



UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia "Galileo Galilei"

Master Degree in Physics

Final Dissertation

Modeling Blazar Broadband Emission with Machine Learning: Toward a Physical Interpretation of the Blazar Sequence

Thesis supervisor

Dott.ssa Elisa Prandini

Thesis co-supervisor

Dott.ssa Ilaria Viale

Candidate

Francesca Bovolon

Academic Year 2023/2024

A mamma, papà, e pure a Cate

Abstract

Blazars, a class of jetted active galactic nuclei, are the most numerous permanent extragalactic gamma-ray sources. Their peculiar double-bumped spectral energy distributions (SEDs) are usually interpreted as non-thermal emission from a relativistic jet of particles closely aligned with the line of sight. Population studies have highlighted a “blazar sequence”, i. e. an anticorrelation between the frequency of the low-energy peak and its bolometric luminosity. Its existence and origin are still unclear, despite the influx of new data, including in the TeV band. This work thus aims at contributing to a new sequence that finally includes very high-energy gamma-ray spectra. A number of representative SEDs from a sample of TeV-detected blazars of the “BL Lac” type, binned according to their low-energy peak frequencies, were modeled based on the standard “Synchrotron Self-Compton” scenario: best-fit parameters were compared to search for trends hinting at the mechanisms underlying the sequence. Different techniques, including analytical tools and machine learning, were used to characterize spectral quantities of the selected sources, and their outcomes and performances were discussed.

Contents

List of Figures	xi
List of Tables	xiii
List of Code Snippets	xvii
List of Acronyms	xix
1 Blazars: Observations	1
1.1 Active Galactic Nuclei	1
1.2 Blazars: Observational Properties	2
1.2.1 Radio Properties	2
1.2.2 Time Evolution	3
1.2.3 Spectral Energy Distribution	7
1.2.4 The Highest Energies	12
1.2.5 Classification	14
1.3 The Blazar Sequence	16
2 Blazars: Theoretical Models	21
2.1 The Spectral Model: Synchrotron Self-Compton	21
2.1.1 Synchrotron Emission	21
2.1.2 Synchrotron Self-Compton	23
2.1.3 Other Scenarios	26
2.2 The Physical Model, Part I: Unification	28
2.3 The Physical Model, Part II: Jets	31
2.3.1 Accretion and the Birth of Jets	31
2.3.2 Jet Structure	32
2.3.3 Acceleration Mechanisms	33
2.4 Physics of the Blazar Sequence	35

CONTENTS

3	Spectral Modeling	39
3.1	Aim: the TeV Blazar Sequence	39
3.2	Data Selection	42
3.2.1	Source Selection	42
3.2.2	Data Points Selection	47
3.3	Tools	54
3.3.1	agnpy	54
3.3.2	MMDC	56
3.4	Modeling	59
3.4.1	Implementation	59
3.4.2	Results	61
3.4.3	Final Considerations	76
4	Estimation of Spectral Quantities in Blazars with Machine Learning	81
4.1	Purpose	81
4.2	Algorithms	86
4.2.1	Random Forest	86
4.2.2	Gradient Boosted Decision Tree	86
4.2.3	Histogram-Based Gradient Boosted Decision Tree	87
4.3	Implementation	88
4.3.1	Tools	88
4.3.2	Data Selection	89
4.3.3	Data Engineering	92
4.3.4	Training	99
4.4	Results	101
4.4.1	Training and Evaluation	101
4.4.2	Predicting on New Data	108
4.5	Final Considerations	112
5	Conclusions	115
	References	117
	Acknowledgments	129
	Ringraziamenti	131
	Appendix	133

A Spectral Modeling Plots and Tables	135
A.1 Bin 4	135
A.1.1 MMDC Fits	135
A.1.2 agnpy Fits	137
A.2 Bin 3	139
A.3 Bin 5	139
B Machine Learning	141
B.1 Code Snippets	141
B.2 Plots	144

List of Figures

1.1	AGN observational classification	2
1.2	PG 1218+304 light curves	4
1.3	PKS 2155-304 light curves	6
1.4	PG 1218+304 archival spectral energy distribution	8
1.5	<i>Fermi</i> -LAT sensitivity and BL Lac detection	11
1.6	Extreme blazars	13
1.7	3C 454.3 spectral energy distribution	15
1.8	Blazar sequence, 1998	17
1.9	Blazar sequence, 2017	18
2.1	Synchrotron self-Compton	25
2.2	Spectra and SSC model parameters	26
2.3	Peak ν and νF_ν against δ_D	27
2.4	Unified models of AGN	29
2.5	Unified scheme, radio-loud	30
2.6	Unification with radio galaxies	30
2.7	Jet model	34
2.8	Compton dominance sequence	36
3.1	<i>Fermi</i> blazar sequence for TeV sources	41
3.2	4LAC-DR2 blazar ν_{syn} frequency distribution	43
3.3	TeV BL Lac ν_{syn} frequency histogram and binning	44
3.4	Observed SEDs	44
3.5	Bin 4 redshift histogram	45
3.6	SEDs at different redshifts	46
3.7	WISE	47
3.8	Swift	48
3.9	BeppoSAX	48
3.10	<i>Fermi</i> -LAT	49
3.11	VHE gamma-ray detection	49

LIST OF FIGURES

3.12	Example of SED data selection	53
3.13	agnpy flowchart	55
3.14	Neural network structure	57
3.15	PG 1218+304 MMDC fit, free parameters	63
3.16	PG 1218+304 fit parameters vs. δ_D	66
3.17	PG 1218+304 agnpy/Sherpa fits, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, $\delta_D = 20, 25$	68
3.18	PG 1218+304 agnpy/Gammapy fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, $\delta_D = 20$	68
3.19	PG 1218+304 agnpy/Gammapy covariance matrix and χ^2 profiles	69
3.20	PG 1218+304 agnpy/Gammapy fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, $\delta_D = 20$	70
3.21	PG 1218+304 models comparison	71
3.22	PKS 2155-304 agnpy/Sherpa fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, $\delta_D = 20$	73
3.23	PKS 0548-322 agnpy/Sherpa fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^7$, $t_{\text{var}} = 1$ d, $\delta_D = 20$	75
4.1	K-Folds	83
4.2	Cross-validation flowchart	84
4.3	Random Forest vs. Gradient Boosting	88
4.4	Covariance matrix of features and target	91
4.5	Correlation coefficients bar chart	92
4.6	ν_{syn} , linear scale	93
4.7	ν_{syn} , logarithmic scale	95
4.8	ν_{syn} , logarithms + power transform	96
4.9	Box plots of features and target	97
4.10	Scatter plots of features against target	98
4.11	Frequency density histogram, F100000, training vs. new data	99
4.12	Performance metrics vs. fit times	102
4.13	Predictions vs. actual values: scatter plot (HGB)	105
4.14	Predictions vs. actual values: frequency histograms (HGB)	106
4.15	Predictions vs. actual values: line plot (HGB)	106
4.16	Predictions vs. actual values: residuals (HGB)	107
4.17	Feature importances	108
4.18	Frequency histograms, new predictions vs. BlaST	110
4.19	Frequency histograms, new predictions vs. old data	111
A.1	PG 1218+304 MMDC fit, $\gamma_{\min} = 10^2$	135
A.2	PG 1218+304 MMDC fit, $\gamma_{\min} = 10^3$	136
A.3	PG 1218+304 MMDC fit, $\gamma_{\min} = 10^3$, Swift-XRT data	136
A.4	PG 1218+304 agnpy/Sherpa fits, alternative best fits	138

A.5	PKS 2155-304 MMDC fit, $\gamma_{\min} = 10^3$	139
A.6	PKS 0548-322 MMDC fit, $\gamma_{\min} = 10^3$	139
B.1	Distributions of features and target, linear scale	144
B.2	Predictions vs. actual values: scatter plot (RF)	145
B.3	Predictions vs. actual values: scatter plot (GB)	146
B.4	Predictions vs. actual values: frequency histograms (RF)	147
B.5	Predictions vs. actual values: frequency histograms (GB)	148

List of Tables

3.1	SEDs and redshift	46
3.2	Instruments	51
3.3	PG 1218+304 MMDC fits	62
3.4	PG 1218+304 agnpy/Sherpa fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$	64
3.5	PG 1218+304 agnpy/Sherpa fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d	65
3.6	PG 1218+304 agnpy/Sherpa fit, $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, fixed δ_D	67
3.7	PG 1218+304 peaks	72
3.8	PKS 2155-304 best-fit parameters	73
3.9	PKS 0548-322 best-fit parameters	75
3.10	Fit summary, BeppoSAX selection, agnpy/Sherpa	77
3.11	Fit summary, BeppoSAX selection, MMDC	77
4.1	Features and target	91
4.2	Skewness and kurtosis of features and target	94
4.3	Performance metrics	103
4.4	Mean performance metrics, train vs. test	104
4.5	Predictions vs. actual values: distribution statistics	105
4.6	New dataset: predictions and BlaST comparison	109
A.1	PG 1218+304 agnpy/Sherpa fit, $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$	137
A.2	PG 1218+304 agnpy/Sherpa fit, $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d	137
A.3	PG 1218+304 agnpy/Sherpa and MMDC fits, $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, fixed δ_D	138

List of Code Snippets

B.1	Dictionaries of implemented models and candidate hyperparameters. . .	141
B.2	Model cross-validation and training loops.	142

List of Acronyms

BL Lac BL Lacertae-type blazar

AGN Active Galactic Nucleus/Nuclei

FSRQ Flat Spectrum Radio Quasar

VHE Very High Energy/Energies

SED Spectral Energy Distribution

IR Infrared

EBL Extragalactic Background Light

LBL Low-synchrotron-peaked BL Lac

IBL Intermediate-synchrotron-peaked BL Lac

HBL High-synchrotron-peaked BL Lac

EHBL Extreme High-synchrotron-peaked BL Lac

SSC Synchrotron Self-Compton

IACCT Imaging Air Cherenkov Telescope

MSE Mean Squared Error

rMSE root Mean Squared Error

1

Blazars: Observations

When BL Lacertae was first detected, at the Sonneberg Observatory in 1929, it was mistaken for a faint variable star, and named according to one of the conventions for that type of object. Almost 40 years would pass before astronomers began to realize it was actually a peculiar non-stellar source for its long series of distinctive features, from the strong optical and radio flux variability on a timescale of days, to the changing polarization in both frequency bands, to the apparently total absence of lines in the optical spectrum. In 1972, these characteristics led researchers to suggest the existence of the “BL Lacertae objects” (Stein, O’Dell, and Strittmatter 1976), or “lacertids”, now shortened to “BL Lacs”: a new class of “quasi-stellar objects”, similar to the group discovered 9 years before when 3C 273 was found to have a redshift and a luminosity too high to be a variable star (Schmidt 1963). Now this other group (the “flat-spectrum radio quasars”, or FSRQs) and the BL Lacs are collectively called “blazars”, and are the under intense scrutiny for being among the most luminous and energetic sources in the Universe, the most numerous population of high-energy gamma ray emitters, and a natural laboratory to investigate on extreme acceleration processes and open issues in cosmology and astroparticle physics.

1.1 ACTIVE GALACTIC NUCLEI

It is not only the redshifts, but also the luminosity¹ of blazars that put these sources outside the realm of variable stars: with values that may approach $10^{48} \text{ erg s}^{-1}$, they cannot be explained by thermonuclear processes (by comparison, a reference luminosity for “normal” galaxies is $10^{11} L_{\odot} \sim 10^{44} \text{ erg s}^{-1}$, with $L_{\odot} = 3.828 \times 10^{33} \text{ erg s}^{-1}$ the luminosity

¹ L_{ν} , which is the *emitted* energy per unit time at a specific frequency ν .

1.2. BLAZARS: OBSERVATIONAL PROPERTIES

of the Sun). The “standard model” therefore invokes radiative processes caused by a accreting supermassive black holes at the center of galaxies (see e.g. Ghisellini 2013; Section 2.2). In this context, blazars are only the 10% of a wider class, *active galactic nuclei* (AGN), that includes a “zoo” of different objects that are, in turn, $\sim 1\%$ of all galaxies, and take different names depending on which frequency band they are observed in (and, arguably, on the researcher’s opinion; see Padovani et al. 2017 for an overview of the nomenclature). Blazars are emblematic gamma-ray selected sources, and they seem to share a very close link to another category of AGN, *radio galaxies*, in both being loud radio sources whose emission is mainly non-thermal in origin, and located principally in external structures (jets) originating from the core, rather than in the nucleus itself. Other AGN, instead, do not feature jets, or at most weak ones; their spectrum is dominated by thermal processes, and their radio emission is weak (“non-jetted” and “radio-quiet” AGN, like optical quasars and Seyfert galaxies). Figure 1.1 gives one example of simplified AGN classification, including BL Lacs and FSRQs, based on observational properties.

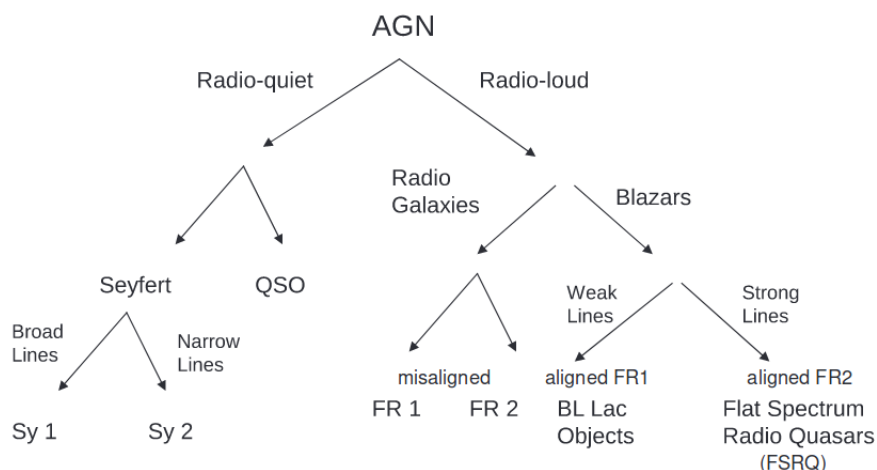


Figure 1.1: Classification of AGN from observational properties, from Dermer and Giebels 2016.

1.2 BLAZARS: OBSERVATIONAL PROPERTIES

1.2.1 RADIO PROPERTIES

The optical band was the window through which BL Lacertae was first discovered, but it was radio observations that led to the identification of the other prototype, 3C 273 and gave the definitive clues to single out the category of radio-loud quasars first, and then of “blazars” once coupled with detections at high energies (Dermer and Giebels 2016). Blazars are characterized by a strong radio emission (“radio-loud” objects) that

usually appears variable in flux and polarization; the spectrum is flat, appearing constant when considering the spectral flux density² $F_\nu \propto \nu^{-\alpha_r}$ with $\alpha_r \sim 0$. The development of very long-baseline radio interferometry in the 1970s, with a resolution of 10^{-4} arcseconds, allowed to detect the existence of individual blobs of enhanced emission, that at times look like they are moving at *superluminal* speed. Far from disproving special relativity, they are actually a spectacular display of relativistic kinematics. Indeed, if a source of electromagnetic radiation is moving toward the observer with a bulk speed V and a very small angle θ , it will have an apparent transverse speed of

$$v_{\text{app}} = \frac{V \sin \theta}{1 - \frac{V}{c} \cos \theta}$$

which peaks for $\cos \theta = \frac{V}{c}$ yielding $v_{\text{app}} = \Gamma V$, with $\Gamma = \frac{1}{\sqrt{1 - \frac{V^2}{c^2}}}$, the bulk Lorentz factor of the radio blob. From this relation, the *Doppler factor*, or *beaming factor*, can be defined,

$$\delta_D = \frac{1}{\Gamma \left(1 - \frac{V}{c} \cos \theta\right)} \quad (1.1)$$

Powers of the beaming factor scale relevant physical quantities between the observer to the blob reference frame, and is therefore a fundamental parameter in blazar modeling. Superluminal motions, being a relativistic effect that well explains radio observations, are a kind of “smoking gun” on the nature of blazars: collimated jets of particles moving at relativistic speeds, observed at a small angle ($\theta \leq \frac{1}{\Gamma}$) from their axis.

1.2.2 TIME EVOLUTION

Blazars were defined by their **strong variability** from the very beginning, when the first sources of their kind were detected in the radio (Dermer and Giebels 2016). Characterizing these sources in the time domain is fundamental, as it allows to **put constraints on the maximum size of the emission region in a given energy band**, becoming a way to inspect the structure of an otherwise unresolved object (Ulrich, Maraschi, and Urry 1997, Spurio 2018). Given the minimum measured variability timescale, t_{var} , and the Doppler factor δ_D , the maximum size of the emission region R in its reference frame will be given by considering t_{var} the *light-crossing time*:

$$R = \frac{c \delta_D t_{\text{var}}}{1 + z} \quad (1.2)$$

²Energy of radiation at a specific frequency crossing a unit surface in a unit time. In terms of the luminosity L_ν of a source at a distance r , the spectral flux density is $F_\nu = \frac{L_\nu}{4\pi r^2}$.

1.2. BLAZARS: OBSERVATIONAL PROPERTIES

such that both the cosmological redshift and the apparent contraction of variability timescales due to relativistic beaming are accounted for. In fact, blazars are known to show **variability on more than one timescale**, depending on the energy band in which time analysis is done: apart from potential **intrinsic differences** in the emission region and variability mechanism at different frequencies, there is an **impact of the instrument performance and observation strategy**, as exposure or integration times may prevent a detector to catch flux variations on “incompatible” timescales. Indeed, blazar *light curves*, plots that show the measured photon flux as a function of time, need to be interpreted as a realization of the **underlying stochastic process as it is sampled by the instrument** (Rieger 2019). One of the bands where variability is best known is high-energy gamma rays: three examples referring to the same blazar (TXS 0506+056) are shown in Figure 1.2, where light curves are built from data taken by the *Fermi* Large Area Telescope (see Section 3.2) in the 100 MeV-100 GeV range over 16 years of activity, integrating measurements over 3 days, one week, or one month. While the first two plots




Figure 1.2: Light curves with different time samplings of the blazar TXS 0506+056, obtained with the *Fermi* LAT Light Curve Repository (Abdollahi et al. 2023, <https://fermi.gsfc.nasa.gov/ssc/data/access/lat/LightCurveRepository/>) in the 100 MeV-100 GeV energy range. Each plot shows the photon flux using: 1.2a time bins of 3 days (August 23, 2024), 1.2b one week (August 8, 2008 to September 20, 2024), 1.2c one month (August 20, 2008 to August 26, 2024). The photon index is left free in spectral fitting. The minimum detection threshold is equivalent to 2σ ; upper limits have been hidden for clarity.

reveal evident variation in the photon flux already in the span of few days, the third one shows that in the last 16 years, the source underwent longer-term periods of enhanced activity, including a very prominent one between 2017 and 2018 (Tanaka, Buson, and Kocevski 2017) where the flux was measured to be about 6 times the catalogued one. These apparently random *flares* are triggers of multiwavelength and, lately, multimessenger observation and data analysis campaigns (The IceCube Collaboration et al. 2018), that often reveal a **coherent increase in flux and hardness on the whole electromagnetic spectrum, hinting at a common emission mechanism**. These coordinated efforts are made possible by alert networks between observatories, like the General Coordinates Network³ and the Astrophysical Multimessenger Observatory Network (AMON:

³<https://gcn.nasa.gov/>

Smith et al. 2013).

Variability Indices

In the specific case of *Fermi*-LAT observations, source variability (but even more, the likelihood of source variability) is described in catalogues using two statistical quantities, that will become relevant in Chapter 4: the variability index,

$$TS_{\text{var}} = 2 \sum_i \log \left[\frac{\mathcal{L}_i(F_i)}{\mathcal{L}_i(F_{\text{glob}})} \right] - \max(\chi^2(F_{\text{glob}}) - \chi^2(F_{\text{av}}), 0)$$

$$\chi^2(F) = \sum_i \frac{(F_i - F)^2}{\sigma_i^2} \quad (1.3)$$

and the fractional variability,

$$\text{Var} = \frac{1}{N_{\text{int}} - 1} \sum_i (F_i - F_{\text{av}})^2$$

$$\delta F = \sqrt{\max\left(\text{Var} - \frac{\sum_i \sigma_i^2}{N_{\text{int}}}, 0\right)} \quad (1.4)$$

$$\frac{\sigma_F}{F} = \max\left(\frac{1}{\sqrt{2(N_{\text{int}} - 1)}} \frac{V_i}{F_{\text{av}} \delta F}, 10\right)$$

such that the fractional variability is $\frac{\delta F}{F}$. F is the assumed flux, F_i is the individual flux value in the i -th time bin (2 months or one year), σ_i is its error, F_{av} is the average flux from the light curve, F_{glob} is the overall flux obtained from the global analysis of the source, N_{int} is the number of time bins, $V_i = (F_i - F)^2$ (Abdollahi et al. 2022)^a. The variability index is then a likelihood ratio of yearly or bi-monthly measured variability compared to the null hypothesis of stationarity and may not reflect actual variability for more “confused” sources. The fractional variability is instead a measure of the excess variance in each time bin on top of statistical and systematic uncertainties.

^aEducated guess. The 4FGL paper seems to have forgotten to define V_i .

The most debated results come from the detection of short-term variability at the highest energies. As it will be explained later more in detail, telescopes covering the TeV energy range are sensitive enough to measure flux variations during a single observation night. Variations over just 2 to 3 min have been observed at very high energies (VHE: $E_\gamma > 100$ GeV) for a number of sources (Rieger 2019 and references therein). Figure 1.3a depicts the intra-night light curve of the BL Lac blazar PKS 2155-304. The fact that some sources are not only able to produce such energetic photons, but with such

1.2. BLAZARS: OBSERVATIONAL PROPERTIES

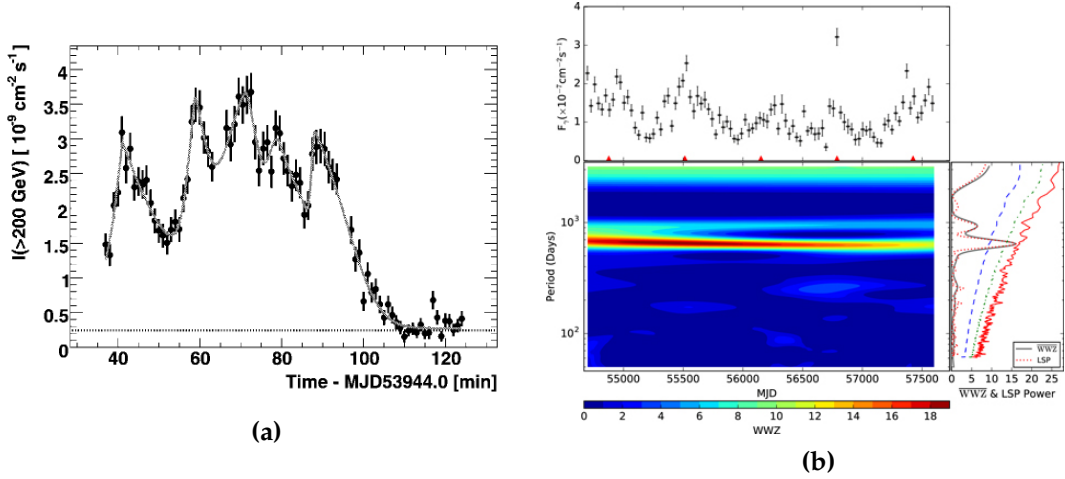


Figure 1.3: Light curves of the BL Lac object PKS 2155-304 in different energy ranges and on different timescales. 1.3a: intra-night VHE light curve measured by H.E.S.S. during a flare on July 28, 2006, from Aharonian et al. 2007. 1.3b: 8-year light curve (2008-2016) measured by *Fermi*-LAT above 100 MeV with 30-day binning, and variability power spectra computed with two different methods, from Zhang et al. 2017.

fast significant flux fluctuations, imposes the strictest constraints on the emission region and on its models, as explained above. For example, taking a timescale of ~ 1 h, the characteristic length scale of the central engine powering the blazar should be comparable to the Schwarzschild radius of a black hole of $10^9 - 10^{10} M_{\odot}$ (Spurio 2018) and lie close to the AGN core. Given that the optical depth of pair production, $\tau_{\gamma\gamma} \sim \frac{L}{R}$ would be large, making the emission region opaque to high-energy radiation, the ability to measure a significant gamma-ray flux from a luminous compact source confirms the relativistic beaming effect encoded in the δ_D parameter (Urry and Padovani 1995). This, however, still does not explain the origin of the shortest variability timescales completely, and a number of models for the source structure and radiative processes have been put forward to reconcile the observations with other quantities found independently: for example, the fact that there frequently seems to be a correlation between the variability in the gamma-ray and radio or optical band, suggesting an emission region a few parsecs away from the active core, or the existence of VHE “orphan flares” without a lower frequency counterpart (Dermer and Giebels 2016, Spurio 2018).

On the other side of the variability spectrum, multi-year trends have been found in some TeV-detected blazars, thanks to the continuous monitoring at high energies provided by *Fermi*-LAT. In particular, there are a few sources, like the aforementioned PKS 2155-304 and PG 1553+113, where potential quasi-periodic oscillations have been detected (Rieger 2019 and references therein; figure 1.3b).

Variability is not limited to fluxes: **polarization** is another feature of blazar emission that shows variability over different energy ranges, sometimes correlating to flux variations (e.g. Tavecchio 2021 and references therein), and that has been recently under the

spotlight since the launch of the Imaging X-ray Polarimetry Explorer (IXPE) in late 2021, with the opening of the high-energy window in blazar polarimetry studies (e.g. Kouch et al. 2024).

1.2.3 SPECTRAL ENERGY DISTRIBUTION

The time evolution of variable sources like blazars is strictly linked to their detectability in different energy bands in different periods and, consequently, the observational characteristics of their spectral energy distributions (SEDs), the plots showing the flux density νF_ν [energy area⁻² time⁻¹] as a function of frequency, or equivalently $E^2 \frac{dN}{dE}$ as a function of photon energy. The 1990s, which saw the simultaneous activity of several X-ray experiments, of the Compton Gamma-Ray Observatory (and its high-energy instrument EGRET in particular: Hartman et al. 1999), and of a growing number of ground-based Cherenkov telescopes, definitively showed the width of the electromagnetic emission from blazars, ushering the era of multiwavelength analysis.

GENERAL PROPERTIES

Figure 1.4 plots the observed SED of PG 1218+304 as an example. Spectral data points were retrieved from the Markarian Multiwavelength Data Center (see box in next Subsection) and the STeVeCat (Gréaux et al. 2023), and were obtained from different observation campaigns and catalogs. General characteristics of blazars, and of BL Lacs specifically, can be noticed:

- As already mentioned, the spectrum can cover 20 decades of frequency, from the radio (down to $\sim 10^7$ Hz) to the VHE gamma rays (10^{27} Hz).
- The overall shape of the SED is dominated by **two broad humps**, and no other evident features. More specifically, in restricted energy bands, it can be **modeled with simple power laws**. The example in the picture peaks at high energies: the lower-energy component, which covers the radio to hard X-ray range, has its maximum in the soft X rays ($O(100$ eV)), while the higher-energy curve is centered on the gamma-ray band, and peaks above 100 GeV. In general, the spectra of these objects see a low-energy peak between the far infrared ($\lambda \sim 1$ mm) and the soft X rays, and a high-energy peak between MeV and TeV gamma rays.
- Understandably, the flux variability is evident in archival data, as they represent observations done in different periods. Variability is one of the defining features of blazars and is strongly present at all frequencies (Ulrich, Maraschi, and Urry 1997), showing correlations especially during periods of strongly increased activity. The plot in Figure 1.4, however, only displays clear variability around the two peaks, especially in the X-ray band, where it exceeds two orders of magnitude at the higher energies. Ultraviolet and gamma-ray variability is less pronounced, but still noticeable, being about a factor 5-10.
- The the slope of the flux density, νF_ν , in the **radio** band corresponds to a constant spectral flux density, F_ν (flat spectrum), as anticipated in the previous paragraphs.

1.2. BLAZARS: OBSERVATIONAL PROPERTIES

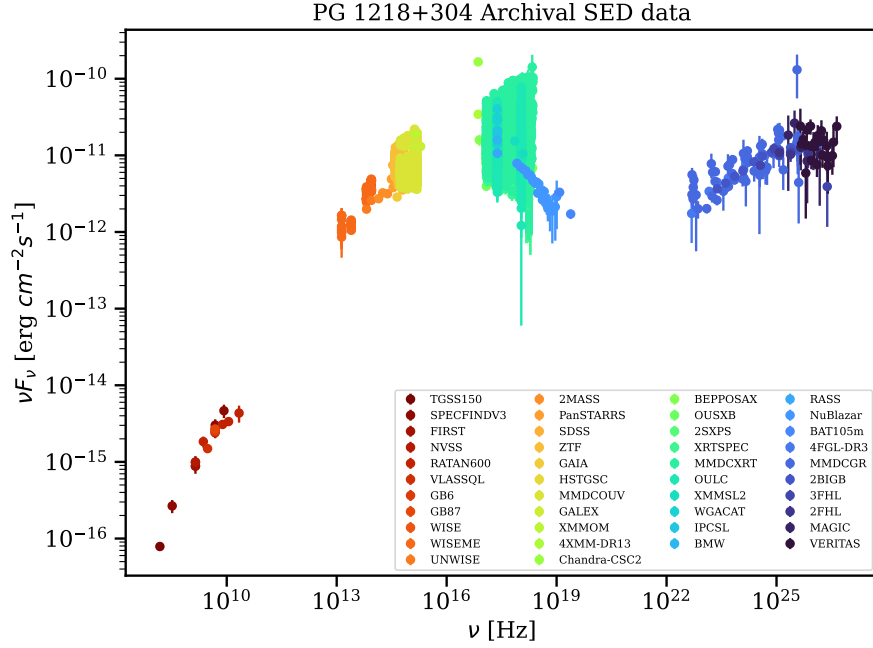


Figure 1.4: SED of the BL Lac PG 1218+304, from archival data downloaded from the Markarian Multiwavelength Data Center and Gréaux et al. 2023. VHE data have been deabsorbed for EBL.

- Observed flux densities, which may need to be K -corrected for redshift to obtain the observed bolometric luminosity⁴, go from less than 10^{-16} erg cm⁻²s⁻¹ in the radio band to a little more than 10^{-10} erg cm⁻²s⁻¹ for the X-ray peak during the highest emission states. Other blazars may instead have the highest peak of their emission in the *second* hump.

Breaking down the description by frequency band allows to notice more features that are characteristic of **BL Lacs** specifically:

Infrared and Optical Data seem to follow the power-law continuum smoothly, up to variability. In the SED of some low-power sources of the same class, however, the profile of the thermal emission from the **host galaxy** (typically a giant elliptical: Urry et al. 2000) can emerge. When going in more detail on this region of the spectrum, no other thermal emission can be discerned, and there are **no (or very weak) emission lines**. This makes it **difficult to obtain a precise estimate of the redshift**: for example, the ‘redshift’ column is non-empty for only $\sim 62\%$ of ascertained BL Lacs in the catalog 4LAC-DR3 (Ajello et al. 2022).

⁴The bolometric luminosity is defined as the quantity of energy *irradiated* per second [erg s⁻¹] integrated over frequency, $L = \int_0^\infty L_\nu d\nu$. See e.g. Ghisellini 2013. The K -correction is applied to flux, luminosity, or magnitude to account for the fact that different sources selected in the same frequency band may have different redshifts, such that the emitted photons actually have different energies from one source to another in their respective rest frames. See Hogg et al. 2002.

Ultraviolet Here variability starts to become more evident, but the spectrum still follows the smooth trend visible at other energies. The slope of the flux density plotted in the SED is decreasing, but still positive. Other sources of the same class may peak here, or show a decreasing trend, meaning they have peaked in the IR-Optical range already.

X rays This is the region where the specific source in the example seems to have its first peak, and the variability is strongest. Higher emission states see a rising spectrum (in νF_ν) even here, while lower emission states are characterized by a decreasing trend. For other sources, the flux density rises, but only after a decline at lower energies, meaning that this region already belongs to the second hump.

Gamma rays The behavior is again smooth, and variability is less pronounced. Data clearly show a second peak, of comparable height to the average of the lower-energy maximum. **Emission is able to reach the realm of VHE gamma rays in individual cases like this.**

MULTIWAVELENGTH OBSERVATIONS

The observed SED is not just a product of the intrinsic spectrum of the source, but also highlights the effect of **detection and monitoring by different instruments, working in a variety of configurations that are not always perfectly compatible in terms of energy ranges, sensitivity and observation mode.**

Online Databases

The broadband study of blazar SEDs is now made easier by a number of interconnected online platforms, that host catalogs at different energies and make their data easily reachable. Apart from more “classic” archives and related tools (HEASARC^a and Xamin, VizieR^b, SIMBAD^c), there are a few interactive services that provide the desired SEDs directly, almost ready to use, after gathering and uniforming data from the archived catalogs at all energy bands. The Space Science Data Center SED Builder of the Italian Space Agency^d, which provides spectra and light curves for astronomical objects of any class, is now joined by other platforms created specifically for blazar astrophysics in the context of a growing need for coordination between observatories and public data availability, as exemplified by the Virtual Observatories network^e and the development of the Open Universe for Blazars tool (VOU-Blazars: Chang, Brandt, and Giommi 2020). The Markarian Multiwavelength Data Center (MMDC: Sahakyan et al. 2024a), created by ICRANet Armenia^f, produces SEDs with archival data obtained with VOU-Blazars and combines them with time-resolved data from var-

ious frequency bands. Like the SSDC SED Builder, it allows to filter data by observation period and instrument, and even provides animated SEDs and light curves for some bright, variable objects. Another Open Universe offshoot is **Firmamento** (Tripathi et al. 2024), hosted by the New York University, Abu Dhabi^g. Originally conceived as a mobile-friendly portal for citizen science, it has now turned to a scientific platform to search for blazars with a multimessenger perspective, retrieve their SEDs from over 90 catalogs, and characterize them through the machine learning models, *wpeak* and *BlaST* (see Section 4.4.2). It does not allow to filter data by frequency, instrument or observation period. Both MMDC and Firmamento, and the VOU-Blazars catalogs, were fundamental tools in this work.

^a<https://heasarc.gsfc.nasa.gov/>

^b<https://vizier.cds.unistra.fr/viz-bin/VizieR>

^c<https://simbad.cds.unistra.fr/simbad/>

^d<https://tools.ssdsc.asi.it/SED/>

^e<https://www.ivoa.net/>

^f<https://mmdc.am/>

^g<https://firmamento.hosting.nyu.edu/home>

Gaps in the SED are probably the first thing that catches the eye. A first reason for this could be that the databases from which points have been retrieved do not include all available catalogs: for example, data are taken from two separate repositories, MMDC up to high-energy gamma rays, and *STeVeCat* for VHE gamma rays (see Section 3.2). The gap in the hard X-ray to low-energy gamma-ray region is instead due in part to the fact that the “valley” between the two humps falls here for the example source, such that the flux is low ($\lesssim 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$) and may go below the sensitivity of current and past observatories. For example, *Swift*-BAT, *Fermi*-GBM (see 3.2) and INTEGRAL-SPI (Kuulkers et al. 2021) should be able to “see” part of that range, but the sensitivities of all three instruments are generally not enough; one of the selected sources in this work features *Swift*-BAT data, but at high fluxes close to the lower-energy peak. A similar discussion about the impact of instrument performance on the shape and the knowledge of the SED also applies to detectors that are even more renown for their unprecedented sensitivity in their energy range, which led them to discover and characterize thousands of blazars. For example, the latest *Fermi*-LAT (see Section 3.2) source catalog, 4FGL-DR4 (Ballet et al. 2024), lists 7194 different objects, 3935 of which are blazars; 1490 are BL Lacs. As already explained, in most cases BL Lacs have their second peak in the gamma-ray band, so the observed flux in the interval covered by the telescope (50 MeV to 1 TeV) is usually above its sensitivity, and data are of good quality. However, as for any instrument, sensitivity is described by an energy-dependent curve (Figure 1.5a): **the detectability of a source will therefore change depending on the energy**, and on the

shape of the spectrum compared to the sensitivity curve.

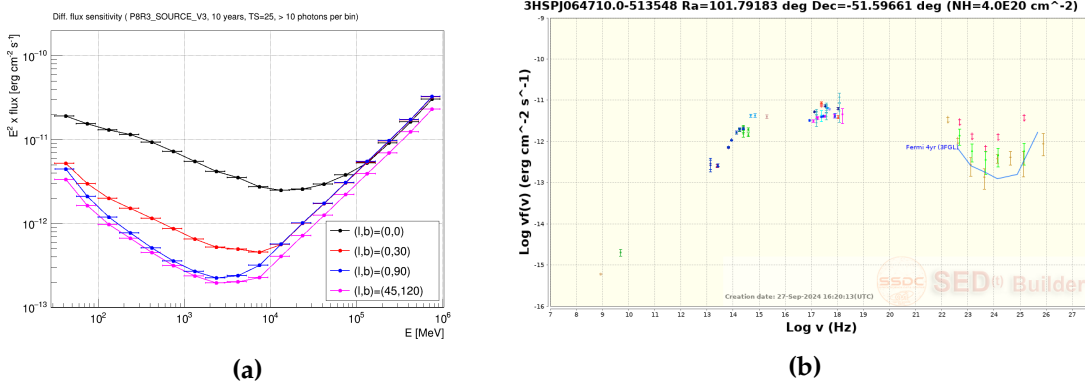


Figure 1.5: 1.5a: 10-year differential sensitivity of the *Fermi* Large Area Telescope, from the LAT Performance web page (https://www.slac.stanford.edu/exp/glast/groups/canda/lat_Performance.htm). 1.5b: SED of the BL Lac object 3HSP J064710.0-513548, downloaded from SSDC SED Builder and compared to the *Fermi*-LAT 4-year sensitivity. Lime green points at gamma-ray energies come from the 3FGL catalog (2008-2012 averages: Acero et al. 2015). Brown points come from the 4FGL-DR3 catalog (2008-2020: Abdollahi et al. 2022). Magenta arrows are upper limits from the 2FGL catalog (2008-2010: Nolan et al. 2012).

Selection biases, as pointed out by Padovani et al. 2017, are a consequence of this: blazars of different subclasses will be selected more or less frequently than others depending on the instrument and the energy range—potentially affecting population studies, especially if based on just one or few detectors as it happens for *Fermi*-LAT in the high-energy gamma ray band. There are BL Lacs, for example (see Figure 1.5b), whose observed flux is in general so low, and whose high-energy peak in the SED is shifted to such high frequencies, that the “valley”, or the rising side of the second hump, falls in the *Fermi*-LAT range and is barely detected, yielding an observed spectrum that suspiciously follows the sensitivity curve despite not being flagged as upper limits.

Another issue that emerges when inspecting the broadband SED of blazars, and is related to the instrument performance, is the variability. The fact blazars seem to be more widely variable at high energies, especially higher-peaked sources whose first flux maximum is in the X-ray band, is known since a long time (see e.g. Ulrich, Maraschi, and Urry 1997), and this is also evident from the example SED. However **how different instruments at different energies take data also has a role**: this is covered more in detail in Section 3.2. Each instrument, depending on the sensitivity and field of view, will adopt its own observation strategy: single pointings with short exposures that are recorded one by one in catalogs, or wide surveys whose data are averaged over several months of years, reducing the effect of variability in the observed SED. In other words, data are so heterogeneous they may pose a problem to analysis, as illustrated in the next Chapter.

1.2.4 THE HIGHEST ENERGIES

A few tens BL Lacs are currently detected at TeV energies⁵. There are several reasons for this: the first is, of course, the fact that only a small fraction of astrophysical sources would be able to emit such energetic photons. There are others, however, that pertain to observation conditions. As these objects are looked at from a distance of hundreds of megaparsecs, the extragalactic background light (EBL) that pervades the Universe starts to have an impact on the VHE observations, attenuating the flux in an energy- and redshift-dependent way (e.g. Spurio 2018):

$$I(E, z) = I_0 \exp^{-\tau_{\gamma\gamma}(E, z)}$$

where $I(E, z)$ is the measured intensity, I_0 is the intensity at the source, and $\tau_{\gamma\gamma}(E, z)$ is the optical depth for the interaction between VHE gamma rays (γ_E) and EBL photons (γ_ϵ):

$$\gamma_E + \gamma_\epsilon \rightarrow e^+ + e^-$$

The maximum happens at

$$\epsilon \simeq \frac{520 \text{ GeV}}{E} \text{ eV} \quad (1.5)$$

so the VHE band suffers the effect of the infrared-optical range of the background light, coming from all the stars, galaxies and AGN across the history of the Universe. The highest-energy data in Figure 1.4 had already been de-absorbed before plotting. **EBL makes the Universe opaque to VHE gamma rays, reducing the horizon to low redshifts**, but observing blazars in turn helps to better characterize this relic of past epochs.

Imaging air Cherenkov telescopes (Section 3.2) are the main type of instruments for the detection of VHE gamma rays today, through the reconstruction of the electromagnetic particle showers induced by the interaction of energetic photons with the atmosphere. Their working principle requires them to have excellent timing and repositioning abilities, but they have a small field of view (\sim msr) both when considering single mirrors and arrays. Consequently, they need to perform specific observation campaigns toward a target, rather than surveying the whole sky: this produces a **bias** in the blazar population covered by these instruments in favor of brighter objects, often in the context of multiwavelength campaigns during flares, or selected as potential TeV sources by extrapolating from the information in *Fermi*-LAT catalogs (Prandini and Ghisellini 2022). These campaigns have a duration of the order of tens of hours, comparable to the longest exposures of some X-ray space telescopes, such that sensitivities are typically expressed for integration times of \sim 25 h to 50 h and reach $\sim 10^{-10}$ erg cm⁻²s⁻¹ in

⁵<http://tevcat2.uchicago.edu/>

that time span; **catalogued spectra that appear in SEDs from online databases also refer to multi-hour and multi-night observations.** However, for bright enough sources, the good sensitivity and the extended exposure times together make this type of observatories the ones able to detect the *shortest* variability timescales (\sim few minutes) and, consequently, to put the strictest constraints on emission models, as already underlined in Section 1.2.2.

Extreme-TeV BL Lacs TeV observations have confirmed the existence of a subclass of BL Lacs dubbed *extreme-TeV blazars* by virtue of their high-energy hump peaking above 1 TeV. *Extreme blazars* in general are not a homogeneous category (Figure 1.6): there are *extreme-synchrotron* sources (see Section 2.1 for the discussion on the synchrotron emission model) whose lower-energy peak is found at either $\nu_{\text{syn}} \geq 10^{17}$ Hz or $E_{\text{syn}} \geq 1$ keV, depending on the definition, but they do not show a second maximum at TeV energies; some others become extreme for both peaks during flares; and others that are both extreme-synchrotron and extreme-TeV sources steadily, and are characterized by a hard GeV spectrum, with an unusual photon index $\Gamma_\gamma < 2$, that often make them candidates for TeV detection by air Cherenkov telescopes already before discovery.

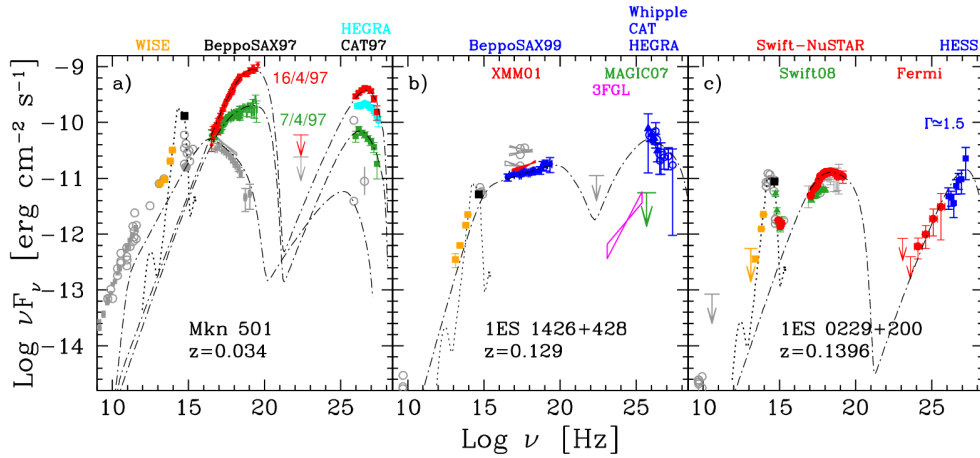


Figure 1.6: Examples of the three types of extreme blazars, from Biteau et al. 2020. a): Markarian 501 becomes an extreme blazar during flares. b): 1ES 1426+428 is an extreme-synchrotron blazar, but not extreme-TeV, during quiescence. c): 1ES 0229+200 is the prototype of extreme-TeV blazars.

The latter group in particular is characterized by a relatively **low luminosity**, such that they are sometimes barely detectable by *Fermi*-LAT (the main reference for blazar-linked observation proposals to Cherenkov telescopes), as explained in the previous Subsection: there is a possibility that a population of *extremely* low-power, high-energy BL Lacs is being missed by our instruments. Extreme blazars are particularly interesting not only because they may shed light on the most energetic acceleration and radiative processes, but also because the detection of TeV photons from extragalactic sources may

give experimental contributions to open issues in cosmology and fundamental physics (Biteau et al. 2020). TeV photons allow to inspect the lowest-frequency range of the infrared EBL (see Equation 1.5); the outgoing electron-positron pairs may be deflected by the intergalactic magnetic field and upscatter photons of the cosmic microwave background to the gamma-ray band, creating image and spectral signatures that can be used to put constraints on the intergalactic magnetic field. More exotic applications include the search for violations of the Lorentz invariance and for mixing with axion-like particles.

1.2.5 CLASSIFICATION

The observational description of blazars so far focused on BL Lacs, the subclass this work is centered on. As already mentioned, they are not the only category of these intriguing objects: there are also *flat-spectrum radio quasars* (FSRQs). Being blazars, FSRQs obviously share most fundamental features with BL Lacs: as the name suggests, the flat spectral flux density (F_ν) in the radio band; the overall shape of the spectrum, with the usual two humps; the polarization of most electromagnetic emission; the fact that these sources are variable. Other important characteristics tell them apart from their BL Lac cousins, starting from the definition itself. The *equivalent width* of spectral emission lines is

$$EW = \int \frac{F_0 - F_\lambda}{F_0} d\lambda$$

where F_0 is the flux of the power-law continuum, F_λ is the total flux of the line and the continuum at the wavelength λ . FSRQs are defined as blazars with broad emission lines in the optical-IR band, $EW \geq 5 \text{ \AA}$; BL Lacs are characterized by very weak or absent emission lines, as already mentioned in Subsection 1.2.3: $EW < 5 \text{ \AA}$. The FSRQ class also seems to correlate to higher redshifts, higher bolometric luminosities (dominated by the higher-energy peak) coupled with steeper spectra (the “blazar divide”) corresponding to lower peak frequencies, stronger flux variability in the GeV range (Dermer and Giebels 2016 and references therein), weaker variability in the first peak frequency, which generally sits at lower energies (Giommi and Padovani 2021), and a more complex SED than the almost featureless one of BL Lacs, with contributions from thermal emitters from the infrared to the ultraviolet band (Figure 1.7). BL Lacs are in turn classified into three main categories, depending on the position of their lower-energy peaks:

LBL (Low-synchrotron peaked BL Lacs): $\nu_{\text{syn}} < 10^{14} \text{ Hz}$

IBL (Intermediate-synchrotron peaked BL Lacs): $10^{14} \text{ Hz} \leq \nu_{\text{syn}} < 10^{15} \text{ Hz}$

HBL (High-synchrotron peaked BL Lacs): $\nu_{\text{syn}} \geq 10^{15} \text{ Hz}$

Extreme-synchrotron BL Lacs ($\nu_{\text{syn}} \geq 10^{17} \text{ Hz}$) may be singled out as a subcategory of HBLs. Some sources, like Giommi and Padovani 2021, adopt a similar nomenclature




Figure 1.7: SEDs of 3C 454.3, a FSRQ, in different emission states, from Bonnoli et al. 2010. The dashed lines represent the thermal contribution from other structures (an accretion disk, a dust torus, and a hot corona). Notice the higher luminosity of the high-energy hump and the strong, fast variability.

without distinction between BL Lacs and FSRQs. This reproduces the alternative classification introduced by Abdo et al. 2010, where the three groups are called LSP, ISP and HSP (low-, intermediate-, high-synchrotron peaked blazars).

FSRQ-like Blazars

The picture of BL Lac classification, which may look like a mere matter of conventions, is made more compelling by the fact that the BL Lac/FSRQ divide is indeed blurred. Ghisellini et al. 2011 and Giommi and Padovani 2021, for example, discuss about how there is a number of sources that present optical and gamma-ray features of both BL Lacs (typically LBL or IBL) and FSRQs, but are classified as BL Lacs under the definition; in this scenario, these sources would be described as “masquerading” FSRQs, whose non-thermal continuum at lower energies is so bright that it obscures optical emission lines. Another family of objects, steep-spectrum “true” BL Lacs (mostly LBLs), have gamma-ray luminosities spanning a wide range that may come close to that of “classic” FSRQs and “masquerading” ones, and are considered transitional between the two main blazar classes (Ghisellini et al. 2011). Examples of potentially “transitional” or “masquerading” BL Lacs can be found, for example, in Ahnen et al. 2018 and Padovani et al. 2019. Papers like Giommi and Padovani 2021 even suggest to completely abandon the classic dichotomy in favor of another one that groups FSRQs and LBLs on one side, and IBLs and HBLs on the other, while Ghisellini et al. 2011 is more cautious and simply warns not to trust the emission line-based definition too much, as the strong variability of blazars may lead to wrongly classify a FSRQ as a BL Lac

because it was observed during a period of strong brightening of the non-thermal continuum, or to mistake a BL Lac for a FSRQ during a very low emission state when intrinsically weak lines may stand out more. **In this work, sources will be selected according to the classical definition, but the issue of “transitional” objects will have an impact**, as better explained in Section 3.4.

1.3 THE BLAZAR SEQUENCE

Population studies on blazar SEDs reveal that they may follow a general trend, called *blazar sequence*, that relates the bolometric luminosity reconstructed from observations in specific energy bands and the frequency of the lower-energy peak in the spectrum.

The first work to realize this, and to perform a systematic study of blazar SEDs including gamma-ray data above 30 MeV, was Fossati et al. 1998, at a time where the availability of multiwavelength data was much more limited than today. The authors selected their sources from 3 samples: one in the X-ray band (the *Einstein* Slew Survey of 1992) and two in the radio band (one for BL Lacs and one for FSRQs). The result was a dataset on 126 blazars, of which 6 BL Lacs selected in both frequency intervals, and 33 objects detected in gamma rays by EGRET. The average SED was reconstructed for each source by computing the logarithmic mean of archival data points at each frequency. Taking the maximum fluxes to include gamma rays was considered as an alternative, but discarded to avoid biases in favor of higher-energy states detected by EGRET (see discussion on a similar problem in Section 3.2). After performing the *K*-correction on observed fluxes to account for the redshift, the frequencies of the lower-energy peaks were computed with polynomial fits of the SEDs and plotted against a number of other quantities: luminosities νL_ν in the radio and gamma-ray bands and at the peak, broad-band photon indices, and *gamma-ray dominance* (defined either as the ratio of gamma-to-low-energy peak luminosities, what is now more often called *Compton dominance*, or as the ratio between gamma-ray and optical luminosities). This analysis was completed by building the sequence in its most famous form: sources were classified in 5 GHz radio luminosity bins, as luminosity seems to be the quantity that correlates with the others the most, and their SEDs were averaged in each bin irrespective of their original classification. Part of the result is shown in Figure 1.8. The main findings of this study were:

- When accounting for redshift biases, there is an **anti-correlation between the first peak frequency**, ν_{syn} and the **radio luminosity** at 5 GHz, $L_{5\text{GHz}}$, while the correlation with the peak luminosity is weak; when considering the average SEDs, the $\nu_{\text{syn}}-L_{5\text{GHz}}$ relation becomes more complex, such that “high-luminosity” and “low-luminosity” blazars need to be analyzed separately: the distinction roughly matches the one between FSRQs and BL Lacs;

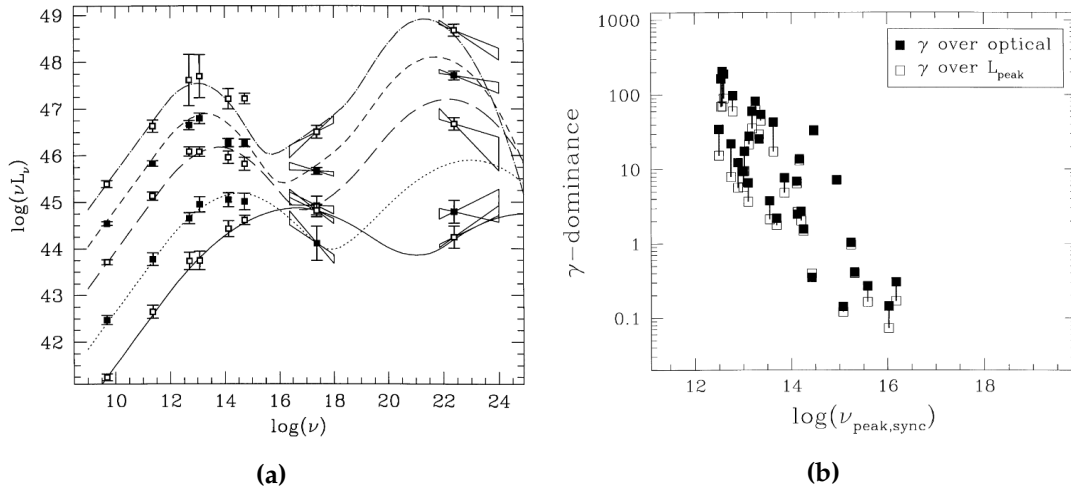


Figure 1.8: Two realizations of the first blazar sequence, from Fossati et al. 1998. 1.8a: average SEDs of sample sources binned by luminosity, irrespective of classification. The curves are phenomenological approximations computed by glueing power-laws and parabolae together. 1.8b: Scatter plot of the low-energy peak frequencies against the gamma-ray dominance, according to two different definitions.

- There is a strong anti-correlation between ν_{syn} and the broadband spectral indices α_{RO} and α_{RX} (optical/radio and X-ray/radio spectral slope in logarithmic scale, as $F(\nu) \propto \nu^{-\alpha}$);
- There is a **strong anti-correlation** between ν_{syn} and the **gamma-ray dominance**;
- There is a **positive correlation between the frequency of the first and the second peak**: the X-ray spectrum becomes harder, and the gamma-ray spectrum becomes softer, with increasing luminosity, indicating a shift of the higher-energy hump, too, toward lower frequencies;
- No source peaking in the X-ray band and emitting at TeV energies is predicted to have high luminosity;
- Not only high-peaked sources are predicted to be detectable at TeV energies, but also some low- and intermediate-peaked ones.

The original sequence was later updated in a paper by Donato et al. 2001, confirming and strengthening the previous findings with an extension of the analysis to the average spectral quantities (fluxes and spectral indices) in the 2 keV to 10 keV. The new study gave weight to the *physical model* according to which the two portions of the SED are generated by different, but related, radiative processes (the synchrotron self-Compton scenario: see Section 2.1 in the next Chapter).

Many years and discussions later, in 2017, a revision of the sequence (Ghisellini et al. 2017) was published to include high-energy gamma rays in the picture. In the late 1990s, only scarce (but fundamental) data from EGRET were available; in the late 2010s, the third catalog of AGN detected by *Fermi*-LAT, the successor of EGRET, had been published (Ackermann et al. 2015), listing 1563 AGN of any type and their average

1.3. THE BLAZAR SEQUENCE

fluxes over 4 years of observations. Blazars with known redshift and no identification issues were selected, producing a sample of 747 objects. A similar approach to the original 1998 work was adopted: SEDs were reconstructed from archival data, averaging different emission states of the same source, and classified in 6 bins according to the K -corrected *gamma-ray* luminosity in the 100 MeV–100 GeV range. The average SEDs for each bin were then interpolated with a piecewise phenomenological model. Results

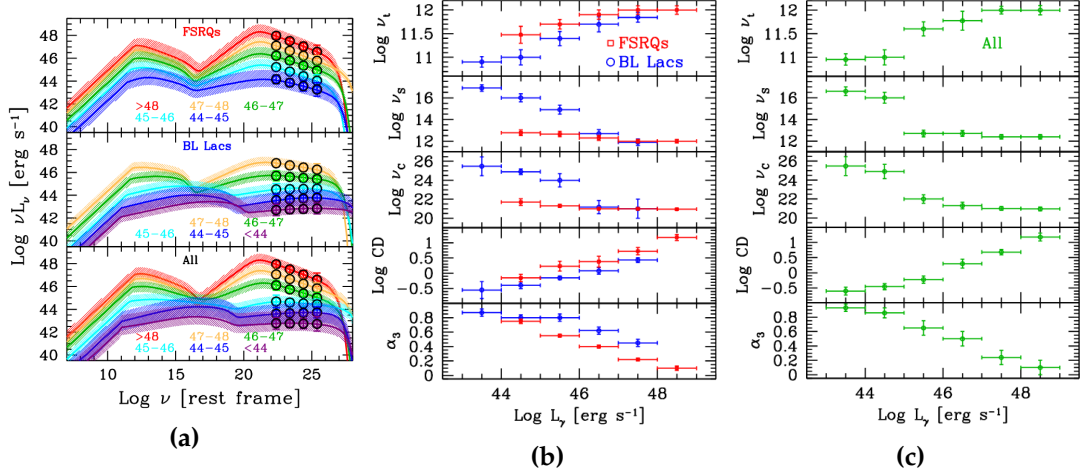


Figure 1.9: The blazar sequence, from Ghisellini et al. 2017. 1.9a: the sequence of average SEDs. 1.9b: self-absorption frequency ν_t , lower-energy peak frequency ν_s , higher-energy peak frequency ν_c , Compton dominance CD and main slope of the high-energy component α_3 against gamma-ray luminosity in logarithmic scale, $\log_{10} L_\gamma$, for FSRQs and BL Lacs separately. 1.9c: same as the previous plots, but for all blazars together.

(see e.g. Figure 1.9) were discussed for all the blazars together, or for FSRQs and BL Lacs separately:

- **BL Lacs** have generally **lower gamma-ray luminosities and redshifts** than FSRQs.
- The photon index in the radio band, $\alpha_R = -0.1$ works for all sources.
- For **FSRQs**, the sequence mostly regards the “**Compton dominance**”, which increases smoothly with luminosity. The change in the lower-energy peak frequency is slight compared to the increase in gamma-ray luminosity.
- For BL Lacs, there is an overall trend of **both peak frequencies being anti-correlated with the gamma-ray luminosity**. The low-energy peak frequency ν_{syn} increases by 5 orders of magnitude with a 4-orders of magnitude decrease in luminosity; almost the same happens for the high-energy peak. The Compton dominance changes by one order of magnitude across the sequence, indicating that the two components of the spectrum are almost equal for BL Lacs.
- For blazars all together, the sequence becomes more evident, as high luminosities are dominated by FSRQs, and low luminosities by BL Lacs. There is an abrupt change in ν_{syn} at $L_\gamma \sim 10^{45} \text{ erg s}^{-1}$.
- There is a large dispersion of points around the median spectrum for low- and intermediate-luminosity bins of the overall sequence, probably caused by the co-

existence of objects of different nature (“low”-luminosity FSRQs and intermediate-luminosity BL Lacs) in the same bin, which did not happen in the original sequence due to the sample being more biased in favor of brighter sources; there is a possibility that the sequence does not really involve luminosity *per se*, but rather the luminosity in units of the Eddington limit.

The new sequence therefore confirmed the old one in several aspects, but also brought some surprises, especially when analyzing the two blazar subclasses separately. The point about dispersion is mainly a consequence of more sensitive instruments being available, that can detect objects with a wider range of luminosities and black hole masses. Considering that *Fermi*-LAT has a limited energy range, despite the smaller bias in the sample there are still at least two groups of sources that have probably been excluded: powerful FSRQs whose high-energy peak falls in the MeV range, and the most extreme-TeV BL Lacs, at the opposite end of the sequence. These two limit cases are, reasonably, the next step: this work will focus on the latter and on other TeV-detectable sources.

2

Blazars: Theoretical Models

Decades of ever-improving observational evidence have allowed to build a sort of “standard model” of the general structure and radiative processes happening in AGN and producing the fluxes of photons (and, lately, neutrinos) that we observe from Earth. However, it is this same wealth of increasingly accurate, and sometimes surprising information that is now leading to a growing variety of theoretical models and making these objects as intriguing as they were 60 years ago. Understanding the blazar sequence is one of the most debated ways to investigate on the physics of these astrophysical sources.

2.1 THE SPECTRAL MODEL: SYNCHROTRON SELF-COMPTON

2.1.1 SYNCHROTRON EMISSION

The peculiar shape of the blazar SED, along with radio imaging and analyses in the time domain, paints a picture where **particles are accelerated to high energies along a direction close to our line of sight, and their radiation undergoes relativistic beaming, such that it ends up dominating the spectrum of these sources**. In explaining the SED of BL Lacs specifically, there is consensus on the origin of the *lower-energy* component: **electron synchrotron**. Even in a classical context a charged particle spiraling in a magnetic field, being subject to centripetal acceleration under the Lorentz force, will radiate power. In case of relativistic electrons in an astrophysical environment, the gyration radius will be the Larmor radius $r_L = \gamma \frac{m_e c^2}{eB}$, depending by the modulus of the magnetic field B , and the distribution of pitch angles will be isotropic due to particle scattering and irregularities in the field (Spurio 2018). Each electron emits strongly

beamed radiation, according to:

$$P_e = \frac{4}{3} \sigma_T c u_B \gamma^2 \beta^2 \quad (2.1)$$

in the reference frame of the observer; of course, $\beta = \frac{v}{c}$ for the electron, and $\sigma_T = \frac{8\pi}{3} \frac{e^4}{m^2 c^4}$ is the Thomson cross section; $u_B = \frac{B^2}{8\pi}$ is the magnetic field energy density. Most of this power will be radiated at a characteristic *synchrotron frequency* $\nu_S = \gamma^2 \nu_L = \gamma^2 \frac{eB}{2\pi m_e c}$, with ν_L the Larmor frequency. Due to relativistic beaming, the observer will only see the radiation when the electron velocity vector lies at an angle $\theta \lesssim \frac{1}{\gamma}$. In the **“standard” leptonic model for blazars, electrons move at relativistic speeds in a “blob” immersed in a magnetic field appearing homogeneous and random in its reference frame**, so one may expect this relation does not really hold, as synchrotron radiation will be emitted in all directions; however, switching to the observer frame, the blob has a bulk relativistic motion with Lorentz factor Γ , that beams the collective synchrotron radiation in an angle of aperture $\frac{2}{\Gamma}$. When the jet is formed by the nucleus of an active galaxy along an axis lying at a smaller angle than $\frac{1}{\Gamma}$ from the line of sight, the observer will see a blazar (see Section 2.2).

When, instead of a single particle, the radiation source is a distribution of electrons accelerated with Doppler factor δ_D in a blazar jet at a luminosity distance d_L , one may write for the *flux density*, following Finke, Dermer, and Böttcher 2008¹:

$$f_\epsilon = \frac{\sqrt{3}}{h} \frac{\delta_D^4 \epsilon' e^3 B}{4\pi d_L^2} \int_1^\infty d\gamma' \frac{dN'_e}{d\gamma'}(\gamma') R(x) \quad (2.2)$$

where all the primed quantities ϵ' , $N'(\gamma)$ are in the reference frame of the electron distribution; $\epsilon = \frac{h\nu}{m_e c^2}$ is the dimensionless energy of emitted synchrotron photons in units of electron mass-energy, $\frac{dN_e}{d\gamma}$ is the differential number of particles in the total volume of the source; $x = \frac{4\pi\epsilon' m_e^2 c^3}{3eBh\gamma'^2}$, and $R(x)$ is the synchrotron spectral power of a single electron averaged over the pitch angle (see also Ghisellini 2013, with a different notation).

Once the **electron distribution** is defined, a more concrete description of the spectral shape can be done. For **particles accelerated to relativistic energies by non-thermal processes**, the trend is a **power law**,

$$\frac{dN'_e}{d\gamma} = k_e \gamma^{-p}$$

Combining this with Equation 2.1, one can obtain that the spectral flux density in the

¹Chosen and reproduced here as it is the formalism adopted by `agnpy`, one of the SED modeling tools used in this work.

observer frame is

$$F(\nu) \propto k_e B^{\frac{p+1}{2}} \nu^{-\frac{p-1}{2}} \quad (2.3)$$

So **the slope depends from that of the electron distribution**, that can therefore be measured giving precious information on the acceleration process.

Synchrotron photons may be re-absorbed by their parent electrons themselves (*synchrotron self-absorption*). This mechanism will not be included in the models in Section 3.4, as it is not observable in blazars, but it is remarkable in that the self-absorbed flux

$$F(\nu) \propto \nu^{\frac{5}{2}} B^{-\frac{1}{2}} \quad (2.4)$$

has a fixed slope, and may constrain the intensity of the magnetic field just by knowing the angular size of the source. Self-absorption appears at frequencies below a certain characteristic limit (the *self-absorption frequency*) marking the transition between optically thin and thick regimes, and usually falls close to the synchrotron peak frequency ν_{syn} (Ghisellini 2013).

2.1.2 SYNCHROTRON SELF-COMPTON

The interpretation of the second hump in the SED of blazars is not univocal, and depends from the nature of the source and from the physical model that seems the most suited to observations beyond the mere spectrum (imaging, multimessenger). It has been found that BL Lacs are generally well described by the *homogeneous one-zone synchrotron self-Compton* (SSC) model. This clever scenario stems, among other things, by the fact that variability in different frequency bands seems to be correlated, so it is reasonable to suppose that all radiative processes happen in a single, limited volume of the jet: the same magnetized blob that emits the synchrotron radiation. In this context, the electrons cool down by transferring part of their energy by inverse Compton scattering to a fraction of the synchrotron photons they produced themselves.

Inverse Compton scattering may happen in two regimes: Thomson, when $\epsilon' < m_e c^2$ in the electron rest frame, and Klein-Nishina, when seed photons are more energetic. In the first case, the radiated power by a single electron is given by

$$P_C = \frac{4}{3} \sigma_{\text{TC}} \gamma^2 \beta^2 u_r \quad (2.5)$$

with $u_r = \int \epsilon n(\epsilon) d\epsilon$ is the energy density of the target radiation before scattering, in the observer frame. The similarity of this expression with the equivalent one for synchrotron allows to write a simple relation:

$$\frac{L_{\text{IC}}}{L_{\text{syn}}} = \frac{P_{\text{IC}}}{P_{\text{syn}}} = \frac{u_r}{u_B} \quad (2.6)$$

which holds as long as either process is not inhibited: for example, it does not work in the Klein-Nishina regime, when the scattering cross section becomes small and energy-dependent ($\sigma_{KN} \propto \epsilon^{-1}$ at high energies), hindering electron cooling. It is therefore possible to estimate the relative weight of the radiation and the magnetic field by measuring the observed luminosity ratio. In the Thomson regime, the maximum energy of scattered photons is $E_\gamma^{\max} \simeq 4\gamma^2\epsilon$, meaning that, even with an electron Lorentz factor γ close to the usual *minimum* value for high-synchrotron peaked BL Lacs ($\sim 10^2$), the scattered photons can reach energies that are a factor 10^4 higher than the original: for example, a near-infrared photon of frequency $\nu = 10^{14}$ Hz becomes a 4 keV X-ray. In the Klein-Nishina regime, the photon energy is boosted by a factor $\leq \gamma$. The spectral flux density of inverse Compton radiation when the population of seed photons is non-monochromatic, and the electron distribution is a power law is

$$F(\nu_C) \propto \nu_C^{-\alpha} \int_{\nu_{\min}}^{\nu_{\max}} u_r(\nu) \nu^{\alpha-1} d\nu \quad (2.7)$$

where ν_C is the frequency of the scattered photon. The integral is non-trivial, as the integration limits depend on ν_C , but if the process is **SSC**, $u_r \propto \nu^{-\alpha}$ and the integral is a constant, $\ln \Lambda$:

$$F(\nu_C) \propto \nu_C^{-\alpha} \tau_C \ln \Lambda \quad (2.8)$$

where $\tau_C = \sigma R k_e$ is the inverse Compton optical depth, and depends from the interaction cross section σ , the source size R , and the electron number density k_e . Figure 2.1 shows a schematic view of a SSC spectrum.

When accounting for Doppler boosting of the radiation emitted by a blob moving toward the observer at relativistic speed with a Lorentz factor Γ , as it happens in blazars, the complete formula for the SSC observed flux density is (Nigro et al. 2022 and references therein):

$$\begin{aligned} \epsilon_c F_{\epsilon_c} &= \frac{3}{4} c \sigma_T \epsilon_C^2 \frac{\delta_D^4}{4\pi d_L^2} \int_0^\infty \frac{d\epsilon'}{\epsilon'^2} \frac{du'_{\text{syn}}(\epsilon')}{d\epsilon'} \int_0^\infty \frac{d\gamma'}{\gamma'^2} \frac{dN'_e(\gamma')}{d\gamma'} F_C(q', \Gamma'_e) \\ \frac{du'_{\text{syn}}(\epsilon')}{d\epsilon'} &= \frac{3}{4} \frac{3d_L^2 \nu F_{\nu, \text{syn}}}{cR^2 \delta_d^4 \epsilon'} \end{aligned} \quad (2.9)$$

where u'_{syn} is the spectral energy density of synchrotron photons, and $F_C(q', \Gamma'_e)$ is an integration kernel that encodes the Compton cross section for electrons and photons with uniform spatial distribution.

Notice how several physical parameters of the emission region appear in the two equations, 2.2 and 2.9: fitting these formulae to observational data may yield information on the **source radius**, its **beaming factor** (that includes both the observing angle and the bulk Lorentz factor in its definition), the **number density** and **energy of elec-**




Figure 2.1: A SSC spectrum, from Ghisellini 2013. The red curve is the synchrotron component, the blue one is produced by the inverse Compton process. ν_t is the self-absorption frequency: in blazars, it is not visible and the low-energy cutoff is dictated by the minimum Lorentz factor of the electron distribution. The $\sim \nu^{-\alpha}$ trend is recognizable in both components. The high-energy cutoff for both humps is due to the maximum Lorentz factor of the electron distribution, joined by $\ln \Lambda$ in the inverse Compton curve.

trons, the **slope** of their distribution, and the **magnetic field**. In practice, depending on the chosen electron distribution, there are at least 8 fit parameters, a relatively small number (one of the great advantages of the SSC model) but still enough to introduce degeneracy. The dependence of the SED from single parameters is illustrated in Figure 2.2, comparing a plot from literature with one produced specifically for this work to show the relation of the SED with the beaming factor, and based on a preliminary fit of an example source, the high-synchrotron peaked BL Lac PG 1218+304. The selected electron density distribution is a broken power law (see Section 3.3.1), a standard choice to describe the stationary broadband emission by SSC: the break in the electron distribution is introduced to account for radiative cooling when building a time-independent model. All plots are obtained by varying one parameter, and keeping all the others fixed. The frequency and flux density at both peaks was computed for different values of δ_D , the quantity that impacted the data selection in Chapter 3 the most. Peaks of the SEDs in Figure 2.2b were approximated locally by parabolic fits to obtain maximum precision in the estimate of frequencies and flux densities². These quantities were plotted against the respective value of δ_D , showing precise trends: the expected linear increase

²The package used for spectral modeling, *agnpy*, includes methods to compute these quantities directly, but attempts at using it showed that the outputs are binned and produce step-like plots (see Section 3.2).

2.1. THE SPECTRAL MODEL: SYNCHROTRON SELF-COMPTON

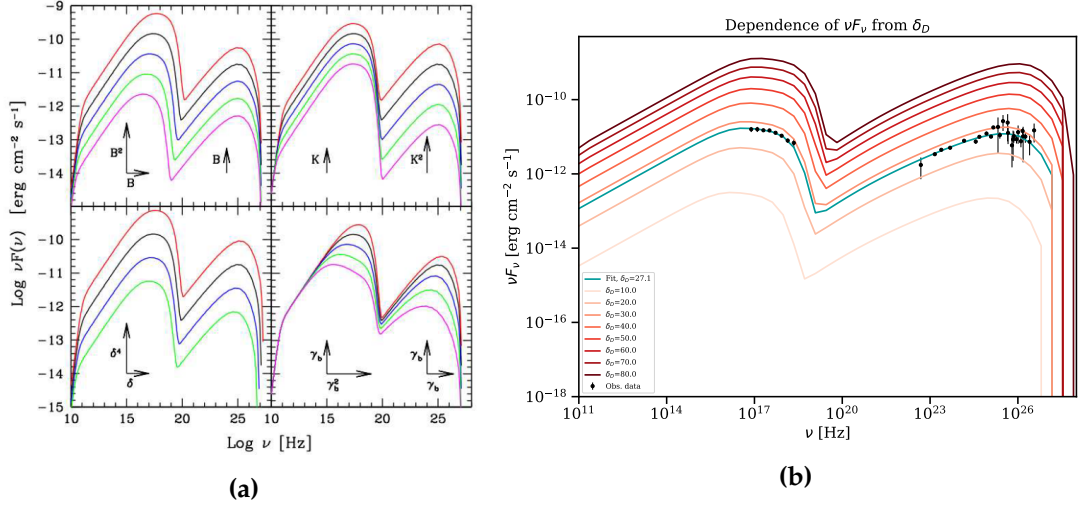


Figure 2.2: Scaling of the SED when varying the SSC model parameters. 2.2a: variation of magnetic field B , electron density normalization k_e , Doppler factor δ_D , and break Lorentz factor γ_b . From Stamerra and Bonnoli 2010. 2.2b: SEDs when varying δ_D between 10 and 80 while keeping all the other parameters fixed. Based on a preliminary fit (teal curve) of real data (black dots) from PG 1218+304: $\log k_e = -7.73 \pm 0.25$ [cm $^{-3}$], $p_1 = 2.12 \pm 0.04$, $p_2 = p_1 + 1 = 3.12 \pm 0.04$, $\log \gamma_b = 5.03 \pm 0.10$, $\log \gamma_{\min} = 2$ (fixed), $\log \gamma_{\max} = 6$ (fixed), $\log B = -1.73 \pm 0.06$ [G], $t_{\text{var}} = 1$ d.

of the peak frequency with δ_D for both components, and the $\nu F_\nu^{\text{syn}} \propto \delta_D^4$ relation that can be found when applying the Lorentz transforms from the blob frame to the observer frame. For what concerns the inverse Compton peak flux density, the relation found empirically seems to be exactly the same of the synchrotron peak, while in Figure 2.2a no dependence is indicated explicitly, as the second component may be governed by other processes than SSC and there could be an impact of the transition from the Thomson to the Klein-Nishina regime (Stamerra and Bonnoli 2010). What emerges is that if slightly different data for a fit are selected, the Doppler factor should change considerably. In fact, this is not *exactly* the case, as more parameters are left free to vary for fitting and “compensate” for the changing δ_D (Section 3.2).

2.1.3 OTHER SCENARIOS

The one-zone SSC model is the standard for BL Lacs, but it is not the only choice when considering blazars at large, or even some BL Lacs themselves. A *dual-zone SSC* has been invoked for a growing number of sources in the last years, to account for mismatches in variability timescales and inferred Doppler factors across different energy bands (Aleksić et al. 2014). The scenario described in the paper for a TeV-detected blazar requires a larger, parsec-scale emission zone for the radio and the majority of the optical emission, while X rays and gamma rays would be produced in a smaller region (with a minor contribution from the larger region).

The existence of transitional sources with the FSRQ class, and the association of the




Figure 2.3: Dependence of both synchrotron (top) and inverse Compton (bottom) peak frequency and flux density from the Doppler factor, δ_D . Points were obtained from parabolic fits of the SEDs in Figure 2.2b around the peaks, and interpolated with a straight line ($\nu_{\text{peak}}(\delta_D)$) or a fourth-degree polynomial ($\nu F_{\nu}^{\text{peak}}(\delta_D)$).

intermediate synchrotron-peaked BL Lac TXS 0506+056 with high-energy neutrinos introduces additional radiative processes to consider alongside the SSC for a finer-tuned model. The gamma-ray hump of FSRQs is usually described with an *external Compton* model. The presence of thermal features like broad emission lines, or the typical profile of an accretion disk, reveals that the FSRQ environment is rich in material surrounding the central engine, which produces abundant low-energy *seed photons* that are upscattered by the energetic electrons in the relativistic jet. The external contribution to the emission is thought to be at the origin of the high luminosity and Compton dominance of FSRQs and some transitional, “red” BL Lacs, since the large quantity of interaction targets makes electron cooling more efficient.

Neutrino emission is also deemed to be more efficient for “FSRQ-like” sources, as more external material and photon fields means more targets for *hadronic processes* (Prandini and Ghisellini 2022 and references therein). In that case, the neutrino flux, and the related high-energy photon emission, would most probably be a byproduct of photo-meson interactions between accelerated protons and target photons: e.g.,

$$p + \gamma \rightarrow \pi + h$$

$$\pi^0 \rightarrow \gamma\gamma$$

2.2. THE PHYSICAL MODEL, PART I: UNIFICATION

$$\pi^- \rightarrow \mu^- \bar{\nu}_\mu$$

$$\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$$

High-energy photons could also be produced via proton synchrotron in presence of a strong magnetic field ($B \sim 30$ G to 100 G) and ultra-high particle energies ($E_p \geq 10^{19}$ eV; Aharonian 2000). In any case, the hadronic component would be limited to the high-energy region, while the leptonic one would still prevail at low energies, or even contribute to the gamma-ray band too, in the so-called *lepto-hadronic scenario*.

Taking into account all possible components, including the hadronic ones, the total power of the relativistic jet would be:

$$P = \sum_i \pi R^2 \Gamma^2 \beta c u'_i \quad (2.10)$$

where u'_i are the energy densities of each contribution in the emission region comoving frame: radiation $u'_r = \frac{L}{4\pi R^2 c}$ for radiation fields, $u'_p = u'_e \frac{m_p}{m_e} \frac{n'_p}{\langle \gamma \rangle n'_e}$ for *cold* protons, and $u'_B = \frac{B^2}{8\pi}$ for the magnetic field.

2.2 THE PHYSICAL MODEL, PART I: UNIFICATION

Now that it is clear that a relativistic jet is involved, the question is what is it powered by. Over the years, observational evidence lent support to general *unified models* (Urry and Padovani 1995), which explain the labyrinth of AGN classes in terms of **viewing angle** and few other features (Figure 2.4). The fundamental components of a blazar are:

- A **supermassive, probably rotating black hole**, with $10^7 M_\odot \leq M_{\text{BH}} \leq 10^{10} M_\odot$, located at the center of a giant elliptical galaxy;
- An **accretion disk** of matter falling onto the black hole;
- A pair of **highly collimated relativistic jets** of matter expelled in from the vicinity of the black hole and extending for several kiloparsecs in opposite directions, along the rotational axis of the central engine.

A rotating black hole orders of magnitude more massive than Sagittarius A* (at the center of the Milky Way, $M \sim 10^6 M_\odot$) would therefore be the “central engine” of AGN, fueled by accretion. As well summarized by Dermer and Giebels 2016, the engine needs to account for a bolometric luminosity above 10^{43} W during powerful flares, that develop in a short amount of time inside jets of particles at relativistic bulk speeds, which are kept collimated for thousand of parsecs: only accreting supermassive black holes would be able to do all that. This picture is supported by an important body of evidence, not least the spectacular result of the Event Horizon Telescope collaboration (The Event Horizon




Figure 2.4: Classification of AGN according to mass accretion rate in Eddington units and black hole spin parameter, from Dermer and Giebels 2016

Telescope Collaboration et al. 2019), who imaged the inner accretion disk of the nearby radio galaxy M87, showing its connection with the spin of the central black hole.

Even if they have a common general structure, the two classes of blazars still differ physically, as the existence of broad emission lines and sub-dominating thermal components in the FSRQ spectrum indicate the prominent presence of abundant circumnuclear material (hot corona, broad-line region, molecular torus, narrow-line region: see Figure 2.5) of which there are no signs in BL Lacs; this translates to different dominating emission processes in the high-energy band (external Compton and SSC, respectively), different accretion powers, and different luminosity, as highlighted in the blazar sequence. Jets appear in <10% of AGN, and are not only present in blazars, but also in *radio galaxies*. Together, they constitute the class of “radio-loud” AGN, now often called “jetted” AGN as their emitted energy is indeed dominated by the jets in both cases (Padovani et al. 2017). According to the unified model, blazars and radio galaxies are the same type of object, and differ in their observational properties by their viewing angle, and the consequent beaming effect on their emission. In particular, BL Lacs are usually considered to be the *aligned* counterparts of “Fanaroff-Riley I” radio galaxies, characterized by low luminosity, and a higher radio jet surface luminosity towards the central engine. FSRQs would be the aligned versions of “Fanaroff-Riley II” radio galaxies, which are more luminous, and feature powerful collimated jets whose surface luminosity is highest on “hotspots” at the far end of extended radio lobes (Figure 2.6a). The beaming effect is therefore most evident in the radio band, where all types of

2.2. THE PHYSICAL MODEL, PART I: UNIFICATION




Figure 2.5: Unified scheme of a radio-loud AGN from Venturi et al. 2020, with an indication of the possible observational counterparts of some of the regions at different energies, from radio to VHE. The scheme follows the “spine-layer” model of the jet, as explained in 2.3.

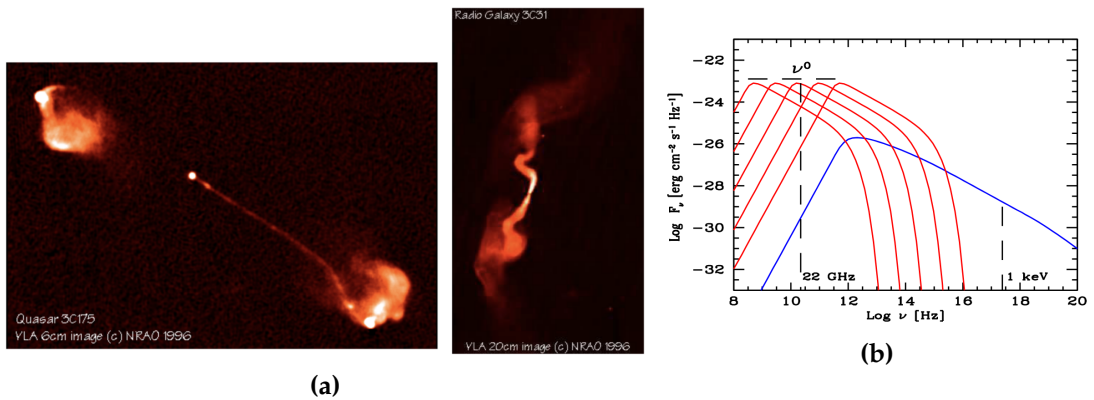


Figure 2.6: Blazars and radio galaxies. 2.6a: Fanaroff-Riley II-like radio quasar (left) and Fanaroff-Riley I radio galaxy (credit: NRAO). 2.6b: composite spectrum of a FSRQ, showing the flatness in the radio band, from Ghisellini 2013.

sources can be observed. At the largest angles the de-beaming is total, and only the lobes are well visible (radio galaxies); as the angle decreases, other regions of the jet start to appear at higher frequencies (steep spectrum radio quasars); for small angles, all radio components are visible, such that their synchrotron spectra, whose peak flux is constant, are superimposed, yielding a flat spectrum up to the self-absorption frequency of the most compact component (Figure 2.6b). The reason for this peculiar behavior is thought to be linked to the conservation of the number of emitting particles and of the Poynting flux, requiring $B(R) \propto R^{-1}$ and $K(R) \propto R^{-2}$. **Since it is the product of more than one emission zone, the radio spectrum has nothing to do with the one-zone synchrotron profile dominating at higher energies, so radio data at $\nu < 10^{11}$ Hz are excluded from the model fits.**

2.3 THE PHYSICAL MODEL, PART II: JETS

The emission models considered so far all require non-thermal distributions of energetic particles: it is clear that **blazar jets, and the underlying mechanism that launches them from the cores of the host galaxies, are powerful natural accelerators.** Among the main reasons for which these sources are studied are, indeed, **understanding what are the exact acceleration mechanisms happening in the extreme conditions of relativistic jets, and evaluating whether they could be the source of high-energy cosmic rays.** Acceleration processes are still an open problem, and their solution is inevitably linked to the modeling of the jet structure on various scales. Two related issues can be singled out: a large-scale one, on how the jet itself is generated and kept collimated on kiloparsec scales, what is its structure, and how are the emitting blobs launched at relativistic bulk speed along the rotation axis of the engine; and a smaller-scale one, related to how can the particles in an emission region be relativistic even in the blob reference frame.

2.3.1 ACCRETION AND THE BIRTH OF JETS

One of the biggest puzzles in astrophysics is not only why are some supermassive black hole at the center of galaxies active and others not, but why do some AGN have strong jets and others not, showing instead a SED dominated by thermal processes in the nuclear region. As shown in Figure 2.4, the key may be in the most fundamental properties of the central engine: the accretion rate and mechanism, and the black hole spin, with a fundamental role played by magnetic fields produced in this context. Imaging and polarization radio observations (e.g. The Event Horizon Telescope Collaboration et al. 2019, Issaoun et al. 2022 for recent examples) show that the jets are originated close to the nucleus, and that they are associated to large-scale ordered mag-

netic fields of toroidal (in the core) or helical (in the jet) shape. The two most popular jet production mechanisms both require similar environments. In the Blandford-Znajek mechanism (Blandford and Znajek 1977), the magnetized inflow from the accretion disk threads a spinning Kerr black hole with magnetic field lines, inducing the formation of a magnetosphere; energy and momentum can be extracted electromagnetically from the black hole spin, powering an electron-dominated jet. The Blandford-Payne mechanism (Blandford and Payne 1982) is instead based on the magnetic extraction of energy and momentum from the accretion disk: in certain conditions, the magnetic field is able to induce an outflow by centrifugal acceleration of the charged particles, creating a jet that is kept collimated by the toroidal component of the field.

The role of the accretion mechanism is also invoked to explain the difference in the luminosity and structure between BL Lacs and FSRQs. **The low luminosity and apparent absence of emission lines and of a molecular torus in BL Lacs can be explained by inefficient matter accretion**, such that the accretion luminosity of the disk is smaller than a certain threshold in Eddington units ($\frac{L_{\text{disk}}}{L_{\text{Edd}}} \leq 10^{-2} - 10^{-3}$; Prandini and Ghisellini 2022 and references therein). BL Lac accretion is thought to occur in the **Advection-Dominated Accretion Flow regime**, a radiatively inefficient mechanism that on one hand, is unable to produce enough photo-ionizing UV radiation to induce line emission in the inner gas clouds (the broad-line region of FSRQs); on the other, it can be associated to the emission of jets and winds (Yuan and Narayan 2014).

2.3.2 JET STRUCTURE

The scenario described in the previous Sections for BL Lacs, the **one-zone SSC**, is quite successful in spectral modeling despite being quite simplistic: the broadband emission is almost completely dominated by radiative processes happening in a **single blob** of relativistic electrons, with a uniform and tangled magnetic field. This assumption is actually **well justified by a number of experimental data**: the aforementioned coherence of emission enhancement during flares across the electromagnetic spectrum, but also the observation of moving “**radio knots**” and apparent **standing shocks** in the jet, where matter should be crossing at relativistic bulk speed, a situation that can be roughly described with the “one-zone” prescription. The problem is how heavy this approximation is: the state is steady, with no time evolution, no particles of other species intervening, and no regard to what may be happening outside of the emitting volume, especially large-scale phenomena. Other more realistic models of an inhomogeneous jet structure have therefore been put forward, often with the support of observational findings, but at the price of more parameters and degeneracy (Sol and Zech 2022). *Two-zone* SSC, or even *multi-zone* models have been invoked to explain peculiar spectra and variability patterns, especially when explaining flares (Abe et al. 2023), but also to account for long-term multiwavelength behavior (Aleksić et al. 2014).

In case of multi-zone models, particles are accelerated and radiate energy all the way along the jet, whose geometry becomes fundamental (e.g., the “conical jet” model). An inhomogeneous scenario that has been applied to some interesting cases, like TXS 0506+056 (Ansoldi et al. 2018) is the *spine-layer*, or *spine-sheath* model (e.g. the jet in Figure 2.5): the jet proper, with a higher bulk Lorentz factor ($\Gamma \sim 13$) and sometimes described as purely leptonic, is surrounded by a slower baryonic envelope ($\Gamma \sim 3$). The two components “see” each other’s radiation beamed due to the difference in speed, enhancing mutual external Compton emission. This picture would also help explaining fast gamma-ray variability in sources that appear more stationary at smaller frequencies.

It was, indeed, the extremely fast variability of some energetic blazars in the TeV band (Section 1.2.2) that inspired even more complex models: from “jets-in-jet” and “mini-jets”, where the jet is populated by several compact sub-components, to electron-positron pair cascades induced in black hole magnetospheres, to jets encountering obstacles like clouds or stars (Sol and Zech 2022). According to Rieger 2019, the log-normality of flux distributions in variable blazars favors a disk-jet connection, and disfavors “additive” models like jets-in-jet.

Recent polarization studies with data from the IXPE mission (e.g. Kouch et al. 2024) lend credibility to the *energy-stratified shock-acceleration* scenario. Particles, after being accelerated by a compact shock front, cool down in the more turbulent, magnetized downstream regions, emitting less and less energetic photons as they travel away from the shock. The result is that a one-zone model, where the entire SED is produced in the same blob, is no more valid, as there are different regions of the jet that emit photons of different energies and polarization degrees, with the highest-frequency flux being more polarized as it is produced closer to the shock front, where the magnetic field is more ordered. Adding to this is the structure of the innermost part of the jet, as it has been reconstructed with the help of numerical magnetohydrodynamical simulations: jets are launched as magnetically dominated outflows (Poynting flux region) and then the magnetic energy is progressively converted to kinetic energy as the jet is collimated and accelerates outwards, as shown in Figure 2.7.

2.3.3 ACCELERATION MECHANISMS

Shock Acceleration Shock acceleration is, indeed, the classic explanation for the emission of energetic massive particles in astrophysical environments, and blazar jets would not be an exception. Shock fronts can form easily, especially in supersonic plasma flows: for example, in case of internal instability, or recollimation of the expanding jet in outer regions by the pressure of the intergalactic medium. Bow shocks due to collision with the external medium form at the far end of the extended lobes in Fanaroff-Riley II radio galaxies (FSRQs), producing characteristic radio hotspots. In such cases, the **Fermi I** process (**diffusive shock acceleration**) happens. Fast particles crossing the shock into

2.3. THE PHYSICAL MODEL, PART II: JETS

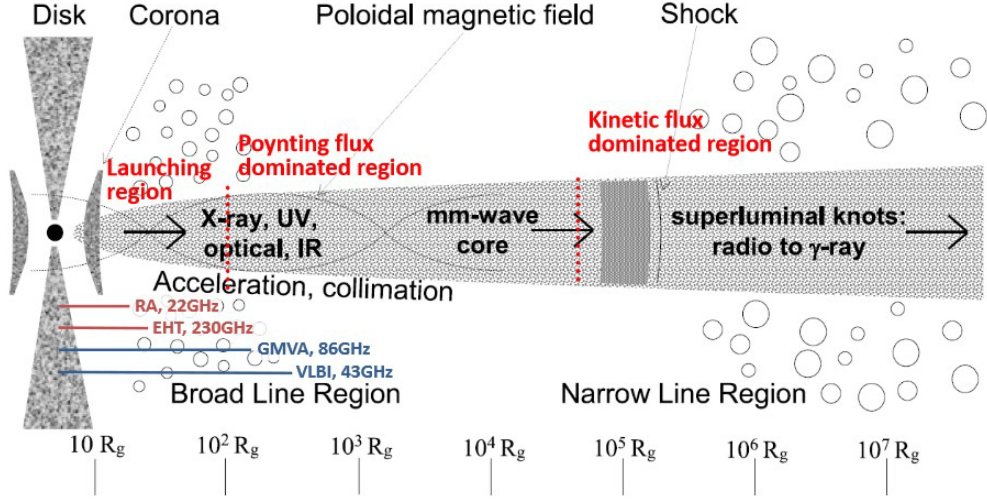


Figure 2.7: Model of the jet structure and radiative emission at different energies, from Venturi et al. 2020.

the downstream region are trapped by the magnetic turbulence, which themselves amplify by inducing Alfvén waves, and cross the shock several times:

$$\langle \Delta E \rangle \propto \frac{r-1}{r} \frac{u_{\text{shock}}}{c} \quad (2.11)$$

is the energy gain per crossing, where r is the *compression factor* of the shock and u_{shock} is the shock speed in the reference frame of the upstream medium (Sol and Zech 2022). This mechanism produces power-law particle distribution with an index

$$p = \frac{r+2}{r-1}$$

which is about 2 for strong shocks. The problem with this process is that it is not efficient for strongly magnetized ultra-relativistic shocks, as it is difficult for particles to cross back to the upstream region: in that case, the power index is softer, $p = 2.2 - 2.3$, and other acceleration mechanisms may gain more weight: for example, “fast Fermi processes”, that exploit the electric fields induced by particle moving toward the shock and may allow acceleration to high energies with a single crossing. In that case, power indices may assume values as low as $p \simeq 1$. For extreme-TeV BL Lacs, multiple shock re-acceleration has also been invoked.

Turbulence Turbulence and instabilities are fundamental for the efficiency of diffusive shock acceleration, but they may induce themselves stochastic acceleration mechanisms of the **Fermi II** type that, despite being second-order, are expected to produce very en-

ergetic particles. The energy gain per interaction is, on average,

$$\langle \Delta E \rangle \propto E \left(\frac{u_c}{c} \right)^2 \quad (2.12)$$

where u_c is the velocity of the scattering center. Turbulence can happen at the interfaces between layers in a scenario like the spine-sheath, and induce *shear acceleration*, with an average particle gain per interaction

$$\langle \Delta E \rangle \propto E \left(\frac{u}{c} \right)^2 \quad (2.13)$$

where u is the effective speed change of scattering centers from the reference frame of a particle crossing the flow along the x direction (transversely) in one mean free path. For an outer layer of thickness $\sim 0.1R_{\text{jet}}$, electrons can be accelerated to PeV energies, and protons to EeV energies, without jet disruption (Sol and Zech 2022 and references therein).

Magnetic Reconnection The main competitor of shock acceleration, as shown by magnetohydrodynamical simulations, is **magnetic reconnection**. Favorable conditions are present in blazar jets, especially the innermost regions closer to the central engine, where the jet is dominated by the Poynting flux and is therefore particularly prone to instabilities. Here “current sheets” may form, where magnetic energy is dissipated and converted to particle kinetic energy in a fast and efficient way (Tavecchio 2021). Reconnecting magnetic fields in a relativistic context accelerate particles to very high energies with a hard power-law distribution of index $p \simeq 1 - 1.5$ (Sol and Zech 2022 and references therein).

2.4 PHYSICS OF THE BLAZAR SEQUENCE

As already highlighted, modeling source spectra and studying the blazar sequence has the clear aim of understanding the physical processes behind particle acceleration and radiation emission in these powerful sources. This is the main concern of the first part of this work (Chapter 3), which was done as part of a **wider study on the physical roots of the sequence**.

The first papers on the topic already tried to give a physical interpretation. Fosfati et al. 1998 underlines how the sequence could give a contribution to understanding whether the high- and low-luminosity blazars are emitting at high energies according to different flavors of inverse Compton scattering (external Compton versus SSC, respectively), with a smooth transition between the two mechanisms, or if the same process is behind the spectra at all peak frequencies, favoring the former idea; they also ad-




Figure 2.8: The Compton dominance of the 2LAC clean sample of blazars against ν_{syn} , and theoretical predictions, from Finke 2013. Black dots: FSRQs; empty red circles: BL Lacs; pink triangles: objects of unknown redshift; green squares: ambiguous AGN.

mit they have no immediate justifications for the phenomenologically dominant role of the luminosity. The 2017 update (Ghisellini et al. 2017), coming at a time when the EC-SSC dichotomy for FSRQs and BL Lacs was consolidated, discusses the findings at length, focusing mostly on the behavior of the FSRQ population. The idea is that **the cooling timescale** is the most relevant quantity in characterizing radiative emission: for FSRQs, it turns out it is constant, while for BL Lacs, the direct proportionality of the synchrotron radiation energy density to the observed luminosity (Eq. 2.9) means that **electron cooling is stronger in more powerful objects**. This reduces the electron energy at the synchrotron peak, whose frequency (alongside the inverse Compton peak frequency) becomes in turn smaller.

A few years before, Finke 2013 analyzed the sequence with a slightly different approach, focusing on some observational quantities extracted from literature and archival data and trying to correlate them to a simplified physical model. A correlation between the synchrotron peak frequency and luminosity is confirmed, as it is the analogous relation between ν_{syn} and the Compton dominance: the latter, in particular, is the focus of the study as it is a redshift-independent quantity that can reduce selection bias. The simplified model is able to predict both trends (Figure 2.8) and depicts a scenario where the **viewing angle**, encoded by δ_D (as $\Gamma = 30$ is fixed) is the main parameter governing the synchrotron peak luminosity, but the overall trend actually exists, and seems to be mostly driven by the **magnetic field** B , which is larger for FSRQs than BL Lacs, and by the **external photon field energy density** u_{ext} . Together, these two quantities determine the **cooling regime**: slow cooling prevails for a minimum electron Lorentz factor $\gamma_{\text{min}} < \gamma_c$, where γ_c is a characteristic cooling electron Lorentz factor, and the synchrotron peak occurs at $\gamma = \gamma_c$; if γ_{min} is larger, fast cooling prevails, and the syn-

chrotron peak happens at $\gamma = \gamma_{\min}$. In any case, electrons lose energy by synchrotron radiation and inverse Compton scattering in the Thomson regime. Higher ν_{syn} matches low values of B and u_{ext} , due to slow cooling; as the two quantities increase, the amount of cooling also increases, and γ_c decreases, until the transition to fast cooling at about $\nu_{\text{syn}} \lesssim 10^{13}$ Hz. A second break at $\nu_{\text{syn}} \sim 10^{14}$ Hz to 10^{15} Hz marks the transition between the external Compton and the SSC. In other words, up to the effect of the observing angle, the sequence really exists and has a physical reason: low-synchrotron peaked sources, especially most FSRQs, emit by synchrotron and external Compton by fast cooling, while **the higher-peaked sources, mostly BL Lacs, follow the SSC radiation model in the slow regime.**

This paper is now quite old, and a wealth of new data have come from all current high-energy observatories. Other approaches can also be adopted: for example, modeling the single SEDs. A recent paper (Kerby and Falcone 2023) is particularly interesting, because it has several elements in common with this thesis as illustrated in Chapter 3. The authors analyze a small sample of dim gamma-ray blazars discovered recently, and discuss their role in the sequence after fitting their SEDs one by one with *agnpy* and evaluating their ν_{syn} with *BlaST*, two of the modeling tools that were employed here. Apart from the anomalous Compton dominance, the average results seem to be coherent with the blazar sequence. The paper, however, is limited to a small population and does not try to do for low-luminosity blazars what is **attempted here with high-energy BL Lacs: rebuild the sequence with physical model parameters.**

Criticisms of the Blazar Sequence Objections to the blazar sequence are not rare. Several researchers argue that the observed trend does not exist physically, and suggest alternative models. According to Fan et al. 2017 and, more recently, Ouyang et al. 2023 (described in Section 3.1) the key is not the cooling mechanism, but rather the relativistic beaming: since δ_D boosts both the intrinsic ν_{syn} and the intrinsic synchrotron peak luminosity, the sequence would disappear, or even show the opposite trend, after correcting for beaming. A radical proposal was advanced by Giommi et al. 2012: in the so-called “simplified scenario”, the sequence is nothing more than a selection artifact related to the depth of radio and X-ray surveys. Monte Carlo simulations based on rather strong assumptions (among others, that all blazars have the same magnetic field, $B = 0.15$ G, and the Doppler factor is $\langle \delta_D \rangle = 15 \pm 2$) show that both high-luminosity, high-peaked blazars and low-luminosity, low-peaked blazars should exist. The latter class, in particular, still would have not been recognized because it would be difficult to detect in the X ray band. Both Finke 2013 and Ghisellini et al. 2017 challenge this model, but future X-ray observations will probably have the last word.

3

Spectral Modeling

TeVCat¹, the online collection of astrophysical sources detected at teraelectronvolt energies (Wakely and Horan 2008), currently lists 334 objects, 85 of which are blazars. Understanding the physical mechanisms behind the emission of such high-energy photons, especially when it peaks in that energy range, may happen by understanding, in turn, their place in the blazar sequence. A way to do this, that is sketched in this Chapter, is by modeling their SEDs: however, as pointed out by Bégué et al. 2023, fitting blazar spectra has progressively become harder from the computational point of view, as it has become clearer, thanks to an increasing number and quality of multiwavelength and multimessenger data, that the simplest spectral models are not enough. Even if reforming the sequence, being a population study, is a task that may in principle be limited to a rougher evaluation of physical parameters, complexity does not disappear, and computational costs rise again if the SEDs of all the sample sources are fitted to avoid indiscriminate averaging. Machine learning algorithms may therefore become useful to tackle these issues.

3.1 AIM: THE TEV BLAZAR SEQUENCE

Most studies on the blazar sequence, as already shown in the previous chapters, have been carried out by selecting sources in specific energy bands: radio and X rays in the original version (Fossati et al. 1998), high-energy gamma rays in the 2017 update (Ghisellini et al. 2017). The latter work was encouraged by the great achievements of the *Fermi* Large Area Telescope, which had been able to detect more than 1700 blazars

¹<http://tevcats2.uchicago.edu/>

from the start of its operations in 2008 (Acero et al. 2015). The 2000s and early 2010s also saw the commissioning of the three current imaging atmospheric Cherenkov telescopes, H.E.S.S., VERITAS and MAGIC (see next Section), for ground-based indirect detection of gamma rays at TeV energies. The completion of the three arrays marked a sharp increase in the total number of jetted AGN, mostly BL Lacs (but also FSRQs and radio galaxies), observed at very high energies. The natural consequence is to check what is the role of these objects in the sequence, if they tend to confirm or disprove its existence in particular, and to discuss their features in this context to have a clearer physical picture of the highest-energy radiative processes.

Of course, this is not the first attempt at including TeV data in the sequence: discussions have been ongoing at least since the discovery of the first extreme synchrotron-peaked blazars in the late 1990s (Costamante et al. 2001), which seemed to be anomalous due to their relative brightness. The last few years saw the publication of several papers focusing on the observational properties and physical modeling of VHE-emitting blazars, encouraged by the fact that the number of detections is now becoming close to that of objects in the original 1998 study. Qin et al. 2018 found consistent results with the blazar sequence in the physical parameters from fits of a sample of TeV-selected BL Lac SEDs; Zhou et al. 2021, working with a larger sample of TeV blazars, found an anti-correlation between the peak Lorentz factor of the electron distribution, γ_p , and the magnetic field, among other things. Prandini and Ghisellini 2022 applied the same analysis as Ghisellini et al. 2017 to a sample of 67 TeV-detected blazars with known redshift, building their SEDs from SSDC data, singling out BL Lacs, computing their luminosity and classifying them in the same gamma-ray luminosity bins adopted by the 2017 paper. The result indicates no strong differences in the overall SEDs from the sources selected from lower-energy gamma rays (Figure 3.1), but there seems to be a higher X-ray luminosity in TeV BL Lacs: a slight deviation from the sequence that, according to the authors, may be explained by wrong redshift estimations, or a selection effect linked to a strong bias in favor of bright X-ray sources, or the existence of “blue quasars”, i.e. FSRQs whose emission happens outside the broad line region (Ghisellini and Tavecchio 2008). Another recent paper (Ouyang et al. 2023), which tries to rebuild the sequence by focusing on the observed TeV luminosity, instead puts it into question following the idea that the trend is actually an effect of different Doppler beaming factors, δ_D : this is based on the comparison between a sample of 48 quiescent blazars and 21 flaring ones (with 11 sources known in both states), which seems to show that the expected anticorrelation exists only during flares, and disappears once correcting for beaming.

The work described in this Chapter was carried out as a contribution to an ongoing effort to reform the blazar sequence focusing on TeV-detected sources, and adopting a different approach than the previous ones. The “traditional” *modus operandi*, followed by the majority of the papers mentioned in the above paragraph, starts from




Figure 3.1: The *Fermi* blazar sequence reproduced with TeV-detected blazars only, from Prandini and Ghisellini 2022.

considering the K -corrected luminosity of sources in a specific frequency band, and classifying it in *luminosity* bins in which spectra are averaged or interpolated to find a single “average” SED of that bin. This approach, while reasonable (especially in a period when data were not as abundant at most frequencies as today), still needs working on different energy bands separately, making assumptions on the spectral shapes, and is “dangerous” as it means blending together sources with possibly different characteristics, all with large flux variabilities. Now that high-quality data are available in most frequency bands, and many known blazars have their synchrotron peak frequencies estimated (e.g. in the 4LAC catalogue, see next Section), the new endeavor this work is part of will proceed like this:

1. select only **TeV-detected BL Lacs**;
2. classify them by their **synchrotron peak frequency** in 5 bins;
3. choose **one** representative source per each bin, and select available spectral data in a **average** emission state, that is most emblematic of the general behavior of the single object and of the sources in the same subgroup;
4. fit the observed SEDs of the representative sources, using the same model for all;
5. discuss potential trends shown by the physical model parameters across the synchrotron peak frequency bins.

In this way, there is no averaging over different sources and emission states, nor the need to reconstruct K -corrected or even intrinsic, beaming-corrected luminosity: the

3.2. DATA SELECTION

final result will hopefully be a sequence of *physical quantities*, which may cast light on the physics behind the phenomenology.

3.2 DATA SELECTION

3.2.1 SOURCE SELECTION

To perform the analysis outlined above, the following information is needed about the objects in the sample:

- the source **class**, in order to select only BL Lacs and discard FSRQs and blazars of uncertain type;
- the **synchrotron peak frequency** ν_{syn} , to classify them;
- the **redshift** z , which is needed not only to discuss its impact on measured fluxes and its possible link to the sequence, but also to perform the correction for absorption by the EBL on VHE gamma-ray spectral data;
- the **spectral points**: frequencies or energies with their respective flux densities and errors;
- related information like the instrument, the observation time, the exposure time, the systematic errors, the emission state of the source when it produced that SED.

The first step was **selecting the sample** over which rebuilding the TeV blazar sequence². Sources were selected in the high-energy gamma-ray band, covered by the 4FGL (Abdollahi et al. 2020) and 4LAC (Ajello et al. 2020) catalogs. As better explained in the next paragraph, they were both built from data collected and averaged over the years by the *Fermi*-LAT instrument in the range 50 MeV to 1 TeV; FGL catalogs include all persistent sources detected by LAT, of which the large majority are blazars; LAC catalogs are instead built from the corresponding FGL to only include AGN. 4LAC-DR2 (Lott, Gasparrini, and Ciprini 2020) was first merged with an updated version that features redshift estimates that are deemed more reliable, as they were found in recent dedicated spectroscopic campaigns for potential target blazars of the future VHE gamma-ray observatory CTA (Goldoni et al. 2021, Kasai et al. 2022, D’Ammando et al. 2024). After filtering out the 259 sources still without a redshift, a new dataset was created by merging the result with 4FGL-DR3 (Abdollahi et al. 2022). This operation, that produced a table of 1703 entries, was done to exploit the information from both catalogs, whose columns do not coincide: more specifically, the source name and TeV counterpart were taken from 4FGL, while ν_{syn} and the radio counterpart from 4LAC. Small corrections were then applied to this dataset: 4 sources were flagged as TeV-detected based

²This part of the work, which was already underway when this thesis was begun, was carried out by dr.ssa Elisa Prandini, dr.ssa Ilaria Viale and dr.ssa Chiara Righi. The collective plot of the selected SEDs, the discussion about redshifts, and the fitting test with altered X-ray data are instead personal work.

on new information from the aforementioned TeVCat, and other 3 TeV sources had their unreliable ν_{syn} corrected with the value from an external tool, BlaST (see Section 4.4.2). At this point, based on values in the TEVCAT_FLAG and CLASS columns, only the **TeV-detected BL Lacs** were selected. Choosing BL Lacs only is based on the fact that the physical model, the synchrotron self-Compton (Section 2.1), does not work for FSRQs, which in addition are rarely detected at very high energies. The final sample consists of **56** sources, whose ν_{syn} distribution is illustrated in Figure 3.2 in comparison with the other BL Lacs and FSRQs of known redshift. Selected sources were classified in syn-




Figure 3.2: Frequency distributions of ν_{syn} for BL Lacs, FSRQs and TeV detected BL Lacs of known redshift in the 4LAC-DR2 catalog. Courtesy of dr.ssa Chiara Righi and dr.ssa Ilaria Viale.

chrotron peak frequency bins, recalling the BL Lac categorization in LBL, IBL, HBL and EHBL objects and its thresholds; one representative source per bin was then chosen on the basis of the abundance and quality of available spectral data:

Bin 1 LBLs, $\nu_{\text{syn}} < 10^{14}$ Hz, **BL Lacertae** ($\nu_{\text{syn}} = 3.86 \times 10^{13}$ Hz)

Bin 2 IBLs, $10^{14} \leq \nu_{\text{syn}} < 10^{15}$ Hz, **TXS 0506+056** ($\nu_{\text{syn}} = 3.55 \times 10^{14}$ Hz)

Bin 3 HBLs, $10^{15} \leq \nu_{\text{syn}} < 10^{16}$ Hz, **PKS 2155-304** ($\nu_{\text{syn}} = 5.69 \times 10^{15}$ Hz)

Bin 4 HBLs, $10^{16} \leq \nu_{\text{syn}} < 10^{17}$ Hz, **PG 1218+304** ($\nu_{\text{syn}} = 1.86 \times 10^{16}$ Hz)

Bin 5 EHBLs, $\nu_{\text{syn}} \geq 10^{17}$ Hz, **PKS 0548-322** ($\nu_{\text{syn}} = 2.29 \times 10^{17}$ Hz)

The frequency histogram of the sample is shown in Figure 3.3 along with an approximate graphical indication of the synchrotron peak frequency classes and the position of each representative object. The average spectra of the five emblematic sources, built from archival data according to the criteria explained in the next Subsection, are plotted together in Figure 3.4, except for BL Lacertae. Clearly, no trend is obviously visible

3.2. DATA SELECTION




Figure 3.3: Frequency histogram of ν_{syn} for the sample of TeV-detected BL Lacs of known redshift. Gray dashed lines indicate the limits of each frequency class. Green solid lines indicate the positions of the 5 representative sources. Courtesy of dr.ssa Ilaria Viale.




Figure 3.4: Observed SEDs of the selected representative sources (except BL Lacertae for Bin 1).

when plotting the observed SEDs, as they do not represent the luminosity on which the original sequence was built: observed fluxes and frequencies would need to be corrected for redshift in each energy band (*K*-correction). To inspect the impact of z on measured spectra, its frequency distribution was plotted in Figure 3.5, limiting the scope to Bin 4 for simplicity. There is about one order of magnitude difference between the nearest

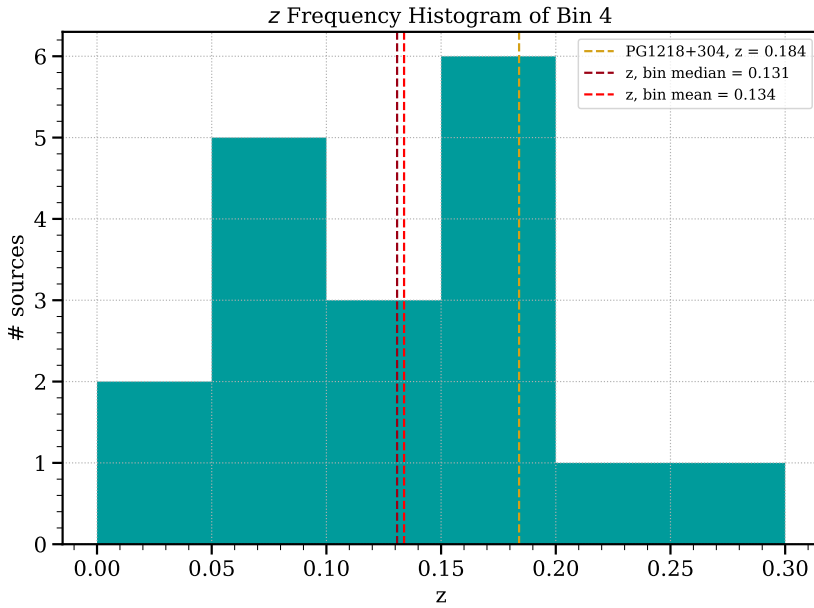


Figure 3.5: Frequency histogram of source redshifts in Bin 4 ($\nu_{\text{syn}} \in [10^{16}, 10^{17}]$ Hz). The vertical lines show the mean (light red) and median (dark red) redshift, and the redshift of the representative source, PG 1218+304 (gold).

and farthest source, with the representative object chosen for this work, PG 1218+304, being slightly farther from Earth than the average. Figure 3.6 displays graphically how the observed SED of PG 1218+304 changes when the redshift is made to vary over the range of values in Bin 4. The spectrum is fitted with a classical SSC model (Section 2.1), obtaining a set of physical parameters that are then kept fixed while the redshift is changed to the two extreme values in Bin 4 ($z = 0.03$, $z = 0.287$). Table 3.1 shows a rough comparison of the synchrotron and SSC peak frequencies and fluxes obtained in this analysis with the *agnpy* package (see Subsection 3.3.1). Notice how all frequencies, apart from ν_{syn} for $z = 0.03$, seem to remain the same for the three values of z . This is due to an issue with discretization in the *agnpy* functions used for computing SED peak frequencies: “true” values of ν_{syn} that are close enough are returned as the same value by *agnpy*. This is also visible in Figure 3.6, whose vertical lines indicate the peak frequencies evaluated by *agnpy*. In any case, the behavior is clear: given the same physical model, observed peak frequencies and fluxes decrease as z increases, as expected; νF_{ν} increases by more than two orders of magnitude when redshift increases by one order. The two peak fluxes vary by the same amount, such that the Compton dominance re-

3.2. DATA SELECTION

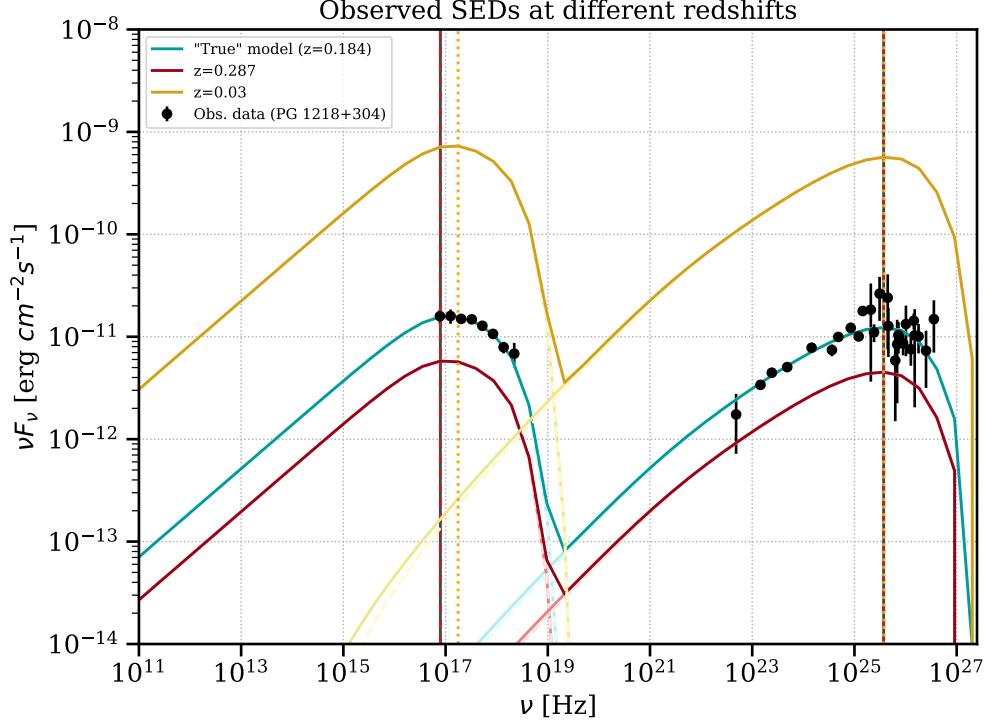


Figure 3.6: Data and SED of PG 1218+304 compared to the SED the same physical model would produce, but at the two extreme redshifts of Bin 4. Vertical lines show the frequencies of the two peaks for the original case (blue), the “lower redshift” model (orange), and the “higher redshift” model (brown).

	z	0.03	0.184	0.287
$\nu_{\text{syn}} [\text{Hz}]$		1.76×10^{17}	7.91×10^{16}	7.91×10^{16}
$\nu F_{\nu}^{\text{syn}} [\text{erg cm}^{-2} \text{s}^{-1}]$		7.28×10^{-10}	1.57×10^{-11}	5.79×10^{-12}
$\nu_{\text{SSC}} [\text{Hz}]$		3.73×10^{25}	3.73×10^{25}	3.73×10^{25}
$\nu F_{\nu}^{\text{SSC}} [\text{erg cm}^{-2} \text{s}^{-1}]$		5.66×10^{-10}	1.23×10^{-11}	4.52×10^{-12}
Compton dominance		0.786	0.786	0.786

Table 3.1: Rough comparison of the synchrotron and SSC peak frequencies of BL Lac SEDs from the same physical model, but observed at different redshifts. The model parameters are from one of the best fits for the Bin 4 representative source PG 1218+304 (spectral points in black): $k_e = 9.91 \times 10^{-9} \text{ cm}^{-3}$, $p_1 = 2.14$, $p_2 = 3.28$, $\gamma_{\text{min}} = 10^2$, $\gamma_{\text{max}} = 10^6$, $\gamma_b = 1.95 \times 10^5$, $t_{\text{var}} = 1 \text{ d}$, $\delta_D = 20$, $B = 2.99 \times 10^{-2} \text{ G}$. The original redshift is $z = 0.184$. The other two are the lowest found in Bin 4, $z = 0.03$, and the highest, $z = 0.287$. The results are subject to the uncertainty due to agnpy estimating binned peak frequencies.

mains the same (leading, e.g., Finke 2013 to focus the discussion on that, as it allows to include BL Lacs with undefined redshift in the blazar sequence).

3.2.2 DATA POINTS SELECTION

INSTRUMENTS AND CATALOGS

At this point, SEDs for modeling were built for each of the 5 representative sources. Points from the infrared to the high-energy gamma-ray band were taken from the **Markarian Multiwavelength Data Center** database (MMDC; see box in Section 1.2.3). VHE gamma-ray data were instead retrieved from the Spectral TeV Extragalactic Catalog (**STeVeCat**, Gréaux et al. 2023), a unified repository of spectral data and other useful information (observation periods, livetime, significance, coordinates, types, redshifts) from 173 papers that include at least 2 spectral points from imaging air Cherenkov telescopes between 1991 and 2021, all re-formatted homogeneously. All data selected for now come from the instruments described below.




Figure 3.7: WISE.
Credit: NASA/JPL.

WISE The Wide-Field Infrared Survey Explorer (WISE; Wright et al. 2010) was a NASA space telescope designed to provide a sensitive all-sky survey in the mid-infrared band. Launched in 2009, it ceased its main mission in 2010 and, after a renewed period of activity (NEOWISE) focused on the near-Earth asteroids, it was shut down in August 2024 and is expected to deorbit by the end of the year³. It worked by scanning along great circles, freezing each frame for 11.002 s such that at least 99% of the sky could be covered by at least 8 frames in 6 months. Data were taken in 4 bands (W1 – 4) centered at 3.4, 4.6, 12 and 22 μm ; the selected data points for this work belong to the W3 and W4 bands, where the integration time per frame was 8.8 s.

Swift The *Neil Gehrels Swift* Observatory (Gehrels et al. 2004) is a NASA mission, launched in 2004 and still active, whose main goal is gamma-ray burst astrophysics. It features three instruments: the Burst Alert Telescope (BAT) in the 15 keV to 195 keV range, the X-Ray Telescope (XRT) in the 0.3 keV to 10 keV interval, and the Ultra-Violet/Optical Telescope (UVOT) in the 170 nm to 600 nm range. BAT is a wide-field of view coded mask telescope with the role of detecting GRBs first: in the process, it is doing an all-sky survey at hard X-ray energies. Lower-energy instruments have a narrower field of

³<https://www.nasa.gov/news-release/nasa-mission-concludes-after-years-of-successful-asteroid-detections/>

3.2. DATA SELECTION

view, and are usually employed for BAT alert follow-ups, but they also dedicate time to individual pointings of persistent sources, including blazars.

UVOT records fluxes in 3 visible (U, B, V) and 3 near-ultraviolet bands: UVW2 (192.8 nm, FWHM=65.7 nm), UVM2 (224.6 nm, FWHM=49.8 nm), UVW1 (260.0 nm, FWHM=69.3 nm). Data selected here were recorded through the UV filters and are taken from the *Swift*/UVOT Serendipitous Source Catalog (Yershov 2014), published to describe single multi-exposure observations of objects detected by the telescope while pointing at other targets between 2005 and 2010; differential fluxes in $\text{erg cm}^{-2}\text{s}^{-1}\text{\AA}^{-1}$ and their errors are included. **XRT** measurements are taken from the 1OUSXB catalog (Giommi et al. 2019), published to collect individual multi-exposure targeted observations of blazars performed between 2004 and 2018. The catalog, produced in the context of Open Universe for Blazars, records fluxes ($\text{erg cm}^{-2}\text{s}^{-1}$) at 4 energies (0.5, 1.5, 3.0 and 4.5 keV) and the interpolated flux at 1 keV from measurements on the whole 0.3 keV to 10 keV range. **BAT** spectra come from the 105-Month *Swift*-BAT All-sky Hard X-Ray Survey (Oh et al. 2018): fluxes in that catalog were extracted in 8 channels from the region of the identified source counterpart in an all-sky mosaic image, obtained in turn by combining all 20 min snapshots (total coverage of at least 7.25×10^6 s of 90% of the sky).




Figure 3.8: *Swift* diagram. Credit: NASA/GSFC.

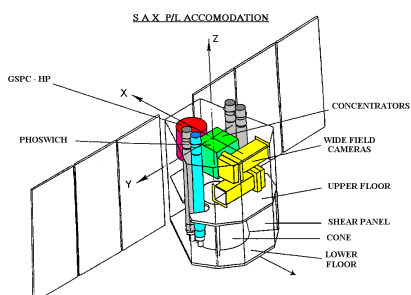


Figure 3.9: BeppoSAX diagram. Credit: ASI, BeppoSAX SDC.

BeppoSAX BeppoSAX (formerly SAX, Satellite per Astronomia a raggi X; Boella et al. 1997) was an ASI/NIVR mission for the observation of the X-ray sky that operated from 1996 to 2002. The spacecraft consisted of two instrument arrays: 2 Wide-Field Cameras and 6 co-aligned Narrow-Field Instruments, including one Low Energy Concentrator Spectrometer (LECS), 3 Medium Energy Concentration Spectrometers (MECS), one High Pressure Gas Scintillation Proportional Counter (HPGSPC) and one Phoswich Detector System (PDS), collectively covering the 0.1 keV to 200 keV range. This work uses a catalog (Giommi et al. 2002) of 157 single observations made by LECS, MECS and PDS in 1996-2001; each exposure lasted $\sim 10^4$ s to 10^5 s.

Fermi The *Fermi* Gamma-ray Space Telescope is a NASA mission, launched in 2008 and still ongoing, devoted to the surveying of the high-energy gamma-ray sky with two instruments: the Large Area Telescope (LAT; Atwood et al. 2009), for scanning and pointings of both persistent and transient sources at higher energies, and the Gamma-ray Burst Monitor (Meegan et al. 2009), for the fast detection of gamma-ray bursts in the 8 keV to 40 MeV range. Fermi-LAT is a pair-conversion telescope featuring a silicon-strip tracker, a scintillation calorimeter and a scintillation anticoincidence detector; it is optimized for the detection of gamma-ray photons in the range 50 MeV to ~ 1 TeV and it mostly works in scanning mode, covering the whole sky in about 3 hours. Spectra for this work were taken from two catalogs of LAT sources: 4FGL-DR3 (Abdollahi et al. 2022), listing persistent objects of any class, and the older 3FHL (Ajello et al. 2017), focusing on sources detected above 10 GeV. In both databases, SEDs are reconstructed from multi-year photon counts (12 years from 2008 to 2020 in 4FGL-DR3, 7 years from 2008 to 2015 for 3FHL), by freezing the spectral index obtained in the global fit over the whole energy interval and the region of interest of each source, and adjusting the normalization in 8 (4FGL-DR3) or 5 (3FHL) energy bands.




Figure 3.10: *Fermi*-LAT diagram, from Atwood et al. 2009.




Figure 3.11: VHE gamma-ray detection, from Sitarek 2022.

IACTs: H.E.S.S., VERITAS and MAGIC

The highest-energy photons can be detected indirectly, among other techniques, by imaging air Cherenkov telescopes (IACTs; see Sitarek 2022). As very-high-energy gamma rays interact with the atmosphere, producing an electromagnetic cascade, secondary particles in the shower produce in turn a flash of Cherenkov light, which can be recorded to obtain information about the energy and arrival direction of the original photon.

Typical IACTs are tessellated optical telescopes that focus radiation on a segmented photomultiplier camera, whose sensitivity and fast response allow to detect Cherenkov light (signals of \sim ns duration and $\sim 100 \text{ cm}^{-2} \text{ s}^{-1} \text{ TeV}^{-1}$ photon differential flux) and whose imaging capabilities allow to reconstruct the electromagnetic shower and the properties of the parent photon. The Cherenkov light cone from electromagnetic cascades usually has a cross section (light pool) of about 120 m radius: several telescopes,

3.2. DATA SELECTION

each with a mirror of ~ 10 m diameter or more, are therefore set up in arrays, that reach a total effective area of $\sim 10^5$ m² and the ability for stereoscopic reconstruction of the cascade. The angular resolution is typically $\sim 0.1^\circ$, the energy resolution is $\sim 15\%$. Systematic uncertainties of the order of 30% are typically considered for measured energies, fluxes and thresholds, due to the calibration and response of the instruments, the reconstruction of the shower (based on complex simulations of the involved particle physics and the conditions of the atmosphere), the simulation and subtraction of the cosmic ray background, and other aspects of the data analysis (Albert et al. 2006, H.E.S.S. Collaboration et al. 2018).

Three major sites are currently operating: H.E.S.S.⁴ (Ohm and Wagner 2023) in Namibia, VERITAS⁵ (Adams et al. 2022) in the United States, and MAGIC (Aleksić et al. 2016a, Aleksić et al. 2016b) on the island of La Palma, Spain. **H.E.S.S.** consists of four 12 m telescopes, active since 2004, and a fifth 28 m telescope added in 2012 to extend the energy range from a few tens GeV to ~ 100 TeV. **VERITAS** was inaugurated in 2007 and features four 12 m telescopes covering the 85 GeV – 30 TeV energy interval. **MAGIC** is a pair of 17 m instruments: MAGIC-I began activity in 2004, MAGIC-II in 2009; together, they inspect gamma-ray sources in the interval between ~ 50 GeV to more than 50 TeV.

These instruments have excellent sensitivities compared to other ground-based gamma-ray telescopes (like surface arrays), but their field of view is limited ($\sim 3^\circ - 5^\circ$), and their duty cycle is only $\sim 10\%$ as they need a perfectly clear and dark sky. Pointed observations of single sources are therefore carried out over several nights, until a certain total exposure time (usually on the order of tens of hours) is reached. Data from IACTs can be found in publications referring to individual observation campaigns, collected in unified repositories like the STeVeCat. The H.E.S.S. website features a page⁶ listing the main detections; VERITAS published the VTSCat catalog in 2023 (Acharyya et al. 2023); a repository of MAGIC data⁷ is hosted by the Port d'Informació Científica in Barcelona. As already mentioned in Section 1.2.4, VHE data need to be deabsorbed to account for the EBL: here⁸ the model in Saldana-Lopez et al. 2021 was employed.

Instrument specifications are summarized in Table 3.2. Sensitivities are included, as they govern the minimum exposure time needed for detection (and is therefore linked to the observation mode and the time sampling of the flux from a source). Future developments may include other observatories that have so far been ignored, like Planck or NuSTAR.

⁴<https://www.mpi-hd.mpg.de/HESS/pages/about/telescopes/>

⁵<https://veritas.sao.arizona.edu/about-veritas/veritas-specifications>

⁶<https://www.mpi-hd.mpg.de/HESS/pages/home/sources/>

⁷<https://magic.pic.es/>

⁸Work done by dott.ssa Ilaria Viale.

Energy/Wavelength range	Instrument	Catalog/Reference	Field of view	Observation type	Sensitivity	No. sources (any class)
$\sim 3 \mu\text{m}$ to $30 \mu\text{m}$ = 0.04 eV to 0.41 eV	WISE	All-Sky Release (2012)? Wright et al. 2010, Cutri et al. 2012	47'	Freeze-frame scanning	0.86 mJy , 5.4 mJy in W3, W4 (5σ)	563921584
170 nm to 600 nm = 2.1 eV to 7.3 eV	Swift-UVOT	UVOTSSC Yershov 2014	17'	Pointing, photon counting mode	$m_B = 24.0$ in white light in 1000 s	6270743
0.3 keV to 10 keV	Swift-XRT	IOUSXB Giommi et al. 2019	23.6'	Pointing, photon counting mode	$10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ (short exposure) $7 \times 10^{-15} \text{ erg cm}^{-2} \text{ s}^{-1}$ (long exposure)	2308 distinct blazars
0.1 keV to 200 keV	BeppoSAX	Giommi et al. 2002	0.5° (LECS; MECS) 1.3° (PDS)	Pointing, photon counting mode	$5 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$ (PDS)	84 blazars, 157 spectra
14 keV to 195 keV	Swift-BAT	BAT105m Oh et al. 2018	$60^\circ \times 100^\circ$	All-sky surveying	$8.40 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1}$	1632
100 MeV-1 TeV	Fermi-LAT	4FGL-DR3 Abdollahi et al. 2022	2.4 sr	All-sky surveying, 12-yr	$4.4 \times 10^{-11} \text{ cm}^{-2} \text{ s}^{-1}$ to $9.5 \times 10^{-11} \text{ cm}^{-2} \text{ s}^{-1}$ ($>10 \text{ GeV}$)	6658
10 GeV-1 TeV		3FHL Ajello et al. 2017		All-sky surveying, 7-yr		1556
$O(10 \text{ GeV})$ -100 TeV	H.E.S.S.	Individual publications; STeVeCat; Hess Source Catalog	5° (H.E.S.S. I) 3.2° (H.E.S.S. II)	Observation campaigns	1% Crab (H.E.S.S. I, 50 h)	95 (138 in TeVCat)
85 GeV-30 TeV	VERITAS	Individual publications; STeVeCat; VTSCat (Acharyya et al. 2023)	3.5°	Observation campaigns	1% Crab in $< 25 \text{ h}$, 10% in 25 min	67 (TeVCat)
50 GeV-50 TeV	MAGIC	Individual publications; STeVeCat	3.5°	Observation campaigns	0.55% Crab (50 h, $\sim 300 \text{ GeV}$)	80 (TeVCat)

Table 3.2: Summary of the instruments and catalogs used to build the SEDs of the selected sources (bins 3-5).

SELECTION CRITERIA

Since the sequence would be the product of a **population study** on the general acceleration and emission mechanisms in blazars, the **average states needed to be selected**. However, this task was not trivial, as blazars are characterized by a strong variability on different timescales (see Section 1.2.2) and this may introduce biases, depending on the features of each instrument. As already mentioned in the previous paragraph, X-ray telescopes, in particular, most often work in “pointing” mode when observing blazars, offering data from single campaigns targeted at a specific source (see, for example, Giommi et al. 2019). High states and X-ray bright objects are therefore favored, both by the *a priori* selection when allocating observation time (Giommi et al. 2002), and by the subsequent data analysis, as their higher signal-to-noise ratios increase their chances to be included in catalogs as standalone sources without confusion. This bias unluckily affects not only the source selection itself (as objects with most and best data were preferred for this study), but it is stronger in the region of the spectrum where the synchrotron peak falls (at least for HBLs): right where variability is more evident, and can influence more heavily the estimate of the Compton dominance, one of the characterizing parameters of the blazar sequence. It would be fundamental, given these issues, to either have fluxes from multi-year surveys by wide field-of-view instruments, which would then be averaged to be included in catalogs, or single observations or campaigns coordinated *simultaneously* among “pointing” instruments at different energies during quiescent states of the blazar. The first case is that of *Fermi*-LAT: as already explained above, both catalogs used here were obtained by averaging fluxes in 5 or 8 energy bins over the total exposure time since the start of the mission, which dilutes the impact of flaring periods and produces a single spectrum. The second case is the trickiest one, as the frequent lack of simultaneity, or even of any information on observation periods⁹, forced to make choices with some degree of arbitrariness. Two criteria were considered:

1. **using a single observation**, as it is. This was the initial approach: selecting TeV and X-ray single observations based on quality, data taking period (for quiescent states), position in the flux distribution for the respective frequency range (to guarantee they represent an average state), and how well they joined with adjacent data points in other bands. This led to selecting data from the IACTs and BeppoSAX.
2. Resorting to a more traditional way: **averaging over different fluxes** measured at the same frequencies. This was applied routinely to *Swift*-UVOT data, and was introduced for *Swift*-XRT measurements as an alternative to single BeppoSAX observations.

The second criterion was adopted after the first SED fits were performed, and it be-

⁹Even if MMDC has a time filtering option, and masking tables by the values in the date columns is an easy task in Python, several VOU-Blazars catalogs (see box in Section 1.2.3) were noticed to only feature placeholder observation times, which are meaningless.

came evident that BeppoSAX data for the X-ray band, despite their high quality, wide frequency range, and clear coherence of the selected observations with points in other bands, constrained the modeling too much right by virtue of their high precision. Their errors are so small that they give them a disproportionate weight in χ^2 -based minimization compared both to other frequency bands, especially VHE gamma rays which are the focus of this work, and X rays themselves, where variability is enormous. Models of spectra of PG 1218+304 taken in simultaneous multiwavelength campaigns during high states¹⁰ produced much smaller values of the Doppler factor δ_D than obtained so far for the data selection as described above. A test was therefore done to check how much a slightly different X-ray data choice, well within the flux variance of archival data in that range, still “well-connecting” to points in other bands, and producing a closer Compton dominance to the flaring spectrum, could change the fit parameters. The result was striking, especially for the δ_D (see, indeed, Section 2.1 for the $\nu F_\nu^{\text{syn}}-\delta_D$ relation), and showed how arbitrary the first approach is. Without discarding the results with BeppoSAX spectral points, a second selection was introduced using the means of *Swift*-XRT fluxes around the peaks of their distributions in each energy bin. In this way the “most

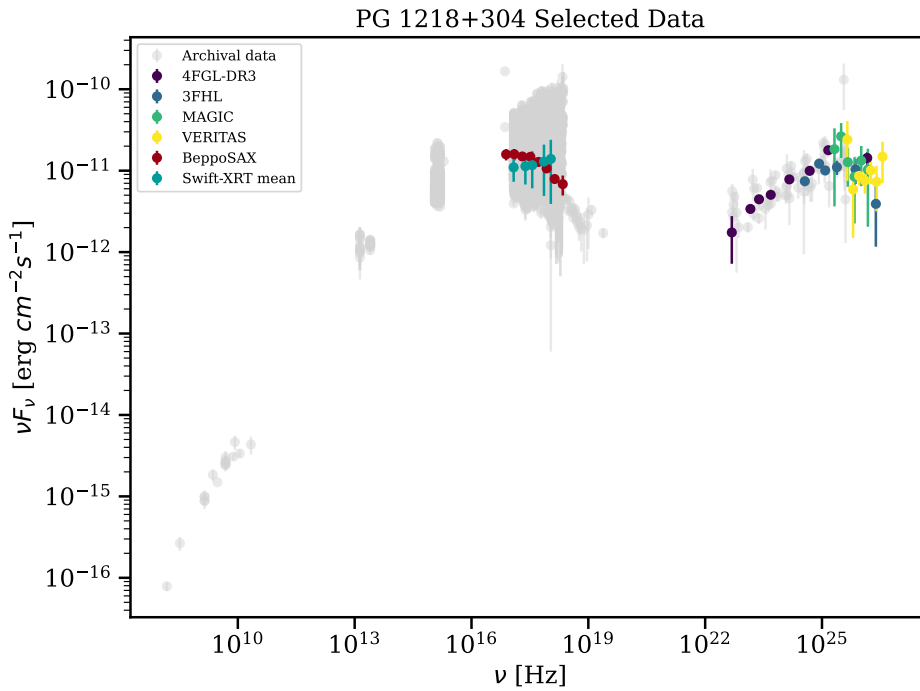


Figure 3.12: Selected SEDs for the example source, PG 1218+304, compared to archival data (gray). BeppoSAX (dark red) and mean *Swift*-XRT (teal) data are mutually exclusive alternatives under the two adopted selection criteria, while higher-energy points are in common.

frequent” emission state, hopefully the average, is approximated. The variance of the

¹⁰Done by Julia Nguyen, who joined the group briefly as an intern.

3.3. TOOLS

averaged sample is attached to each new data point; uncertainties become larger, and more emblematic of variability in that energy range, but the bias in favor of higher flux states is still present. *Swift*-XRT was chosen on account of the abundance of available data from it. Figure 3.12 compares the two data selections for the Bin 4 source, PG 1218+304. Notice the absence of ultraviolet, optical, infrared and radio data in the selection in this case. Mid-infrared and UV measurements from WISE and *Swift*-UVOT, respectively, were included or not depending on the source. In this example, WISE data were discarded because they looked “anomalous” (two very close points, with almost the same flux) and, being the only low-energy data and having very small errors, the constraint they put on the fit was too strong. Besides this, although WISE was a surveying instrument, it took data over a period of few months, comparable to flaring timescales, and the object for which its points were included (an apparently stationary blazar, indeed) was observed in 19 8.8 s-exposures over less than two days. The optical band between 6×10^{13} Hz and 10^{15} Hz was always deleted from the dataset to avoid the contamination of the jet spectrum by the host galaxy or components of the active nucleus. The aforementioned fits of the PG 1218+304 SED during a flare suggested that WISE points could be a low-energy tail of these additional components. The radio band was always ignored, as it needs to be modeled outside of the SSC framework, being the superposition of synchrotron curves from different zones (Chapter 2.2).

3.3 TOOLS

Two tools, based on different working principles, were used to model the selected SEDs: the **agnpy** Python package and the **MMDC SED fitting** online tool. As explained later more in detail, the outputs are not directly comparable except for a few parameters, so they were considered complementary to each other, with cross-checking limited to a small number of fits.

3.3.1 AGNPY

agnpy (Nigro et al. 2022) is a recent open-source Python package specifically built for modeling the broadband SEDs of jetted AGN, both aligned (blazars) and misaligned (radio galaxies). Its focus is on leptonic non-thermal processes happening in the jet, but it is peculiar in providing models for thermal and line emitters, especially when considered as soft photon sources for inverse Compton scattering or $\gamma\gamma$ absorption, including on the cosmic microwave background and the EBL. Recent updates introduced the simplest hadronic process, proton synchrotron. Differently from other similar packages, it is not designed for time evolution, but can only set constraints on parameters after the characteristic timescales. **agnpy** is primarily for theoretical spectral modeling: it can

compute SEDs and some of their properties starting from a given model of the emission region, the non-thermal electron distribution, external photon field sources and absorption processes. The flowchart in Figure 3.13 summarizes this process.

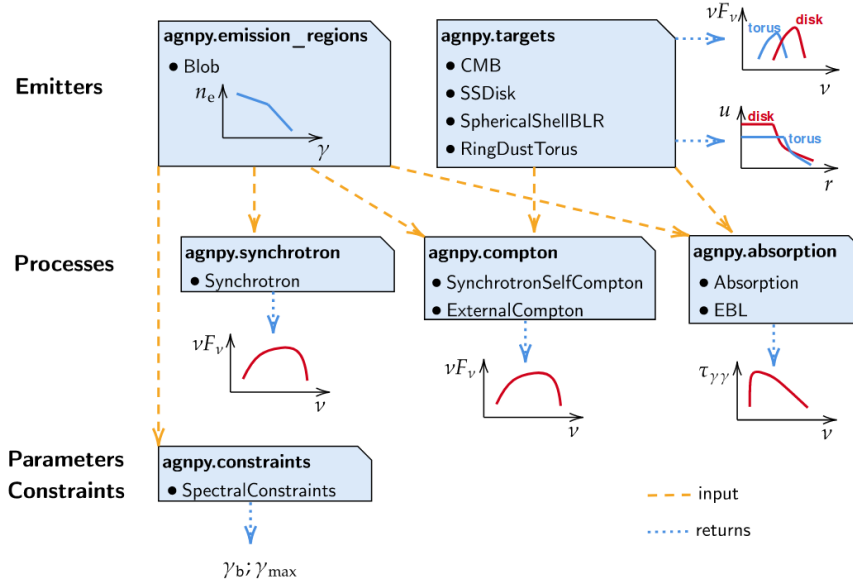


Figure 3.13: Scheme of `agnpy` modules interactions and outputs, from Nigro et al. 2022.

Fits on experimental data cannot be done natively, but it is possible to wrap `agnpy` objects in either `Gammapy` or `Sherpa` code for this purpose. Both were used in this work. **Gammapy** (Donath et al. 2023) is an open-source Python package developed as a prototype, or core library, of the *CTA Science Analysis Tools* software, which will serve as a complete data analysis platform for the upcoming Cherenkov Telescope Array (CTA: The Cherenkov Telescope Array Consortium et al. 2019); `Gammapy` is, however, already in use for existing high- and very high-energy gamma-ray telescopes. **Sherpa** (Burke et al. 2023) is a Python package for modeling and fitting data, originally conceived for analyzing spectral and imaging data from the Chandra X-ray telescope, but again extended to general use in astrophysics. An `agnpy` class called `fit.SynchrotronSelfComptonModel` is initialized on the electron number density distribution (from the `emission_regions` module). To create the experimental dataset, spectral points are read from an `.ecsv` file and passed as an argument to another class of the `.fit` module, either `load_sherpa_flux_points` or `load_gammapy_flux_points`, together with the fit energy range and a dictionary of systematic uncertainties. In the case of `Gammapy`, the `agnpy` spectral model is converted to a `Gammapy SkyModel`, which is then set as the `.models` attribute of the `agnpy` dataset. At this point, after eventually adjusting settings like the fit statistic and minimizer, the fit can be performed with the `Fit` class of either backend.

Considering that this package does not compute the time evolution of particle dis-

3.3. TOOLS

tributions, the **standard choice** to describe the **electron distribution in the stationary state** was adopted: the **broken power law**,

$$\frac{dn_e}{d\gamma} = k_e \left[\left(\frac{\gamma}{\gamma_b} \right)^{-p_1} H(\gamma; \gamma_{\min}, \gamma_b) + \left(\frac{\gamma}{\gamma_b} \right)^{-p_2} H(\gamma; \gamma_b, \gamma_{\max}) \right] \quad (3.1)$$

with H the Heaviside function. In this scenario, the break accounts for synchrotron cooling. Given the SSC model described in Chapter 2.1, the fit parameters will therefore be:

k_e Normalization constant of the electron **number density** distribution [cm^{-3}]

p_1 **Low-energy power-law index** of the electron distribution (before the break)

p_2 **High-energy power-law index** of the electron distribution (after the break)

γ_{\min} **Minimum Lorentz factor** of electrons

γ_b **Lorentz factor** of electrons at the **spectral break**

γ_{\max} **Maximum Lorentz factor** of electrons

δ_D **Doppler factor** of the emission region

B **Magnetic field** of the emission region [G]

t_{var} Observed **variability**/dynamical **timescale** [s]

Except for t_{var} , all the quantities are in the **reference frame of the emission region**.

3.3.2 MMDC

The Markarian Multiwavelength Data Center¹¹, already described in Chapter 1, is not only an online platform for retrieving and filtering spectral data and literature on known blazars, but it also provides tools to analyze said data. Its homepage includes an extremely easy interface through which users can build theoretical models by giving parameters as inputs, compare them manually with observed SEDs uploaded by themselves as a .csv file, or fit such spectra after choosing the model type (SSC or external Compton), providing a value for the redshift, an indication whether VHE gamma ray data need EBL de-absorbing or not, a fixed minimum Lorentz factor of the electron distribution (in the SSC case) optionally, and an e-mail address. After about half a hour, an e-mail message will notify the user that the fit results are ready to be accessed through a dedicated link. The same interface will appear providing best-fit parameters, the final

¹¹<https://mmdc.am/>

plot including an “error region” given by random intermediate realizations of the fitting curve obtained during the optimization process, and the possibility to download a .png image of the plot, a .pdf file of the corner plot, and .csv files of the fitting curve (as a 2-d array of points for plotting) and of the best parameters and their errors.

Apart from this extreme user-friendliness, an innovative aspect of this tool is the modeling procedure itself, described in Bégué et al. 2023 for the synchrotron self-Compton, and Sahakyan et al. 2024b for the external Compton (the hadronic option is in preparation). In order to tackle the issue of building spectral models that are becoming more and more resource-intensive, the developers based their product on a **machine learning** algorithm coupled to Bayesian inference.

Basics of Learning - part 1: Neural Networks

In analytical tools, classes and functions take parameter values and the arrays of the independent variables as arguments, and apply a given formula to return the dependent variable as an output: in other words, they follow instructions and are limited to the exact mathematical relation they need to apply. Machine learning, or statistical learning, instead shifts its focus on the data: a learning model is *a function that uses information from an input dataset to optimize a set of parameters or a “learning structure”, such that it will become able to make predictions that are as close to an expected output as possible* (Michelucci 2022). The MMDC tool is based on a specific type of learning algorithm called *convolutional neural network* (CNN). The structure can be seen in Figure 3.14.




Figure 3.14: Structure of a deep feed-forward neural network. (a): general structure. (b): scheme of the first two layers. From Michelucci 2022.

Neural networks in general need an optimization algorithm to tune a large set of parameters (*weights*) by minimizing a so-called *loss function*, that encodes the departure of model predictions from expected outputs; weights are the coefficients of a linear combination of the input quantities, to which a *non-linear activation function* is applied to predict the output. The weights and the activation function constitute the basic unit of a neural network, called a *neuron*. In

the architecture called *deep feed-forward network*, the training data form an *input layer*, linked through weights to several uncorrelated parallel neurons that form a first *hidden layer*. Their outputs serve in turn as inputs of another layer, and so on, until the last hidden layer feeds the final *output layer*. Convolutional neural networks are a type of feed-forward deep network that is particularly suited to multi-dimensional data types: a technical description can be found in Michelucci 2022.

Implementation in the MMDC Tool To build an unbiased training dataset for the MMDC network, 2×10^5 SSC blazar spectra were simulated by the MMDC team with the help of a package, *SOPRANO* (Gasparyan, Bégué, and Sahakyan 2021), that works by computing the time evolution of photon and electron distributions until reaching the final equilibrium solution, given an electron injection function for electrons and the terms for the main radiative processes at play (synchrotron cooling, inverse Compton cooling, Compton scattering, pair production). The end user therefore does not need to make assumptions on which equilibrium distribution electrons will have. Model parameters are sampled from given intervals in a way that ensures simulations cover uniformly the whole spectrum of plausible combinations. *SOPRANO* outputs and the corresponding model parameters, the input dataset, is then split into three subsets (80% for training proper, 10% for validating, 10% for testing: see Chapter 4) and fed to an 8-layer convolutional neural network, designed to reproduce spectra in 150 energy bins (from 10^{-2} Hz to 10^{30} Hz: Gasparyan, Bégué, and Sahakyan 2021). The advantage of this approach is that, while the training phase takes $\sim 2 \times 10^4$ core hours of computation time (about two weeks), it only needs to be performed once, and subsequent predictions based on new given parameters will only take a few milliseconds: a large improvement over the performance of traditional analytical tools (for example, *SOPRANO* itself takes between a few tens of seconds to a few minutes, similarly to *agnpy*). The MMDC modeling tool is therefore particularly suited to be used iteratively: for example, it can be integrated into a fitting package to boost its performance by computing the intermediate realizations of the SED model. MultiNest (Feroz, Hobson, and Bridges 2009) was chosen for this purpose. Developed specifically for the fields of astrophysics, cosmology and particle physics, MultiNest employs Bayesian inference to estimate model parameters and compute posterior probability distributions and credible intervals. This is particularly convenient in cases like blazar spectral fitting, where the parameter space is multi-dimensional, multimodal and degenerate. This tool would then be able to find the global best fit, without risking to be stuck in a local best-likelihood region around a given initialization as it may happen with other fitters: this has the practical consequence that the end user does not need to provide anything to the web interface, apart from the spectral data and the

redshift of the source.

The version of the network exploited in this work was trained on SSC spectra, evolved by *SOPRANO* from an electron injection function in the form of a **simple power law with exponential cutoff**:

$$Q_e = \begin{cases} Q_{e,0} \gamma^{-p} \exp\left(-\frac{\gamma}{\gamma_{\max}}\right) & \text{if } \gamma \geq \gamma_{\min}, \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

with $Q_{e,0}$ defined such that the electron luminosity is:

$$L_e = \pi R^2 \delta_D^2 m_e c^3 \int_1^\infty \gamma Q_e d\gamma \quad (3.3)$$

The fit parameters are therefore:

L_e Electron **luminosity**, in the range 10^{42} erg s $^{-1}$ to 10^{48} erg s $^{-1}$

p Electron injection function **spectral index**, in the range 1.8-5

γ_{\min} **Minimum Lorentz factor** of the electron injection function, in the range 3.16×10^2 - 10^5

γ_{cut} **Cutoff**/maximum **Lorentz factor** of the electron injection function, in the range 10^2 - 10^8

δ_D **Doppler factor** of the emission region, in the range 3-50

B **Magnetic field** of the emission region, in the range 10^{-3} G to 10^2 G

R **Radius** of the emission region, in the range 10^{15} cm to 10^{18} cm

3.4 MODELING

3.4.1 IMPLEMENTATION

Blazar modeling is notoriously plagued by **parameter degeneracy**. Even a model like the SSC, where the number of actors at play is minimum, has quite a large number of correlated free parameters, whose values may combine differently to produce equally acceptable fits of the experimental data. Degeneracy would make it difficult to understand if there actually are some physical quantities that lead to the blazar sequence, while others stay roughly the same across different values of v_{syn} : it would all depend from which “parameter family” is chosen. For this reason, the consolidated approach of **stabilizing fits by fixing some selected parameters**, and maybe attempting

3.4. MODELING

to keep them across bins, was adopted, even if that would be interpreted as arbitrarily “imposing” *a priori* that they are not drivers of the sequence.

The MMDC tool is built to scan the whole range of parameters to find the global best fit in a situation where degeneracy is an important issue, so it would be the best choice if the actual physical quantities need to be found. Its outcomes, however, showed that some parameters are still very ill-defined (the δ_D in particular), with enormous errors, such that pinpointing a number of them to fix in order to obtain more “stable” results (and less degeneracy, even if the MultiNest fitter is already made for this purpose) could be beneficial. This is not possible for the ordinary user, who can fix γ_{\min} at most. Moreover, MMDC is still not ready for including systematic errors on experimental data, and to apply “composite” models like the ones that could be useful for “transitional blazars”. *agnpy* fitting was therefore used to complement MMDC, mainly because it gives more control over the parameters by allowing to fix or constrain them in various ways. The process was:

1. Fit the SED of the **Bin 4** source, PG 1218+304 using MMDC
2. Repeat using MMDC with fixed γ_{\min}
3. Fit the same spectrum using *agnpy*/Sherpa, with fixed γ_{\min} and γ_{\max}
4. Fit again using *agnpy*/Sherpa after fixing t_{var}
5. Scan over different fixed values of δ_D and choose the best realization
6. Fit the same source (different data selection) keeping the same parameters fixed
7. Compare with MMDC fit of the new data selection (γ_{\min} fixed)
8. Repeat the last two steps with data from **Bin 3** (PKS 2155-304) and **Bin 5** (PKS 0548-322).

Since *agnpy*/Sherpa/Gammapy and the original training tool *SOPRANO* work differently, there are few parameters that can be directly compared between the former and MMDC. *agnpy* was also initialized with a systematic errors dictionary, following the advice from the official documentation: 0.05 for optical-infrared points (WISE), 0.10 for high-energy data (BeppoSAX, *Swift*, *Fermi*-LAT), 0.30 for very-high-energy gamma rays (H.E.S.S., VERITAS, MAGIC). They are relative errors on the fluxes, and are summed in quadrature to the statistical errors. Two fits were done with Gammapy, as well as Sherpa, to cross-check results. The default settings were kept for both backends: Levenberg-Marquardt algorithm, χ^2 statistic and Covariance method to compute errors in Sherpa; MINUIT MIGrad χ^2 minimizer and MINUIT HESse to calculate the error matrix for Gammapy.

3.4.2 RESULTS

BINS 1 AND 2: LBLs AND IBLs

The first two bins, including the lowest-energy peaked BL Lacs, are not covered in this work, and are being analyzed separately by the other people involved in this reform of the blazar sequence. The reason is the issue of “FSRQ-like” BL Lacs, as already outlined in Section 1.2.3. More specifically, this also applies to the sources chosen for fitting in these two bins.

“**Bin 1**” (LBLs, $\nu_{\text{syn}} < 10^{14}$ Hz) is represented by the prototype of the whole subclass of blazars, **BL Lacertae**, which however seems to be best modeled by a two-zone SSC+EC model, instead of the “pure” one-zone SSC usually applied to incontestable BL Lacs, whenever its emission state is characterized by a high Compton dominance or TeV flux. SSC seems unable to decently explain the broadband variability of BL Lacertae, whose spectrum also features $H\alpha$ and $H\beta$ emission lines in different periods, indicative of a prominent broad-line region, an unusual occurrence in BL Lacs (Sahakyan and Giommi 2022 and references therein). For this source, data from the low-state multiwavelength campaign described in Abdo et al. 2011b are going to be used.

The chosen source for “**Bin 2**” (IBLs, 10^{14} Hz $< \nu_{\text{syn}} < 10^{15}$ Hz) is the famous **TXS 0506+056**. Its $\sim 3\sigma$ -level association to the high-energy neutrino IceCube-170922A, observed by IceCube on September 22nd, 2017 (The IceCube Collaboration et al. 2018) induced extensive multiwavelength observation campaigns both during the main flare in the days following the neutrino observation, and during average-low states in the subsequent years (Acciari et al. 2022); moreover, the IceCube collaboration found evidence of a possible neutrino flare in 2014-2015 which seems to not having been accompanied by an increased flux in the electromagnetic spectrum (IceCube Collaboration et al. 2018), complicating the picture. As shown in Acciari et al. 2022, accounting for neutrino emission implies adopting a more complex model than the classic one-zone SSC: that paper uses a spine-layer scenario where the gamma-ray band is dominated by inverse Compton on sheath photons, and lepto-hadronic contributions (Bethe-Heitler cascade, pion decay) take a role in fine-tuning the model. Despite the model in Ansoldi et al. 2018 highlights physical parameters that are compatible to what expected for BL Lacs, up to the external Compton/hadronic “correction”, some authors (Padovani et al. 2019) argue that TXS 0506+056 is not even a BL Lac, but a “masquerading FSRQ”, as explained in Section 1.2.3, on account of several features (radio luminosity, emission lines, Eddington ratio) that also challenge its position in the blazar sequence: fitting the SED of this object with a SSC model would therefore be inadequate in any case, even when not accounting for hadronic processes. Given that simultaneous average-state observations are available from Acciari et al. 2022, that data points will be used to model this source separately from this work.

BIN 4: HIGH-END HBLs

The fourth synchrotron peak frequency bin (“**Bin 4**”) includes sources where 10^{16} Hz $\leq \nu_{\text{syn}} < 10^{17}$ Hz, and are therefore considered the “almost-extreme” HBLs. Being away from the “FSRQ-like” objects of the lower-energy bins, but still not belonging to the extreme BL Lac subfamily with all its peculiarities, these sources are presumably the best starting point to apply the SSC model in fitting; moreover, this bin is the most populated one (18 entries), allowing (sub-)population studies (like in the redshift case of Section 3.2) for a future fine-grained analysis of the sequence limited to that single bin. For this reason, Bin 4 is the first to be treated in this work: its representative source, as already mentioned in previous Sections, is **PG 1218+304**. Appearing in literature for the first time in a 1970 catalogue of radio sources observed at the 408 MHz frequency by the “Croce del Nord” telescope near Bologna (Colla et al. 1970), it was later detected on the full electromagnetic spectrum, including at VHE (> 100 GeV) by MAGIC in 2005 (Albert et al. 2006) and VERITAS in 2006-2007 (Fortin 2008 and Acciari et al. 2009, including the first TeV detection). Its redshift (from NED) is $z = 0.184$.

First Data Selection Data points selected for the average SED of this source are:

- **VHE gamma rays:** discovery papers, Albert et al. 2006 (MAGIC) and Acciari et al. 2009 (VERITAS).
- **HE gamma rays:** *Fermi*-LAT 4FGL-DR3 (Abdollahi et al. 2022) and 3FHL catalogs (Ajello et al. 2017).
- **X rays:** Single BeppoSAX observation from July 12th, 1999 (Giommi et al. 2002).

As already anticipated, the **MMDC** SED fitting tool was first ran as a starting point. Figure 3.15a shows the resulting best-fit spectral model with the original data points and one in 10 randomly selected samples from the Bayesian sampling; best-fit parameters are in the first column of Table 3.3. The corner plot (Figure 3.15b) highlights

MMDC fits	All free	Fix $\log(\gamma_{\text{min}}) = 2$	Fix $\log(\gamma_{\text{min}}) = 3$
$\log(B)$ [G]	-2.1 ± 0.5	-1.3 ± 0.6	-2.0 ± 0.6
$\log(L_e)$ [cm^{-3}]	44.5 ± 0.4	44.4 ± 0.5	44.6 ± 0.4
$\log(\gamma_{\text{min}})$	3.5 ± 0.7	2	3
$\log(\gamma_{\text{cut}})$	6.4 ± 0.5	5.5 ± 0.5	5.9 ± 0.5
$\log(R)$ [cm]	17.6 ± 0.6	16.4 ± 0.7	17.1 ± 0.6
δ_D	24 ± 10	35 ± 10	31 ± 10
t_{var} [s]	$(6 \pm 9) \times 10^5$	$(3 \pm 4) \times 10^4$	$(2 \pm 2) \times 10^5$
p	2.4 ± 0.2	2.1 ± 0.2	2.3 ± 0.2

Table 3.3: Best-fit parameters obtained by the MMDC SED fitting tool on the selected spectral data for PG 1218+304, keeping all parameters free or fixing $\gamma_{\text{min}} = 10^2, 10^3$ (dark red). The t_{var} (teal) is not a fit parameter, but it was computed from R and δ_D to facilitate comparison with the agnpy fits.

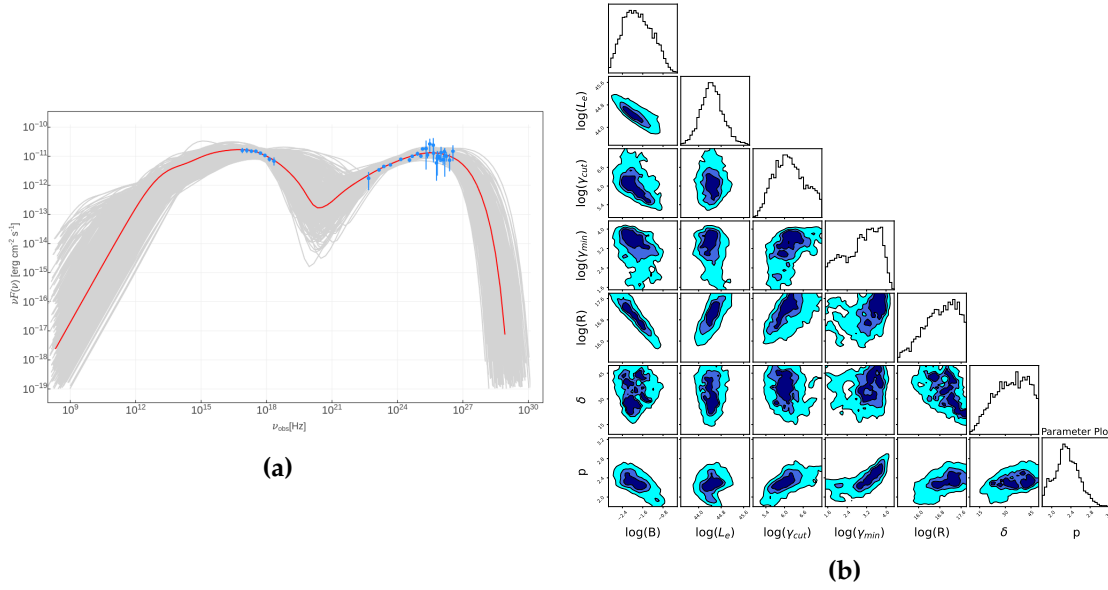


Figure 3.15: Results of the MMDC tool fit on PG 1218+304, keeping all model parameters free. 3.15a: selected SED (blue dots), best-fit model (red line) and random Bayesian samples (gray lines). 3.15b: corner plot of the same fit.

the marginalized posterior probability distributions of the free parameters, to have a first-look estimate of the confidence intervals and pairwise covariances between model quantities. Some pairs of parameters indeed seem to be linearly anti-correlated (for example, $\log(B)$ and $\log(R)$, or $\log(B)$ and $\log(L_e)$). Several 2-d projections of the distributions (off-diagonal contour plots) look “patchy”, underlining parameter degeneracy (existence of more than one family of solutions); this is a known issue in blazar modeling, and is reasonably expected when fitting a large number of free parameters. This is also evident in the 1-d marginalized probability distributions of some parameters, especially $\log(R)$ and δ : one may already expect that t_{var} estimates, when this quantity is left free, will consequently have a large uncertainty, and that they will largely depend on the choice of data points around the synchrotron peak (X-ray band), as evidenced in Section 2.1 while discussing the $\nu F_{\nu}^{\text{syn}} \sim \delta_D^4$ dependence, and anticipated in Section 3.2. After fixing γ_{min} at either 10^2 or 10^3 (typical orders of magnitude found for HBLs and extreme HBLs: see e.g. Abdo et al. 2011a, Sahakyan 2020), free parameters do not look better constrained than before, and some of them change dramatically, as evidenced by the latter two columns of Table 3.3, but the correlation between $\log(R)$ and $\log(L_e)$ (through $\log(B)$) emerges more clearly. Plots can be found in Appendix A.

From the MMDC results it is clear that leaving all the parameters free in fitting, even if it allows to search for the global best fit, will still lead to some instability, with several quantities that have errors so large that they lose meaning: the lower end of the 1σ error interval for t_{var} would even lead to unphysical negative times. This confirmed the need to fix some selected parameters. Doing this needed **switching to agnpy. The**

Sherpa backend was used for most fits, as it was shown to be faster and more stable than Gammapy when free parameters are many and initializations are too far from the data points.

The first step was to fix γ_{\min} and γ_{\max} . The latter parameter was set at 10^6 , more or less halfway (in logarithmic scale) between the values of γ_{cut} in Table 3.3 and often taken as a standard, reference value for that parameter in BL Lac modeling (see e.g. Acciari et al. 2020 for extreme HBLs, and Zhao et al. 2024¹²). Fixing it was deemed necessary, as Sherpa has often shown the weird tendency to treat this parameter as frozen, with $\sigma_{\gamma_{\max}} = 0$, even when it was explicitly thawed. γ_{\min} was set at either 10^2 or 10^3 . The Doppler factor, while left free, was initialized at $\delta_D = 10, 15, 20, 30, 40$. Results are shown in Table 3.4 for $\gamma_{\min} = 10^3$, a choice privileged here due to slightly better *Q-values*¹³ and the fact that, apart from the “problematic” parameters δ_D and t_{var} , there seems to be a general convergence towards similar values for all the other quantities, independently of the initialization. Results from fixing $\gamma_{\min} = 10^2$ can be found in Appendix A. It must be noted that these results are very approximate, since even Sherpa becomes unstable with so many free parameters, and can only compute a rough estimate of the errors.

δ_D initialization	10	15	20	30	40
$\log k_e [\text{cm}^{-3}]$	-8.2 ± 0.5	-8.2 ± 0.5	-8.0 ± 0.6	-7.8 ± 0.6	-7.8 ± 0.9
p_1	2.21 ± 0.07	2.21 ± 0.07	2.24 ± 0.05	2.26 ± 0.09	2.26 ± 0.06
p_2	3.4 ± 0.8	3.4 ± 0.8	3.5 ± 0.8	3.7 ± 0.5	3.7 ± 0.7
$\log \gamma_b$	5.4 ± 0.2	5.4 ± 0.2	5.3 ± 0.2	5.3 ± 0.2	5.3 ± 0.2
$\log \gamma_{\min}$	3	3	3	3	3
$\log \gamma_{\max}$	6	6	6	6	6
δ_D	20.6 ± 0.7	20.6 ± 0.7	27 ± 3	34 ± 12	33 ± 12
$\log B [\text{G}]$	-1.5 ± 0.3	-1.5 ± 0.3	-1.5 ± 0.5	-1.5 ± 0.4	-1.5 ± 0.4
$t_{\text{var}} [\text{s}]$	$(8 \pm 6) \times 10^4$	$(8 \pm 7) \times 10^4$	$(4 \pm 4) \times 10^4$	2.1×10^4	$(2 \pm 3) \times 10^4$
$\chi^2/d.o.f.$	0.761	0.761	0.750	0.744	0.744
Q-value	0.807	0.807	0.820	0.827	0.826

Table 3.4: Fit parameters and statistics for the SED of PG 1218+304, obtained with agnpy and Sherpa by fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$ and initializing δ_D at various values while keeping it free. Values in dark red are fixed. Values in italics only have the first approximate estimate of their errors.

Since these **outcomes are still unstable and depend strongly from the δ_D initialization, something more needs to be fixed. The obvious choice is t_{var}** , as early studies

¹²This recent paper, published while work on this thesis was already underway, also presents population studies on HBLs based on broadband SED fitting, but with a less specific data selection.

¹³Defined as “measure of the probability that one would observe the reduced statistic value, or a larger value, if the assumed model is true and the best-fit model parameters are the true parameter values” (Sherpa online documentation: <https://sherpa.readthedocs.io/en/latest/index.html>). Being a probability, the Q-value is always equal or smaller than 1: simply put, the closer to 1 it is, the better.

on PG 1218+304 TeV emission obtained an observational estimate (more precisely, an evaluation of the characteristic flux doubling time) of $t_{\text{var}} \lesssim 1 \text{ d} = 8.64 \times 10^4 \text{ s}$ (Acciari et al. 2010). Table 3.4 would already hint at a value of $\delta_D \lesssim 20$, luckily matching the most stable fit outcomes. Results are shown in Table 3.5 for fixed $\gamma_{\text{min}} = 10^3$, $\gamma_{\text{max}} = 10^6$ and $t_{\text{var}} = 1 \text{ d}$, and for different initializations of δ_D (again left free). Fits now tend to

δ_D initialization	10	15	20	30	40
$\log k_e [\text{cm}^{-3}]$	-8.2 ± 0.5	-8.2 ± 0.5	-8 ± 1	-8.0 ± 0.6	-8.0 ± 0.5
p_1	2.21 ± 0.07	2.21 ± 0.07	2.2 ± 0.1	2.2 ± 0.1	2.20 ± 0.10
p_2	3.3 ± 0.3	3.3 ± 0.3	3.1 ± 0.2	3.1 ± 0.2	3.08 ± 0.05
$\log \gamma_b$	5.3 ± 0.2	5.4 ± 0.2	5.2 ± 0.7	5.2 ± 0.3	5.2 ± 0.2
$\log \gamma_{\text{min}}$	3	3	3	3	3
$\log \gamma_{\text{max}}$	6	6	6	6	6
δ_D	20.6 ± 0.7	20.6 ± 0.7	25 ± 1	25 ± 1	25 ± 14
$\log B [\text{G}]$	-1.56 ± 0.09	-1.56 ± 0.09	-1.72 ± 0.08	-1.7 ± 0.1	-1.7 ± 0.4
$t_{\text{var}} [\text{s}]$	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4
$\chi^2/d.o.f.$	0.734	0.734	0.738	0.738	0.738
Q-value	0.843	0.843	0.839	0.839	0.839

Table 3.5: Fit parameters and statistics for the SED of PG 1218+304, obtained with *agnpy* and *Sherpa* by fixing $\gamma_{\text{min}} = 10^3$, $\gamma_{\text{max}} = 10^6$ and $t_{\text{var}} = 1 \text{ d}$, and initializing δ_D at various values while keeping it free. Values in dark red are fixed. Values in italics only have the first approximate estimate of their errors.

collapse toward two possible parameter combinations: one with $\delta_D = 20.6 \pm 0.7$ and $\log(B[\text{G}]) = -1.56 \pm 0.09$ ($B = 0.028 \pm 0.006\text{G}$), and another one with $\delta_D = 25 \pm 1$ and $\log(B[\text{G}]) = -1.7 \pm 0.1$ ($B = 0.019 \pm 0.005\text{G}$). All other parameters are well compatible between the two solutions, and Q-values are close. It seems now that this has **definitely become an issue of choosing δ_D to further stabilize the fit**, with due caution in the light of its dependence from the bulk Lorentz factor, Γ and of the existence of blazar sequence models that give a fundamental role to the observing angle, as already discussed in previous sections and chapters.

To explore how the choice of δ_D changes the other parameters more systematically, several fits were performed by fixing different values of δ_D , beyond the already frozen quantities γ_{min} , γ_{max} and t_{var} . Given the results from the previous fits, $\delta_D = 25$ was added to the list of considered values, while $\delta_D = 10$ was excluded as fits were completely off the observed data. Table 3.6 shows the results for $\gamma_{\text{min}} = 10^3$. The behavior of the free fit parameters with respect to the choice of δ_D is shown in Figure 3.16. $\log \gamma_b$ and $\log B$ are perhaps the two parameters that change more clearly with the value of δ_D : in both cases, a stronger beaming and/or higher jet speed (higher δ_D) is compensated by an earlier cooling break and a weaker magnetic field. Considering the error bars are often wide enough to make values of the same quantity generally compatible for different δ_D , other trends are less evident: the electron number density k_e vaguely increases, while p_2 seems to decrease, indicating a slightly harder electron spectrum af-

3.4. MODELING

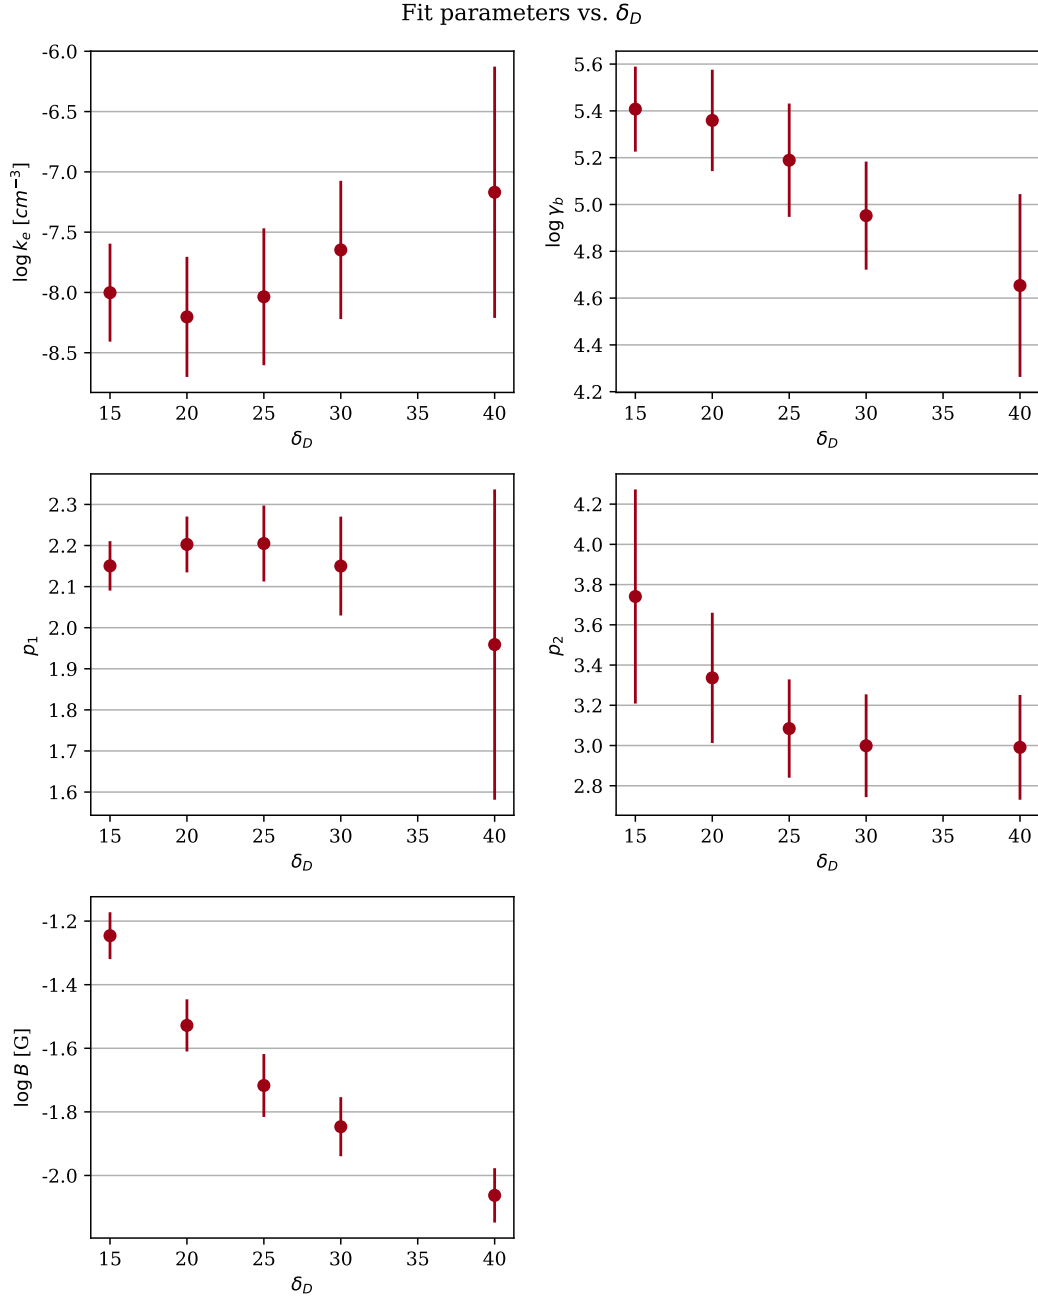


Figure 3.16: Scatter plots of the free fit parameters for PG 1218+304 (Table 3.6) with respect to the fixed values of δ_D .

$\log k_e [\text{cm}^{-3}]$	-8.0 ± 0.4	-8.2 ± 0.5	-8.0 ± 0.6	-7.6 ± 0.6	-7 ± 1
p_1	2.15 ± 0.06	2.20 ± 0.07	2.20 ± 0.09	2.2 ± 0.1	2.0 ± 0.4
p_2	3.7 ± 0.5	3.3 ± 0.3	3.1 ± 0.2	3.0 ± 0.3	3.0 ± 0.3
$\log \gamma_b$	5.4 ± 0.2	5.4 ± 0.2	5.2 ± 0.2	5.0 ± 0.2	4.7 ± 0.4
$\log \gamma_{\min}$	3	3	3	3	3
$\log \gamma_{\max}$	6	6	6	6	6
δ_D	15	20	25	30	40
$\log B [\text{G}]$	-1.25 ± 0.07	-1.53 ± 0.08	-1.72 ± 0.10	-1.85 ± 0.09	-2.06 ± 0.09
$t_{\text{var}} [\text{s}]$	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4
$\chi^2/d.o.f.$	0.767	0.711	0.712	0.723	0.807
Q-value	0.810	0.873	0.872	0.860	0.757

Table 3.6: Fit parameters and statistics for the SED of PG 1218+304, obtained with agnpy and Sherpa by fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$ and $t_{\text{var}} = 1 \text{ d}$, and freezing various values of δ_D . Values in dark red are fixed.

ter the break as δ_D increases. p_1 seems to remain roughly constant, up to the value for $\delta_D = 40$ which has a huge error. Similar trends can be found when looking at the fit parameters obtained after fixing $\gamma_{\min} = 10^2$ (Table A.3 in Appendix A). The same table includes some outputs of the MMDC tool for cross-checking purposes¹⁴.

In any case, fit results are again close both in terms of compatibility between the values of several free parameters, and in terms of fit statistics, which are all similar for different choices of δ_D and γ_{\min} . The reduced χ^2 , in particular, is almost always $\gtrsim 0.7$, meaning that the models are even slightly overfitting experimental data, which indeed have large errors in the TeV region. After looking at the Q-values, **the “reference” fit from now on will be the one with fixed $\delta_D = 20$, the one with the closest value to 1, even if the difference in fit quality with the $\delta_D = 25$ case is almost negligible.** This again highlights parameter degeneracy in blazar modeling, and is clear from the similarity of the plots in Figure 3.17, where the observed SED and the model for both values of δ_D are shown. This choice is consistent with the outcomes of fits where δ_D is not frozen, but t_{var} and the extremes of the electron distribution are (plots in Figure A.4 in Appendix A), and is valid for either $\gamma_{\min} = 10^3$ or $\gamma_{\min} = 10^2$. This latter case also has good test statistics, and very similar best-fit parameters: the difference lies mostly in the lack of a break at infrared energies (at least in the range covered here) due to a smaller γ_{\min} . Parameter degeneracy is therefore still an issue, but, on the other hand, it must be kept in mind that **fine-tuning models excessively may be useless for a population study.** To better visualize the fit quality, Gammapy can now be used, as there are enough fixed parameters to obtain an acceptable outcome¹⁵. Gammapy allows to easily plot data by

¹⁴Kindly provided by Prof. Narek Sahakyan, as the MMDC web interface still does not allow to freeze parameters other than γ_{\min} .

¹⁵Both Gammapy fits presented here actually returned the message “Optimization terminated successfully, but uncertainties are unreliable”. The results were taken as good anyway, since all the reduced χ^2

3.4. MODELING




Figure 3.17: Observed SED (dark red dots) and best-fit models (teal lines) of PG 1218+304, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and δ_D . 3.17a: $\delta_D = 20$. 3.17b: $\delta_D = 25$.

instrument, include error regions (Figure 3.18) and generate the covariance matrix and the test statistic profiles for each free parameter (Figure 3.19). The best-fit model not




Figure 3.18: Fit of the PG 1218+304 SED computed with `agnpy/Gammapy`, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$. Left: Observed SED and best-fit model. Right: Best-fit parameters. Dark red values are frozen.

only follows the data closely, but its error region is very thin, even if compared with the small errors in BeppoSax and *Fermi*-LAT spectral points. More quantitatively, the `Gammapy` best-fit parameters are almost identical to the `Sherpa` ones, but with generally smaller errors; the `Sherpa` fit outcome will be anyway considered the final model for the

profiles are quasi-parabolic inside the 1σ confidence interval, apart for a slight asymmetry, and errors are as good as `Sherpa`'s, if not better.

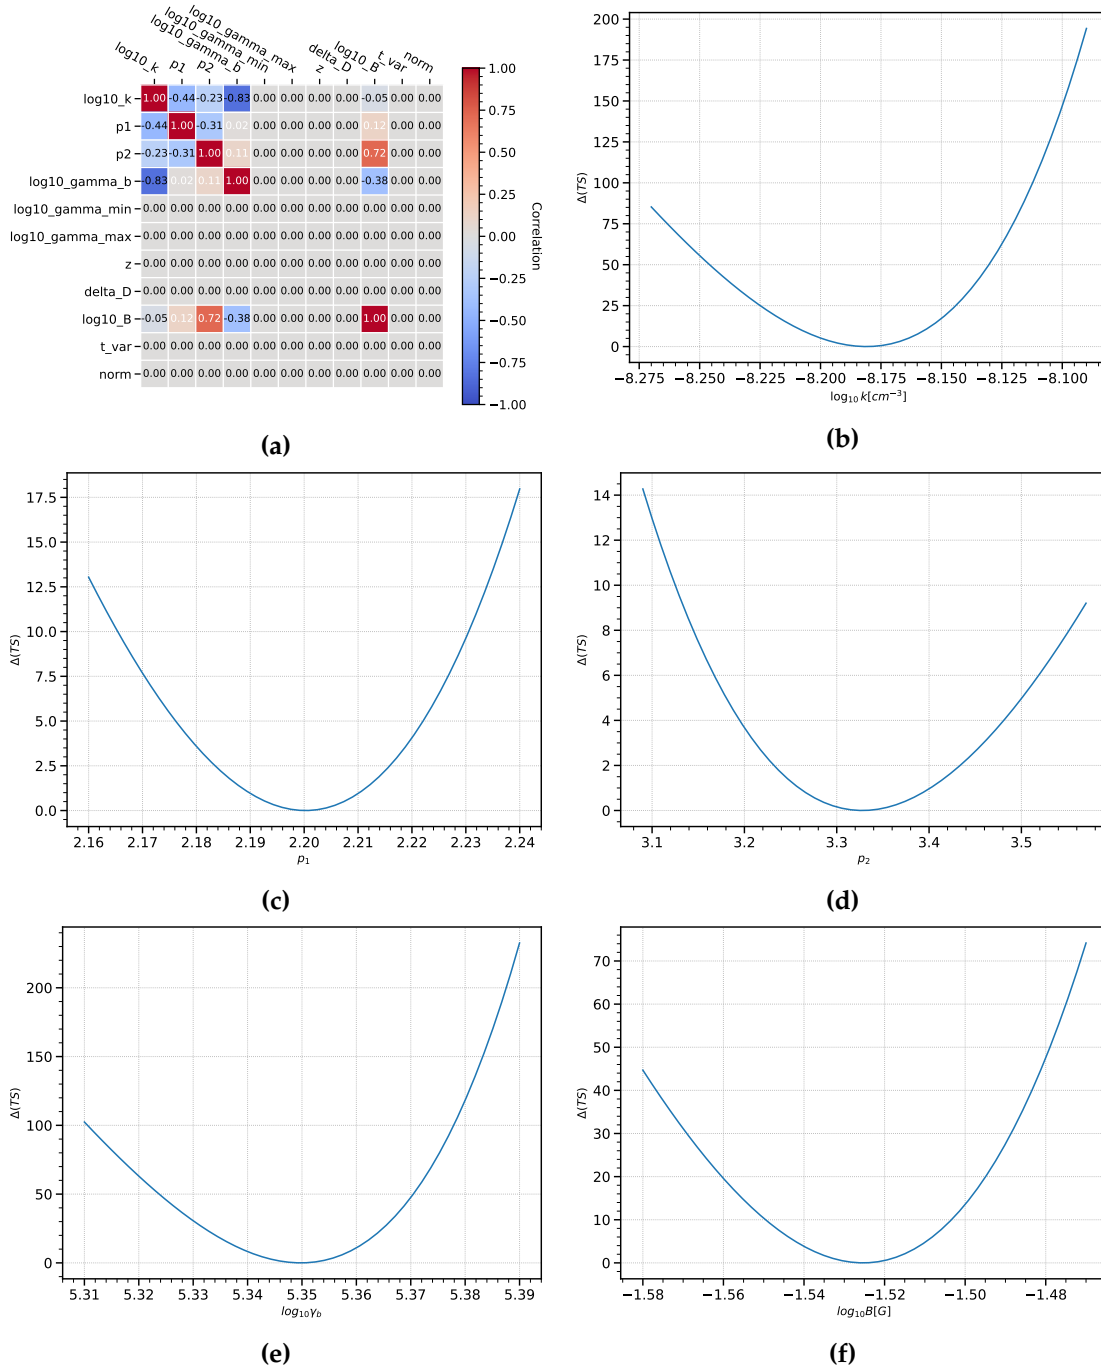


Figure 3.19: agnpy/Gammapy fit statistics for PG 1218+304, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$. 3.19a: Covariance matrix. 3.19b, 3.19c, 3.19d, 3.19e, 3.19f: Reduced χ^2 profiles in the 1σ confidence intervals for each free parameter.

3.4. MODELING

SED of PG 1218+304 with BeppoSAX data to allow for more direct comparisons with other models.

Second Data Selection As already discussed in Section 3.2, these results are actually the product of **one choice of X-ray data**. As an alternative, the means of the most frequent fluxes measured by *Swift*-XRT (from the Open Universe for Blazars catalog, Giommi et al. 2019) for each discrete value of the photon energy, with the square root of the variance of the sample as error, were chosen. The same parameters of the “best” models found in the above discussion were fixed: $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$. Since the Sherpa fit proved very unstable, Figure 3.20 shows the result from the Gammapy fit. Notice the widening of the error region compared to the pre-




Figure 3.20: Fit of the PG 1218+304 SED computed with *agnpy*/*Gammapy*, using the *Swift*-XRT mean instead of BeppoSAX data in the X-ray range, and fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$. Left: Observed SED and best-fit model. Right: Best-fit parameters. Dark red values are frozen.

vious data selection, and a worse (but still roughly acceptable) adherence to observed data. The main changes in parameters are a marked decrease of the magnetic field, and a decrease of the after-break slope p_2 , now not significantly higher than p_1 . The two data selections and the fits from both *agnpy*/*Gammapy* and the MMDC SED fitting tool are compared in Figure 3.21. MMDC results for this new data selection can be seen in Appendix A. “New” data seem to shift the lower-frequency peak to higher energies, making PG 1218+304 a fully extreme-synchrotron BL Lac. Approximate values of the peaks can be found with *agnpy* by building the theoretical emission region and SSC model from the given best-fit parameters, as explained in Section 3.3.1. Table 3.7 compares the values computed by *agnpy* for the two data selections and the same quantities as found in the 4LAC-DR3 catalog (Ajello et al. 2022), which includes estimates

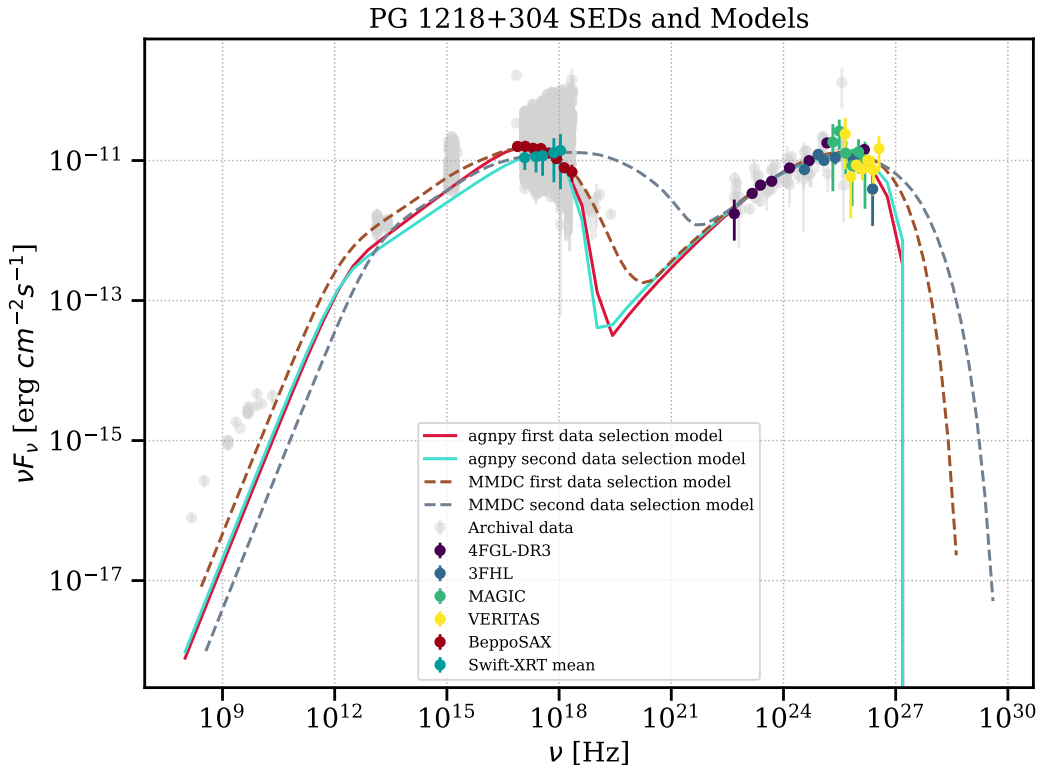


Figure 3.21: Comparison of selected data, archival data and SSC models from both agnpy/Gammapy and the MMDC SED fitting tool, for both data selections in the X-ray band. All models have fixed $\gamma_{\min} = 10^3$. agnpy/Gammapy models also have frozen $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$.

3.4. MODELING

for the inverse Compton peak frequencies and fluxes that are not in 4LAC-DR2. Fluxes

	agnpy, BeppoSAX	agnpy, Swift-XRT	4LAC-DR3
ν_{syn} [Hz]	9.54×10^{16}	2.44×10^{17}	1.86×10^{16}
νF_{ν}^{syn} [erg cm ⁻² s ⁻¹]	1.55×10^{-11}	1.23×10^{-11}	1.25×10^{-11}
ν_{SSC} [Hz]	3.56×10^{25}	3.56×10^{25}	2.33×10^{26}
νF_{ν}^{SSC} [erg cm ⁻² s ⁻¹]	1.20×10^{-11}	1.18×10^{-11}	1.65×10^{-11}
Compton Dominance	0.78	0.96	1.32

Table 3.7: Comparison between the synchrotron and SSC peak frequencies and fluxes of PG 1218+304 computed by agnpy for the two data selections, and as found in the 4LAC-DR3 catalog. The .sed_peak_nu() and .sed_peak_flux() functions of the Synchrotron and SynchrotronSelfCompton classes were used.

found with agnpy are similar to the *Fermi*-LAT estimates, especially νF_{ν}^{syn} found with the *Swift* data selection. There is more discrepancy in the estimates of the frequencies of both peaks, as the *Fermi*-LAT estimate for ν_{syn} is almost one order of magnitude smaller than the agnpy counterparts. The exact opposite happens for ν_{SSC} . This is probably due to differences both in data selection and the technique to evaluate the peaks (in LAC catalogs, local polynomial fits were performed).

BIN 3: LOW-END HBLs

The results from the agnpy fits were applied to the other two HBL bins, starting from “Bin 3” (10^{15} Hz $\leq \nu_{\text{syn}} < 10^{16}$ Hz). This subgroup of selected sources still includes 15 objects, a relatively large number, and **PKS 2155-304** was selected as representative source. This high synchrotron-peaked BL Lac at redshift $z = 0.1167$ (from NED), discovered by a radio survey in 1974, is more famous for its brightness in the X-ray band (Schwartz et al. 1979, Abdalla et al. 2020). The first VHE detection was in 1996, by the University of Durham Mark 6 telescope (Chadwick et al. 1999). H.E.S.S. was the first instrument to observe it at TeV energies during a series of campaigns in 2002-2003 (Aharonian et al. 2005) and to estimate a variability timescale on the order of days: for this reason, the $t_{\text{var}} = 1$ d assumed for PG 1218+304 was also adopted here. It must be noted, however, that this source has a long history of investigations about its strong variability on different timescales, from a few minutes to a potential gamma-ray quasi-periodicity on a \sim year scale (see Section 1.2.2; e.g. Peñil et al. 2023, Zhang et al. 2017). Rieger 2019 discusses the variability power spectrum of PKS 2155-304, confirming a characteristic timescale of ~ 1 d.

The selected SED data for this source were:

- **VHE gamma rays:** H.E.S.S. observation campaign from September 1st to 30th, 2003 (Aharonian et al. 2005)
- **HE gamma rays:** *Fermi*-LAT 4FGL-DR3 and 3FHL catalogs.

- **X rays:** single BeppoSAX observation from November 4th, 1999 (Giommi et al. 2002) or *Swift*-XRT data averaged from the Open Universe for Blazars catalog.

For both data selections the fit was initialized on the “reference” result for Bin 4, with fixed $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d and $\delta_D = 20$. Results are shown in Table 3.8 along with the MMDC best-fit parameters. Figure 3.22 shows agnpy/Sherpa plots;

	BeppoSAX		<i>Swift</i> -XRT	
	MMDC	agnpy	MMDC	agnpy
$\log L_e$ [erg s ⁻¹]	44.6 ± 0.3	-7.95 ± 0.05	44.5 ± 0.4	-7.65 ± 0.09
$\log k_e$ [cm ⁻³]	2.2 ± 0.2	2.40 ± 0.03	1.9 ± 0.2	2.43 ± 0.05
p_1		4.7 ± 0.1		3.8 ± 0.3
$\log \gamma_{\text{cut}}$	5.5 ± 0.2	5.10 ± 0.01	6.1 ± 0.6	5.06 ± 0.03
$\log \gamma_{\text{min}}$	3	3	3	3
$\log \gamma_{\text{max}}$		6		6
δ_D	23 ± 10	20	12 ± 11	20
$\log B$ [G]	-1.7 ± 0.4	-1.14 ± 0.03	-0.7 ± 0.6	-1.40 ± 0.07
$\log R$ [cm]	17.5 ± 0.5		16.5 ± 0.7	
t_{var} [s]	$(5 \pm 6) \times 10^5$	8.64×10^4	$(0.9 \pm 2) \times 10^5$	8.64×10^4
$\chi^2/d.o.f.$		0.767		0.767
Q-value		0.800		0.777

Table 3.8: Best-fit parameters of the PKS 2155-304 SED computed with MMDC and agnpy/Sherpa, using either BeppoSAX data or the *Swift*-XRT mean in the X-ray range, fixing $\gamma_{\min} = 10^3$ for both tools, and $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$ for agnpy. Dark red values are frozen. Teal values were computed later for comparison and are not part of the original fit parameters.

MMDC plots can be found in Appendix A.

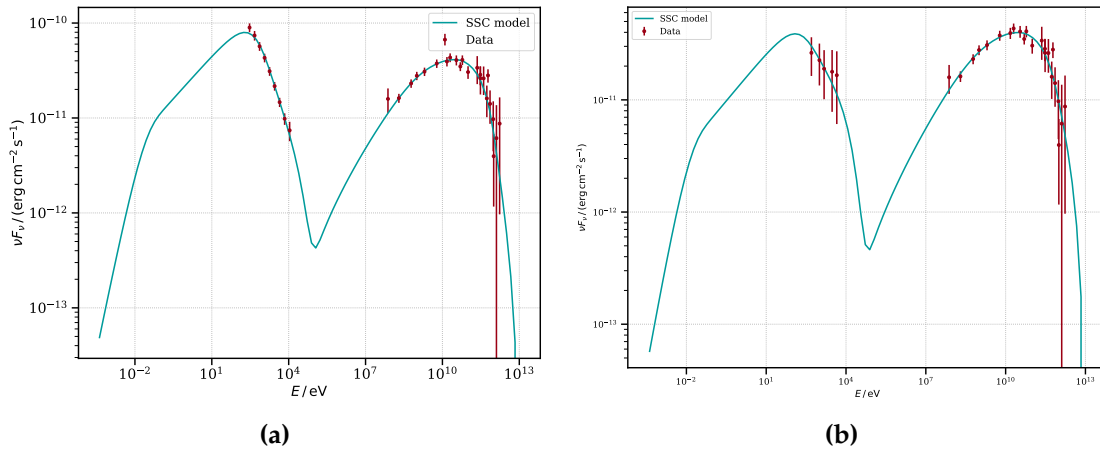


Figure 3.22: Observed SED and agnpy/Sherpa best-fit model of PKS 2155-304, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$, $t_{\text{var}} = 1$ d and $\delta_D = 20$. 3.22a: BeppoSAX X-ray data. 3.22b: mean *Swift*-XRT X-ray data.

Considering the agnpy/Sherpa results, both data choices produce good fit statistics, and an excellent fit of the inverse Compton bump, but it is clear from the plots that the curve is not able to reproduce the slope of the synchrotron curve in the X-ray band

for the *Swift*-XRT data selection. The BeppoSAX version is almost perfect, but the very small errors on that single observation strongly constrain p_2 to an excessively large value which, however, is not unheard of: see, for example, the “doubly-broken” power law model adopted for the emission region of Mrk 421 in Abdo et al. 2011a, where the large Δp is also an issue. The authors underline how such a large p_2 is not compatible with slow synchrotron cooling in a homogeneous model, which would produce a $\Delta p \sim 0.5$, and that constraining $p_2 = p_1 + 1$, the upper limit in this scenario, is unable to produce satisfying outcomes. Concerning PKS 2155-304, it has already been stressed here that the above results also depend from the choice to fix 4 model parameters to stabilize fits and hopefully perform meaningful comparisons with the Bin 4 representative source, which has been shown to be almost extreme. Bin 3 objects are expected to exhibit “milder” features, so more work needs to be done on the possibility to tweak γ_{\min} , whose elected value of 10^3 is very high, and of course t_{var} , as already explained.

BIN 5: EXTREME HBLs

The highest-frequency synchrotron peak bin in this work matches one definition of “extreme-synchrotron” BL Lacs: $\nu_{\text{syn}} \geq 10^{17}$ Hz. **PKS 0548-322** has been chosen to represent this peculiar subgroup of 13 sources. Like PKS 2155-304, it was one of the very first BL Lacs to be detected in the X-ray band (Schwartz et al. 1979) and it was first observed at very high energies by H.E.S.S. during a dedicated campaign between 2004 and 2006 (Superina et al. 2008). Its redshift (from NED) is $z = 0.069$.

Selected data for this work include:

- **VHE gamma rays:** H.E.S.S. discovery campaign from October 1st, 2004 to January 31st, 2008 (Aharonian et al. 2010).
- **HE gamma rays:** *Fermi*-LAT 4FGL-DR3 and 3FHL catalogs.
- **X rays:** *Swift*-BAT 105-month survey and single BeppoSAX observation from February 20th, 1999 (Giommi et al. 2002), or *Swift*-XRT data averaged from the Open Universe for Blazars catalog. Only one *Swift*-BAT flux on the full 14 keV to 195 keV range is recorded in the catalog, but MMDC, Firmamento and SSDC all report two data points (at 50 and 100 keV) for PKS 0548-322.
- **Optical - UV:** median of *Swift*-UVOT observations from May 2005 (Yershov 2014).
- **IR:** WISE data from March 11th to 12th, 2010 as retrievable from the MMDC and Firmamento online platforms (Wright et al. 2010). It is not completely clear if the data release used by MMDC is the 2012 All-Sky (Cutri et al. 2012) or the later AllWISE quoted by VOU-Blazars; in any case, fluxes are given as magnitudes in each band, but the MMDC table already lists them as energy flux densities.

The fit was again initialized on the “reference” result for Bin 4 for both X-ray selections, with fixed $\gamma_{\min} = 10^3$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$, but $\gamma_{\max} = 10^7$ to leave more room for the emission of VHE photons. The source has not been shown to be variable, at least on timescales of less than ~ 1 d (see e.g. Devanand et al. 2022), so t_{var} was kept at the

usual value. Results are shown in Table 3.9 along with the MMDC best-fit parameters. MMDC plots can be found in Figure A.6 in Appendix A. Although the reference object,

	BeppoSAX		Swift-XRT	
	MMDC	agnpy	MMDC	agnpy
$\log L_e$ [erg s ⁻¹]	44.2 ± 0.3	-10.6 ± 0.1	44.2 ± 0.3	-6.6 ± 0.8
$\log k_e$ [cm ⁻³]				
p_1	2.3 ± 0.1	2.32 ± 0.02	2.4 ± 0.1	2.0 ± 0.3
p_2		4.5 ± 0.3		2.49 ± 0.04
$\log \gamma_{\text{cut}}$	6.6 ± 0.3	5.89 ± 0.05	6.6 ± 0.4	4.2 ± 0.3
$\log \gamma_{\text{min}}$	3	3	3	3
$\log \gamma_{\text{max}}$		7		7
δ_D	40 ± 10	20	44 ± 10	20
$\log B$ [G]	-2.8 ± 0.4	-1.62 ± 0.03	-2.8 ± 0.5	-1.61 ± 0.03
t_{var} [s]	(2 ± 2) × 10 ⁵	8.64 × 10 ⁴	(2 ± 2) × 10 ⁵	8.64 × 10 ⁴
$\log R$ [cm]	17.4 ± 0.5		17.3 ± 0.5	
$\chi^2/d.o.f.$		1.013		1.085
Q-value		0.446		0.352

Table 3.9: Best-fit parameters of the PKS 0548-322 SED computed with MMDC and agnpy/Sherpa, using either BeppoSAX and Swift-BAT data or the Swift-XRT mean in the X-ray range, fixing $\gamma_{\text{min}} = 10^3$ for both tools, and $\gamma_{\text{max}} = 10^7$, $t_{\text{var}} = 1$ d, and $\delta_D = 20$ for agnpy. Dark red values are frozen. Teal values were computed at a later moment for comparison and are not part of the original fit parameters.

PG 1218+304, is a “borderline” extreme BL Lac, fits of PKS 0548-322 based on its model parameters do not yield good results. This is clear from the plots in Figure 3.23, where the fitting curve struggles in catching the slope of the observed SED in the infrared and high-energy X-ray bands (an effect amplified by the very small errors of chosen data points), and probably underestimates the inverse Compton peak frequency in a region where errors are large, there is no observational cue on where the peak could be, and the slope seems not to be constant. Mean Swift-XRT fluxes cannot be fitted well if Swift-




Figure 3.23: Observed SED and agnpy/Sherpa best-fit model of PKS 0548-322, fixing $\gamma_{\text{min}} = 10^3$, $\gamma_{\text{max}} = 10^7$, $t_{\text{var}} = 1$ d and $\delta_D = 20$. 3.23a: BeppoSAX and Swift-BAT X-ray data. 3.23b: mean Swift-XRT X-ray data.

3.4. MODELING

BAT data are also considered, but discarding the latter produces a cooling break at low energies (optical-UV band) and an incredible peak at $E_{\text{syn}} > 10$ keV, followed by a steep descent (caused not by cooling, but by γ_{max}), that does not follow the behavior of the actual observed SEDs, if not perhaps during flaring states. This result was chosen anyway over the attempts that included *Swift*-BAT points, as it was produced by the most stable fits with the “least bad” statistics. Notice, in particular, the excellent reduced χ^2 and the reasonable value of p_2 , perfectly in line with the theoretical slow cooling scenario usually invoked in blazar modeling. This is not true when *Swift*-BAT data are considered, and in general for all attempts done with the BeppoSAX selection. In that cases, the issue of $p_2 \sim 4.5$ arises again, like for the Bin 3 source PKS 2155-304.

It seems that for PKS 0548-322, too, adjusting the data selection and fixed parameters could be the starting point for better results, with due caution since extreme BL Lac SEDs are notoriously hard to fit with a one-zone SSC model without invoking, indeed, extreme values for some parameters, including some that were frozen like δ_D and γ_{min} (needed in the range 10^3 - 10^5 : see Biteau et al. 2020). Considering that WISE data points may suggest a change in slope somewhere in the optical band, and that the fitting curve goes systematically above the lowest-energy WISE datum, it may well be that a better fit at low frequencies is achieved by increasing γ_{min} (and the frequency of the associated break in the SED). As pointed out in several papers, like Biteau et al. 2020, alternative scenarios like multi-zone models, external Compton and other radiative processes, including lepto-hadronic ones, are often considered for extreme HBLs in order to “soften” parameters, so it is not granted that the usual one-zone SSC model will really work for PKS 0548-322 and its peers.

3.4.3 FINAL CONSIDERATIONS

A summary of the fit outcomes for all the three bins this work focused on can be seen in Table 3.10 for agnpy/Sherpa, and Table 3.11 for MMDC. The BeppoSAX data selection, which gave the best agnpy/Sherpa results in general and drove the path to fixing parameters, was chosen as an example, but there is actually **no robust way to determine which choice is the best**. As already explained in Section 3.2, BeppoSAX data are of excellent quality, spread over a wide frequency range, and seem to join well spectral points at other energies, but their choice is quite arbitrary, coming from single observations that are non-simultaneous with the TeV ones, and their errors are so small they do not encode the large flux spread in the X-ray band. *Swift*-XRT means could be biased in favor of high states, and they indeed shift synchrotron peaks to higher energies that may not represent the actual average emission; the quality of fits is generally worse, but adopting the sample variances as errors is more meaningful in the light of the strong X-ray variability. For this reason, not many conclusions can be drawn for now. Moreover, even with a robust data selection, it is clear that **attempting to fix the worst**

	Bin 3 (PKS 2155-304)	Bin 4 (PG 1218+304)	Bin 5 (PKS 0548-322)
$\log k_e [\text{cm}^{-3}]$	-7.95 ± 0.05	-8.2 ± 0.5	-10.6 ± 0.1
p_1	2.40 ± 0.03	2.20 ± 0.07	2.32 ± 0.02
p_2	4.7 ± 0.1	3.3 ± 0.3	4.5 ± 0.3
$\log \gamma_b$	5.10 ± 0.01	5.4 ± 0.2	5.89 ± 0.05
$\log \gamma_{\min}$	3	3	3
$\log \gamma_{\max}$	6	6	7
δ_D	20	20	20
$\log B [\text{G}]$	-1.14 ± 0.03	-1.53 ± 0.08	-1.62 ± 0.03
$t_{\text{var}} [\text{s}]$	8.64×10^4	8.64×10^4	8.64×10^4
$\chi^2/d.o.f.$	0.767	0.711	1.013
Q-value	0.800	0.873	0.446

Table 3.10: Best-fit parameters for the three representative sources computed with agnpy/Sherpa, using BeppoSAX (+ *Swift*-BAT) data in the X-ray range, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$ (Bins 3 and 4), $\gamma_{\max} = 10^7$ (Bin 5), $t_{\text{var}} = 1$ d, and $\delta_D = 20$. Dark red values are frozen.

	Bin 3 (PKS 2155-304)	Bin 4 (PG 1218+304)	Bin 5 (PKS 0548-322)
$\log L_e [\text{erg s}^{-1}]$	44.6 ± 0.3	44.6 ± 0.4	44.2 ± 0.3
p	2.2 ± 0.2	2.3 ± 0.2	2.3 ± 0.1
$\log \gamma_{\text{cut}}$	5.5 ± 0.2	5.9 ± 0.5	6.6 ± 0.3
$\log \gamma_{\min}$	3	3	3
δ_D	23 ± 10	31 ± 10	40 ± 10
$\log B [\text{G}]$	-1.7 ± 0.4	-2.0 ± 0.6	-2.8 ± 0.4
$\log R [\text{cm}]$	17.5 ± 0.5	17.0 ± 0.6	17.4 ± 0.5
$t_{\text{var}} [\text{s}]$	$(5 \pm 6) \times 10^5$	$(2 \pm 2) \times 10^5$	$(2 \pm 2) \times 10^5$

Table 3.11: Best-fit parameters for the three representative sources computed with the MMDC SED fitting tool, using BeppoSAX (+ *Swift*-BAT) data in the X-ray range, fixing $\gamma_{\min} = 10^3$, $\gamma_{\max} = 10^6$ (Bins 3 and 4), $\gamma_{\max} = 10^7$ (Bin 5), $t_{\text{var}} = 1$ d, and $\delta_D = 20$. Dark red values are frozen. Teal values were computed later for comparison and are not part of the original fit parameters.

defined parameters for all bins is probably too ambitious at this point: on one hand, most agnpy/Sherpa fits of the PG 1218+304 spectrum (Bin 4) are very good for most of the parameter combinations that were obtained, highlighting the issue of degeneracy; on the other, fit quality declines for the other sources when fixing the same quantities frozen for Bin 4, especially in Bin 5. It is therefore reasonable to think that thawing these parameters, maybe still constraining them to a small interval of plausible values, could produce better fits, but probably at the expense of stability when computing errors, and of opportunities to pinpoint trends across the sequence (as there would be too many changing correlated parameters).

Indeed, those presumed trends are not so easily discernible from the agnpy/Sherpa fits. MMDC ones, on the other hand, seem to be clearer. Outcomes from both tools show that:

- the main electron distribution **power index** p_1 **fluctuates around** $\gtrsim 2.2$, in the expected range for first-order Fermi acceleration in shocks (see e.g. Spurio 2018 and references therein);
- the magnetic field, B **decreases as** v_{syn} **increases**;
- the **radius** of the emission region, R is more or less constant: in the case of agnpy this is trivial, since δ_D and t_{var} were fixed by choice, such that R will vary only because of different redshifts; in the case of MMDC, instead, R is a free fit parameter, but it is still **almost constant**.

According to agnpy/Sherpa results,

- the normalization of the electron distribution, k_e , **decreases strongly with increasing** v_{syn} ;
- p_2 assumes a “reasonable” (under the a homogeneous synchrotron slow cooling model) value for Bin 4 only, while it goes above 4 for both Bin 3 and Bin 5;
- the electron Lorentz factor at the spectral break, γ_b **increases**.

MMDC fits reveal that:

- the cutoff Lorentz factor of the electron energy distribution, γ_{cut} **slowly increases**, and is of the same order of magnitude of typical γ_{max} and γ_b when considering broken power law models;
- δ_D **increases**, and is often large;
- both R and t_{var} are predicted to be quite large;
- the electron luminosity, L_e **is almost perfectly constant**.

This latter consideration, along with the fact that the electron power index seems to be roughly constant, may be a good starting point for advancing this work, provided that either agnpy developers include the possibility to set L_e in place of k_e as a fit parameter, or that MMDC’s make fixing parameters a feature available to end users. On the other hand, it is probably the fact that almost all parameters are left free that makes the fitting

curves from MMDC follow the experimental data closely (with the slight exception of Bin 5 in the infrared and high-energy X-ray bands); errors are also comparable with the ones computed by both Sherpa and Gammapy, except for the δ_D (and, consequently, the t_{var}). Apart from this latter issue and the apparent systematic tendency to predict high values of δ_D , this level of performance, the ease of use, and the speed make this new machine learning-based tool a resource with **good potential** for future studies of the blazar sequence.

4

Estimation of Spectral Quantities in Blazars with Machine Learning

So far in this thesis, machine learning has been exploited from the point of view of the end user of a ready-made tool: other than uploading a data file with the right structure and format, and providing the source redshift and an e-mail address through a user-friendly web interface, nothing needs to be done to receive all the requested results in about half a hour. In principle, the MMDC SED modeling tool does not even require to know what machine learning is, which model is it based on, how was such model trained, and which steps of the fitting procedure employ it. This chapter instead describes an attempt at applying machine learning techniques at a more fundamental level, by writing a code tailored on the user's specific purpose.

4.1 PURPOSE

The attempt at reforming the blazar sequence, as already explained, reversed the usual approach by classifying SEDs by synchrotron peak frequency bins, instead of intrinsic luminosity in a single energy band. For this purpose, the synchrotron peak frequencies from the 4LAC-DR2 catalog were chosen. This, however, implied excluding 226 sources, whose ν_{syn} is not reported in the catalog; two of them are TeV-detected objects, whose ν_{syn} were corrected with the values found with other means (the aforementioned BlaST tool, see later discussion) in order to consider them for the new blazar sequence¹. Computing these missing data would allow to expand the sequence analysis

¹Curiously, they both ended up right in Bin 4, the starting point for spectral modeling.

4.1. PURPOSE

to all BL Lacs, including the ones that were not detected at TeV energies, while adopting the same dataset binning principle.

Considering how ν_{syn} is computed in the 4LAC catalogs, one may suppose that a failed estimate could result from scarcity of data at energies outside the *Fermi*-LAT range, such that the broadband SED has not enough points in the synchrotron peak region, or not at all, to be fitted. On the other hand, while discussing data selection, it became clear that the most suited spectral points for reconstructing average SEDs are the ones from the *Fermi* catalogs, due to them being the result of multi-year integration and therefore a good representation of the average emission state of each source. *Fermi*-LAT spectral data well sample the high-energy peak, and their power-law index has been shown to be correlated to the inverse Compton peak frequency. An one-zone synchrotron/inverse Compton scenario, especially SSC, justifies a correlation between the two peaks, meaning that the gamma-ray slope is correlated to ν_{syn} (Ajello et al. 2020 and references therein). From this came the ambitious idea of **training machine learning models on *Fermi*-LAT data only, to make predicting BL Lac synchrotron peak frequencies as independent from multiwavelength observations as possible**. Moreover, as pointed out in Glauch, Kerscher, and Giommi 2022, creating a good machine learning tool to predict spectral quantities like ν_{syn} would make this task faster and more automatic, reducing the impact of human error.

Basics of Learning - Part 2: Some Useful Definitions

As already explained in the box in Section 3.3.2, machine learning works by taking a dataset as an input and performing an optimization on one or more functions or “learning structures” to make predictions of some kind. More specifically, in a *supervised* model the algorithm will receive a **training dataset** as an input: each entry (**instance**) will be a row of the values taken by a number of variables, of which some are considered *independent* and are called **features**; some others instead are the *dependent* variables that the model will learn to predict, called **targets**. For example, an instance in the training dataset for the models described in this chapter would be a row from the merged 4LAC+4FGL catalog (see Subsection 4.3.2), referring to one specific source, with all its related quantities like fluxes, variability, etc. The model will learn to map the values of the features to the matching targets, according to some unknown relation that the algorithm will “understand” from the data themselves. When presented with a table of new feature values but unknown targets, the model should be able to recognize the same patterns in them and predict the new values of the targets. When the requested output is numerical, as in the cases discussed in this thesis, the task is called **regression**.

Building a machine learning model of this type requires three basic steps: *training*, *validation*, and *testing*. The dataset is split in three parts: the model will learn from the training subset, be evaluated on the validation subset, and then, once the training is considered successful and the best model is selected, there will be a final assessment on the test set. This stages is needed to avoid the conceptual error of training a model on all the input data, and then evaluating the quality of learning on the input data themselves. The central element of the latter two steps is predicting the target values from the features of known data, and comparing them to the actual values to compute one or more metrics (*scores*) that quantify the quality of training. This approach is sometimes risky, as there could be too few input data to create all three subsets, and there is the danger of wrongly using the test set after each validation. The latter mistake would optimize the model on the test set, leading to *overfitting* (see e.g. Hastie, Tibshirani, and Friedman 2009). Overfitting happens when the algorithm starts modeling not only the general patterns, but also the noise: in other words, it “learns too much”, becoming unable to generalize on new data. There is a particularly convenient way to handle the issue: *k-fold cross-validation*, where the input dataset no longer needs to be split in three subsets. It is instead split in k subgroups (*folds*): one is kept out and used for validation, while the model is trained on the remaining $k-1$; a validation score is computed and stored; the procedure is then iterated by choosing another fold as validation and test data, while the first one now joins the training subset, and so on until each subset has been used to evaluate the training once (Figure 4.1).

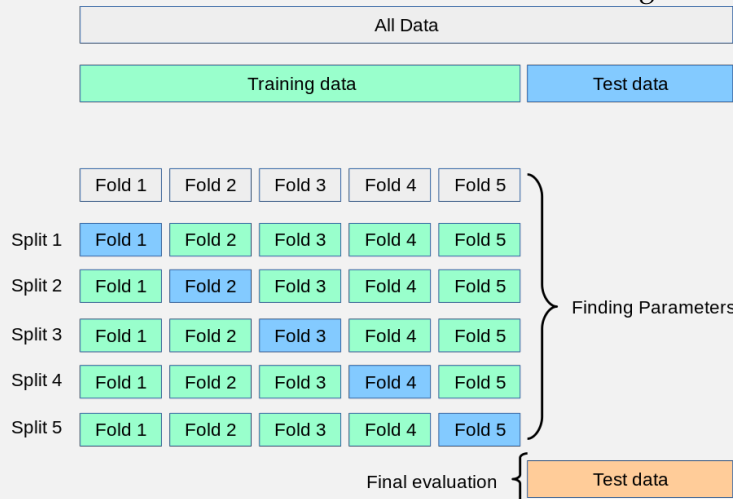


Figure 4.1: k -fold cross-validation scheme, from Scikit-learn Developers 2024. The final performance score will be the average found at each iteration. For validation, a number of metrics can be used: for example, the mean squared error, the mean absolute error, and the *coefficient of determination*, R^2 (Scikit-

4.1. PURPOSE

learn Developers 2024):

$$R^2(y, \hat{y}) = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

where $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$, n is the number of samples, y_i is the true value of the i -th sample and \hat{y}_i is the corresponding predicted value. From this definition, R^2 can be interpreted as an estimate of how much of the predicted variable variance can be explained by the independent variables in the fit; or, in other words, how likely it is that new samples will be predicted well by the model.

Another trick to prevent overfitting, and improve efficiency, is *early stopping*: a portion of each fold is indeed set aside for internal validation after each iteration, and the training is interrupted if the performance on the validation subset has become stable or began worsening in a certain user-defined number of iterations. Models can be refined by choosing some settings, called *hyperparameters*, that regulate some properties of the procedures and structures used by the algorithm, like the number of iterations to perform. In machine learning packages there are objects, like `GridSearchCV`, that scan lists of candidate values of some hyperparameters, given by the user, train and evaluate the model with every possible combination, and use cross-validation to choose the settings that give the best validation metrics. The model with the best hyperparameter combination is then implemented (Figure 4.2).

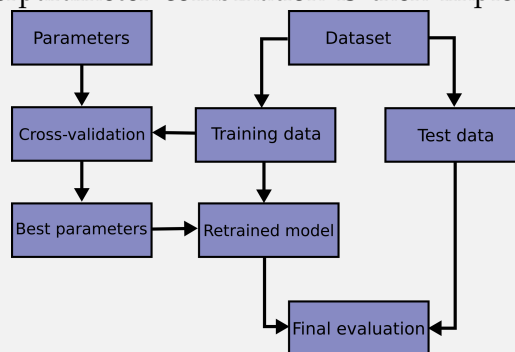


Figure 4.2: Flowchart of cross-validation workflow in model training, from the *Scikit-learn* documentation (Scikit-learn Developers 2024).

This is not the only way to improve model performance: for some models, *pre-processing* data is fundamental. In particular, there are algorithms that work best with data distributed as close as possible to a standardized Gaussian. For this reason, special methods called *transformers* can be used: they are described in Section 4.3.3. To check whether distributions are suitable for optimal model training, *box plots* and *probability plots* can be employed. A box plot summarizes graphically the distribution of a variable: one axis represents the values the

variable may take; for each quantity, a rectangle is drawn extending from the the Q1 to the Q3 quartile of the distribution (such that the length of the side of the box along the axis is the interquartile range). The median is highlighted inside the rectangle, in order to see immediately how skewed is the distribution, depending on its distance from the ends of the box. Two segments called “whiskers” extend from each end of the rectangle for $1.5\times$ the interquartile range: any data point that falls outside is drawn individually and may be considered an outlier. Normal probability plots describe how much a distribution deviates from a Gaussian, by plotting the sorted experimental values of a variable against the matching quantiles of the normal distribution: consequently, the closer are plotted data to the bisector, the closer their distribution is to a Gaussian.

OUTLINE

To implement what has just been described, the following steps were taken:

- Choose which **algorithm**, or class of algorithms, to base the model on.
- Choose which coding **library** to employ.
- Select the **training dataset**: define the **target** variable and the training **features**, and eventually a dataset to make predictions on after training.
- Prepare data for optimal training:
 - Discuss the **distributions** of each feature and of the target in the training dataset.
 - Eliminate possible **outliers** that may disrupt the learning process.
 - If needed, **transform** data distributions to make them closer to a standardized Gaussian, which is handled more effectively by learning algorithms.
 - Discuss possible relations between the features and the target through *bivariate analysis*, by plotting the values of the target variable against each feature in the training dataset.
- Train the model on the known values of the features and target:
 - Select the **performance metrics** to evaluate models.
 - Select the best number of *k-folds* and the best **hyperparameters** with cross-validation methods.
 - Use **cross-validation** to train and evaluate the selected best models on the training dataset.
 - Compare models by back-transforming the results (predictions on training data, performance scores) to the original scale.
- Apply the model to make **predictions** on a previously unseen dataset.
- Compare the result on new data with the output of other tools.

4.2 ALGORITHMS

The choice of which machine learning models to adopt was dictated mainly by practical needs. The basic concept of **decision tree-based ensemble models** and their inner mechanisms, by which they train and predict on data, is relatively easy to understand for users, even beginners: algorithms of this type have a good level of *interpretability* despite being built precisely to capture non-linear, non-trivial interactions between input quantities (Molnar 2022), which indeed is the main purpose of this work. Models may be limited by several factors, like missing entries, outliers in the data distributions, and irrelevant inputs, like quantities the target variable has no relation to. Tree-based algorithms handle well these issues, albeit at the expense of predictive power compared to other popular, but much less interpretable, models like neural networks (see e.g. Hastie, Tibshirani, and Friedman 2009, Table 10.1). They are also already widely used in the very-high-energy astrophysics field for low- and intermediate-level data analysis tasks like event classification and energy reconstruction (e.g. Albert 2008).

4.2.1 RANDOM FOREST

Random Forests are perhaps the most widespread machine learning algorithms for the astrophysical data analysis needs mentioned above, and are an immediate evolution of the basic *decision tree* model. The latter works in a straightforward way: when doing a regression task, learning structures called “**trees**” are **grown by splitting data multiple times** according to some threshold value of one feature per split, such that each input element belongs to one subgroup (*node*). This goes on until terminal nodes (*leaves*) are reached. **A simple model is then fit in each of the partitions** of the feature space: for example, the target for each terminal node may be modeled to a different constant value, like the average target of the training data in that leaf (Molnar 2022; Hastie, Tibshirani, and Friedman 2009).

Random forests are ensembles of **parallel** decision trees: a number of random subsets is created, and a decision tree is built for each of these samples. **The predictions from these decision trees are finally averaged**, reducing the overall variance. In this way, they improve the performance of single decision trees by reducing their weaknesses (the high variance and the tendency to overfit) while retaining their strengths (the low bias).

4.2.2 GRADIENT BOOSTED DECISION TREE

Gradient boosting regression is another way to increase the performance of a “weak learner”, specifically a decision tree, by building an ensemble of **consecutive** trees. In this case, we have *forward stagewise additive modeling*: given an N -dimensional input x ,

the matching prediction \hat{y} will be (see Scikit-learn Developers 2024 and Hastie, Tibshirani, and Friedman 2009):

$$\hat{y} = F_M(x) = \sum_{m=1}^M h_m(x) \quad (4.1)$$

where $h_m(x)$ is the m -th tree, and $m = 1, \dots, M$. At each step,

$$F_m(x) = F_{m-1}(x) + h_m(x) \quad (4.2)$$

without modifying previous terms. The model is fit, like in neural networks, by minimizing a *loss function* that describes how distant are predictions from actual values of the target variables in the training dataset:

$$h_m = \arg \min_h \sum_{i=1}^N l(y_i, F_{m-1}(x_i) + h(x_i)) \quad (4.3)$$

where x_i is the i -th training input, $F_{m-1}(x_i)$ is the prediction of the target value done by the $(m-1)$ -th tree on x_i , y_i is the respective “true” value of the target variable from the training dataset, and $h(x_i)$ is a tree computed for x_i . The basic, default choice for the loss function (which will be used here) is squared errors:

$$l(y, F(x)) = (y - F(x))^2 \quad (4.4)$$

The addends in the last term of Equation 4.3 are expanded to first order of a Taylor series,

$$h_m \approx \arg \min_h \sum_{i=1}^N h(x_i) \left[\frac{\partial l(y_i, F(x_i))}{\partial F(x_i)} \right]_{F=F_{m-1}} \equiv \arg \min_h \sum_{i=1}^N h(x_i) g_i \quad (4.5)$$

with the practical result that the m -th tree, h_m will be fitted to predict the *negative* gradients of the samples, $-g_i$ (*pseudoresiduals*), which are updated at each iteration and are proportional to residuals in the case of squared errors loss. In other words, **the model progressively “corrects” the errors with respect to the previous iterations**. Figure 4.3 shows graphically the difference between gradient boosting and random forest.

4.2.3 HISTOGRAM-BASED GRADIENT BOOSTED DECISION TREE

Histogram-based gradient boosting is an evolution of gradient boosting optimized for very large datasets ($>10^4$ instances) and inspired by the popular LightGBM framework (Ke et al. 2017). Other tree-based ensembles, like the single trees, need to scan all instances for each feature to find the optimal split points: LightGBM tackles the problem of the high computational cost by implementing a number of techniques to smartly reduce the size of the training dataset. Histogram-based gradient boosting regressors

4.3. IMPLEMENTATION

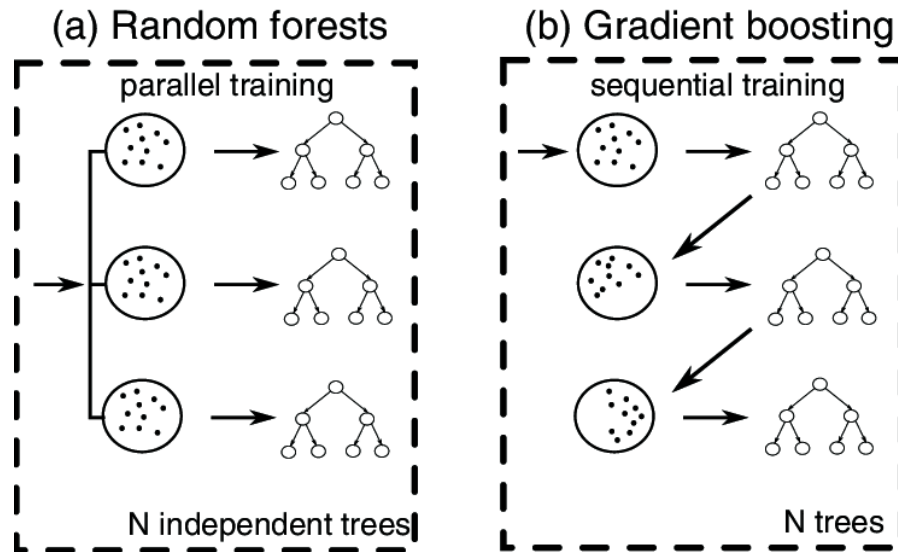


Figure 4.3: Graphical comparison between Random Forests and Gradient Boosting, from Kowalek, Loch-Olszewska, and Szwabiński 2019.

like the one used in this work **bunch the input samples into bins**. In this way, trees are built by handling a smaller number of discrete quantities to scan, and much less splitting points are needed. Sorting values for each feature is done only once, at the beginning, to create the histograms, while usual gradient boosting requires sorting each feature at each split: **histogram-based boosting can therefore be orders of magnitude faster for bigger datasets**. The specific estimator used here also includes native support for missing values, so it is particularly suited for sparse tables of training data.

4.3 IMPLEMENTATION

4.3.1 TOOLS

The Python module *Scikit-learn* (Pedregosa et al. 2011) was chosen to implement the models described above. *Scikit-learn* has several practical advantages that were crucial in its selection for this work:

- it is distributed under a new BSD license, and its development is mainly community-driven, making it popular in the world of free software;
- it was specifically created to be tightly integrated with Python, a widespread high-level programming language, in order to make machine learning more easily understandable and usable even in fields outside data science;
- it is built to maximize computational efficiency as much as possible while keeping simplicity as the main objective;
- it is often the package of choice in online tutorials on the basics of machine learning, and an extensive documentation, both official and not, is available;

- while focusing more on solidity, it still provides a large number of machine learning models of various types, including the three that were chosen for this work, and all the required tools for model tuning and training data engineering, all ready to use.

The design revolves around objects called estimators, that accept data and matching label arrays as arguments and can provide methods such as: `fit` for model training; `predict`, to be applied either on a validation dataset, or on new data after model training; and `score`, to evaluate the goodness of fit. The three selected machine learning models are implemented in the `RandomForestRegressor`, `GradientBoostingRegressor` and `HistGradientBoostingRegressor` classes. A particular type of estimators, called transformers, can modify input data to improve the quality of training. Another type of object, the *cross-validation iterator*, can split input data into training, test and validation sets. `GridSearchCV` (see box) allows to apply **cross-validation** to evaluate models.

Scikit-learn 1.4.2 was conveniently found to be already available on Bigmama, the CloudVeneto machine dedicated to statistical learning:

CPU AMD Ryzen Threadripper 3990X 64-Core Processor

Total Memory (RAM) 258 GB

Operating System AlmaLinux 9.4 (Seafoam Ocelot)

The code developed in this thesis was written in Python 3.11, on a Jupyter 1.0.0 notebook. Imported libraries, other than *Scikit-learn*, include NumPy 1.26.4, pandas 2.2.2, Matplotlib 3.8.4, and seaborn 0.13.2.

4.3.2 DATA SELECTION

The **4LAC-DR3** catalog (Ajello et al. 2022; data: Ajello et al. 2023) was chosen for this work, in place of the previous version. This allows to deal with more recent data, as 4LAC-DR3 was built from 4FGL-DR3 using 12 years of observations, instead of the 10 years covered by 4LAC-DR2. It was retrieved from the VizieR online database² and integrated with the more recent **4FGL-DR4** catalog (Abdollahi et al. 2022; Ballet et al. 2024) using the CDS X-Match service (Boch, Pineau, and Derriere 2012; Pineau et al. 2020). This operation produced a single dataset featuring the 3805 common sources, of 3814 total in 4LAC-DR3, described by the data columns of both catalogs.

As already explained in Section 3.2, LAC catalogs include all FGL sources that can be reliably associated to AGN-like counterparts in other energy bands of their SEDs: not only blazars, but also steep-spectrum radio quasars, compact steep spectrum radio sources, Seyfert galaxies and other objects of unknown nature (including BCUs, blazars

²<https://vizier.cds.unistra.fr/>

4.3. IMPLEMENTATION

of uncertain classification). This work focuses on BL Lacs, so the data table was filtered to select only the sources flagged as either BLL or b11 in the Class column of the 4LAC-DR3 catalog: 1458 objects, **276** of which lack a ν_{syn} estimate (column nu reads 0), leaving **1182** items (*instances*) in the training dataset.

Since some quantities are featured in both catalogs, several columns are, in fact, duplicates of each other, up to slight differences in their names and contents, due to different observation timespans: 12 years (2008-2020) for 4LAC-DR3, 14 years (2008-2022) for 4FGL-DR4. For this reason, when selecting a duplicate quantity as a training feature, the column from 4FGL was preferred, as it is the most up-to-date.

A few simple considerations guided the choice of training features:

- Only columns with purely **numerical** entries were taken into account; some of them, like the ones regarding coordinates and their uncertainties, are thought not to be correlated to the target feature ν_{syn} , and were discarded.
- A few columns have a high incidence of **missing values**, and were set aside to simplify the task. Due to the definition of “BL Lac” itself, the issue of missing values understandably plagues the redshift z column, whose cells are empty for 50.9% of selected sources. Unfortunately, the same scarcity holds for one of the most interesting new additions to the 4LAC catalogs: the energy (EpkHE) and flux (nuFnuPk) of the higher-energy SED peak, which would allow the model to assess the possible role of the Compton dominance, and of a presumed correlation between the two SED peaks in a SSC scenario, in determining ν_{syn} . Other quantities lacking a large fraction of values are the uncertainties on photon fluxes; they were also ignored.
- The **measured photon flux** in each band was preferred over the matching flux density, in order to avoid relying on values that depend on the phenomenological fit of the overall spectrum in the *Fermi*-LAT energy range (see Section 3.5 in Abdollahi et al. 2020). The total energy flux, Energy_Flux100, one of the shared columns between the two catalogs, was included anyway.
- A choice about the spectral model, however, was still necessary to include the most important feature, the **photon index**: the power law model was considered for all entries, so parameters referring to other models were discarded.

The selected quantities are described in Table 4.1. All features are from 4FGL-DR4; the target, ν_{syn} (nu), is from 4LAC-DR3. Figure 4.4 shows the Pearson correlation coefficient matrix³ of the selected quantities, followed by the chosen training features ranked by absolute value of the **correlation coefficient** with ν_{syn} (nu) in Figure 4.5. Values were computed using `pandas.DataFrame.corr`. Notice the small values: this, however, does not mean that the chosen features will not be adequate to reach the stated aim of predicting the target. Coefficients are *pairwise*, assessing the one-to-one correlation between ν_{syn} and *one* other quantity; the complex physics of BL Lac objects makes it reasonable

³Defined for two variables x, y as $\rho = \frac{\text{Cov}(x,y)}{\sigma_x \sigma_y}$, where σ_x, σ_y are the variances, and $\text{Cov}(x, y)$ is the correlation of the two variables. In turn, $\text{Cov}(x, y) = E[(x - E(x))(y - E(y))]$ and E is the expected value.

Feature	Definition	Units	Notes
Pivot_Energy	Energy at which the error on the differential flux is minimal	MeV	
Flux1000	Integral photon flux, 1 – 100 GeV	$\text{cm}^{-2} \text{s}^{-1}$	
Energy_Flux100	Energy flux, 100 MeV to 100 GeV by spectral fitting	$\text{erg cm}^{-2} \text{s}^{-1}$	
PL_Index	Photon index, Power law fit	/	
F50	Integral photon flux from 50 MeV to 100 MeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F100	Integral photon flux from 100 MeV to 300 MeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F300	Integral photon flux from 300 MeV to 1 GeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F1000	Integral photon flux from 1 GeV to 3 GeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F3000	Integral photon flux from 3 GeV to 10 GeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F10000	Integral photon flux from 10 GeV to 30 GeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F30000	Integral photon flux from 30 GeV to 100 GeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
F100000	Integral photon flux from 100 GeV to 1 TeV	$\text{cm}^{-2} \text{s}^{-1}$	Flux_Band
Variability_Index	Variability index, 1-yr time intervals	/	
Frac_Variability	Fractional variability, 1-yr time intervals	/	
nu	Synchrotron peak frequency, <i>observer</i> frame	Hz	

Table 4.1: First choice of training features and target. As highlighted in the Notes, fluxes in individual energy bands have been extracted from the Flux_Band column in 4FGL-DR4, which is actually a vector column with 8 sub-columns (see Sect. 3.4 in Abdollahi et al. 2022).




Figure 4.4: Pearson covariance matrix of the selected features and target.

4.3. IMPLEMENTATION




Figure 4.5: Bar chart of the Pearson correlation coefficients (absolute value) of the selected features with the target, ν_{syn} (ν).

to suppose that the target variable will be determined by a non-trivial combination of more features. Three parameters seem to correlate with ν_{syn} more than the others: two of them, the photon index (PL_Index) and the pivot energy (Pivot_Energy), appear in the definition of the power-law phenomenological model for the gamma-ray spectrum,

$$\frac{dN}{dE} = K \left(\frac{E}{E_0} \right)^{-\Gamma} \quad (4.6)$$

where Γ is the photon index and E_0 is the pivot energy. The apparent link between Γ and ν_{syn} is known, as already explained; the pivot energy is used in the above expression as a sort of “reference energy” of the source gamma-ray spectrum, as it is by definition the energy at which the statistical error on the flux normalization K is minimized (Abdollahi et al. 2020): it can be sensibly supposed that sources with spectra shifted to higher frequencies, as long as they are not too extreme (see Section 1.2.4), will show a higher, better-measured flux at higher energies; the third quantity is the fractional variability (Frac_Variability; see box in Section 1.2.2), and its higher Pearson coefficient is also reasonable as an anti-correlation between variability and ν_{syn} has already been discussed (Ajello et al. 2020).

4.3.3 DATA ENGINEERING

A fundamental step in good machine learning is **pre-processing** input data. There are several circumstances in which the learning performance can be improved by engineering the training dataset in various ways. As already mentioned above, tree-based

models are actually well-built for dealing with “messy” data, especially when there are outliers in their distributions, or unimportant variables are included: however, cleaning and standardizing the training dataset could still help training performance, having a clearer picture, and creating inputs ready to be eventually applied to other models that work less “off-the-shelf”, if needed.

DATA DISTRIBUTIONS

The first step is, obviously, **inspecting the distributions** of each training feature and the target. Almost all of the distributions are very skewed towards smaller values: an example can be seen in Figure 4.6 for the target variable ν_{syn} , the synchrotron peak frequency. The presence of an extreme outlier makes the other entries collapse to one single bin, while the probability plot is almost unreadable. This behavior can




Figure 4.6: Frequency histogram and probability plot of the target variable, ν_{syn} (ν) in linear scale. An outlier makes both plots unreadable.

be evaluated by the *skewness* and the *kurtosis* of the distributions, quantifying the relative weights and the “extremity” of the tails. Their values were computed according to the Fisher-Pearson and Fisher definitions respectively, using the `scipy.stats.skew` and `scipy.stats.kurtosis` functions, and are shown in Table 4.2 together with the values for scaled or transformed data (see later discussion). As expected, apart for `PL_Index`, values are always large and positive, generally indicating “heavy” tails (slower decay of the distribution tails compared to the Gaussian case, a large number of outliers) and a strong asymmetry in favor of the lower tail.

A first step in dealing with data like these was to **check outliers**. Considering it is not a particularly pressing issue with the chosen learning algorithms, the only clear outlier to be eliminated was 4FGL J1111.5+3455, because of its $\nu_{\text{syn}} = 1.798\,870 \times 10^{20}$ Hz. Another machine learning tool, BlaST (Glauch, Kerscher, and Giommi 2022, see later discussion) estimates $\nu_{\text{syn}} \simeq 2 \times 10^{15}$ Hz for this source; both 4FGL-DR4 and 4LAC-

4.3. IMPLEMENTATION

Feature	Linear		Log		Standard		Robust		Power		Quantile	
	Skew.	Kurt.	Skew.	Kurt.	Skew.	Kurt.	Skew.	Kurt.	Skew.	Kurt.	Skew.	Kurt.
Pivot_Energy	2.8	11.1	0.4	0.1	0.4	0.1	0.4	0.1	0.003	-0.1	0.003	0.8
Flux1000	9.5	126.9	0.7	0.6	0.7	0.6	0.7	0.6	0.02	-0.2	0.00009	0.8
Energy_Flux100	9.9	135.1	0.8	0.8	0.8	0.8	0.8	0.8	0.04	-0.3	0.06	1.1
PL_Index	0.3	0.2	-0.01	0.05	-0.01	0.05	-0.01	0.05	0.0006	0.05	-0.002	0.8
Variability_Index	22.1	588.7	1.6	3.2	1.6	3.2	1.6	3.2	-0.003	0.3	-0.0006	0.8
Frac_Variability	1.3	4.9										
F50	9.7	159.0	-1.1	1.2	-1.1	1.2	-1.1	1.2	-0.2	-1.0	-0.00005	0.8
F100	11.4	200.5	-1.9	3.4	-1.9	3.4	-1.9	3.4	-0.04	0.8	-0.0008	0.8
F300	11.4	193.3	-3.4	17.0	-3.4	17.0	-3.4	17.0	0.3	2.9	-0.0003	0.8
F1000	9.9	142.6	-1.0	13.6	-1.0	13.6	-1.0	13.6	0.04	4.4	-0.0005	0.8
F3000	9.2	116.5	-0.1	4.1	-0.1	4.1	-0.1	4.1	-0.05	3.9	0.02	0.9
F10000	10.5	154.4	-4.1	31.5	-4.1	31.5	-4.1	31.5	0.3	4.0	-0.001	0.8
F30000	13.8	268.1	-2.5	6.6	-2.5	6.5	-2.5	6.5	0.05	1.8	-0.0005	0.8
F100000	20.8	541.8	-0.5	-0.4	-0.5	-0.4	-0.5	-0.4	-0.06	-1.3	-0.001	0.8
nu	34.1	1164.4	0.4	-0.2	0.4	-0.2	0.4	-0.4	0.02	-0.6	-0.0001	0.8

Table 4.2: Skewnesses and kurtoses of the features and target, computed in the original scale, after logarithmic transformation and outlier purge, and after further scalings: `StandardScaler`, `RobustScaler`, `PowerTransformer` and `QuantileTransformer`. The column in bold refers to the transform finally selected for this work.

DR3 identify it with RX J1111.5+3452, whose synchrotron peak frequency is given as $\nu_{\text{syn}} \approx 1.26 \times 10^{17}$ Hz by the 3HSP catalog (Chang et al. 2019), and $\nu_{\text{syn}} \approx 3.98 \times 10^{17}$ Hz by BlaST. It is sensible to suppose that the 4LAC-DR3 peak frequency estimate is not reliable, so the matching instance was deleted from the training dataset.

TRANSFORMATION

A positive skewness is convenient, as it allows to **make the distributions more symmetric by simply applying a (base-10) logarithmic transform**, a standard procedure. This, unfortunately, forced **discarding the fractional variability**, one of the most correlated quantities to ν_{syn} : `Frac_Variability` is exactly 0 for 261 sources, so converting to logarithms is impossible. Considering the definition of variability and fractional variability in 4FGL, a null entry could be conceived as a sort of “upper limit”, and would need the same treatment reserved to missing values, which is beyond the scope of this thesis.

Both the features and the target were converted with the standard method in *Scikit-learn*, by using `FunctionTransformer`, the class in the `sklearn.preprocessing` module to apply custom transforms. Table 4.2 shows how skewness and kurtosis, while remaining often large, greatly improved after discarding the worst outlier and converting to logarithms. Figure 4.7 shows the new frequency histogram and probability plot for ν_{syn} . A further step, at this point, would be making input data as “training-friendly” as possible, by further **scaling or transforming their distributions to make them closer to standardized Gaussians**. Four classes from the `sklearn.preprocessing` module were applied and compared in this work (see Scikit-learn Developers 2024):

StandardScaler It standardizes data by removing the mean and scaling data to unit




Figure 4.7: Frequency histogram and probability plot of the target variable, ν_{syn} (ν) in logarithmic scale.

variance based on the relevant statistics of each quantity.

RobustScaler It removes the mean and scales the data according to the interquartile range of each single distribution. For this reason, it is more robust to the presence of outliers, but it is not suited to sparse data tables.

PowerTransformer Transforms, in general, alter data by applying some non-linear function to them. Power transforms are strictly monotonic, and attempt to map input distributions to standard Gaussians; there are several types, including the Yeo-Johnson, which was chosen here for being usable on negative data.

QuantileTransformer It maps data with a continuous cumulative distribution function, $F(x)$ to another distribution with quantile function $G(u)$ using the formula $G^{-1}(F(x))$. It is powerful in smoothing unusual distributions, but it heavily distorts correlations and distances (for example, between the main “body” of the distribution and outliers, which are “absorbed” and lose their “anomalous” nature, potentially producing biases).

The features and the target were transformed separately: the target directly, with the `.fit_transform` function of the chosen scaler or transformer, while the feature dataframe was transformed with the help of `sklearn.compose.ColumnTransformer`, a class built specifically to apply scalers and transforms to each column separately. The same class was applied to all features and the target here, but each data column was fit and transformed separately according to the statistics of its distribution. Table 4.2 shows the skewnesses and kurtoses computed for each scaler or transform tested in this work. Applying these data processing techniques improves the values sensibly, making them closer to 0 as expected for a standardized Gaussian distribution, but outcomes depend

4.3. IMPLEMENTATION

on the type: while there is not much difference between the two scalers and between them and the simple logarithmic transformation (as scalers only manipulate the mean and the variance), the two transformers are distinct in the fact that `QuantileTransformer` works to deform all the distributions to the same new one in a much stricter way than `PowerTransformer`, producing almost identical values of the two coefficients for all data columns. For simplicity, **from now on the transformer that gave the best results, `PowerTransformer`, will be discussed** on most of the obtained results. Figure 4.8 shows the distribution and probability plot of `log_nu` when using `PowerTransformer`. Figure 4.9

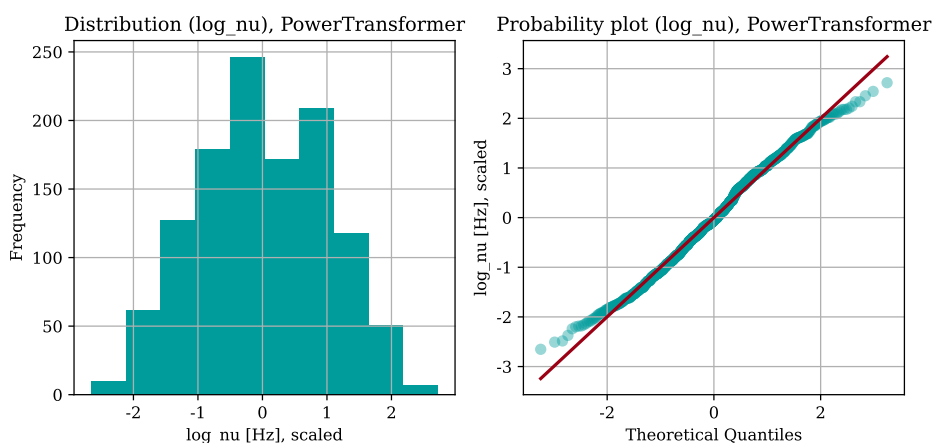




Figure 4.8: Histogram and probability plot for the target variable, ν_{syn} (ν) transformed with logarithms + `PowerTransformer`.

compares feature and target box plots for the logarithmic transformation and the logarithmic transformation plus the `PowerTransformer`; they were obtained with the default parameters of the `pandas.DataFrame.boxplot` method. Notice how rescaling or transforming data allows to restrict the value ranges to similar orders of magnitude, helping both users and machine learning models to handle data more easily. The fact that there are so many outliers, and that they tend to form clusters, indicates how some features (fluxes, in particular) seem to have bimodal distributions, with a second peak that appears as a clump of outliers. This “bizarre” behavior is even more evident in Figure 4.10, that shows the scatter plots of all the features versus ν_{syn} , to evaluate visually the possible existence of one-to-one correlations of some type that may lead the learning phase. On the other hand, the scatter plots of two much more “standard” features, the pivot energy and the photon index, seem to hint at a roughly linear relation to ν_{syn} , coherently with the correlation matrix outcome. Other correlations are up to the model to discover.

The “new” dataset, whose instances are all the 4LAC-DR3 sources without a ν_{syn} estimate, **will need to be transformed in the same way as the training dataset to allow for target prediction**. After converting to base-10 logarithms, the new data columns were transformed with the `.transform` function of the same `ColumnTransformer` used on the training dataset, clearly without refitting it to guarantee all data related to the



(a)



(b)

Figure 4.9: Box plots of features and target. 4.9a: logarithmic transformation. 4.9b: logarithmic transformation + PowerTransformer. Red solid lines: medians; red dashed lines: means.

4.3. IMPLEMENTATION

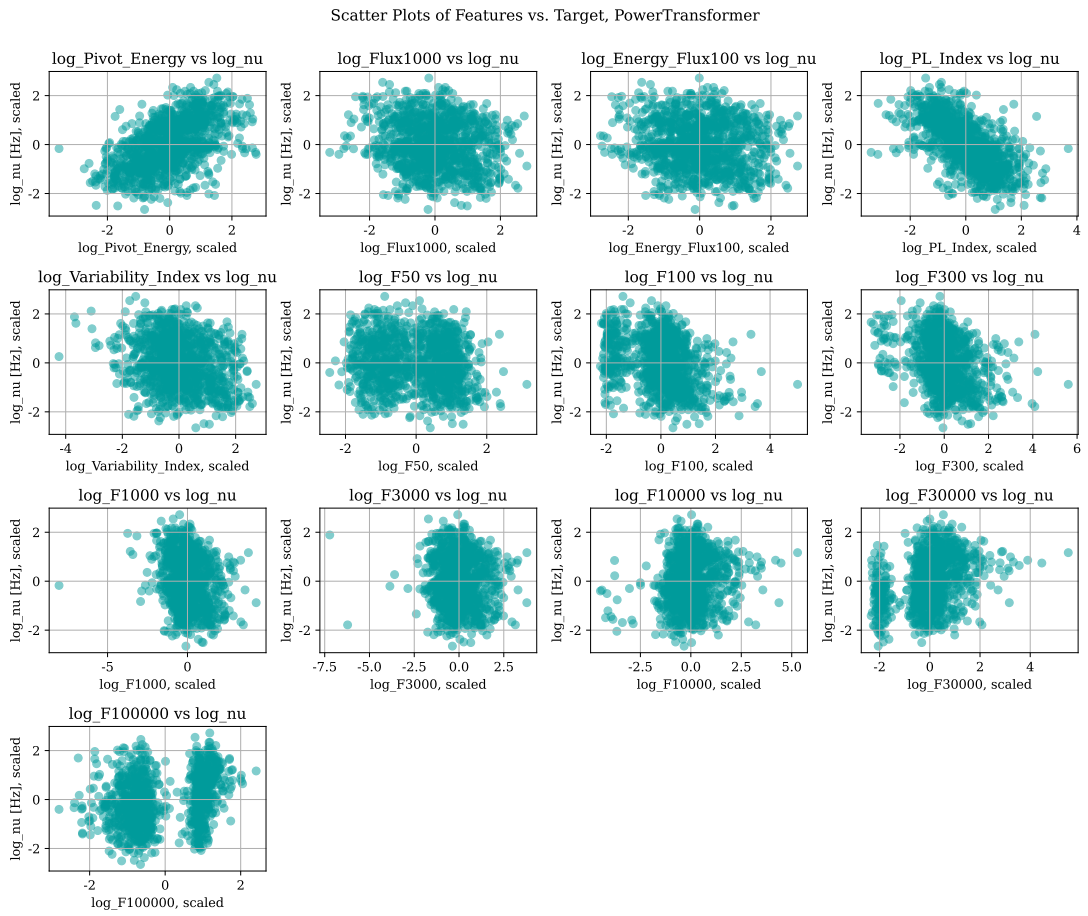


Figure 4.10: Scatter plots of the features against the target ν_{syn} (ν) for transformed data (logarithms + PowerTransformer).

same quantity, both training and new, are scaled in the same way. This would ensure that predictions are done correctly, and allow to compare the distributions of the same feature in the two datasets, to check their similarity and therefore estimate what could come out of the prediction beforehand. In the specific case treated here, there is a very good agreement between training and new data, so it is reasonable to expect for the predicted ν_{syn} to be distributed similarly to the ones in the training dataset: an example is given in Figure 4.11 for the “weirdest” distribution, that of the integral photon flux from 100 GeV to 1 TeV. The values of the target variable predicted on the new dataset




Figure 4.11: Frequency density histograms of F100000 from training and new datasets, transformed (logarithms + PowerTransformer).

will be, of course, scaled according to the transforms applied to the training data, so they need to be back-transformed to the original scale, using the `.inverse_transform` function of the target transformer.

4.3.4 TRAINING

As already mentioned, **learning** correctly consists in three phases: training, validation, and test. *Scikit-learn* provides a class to easily implement *k-fold cross-validation* (see box in Section 4.2): `KFold`. There is not a general rule to choose the number of *k*-folds: in general, as Hastie, Tibshirani, and Friedman 2009 point out, using a smaller number of splits would introduce a bias for smaller datasets, as there would be too few training data, but using too many folds would make training subsets too close to the entire

4.3. IMPLEMENTATION

dataset, and variance would increase. Different possibilities were therefore tested: for this work, k -fold cross-validation was implemented by looping on a list of possible numbers of splits: 5, 10, or 20 folds, clearly equivalent to using 20%, 10% or 5% of the input data for validation.

The training procedure was implemented in this way (Code B.2 in Appendix B):

1. The loop over the k -folds was nested inside another loop over a dictionary of the three chosen models, in order to **select the best number of dataset splits for each model separately**.
2. Another nested dictionary instead included lists of possible hyperparameters for each model, to be scanned by the `GridSearchCV`. The training results (found in the `GridSearchCV.cv_results_` attribute) were written to a list, that consequently includes the training and testing times and scores and their averages for each model, number of k -folds tested for that model, and single fold.
3. Once `GridSearchCV` found the best hyperparameter combination for the model and that number of folds, its `.fit` method was invoked to refit the best estimator on the entire dataset.
4. At this point, the `cross_val_score` function took the dataset, the given `KFold`, the chosen scorer and the best model (`GridSearchCV.best_estimator_`) as arguments to give a final evaluation across the k -folds, computing values that could be finally averaged to give the overall performance score of the estimator.
5. An `if-else` statement was used to compare it with the one from the previous iteration of the loop on the number of k -folds; if the score was better than the previous one for that model, then **the number of k -folds, the hyperparameter combination, and the best estimator were recorded in the respective lists of the best settings and results for each of the three machine learning models**.
6. Another loop, this time over the `best_models` dictionary, allowed to predict the values of the target variable, $v_{\text{syn}}(\text{nu})$ on the training dataset, make a **final evaluation of the performance metrics for each best model**, and plot the results, which are presented in the next Section.

The performance metric chosen to evaluate the model was the **mean squared error** (MSE), one of the available parameters of the `sklearn.metrics.make_scorer` method. Since it is a loss metric, it needs to be minimized; `GridSearchCV` and `cross_val_score`, however, work by maximizing the score, so the output of these classes will be a negative number: the root mean squared error (rMSE) will be therefore computed by extracting the square root of -MSE. The coefficient of determination R^2 was also computed, but not used for validation.

User-defined hyperparameters (Code B.1 in Appendix B) were:

random_state The seed of random numbers generation that regulates data splitting and feature permutation in decision trees. Fixing it ensures reproducibility.

n_estimators Either the number of trees (random forest) or the number of boosting stages to perform (gradient boosting regressor).

`validation_fraction` The fraction of data per fold that is set aside for validation, in order to implement *early stopping* (see box in 4.2) in boosted tree models.

`no_iter_no_change` The number of iterations over which the result should not improve before implementing early stopping.

`max_iter` Only for the histogram-based gradient boosting regressor, it is the maximum number of iterations of the boosting process and is equivalent to the maximum number of trees.

`max_depth` The maximum number of nodes in each tree: None means that trees can split indefinitely, until all leaves are pure or contain less than a certain number of items; this may induce the model to overfit, so the other values in the list were chosen to be relatively small.

Other hyperparameters were kept at their default values.

4.4 RESULTS

4.4.1 TRAINING AND EVALUATION

The training part was actually performed a first time without early stopping: it produced a results dictionary that includes all the **training and test times** for all the values of `n_estimators` or `max_iter`. This object was filtered to select, for each model, the estimators with the best number of folds and hyperparameters, except for the ones just mentioned. Generally, the more trees or iterations there are, the longer it will take to train and test the model: each item in the `n_estimators` or `max_iter` matches a mean fit time, a mean score time (usually much smaller, and therefore ignored) and, of course, the corresponding mean values of the performance metrics. Plotting the fit times against the respective MSE and R^2 scores, as in Figure 4.12, will show how quickly each model reaches its optimal performance, and how it behaves if the number of trees or iterations is increased. As expected, the **gradient boosting models**, especially the histogram-based one which is optimized for large datasets, **reach their best much faster than the random forest**, but the latter is more stable, as the scores tend to saturate with increasing numbers of trees. Gradient boosting estimators instead see their metrics worsen after an early peak, indicating overfitting.

Of course, gradient boosted trees were re-trained with early stopping. The Grid SearchCV choice of best hyperparameters for all three models **did not return much different results across data scalers or transforms**: this is expected for tree-based models, which are robust to variations in data quality and scaling. An exception may be the `QuantileTransformer`, which needed double the number of estimators (trees) in the

4.4. RESULTS




Figure 4.12: Fit times vs. negative mean squared error and R^2 for the three models, each fit with the best hyperparameters, without early stopping, on data processed with a logarithmic transformation and PowerTransformer. Points were plotted without error bars for clarity: $\sigma_{\text{MSE}} \sim 0.08$, $\sigma_{R^2} \sim 0.09$.

random forest and the gradient boosting regressor, and a tree depth of 5 nodes instead of 3; for all other transforms, the **best settings** are:

Random Forest: `n_splits=20 (KFolds), max_depth=5, n_estimators=500`

Gradient Boosting: `n_splits=5, max_depth=3, n_estimators=50`

Histogram-based Gradient Boosting: `n_splits=20, max_depth=3, max_iter=50`

Notice that the number of trees or iterations for gradient boosting models is half the one needed by the random forest, confirming what has already been seen from the fit times - scores plot. Table 4.3 shows the MSEs and R^2 coefficients for all the three models and all the data transforms. It can be noticed immediately how the scores across different models and data scalings are almost identical, with differences that sometimes appear at the fourth decimal digit only. This again shows the relatively little impact of data pre-

	Random Forest		Gradient Boosting		Histogram-b. Grad. Boost.	
	MSE	R^2	MSE	R^2	MSE	R^2
Log	0.837	0.487	0.832	0.490	0.832	0.490
Standard	0.837	0.487	0.833	0.490	0.832	0.490
Robust	0.837	0.488	0.834	0.489	0.832	0.490
Power	0.843	0.484	0.846	0.482	0.826	0.494
Quantile	0.854	0.477	0.839	0.486	0.823	0.492

Table 4.3: Performance metrics of the best estimators for each machine learning model and data scaler. All scores were computed after back-transforming predictions on the training dataset to the original logarithmic scale, to make them comparable.

processing (other than the logarithmic transformation) in this specific case, and for tree-based models in general. Curiously, the same data transform, `QuantileTransformer`, yields both the worst result (with the random forest) and the best one in terms of mean squared error (with the histogram-based gradient boosting); the best result in terms of R^2 is obtained by the histogram-based gradient boosting regressor applied to data pre-processed with `PowerTransformer`. There is a **hint at some predictive power**, as the R^2 coefficients are always positive, but it is already clear from here that the quality of the selected estimators is not high and they are **not able to generalize data correctly**. Obtaining relatively bad score values ($R^2 < 0.5$ always in this specific case) was anyway expected, given the chosen models (Hastie, Tibshirani, and Friedman 2009).

To better understand the nature of the problem, whether it is a case of overfitting in particular, a standard check is to compare the fit and score performance metrics from the `results` dictionary. Mean results from the example estimators (`PowerTransformer`) are described in Table 4.4. Note that mean squared errors here are different from the final ones in the table above due to being computed at different steps of the procedure:

4.4. RESULTS

	Random Forest	Gradient Boosting	Histogram-b. Grad. Boost.
MSE (train)	0.365 ± 0.003	0.37 ± 0.01	0.401 ± 0.010
MSE (test)	0.51 ± 0.08	0.51 ± 0.04	0.50 ± 0.09
R^2 (train)	0.635 ± 0.003	0.63 ± 0.01	0.599 ± 0.009
R^2 (test)	0.48 ± 0.09	0.49 ± 0.03	0.49 ± 0.09

Table 4.4: Comparison of mean performance metrics for the train and test stages of the three best models (data transformed with logarithms + PowerTransformer).

Table 4.4 gives the MSEs averaged across the k -folds during training, while Table 4.3 computed a single final value, after refitting the best models on the *whole* dataset, based on the deviations of the predicted target values from the actual ones, without averaging over folds and after back-transforming to the original logarithmic scale. A natural slight increase in mean squared errors from training to testing can be seen, while there is a noticeable decrease in the R^2 coefficient. Even if the latter score is not good in any case, so the model is not excessively optimized on training data, this behavior is usually considered a sign of slight overfitting, as the given estimator is shown to be worse at generalizing what it has learned from the training subsets to new data.

The quality of training can be visually inspected by plotting the predicted target variables against the actual ones (Figure 4.13), the frequency histogram of predictions against real values (Figure 4.14), the predicted and actual values against the sample index (Figure 4.15) and the residuals, from which the final mean squared errors are computed (Figure 4.16). For simplicity, only the results from the histogram-based gradient boosting regressor trained on data pre-processed with logarithmic transformation and PowerTransformer are presented; the results from the other models trained on the same dataset are in Appendix B. Notice how, in the scatter plots, samples are roughly clustered around the bisector (representing a perfect match of predictions and actual values), but there is a large dispersion in the horizontal direction. This means that all the estimators have some predictive power, but they all have the tendency to **“squeeze” predictions around the average values**, being unable to reproduce the correct values for the lowest- and the highest-synchrotron peaked BL Lacs. This is more evident in the “line plot” and the histograms, while the residuals show that the worst-predicted values of v_{syn} are more than three orders of magnitude from the actual values, but most other samples are grouped around zero. Table 4.5 compares the main statistics of the “true” and predicted values distributions, computed with `pd.DataFrame.describe()`; it confirms that the selected estimators tend to give predictions that are less dispersed around the (almost coincident) means and medians of the respective distributions. Some scatter plots, especially the random forest one (Figure B.2), show visible gaps in the points distribution, which is also present, less evidently, in the frequency histograms of the predictions. This may be due to the usage of tree-based models, that work by splitting




Figure 4.13: Scatter plot of predictions against actual values of the target, \log_{ν} [Hz]. Histogram-based gradient boosting fit on a transformed dataset (logarithms + Power Transformer).

	Actual values	Random forest	Gradient boosting	Histogram-b. Grad. Boost.
Mean	14.71	14.66	14.66	14.66
Standard deviation	1.28	0.89	0.89	0.89
Minimum value	11.92	13.00	12.64	12.68
25%	13.72	13.90	13.91	13.88
50% (median)	14.55	14.64	14.66	14.68
75%	15.65	15.38	15.35	15.40
Maximum value	19.05	16.66	16.96	16.80

Table 4.5: Statistics of the actual values of \log_{ν} [Hz] distribution and the prediction distributions for the three best models fit on the transformed dataset without outlier (logarithms + PowerTransformer). Predictions were back-transformed to the logarithmic scale.

4.4. RESULTS




Figure 4.14: Frequency histograms of predictions against actual values of the target, \log_{ν} [Hz]. Histogram-based gradient boosting fit on a transformed dataset (logarithms + Power Transformer).




Figure 4.15: Line plot of the predicted and actual values of the target, \log_{ν} [Hz], for the histogram-based gradient boosting regressor fit on transformed data (logarithms + Power Transformer).




Figure 4.16: Line plot of the fit residuals of the target, \log_{ν} [Hz], for the histogram-based gradient boosting regressor fit on transformed data (logarithms + `PowerTransformer`).

data until a discrete number of leaf nodes is reached: even if random forests average over several trees, it is reasonable to suppose that when the number of nodes and trees is too small, the smoothing effect of averaging is hindered by the fact that several trees would classify the same samples in the same leaves.

An advantage of tree-based models, on the other hand, is that data splitting relies on feature values meeting certain thresholds. The depth of the node defined by a certain feature (i.e. what fraction of the training dataset is involved) and the decrease in impurity produced by the split are used by *Scikit-learn* to compute a normalized estimate of the predictive power of that variable (*feature importance*: see Hastie, Tibshirani, and Friedman 2009 for a more general and formal definition). The results are plotted in Figure 4.17 for the three models. They need to be taken with a grain of salt, as importances for the random forest and the gradient boosting regressors are evaluated on statistics computed on the entire training dataset (Scikit-learn Developers 2024). The importances for the histogram-based gradient boosting regressor were estimated by *permutation importance*, which is a technique that should solve this issue by shuffling values for a given variable and assessing how the model performance worsens, but applying it rigorously would have required complicating the training loop beyond manageability, so this case, too, may suffer the same problem (Scikit-learn Developers 2024; Molnar 2022). What emerges, in any case, is that **all three models give an enormous relative weight to the photon index Γ (`log_PL_Index`)**. This was expected, as shown by the target-features scatter plots. The pivot energy, which had the highest correlation coefficient with ν_{syn} and whose scatter plot also hinted at a linear relation with the target, comes in second

4.4. RESULTS

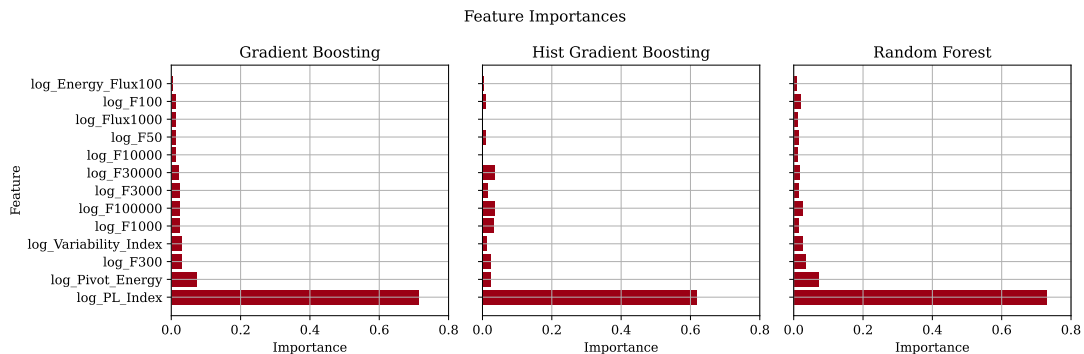


Figure 4.17: Feature importances bar chart for the three models trained on transformed data (logarithms + `PowerTransformer`). From left to right: gradient boosting, histogram-based gradient boosting, random forest.

according to the random forest and gradient boosting models, but with a much smaller importance. In the histogram gradient boosting case, it even sits behind some of the fluxes, especially the ones at the highest energies. In general, the three models show the expected behavior, with the random forest distributing importances more equally to the various features, at least the “secondary” ones (Hastie, Tibshirani, and Friedman 2009).

4.4.2 PREDICTING ON NEW DATA

The three models, now ready for use, were applied to a new, previously **unseen dataset** to finally predict the missing BL Lac ν_{syn} values in 4LAC-DR3. The new data were scaled according to the transforms fit on the training dataset, as explained in Section 4.3.3 before being fed to the best estimators found in the previous steps; after prediction, the output was back-transformed to the logarithmic scale by applying the inverse target transform. As usual, the results obtained by training on logarithmic data scaled with the `PowerTransformer` are chosen as an example.

The second, third and fourth column of Table 4.6 show the distribution statistics of the predictions from the new dataset (scaled back from `PowerTransformer`). Values seem to have a relatively small dispersion, with no ν_{syn} below 10^{13} Hz or above 10^{17} Hz (no lower-peaked LBLs and no extreme HBLs), coherently with the predictions on the training dataset done during learning. This is even more evident when looking at the frequency histograms in Figures 4.18 and 4.19, where the “comb-like” behavior, with wide gaps in the distribution, also reappears.

	Training data	Random forest	Gradient boosting	Histogram-b. Grad. Boost.	BlaST
Mean	14.71	14.86	14.85	14.83	15.58
Standard deviation	1.28	0.80	0.80	0.80	1.36
Minimum value	11.92	13.30	13.21	13.27	12.50
25%	13.72	14.18	14.14	14.11	14.58
50% (median)	14.55	14.88	14.93	14.88	15.60
75%	15.65	15.66	15.50	15.53	16.70
Maximum value	19.05	16.57	16.52	16.67	18.50

Table 4.6: Statistics of the $\log_{10} \nu$ [Hz] distributions from the training dataset, the prediction distributions for the three best models fit on the new dataset (logarithms + PowerTransformer), back-transformed to the logarithmic scale, and the BlaST predictions.

Comparison with Other Tools To have a better assessment of the quality of these results, a comparison was made⁴ with another machine learning tool with the same task: the Blazar Synchrotron Tool⁵, or **BlaST** (Glauch, Kerscher, and Giommi 2022), already mentioned in Section 4.3.3. It was developed to compute unbiased estimates of ν_{syn} and confidence intervals, based on broadband SED data that may include spurious components from the host galaxy or its nucleus, with limited need to clean data beforehand. Its working principle is an ensemble of **neural networks**, built with the PyTorch package (Paszke et al. 2019), each trained on a dataset of 3793 blazar SEDs from a subsample of the Open Universe for Blazars master list (v. 2)⁶. Actually, in order to avoid bias due to some ν_{syn} being much more frequent than others, oversampling was applied to fill “emptier” bins with altered copies of other real SEDs of the same bin, bringing the total number of instances in the training dataset (including the validation subset) to 10400 real or augmented samples.

Figure 4.18 shows the frequency distribution of values predicted by the three example models in this work against the BlaST outcomes. As revealed in Table 4.6, the predicted ν_{syn} from this work are spread over a smaller range, with a taller peak close to the mean and the median of the distribution, while BlaST values follow a wider distribution, with higher mean and median and dozens of predicted extreme BL Lacs.

The agreement with the known 4LAC-DR2 values seems to be high in general, according to the BlaST paper, so there is a possibility that the “new” sources tend to be intrinsically shifted to higher frequencies and that the lower values for ν_{syn} found by this work are a model quality issue; however, the paper also underlines that BlaST tends to overestimate peak frequencies when $\nu_{\text{syn}} > 10^{16}$ Hz, and that there is some disagreement with 4LAC in the 10^{14} Hz to 10^{15} Hz interval, right where means and medians of the new predictions fall.

⁴The BlaST predictions were kindly provided by Dott.ssa Elisa Prandini.

⁵<https://github.com/tkerscher/blast>

⁶<https://openuniverse.asi.it/OU4Blazars/MasterListV2/>

4.4. RESULTS




Figure 4.18: Frequency histograms of \log_{ν} [Hz] values predicted on the new dataset by the three models (trained on data scaled with logarithms + PowerTransformer; colored line histograms), and the BlaST predictions on the same sources (gray filled histogram).

Comparison with the Input Dataset Figure 4.19 shows the the frequency densities of the new values of ν_{syn} predicted by the three models, compared with the frequency density histogram from the input dataset: as highlighted in Section 4.3.3, the distributions of features in the training and the “new” dataset are very similar, so it is reasonable to check the quality of the new predictions by comparing them with the values of ν_{syn} found in the input dataset, even if it is not conceptually correct as new data may have a different distribution. Table 4.6 also includes a summary of the distribution statistics




Figure 4.19: Frequency densities of \log_{ν} values predicted by the three models (colored line histograms) and as found in the training dataset (gray filled histogram). Models were trained on data scaled with logarithms + `PowerTransformer`.

of training targets for comparison. Predictions on new data are, again, spread over a smaller interval of synchrotron peak frequencies; on the other hand, they are in good agreement with the target predictions computed on the training dataset for model evaluation, already shown in Table 4.5. This lends weight to the idea that the models built in this work are not optimal for prediction.

4.5 FINAL CONSIDERATIONS

Summary

Target variable: synchrotron peak frequency ν_{syn} [Hz] of BL Lac blazars as found in the 4LAC-DR3 catalog (nu).

Training features: observed quantities on BL Lac blazars as found in the 4FGL-DR4 catalog, from *Fermi*-LAT data in the energy range 50 MeV-1 TeV:

- **Pivot energy, E_0** [MeV]: energy at which the statistical error on the differential flux is minimal (`Pivot_Energy`).
- **Flux1000** [$\text{cm}^{-2}\text{s}^{-1}$]: integral photon flux measured from 1 GeV to 100 GeV.
- **Energy_Flux100** [$\text{erg cm}^{-2}\text{s}^{-1}$]: energy flux from 100 MeV to 100 GeV obtained by spectral fitting.
- **Photon index, Γ** : photon index when fitting the spectrum with a power law (`PL_Index`).
- **Photon fluxes in 8 energy bands** [$\text{erg cm}^{-2}\text{s}^{-1}$]: 50 MeV to 100 MeV (F50), 100 MeV to 300 MeV (F100), 300 MeV to 1 GeV (F300), 1 GeV to 3 GeV (F1000), 3 GeV to 10 GeV (F3000), 10 GeV to 30 GeV (F10000), 30 GeV to 100 GeV (F30000), 100 GeV to 1 TeV (F100000).
- **Variability index TS_{var}** : likelihood test statistic on the variability of a source (Section 1.2.2).

Best Data Transformations: \log_{10} + Yeo-Johnson Power Transformer on all quantities separately.

Best Learning Model: Histogram-based Gradient Boosting with custom hyperparameters:

- **No. of k-folds:** 20
- **Max. tree depth:** 3
- **Max. number of iterations:** 50
- **Validation fraction:** 0.2
- **No. of iterations with no score improvement before early stopping:** 10

Validation metrics: $\text{MSE} \sim 0.8$, $R^2 \sim 0.5$ for all models.

The outcome of the machine learning models implemented here was not of exceptional quality if compared with the training dataset and a more complex tool like BlaST, but it still showed some potential for future development, and the shortcomings of the selected estimators can be explained by the characteristics of tree-based models in gen-

eral. There could be several roads to learning improvement:

- **Test alternative models.** As already explained, tree-based algorithms are based on a weak estimator that has a limited predictive power and that, while built to catch non-linear relations, is lacking in presence of linearity or linear combinations among few variables (see e.g. Hastie, Tibshirani, and Friedman 2009). The *Scikit-learn* documentation (Scikit-learn Developers 2024) is clear: ensemble regressors should be considered as a last resort when datasets are relatively small (< 100000 instances) and there are reasons to suppose that the target values will depend from a small number of instances; in that case, using linear regressors like Lasso and ElasticNet would be preferable. Scatter plots in Fig. 4.10 and, in part, feature importances (Fig. 4.17) seem to suggest this is the case. Neural networks are another popular choice, and are the type of estimator BlaST, which also takes spectral points (but not variability and photon indices) as an input, was built on, after attempts with a random forest and a gradient boosting regressor (Glauch, Kerscher, and Giommi 2022);
- **Hyperparameter optimization.** Values scanned here are either the defaults, or quite arbitrary. A more careful choice of which values to scan with GridSearch CV, or setting hyperparameters at first kept as defaults (like the *learning rate* for gradient boosting regressors) would benefit efficiency and accuracy;
- **Refine the input dataset.** The three models considered in this work belong to a class that works decently “out-of-the-box”, such that data pre-processing does not have a great impact, but machine learning in general benefits from good data engineering before the learning phase. Potential improvements would include:
 - **Treating outliers more carefully.** Only one source was discarded here, as it was clearly off and cleaning the dataset from outliers could be unimportant for tree-based models. In case the estimator were changed, handling outliers could become much more relevant;
 - **Refining feature scalings and transforms,** tuning their parameters and applying the best transform to each variable separately;
 - **Considering bias in the training dataset.** Instances are just a subset of all possible BL Lacs, the ones that are bright enough in the *Fermi*-LAT energy range to be observed by that instrument. The BlaST paper points out how the target variable has a distribution, that often leaves some ν_{syn} bins underpopulated; models trained on biased datasets will propagate that bias in their predictions, as they could not learn some combinations of feature and target values that may exist in new data, but are underrepresented in the input dataset. An idea that emerged while working on this Chapter was not to train the models on real data, but on a large sample of simulated *Fermi*-LAT spectra covering all plausible spectral model parameter combinations, maybe repurposing the MMDC neural network training dataset.
- **Work on features.** As already mentioned, the models in this work seem to rely on few of the features only. A different choice of independent variables could potentially alter the result, maybe positively: for example, the features that could be the less useful may be filtered out, or others that were initially left out could be added back. A prominent example is the fractional variability, *Frac_Variability*, which looked like one of the most strongly correlated features to ν_{syn} (ν) but was discarded due to often being exactly zero (and therefore not being convertible to logarithmic scale). Another idea could be including the parameters of the alternative model for curved spectra (LogParabola);

4.5. FINAL CONSIDERATIONS

- **Consider the underlying physics.** As already mentioned, the relation between gamma-ray photon index and ν_{syn} , which seems to drive most of the learning, has been known and discussed for several years and may have a physical justification for both BL Lacs and FSRQs in the context of a one-zone synchrotron + inverse Compton model (see e.g. Dermer et al. 2015). The 4LAC paper suggests to study the correlation “more physically”, correcting ν_{syn} for redshift by the (small) factor $1 + z$.

In any case, **trying to obtain a machine learning model that only relies on high-energy gamma ray data is probably a worthy task**, for the reason explained in the “Purpose” section: *Fermi* data are ubiquitous in the blazar world, while observations at other energies are not granted, or retrieving them is not straightforward due to defects of on-line catalogs. The outlier source, for example, despite having an association in both 4LAC and 4FGL catalogs, has separate entries in Firmamento: 4FGL J1111.5+3455 has *Fermi*-LAT data only, the associated object RX J1111.5+3452 instead features a full broadband SED. Consequently, there are BlaST estimates of ν_{syn} for both: needless to say, the two values are very different, as already shown above, and the one based on 4FGL data only is probably unreliable. BlaST, with its apparent accuracy and robustness, is not infallible: apart from the issue just mentioned, it turned out that 36 of the 276 BL Lacs without ν_{syn} estimate in 4LAC-DR3 (on which predictions were done in this work) lack even a BlaST evaluation of this quantity. Checking which sources were, and which observations are available of them, to understand if it is an issue of spectral data scarcity, was out of the scope of this thesis. If that is the case, creating a new tool like the one sketched in this work could be important.

5

Conclusions

This work was carried out as part of an effort to integrate novel machine learning techniques into high-level data analysis performed by researchers of the High-Energy Astrophysics group in Padua, especially when concerning the physical and phenomenological characterization of blazars. Machine learning, in particular, could be suited for tackling issues like fit instability and parameter degeneracy in modeling blazar spectral energy distributions, and to perform population studies that require handling large amounts of data and discovering non-linear relations among them.

Investigating on the so-called *blazar sequence*, a possible anti-correlation between synchrotron peak frequency and bolometric luminosity of blazars, is exactly a task of this type, and the main topic of this thesis. The primary aim was to discuss the physical origin of the sequence and the role of “BL Lac” blazars that have been detected at TeV energies, by adopting a new approach based on the modeling of the spectral energy distribution of individual sources. A recent learning-based fitting tool, available on the Markarian Multiwavelength Data Center (MMDC) web page, was tested alongside more traditional analytical packages to fit the broadband spectra of three TeV BL Lacs. This first part of the work, illustrated in Chapter 3, showed the potential of the MMDC tool for blazar modeling, especially in terms of user-friendliness, speed, and ability to look for the global best fit over the whole parameter space. The experience with the other analysis package, *agpy* and its fitting backends *Sherpa* and *Gammapy*, revealed very clearly the issue of selecting non-simultaneous data in the X-ray band. No good compromise between the two alternative data choices was found, so no firm conclusions can be drawn on the TeV blazar sequence yet, even though some patterns were detected in MMDC results. Longer-term perspectives for the future also include extending the modeling to all the blazars in the sample, hopefully after having found a reliable rule for unbiased data selection, or working with new information from upcoming high-energy

instruments. A need which emerged in this study is that of either X-ray and VHE simultaneous data gathered during average or low emission states of the sources of interest, or time-integrated spectra from large-field of view X-ray instruments. The first solution could be one of the objectives of the Cherenkov Telescope Array (CTA: The Cherenkov Telescope Array Consortium et al. 2019), the next-generation very-high-energy gamma-ray observatory that has recently began its main development phase, with the aim of being operative in a few years. Covering the energy range from 20 GeV to 300 TeV, CTA is expected to have 10 times better sensitivity than current IACT arrays, but will not be able to work in the energy interval currently surveyed by *Fermi*-LAT, which is probably going to be decommissioned in the near future. In the X-ray band, a perfect instrument for of this study would have been eROSITA (Predehl et al. 2021), which, however, was prematurely turned off in 2022 in protest against the Russian invasion of Ukraine. New instruments with improved sensitivity could also be fundamental in providing the necessary information to either confirm or disprove the actual existence of the blazar sequence, which some researchers think is driven more by beaming effect or selection bias than physics. In absence of recent, in-depth surveys of the high-energy sky outside the *Fermi*-LAT range, the alternative is to go back to archival simultaneous observations, a fine-tuned approach that this work intended to avoid. The need for a more careful recording of observation times in Open Universe for Blazars catalogs, and of easy time filtering options in online multiwavelength databases, became apparent.

A further issue emerged during source selection: the lack of reliable estimates of the synchrotron peak frequency for several gamma-ray blazars, the starting point in the new approach to the study of the sequence. As it is reasonably caused by scarcity of multi-wavelength data at low energies, which makes the traditional evaluation of this quantity by fitting spectral points difficult, a machine learning code was developed and described in Chapter 4 to predict the synchrotron peak frequency of BL Lac blazars based exclusively on *Fermi*-LAT data in the high-energy gamma-ray band. An assessment of the model performance based on evaluation statistics and comparison with the prediction of another machine learning-based tool, BlaST, revealed that the algorithm is promising, but has large margins of improvement in terms of predictive power and overfitting prevention. The fact that BlaST, which is trained on broadband spectra, is unable to make predictions for a few *Fermi*-detected sources, shows that carrying on developing a machine learning tool based on high-energy gamma-ray data only is probably worth the effort.

References

- Abdalla, H. et al. (2020). “Simultaneous observations of the blazar PKS 2155–304 from ultra-violet to TeV energies”. en. In: *Astronomy & Astrophysics* 639, A42. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201936900.
- Abdo, A. A. et al. (2010). “The Spectral Energy Distribution of Fermi Bright Blazars”. In: *ApJ* 716.1, pp. 30–70. DOI: 10.1088/0004-637X/716/1/30. arXiv: 0912.2040 [astro-ph.CO].
- Abdo, A. A. et al. (2011a). “FERMI LARGE AREA TELESCOPE OBSERVATIONS OF MARKARIAN 421: THE MISSING PIECE OF ITS SPECTRAL ENERGY DISTRIBUTION”. en. In: *The Astrophysical Journal* 736.2, p. 131. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/736/2/131.
- Abdo, A. A. et al. (2011b). “THE FIRST FERMI MULTIFREQUENCY CAMPAIGN ON BL LACERTAE: CHARACTERIZING THE LOW-ACTIVITY STATE OF THE EPONYMOUS BLAZAR”. en. In: *The Astrophysical Journal* 730.2, p. 101. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/730/2/101.
- Abdollahi, S. et al. (2020). “Fermi Large Area Telescope Fourth Source Catalog”. In: *The Astrophysical Journal Supplement Series* 247.1. Publisher: The American Astronomical Society, p. 33. ISSN: 0067-0049. DOI: 10.3847/1538-4365/ab6bcb.
- Abdollahi, S. et al. (2022). “Incremental Fermi Large Area Telescope Fourth Source Catalog”. In: *The Astrophysical Journal Supplement Series* 260.2. Publisher: The American Astronomical Society, p. 53. ISSN: 0067-0049. DOI: 10.3847/1538-4365/ac6751.
- Abdollahi, S. et al. (2023). “The Fermi-LAT Light Curve Repository”. en. In: *The Astrophysical Journal Supplement Series* 265.2. arXiv:2301.01607 [astro-ph], p. 31. ISSN: 0067-0049, 1538-4365. DOI: 10.3847/1538-4365/acbb6a.
- Abe, H. et al. (2023). “Multimessenger Characterization of Markarian 501 during Historically Low X-Ray and γ -Ray Activity”. en. In: *The Astrophysical Journal Supplement Series* 266.2. Publisher: The American Astronomical Society, p. 37. ISSN: 0067-0049. DOI: 10.3847/1538-4365/acc181.
- Acciari, V. A. et al. (2009). “VERITAS OBSERVATIONS OF THE BL LAC OBJECT 1ES 1218+304”. en. In: *The Astrophysical Journal* 695.2. Publisher: American Astronomical

REFERENCES

- Society, pp. 1370–1375. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637x/695/2/1370.
- Acciari, V. A. et al. (2010). “DISCOVERY OF VARIABILITY IN THE VERY HIGH ENERGY γ -RAY EMISSION OF 1ES 1218+304 WITH VERITAS”. en. In: *The Astrophysical Journal* 709.2. Publisher: American Astronomical Society, pp. L163–L167. ISSN: 2041-8205, 2041-8213. DOI: 10.1088/2041-8205/709/2/1163.
- Acciari, V. A. et al. (2020). “New Hard-TeV Extreme Blazars Detected with the MAGIC Telescopes*”. en. In: *The Astrophysical Journal Supplement Series* 247.1, p. 16. ISSN: 0067-0049, 1538-4365. DOI: 10.3847/1538-4365/ab5b98.
- Acciari, V. A. et al. (2022). “Investigating the Blazar TXS 0506+056 through Sharp Multiwavelength Eyes During 2017–2019”. en. In: *The Astrophysical Journal* 927.2. Publisher: The American Astronomical Society, p. 197. ISSN: 0004-637X. DOI: 10.3847/1538-4357/ac531d.
- Acerro, F. et al. (2015). “FERMI LARGE AREA TELESCOPE THIRD SOURCE CATALOG”. en. In: *The Astrophysical Journal Supplement Series* 218.2, p. 23. ISSN: 1538-4365. DOI: 10.1088/0067-0049/218/2/23.
- Acharyya, A. et al. (2023). “VTSCat: The VERITAS Catalog of Gamma-Ray Observations”. In: *Research Notes of the AAS* 7.1, p. 6. ISSN: 2515-5172. DOI: 10.3847/2515-5172/acb147.
- Ackermann, M. et al. (2015). “The Third Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope”. In: *ApJ* 810.1, 14, p. 14. DOI: 10.1088/0004-637X/810/1/14. arXiv: 1501.06054 [astro-ph.HE].
- Adams, C. B. et al. (2022). “The throughput calibration of the VERITAS telescopes”. en. In: *Astronomy & Astrophysics* 658. arXiv:2111.04676 [astro-ph], A83. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202142275.
- Aharonian, F. et al. (2005). “H.E.S.S. observations of PKS 2155-304”. In: *A&A* 430, pp. 865–875. DOI: 10.1051/0004-6361:20041853. arXiv: astro-ph/0411582 [astro-ph].
- Aharonian, F. et al. (2007). “An Exceptional Very High Energy Gamma-Ray Flare of PKS 2155-304”. In: *ApJ* 664.2, pp. L71–L74. DOI: 10.1086/520635. arXiv: 0706.0797 [astro-ph].
- Aharonian, F. et al. (2010). “Discovery of VHE γ -rays from the BL Lacertae object PKS 0548–322”. en. In: *Astronomy & Astrophysics* 521. Publisher: EDP Sciences, A69. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/200912363.
- Aharonian, F. (2000). “TeV gamma rays from BL Lac objects due to synchrotron radiation of extremely high energy protons”. In: *New Astronomy* 5.7, pp. 377–395. ISSN: 1384-1076. DOI: 10.1016/s1384-1076(00)00039-7.
- Ahnen, M. L. et al. (2018). “Detection of the blazar S4 0954+65 at very-high-energy with the MAGIC telescopes during an exceptionally high optical state”. en. In: *Astronomy*

- & Astrophysics* 617. Publisher: EDP Sciences, A30. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201832624.
- Ajello, M. et al. (2017). “3FHL: The Third Catalog of Hard Fermi-LAT Sources”. en. In: *The Astrophysical Journal Supplement Series* 232.2. Publisher: American Astronomical Society, p. 18. ISSN: 0067-0049, 1538-4365. DOI: 10.3847/1538-4365/aa8221.
- Ajello, M. et al. (2020). “The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope”. In: *The Astrophysical Journal* 892.2, p. 105. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/ab791e.
- Ajello, M. et al. (2022). “The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope: Data Release 3”. In: *The Astrophysical Journal Supplement Series* 263.2, p. 24. ISSN: 0067-0049, 1538-4365. DOI: 10.3847/1538-4365/ac9523.
- Ajello, M. et al. (2023). *The 4LAC-DR3 catalog*. Centre de Données astronomique de Strasbourg (CDS). DOI: <https://doi.org/10.26093/cds/vizier.22630024>.
- Albert, J. (2008). “Implementation of the Random Forest Method for the Imaging Atmospheric Cherenkov Telescope MAGIC”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 588.3, pp. 424–432. ISSN: 01689002. DOI: 10.1016/j.nima.2007.11.068. arXiv: 0709.3719[astro-ph].
- Albert, J. et al. (2006). “Discovery of Very High Energy Gamma Rays from 1ES 1218+30.4”. In: *ApJ* 642.2, pp. L119–L122. DOI: 10.1086/504845. arXiv: astro-ph/0603529 [astro-ph].
- Aleksić, J. et al. (2014). “MAGIC long-term study of the distant TeV blazar PKS 1424+240 in a multiwavelength context”. en. In: *Astronomy & Astrophysics* 567. Publisher: EDP Sciences, A135. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201423364.
- Aleksić, J. et al. (2016a). “The major upgrade of the MAGIC telescopes, Part I: The hardware improvements and the commissioning of the system”. In: *Astroparticle Physics* 72, pp. 61–75. ISSN: 0927-6505. DOI: <https://doi.org/10.1016/j.astropartphys.2015.04.004>.
- Aleksić, J. et al. (2016b). “The major upgrade of the MAGIC telescopes, Part II: A performance study using observations of the Crab Nebula”. en. In: *Astroparticle Physics* 72, pp. 76–94. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2015.02.005.
- Ansoldi, S. et al. (2018). “The Blazar TXS 0506+056 Associated with a High-energy Neutrino: Insights into Extragalactic Jets and Cosmic-Ray Acceleration”. en. In: *The Astrophysical Journal Letters* 863.1, p. L10. ISSN: 2041-8205, 2041-8213. DOI: 10.3847/2041-8213/aad083.
- Atwood, W. B. et al. (2009). “THE LARGE AREA TELESCOPE ON THE FERMI GAMMA-RAY SPACE TELESCOPE MISSION”. en. In: *The Astrophysical Journal* 697.2, pp. 1071–1102. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/697/2/1071.

REFERENCES

- Ballet, J. et al. (2024). *Fermi Large Area Telescope Fourth Source Catalog Data Release 4 (4FGL-DR4)*. arXiv: 2307.12546 [astro-ph].
- Bégué, D. et al. (2023). *Modeling blazar broadband emission with convolutional neural networks – I. Synchrotron self-Compton model*. en. arXiv:2311.02979 [astro-ph]. URL: <http://arxiv.org/abs/2311.02979>.
- Biteau, J. et al. (2020). “Progress in unveiling extreme particle acceleration in persistent astrophysical jets”. en. In: *Nature Astronomy* 4.2. arXiv:2001.09222 [astro-ph, physics:gr-qc], pp. 124–131. ISSN: 2397-3366. DOI: 10.1038/s41550-019-0988-4.
- Blandford, R. D. and D. G. Payne (1982). “Hydromagnetic flows from accretion discs and the production of radio jets”. en. In: *Monthly Notices of the Royal Astronomical Society* 199.4, pp. 883–903. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/199.4.883.
- Blandford, R. D. and R. L. Znajek (1977). “Electromagnetic extraction of energy from Kerr black holes”. In: *Monthly Notices of the Royal Astronomical Society* 179.3, pp. 433–456. ISSN: 0035-8711. DOI: 10.1093/mnras/179.3.433.
- Boch, T., F. Pineau, and S. Derriere (2012). “The CDS Cross-Match Service”. In: *Astronomical Data Analysis Software and Systems XXI*. Ed. by P. Ballester, D. Egret, and N. P. F. Lorente. Vol. 461. Astronomical Society of the Pacific Conference Series, p. 291.
- Boella, G. et al. (1997). “BeppoSAX, the wide band mission for X-ray astronomy”. en. In: *Astronomy and Astrophysics Supplement Series* 122.2, pp. 299–307. ISSN: 0365-0138, 1286-4846. DOI: 10.1051/aas:1997136.
- Bonnoli, G. et al. (2010). “The γ -ray brightest days of the blazar 3C 454.3”. In: *Monthly Notices of the Royal Astronomical Society* 410.1, pp. 368–380. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2010.17450.x.
- Burke, D. et al. (2023). *sherpa/sherpa: Sherpa 4.15.1*. DOI: 10.5281/zenodo.7948720. URL: <https://doi.org/10.5281/zenodo.7948720>.
- Chadwick, P. M. et al. (1999). “PKS 2155-304 - a source of VHE γ -rays”. In: *Astroparticle Physics* 11.1-2, pp. 145–148. DOI: 10.1016/S0927-6505(99)00039-0. arXiv: astro-ph/9812122 [astro-ph].
- Chang, Y.-L., C. Brandt, and P. Giommi (2020). “The Open Universe VOU-Blazars tool”. In: *Astronomy and Computing* 30, p. 100350. ISSN: 2213-1337. DOI: 10.1016/j.ascom.2019.100350.
- Chang, Y.-L. et al. (2019). “The 3HSP catalogue of extreme and high-synchrotron peaked blazars”. en. In: *Astronomy & Astrophysics* 632, A77. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201834526.
- Colla, G. et al. (1970). “A catalogue of 3235 radio sources at 408 MHz”. In: *A&AS* 1.3, p. 281.
- Costamante, L. et al. (2001). “Extreme synchrotron BL Lac objects: Stretching the blazar sequence”. en. In: *Astronomy & Astrophysics* 371.2, pp. 512–526. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361:20010412.

- Cutri, R. M. et al. (2012). *Explanatory Supplement to the WISE All-Sky Data Release Products*. Explanatory Supplement to the WISE All-Sky Data Release Products.
- D’Ammando, F. et al. (2024). “Optical spectroscopy of blazars for the Cherenkov Telescope Array – III”. en. In: *Astronomy & Astrophysics* 683, A222. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202348507.
- Dermer, C. D. et al. (2015). “NEAR-EQUIPARTITION JETS WITH LOG-PARABOLA ELECTRON ENERGY DISTRIBUTION AND THE BLAZAR SPECTRAL-INDEX DIAGRAMS”. en. In: *The Astrophysical Journal* 809.2, p. 174. ISSN: 1538-4357. DOI: 10.1088/0004-637X/809/2/174.
- Dermer, C. D. and B. Giebels (2016). “Active galactic nuclei at gamma-ray energies”. en. In: *Comptes Rendus. Physique* 17.6. arXiv:1602.06592 [astro-ph], pp. 594–616. ISSN: 1878-1535. DOI: 10.1016/j.crhy.2016.04.004.
- Devanand, P. U. et al. (2022). “Study of X-Ray Intraday Variability of HBL Blazars Based on Observations Obtained with XMM-Newton”. en. In: *The Astrophysical Journal* 939.2, p. 80. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/ac9064.
- Donath, A. et al. (2023). “Gammapy: A Python package for gamma-ray astronomy”. In: *A&A* 678, A157. DOI: 10.1051/0004-6361/202346488.
- Donato, D. et al. (2001). “Hard X-ray properties of blazars”. en. In: *Astronomy & Astrophysics* 375.3, pp. 739–751. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361:20010675.
- Fan, J. H. et al. (2017). “Intrinsic Correlations for Flaring Blazars Detected by Fermi”. In: *ApJ* 835.2, L38, p. L38. DOI: 10.3847/2041-8213/835/2/L38. arXiv: 1702.05692 [astro-ph.HE].
- Feroz, F., M. P. Hobson, and M. Bridges (2009). “MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics”. In: *Monthly Notices of the Royal Astronomical Society* 398.4, pp. 1601–1614. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2009.14548.x.
- Finke, J. D. (2013). “COMPTON DOMINANCE AND THE BLAZAR SEQUENCE”. en. In: *The Astrophysical Journal* 763.2, p. 134. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/763/2/134.
- Finke, J. D., C. D. Dermer, and M. Böttcher (2008). “Synchrotron Self-Compton Analysis of TeV X-Ray-Selected BL Lacertae Objects”. en. In: *The Astrophysical Journal* 686.1, pp. 181–194. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/590900.
- Fortin, P. (2008). “VERITAS Observations of the BL Lac Object 1ES 1218+304”. en. In: arXiv:0810.0301 [astro-ph], pp. 565–568. DOI: 10.1063/1.3076735.
- Fossati, G. et al. (1998). “A unifying view of the spectral energy distributions of blazars”. en. In: *Monthly Notices of the Royal Astronomical Society* 299.2, pp. 433–448. ISSN: 0035-8711, 1365-2966. DOI: 10.1046/j.1365-8711.1998.01828.x.

REFERENCES

- Gasparyan, S., D. Bégué, and N. Sahakyan (2021). “Time-dependent lepto-hadronic modeling of the emission from blazar jets with SOPRANO: the case of TXS 0506+056, 3HSP J095507.9+355101 and 3C 279”. en. In: *Monthly Notices of the Royal Astronomical Society* 509.2. arXiv:2110.01549 [astro-ph], pp. 2102–2121. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stab2688.
- Gehrels, N. et al. (2004). “The Swift Gamma-Ray Burst Mission”. In: *ApJ* 611.2, pp. 1005–1020. DOI: 10.1086/422091. arXiv: astro-ph/0405233 [astro-ph].
- Ghisellini, G. and F. Tavecchio (2008). “The blazar sequence: a new perspective”. en. In: *Monthly Notices of the Royal Astronomical Society* 387.4, pp. 1669–1680. ISSN: 00358711, 13652966. DOI: 10.1111/j.1365-2966.2008.13360.x.
- Ghisellini, G. et al. (2011). “The transition between BL Lac objects and flat spectrum radio quasars”. In: *Monthly Notices of the Royal Astronomical Society* 414.3, pp. 2674–2689. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2011.18578.x.
- Ghisellini, G. et al. (2017). “The Fermi blazar sequence”. In: *Monthly Notices of the Royal Astronomical Society* 469.1, pp. 255–266. ISSN: 0035-8711. DOI: 10.1093/mnras/stx806.
- Ghisellini, G. (2013). *Radiative Processes in High Energy Astrophysics*. en. Vol. 873. arXiv:1202.5949 [astro-ph]. DOI: 10.1007/978-3-319-00612-3.
- Giommi, P. et al. (2012). “A simplified view of blazars: clearing the fog around long-standing selection effects: A simplified view of blazars”. en. In: *Monthly Notices of the Royal Astronomical Society* 420.4, pp. 2899–2911. ISSN: 00358711. DOI: 10.1111/j.1365-2966.2011.20044.x.
- Giommi, P. et al. (2019). “Open Universe for Blazars: a new generation of astronomical products based on 14 years of *Swift*-XRT data”. en. In: *Astronomy & Astrophysics* 631. Publisher: EDP Sciences, A116. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201935646.
- Giommi, P. and P. Padovani (2021). *Astrophysical Neutrinos and Blazars*. en. arXiv:2112.06232 [astro-ph]. URL: <http://arxiv.org/abs/2112.06232>.
- Giommi, P. et al. (2002). *A Catalog of 157 X-ray Spectra and 84 Spectral Energy Distributions of Blazars observed with BeppoSAX*. en. arXiv:astro-ph/0209596.
- Glauch, T., T. Kerscher, and P. Giommi (2022). *BlaST – A Machine-Learning Estimator for the Synchrotron Peak of Blazars*. arXiv: 2207.03813 [astro-ph, physics:hep-ex].
- Goldoni, P. et al. (2021). “Optical spectroscopy of Blazars for the Cherenkov Telescope Array”. en. In: *Astronomy & Astrophysics* 650. arXiv:2012.05176 [astro-ph], A106. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202040090.
- Gréaux, L. et al. (2023). *STeVECAt, the Spectral TeV Extragalactic Catalog*. en. arXiv:2304.00835 [astro-ph]. URL: <http://arxiv.org/abs/2304.00835>.
- H.E.S.S. Collaboration et al. (2018). “The H.E.S.S. Galactic plane survey”. en. In: *Astronomy & Astrophysics* 612, A1. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201732098.

- Hartman, R. C. et al. (1999). “The Third EGRET Catalog of High-Energy Gamma-Ray Sources”. In: *ApJS* 123.1, pp. 79–202. DOI: 10.1086/313231.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The elements of statistical learning: data mining, inference and prediction*. 2nd ed. Springer.
- Hogg, D. W. et al. (2002). *The K correction*. en. arXiv:astro-ph/0210394. URL: <http://arxiv.org/abs/astro-ph/0210394>.
- IceCube Collaboration et al. (2018). “Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert”. In: *Science* 361.6398, pp. 147–151. DOI: 10.1126/science.aat2890.
- Issaoun, S. et al. (2022). “Resolving the Inner Parsec of the Blazar J1924–2914 with the Event Horizon Telescope”. en. In: *The Astrophysical Journal* 934.2, p. 145. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/ac7a40.
- Kasai, E. et al. (2022). “Optical Spectroscopy of Blazars for the Cherenkov Telescope Array – II”. en. In: *Monthly Notices of the Royal Astronomical Society* 518.2. arXiv:2211.04689 [astro-ph], pp. 2675–2692. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stac3167.
- Ke, G. et al. (2017). “LightGBM: A Highly Efficient Gradient Boosting Decision Tree”. In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc.
- Kerby, S. and A. D. Falcone (2023). “Testing the Blazar Sequence with Spectra of Recently Discovered Dim Blazars from the Fermi Unassociated Catalog”. en. In: *The Astrophysical Journal* 951.2, p. 133. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/acd4c0.
- Kouch, P. M. et al. (2024). “IXPE observation of PKS 2155–304 reveals the most highly polarized blazar”. en. In: *Astronomy & Astrophysics* 689, A119. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202449166.
- Kowalek, P., H. Loch-Olszewska, and J. Szwabiński (2019). “Classification of diffusion modes in single-particle tracking data: Feature-based versus deep-learning approach”. In: *Physical Review E* 100. DOI: 10.1103/PhysRevE.100.032410.
- Kuulkers, E. et al. (2021). “INTEGRAL reloaded: Spacecraft, instruments and ground system”. In: *New Astronomy Reviews* 93, p. 101629. ISSN: 1387-6473. DOI: <https://doi.org/10.1016/j.newar.2021.101629>.
- Loreti, M. (2006). *Teoria degli errori e fondamenti di statistica*. ita. Padova.
- Lott, B., D. Gasparrini, and S. Ciprini (2020). *The Fourth Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope – Data Release 2*. arXiv: 2010.08406 [astro-ph].
- Meegan, C. et al. (2009). “The Fermi Gamma-ray Burst Monitor”. In: *ApJ* 702.1, pp. 791–804. DOI: 10.1088/0004-637X/702/1/791. arXiv: 0908.0450 [astro-ph.IM].

REFERENCES

- Michelucci, U. (2022). *Applied Deep Learning with TensorFlow 2: Learn to Implement Advanced Deep Learning Techniques with Python*. en. Berkeley, CA: Apress. ISBN: 978-1-4842-8019-5 978-1-4842-8020-1. DOI: 10.1007/978-1-4842-8020-1.
- Molnar, C. (2022). *Interpretable Machine Learning. A Guide for Making Black Box Models Explainable*. 2nd ed. URL: <https://christophm.github.io/interpretable-ml-book>.
- Nigro, C. et al. (2022). “agnpy: An open-source python package modelling the radiative processes of jetted active galactic nuclei”. en. In: *Astronomy & Astrophysics* 660, A18. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202142000.
- Nolan, P. L. et al. (2012). “FERMI LARGE AREA TELESCOPE SECOND SOURCE CATALOG”. In: *The Astrophysical Journal Supplement Series* 199.2, p. 31. ISSN: 1538-4365. DOI: 10.1088/0067-0049/199/2/31.
- Ochsenbein, F., P. Bauer, and J. Marcout (2000). “The VizieR database of astronomical catalogues”. en. In: *Astronomy and Astrophysics Supplement Series* 143.1, pp. 23–32. ISSN: 0365-0138, 1286-4846. DOI: 10.1051/aas:2000169.
- Ochsenbein, F. e. a. (n.d.). *The VizieR database of astronomical catalogues*. DOI: 10.26093/cds/vizier.
- Oh, K. et al. (2018). “The 105-Month Swift-BAT All-sky Hard X-Ray Survey”. en. In: *The Astrophysical Journal Supplement Series* 235.1, p. 4. ISSN: 0067-0049, 1538-4365. DOI: 10.3847/1538-4365/aaa7fd.
- Ohm, S. and S. Wagner (2023). “Current status and operation of the H.E.S.S. array of imaging atmospheric Cherenkov telescopes”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1055, p. 168442. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2023.168442>.
- Ouyang, Z. et al. (2023). “The “Blazar Sequence” in the TeV Band”. en. In: *The Astrophysical Journal* 949.2. Publisher: The American Astronomical Society, p. 52. ISSN: 0004-637X. DOI: 10.3847/1538-4357/acc329.
- Padovani, P. et al. (2019). “TXS 0506+056, the first cosmic neutrino source, is not a BL Lac”. en. In: *Monthly Notices of the Royal Astronomical Society: Letters* 484.1, pp. L104–L108. ISSN: 1745-3925, 1745-3933. DOI: 10.1093/mnrasl/slz011.
- Padovani, P. et al. (2017). “Active Galactic Nuclei: what’s in a name?” en. In: *The Astronomy and Astrophysics Review* 25.1. arXiv:1707.07134 [astro-ph], p. 2. ISSN: 0935-4956, 1432-0754. DOI: 10.1007/s00159-017-0102-9.
- Paszke, A. et al. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. arXiv: 1912.01703 [cs.LG].
- Pedregosa, F. et al. (2011). “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12, pp. 2825–2830.

- Peñil, P. et al. (2023). “Multiwavelength analysis of Fermi-LAT blazars with high-significance periodicity: detection of a long-term rising emission in PG 1553+113”. In: *Monthly Notices of the Royal Astronomical Society* 527.4, pp. 10168–10184. ISSN: 0035-8711. DOI: 10.1093/mnras/stad3246.
- Pineau, F.-X. et al. (2020). “The CDS Cross-match Service: Key Figures, Internals and Future Plans”. In: *Astronomical Data Analysis Software and Systems XXVII*. Ed. by P. Ballester et al. Vol. 522. Astronomical Society of the Pacific Conference Series, p. 125.
- Prandini, E. and G. Ghisellini (2022). “The Blazar Sequence and Its Physical Understanding”. en. In: *Galaxies* 10.1. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute, p. 35. ISSN: 2075-4434. DOI: 10.3390/galaxies10010035.
- Predehl, P. et al. (2021). “The eROSITA X-ray telescope on SRG”. en. In: *Astronomy & Astrophysics* 647, A1. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/202039313.
- Qin, L. et al. (2018). “Using the Markov Chain Monte Carlo method to study the physical properties GeV-TeV BL Lac objects”. en. In: *Publications of the Astronomical Society of Japan* 70.1. arXiv:1711.10625 [astro-ph], p. 5. ISSN: 0004-6264, 2053-051X. DOI: 10.1093/pasj/psx150.
- Rieger, F. M. (2019). *Gamma-Ray Astrophysics in the Time Domain*. en. arXiv:1901.10216 [astro-ph]. URL: <http://arxiv.org/abs/1901.10216>.
- Sahakyan, N. (2020). “Broad-band study of high-synchrotron-peaked BL Lac object 1ES 1218+304”. In: *Monthly Notices of the Royal Astronomical Society* 496.4, pp. 5518–5527. ISSN: 0035-8711. DOI: 10.1093/mnras/staa1893.
- Sahakyan, N. and P. Giommi (2022). “A 13-yr-long broad-band view of BL Lac”. In: *Monthly Notices of the Royal Astronomical Society* 513.3, pp. 4645–4656. ISSN: 0035-8711. DOI: 10.1093/mnras/stac1011.
- Sahakyan, N. et al. (2024a). *Markarian Multiwavelength Data Center (MMDC): A Tool for Retrieving and Modeling Multi-temporal, Multi-wavelength and Multi-messenger Data from Blazar Observations*. en. arXiv:2410.01207 [astro-ph].
- Sahakyan, N. et al. (2024b). *Modeling blazar broadband emission with convolutional neural networks – II. External Compton model*. en. arXiv:2402.07495 [astro-ph]. URL: <http://arxiv.org/abs/2402.07495>.
- Saldana-Lopez, A. et al. (2021). “An observational determination of the evolving extragalactic background light from the multiwavelength HST/CANDELS survey in the Fermi and CTA era”. In: *MNRAS* 507.4, pp. 5144–5160. DOI: 10.1093/mnras/stab2393. arXiv: 2012.03035 [astro-ph.CO].
- Schmidt, M. (1963). “3C 273: A Star-Like Object with Large Red-Shift”. In: *Nature* 197.4872. Publisher: Nature Publishing Group, pp. 1040–1040. ISSN: 1476-4687. DOI: 10.1038/1971040a0.

REFERENCES

- Schwartz, D. A. et al. (1979). “X-ray emitting BL Lacertae objects located by the scanning modulation collimator experiment on HEAO 1”. en. In: *The Astrophysical Journal* 229, p. L53. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/182929.
- Scikit-learn Developers (2024). *Scikit-learn: Machine Learning in Python*. Accessed: August 29, 2024. URL: https://scikit-learn.org/stable/user_guide.html.
- Sitarek, J. (2022). *TeV Instrumentation: current and future*. en. arXiv:2201.08611 [astro-ph]. URL: <http://arxiv.org/abs/2201.08611>.
- Smith, M. W. E. et al. (2013). “The Astrophysical Multimessenger Observatory Network (AMON)”. In: *Astroparticle Physics* 45, pp. 56–70. DOI: 10.1016/j.astropartphys.2013.03.003. arXiv: 1211.5602 [astro-ph.HE].
- Sol, H. and A. Zech (2022). “Blazars at Very High Energies: Emission Modelling”. en. In: *Galaxies* 10.6. Number: 6 Publisher: Multidisciplinary Digital Publishing Institute, p. 105. ISSN: 2075-4434. DOI: 10.3390/galaxies10060105.
- Spurio, M. (2018). *Probes of Multimessenger Astrophysics: Charged cosmic rays, neutrinos, γ -rays and gravitational waves*. en. Astronomy and Astrophysics Library. Cham: Springer International Publishing. ISBN: 978-3-319-96853-7 978-3-319-96854-4. DOI: 10.1007/978-3-319-96854-4.
- Stamerra, D. A. and G. Bonnoli (2010). “Very High Energy emission from blazars interpreted through simultaneous multiwavelength observations”. en. In.
- Stein, W. A., S. L. O’Dell, and P. A. Strittmatter (1976). “The BL Lacertae Objects”. In: *Annual Review of Astronomy and Astrophysics* 14. Volume 14, 1976. Publisher: Annual Reviews, pp. 173–195. ISSN: 0066-4146, 1545-4282. DOI: 10.1146/annurev.aa.14.090176.001133.
- Superina, G. et al. (2008). “Discovery of VHE gamma-rays from the BL Lac object PKS 0548-322 with H.E.S.S.”. In: *International Cosmic Ray Conference*. Vol. 3. International Cosmic Ray Conference, pp. 913–916.
- Tanaka, Y. T., S. Buson, and D. Kocevski (2017). “Fermi-LAT detection of increased gamma-ray activity of TXS 0506+056, located inside the IceCube-170922A error region.” In: *The Astronomer’s Telegram* 10791, p. 1.
- Tavecchio, F. (2021). *Probing Magnetic Fields and Acceleration Mechanisms in Blazar Jets with X-ray Polarimetry*. en. arXiv:2105.08401 [astro-ph].
- The Cherenkov Telescope Array Consortium et al. (2019). *Science with the Cherenkov Telescope Array*. en. arXiv:1709.07997 [astro-ph, physics:hep-ex]. DOI: 10.1142/10986.
- The Event Horizon Telescope Collaboration et al. (2019). “First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring”. In: *The Astrophysical Journal Letters* 875.1, p. L5. DOI: 10.3847/2041-8213/ab0f43.
- The IceCube Collaboration et al. (2018). “Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A”. In: *Science* 361.6398, eaat1378. DOI: 10.1126/science.aat1378.

- Tripathi, D. et al. (2024). “Firmamento: A Multimessenger Astronomy Tool for Citizen and Professional Scientists”. en. In: *The Astronomical Journal* 167.3. Publisher: The American Astronomical Society, p. 116. ISSN: 1538-3881. DOI: 10.3847/1538-3881/ad216a.
- Ulrich, M.-H., L. Maraschi, and C. M. Urry (1997). “VARIABILITY OF ACTIVE GALACTIC NUCLEI”. en. In.
- Urry, C. M. and P. Padovani (1995). “Unified Schemes for Radio-Loud Active Galactic Nuclei”. en. In: *Publications of the Astronomical Society of the Pacific* 107. arXiv:astro-ph/9506063, p. 803. ISSN: 0004-6280, 1538-3873. DOI: 10.1086/133630.
- Urry, C. M. et al. (2000). “The Hubble Space Telescope Survey of BL Lacertae Objects. II. Host Galaxies”. In: *The Astrophysical Journal* 532.2, p. 816. DOI: 10.1086/308616.
- Venturi, T. et al. (2020). *VLBI20-30: a scientific roadmap for the next decade – The future of the European VLBI Network*. arXiv: 2007.02347 [astro-ph.IM].
- Wakely, S. P. and D. Horan (2008). “TeVcat: An online catalog for Very High Energy Gamma-Ray Astronomy”. In: *International Cosmic Ray Conference*. Vol. 3. International Cosmic Ray Conference, pp. 1341–1344.
- Wright, E. L. et al. (2010). “THE WIDE-FIELD INFRARED SURVEY EXPLORER (WISE): MISSION DESCRIPTION AND INITIAL ON-ORBIT PERFORMANCE”. en. In: *The Astronomical Journal* 140.6, pp. 1868–1881. ISSN: 0004-6256, 1538-3881. DOI: 10.1088/0004-6256/140/6/1868.
- Yershov, V. N. (2014). “Serendipitous UV source catalogues for 10 years of XMM and 5 years of Swift”. en. In: *Astrophysics and Space Science* 354.1, pp. 97–101. ISSN: 1572-946X. DOI: 10.1007/s10509-014-1944-5.
- Yuan, F. and R. Narayan (2014). “Hot Accretion Flows Around Black Holes”. en. In: *Annual Review of Astronomy and Astrophysics* 52.1. arXiv:1401.0586 [astro-ph], pp. 529–588. ISSN: 0066-4146, 1545-4282. DOI: 10.1146/annurev-astro-082812-141003.
- Zhang, P.-f. et al. (2017). “Revisiting Quasi-periodic Modulation in γ -Ray Blazar PKS 2155-304 with Fermi Pass 8 Data”. In: *The Astrophysical Journal* 835.2, p. 260. DOI: 10.3847/1538-4357/835/2/260.
- Zhao, X. Z. et al. (2024). “The Energy Budget in the Jet of High-frequency Peaked BL Lacertae Objects”. en. In: *The Astrophysical Journal* 967.2, p. 104. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/ad3ba9.
- Zhou, R. X. et al. (2021). “The Intrinsic Properties of Multiwavelength Energy Spectra for Fermi Teraelectronvolt Blazars”. en. In: *The Astrophysical Journal* 915.1, p. 59. ISSN: 0004-637X, 1538-4357. DOI: 10.3847/1538-4357/abfe69.

Acknowledgments

The author thanks her supervisor, Dott.ssa Elisa Prandini, her co-supervisor, Dott.ssa Ilaria Viale, and their former colleague, Dott.ssa Chiara Righi (at the time INAF, Osservatorio Astronomico di Merate), for including her in their ongoing project on the blazar sequence, and carrying out most of the source and data selection work described in Chapter 3. The candidate also thanks Prof. Narek Sahakyan (ICRANet Armenia, Yerevan) for help in understanding his team's fitting tool and himself providing some of the fits, that could not be requested through the user interface.

Some of the data used in this work were obtained from the MMDC.

Part of this work is based on archival data, software or online services provided by the Space Science Data Center - ASI.

This research made use of Firmamento, a tool developed at The Center for Astrophysics and Space Science (CASS) and at the office of Research Outreach of the New York University Abu Dhabi, in the framework of the Open Universe initiative (Tripathi et al. 2024, AJ 167 116).

This research has made use of the VizieR catalogue access tool, CDS, Strasbourg, France (Ochsenbein n.d.). The original description of the VizieR service was published in Ochsenbein, Bauer, and Marcout (2000). This research made use of the cross-match service provided by CDS, Strasbourg.

This research has made use of the NASA/IPAC Extragalactic Database (NED), which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology.

This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration.

Ringraziamenti

“Tempo e paja, e maura anca le nespole”. Di tempo ce ne ho messo io, tanto. La paglia invece l’hanno messa tante persone che in modi differenti mi hanno permesso di arrivare, finalmente, a scrivere le ultime parole di questa tesi.

Ringrazio ancora una volta la mia relatrice, la dott.ssa Elisa Prandini e la mia correlatrice, la dott.ssa Ilaria Viale, per avermi dato l’opportunità di contribuire al loro bel progetto sulla blazar sequence, che spero un giorno porterà a dei risultati importanti. Mi hanno sempre fatta sentire valorizzata e coinvolta, senza mai perdere né il rigore, né l’affabilità, nemmeno quando c’era da criticare la mia lentezza e la mia prolissità. È stato bello provare l’ebbrezza di sentirmi un po’ ricercatrice per qualche mese (anche nelle piccole cose, dal badge che apre tutte le porte alla partecipazione al congresso nazionale della SIF), con la soddisfazione in più di aver partecipato a una squadra tutta femminile. Un ringraziamento anche al mio controrelatore, il dott. Roberto Taverna, per le conversazioni costruttive. Ovviamente estendo la mia gratitudine anche chi ha dato un contributo non in veste ufficiale, ma comunque fondamentale: la dott.ssa Chiara Righi, il prof. Narek Sahakyan, e Julia Nguyen; e certamente è stata di aiuto anche la cordialità che ho trovato nell’intero gruppo di ricerca di High-Energy Astrophysics.

Grazie alle mie amiche “collega Kaputt” e “collega Su”, per essere rimaste. No, non vuol dire che siete un ripiego, ma che avete resistito a me e a come sono: è un complimento! La guerra continua.

Avrei voluto esprimere gratitudine anche a tante sciocchezze, dalla musica agli uccelli, dai disagi di Trenitalia alla terrificante quinta puntata di “Viaggio nel cosmo”, ma forse non è il caso.

La mia gratitudine eterna ai miei genitori e mia sorella, per non avermi fatto mancare il sostegno materiale e morale, nei giorni di gioia per un esame andato bene e negli interminabili periodi di sconforto, quando li trasformavo loro malgrado nella condotta di scarico delle mie frustrazioni, e loro continuavano lo stesso a supportarmi. So che è brutto dedicare così poche righe alla propria famiglia, ma non si possono certo riassumere qui, per giunta su pagine pubbliche, così tanti anni di vita, e tanto meno gli ultimi mesi, che purtroppo sono sembrati altrettanto lunghi. Mamma, papà, e anche Cate, sappiate solo che vi voglio bene, ciascuno in qualche mio modo strano.

Appendix

A

Spectral Modeling Plots and Tables

A.1 BIN 4

A.1.1 MMDC FITS

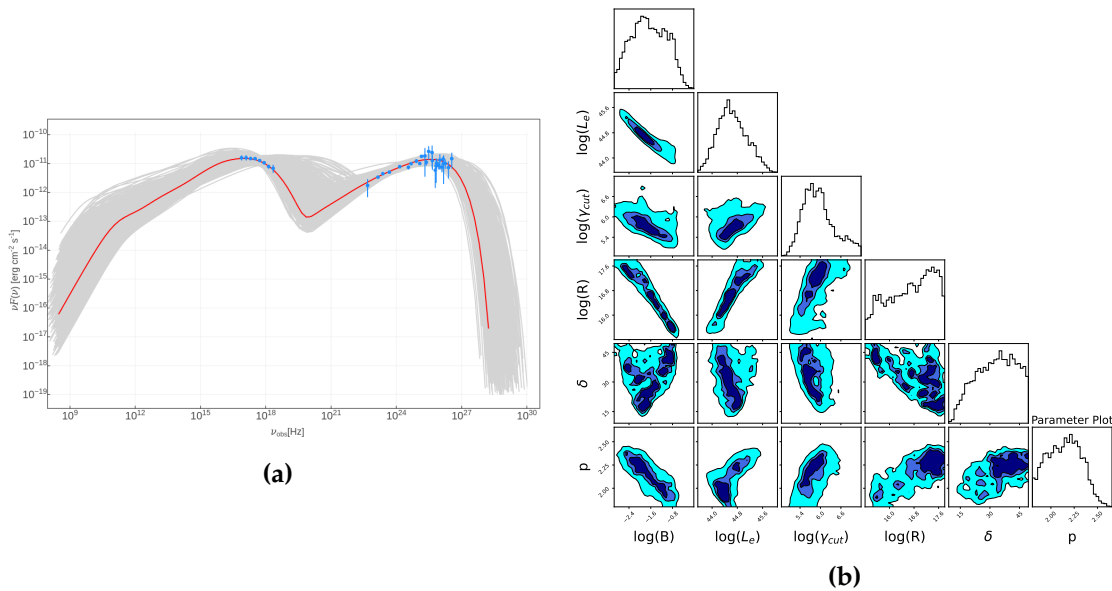


Figure A.1: Results of the MMDC tool fit on PG 1218+304, keeping all model parameters free except for $\gamma_{\text{min}} = 10^2$. A.1a: observed SED (blue dots), best-fit model (red line) and random MCMC samples (gray lines). A.1b: corner plot of the same fit.

A.1. BIN 4

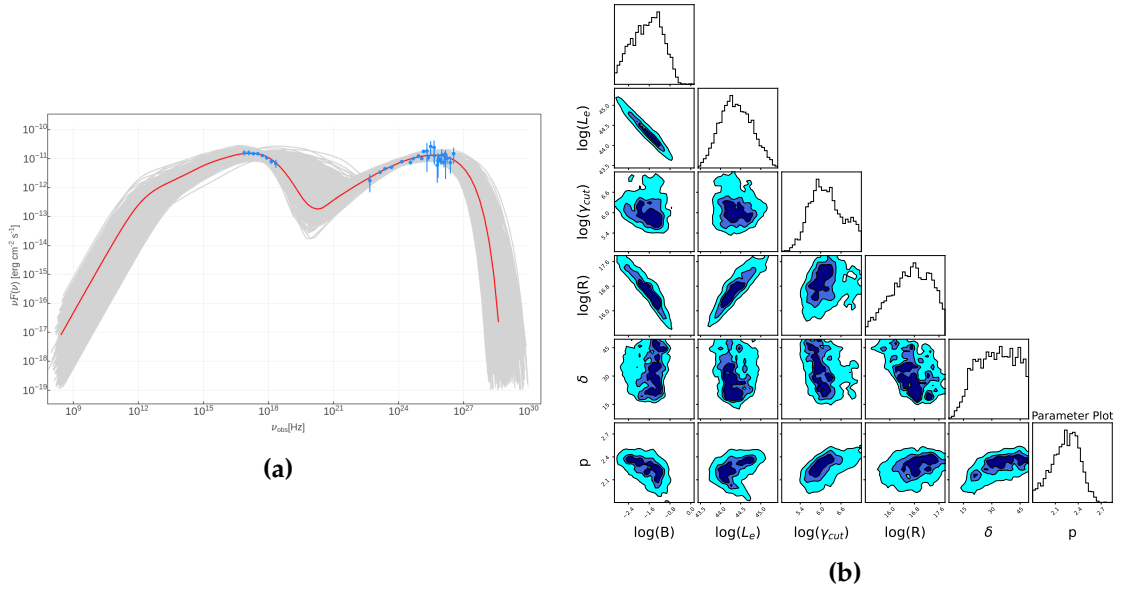


Figure A.2: Results of the MMDC tool fit on PG 1218+304, keeping all model parameters free except for $\gamma_{\min} = 10^3$. A.2a: observed SED (blue dots), best-fit model (red line) and random MCMC samples (gray lines). A.2b: corner plot of the same fit.




Figure A.3: Fit of the PG 1218+304 SED with the mean Swift-XRT data, computed with the MMDC SED fitting tool, fixing $\gamma_{\min} = 10^3$. Left: observed SED (blue dots), best-fit model (red line) and random MCMC samples (gray lines). Right: Best-fit parameters. Dark red values are frozen. t_{var} (teal) was computed at a later moment and is not one of the original fit parameters.

A.1.2 AGNPFY FITS

δ_D initialization	10	15	20	30	40
$\log k_e$ [cm ⁻³]	-8.4 ± 0.4	-8.1 ± 0.5	-7.7 ± 0.4	-7.4 ± 0.2	-7.4 ± 0.2
p_1	2.07 ± 0.06	2.12 ± 0.06	2.13 ± 0.10	2.10 ± 0.06	2.10 ± 0.06
p_2	3.1 ± 0.9	3.4 ± 0.8	3.1 ± 0.3	3.1 ± 0.3	3.1 ± 0.3
$\log \gamma_b$	5.4 ± 0.2	5.4 ± 0.2	5.1 ± 0.2	4.90 ± 0.04	4.90 ± 0.04
$\log \gamma_{\min}$	2	2	2	2	2
$\log \gamma_{\max}$	6	6	6	6	6
δ_D	10 ± 3	16.6 ± 0.8	27 ± 11	35 ± 2	35 ± 2
$\log B$ [G]	-1.36 ± 0.04	-1.4 ± 0.3	-1.7 ± 0.2	-1.8 ± 0.2	-1.8 ± 0.2
t_{var} [s]	$(33 \pm 10) \times 10^4$	$(10 \pm 10) \times 10^4$	$(8 \pm 5) \times 10^4$	$(6 \pm 2) \times 10^4$	$(6 \pm 2) \times 10^4$
$\chi^2/d.o.f.$	0.857	0.782	0.779	0.771	0.771
Q-value	0.677	0.780	0.785	0.795	0.795

Table A.1: Fit parameters and statistics for the SED of PG 1218+304, obtained with agnpy and Sherpa by fixing $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$ and initializing δ_D at various values while keeping it free. Values in dark red are fixed. Values in italics only have the first approximate estimate of their errors.

δ_D initialization	10	15	20	30	40
$\log k_e$ [cm ⁻³]	-8.62 ± 0.10	-8.0 ± 0.4	-7.7 ± 0.4	-7.6 ± 0.5	-7.7 ± 0.4
p_1	2.15 ± 0.02	2.14 ± 0.06	2.12 ± 0.08	2.10 ± 0.10	2.12 ± 0.08
p_2	5 ± 10	3.3 ± 0.3	3.1 ± 0.2	3.0 ± 0.2	3.1 ± 0.2
$\log \gamma_b$	5.76 ± 0.06	5.3 ± 0.2	5.0 ± 0.2	5.0 ± 0.2	5.0 ± 0.2
$\log \gamma_{\min}$	2	2	2	2	2
$\log \gamma_{\max}$	6	6	6	6	6
δ_D	14 ± 4	19 ± 2	26 ± 2	28 ± 2	26 ± 2
$\log B$ [G]	-1.4 ± 0.3	-1.5 ± 0.1	-1.7 ± 0.1	-1.8 ± 0.1	-1.7 ± 0.1
t_{var} [s]	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4	8.64×10^4
$\chi^2/d.o.f.$	1.132	0.749	0.752	0.755	0.752
Q-value	0.287	0.827	0.823	0.820	0.823

Table A.2: Fit parameters and statistics for the SED of PG 1218+304, obtained with agnpy and Sherpa by fixing $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$ and $t_{\text{var}} = 1$ d, and initializing δ_D at various values while keeping it free. Values in dark red are fixed. Values in italics only have the first approximate estimate of their errors.

A.1. BIN 4

agnpy/Sherpa	MMDC	agnpy	agnpy	MMDC	agnpy	agnpy	MMDC	agnpy	MMDC
$\log k_e$ [cm ⁻³]	$\log L_e$ [erg s ⁻¹]	-7.9 ± 0.1	-8.0 ± 0.4	44.6 ± 0.2	-7.8 ± 0.4	-7.5 ± 0.5	44.5 ± 0.1	-7.1 ± 0.9	44.0 ± 0.1
	p_1	2.10 ± 0.04	2.14 ± 0.06	1.85 ± 0.09	2.13 ± 0.08	2.1 ± 0.1	2.1 ± 0.1	1.9 ± 0.3	2.3 ± 0.1
	p_2	3.7 ± 0.3	3.3 ± 0.3		3.1 ± 0.2	3.0 ± 0.3		3.0 ± 0.3	
	$\log \gamma_b$	5.38 ± 0.05	5.3 ± 0.2		5.1 ± 0.2	4.9 ± 0.2		4.6 ± 0.3	
	$\log \gamma_{\min}$	2	2		2	2		2	
	$\log \gamma_{\max}$	6	6	5.8 ± 0.4	6	6	5.7 ± 0.3	6	5.7 ± 0.3
	$\log \gamma_{\text{cut}}$	15	20		25	30		40	
	δ_D								
	$\log B$ [G]	-1.26 ± 0.05	-1.52 ± 0.08	-1.3 ± 0.1	-1.69 ± 0.09	-1.84 ± 0.09	-1.6 ± 0.1	-2.06 ± 0.08	-2.1 ± 0.1
	t_{var} [s]	8.64×10^4	8.64×10^4		8.64×10^4	8.64×10^4		8.64×10^4	
$\chi^2/d.o.f.$		0.745	0.722		0.727	0.729		0.806	
Q-value		0.837	0.861		0.856	0.854		0.760	

Table A.3: Fit parameters and statistics for the SED of PG 1218+304, obtained with agnpy/Sherpa and the MMDC SED fitting tool by fixing $\gamma_{\min} = 10^2$, $\gamma_{\max} = 10^6$ and $t_{\text{var}} = 1$ d, and freezing various values of δ_D . Values in dark red are fixed.

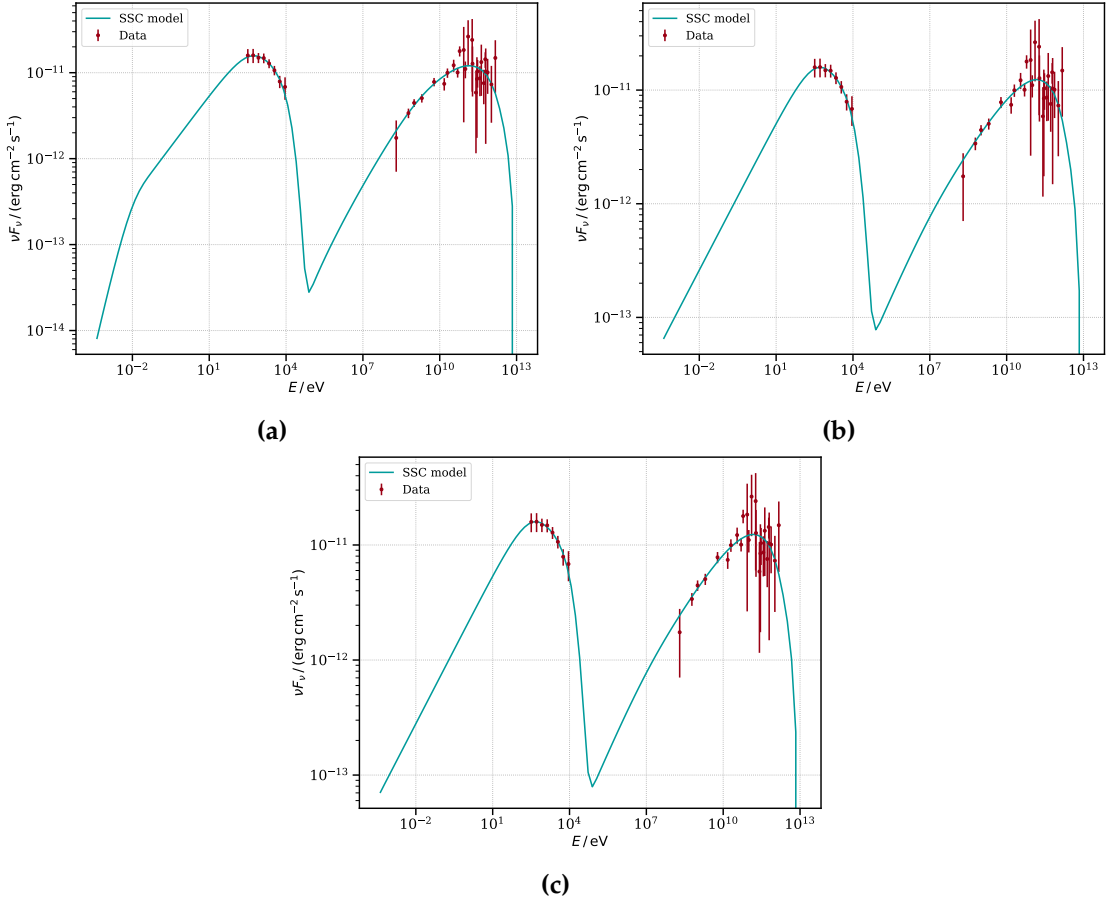


Figure A.4: Observed SED (dark red dots) and best-fit models (teal lines) of PG 1218+304, fixing γ_{\min} , $\gamma_{\max} = 10^6$, and $t_{\text{var}} = 1$ d, with $\delta_D \sim 20$ by fixing or convergence. A.4a: fixed $\gamma_{\min} = 10^3$, initialization with $\delta_D = 10$ left free. A.4b: fixed $\gamma_{\min} = 10^2$, initialization with $\delta_D = 15$ left free. A.4c: fixed $\gamma_{\min} = 10^2$ and $\delta_D = 20$.

A.2 BIN 3

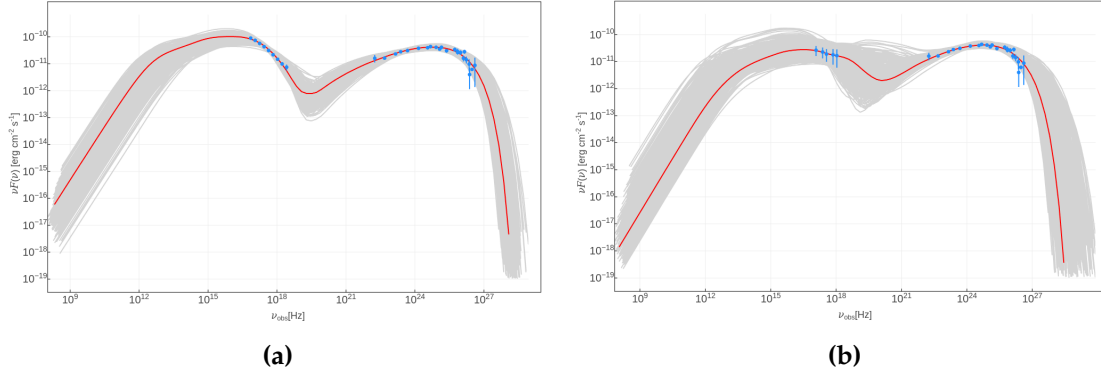


Figure A.5: Observed SED (blue dots), MMDC best-fit models (red line) and random MCMC samples (gray lines) of PKS 2155-304, fixing $\gamma_{\min} = 10^3$. A.5a: BeppoSAX X-ray data. A.5b: mean Swift-XRT X-ray data.

A.3 BIN 5




Figure A.6: Observed SED (blue dots), MMDC best-fit models (red line) and random MCMC samples (gray lines) of PKS 0548-322, fixing $\gamma_{\min} = 10^3$. A.6a: BeppoSAX X-ray data. A.6b: mean Swift-XRT X-ray data.

B

Machine Learning

B.1 CODE SNIPPETS

```
1 import numpy as np
2 from sklearn.ensemble import HistGradientBoostingRegressor,
  RandomForestRegressor, GradientBoostingRegressor
3
4 models = {
5     "Random Forest": RandomForestRegressor(
6         random_state=23,),
7     "Gradient Boosting": GradientBoostingRegressor(
8         random_state=23,),
9     "Hist Gradient Boosting": HistGradientBoostingRegressor(
10        random_state=23, early_stopping=True),
11 }
12 param_grids = {
13     "Random Forest": {"n_estimators": [10, 20, 50, 100, 200, 500,
14     1000], "max_depth": [None, 3, 5, 10]},
15     "Gradient Boosting": {"n_estimators": [10, 20, 50, 100, 200, 500,
16     1000], "validation_fraction": [0.2], "n_iter_no_change": [10], "
17     max_depth": [None, 3, 5, 10]},
18     "Hist Gradient Boosting": {"max_iter": [10, 20, 50, 100, 200, 500,
19     1000], "validation_fraction": [0.2], "n_iter_no_change": [10], "
20     max_depth": [None, 3, 5, 10]},
21 }
```

Code B.1: Dictionaries of implemented models and candidate hyperparameters.

B.1. CODE SNIPPETS

```
1 import numpy as np
2 import pandas as pd
3 from sklearn.ensemble import HistGradientBoostingRegressor,
  RandomForestRegressor, GradientBoostingRegressor
4 from sklearn.model_selection import KFold, cross_val_score,
  GridSearchCV
5 from sklearn.metrics import mean_squared_error, r2_score, make_scorer
6
7 folds = [5, 10, 20]
8
9 results = []
10 best_models = {}
11 best_params = {}
12 best_kfold = {}
13
14 scoring = {
15     'MSE': make_scorer(mean_squared_error, greater_is_better = False),
16     'R2': make_scorer(r2_score),
17 }
18
19 for name, model in models.items():
20     best_score = np.inf
21     best_kf = None
22     best_params = None
23     best_model = None
24
25     for n_folds in folds:
26         kf=KFold(n_splits=n_folds, shuffle=True, random_state=23)
27         print(f"Number of folds: {n_folds}")
28         grid_search = GridSearchCV(
29             estimator=model,
30             param_grid=param_grids[name],
31             return_train_score=True,
32             cv=kf,
33             n_jobs=N_CORES, #parallelize job on all Bigmama cores
34             scoring=scoring,
35             refit='MSE'
36         ).fit(X, y)
37
38         # Store the cross-validation results for this fold
39         result = {
40             "model": name,
41             "n_folds": n_folds,
42             "cv_results": pd.DataFrame(grid_search.cv_results_)
43         }
44         results.append(result)
```



```
45
46     # Perform cross-validation
47     mse_scores = cross_val_score(grid_search.best_estimator_, X, y,
48     cv=kf, scoring=scoring['MSE'])
49     rmse_scores = np.sqrt(-mse_scores)
50     mean_rmse = np.mean(rmse_scores)
51
52     #Track the best model based on MSE
53     if mean_rmse < best_score:
54         best_score = mean_rmse
55         best_kf = kf
56         best_params = grid_search.best_params_
57         best_model = grid_search.best_estimator_
58     r2_scores = cross_val_score(grid_search.best_estimator_, X, y,
59     cv=kf, scoring=scoring['R2'])
60
61     best_models[name] = best_model
62     best_params[name] = best_params
63     best_kfold[name] = best_kf
```

Code B.2: Model cross-validation and training loops.

B.2. PLOTS

B.2 PLOTS

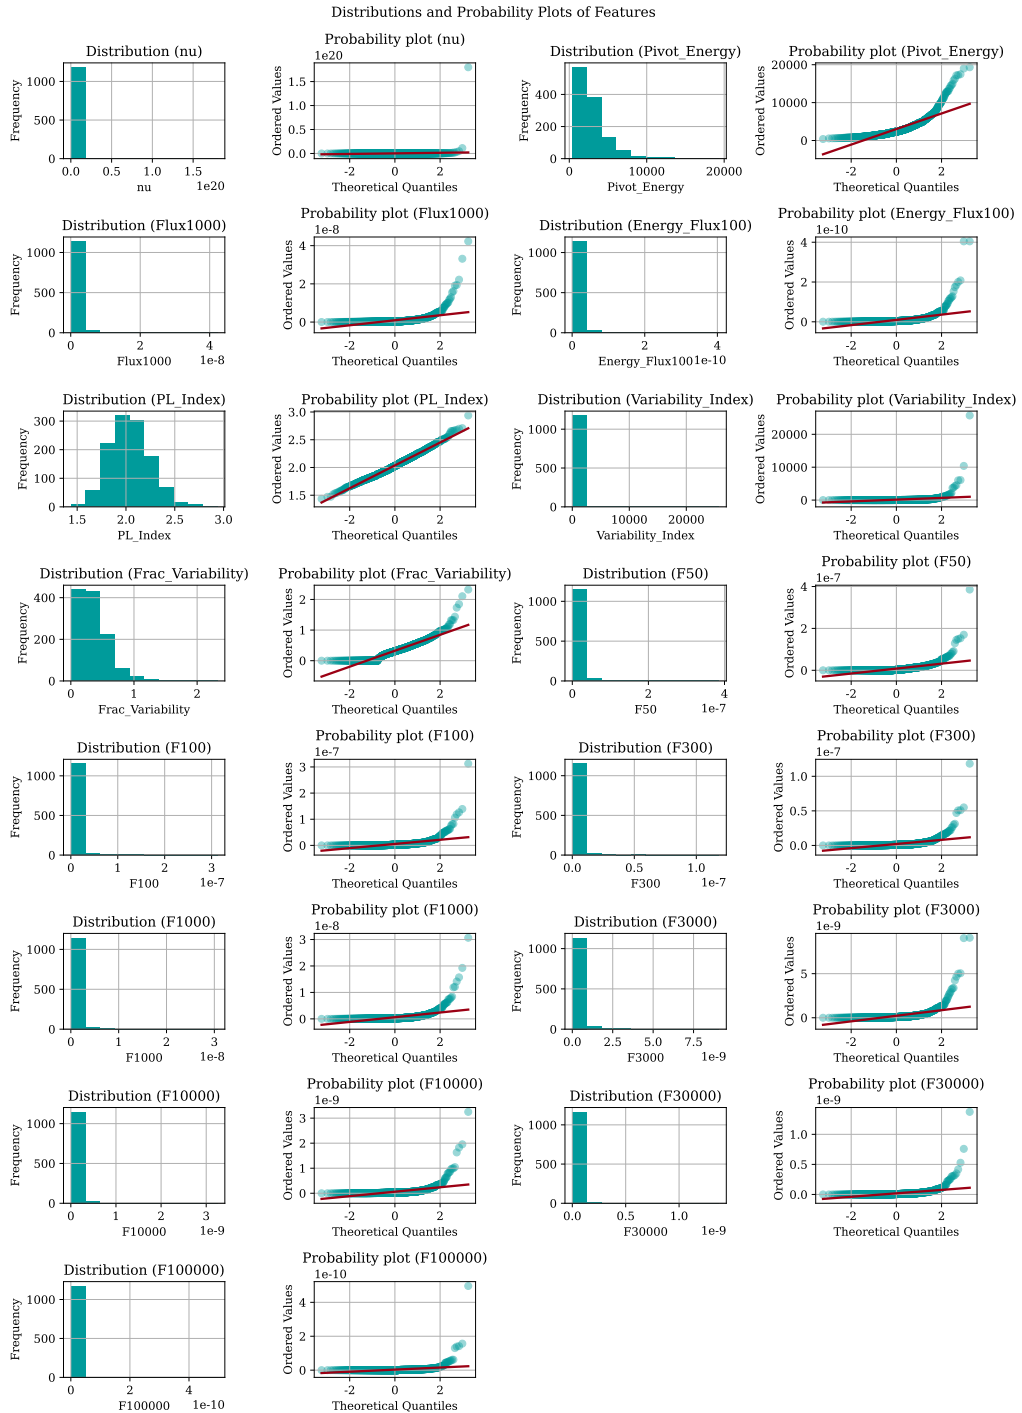


Figure B.1: Frequency histograms and probability distributions of the features and target. In the probability distribution plots, the values and distribution quantiles of the experimental data points (teal) are compared to the behavior of a Gaussian distribution (dark red).




Figure B.2: Scatter plot of predictions against actual values of the target, \log_{ν} [Hz]. Random forest fit on a transformed dataset (logarithms + PowerTransformer).

B.2. PLOTS

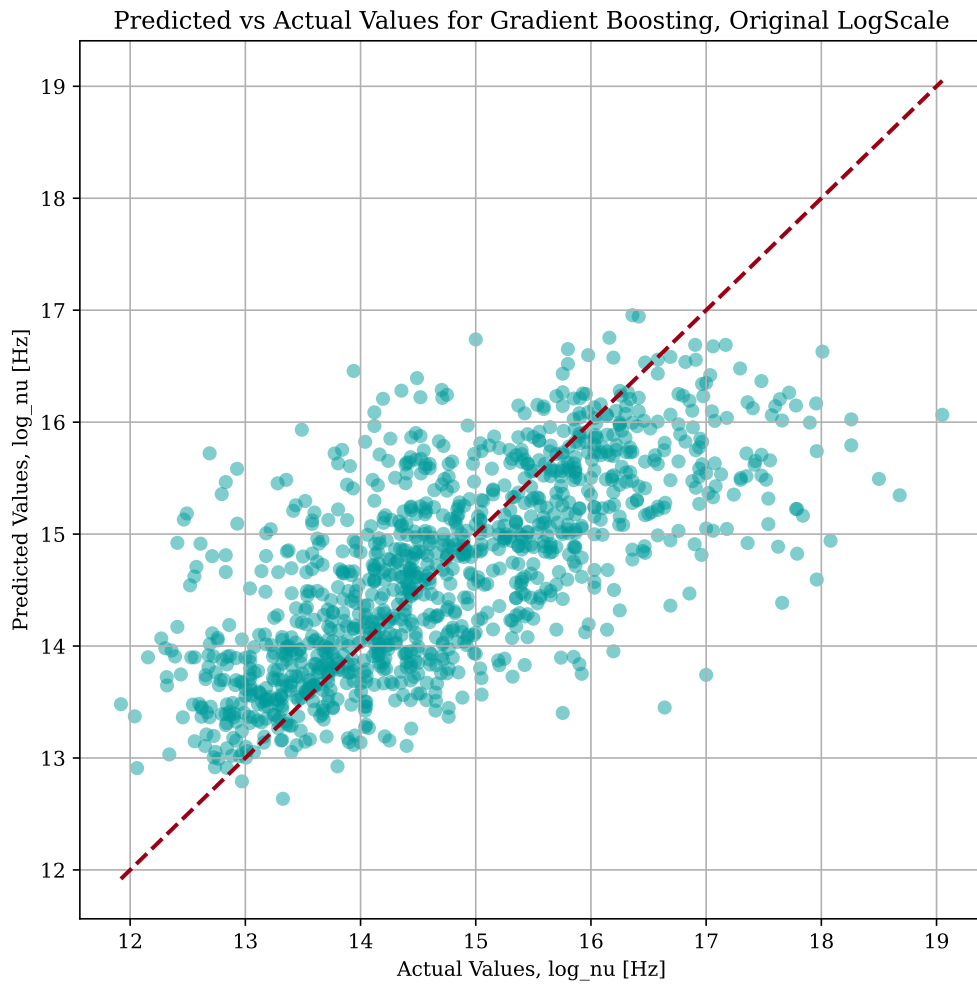


Figure B.3: Scatter plot of predictions against actual values of the target, \log_{ν} [Hz]. Gradient boosting fit on a transformed dataset (logarithms + PowerTransformer).




Figure B.4: Frequency histograms of predictions against actual values of the target, \log_{ν} [Hz]. Random forest fit on a transformed dataset (logarithms + PowerTransformer).

B.2. PLOTS




Figure B.5: Frequency histograms of predictions against actual values of the target, \log_{ν} [Hz]. Gradient boosting fit on a transformed dataset (logarithms + PowerTransformer).