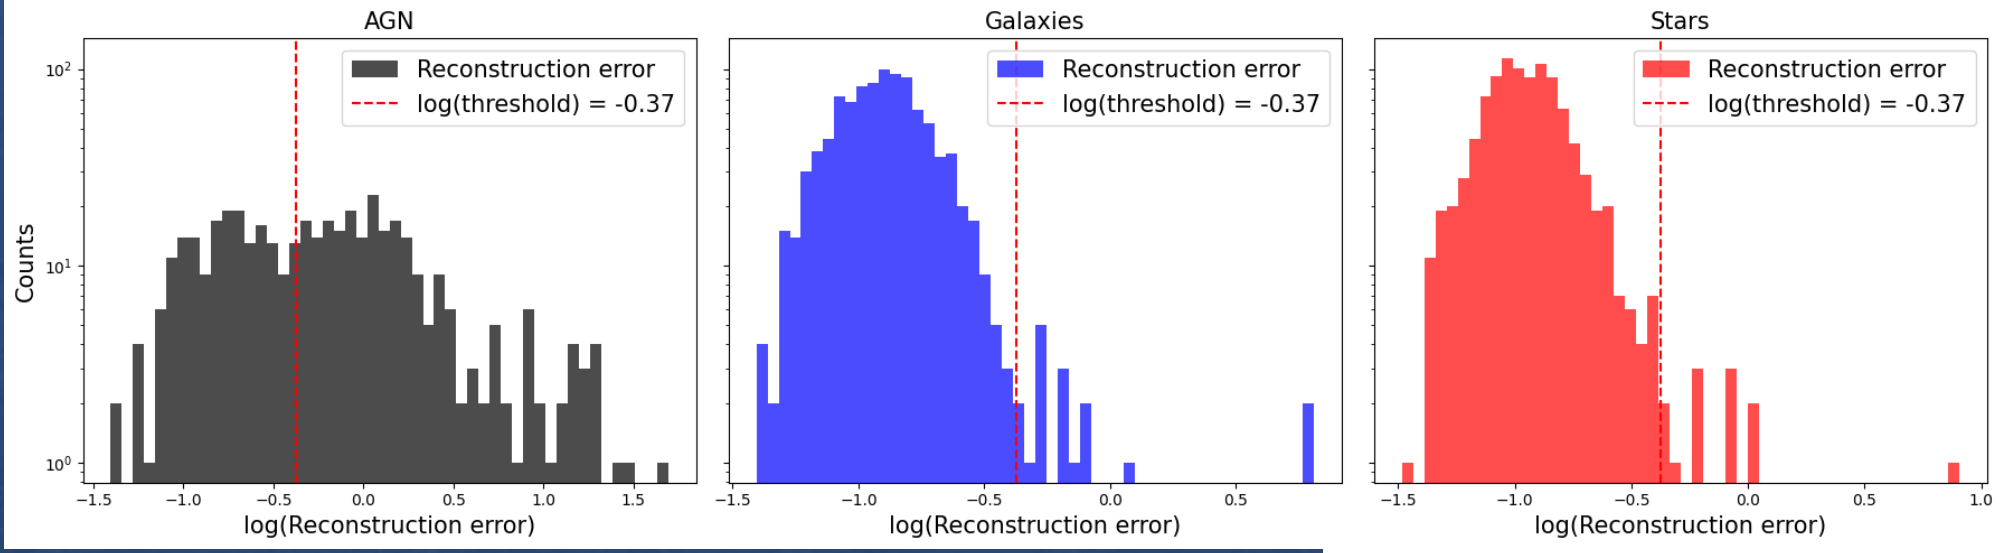

After the training on the US we made predictions on the LS.

11.8% of the sample is flagged as anomalous by the AE. Which is expected since the LS is AGN-enriched compared to the US and AGN (especially unobscured ones) are harder to reconstruct

Out of the 272 sources flagged as anomalous 242 are AGN, 16 galaxies and 14 stars

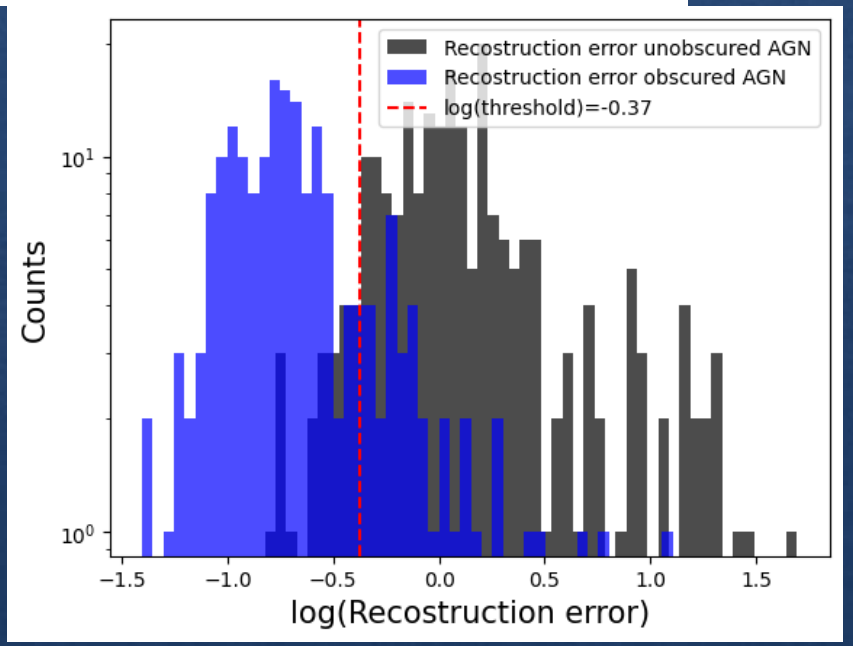

Out of the 242 AGN 204 are unobscured AGN, with only 38 obscured

Feature –driven results interpretation

While powerful, with unsupervised methods is hard to have a concrete explanation of why a specific point is being flagged as an anomaly.

And it is even more important with obscured AGN, since they should behave more or less like a normal galaxy in terms of their variability properties.

To add a layer of interpretability, we used the Shapley Additive exPlanations (SHAP) Python library.

SHAP assigns each feature a contribution value, known as a SHAP value, based on its marginal impact on the model's prediction.

To use SHAP, we need some kind of predictive model; to overcome this limitation, we trained a Light Gradient Boosting Regressor (LGBRegressor) by using as target values the reconstruction errors.

In this way, we can know the contribution of each feature to the RE observed

For each obscured AGN classified as anomalous we tested if the most influential feature was either on the tails of its own distribution (<5° percentile or >95° percentile) or if it was in the center

For the point for which the most influential feature was not in the tails we tested the second most influential and so on with the same criterion

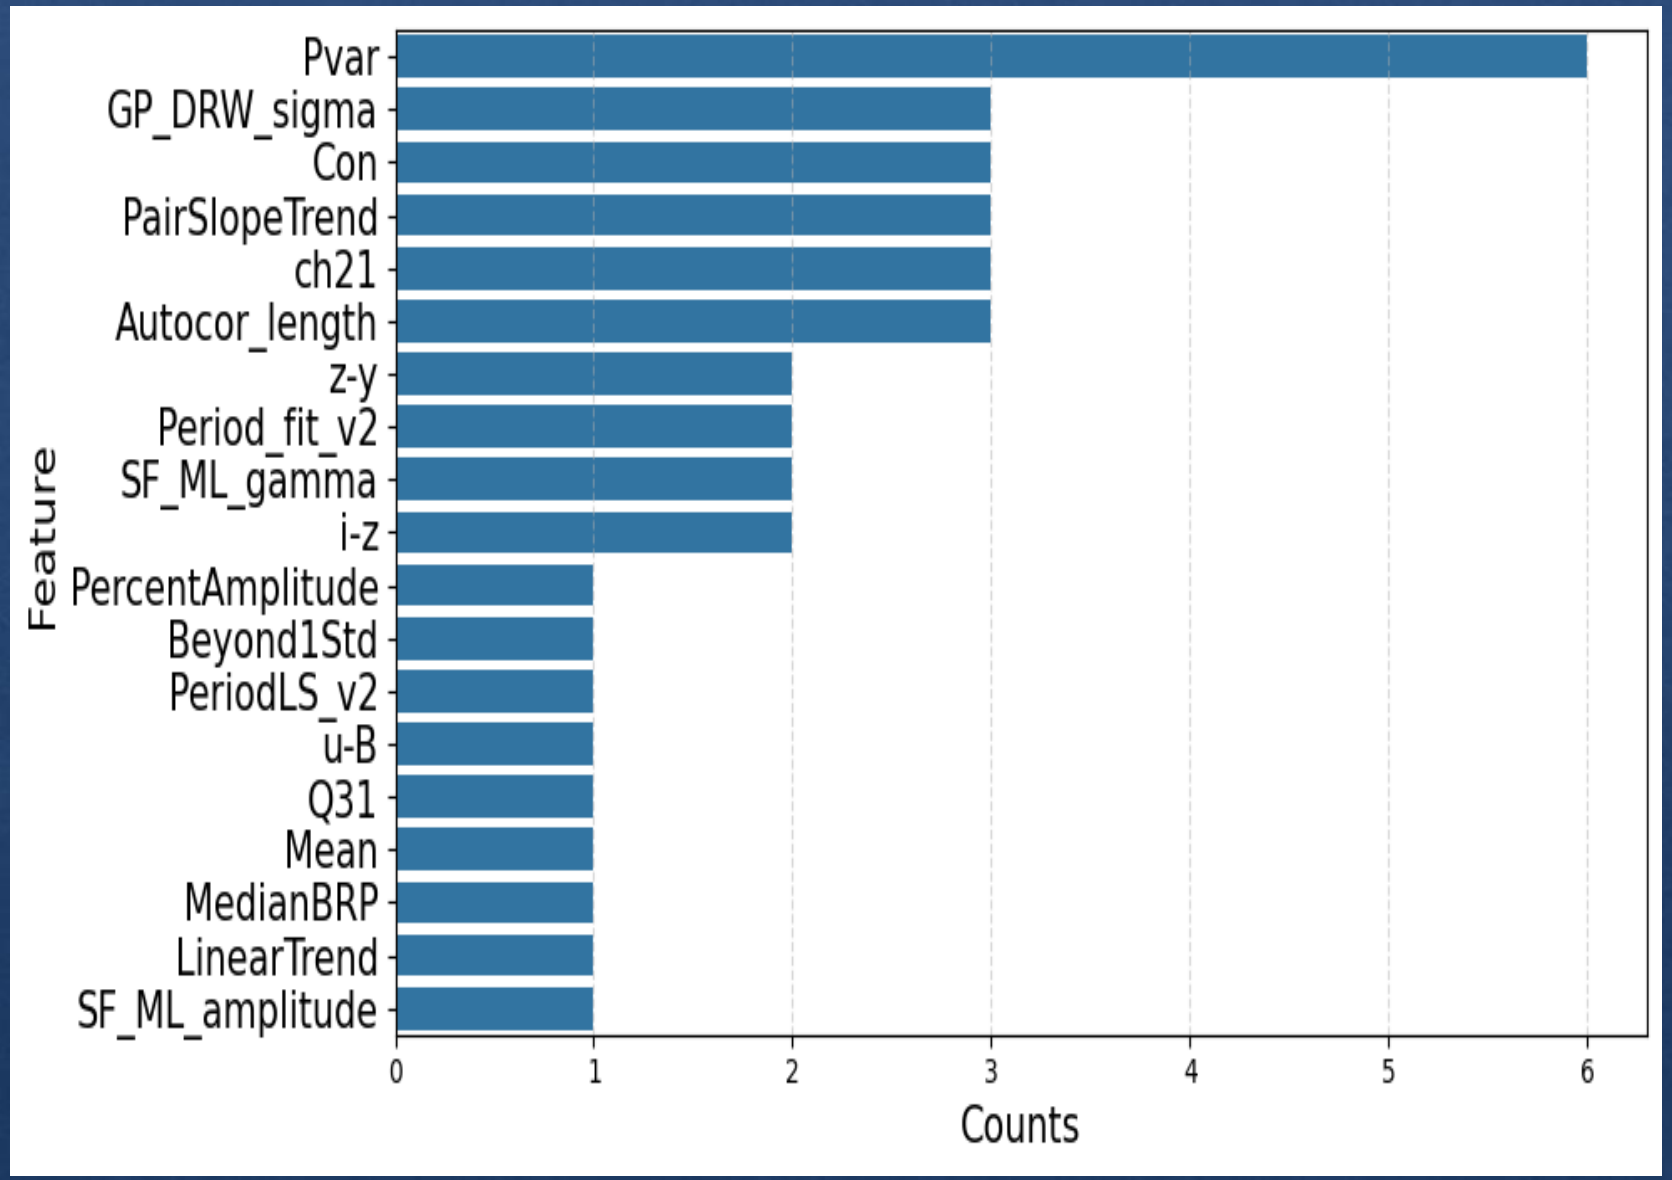

Out of the list obtained, to keep the features with the most discriminative power, we used the Kolmogorov-Smirnov (KS) test.

Only the features exhibiting statistically significant divergence (e.g., p-value < 0.05) were retained for further analysis.

From the starting 37 features only 12 survived to this stage

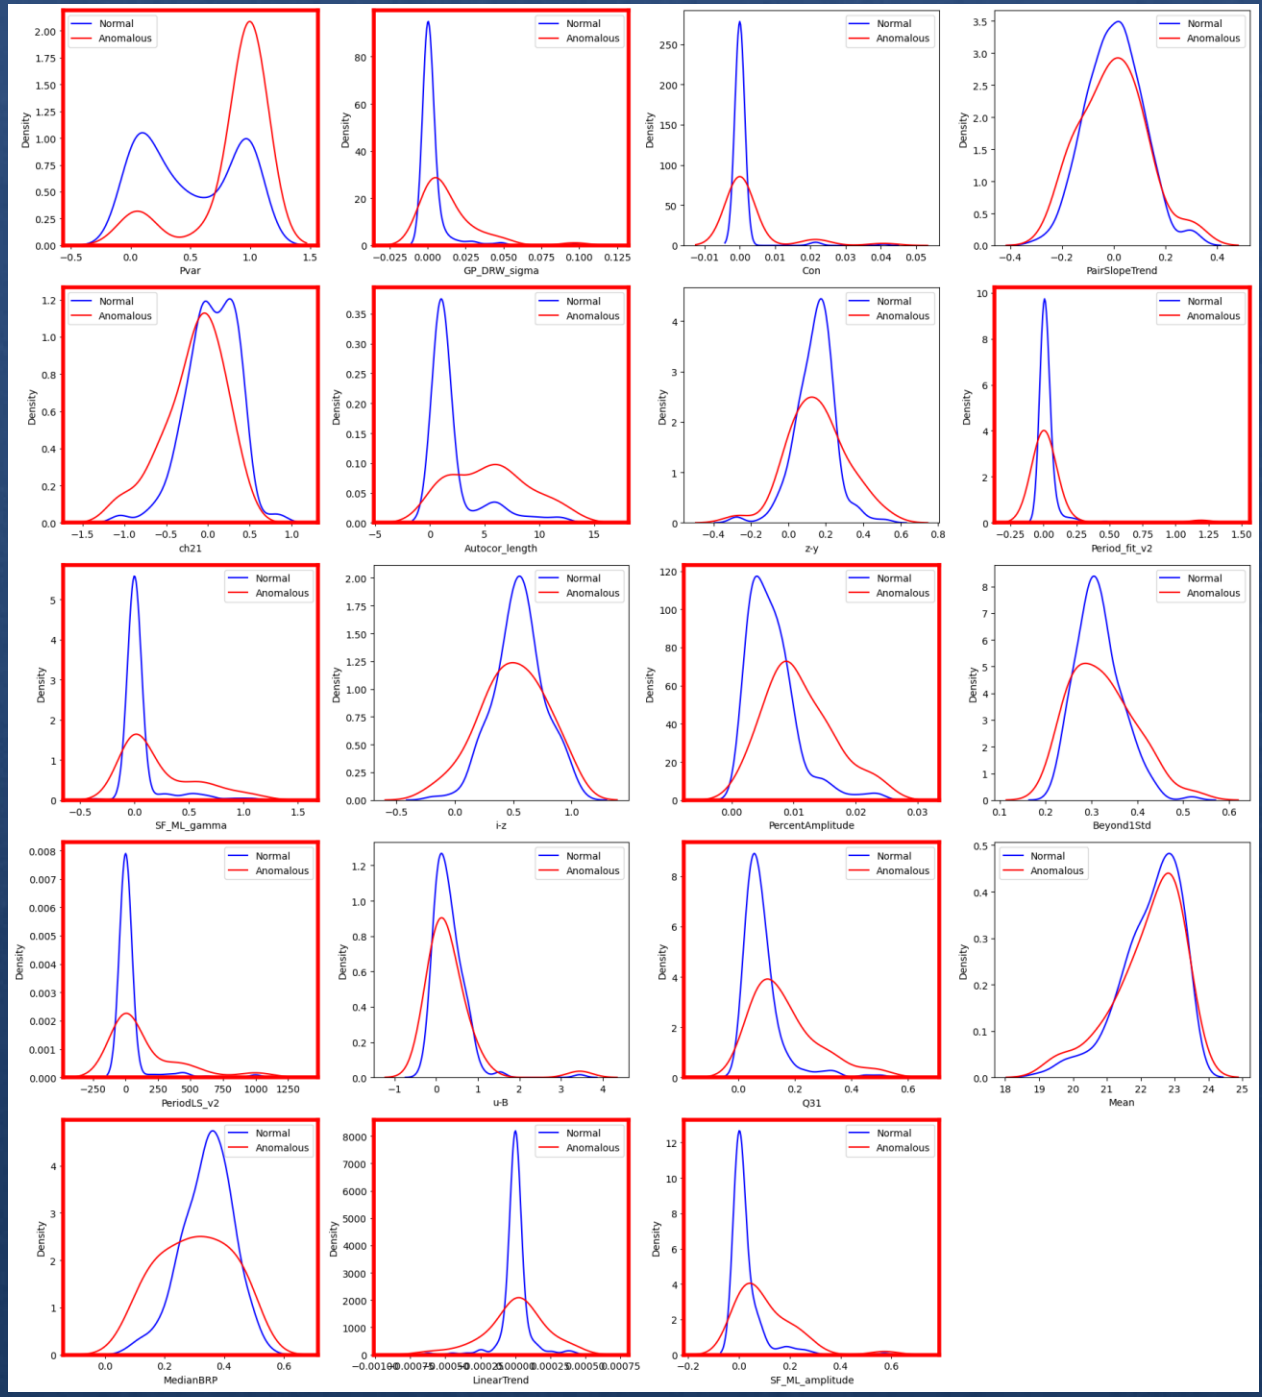

Comparing the feature sets

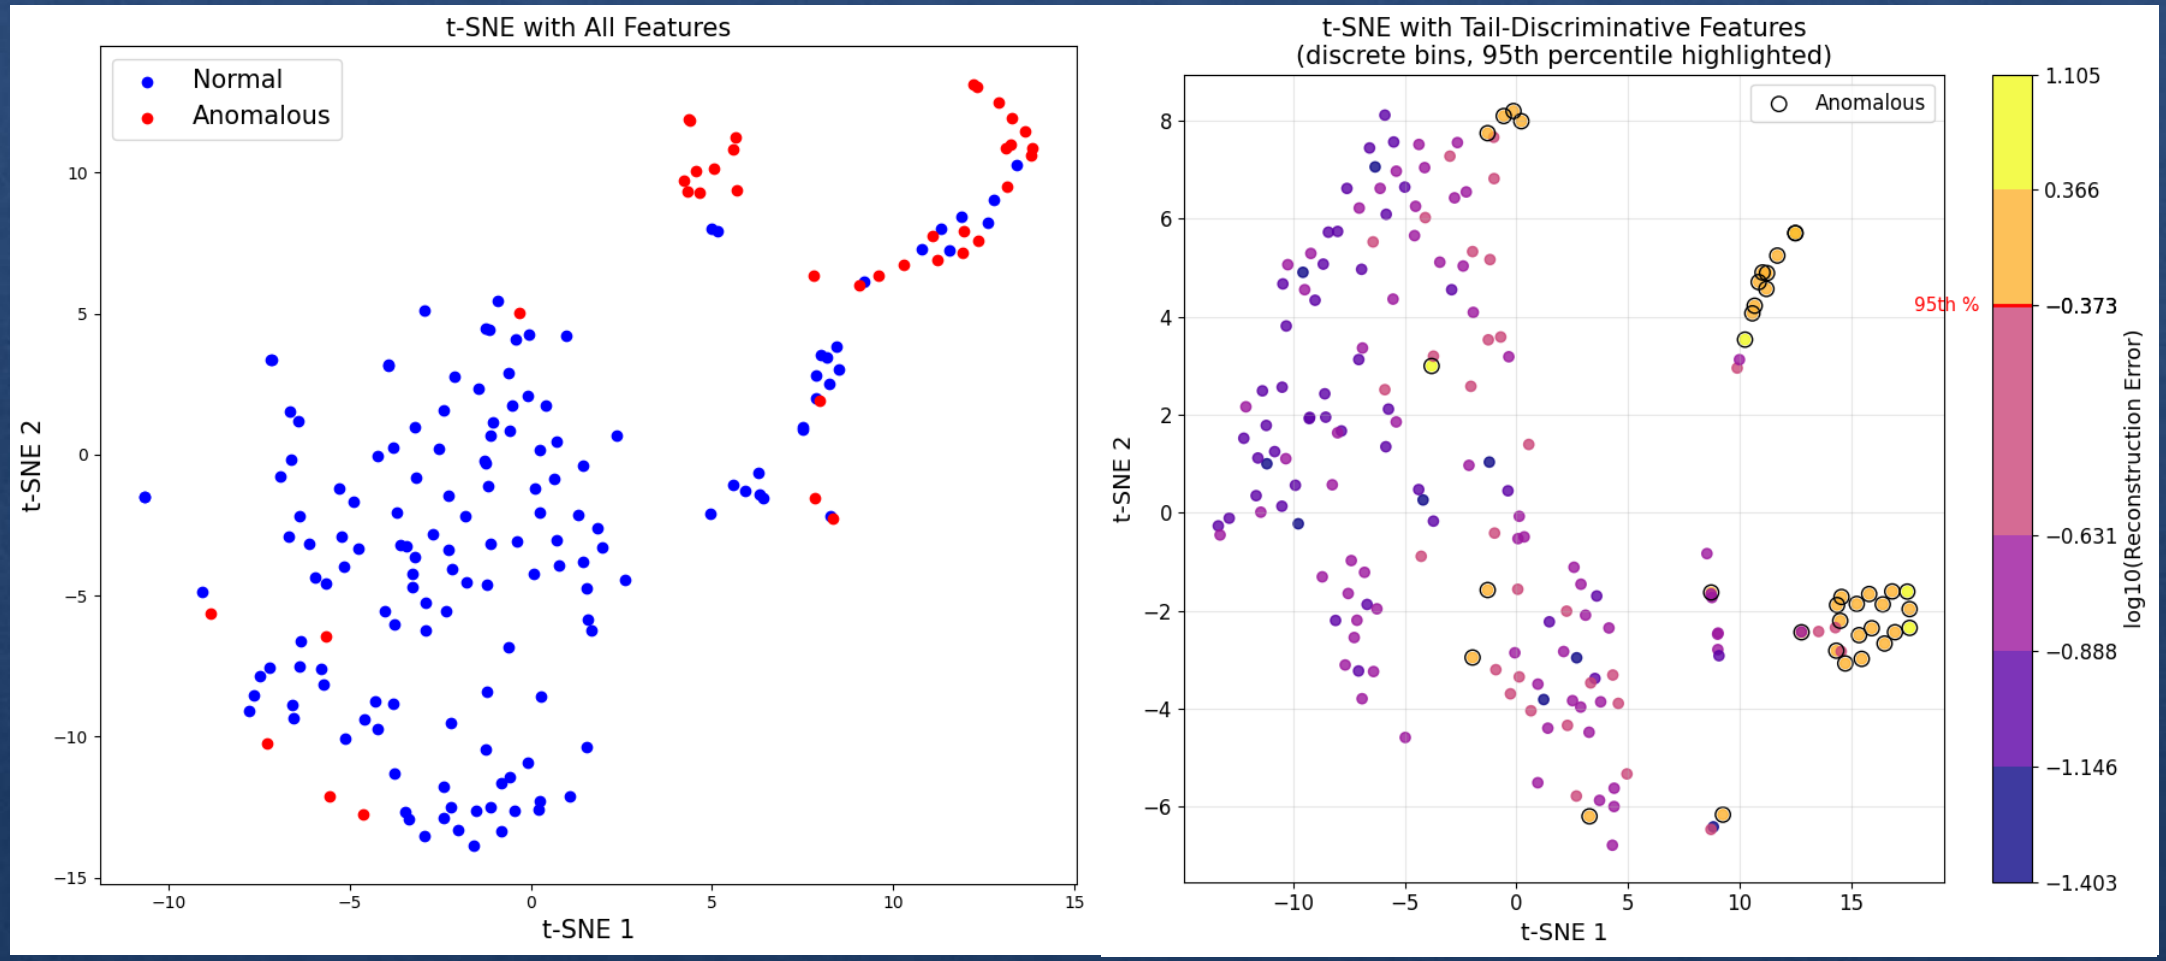

To test the discriminative power we made two different projection using t-SNE one with the whole feature set (left panel) one with the «most discriminative» set (right panal)

To quantify the actual improvement, we used two different classifiers: Random Forest Classifier (RFC) and a Gradient Boosting Classifier (GBC).

We trained these algorithms on both feature sets and evaluated performance using stratified 5-fold cross-validation.

This ensures more robust and unbiased performance estimates, specifically when having small or imbalanced datasets, such as in our case.

At each of the 5 iterations, we computed the following three metrics:

- **ROC AUC (Receiver Operating Characteristic Area Under the Curve):** quantifies the model's ability to distinguish between classes across varying decision thresholds.
- **Precision**: the proportion of true positive predictions among all predicted positives.
- **Recall**: the proportion of true positive predictions among all actual positive instances.

Model (Features)	ROC AUC	Precision	Recall
RFC (whole feature set)	0.962 ± 0.021	0.852 ± 0.128	0.686 ± 0.051
GBC (whole feature set)	0.952 ± 0.034	0.828 ± 0.103	0.764 ± 0.093
RFC (most discriminative feature set)	0.956 ± 0.023	0.843 ± 0.091	0.661 ± 0.044
GBC (most discriminative feature set)	0.943 ± 0.028	0.880 ± 0.062	0.736 ± 0.081

In the following, we will call tail points the obscured AGN for which the most discriminative feature is on the tail of its own distribution.

Out of the 38 obscured AGN flagged as anomalous, 14 of them are tail points.

Since these points deviate significantly from the expected behavior, we did an in-depth analysis of them.

6 of them are color-based anomalies, meaning that their most influential feature is the color

8 of them are variability-based anomalies, meaning that their most influential feature is a variability one

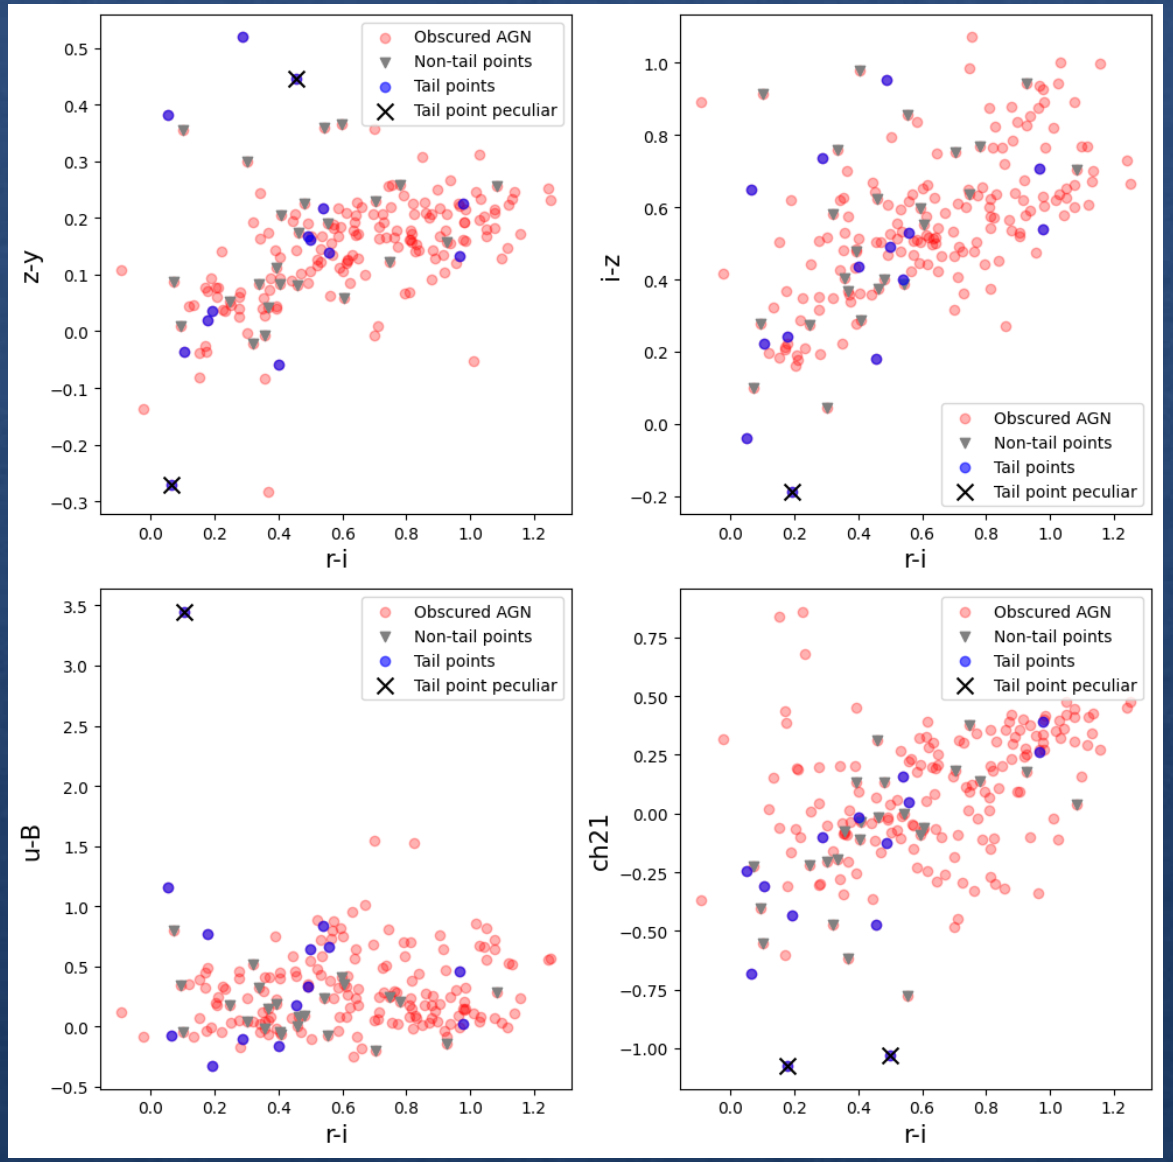

Vera C. Rubin Observatory is a brand-new astronomy and astrophysics facility on Cerro Pachón in Chile. It hosts the Simonyi Survey Telescope (8.4 meters aperture) and the world's largest digital camera, a 3.2-gigapixel instrument.

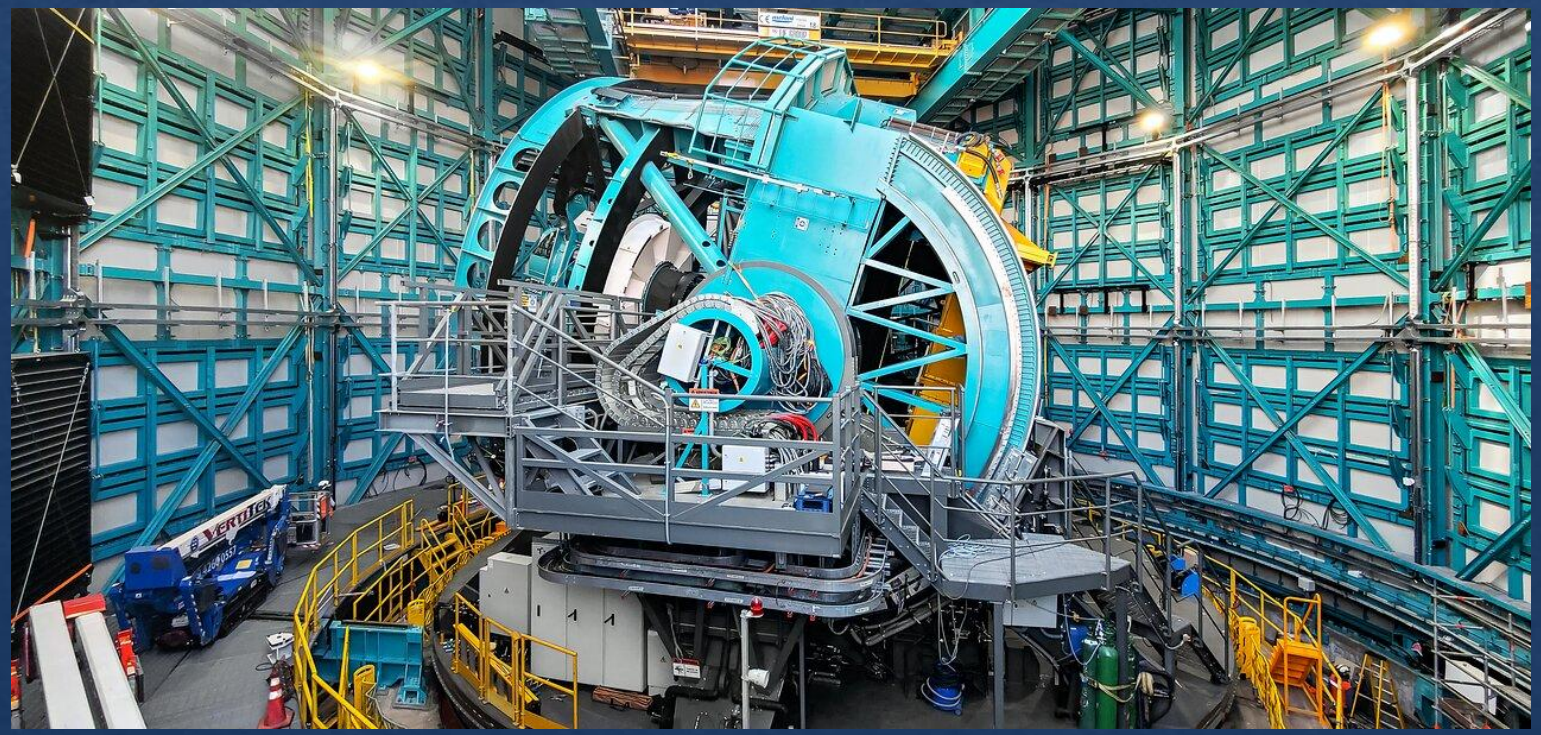

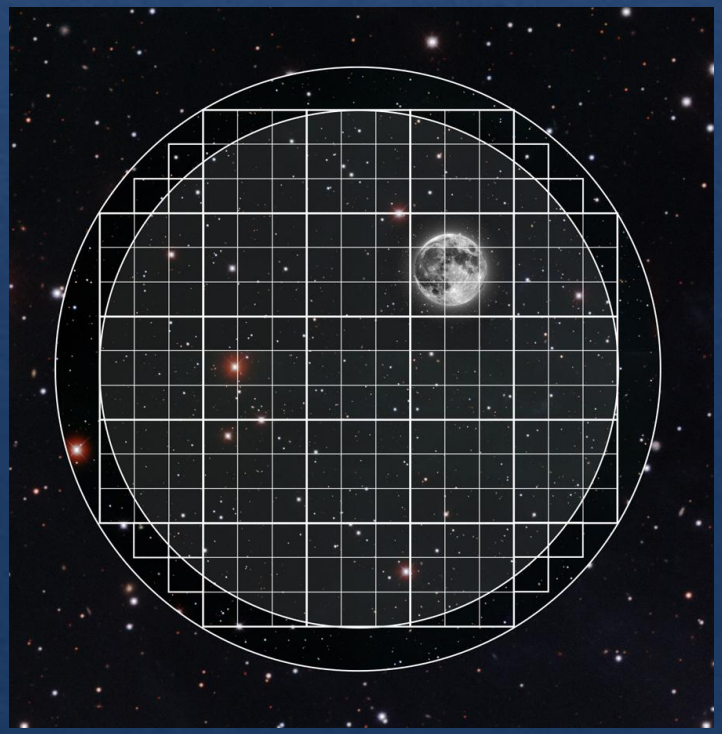

Vera C. Rubin Observatory will take hundreds of images of the Southern Hemisphere sky, every night for ten years, for a survey called the Legacy Survey of Space and Time (LSST).

The data from these images will be used by astronomers around the world to make countless discoveries, but Rubin Observatory was specifically designed to advance four science areas:

- Understanding the nature of dark matter and dark energy
- Creating an inventory of the Solar System
- Mapping the Milky Way
- Exploring the transient optical sky, i.e., studying objects that move or change in brightness (in which of course, AGN fall)

Rubin Observatory will produce about **20 terabytes** of data every night during the ten-year survey! (a lot of possibilities for scientific discovery!)

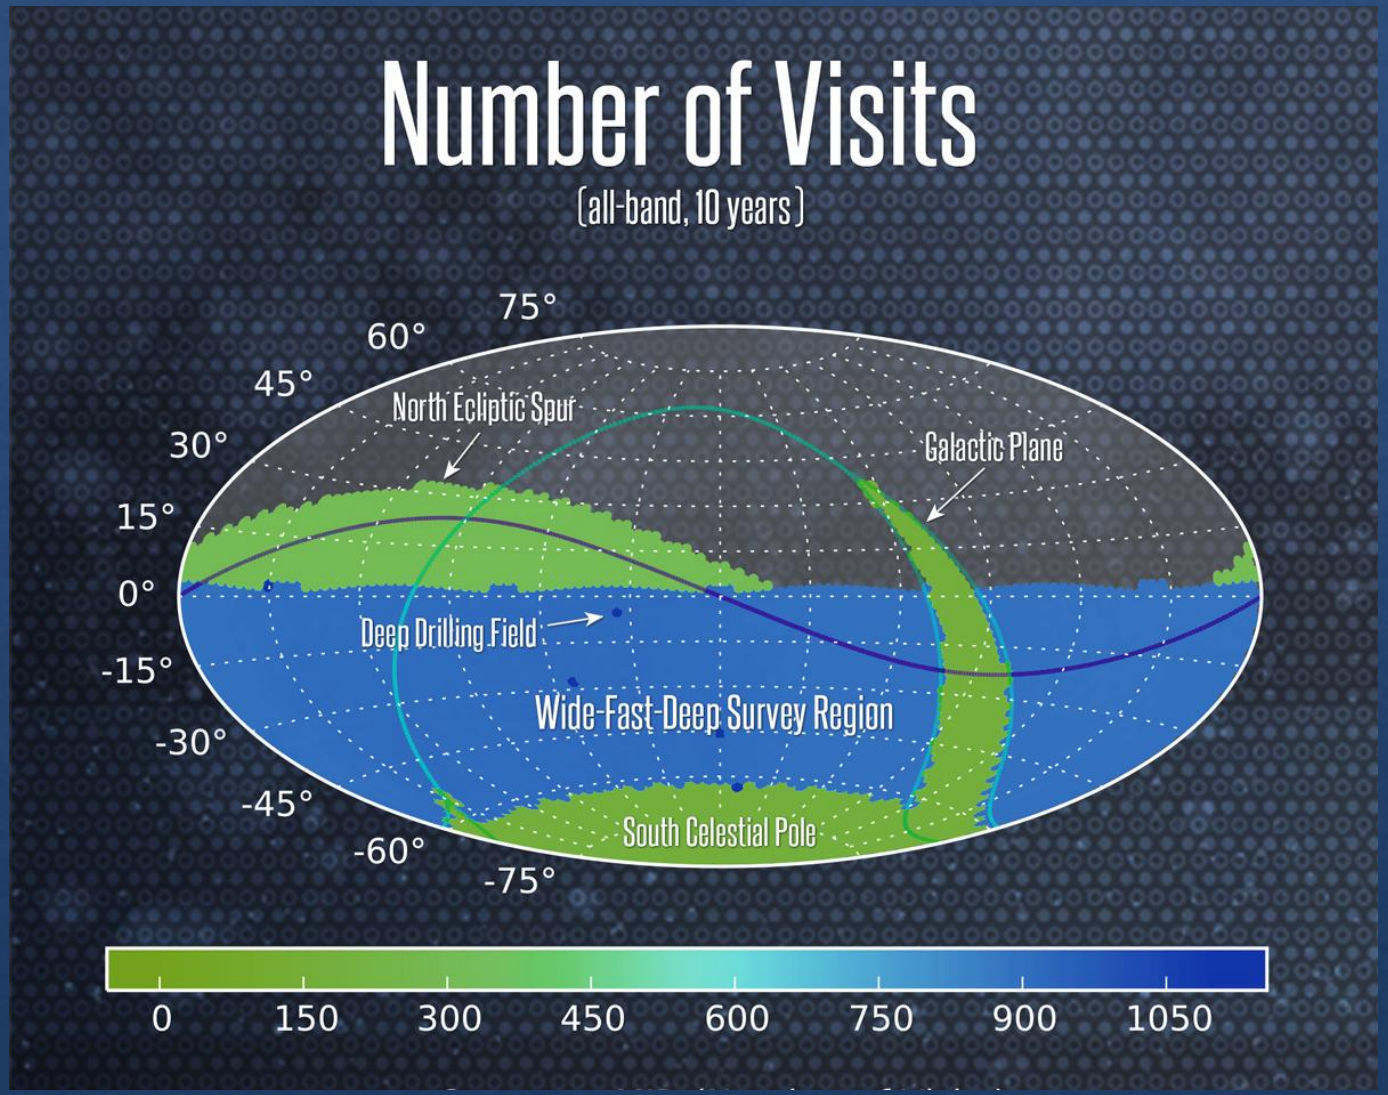

Future prospects: **TALES**

The **TALES** Marie Skłodowska-Curie Actions Doctorate Network is a consortium of 10 astrophysics research groups, 8 industrial and 3 academic partners spread across Europe that aims to study the feeding and feedback cycle of supermassive black holes. The 11 TALES candidates (I'm in there!) will:

- leverage time-domain astronomy observations from state-of-the-art facilities to map the environment in the close vicinity of supermassive black holes;
- use novel analysis methods from the discipline of data science to maximise the information gain from the observations;
- develop new theories and models to interpret the data and learn about the physics of the life-cycle of black holes at the centres of galaxies.

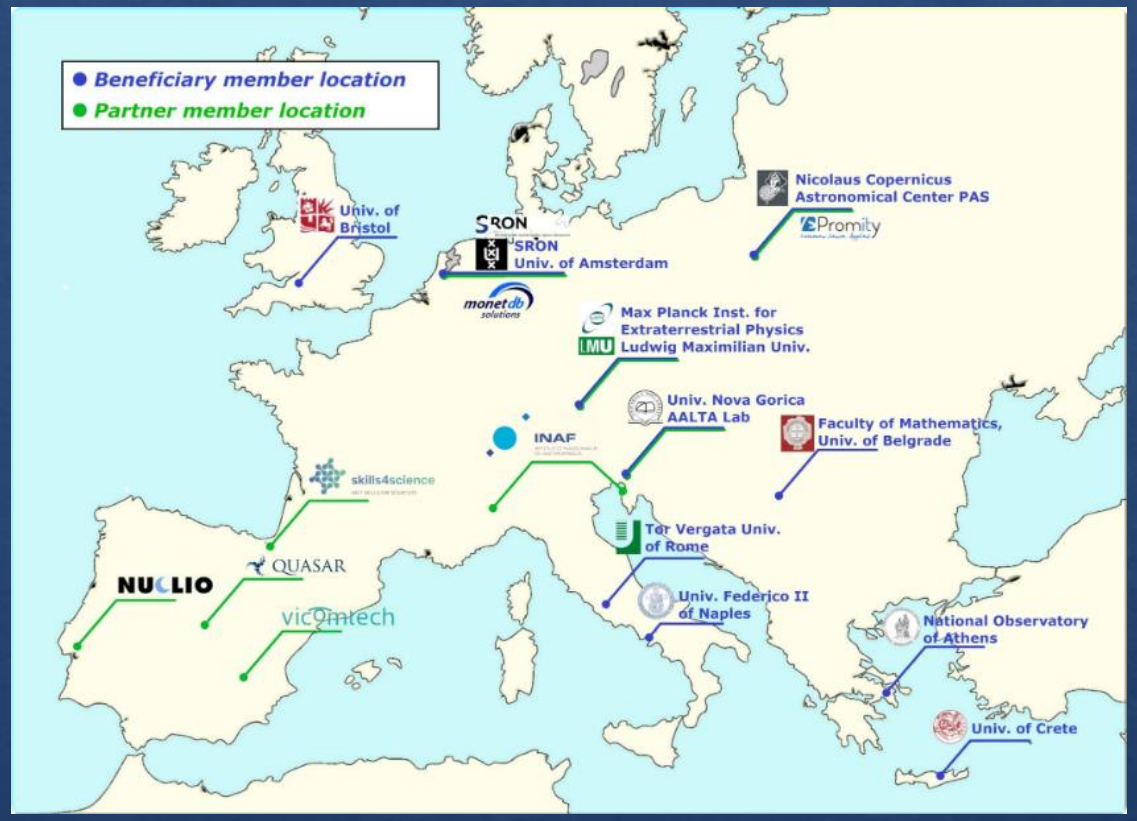

I'm starting my PhD here in Belgrade within the Tales project, and I'm going to be working on decomposing optical accretion disks through time-series modelling, under the supervision of Professors Ilić and Kovačević.

Thank you for the attention!

For any further questions, I'm here:

natale.de.bonis@matf.bg.ac.rs