



UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Fisica e Astronomia “Galileo Galilei”

Master degree in Astrophysics and Cosmology

Master thesis

Modeling and Retrieval of Exoplanetary Atmospheres

Using Autoencoders

Relatore

Prof. Zingales Tiziano

Laureando

Riccardo Simonetto

Accademic year 2024/2025

Contents

1	Introduction	2
1.1	Context of the research	2
1.2	Exoplanetary atmospheres	2
1.3	Machine learning	3
1.4	Work outline	4
2	Transmission spectroscopy	6
2.1	Theoretical basis	8
2.1.1	Atomic and Molecular Absorption	8
2.1.2	Scattering	11
2.1.3	Radiative Transfer Equation	12
2.1.4	Atmospheric Properties to Observables	15
2.2	Hot Jupiters	16
2.2.1	Definition and Formation Mechanism	16
2.2.2	Atmospheric Properties	18
2.2.3	Transmission Spectroscopy of Hot Jupiters	18
2.3	Instruments and analysis	19
2.3.1	Hubble Space Telescope - WFC3	19
2.3.2	JWST - NIRSpec	20
2.3.3	Future Missions and Observatories	20
2.3.4	Data Analysis	20
3	Machine Learning in Astrophysics	22
3.1	Machine Learning Basics	22
3.1.1	Deep Feedforward Neural Networks	23
3.1.2	Convolutional Neural Networks	28
3.1.3	Strengths and Limitations of Machine Learning Models	30
3.2	Autoencoders	31
3.3	Application to astrophysical problems	32
3.3.1	ExoGAN	33
4	Developing the Autoencoder	36
4.1	Dataset structure and Preprocessing	36
4.2	Training	38
4.3	Best models	39

4.3.1	AURA	39
4.3.2	AURA-Small	40
4.3.3	AURA-BCE-Loss	41
4.4	Models Summary and Comparison	42
5	Retrieval and Inpainting algorithm	44
5.1	Description of the Algorithm	44
5.2	Retrieval and Inpainting using AURA model	45
5.2.1	Retrieval	45
5.2.2	Inpainting	46
5.3	Case study	48
5.3.1	Inpainting test with synthetic noise	49
5.3.2	WASP-39 b case study	50
5.4	Discussion of the results	54
6	Conclusions	57
	Appendix	58
A	Derivation of the Atmospheric Scale Height and Its Impact on Transit Depth	59
B	Variational Autoencoders	62
C	Distribution of the retrieved parameters of the test example 3992	63
D	Posterior distributions of WASP-39 b	64

Abstract

The study of exoplanetary atmospheres is critical for understanding the formation, evolution, and potential habitability of planets beyond our Solar System. Transmission spectroscopy, especially with advanced instruments like the James Webb Space Telescope (JWST) and the upcoming ARIEL mission, has become a leading technique for probing these atmospheres. In this work, I develop a machine learning framework based on convolutional autoencoders to retrieve key atmospheric and planetary parameters—such as the mixing ratios of CH₄, CO₂, CO, and H₂O, alongside the planet’s mass, radius, and isothermal temperature—from exoplanet transmission spectra. The model also reconstructs the full transmission spectrum over a wide wavelength range (0.3 to 50 μm), enabling an inpainting approach to recover missing spectral regions caused by instrumental limitations. This method offers a fast and efficient alternative to traditional Bayesian retrieval techniques by capturing complex nonlinear relationships within spectral data. While the approach demonstrates strong performance on synthetic data, particularly when broad spectral coverage is available, its accuracy decreases when dealing with limited spectral ranges or noisy observations. These challenges highlight the need for further refinements, such as including realistic noise in the training data and improved model architectures. Despite these limitations, the proposed framework provides a promising groundwork to accelerate atmospheric characterization and enhance our understanding of the diversity and formation of exoplanetary atmospheres.

Chapter 1

Introduction

1.1 Context of the research

In recent years, the use of Artificial Intelligence (AI), and in particular Machine Learning (ML) techniques in astrophysical studies, has significantly improved. These techniques, due to their strong pattern recognition capabilities, have proven to be extremely valuable across a wide range of applications: from galaxy classification (Tadaki et al., 2020) to exoplanet detection and characterization of their atmospheres (Zingales, Waldmann, 2018). In many cases, ML models, once trained, offer substantial advantages over traditional algorithms, especially when working with large and complex datasets. Once trained, they can achieve comparable accuracy while being significantly less computationally intensive. A clear example of this ability is presented in Zingales, Waldmann (2018), where, thanks to a Deep Convolutional Generative Adversarial Network (DCGAN)¹, they obtained a 300-fold speed-up in exoplanetary atmospheric retrieval compared to a standard Bayesian retrieval, without compromising accuracy.

1.2 Exoplanetary atmospheres

The study of exoplanets is a rapidly evolving field. Today we now more than 6000 exoplanets discovered, mostly, by transit and radial velocity methods, thanks to numerous ground-based observation programs and space missions. Some of these programs focus on detecting new exoplanets, i.e. Kepler/K2 (Borucki et al., 2010; Howell et al., 2014), HATS (Bakos, 2018), WASP (Pollacco et al., 2006), while thanks to the PLATO space mission (Rauer et al., 2024), scheduled for launch in 2026, we will discover more than 10 thousands new exoplanets. Others aim to characterize planetary properties such as mass (i.e. HARPS (Pepe et al., 2000), HARPS-N (Cosentino et al., 2012), Espresso (Pepe et al., 2021), HIRES (Marconi et al., 2021)), radius, orbital elements (i.e. CHEOPS, (Benz et al., 2020)). Other missions are dedicated to the study of the chemical composition and physical structure of their atmospheres, i.e. JWST (Gardner et al., 2023) and Ariel (Tinetti et al., 2021), scheduled for launch in 2029.

Several observational techniques have been developed to probe exoplanetary atmospheres. One method is phase curve analysis, which involves measuring the change in brightness of the planet–star system over time, during the planetary orbital motion. Once stellar variability is removed, variations in the detected signals can

¹GANs are a type of neural network architecture composed of two competing networks: a generator and a discriminator. These networks are trained simultaneously: the generator tries to produce realistic data, while the discriminator evaluates their authenticity. At the end of the training, the generator is able to create new realistic data.

be explained by the planet's changing illumination and thermal emission throughout its orbit. This provides insights into properties such as albedo, cloud coverage, and in some cases even surface features. Another technique is direct imaging, which attempts to isolate the light emitted by the planet from that of its host star. This approach allows for emission and reflection spectroscopy of the planet's atmosphere, but it is currently limited to young gas giants with large semi-major axes, orbiting stars close to the Sun, due to the technical difficulty of separating the planet's faint flux from the much brighter flux of the star.

However, the most widely used and successful technique for atmospheric characterization today is transmission spectroscopy. This method analyzes the star light that filters through a planet's atmosphere during a transit event. As the light passes through the atmosphere, specific wavelengths are absorbed by different molecular species, producing characteristic features in the observed spectrum. These absorption features allow to identify the presence of molecules and to infer additional information about atmospheric temperature profiles, cloud layers, and overall structure. The James Webb Space Telescope (JWST, [Gardner et al. 2023](#)) has marked a major advancement in this field. Thanks to its high sensitivity and broad spectral coverage, JWST is the most technologically advanced space mission to perform transmission spectroscopy, allowing for detailed atmospheric studies of both gas giants and smaller, potentially rocky, exoplanets.

Characterizing exoplanetary atmospheres is crucial for understanding a planet's formation history, climate, and evolution. Atmospheric studies allow to determine the chemical composition, temperature distribution, cloud coverage, and potential surface conditions of a planet. In the case of rocky exoplanets, such investigations also play a fundamental role in the search for biosignatures: molecular species that could indicate biological activity ([Schwieterman, Leung, 2024](#)). The most important space mission, that is expected to revolutionize this field, is Ariel ([Tinetti et al., 2021](#)), programmed to be launched in 2029. During its 4-year mission, Ariel will study exoplanets' composition, their formation history, and their evolution, by surveying a diverse sample of about 1000 extrasolar planets, simultaneously in visible and infrared wavelengths. It is the first mission dedicated to measuring the chemical composition and thermal structures of hundreds of transiting exoplanets. It will enable planetary science studies far beyond the Solar System.

Ariel, together with the JWST and other upcoming missions like PLATO ([Rauer et al., 2024](#)) and ground-based observatories like the Extremely Large Telescope (ELT, [Padovani, Cirasuolo 2023](#)), will generate a vast amount of data, significantly advancing research and expanding our understanding in this field. In this context, it is crucial to develop tools that can efficiently and accurately analyze this data.

1.3 Machine learning

Artificial Intelligence (AI) refers to the development of machines capable of exhibiting intelligent behavior, simulating human cognitive functions such as problem-solving and decision-making². Machine Learning (ML) is a subdomain of AI focused on algorithms that learn directly from data, identifying patterns and making predictions or decisions without being explicitly programmed for each specific task. The growing popularity of AI and ML can be attributed to the continuous reduction in computational power costs, coupled with significant improvements in performance, particularly through the use of GPUs. Additionally, the availability of high-level libraries in recent years, such as TensorFlow ([Abadi et al., 2015](#)), has simplified the development of these algorithms. TensorFlow, for instance, is a Python package that enables users to easily build and train Machine Learning algorithms with remarkable flexibility.

²Definition taken from the [IBM website](#)

The fundamental strength of these algorithms lies in their ability to recognize patterns, allowing them to capture subtle, complex, and often non-linear correlations within vast, high-dimensional datasets. However, a significant challenge, especially with complex models like deep neural networks, is the difficulty in understanding exactly how they process and interpret the data. This is referred to as the “black box” problem (Hassija et al., 2023), which is a major barrier in fields that require interpretability. Furthermore, the reliability of ML models is critically dependent on the quality of the data they learn from. If the training data itself contains biases due to limitations in data collection methods, incomplete sampling, or underrepresented classes, the model will inevitably learn and potentially amplify these biases during training. This can lead to models that perform poorly or inaccurately when applied to new data. Consequently, it is not only essential to have large datasets, which enable complex models to generalize patterns rather than just memorize training examples (overfitting), but also to ensure that the data accurately represent the phenomena we aim to study. Insufficient data, poor data quality, or biased data undermine a model’s accuracy and its ability to generalize when applied to new data, regardless of the algorithm used.

Despite this, their efficiency in feature extraction makes ML exceptionally useful in fields like astrophysics, which deals with enormous datasets from sky surveys. In astrophysics, ML can be used for tasks such as classifying galaxies and stars (Zeraatgari et al., 2023), detecting and analyzing faint or transient events like supernovae (Villar et al., 2020) or exoplanet transits within noisy data (Chintarungruangchai, Jiang, 2019), identifying gravitational wave signals (George, Huerta, 2018), and finding anomalies that could lead to new discoveries (Etsebeth et al., 2024).

In this thesis, I focus on the development of a Machine Learning model designed to efficiently retrieve atmospheric parameters from exoplanet transmission spectra. The target parameters include the mixing ratios of CH_4 , CO_2 , CO , and H_2O , as well as the planet’s mass (M_p), radius (R_p), and isothermal temperature (T_p). To achieve this, we utilize a specific type of neural network architecture known as an autoencoder. An autoencoder (AE) is an unsupervised learning model composed of an encoder that compresses the input data into a lower-dimensional latent representation (a compressed ‘encoding’), and a decoder that attempts to reconstruct the original input from this compressed representation. Autoencoders are particularly effective for tasks such as dimensionality reduction, feature learning, data denoising, and anomaly detection. These capabilities are essential for correlating spectral features to chemical abundances and planetary parameters. Additionally, as a byproduct of the retrieval process, the model is trained to reconstruct the entire planetary spectrum. This reconstruction capability is leveraged in an inpainting³ framework to recover missing spectral regions that may not have been observed due to filters and instrumental restricted wavelength coverage.

1.4 Work outline

This work is structured as follows. Chapter 2 introduces the theoretical background on transmission spectroscopy, focusing on hot Jupiters, their properties, formation, and atmospheric characteristics, as they are the primary targets within my thesis. Chapter 3 covers the theoretical foundations of machine learning, its applications in astrophysics, particularly in exoatmospheres, and introduces autoencoders, key tools for this study. In Chapter 4, I detail the process of training our convolutional autoencoder, including dataset preprocessing, as well as the description of the best-performing model. Chapter 5 discusses the implementation of the retrieval and inpainting algorithm and its results. In Chapter 6, I summarize the findings and suggest future research

³Inpainting is a technique used to reconstruct missing or corrupted parts of data (such as images) by inferring the missing information based on the surrounding context and its training.

directions.

Chapter 2

Transmission spectroscopy

Among the 7442 candidate and confirmed exoplanets discovered, 4532 are discovered with the transit method¹. The transit method consists in observing the decrease in light of a star caused by the passage of a planet in front of the stellar disk. This detection method allows to obtain different planetary parameters such as the ratio between the radius of the planet and the radius of the star (R_p/R_s), the inclination of the orbit (i), and the period (P) if more than one transit is observed. Assuming the value of R_s (thanks to theoretical models constrained on spectroscopic observations of the host star), it is possible to measure the planetary radius directly from the transit depth. Combining the radius R_p with the mass (M_p), i.e. obtained with radial velocity² observations of the system, we can compute the planet's mean density and, therefore, infer its bulk composition. This is the main goal of the upcoming PLATO space mission (Rauer et al., 2024). In particular, it will focus on the detection and characterization of Earth-like exoplanets in the habitable zone³ of their stars.

The extreme importance of this method lies in its simplicity since it is based only on photometric observations, and in its ability to estimate R_p . In particular:

$$\frac{\Delta F}{F} = \left(\frac{R_p}{R_s} \right)^2 \quad (2.1)$$

where ΔF is the measured flux decrease caused by the planet passage, and F is the out-of-transit star flux. $\Delta F/F$ is also referred to as transit depth and indicated with δ . The shape of a transit light curve depends on several parameters such as the spectral type of the host star, the limb darkening coefficients⁴, the observation wavelength, the radius of the planets and the inclination of its orbit. The light curve of a giant planet with $R_p = 0.97R_J$ at separation $a = 0.7au$, orbiting a solar-type star (G2V spectral type) is shown in Figure 2.1. The transit depth is the maximum dimming in flux, measured at the central time of the transit.

The transit light curve is not constant but varies with the observation wavelength. This fundamental concept

¹These numbers are updated to April 3, 2025. Obtained from the website [EXOplanet.eu](https://exoplanet.eu).

²Radial velocity detection measures periodic Doppler shifts in a star's spectral lines, caused by the reflex motion induced by an orbiting exoplanet. Analysis of these shifts allows determination of the planet's minimum mass ($M_p \sin(i)$), orbital period (P), and eccentricity (e).

³The habitable zone (HZ) is the region around a star where conditions are suitable for liquid water to exist on a planet's surface, assuming an atmosphere capable of supporting it. This zone depends primarily on stellar flux. The HZ is a key criterion in assessing a planet's potential habitability.

⁴Limb darkening is the optical effect seen in stars where the center of the stellar disk appears brighter than the edge (limb), due to variations in temperature and optical depth along different lines of sight. Limb darkening coefficients are parameters in empirical or theoretical laws that describe this brightness variation across the disk and are used in stellar modeling and transit light curve analyses.




Figure 2.1: Simulated normalized transit light curve obtained using [BATMAN](#) python package ([Kreidberg, 2015](#)). The light curve is obtained simulating a transiting giant planet with $R_p = 0.97R_J$ at separation $a = 0.7au$, orbiting a solar-type star (G2V spectral type).

is the basis of transmission spectroscopy, a technique that involves observing a transit event simultaneously across a range of wavelengths using a spectrograph. By doing so, it is possible to obtain a distinct transit light curve for each wavelength bin, allowing the measurement of the transit depth as a function of wavelength.

This wavelength-dependent variation in transit depth is directly linked to the atmospheric composition of the transiting planet. At wavelengths where specific atmospheric gases absorb more stellar radiation, the planet appears larger due to the increased atmospheric opacity, resulting in a deeper transit (see Figure 2.2). By analyzing these subtle changes in the transit depth, it is possible to infer key atmospheric properties such as chemical composition, temperature structure, and the presence of clouds. Transmission spectroscopy has therefore become one of the most powerful techniques for the characterization of exoplanetary atmospheres since it enables the detection of molecular features, and provides critical insight into atmospheric dynamics and structure.




Figure 2.2: Schematic representation of the wavelength dependence of the transit depth. In this image we can clearly observed that at wavelengths where specific atmospheric gases absorb more stellar radiation, the planet appears larger due to the increased atmospheric opacity, resulting in a deeper transit. Image taken from [NASA Visualization Studio](#) about exoatmospheres.

In this chapter, I present the theoretical foundations of this technique, outline its primary scientific targets, and describe the key instruments and data reduction methods used in its application.

2.1 Theoretical basis

The interaction between stellar radiation and a planetary atmosphere during transit is governed by fundamental physical processes such as molecular absorption, scattering, and radiative transfer. This section introduces the microphysical mechanisms responsible for shaping transmission spectra, and then connects them to the formalism used to model the observed signal. A more in-depth treatment of the concepts presented in this section can be found in [Liou \(2002\)](#).

2.1.1 Atomic and Molecular Absorption

The key process underlying transmission spectroscopy is the absorption of stellar radiation by atoms and molecules in the planetary atmosphere.

Atomic Transitions

According to Bohr's model of the hydrogen atom (1913), electrons can occupy only specific, quantized energy levels around the nucleus. The level with the lowest energy is the closest to the nucleus and is known as the ground state, while all higher-energy levels are referred to as excited states. While in these states, electrons do not emit radiation. Emission occurs only when the electron transitions from one state with energy E_k to a state with lower energy E_j , releasing a photon with energy equal to the difference between the two levels. We can write a simple relation between the energy of the levels and the energy of the electromagnetic wave emitted by the transition:

$$E_k - E_j = h\nu \quad (2.2)$$

where h is the Planck's constant, and ν is the frequency of the electromagnetic wave, so $h\nu$ is the energy of the emitted radiation. Similarly, if an atom is placed in a radiation field, an electron may absorb a photon whose energy ($h\nu$) exactly matches the energy difference between two atomic levels. This absorption can cause the electron to jump from a lower energy level E_k to a higher energy level E_j , according to the relation:

$$E_k + h\nu = E_j \quad (2.3)$$

These two effects combined are responsible for the atomic emission and absorption spectra respectively. In particular, each quantum jump between fixed energy levels results in the emission or absorption of a characteristic frequency or wavelength. These quanta appear in the spectrum as emission or absorption lines.

Molecular Transitions

Like atoms, molecules can undergo electronic transitions. These transitions typically involve high energies, as molecular electrons are influenced by the potential of multiple nuclei.

In addition to electronic energy, molecules can store energy in two other quantized forms: rotational and vibrational energy. Transitions between these discrete energy levels are responsible for spectral line formation

known as rotational-vibrational transitions. It is worth noting that for a transition to be radiatively active, the molecule must have a permanent electric dipole moment. Symmetric molecules like N_2 and O_2 are inactive, while polar molecules like H_2O and O_3 are active.

Rotational energy arises from the molecule's rotation about an axis through its center of mass. From the solution of the time-independent Schrödinger equation for a rigid rotor, the rotational energy levels are quantized and given by:

$$E_J = BhcJ(J + 1) \quad (2.4)$$

where h is Planck's constant, c is the speed of light, J is the rotational quantum number, and B is the rotational constant which depends on the moment of inertia I of the molecule.

The permitted (most probable) transitions are described by the selection rule for rotational transitions:

$$\Delta J = \pm 1 \quad (2.5)$$

This leads to equally spaced spectral lines (in frequency space) in the microwave or far-infrared regions, separated by $2B$ in wavenumber⁵.

Molecular vibrational energy arises from the fact that the atoms within a molecule are bound by interatomic forces, often approximated as springs, that allow them to oscillate about their equilibrium positions. Vibrational energy levels for a harmonic oscillator are:

$$E_v = hk \left(v_k + \frac{1}{2} \right) \quad (2.6)$$

where h is Planck's constant, v_k is the vibrational quantum number, and k is the vibrational constant.

Selection rules are given by:

$$\Delta v_k = \pm 1 \quad (2.7)$$

Vibrational transitions typically occur in the mid-infrared due to their higher energy compared to rotational transitions. Because of this, vibrational transitions never occur alone, but they are usually accompanied by simultaneous rotational transitions. This coupling results in a group of closely spaced spectral lines, known as rotational-vibrational bands, which appear in the infrared part of the spectrum.

The total energy of this type of transition is the sum of both contributions:

$$E_{v,J} = BhcJ(J + 1) + hk \left(v_k + \frac{1}{2} \right) \quad (2.8)$$

⁵The wave number is an alternative way to express the wavelength or the frequency of a line. It is defined as $\tilde{\nu} = 1/\lambda = \nu/c$ (cm^{-1}). It is often used in the infrared spectral range and in molecular spectroscopy.

In this way, for example, a vibrational transition to an upper vibrational level ($\Delta v_k = +1$) produces different series of lines based on the change in the rotational quantum number ΔJ :

- **P-branch:** $\Delta J = -1$
- **R-branch:** $\Delta J = +1$
- **Q-branch:** $\Delta J = 0$

Figure 2.3 presents a schematic representation of a rotational-vibrational transition.




Figure 2.3: Simultaneous vibrational and rotational transitions: $J = -1$ produces the P-branch and $J = +1$ generates the R-branch, while $J = 0$ corresponds to the Q-branch. Image taken from [Liou \(2002\)](#).

Line Broadening Mechanisms

Monochromatic emission or absorption is practically never observed. Energy levels during energy transitions are normally changed slightly due to both external influences on atoms and molecules, and to intrinsic mechanisms. As a consequence, radiation emitted during repeated energy transitions is non-monochromatic, and spectral lines of finite widths are observed. I now describe the three main broadening processes: natural, collisional, and Doppler broadening.

Natural broadening derives directly from the Heisenberg uncertainty principle. Specifically, an excited state with energy E and a finite lifetime Δt cannot have a perfectly defined energy. The uncertainty in energy and time is related by:

$$\Delta E \Delta t \sim \frac{h}{2\pi} \quad (2.9)$$

Using the relation $E = h\nu$, we can express the resulting frequency uncertainty of a transition between this excited level and the ground state (set as the reference level for the energy) as:

$$\Delta\nu \sim \frac{1}{2\pi\Delta t} \quad (2.10)$$

Therefore, even in the absence of external effects, emission and absorption lines are not perfectly monochromatic, but have a finite width due to this fundamental quantum limitation. The natural broadening causes a very narrow Lorentzian profile that is usually not detected in astrophysical spectra due to the limited resolution of the instruments and the much stronger broadening caused by external processes.

Pressure broadening, also known as collisional broadening, is one of the main broadening mechanisms in high density environments. It depends directly on the gas density (or pressure), on the cross section of the collisional process, and on the thermal velocities of the atoms or molecules. The line profile is a Lorentzian function described by:

$$k_{\tilde{\nu}} = \frac{S}{\pi} \frac{\alpha}{(\tilde{\nu} - \tilde{\nu}_0)^2 + \alpha^2} \quad (2.11)$$

where $k_{\tilde{\nu}}$ is the absorption coefficient, $\tilde{\nu}_0$ is the wavenumber of the ideal monochromatic line, α is the full width at half maximum (FWHM) of the line, and S is the line strength or intensity.

The other main broadening mechanism is the Doppler broadening, which arises from the thermal motion of emitting or absorbing atoms and molecules in a gas. If these particles have a non-zero velocity component along the observer's line of sight, the observed frequency ν will be shifted relative to the rest-frame frequency ν_0 by a factor $1 + v/c$, where v is the velocity component along the line of sight. If the transitional state is in thermodynamic equilibrium, the particles of the gas follow the Maxwell-Boltzmann velocity distribution, from which it is possible to derive that the line profile is a Gaussian described by:

$$k_{\nu} = \frac{S}{\alpha_D \sqrt{\pi}} \exp \left[- \left(\frac{\nu - \nu_0}{\alpha_D} \right)^2 \right] \quad (2.12)$$

where k_{ν} is the absorption coefficient, K is the Boltzmann constant, T is the temperature of the gas, and $\alpha_D = \nu_0 (2kT/mc^2)^{1/2}$ is a measure of the Doppler width of the line, with m being the mass of a single emitter. Doppler effect is the main broadening mechanism in low density environments.

In real astrophysical scenarios, pressure and Doppler broadening coexist, and the actual line profile is a convolution between the Lorentzian ($L(\nu)$) and the Gaussian profile ($G(\nu)$).

$$V(\nu) = G(\nu) * L(\nu) = \int_{-\infty}^{+\infty} G(\nu') L(\nu - \nu') d\nu' \quad (2.13)$$

The result is called Voigt profile ($V(\nu)$) and it is characterized by the fact that the Doppler broadening dominates the central part of the line, while the pressure broadening dominates the wings.

Since these broadening mechanisms are directly linked to both the physical properties of the gas and the shape of the spectral line, analyzing the line profile allows us to retrieve valuable information about the gas's physical conditions, such as temperature and pressure.

2.1.2 Scattering

Most of the light we observe in the natural world is not seen through direct emission from its source, but rather through the process of scattering. Scattering occurs when particles in the path of electromagnetic radiation

deviate the energy in different directions. This process is fundamental to many atmospheric optical phenomena, including the blue color of the sky.

In the atmosphere, scattering is caused by particles ranging in size from gas molecules ($\sim 10^{-4} \mu\text{m}$), to aerosols (ranging from $\sim 1 \mu\text{m}$ to several tenths of μm). The type and efficiency of scattering depend strongly on the size parameter x , defined as:

$$x = \frac{2\pi a}{\lambda} \quad (2.14)$$

where a is the particle radius and λ is the wavelength of the incident light. Depending on the value of x , different scattering regimes apply:

- **Rayleigh Scattering** ($x \ll 1$): Occurs when particles are much smaller than the wavelength of light (e.g., molecules). It is dominant in the visible range. The scattered intensity follows a λ^{-4} dependence:

$$I_\nu \propto \frac{1}{\lambda^4} \quad (2.15)$$

This means that shorter wavelengths (e.g., blue light) are scattered much more efficiently than longer wavelengths (e.g., red light), explaining why the sky appears blue.

- **Mie Scattering** ($x \sim 1$ or $x > 1$): Occurs when particles are comparable to or larger than the wavelength of light. It is typically caused by aerosols, cloud droplets, and dust. Mie scattering tends to be less wavelength-dependent and is more strongly forward-scattered. It contributes to the white appearance of clouds and the bright glare around the Sun.

In a scattering medium, light may interact with one particle (single scattering) or undergo multiple interactions with several particles (multiple scattering). While single scattering dominates in clear atmospheres, multiple scattering becomes significant in optically thick regions such as clouds or dense haze layers. In multiple scattering, light initially scattered out of a line of sight can be redirected back into it by further interactions, affecting the net radiative transfer.

Scattering, especially from high-altitude aerosols, hazes, or clouds, can significantly affect the observed transmission spectrum of an exoplanet by masking or muting molecular absorption features. When scattering dominates, it introduces a nearly wavelength-independent opacity or a sloped continuum, reducing the contrast of spectral lines. This may obscure key diagnostic features, making it more difficult to retrieve accurate atmospheric compositions. In extreme cases, a flat or featureless spectrum may indicate the presence of optically thick clouds or dense haze layers.

2.1.3 Radiative Transfer Equation

I now describe how the intensity of a beam of light is affected by its interaction with a medium. As discussed previously, absorption features depend on temperature, pressure, and molecular composition of the medium. The cumulative absorption along the line of sight determines how much of the incident radiation is absorbed at each wavelength. In addition to absorption, scattering processes, such as Rayleigh scattering by molecules and Mie scattering by aerosols, also contribute to the attenuation of the beam, further reducing its observed intensity.

In this context, the term cross section is used to denote the amount of energy removed from the original beam

by the different particles composing the medium. The total extinction cross section is defined as the sum of the scattering and absorption cross sections and has the unit of an area⁶ (cm^2). In some cases, the cross section is expressed in unit mass (cm^2g^{-1}), resulting in the mass extinction cross section. When this quantity is multiplied by the mass density of the medium, it results in the absorption coefficient, whose units are the inverse of a length (cm^{-1}). The absorption coefficient is typically denoted by k_ν . It quantifies the rate at which the intensity of radiation is diminished due to absorption as it propagates through the medium.

Understanding how light is attenuated by absorption and scattering is essential in transmission spectroscopy, as the observed transit depth at each wavelength reflects how much stellar light is filtered through the planet's atmosphere. Central to this interpretation is the Radiative Transfer Equation, which provides the mathematical framework for modeling how radiation intensity changes as it passes through an absorbing and scattering medium. The absorption coefficient k_ν , introduced above, plays a key role in this equation by determining the amount of attenuation per unit length.

To express this quantitatively, consider a pencil of monochromatic radiation with intensity I_ν traversing an infinitesimal distance ds in the direction toward the observer (taken as the positive direction) through an absorbing medium. The differential change in intensity is given by:

$$dI_\nu = -k_\nu I_\nu ds, \quad (2.16)$$

where k_ν is the absorption coefficient previously defined. It accounts for both absorption and scattering contributions. This equation represents the attenuation of intensity due to extinction.

However, radiation may also gain intensity from emission by the medium itself and from scattering into the beam from other directions. To account for this effect, we introduce the emission coefficient j_ν , such that:

$$dI_\nu = j_\nu ds. \quad (2.17)$$

j_ν has the same physical meaning as the absorption coefficient k_ν .

Combining both extinction and emission, the full radiative transfer equation becomes:

$$\frac{dI_\nu}{ds} = -k_\nu I_\nu + j_\nu \quad (2.18)$$

It is convenient to define the source function J_ν as the ratio of the emission and extinction coefficients:

$$J_\nu = \frac{j_\nu}{k_\nu}, \quad (2.19)$$

J_ν is an intrinsic characteristic, fairly independent from the properties of the radiation source.

The radiative transfer equation simplifies to:

$$-\frac{dI_\nu}{k_\nu ds} = I_\nu - J_\nu. \quad (2.20)$$

⁶The extinction cross section has units of area, analogous to the classical concept of a particle's geometrical cross-sectional area.

Equation 2.20 is the general form of the Radiative Transfer Equation and serves as the foundation for modeling light propagation through a medium.

To solve this equation, we introduce the monochromatic optical depth τ_ν through the definition:

$$d\tau_\nu = -k_\nu ds \quad (2.21)$$

Rewriting the Radiative Transfer Equation in terms of optical depth:

$$\frac{dI_\nu}{d\tau_\nu} = I_\nu - J_\nu \quad (2.22)$$

by rearranging some terms and multiplying both sides by $e^{-\tau_\nu}$ we obtain:

$$e^{-\tau_\nu} \left[\frac{dI_\nu}{d\tau_\nu} - I_\nu \right] = -J_\nu e^{-\tau_\nu} \quad (2.23)$$

that can be rewritten as:

$$\frac{d}{d\tau_\nu} (I_\nu e^{-\tau_\nu}) = -J_\nu e^{-\tau_\nu} \quad (2.24)$$

This first-order differential equation, when integrated between τ_ν^* (the optical depth at the boundary of the medium closest to the radiation source) and $\tau_\nu = 0$ (the optical depth at the boundary of the medium closest to the observer), yields the formal solution:

$$I_\nu(\tau_\nu = 0) = I_{\nu 0} e^{-\tau_\nu^*} + \int_{\tau_\nu^*}^{\tau_\nu=0} -J_\nu e^{-\tau_\nu} d\tau_\nu \quad (2.25)$$

Where $I_{\nu 0}$ is the intensity of the beam of radiation before entering the medium. The first term in Equation 2.25 represents the attenuation of the initial intensity, and the second term accounts for the emitted and scattered radiation accumulated along the path.

To solve the integral in Equation 2.25, it is necessary to know how the source function J_ν varies with optical depth τ_ν . This, in turn, requires knowledge of how the absorption coefficient k_ν and the emission coefficient j_ν depend on τ_ν , which reflects the physical properties of the medium along the line of sight. In most practical cases, this level of information is not available, and therefore the homogeneous medium approximation is often adopted.

In the special case of a non-scattering medium in local thermodynamic equilibrium (LTE), the source function J_ν is equal to the Planck function⁷ $B_\nu(T)$. Thus, Equation 2.20 can be rewritten as:

⁷The Planck function describes the amount of energy emitted per unit area, per unit time, per unit solid angle, and per unit frequency by an ideal blackbody at temperature T . In frequency space, it is given by:

$$B_\nu(T) = \frac{2h\nu^3}{c^2} \cdot \frac{1}{e^{\frac{h\nu}{kT}} - 1}$$

where $B_\nu(T)$ is the spectral radiance ($\text{W m}^{-2} \text{Hz}^{-1} \text{sr}^{-1}$), h is Planck's constant, ν is the frequency of the radiation, c is the speed of

$$-\frac{dI_\nu}{k_\nu ds} = I_\nu - B_\nu(T), \quad (2.26)$$

This equation is also known as Schwarzschild's equation.

Equation 2.25 is fundamental in atmospheric and astrophysical radiative transfer, allowing us to model the emission and absorption properties of the medium according to its physical properties. In the context of transmission spectroscopy, solving the Radiative Transfer Equation allows us to relate the observed spectrum to the planet's atmospheric properties such as composition, temperature, and pressure. Without this formalism, it would not be possible to quantitatively retrieve the physical conditions of exoplanetary atmospheres from the measured data.

2.1.4 Atmospheric Properties to Observables

During an exoplanetary transit, a fraction of the light emitted by the star is attenuated by passing through the thin annulus of the planet's atmosphere surrounding its limb. At wavelengths where the planet's atmosphere absorbs more strongly the star's light, the transit depth increases. The wavelength dependence of the transit depth effectively provides a transmission spectrum.

From Equation 2.1, this transit depth wavelength dependence results in subtle variations of the apparent planetary radius $R_{p\lambda}$. This radius corresponds to the altitude at which the atmosphere becomes optically thick⁸ at that specific wavelength.

The wavelength-dependent transit depth also depends on the extent of the atmospheric annulus. In first approximation, the width of the annulus can be taken as $5 \times H$. H denotes the scale height of the atmosphere and it is defined as the altitude at which the pressure has decreased by a factor e with respect to the surface pressure (or the reference pressure). It can be described by:

$$H = \frac{kT}{\mu_m g_p} \quad (2.27)$$

where k is Boltzmann's constant, T is the atmospheric temperature, μ_m is the mean molecular weight, and g_p is the planet's surface gravity. (Seager et al., 2009). A larger scale height implies a more extended atmosphere and, consequently, more prominent spectral features. Following Perryman (2018), the fractional contribution of the atmosphere to the transmission signal is given by:

$$\delta \simeq 5 \times \frac{2R_p H}{R_s^2} \quad (2.28)$$

This expression highlights how the transmission signal depends on the planet's radius R_p , the atmospheric scale height H , and the stellar radius R_s . It provides a useful approximation for estimating the strength of atmospheric features based on the atmosphere's extent. A derivation of Equations 2.27 and 2.28 is provided in Appendix A.

Transmission spectroscopy measurements probe the planet's atmosphere at the terminator⁹, sampling both the

light, k is Boltzmann's constant, T is the absolute temperature of the blackbody.

⁸A medium is said to be optically thick at wavelength λ when the optical depth $\tau_\lambda \gg 1$. This means that no light filters through the medium.

⁹The planetary terminator is the dividing line between the day side and the night side of a planet.

morning and evening limbs. This allows to study potential differences in temperature, composition, and scale height between the two sides. Such differences are expected to be most pronounced in tidally locked hot Jupiters, where the intense stellar irradiation on the permanent day side leads to strong temperature gradients. These gradients can drive high-speed winds that transport photochemically produced species from the hot day side to the cooler night side, potentially resulting in distinct spectral signatures between the two limbs.

Additionally, an asymmetric cloud or haze coverage between the two limbs may arise from the temperature difference. This asymmetry can lead to variations in absorption strength between the two sides. As already mentioned, the presence of clouds and hazes tends to flatten spectral features, as stellar photons are prevented from probing deeper atmospheric layers. This occurs because aerosols scatter incoming light, effectively obscuring molecular absorption signatures originating below the cloud or haze layers.

In conclusion, detailed characterization of exoplanetary atmospheres is a complex challenge due to the intrinsic diversity and physical complexity of planetary environments. Unlike stellar atmospheres of Solar-type stars, which are primarily composed of atoms and ions and can be reasonably described by a small number of parameters (e.g., mass and metallicity), planetary atmospheres are dominated by molecules, clouds, and hazes. Their composition is shaped by a variety of processes, including their formation environment, accretion history, photochemistry driven by intense stellar irradiation, and even geophysical evolution. The lack of spherical symmetry, due to anisotropic stellar irradiation and planetary rotation, introduces further complexity.

Moreover, cloud formation and haze layers are difficult to model, adding additional extinction sources and degeneracies in spectral interpretation. Furthermore, molecular spectra are more complex than atomic ones, with opacity and cross sections archives often incomplete (Perryman, 2018), especially at the high temperatures typical of close-in exoplanets, introducing even more uncertainties in the models.

Despite these challenges, transmission spectroscopy remains a powerful tool for inferring a planet's chemical composition, thermal structure (pressure-temperature profile), and energy transport mechanisms. In particular, it enables the detection of atomic species, such as H, C, O, Ca, Na, and K, in the optical spectral range, while in the infrared, H_2O , CO_2 and CH_4 are among the most prominent molecules observed. To date, most atmospheric detections have been achieved for strongly irradiated hot Jupiters, which are favorable targets due to their large radii, extended atmospheres, and low mean molecular weight, all of which enhance their transmission signals, as can be seen from Equation 2.28 and Equation 2.27. Some limited measurements have also been made for smaller, lower-mass exoplanets like hot Neptunes, although their spectral signatures are fainter and more challenging to interpret.

2.2 Hot Jupiters

2.2.1 Definition and Formation Mechanism

Hot Jupiters are giant planets¹⁰ with a semi-major axis $a < 0.1\text{au}$ or a period $P = 3 - 9d$ (Perryman, 2018). Hot Jupiters with $P < 1d$ are called ultra-short-period hot Jupiters or very hot Jupiters.

Their detection probability (Winn, 2014), using the transit method, is high because of their small separation. The detection probability in the transit method is indeed defined as the ratio between the side area of a cylinder of radius equal to the semimajor-axis a of the planet and height equal to the stellar disk diameter ($2R_s$), and the

¹⁰Giant planets are gas planets with mass bigger than $0.1M_J$. M_J stands for Jupiter masses, and $1M_J = 1.898 \times 10^{27}\text{kg}$.

area of a sphere of radius a :

$$p_{\text{transit}} = \frac{2\pi a \cdot 2R_s}{4\pi a^2} = \frac{R_s}{a} \quad (2.29)$$

However, their occurrence rate is still uncertain with values ranging from 0.4% from Kepler observations to 1.2% of Radial velocity surveys (Beleznay, Kunimoto, 2022). This very low occurrence rate is probably related to their formation and evolution mechanism.

Giant planets are characterized by the presence of a massive gaseous envelope, accounting for most of the planet's total mass. This property imposes constraints on the formation timescale of these planets. This envelope is captured during the early evolutionary stages, from the primordial gas in the protoplanetary disk. Since this gas is dispersed by photoevaporation over a few million years, giant planets must form within that time window (Perryman, 2018).

The most widely accepted formation theory is the Core-Migration-Accretion model (Raymond, Morbidelli, 2022), which unfolds in two main phases: core accretion and atmospheric capture. Initially, dust particles in the disk grow through collisions and electrostatic interactions. Beyond the snow line, where giant planets are expected to form, this process is limited by fragmentation, which prevents particles from growing beyond a few centimeters. However, the streaming instability can overcome this barrier by concentrating particles in high-density regions, allowing them to gravitationally collapse into planetesimals, typically around 100 km in size.

These large bodies, especially near the snow line, rapidly grow via pebble accretion, efficiently capturing smaller particles drifting inward due to gas drag. Once planetesimals reach masses of 1–10 Earth masses, their growth slows as they start opening gaps in the disk. At around 10 Earth masses, the core begins to accrete significant amounts of gas, marking the transition to the giant planet phase.

Once the mass of the envelope matches that of the core, runaway gas accretion begins. In just a few million years, these protoplanets can grow from ~ 10 Earth masses to several Jupiter masses. At this stage, planetary migration becomes relevant. The interaction with the disk generates torques that typically cause inward migration, pushing these young gas giants toward the inner regions of the system.

This migration process can give rise to Hot Jupiters. However, the low occurrence rates suggest that, in most cases, migration is halted due to photo-evaporation of the gas disk (Gratton et al., 2023). Another migration mechanism can be caused by the chaotic interaction with other giant planets in the system. This chaotic interaction, known as Planet-Planet Scattering, often ends with the ejection of one planet from the system, while the other is injected in a close orbit (Teyssandier et al., 2019). After the migration phase, Hot Jupiters settle into close orbits around their host stars, leading to strong tidal interactions. As a result, their orbits become circularized, and they become tidally locked, always showing the same face to the star (Perryman, 2018; Teyssandier et al., 2019). Moreover, the intense stellar irradiation raises their atmospheric temperatures, often causing thermal expansion and resulting in inflated radii. In some cases, this extreme heating can even lead to atmospheric escape (Murray-Clay et al., 2009).

An alternative formation theory is the gravitational disk instability (Perryman, 2018) that affirms that in some situations, some overdensities in the outer regions of the gas disk may reach a critical mass and start a gravitational collapse, capturing the surrounding gas. This mechanism is able to produce giant planets, skipping the uncertain planetesimal accretion phase. Once formed, planets interact with the protoplanetary disk and may undergo a phase of orbital migration.

2.2.2 Atmospheric Properties

Hot Jupiters share some common atmospheric properties with cold giant planets, also known as Jupiter-like planets. However, their close proximity to the host star and the resulting extreme stellar irradiation lead to a number of distinct and complex atmospheric phenomena.

Hot Jupiters typically possess thick H - He-rich envelopes surrounding dense cores of heavier elements. The main atmospheric chemical species are H_2O , CO , CH_4 , and CO_2 . At high temperatures and low pressures, CO dominates as the main carbon carrier, whereas at lower temperatures and higher pressures, CH_4 is more abundant. In some situations, even atomic species like Na and K have been detected ([Sudarsky et al., 2003](#)). The bulk metallicity of their atmospheres is strongly related to their formation history. Specifically, the location where a planet forms within the protoplanetary disk can influence the abundance of volatile elements, with planets forming in the outer regions typically enriched in these elements. Additionally, during the migration phase, giant planets may capture planetesimals, that contribute to enriching their atmospheres with more refractory elements that are typically more abundant in the inner parts of the disk.

Moreover, stellar irradiation can significantly modify the chemical and physical properties of a planetary atmosphere. Incident energy heats the outer layers, affecting the absorbing/emission properties of the atmosphere and its scale height. This heating may lead to a thermal inversion if optical or UV absorbers (e.g., TiO , VO , or photochemical hazes) are present in the upper atmosphere, resulting in a layer where temperature increases with altitude. Such inversions have been observed in different planets ([Lothringer et al., 2018](#); [Rajpurohit, A. S. et al., 2020](#)).

Due to their tidal locking, hot Jupiters experience permanent day and night-sides. The vertical penetration depth of stellar energy influences the day-night temperature contrast: shallow absorption leads to strong contrast, while deep absorption allows more efficient heat redistribution. The day-side receives constant stellar flux, driving intense atmospheric circulation. This can result in high-velocity zonal winds, with wind speeds reaching several km/s, transporting heat and chemical species from the day to the night-side ([Carone et al., 2023a](#)). In particular, this chemical species transport can lead to heterogeneous distributions of molecules across the planetary limb.

The heat distribution is extremely important in shaping condensate clouds and hazes, that are common in the atmospheres of hot Jupiters. Clouds may form from condensation of refractory species such as silicates ([Carone et al., 2023a](#)), while hazes are generated through UV photolysis. In extremely irradiated hot Jupiters (with $T_{eq} > 2000$ K), clouds are expected to form only on the night-side or the morning terminator, where the temperature is low enough to allow the condensation of chemical species, while photochemical hazes dominate on the hotter day-side and night terminator. These asymmetries can result in ingress and egress spectra probing different atmospheric conditions. Indeed, the aerosols composing these atmospheric structures can significantly flatten transmission spectra, obscure molecular features, and contribute to albedo variations. Differences in morning and night terminator spectra are extremely important since they give information about atmospheric heat transfer ([Perryman, 2018](#)).

2.2.3 Transmission Spectroscopy of Hot Jupiters

Hot Jupiters, due to their large atmospheric scale height, are ideal targets for transmission spectroscopy observations. Their large radii and thick atmospheres lead to strong transit signals, while their short orbital periods make them favorable for observation of multiple and significant transits. Additionally, they present low mean

molecular weight μ_m , being mostly composed by light elements like hydrogen and helium, and low surface gravity g_p because of their inflated radii, caused by intense stellar irradiation. Recalling Equations 2.27 and 2.28, these two characteristics enhance the atmospheric scale heights H , and thus the detectability of spectral features during transits. Figure 2.4 shows an example of transmitted spectrum of the hot Jupiter WASP-39b observed with JWST NIRSpec instrument in PRISM observing mode (Rustamkulov et al., 2023).




Figure 2.4: JWST/NIRSpec PRISM transmission spectrum of WASP-39b obtained by Rustamkulov et al. (2023).

2.3 Instruments and analysis

As mentioned in previous sections, the most important spectral features found in exoplanetary transmission spectra are observed in the infrared spectral range, since they are originated mostly from molecular transitions. Because of Earth's atmospheric extinction, this wavelength range is not easily accessible from ground observations. The most important results in this field are obtained through space observation, mainly using the Hubble Space Telescope (HST), and recently also the James Webb Space Telescope (JWST).

2.3.1 Hubble Space Telescope - WFC3

HST, fully operative since 1993, has been one of the most important space observatory ever made, contributing in almost every astrophysical branch of research. In particular, it had a major role in advancing the knowledge in exoatmospheres thanks to transmission spectroscopy (Gibson et al., 2012; Gressier, A. et al., 2022; Zhou et al., 2022). For this technique, it was mainly used the Wide Field Camera 3 (WFC3). This camera is a versatile imaging instrument designed to capture high-resolution images across a broad range of wavelengths. It features two independent channels: the UVIS channel, covering ultraviolet and visible light (200-1000 nm), and the IR channel, operating in the near-infrared range (850-1700 nm). WFC3 also contains a spectrograph, which, combined with its high sensitivity and broad spectral range, makes it ideal for observing the faint signals coming from exoplanetary atmospheres. It replaced the Wide Field and Planetary Camera 2 (WFPC2) during Servicing Mission 4 in 2009.

2.3.2 JWST - NIRSpec

JWST, launched in 2021, is one of the most ambitious space mission ever realized. It is designed to observe in the infrared spectral range with unprecedented sensitivity. It has a segmented hexagonal primary mirror with a total diameter of $6.5m$ allowing it to investigate faint sources. The [Near-Infrared Spectrograph](#) (NIRSpec) is a spectrograph, designed to work in the near-infrared range from 0.6 to 5.3 microns. It supports multiple observing modes, including multi-object spectroscopy, fixed slits, and an integral field unit (IFU), enabling simultaneous observation of up to 100 sources. NIRSpec offers low, medium, and high spectral resolutions ($R \sim 100$ to $R \sim 2700$), making it ideal for studying the chemical composition, temperature, and dynamics of exoplanetary atmospheres. Currently, the NIRSpec is the most important and most used instrument for investigating exoplanetary atmospheric properties ([Birkmann, S. M. et al., 2022](#); [Rustamkulov et al., 2023](#)).

2.3.3 Future Missions and Observatories

The most important future space mission, in the transmission spectroscopy field is the Ariel mission ([Tinetti et al., 2021](#)), scheduled for launch in 2029, aims to significantly advance the field of transmission spectroscopy by systematically characterizing the atmospheres of approximately 1000 exoplanets. Utilizing a 1-meter class telescope equipped with low-resolution spectrographs, Ariel will observe transiting and eclipsing exoplanets across infrared wavelengths from 1.1 to $7.8 \mu m$. This extensive survey will enable the detection of key molecular species and the assessment of thermal structures within these distant atmospheres, thereby addressing fundamental questions about planetary formation and evolution.

Another very important observatory will be the Extremely Large Telescope ([Padovani, Cirasuolo \(2023\)](#)), with first light expected in 2029. It will feature a 39-meter segmented primary mirror and cutting-edge instruments designed for high-resolution spectroscopy. Its unprecedented light-gathering power will significantly advance transmission spectroscopy, enabling detailed studies of exoplanet atmospheres, including potentially rocky, Earth-like worlds.

2.3.4 Data Analysis

Once preprocessed, transmission spectroscopic data are analyzed using the so-called retrieval method. It consists in iteratively comparing a model transmission spectrum with the observed one within a Bayesian framework. In particular, the Bayes' theorem states:

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)} \quad (2.30)$$

where $P(\theta|D)$ is the probability of having the set of model parameters θ given the data D . This quantity is referred to as the posterior distribution. $P(D|\theta)$ is the probability of observing the data D given a model described by the set of parameters θ . It is known as the likelihood. $P(\theta)$ is the prior distribution representing the prior knowledge about the system, i.e., the probability that the system is described by the parameters θ among all the possible parameter configurations. $P(D)$ is the Bayesian evidence.

By varying the parameters θ describing the model (e.g., equilibrium temperature, chemical species mixing ratios, planetary radius and mass), the posterior distribution is sampled over the parameter space. This process allows identification of the most probable set of parameters that best describe the observed data under the chosen model.

The two most commonly used algorithms implemented for sampling the posterior distribution are the Markov Chains Monte Carlo ([Foreman-Mackey et al., 2013](#)) and the Nested Sampling ([Skilling, 2004](#)). The former, while computationally simpler to implement, only samples the unnormalized posterior, i.e., $P(\theta|D) \propto P(D|\theta)P(\theta)$. In contrast, Nested Sampling is particularly powerful for efficiently computing the evidence $P(D)$, while also providing the posterior as a byproduct.

In the context of exoatmospheres, `TauREX 3` (Tau Retrieval for Exoplanets) is one of the most used tools. It is an open-source, Bayesian atmospheric retrieval framework available as a Python package. It is designed for analyzing exoplanetary atmospheres using transmission, emission, and direct imaging spectroscopy ([Al-Refaie et al., 2021](#)). It employs a nested sampling approach to explore the posterior distributions of atmospheric parameters, allowing users to fit observational data and infer chemical compositions, and temperature profiles. `TauREX 3` is highly modular and allows users to define different atmospheric models by specifying the chemistry, and the physics of the atmosphere, together with the retrieval parameters. It can be run both from Python scripts and directly from the command line (Bash). To improve accuracy in exoplanet transit modeling, `TauREX 3` supports stellar models, such as PHOENIX spectra, for host star characterization. Additionally, it efficiently handles molecular and atomic cross-section data (e.g., from ExoMol, HITRAN, HITEMP) by loading them into cache, reducing computational overhead during the retrieval. To enhance performance, `TauREX 3` supports parallel computing, allowing to distribute the computation required in the retrieval process across multiple cores. This significantly speeds up computations, especially when performing nested sampling over large parameter spaces (e.g., when retrieving many parameters at the same time).

Several other similar and compatible Bayesian retrieval codes are available such as NEMESIS ([Irwin et al., 2008](#)), ARCiS ([Min, Michiel et al., 2020](#)), CHIMERA ([Line et al., 2013](#)), and petitRADTRANS ([Mollière, P. et al., 2019](#)).

While extremely powerful, retrieval techniques are also computationally intensive, often requiring significant computational resources to achieve convergence within a reasonable time frame. Moreover, the choice of the model used to interpret the data is critical. In particular, over simplified models (e.g., isothermal temperature profiles or those neglecting the presence of clouds and hazes) can introduce strong biases, potentially compromising the reliability and physical accuracy of the retrieved atmospheric properties.

Chapter 3

Machine Learning in Astrophysics

Artificial Intelligence (AI) refers to the development of machines capable of exhibiting intelligent behavior, simulating human cognitive functions such as problem solving and decision-making¹. Machine Learning (ML) is a subdomain of AI that focuses on algorithms that learn directly from data by identifying patterns. Although widely used in industry, these algorithms have been increasingly adopted by the scientific community due to their efficiency in handling large and complex datasets.

Modern astrophysics deals with massive and complex datasets coming from observational surveys and simulations. This poses significant challenges for traditional tools, which not only struggle to generalize across different astrophysical scenarios but also tend to be slower and computationally expensive, especially when dealing with the vast volumes of data generated by modern instruments. Machine learning offers powerful methods to address these challenges, enabling efficient data analysis and feature extraction. Moreover, the flexibility and the variety of ML algorithms allow them to be easily adapted to many astrophysical studies (see Section 3.3). Another key advantage of ML is its ability to recognize patterns, making it particularly useful for identifying correlations within large datasets. This capability allows the identification of complex relationships, helping uncover hidden structures and insights in the data.

In this chapter, I first introduce the basic building blocks of machine learning, including neural networks and convolutional architectures. I then focus on autoencoders, a class of unsupervised models particularly suited for dimensionality reduction and data reconstruction. Finally, I discuss the application of ML to astrophysical problems, with a focus on ExoGAN (Zingales, Waldmann, 2018), a deep learning model used for exoplanetary atmospheric retrievals.

3.1 Machine Learning Basics

A machine learning algorithm is an algorithm that is able to learn from data. In particular, adopting the definition used by Mitchell (1997): *A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .*

The ability to learn from experience is the main strength of these machines. It enables the solution of tasks that are too complex for traditional, fixed programs. Examples include classification, language translation, anomaly

¹Definition taken from the IBM website

detection, imputation of missing values (also known as inpainting), and denoising.

We can distinguish between two learning approaches: supervised and unsupervised learning. They differ in the kind of experience that the algorithm is allowed to have during the training process ([Goodfellow et al., 2016](#)). In supervised learning, the dataset given to the algorithm contains not only the features from which the algorithm has to learn, but also the target output it is expected to produce. In unsupervised learning, the algorithm learns directly from the data the properties of the structure of the dataset, i.e., the probability distribution that generated the dataset itself.

To evaluate the performance of a machine learning algorithm, the key criterion is how well it performs on data that it has never seen before. This determines how well it will work when applied to real-world scenarios. To perform this evaluation, the dataset is typically split into two subsets: the train set, used to train the algorithm, and the test set, which contains unseen examples, used to simulate real-world data.

The main objective of machine learning is indeed to develop an algorithm that performs well on new, previously unseen inputs, and not just on the data on which it was trained. The ability to perform well on data never analyzed before is referred to as generalization. During training, our goal is to minimize the training error to allow the algorithm to learn. However, achieving a low training error alone is not sufficient. It is equally important to maintain a low test error, which serves as a proxy for the generalization performance of the model.

In practice, the factors determining how well a machine learning algorithm will perform are its ability to minimize the training error, and its ability to make the gap between training and test error small. These two factors correspond to the two fundamental challenges in machine learning: underfitting and overfitting.

- **Underfitting** occurs when the model is too simple to capture the underlying structure of the data, resulting in high training error.
- **Overfitting** occurs when the model learns the training data too well (including its noise), leading to a low training error but a much higher test error. This results in a poor generalization of the model.

A related concept is the capacity of a model, which refers to its ability to learn complex patterns in the data. Models with low capacity may not fit the train set as they cannot capture enough complexity. Models with high capacity may overfit by memorizing specific training examples rather than learning general patterns, leading to poor generalization.

In the following sections, I present the formalism and the fundamental tools used in machine learning in more detail. I also discuss their key strengths and limitations. A more in-depth treatment of these topics can be found in [Bishop \(2006\)](#) and [Goodfellow et al. \(2016\)](#).

3.1.1 Deep Feedforward Neural Networks

Deep-feedforward neural networks, often simply referred to as Neural Networks (NNs), are considered the foundation of Artificial Intelligence and, in particular, Machine Learning. The term "Neural Network" has its origin in attempts to find a mathematical representation of information processing in biological systems. In fact, the structure and logic behind these models are inspired by biological neural networks. They are organized into interconnected nodes that represent neurons and connections in the network. Like biological neurons, each node receives inputs (stimuli) from the others, processes them, and then sends its output forward. The term "Feedforward" arises from the fact that information flows in one direction only: from the input to the output.

These networks are organized into groups of nodes called layers. The layer that takes in the data is referred to as the input layer. The final layer of the network is referred to as the output layer. All the layers between these two are referred to as hidden layers. In the literature, nodes are also called perceptrons, consequently deep NNs are referred to as Multilayer Perceptrons (MLPs).

The "length" of the sequence of layers is referred to as the depth of the network, which gives rise to the term "Deep" in Deep Feedforward Neural Networks. The number of nodes in the Hidden Layers determines the width of the model.




Figure 3.1: Schematic representation of a Neural Network with 2 hidden layers. Image created adapting the Latex code available in the [TikZ](#) website.

Figure 3.1 shows the schematic representation of a Feedforward Neural Network, with an Input layer of dimension N , an Output layer of dimension K , and two Hidden layers, one of dimension D and the other of dimension M . The connections between the input nodes and the first node of the first hidden layer are highlighted in dark blue. In particular, this network is classified as a dense Neural Network, as each node in one layer is fully connected to every node in both the previous and subsequent layers. If this condition is not met, the model is referred to as sparse.

Mathematical formalism

In this section, I consider a feed-forward neural network model with two hidden layers, as shown in Figure 3.1. All the formulas presented can be extended to deeper networks. As previously seen, the basic structure of a neural network consists of multiple layers of processing units that can be described as a series of functional transformations. Each unit applies a nonlinear transformation to a weighted sum of its inputs. The input variables $\mathbf{x} = (x_1, \dots, x_N)$ are first linearly combined with the corresponding weights $w_{ji}^{(1)}$ and biases $w_{j0}^{(1)}$ in the first layer of the network (as shown for the first node in the first hidden layer). Specifically, for each hidden node j , the activation $a_j^{(1)}$ is computed as:

$$a_j^{(1)} = \sum_{i=1}^N w_{ji}^{(1)} x_i + w_{j0}^{(1)} \quad (3.1)$$

where $w_{ji}^{(1)}$ are the weights connecting the i -th input to the j -th hidden node, and $w_{j0}^{(1)}$ is the bias term. The (n) notation refers to the n -th layer with which the symbol is associated. The activation a_j is then transformed by a nonlinear activation function $h(\cdot)$, resulting in the hidden unit z_j :

$$z_j^{(1)} = h(a_j^{(1)}) \quad (3.2)$$

Common choices for the activation function $h(\cdot)$ include sigmoidal functions² and the Rectified Linear Unit function³ (ReLU). If the activation functions of all the hidden units in a network are taken to be linear, then, for any such network, we can always find an equivalent network without hidden units. This follows from the fact that the composition of successive linear transformations is itself a linear transformation (Bishop, 2006). This would severely limit the ability of the network to learn complex datasets. Nonlinearity is what gives neural networks the potential to approximate virtually any function, making them capable of modeling complex, high-dimensional data.

The hidden units $z_j^{(1)}$ are then linearly combined in the second layer, where each unit m receives the weighted sum of the hidden units. The activation $a_m^{(2)}$ for the m -th node is given by:

$$a_m^{(2)} = \sum_{j=1}^D w_{mj}^{(2)} z_j^{(1)} + w_{m0}^{(2)} \quad (3.3)$$

where $w_{mj}^{(2)}$, and $w_{m0}^{(2)}$ are respectively the weights and the biases in the second layer. Finally, the output y_k is obtained by applying an appropriate activation function $\sigma(\cdot)$ to the hidden units of the second hidden layer.

The overall neural network model can be expressed as:

$$y_k(\mathbf{x}, \mathbf{w}) = \sigma \left(\sum_{j=1}^D w_{mj}^{(2)} h \left(\sum_{i=1}^N w_{ji}^{(1)} x_i + w_{j0}^{(1)} \right) + w_{m0}^{(2)} \right) \quad (3.4)$$

where the set of all weights and biases, $\mathbf{w} = \{w_{ji}^{(1)}, w_{j0}^{(1)}, w_{mj}^{(2)}, w_{m0}^{(2)}\}$, are the adjustable parameters of the network, and the network output y_k is a function of the input variables $\mathbf{x} = (x_1, \dots, x_N)$ and the parameters $\mathbf{w}$.

Although I have presented a simple example, the concepts discussed are fundamental and play a crucial role as they are applied across nearly all machine learning algorithms.

Learning Process: Gradient Descent and Backpropagation

Deep Forward Neural Networks can be used as supervised models: each example in the dataset is associated with a target value $\mathbf{t}$ that the networks should reproduce. During the learning process, the parameters $\mathbf{w}$ are

²The sigmoid is a widely used nonlinear activation function in neural networks, particularly for binary classification tasks. It is defined as:

$$\sigma(x) = \frac{1}{1 + \exp(-x)}$$

The sigmoid function outputs values between 0 and 1, making it useful for modeling probabilities.

³The ReLU (Rectified Linear Unit) is a widely used nonlinear activation function in neural networks. It is defined as:

$$\text{ReLU}(x) = \max(0, x)$$

iteratively adjusted to maximize the performance of the model. This is typically achieved by minimizing an error function, often defined as the sum over individual data point errors:

$$E(\mathbf{w}) = \sum_{n=1}^N E_n(\mathbf{w}), \quad (3.5)$$

where $E_n(\mathbf{w})$ denotes the contribution to the total error of the n -th data point. Data points are also referred to as data examples.

The approach used to optimize a neural network is known as gradient descent. It involves updating the weights of the network based on the gradient of the error function with respect to the weights. The weights update rule for gradient descent is given by

$$\mathbf{w}^{(\tau+1)} = \mathbf{w}^{(\tau)} - \eta \nabla E(\mathbf{w}^{(\tau)}) \quad (3.6)$$

where $\eta > 0$ is the learning rate, and $\nabla E(\mathbf{w}^{(\tau)})$ is the gradient of the error function $E(\mathbf{w})$ at the current weight vector $\mathbf{w}^{(\tau)}$. This update is repeated iteratively, with the gradient being re-evaluated at each step τ . The steps τ of the learning process are also referred to as epochs.

In the case of batch gradient descent, the gradient is computed over the entire training set, making it computationally expensive. This approach can be quite slow because it requires evaluating the error function for all data examples at each iteration. Furthermore, gradient descent can be prone to get stuck in local minima or saddle points, leading to suboptimal solutions.

An alternative to batch gradient descent is stochastic gradient descent (SGD), also known as on-line gradient descent. In SGD, the weight update rule is applied after processing randomly selected data points one at a time, as opposed to the entire dataset:

$$\mathbf{w}^{(\tau+1)} = \mathbf{w}^{(\tau)} - \eta \nabla E_n(\mathbf{w}^{(\tau)}) \quad (3.7)$$

where $\nabla E_n(\mathbf{w}^{(\tau)})$ is the gradient of the error function for a single data point at epoch τ . This results in faster updates, as it avoids the need to process the entire dataset at each step. Another advantage of SGD is its stochastic nature that helps avoid local minima, increasing the efficiency in finding the global minimum. This makes SGD a better choice for handling large datasets. Indeed, stochastic gradient descent can update the weights based on just one data point. This makes it particularly suitable for training neural networks on large-scale datasets, where processing the entire dataset at once would be computationally prohibitive. Additionally, because SGD involves processing data points in a sequential manner, it can be more adaptable to online learning scenarios, where data may arrive over time.

As we have seen, the gradient of the error function is the key quantity used in updating the weights of the model. Error backpropagation, or backprop, is a key algorithm used to efficiently compute the gradient of the error function with respect to the weights of the network. Backpropagation involves two stages: forward propagation and backward propagation.

Forward propagation refers to the process of feeding the input data through the network to compute the activations of all the hidden and output units. As we have seen, each unit computes a weighted sum of its inputs and

then applies a nonlinear activation function:

$$a_j = \sum_i w_{ji} z_i + w_{j0}, \quad z_j = h(a_j) \quad (3.8)$$

where z_i represents the activation of an input unit, w_{ji} is the weight associated with the connection from the input unit i to the hidden unit j , and w_{j0} is the bias associated with the j -th unit. The activation function $h(\cdot)$ transforms the weighted sum a_j into the output z_j of the unit j .

In backward propagation, the goal is to compute the derivatives of the error function with respect to the weights in an iterative way. This is done using the chain rule knowing that E_n depends on the weights w_{ji} only via the summed input a_j to unit j :

$$\frac{\partial E_n}{\partial w_{ji}} = \frac{\partial E_n}{\partial a_j} \frac{\partial a_j}{\partial w_{ji}} \quad (3.9)$$

We introduce the notation:

$$\delta_j \equiv \frac{\partial E_n}{\partial a_j} \quad (3.10)$$

where δ 's are often referred to as errors. From equation 3.8 we easily obtain:

$$\frac{\partial a_j}{\partial w_{ji}} = z_i \quad (3.11)$$

Substituting Equations 3.10 and 3.11, we obtain:

$$\frac{\partial E_n}{\partial w_{ji}} = \delta_j z_i \quad (3.12)$$

Thus, in order to evaluate the derivatives, we need only to calculate the value of δ_j for each hidden and output unit, since the z_i have already been computed during forward propagation.

For the output units, we define the error as:

$$\delta_k = y_k - t_k \quad (3.13)$$

where t_k is the target value of the k -th output unit. To evaluate the δ 's for hidden units, we again make use of the chain rule for partial derivatives:

$$\delta_j \equiv \frac{\partial E_n}{\partial a_j} = \sum_k \frac{\partial E_n}{\partial a_k} \frac{\partial a_k}{\partial a_j} \quad (3.14)$$

where the sum runs over all units k to which unit j sends connections. This allows us to express how a change in a_j influences E_n via the units in the next layer a_k (the ones that receive input from units j).

If we now substitute the definition of δ given by Equation 3.10 into Equation 3.14, and make use of Equations 3.8, we obtain the following backpropagation formula:

$$\delta_j = h'(a_j) \sum_k w_{kj} \delta_k \quad (3.15)$$

where $h'(a_j)$ is the derivative of $h(a_j)$ with respect to a_j . This tells us that the value of δ for a particular hidden unit can be obtained by propagating the δ 's backwards from units higher up in the network. Since we already know the values of the δ 's for the output units, it follows that by recursively applying Equation 3.15 we can evaluate the δ 's for all hidden units in a feedforward network, regardless of its topology.

The backpropagation procedure can therefore be summarized as follows:

1. Apply an input vector $\mathbf{x}$ to the network and forward propagate through the network to find the activations of all the hidden and output units.
2. Evaluate δ_k for all output units using Equation 3.13.
3. Backpropagate the δ 's using Equation 3.15 to obtain δ_j for each hidden unit in the network.
4. Use Equation 3.12 to evaluate the required derivatives.

Backpropagation provides an efficient algorithm for computing the derivatives of the error function with respect to the network's parameters. These derivatives form the gradient ∇E used in the gradient descent optimization algorithm. The combination of these two techniques gives a practical and efficient way to iteratively adjust the weights w to minimize the error function and effectively train the neural network.

Although backpropagation involves a significant number of computations, especially in large and deep networks, its structure is parallelizable. This makes it particularly well-suited for implementation on modern Graphics Processing Units (GPUs), which are designed to handle large-scale parallel computations efficiently.

GPUs offer a major performance advantage in training neural networks due to their ability to execute matrix and element-wise operations. Since the key steps in forward and backpropagation, such as computing activations, propagating errors, and updating weights, can all be expressed as matrix operations, GPUs can process them simultaneously across multiple data examples.

As a result, GPUs enable the training of neural networks with millions of parameters on large datasets within feasible time frames, something that would be computationally impossible on traditional CPUs. This computational efficiency has been a critical factor in the success and rapid advancement of deep learning in recent years, enabling breakthroughs in fields such as computer vision and natural language processing.

3.1.2 Convolutional Neural Networks

Convolutional Neural Networks, or CNNs (LeCun, 1989), are a particular type of neural network specialized in processing data that have a grid-like topology, such as time-series data (1D grid) or image data (2D grid of pixels). Their name originates from the fact that the convolution operation is employed in place of general matrix multiplication in at least one layer (Goodfellow et al., 2016). In machine learning, the convolution $(x * w)(t)$ between two functions x and w is defined as:

$$s(t) = \int x(a)w(t-a)da \quad (3.16)$$

or, in its discrete formulation as:

$$s(t) = \sum_{a=-\infty}^{\infty} x(a)w(t-a) \quad (3.17)$$

The first argument is referred to as the input (I), the second as the kernel (K), and the output is commonly called the feature map. Convolution is a commutative operation, so, in some cases, the input and the kernel can be exchanged.

When working with image data, convolution is often applied over multiple axis at a time using a multidimensional kernel:

$$S(i, j) = (I * K)(i, j) = \sum_m \sum_n I(m, n)K(i - m, j - n) \quad (3.18)$$

In Figure 3.2, I present a schematic example of the application of a 2D convolution to an image. The output is restricted to only positions where the kernel lies entirely within the image. This type of convolution is sometimes referred to as valid convolution.




Figure 3.2: Schematic representation of a 2D convolution. Image taken from [Goodfellow et al. \(2016\)](#).

Typically, a layer in a convolutional network is divided into three stages. In the first stage, several convolutions are performed in parallel to produce a set of linear activations. In the second stage, a nonlinear activation function is applied to each linear activation. In the third stage, a pooling function is applied to further modify the output of the layer. Pooling functions replace the output of the net at a certain location with a summary statistic of the nearby outputs such as the maximum value (max pooling) or the average value (average pooling). The pooling operation is used to make the model robust to small translations of the input. This is particularly

useful in image feature extraction, where the presence of a feature is more important than its exact position (e.g., the presence of the eyes in a face recognition task).

CNNs leverage several fundamental concepts: sparse interactions, parameter sharing, and equivariant representation, which collectively enhance efficiency and effectiveness in image analysis. Unlike traditional fully connected networks, where each output unit interacts with every input unit, CNNs employ sparse connectivity by using kernels smaller than the input data. This drastically reduces the number of parameters and computational cost, while maintaining the ability to detect meaningful local features. Through parameter sharing, the same set of weights is applied across different regions of the input, enabling the network to detect similar features regardless of their spatial location. This further improves statistical efficiency and reduces storage requirements. Moreover, stacking multiple convolutional layers allows CNNs to progressively expand the receptive field, meaning that deeper layers can integrate information from increasingly larger portions of the input. This hierarchical structure enables the network to capture complex patterns and relationships by combining simple, local features detected by the kernel in early layers into higher-order abstractions, effectively analyzing features that are spatially distant in the original image. Equivariant representation is achieved through mechanisms such as pooling, which introduces invariance to small translations and distortions, enhancing robustness.

Overall, CNNs exploit the local correlation of image data, achieving efficient parameter storage through sparse connectivity and parameter sharing. Through the layered composition of convolutions, they can model intricate dependencies across spatially distant regions of the input, making them exceptionally powerful for pattern recognition, complex relationship modeling, and robust feature extraction in a wide range of applications.

3.1.3 Strengths and Limitations of Machine Learning Models

As already mentioned, the rapid advancement in ML in recent years is partly due to the improvement in GPU technology. GPUs allow for efficient training of large and complex networks, with millions of parameters, over large datasets. Once trained, these networks can achieve results comparable to classical algorithms in a much shorter time frame (Zingales, Waldmann, 2018), as they perform only simple calculations. In some cases, NNs can outperform traditional algorithms thanks to their flexibility in handling complex previously unseen data. This ability comes from the strong pattern recognition and feature extraction capabilities that some deep neural networks possess. These networks can detect subtle and intricate patterns within data, enabling them to effectively address the tasks for which they are designed.

However, these patterns are often unclear due to the complexity of the network and the dataset. As a result, it is difficult to understand exactly how these NNs process and interpret the data, leading to the so-called black-box problem (Hassija et al., 2023). It represents a major challenge in fields such as astrophysics, where interpretability is essential. This significantly limits the scientific reliability of the output of these models, introducing the need for a validation phase through comparison with classical algorithms. To address the black-box nature of complex machine learning models, especially deep neural networks, a range of interpretability methods has been developed. These techniques aim to provide local or global explanations by approximating the model's behavior or attributing feature importance to individual predictions. In the context of image data and CNNs, methods such as Grad-CAM (Selvaraju et al., 2019) highlight the regions of the input that most influence the model's decisions. While these approaches offer valuable insights, they often provide approximations rather than full transparency, and ongoing research seeks to improve their reliability and applicability.

As previously mentioned, one of the most critical challenges in ML is overfitting, a phenomenon in which a model learns the training data too well, including its noise and anomalies, resulting in poor generalization to

unseen data. Overfitting typically occurs when models have high capacity, leading them to memorize specific patterns rather than capturing the underlying trends. To mitigate this issue, it is essential to carefully select appropriate hyperparameters, such as the learning rate η , and the network architecture, to control the learning process and the complexity of the model. Proper hyperparameter tuning helps balance the trade-off between underfitting and overfitting, ensuring that the model captures meaningful patterns without becoming too tailored to the training data. Other common methods used to contrast overfitting are regularization, dropout, and early stopping. Regularization introduces additional constraints that penalize excessive model complexity. In particular, common regularization methods add a penalty term to the loss function based on the magnitude of model weights, encouraging simpler models that generalize better to unseen data. Dropout involves randomly deactivating neurons during training to encourage generalization. Early stopping halts the training when performance on a validation set⁴ starts to degrade. These strategies help ensure that the model captures the underlying patterns in the data without fitting noise or anomalies, thereby improving its robustness and predictive performance.

Furthermore, the performance and reliability of ML models heavily depend on the quality and representativeness of the datasets they are trained on. If the training data contains biases, the model will likely inherit and even amplify these biases, leading to inaccurate predictions when applied to new data. These biases may arise from incomplete understanding of the physical scenario, underrepresented classes, or limited data collection. Therefore, relying solely on large datasets is not sufficient to improve generalization; it is equally important to ensure that the data accurately reflect the true characteristics of the phenomena under study. Poor-quality, insufficient, or biased data can severely compromise a model's ability to generalize, regardless of how well the algorithm and its hyperparameters are optimized. Addressing both model complexity through hyperparameter tuning and data integrity is fundamental to developing robust and reliable ML systems.

3.2 Autoencoders

An autoencoder is an unsupervised neural network that is trained to copy its input to its output (Goodfellow et al., 2016). It can be described by two functions: an encoder and a decoder. The encoder takes the input x and encodes its main properties in a vector h often referred to as the latent vector. The decoder takes h , as input and decodes it to reproduce x , the original input of the model. A schematic representation of the structure of an autoencoder is presented in Figure 3.3.

Autoencoders are trained using standard methods such as gradient descent and backpropagation. Typically, the error function is a function that penalizes the reconstructed output r for being dissimilar from x . A typical choice for the loss function is the Mean Squared Error (MSE).

When the dimension of h is smaller than that of the input data x , the autoencoder is said to be undercomplete. In this case, the model learns to schematically describe the main properties of the data, thus learning a general representation of the dataset and the distribution of its properties. Due to this powerful ability, they are widely used in many tasks such as denoising, inpainting, and anomaly detection. Furthermore, if the model is trained not to perfectly reproduce x , it can also be used as a generative model.

Careful consideration must be given to the design of the network topology to avoid excessive model capacity. In fact, if the model is too complex, it may simply copy the input without extracting meaningful information from the dataset, thus performing poorly when applied to previously unseen data.

⁴The validation set is a subset of the train set used to estimate generalization performance during training.




Figure 3.3: Schematic representation of an Autoencoder. Image created adapting the Latex code available in the [TikZ](#) website.

An important subtype of autoencoders is Variational Autoencoders (VAEs). They are specifically designed for generative purposes. Thanks to their topology and their training, VAEs encode inputs into a probability distribution over the latent space, typically modeled as a multivariate Gaussian distribution. This structured latent space makes VAEs powerful generative models capable of producing new data samples similar to those in the training set by sampling from the prior distribution and passing these samples through the decoder. This makes VAEs widely used in tasks such as image generation. A more in-depth description of VAEs can be found in Appendix B.

In this work, I mainly focus on convolutional autoencoders to leverage their reconstructive and generalization abilities. This allows us to develop an efficient retrieval framework for extracting planetary and atmospheric parameters from transmission spectra of Hot Jupiters. Moreover, their reconstruction capability is used to develop an inpainting module to obtain missing spectral regions in observational data. VAEs offer potential future extensions to this work for generating synthetic spectra or exploring the variability within the parameter space of exoplanetary atmospheres.

3.3 Application to astrophysical problems

As we have seen, although machine learning models exhibit several limitations, they can offer significant advantages in scientific tasks when properly applied. ML models have proven to be highly time-efficient. Although the training phase can be very time-consuming, often requiring several days, the subsequent application to individual data points is remarkably fast. This is because, once trained, ML models rely on performing numerous simple calculations, unlike traditional algorithms, which must execute complex logic and computations for every data instance. This characteristic makes ML models particularly well-suited for managing large and complex datasets, that are common in astrophysical research. Furthermore, their ability to identify patterns and subtle correlations within data enhances their effectiveness in handling noisy or low-quality data, which represents one of the primary challenges in traditional data reduction techniques.

So far, the primary astrophysical applications of ML have been in star and galaxy classification ([Zeraatgari et al., 2023](#)), as well as in the detection of transiting exoplanets ([Chintarungruangchai, Jiang, 2019](#)). Both fields involve processing massive datasets generated by observational surveys and continuous monitoring from space

missions. In the case of galaxy classification, ML techniques are used to extract features from photometric and spectroscopic data to facilitate accurate automated classification. For transiting exoplanet detection, ML is particularly valuable for modeling and removing the effects of noise and stellar activity, which can introduce variations in light curves that either mimic planetary transits or obscure genuine transit signals.

Thanks to advances in technology and a growing understanding of its strengths and limitations, ML is beginning to expand into many other areas of astrophysics. The development of new ML-based tools, when integrated with traditional methodologies, is expected to significantly enhance scientific research.

In the following section, I present, as an example, a summary of the work described in [Zingales, Waldmann \(2018\)](#), which serves as the inspiration for this study, representing one of the first attempts to apply deep neural networks to retrievals of exoplanetary transmission spectra.

3.3.1 ExoGAN

[Zingales, Waldmann \(2018\)](#) introduces ExoGAN, a deep convolutional generative adversarial network (DCGAN) designed for atmospheric retrievals on exoplanets. Traditional atmospheric retrieval methods, like Bayesian techniques, require heavy computation, making complex atmospheric models impractical. ExoGAN uses unsupervised deep learning to efficiently recognize molecular signatures, trace-gas abundances, and planetary parameters from exoplanetary transmission spectra, significantly accelerating the retrieval process compared to traditional methods.

Generative Adversarial Networks (GANs) consist of two neural networks: the Generator (G) and the Discriminator (D), trained simultaneously through adversarial competition. The Generator tries to produce realistic data from random noise, while the Discriminator tries to distinguish real data from generated (fake) data. Training progresses until the Generator produces data indistinguishable from real data, and equilibrium is reached. At this point, the Generator can be used to generate new realistic data.

In ExoGAN, this structure is specifically realized as a Deep Convolutional Generative Adversarial Network (DCGAN), which utilizes convolutional neural networks for both the Generator and the Discriminator. Convolutional layers efficiently capture spatial correlations and complex patterns in spectral data, essential for analyzing atmospheric spectra.

The training dataset was generated using TauREx ([Al-Refaie et al., 2021](#)), consisting of 10^7 transmission spectra covering the spectral range between 1 and 50 μm . Each model varied 7 parameters: H_2O , CO_2 , CO , and CH_4 mixing ratios (with values ranging from 10^{-1} to 10^{-8}), together with planet mass M_p (with values ranging from 0.8 to 2 M_J), radius R_p (with values ranging from 0.8 to 1.5 R_J), and temperature T_p (with values ranging from 1000 to 2000 K). Both parameters and spectral data were normalized between 0 and 1 to make them machine-readable. In particular, spectral data normalization involved dividing the spectral range based on prominent water absorption bands and instrument-specific pass-bands (JWST NIRISS, NIRCams, MIRI, and Hubble WFC3), followed by scaling each spectral band between 0 and 1.

Parameters and spectral data were organized in the so-called Atmospheric Spectrum and Parametrization Array (ASPA). The ASPA is a 33×33 pixel array organized as presented in [Figure 3.4](#)

Region 1 is the normalized spectrum. Regions 2 to 5 contain the maximum values (regions 2 and 4) and the minimum values (regions 3 and 5) of the spectral segments used for the normalization. Regions 6 to 8 contain normalized atmospheric trace-gases mixing ratios of CO_2 , CO , and CH_4 respectively. Regions 9 to 11 are the




Figure 3.4: Representation of how spectral data, parameters, and normalization factors are organized in the ASPA matrix. Region 1 is the normalized spectrum. Regions 2 to 5 contain the maximum and the minimum values of the spectral segments used for the normalization. Regions 6 to 8 contain normalized atmospheric trace-gases mixing ratios of CO_2 , CO , and CH_4 . Regions 9 to 11 are the normalized Planetary parameters. Region 12 contains the normalized H_2O mixing ratio. Image taken from [Zingales, Waldmann \(2018\)](#).

normalized M_p , R_p and T_p respectively. Region 12 contains the normalized H_2O mixing ratio.

ExoGAN significantly improves computational speed compared to traditional methods like Nested Sampling or Markov Chain Monte Carlo. In fact, it reduces retrieval time from hours to approximately 2 minutes, regardless of the number of parameters. This corresponds to 300-fold speed-up.

ExoGAN retrievals closely match traditional TauREx retrievals within the error margins. Robustness tests confirm that, despite being less accurate, it reliably retrieves atmospheric parameters even in challenging scenarios, such as the presence of clouds, unidentified molecules (correctly identifying presence of known gases without including unknown molecular features), and parameters extending beyond its training range.

In conclusion, ExoGAN provides rapid, robust, and accurate atmospheric retrievals that can be coupled with traditional tools, like TauREx, to improve both the efficiency and the scientific output of exoplanetary atmospheric studies.

Extending ExoGAN with Autoencoders

In the context of ExoGAN's work on exoplanetary atmospheric retrievals, an autoencoder could significantly enhance both performance and accuracy. Autoencoders are particularly effective at compressing large datasets into lower-dimensional latent spaces, from which they can generate highly accurate reconstructions. This capability has been successfully applied across various industries, particularly in the field of personalized content suggestion such as music and movies (e.g., [Spotify](#), [Netflix](#)). In these applications, autoencoders analyze user preferences and behaviors, encoding this information into a simple representation within the latent space. The decoder can then use this compact representation to generate tailored content suggestions that closely match the user's tastes, creating a highly personalized experience.

In astrophysics, autoencoders can offer similar benefits. By training an autoencoder on exoplanetary transmission spectra, we can compress the data into a latent space that captures the underlying patterns in the spectrum,

such as molecular signatures and planetary parameters. Once trained, the decoder can be used to generate new, highly realistic synthetic spectra from the learned latent space representation. Additionally, autoencoders are particularly well suited at handling missing or incomplete data through inpainting techniques. For instance, in the case of missing spectral regions in observational data, the autoencoder's decoder can reconstruct the missing parts of the spectrum based on the learned representation from the encoder. This inpainting capability allows for the reconstruction of partial spectra and the recovery of the parameter space, expanding the capabilities of classical atmospheric retrieval algorithms. These techniques will be discussed in the following chapters. Thus, while ExoGAN offers rapid and reliable retrievals, autoencoders could enhance the accuracy of the reconstructed spectral regions and parameter space.

Chapter 4

Developing the Autoencoder

When developing a Machine Learning model, two key aspects must be considered: the model’s architecture and its training process (see Chapter 3). The architecture determines the model’s capacity (its ability to learn complex patterns), while the training process is related to how effectively the model learns from the data. The optimal choice of the model structure and training hyperparameters, such as the learning rate and the total number of training iterations (also referred to as epochs), is often determined through trial-and-error or grid search. Another important aspect is preprocessing the input data to make them machine readable and facilitate the learning process. In this chapter, I present the development of three convolutional autoencoders, focusing on the aspects mentioned above.

4.1 Dataset structure and Preprocessing

The dataset used to train the model corresponds to 1% of the dataset used in [Zingales, Waldmann \(2018\)](#). It is composed of 10^5 transmission spectra in the spectral range between $0.3\mu m$ and $50\mu m$ at resolution $R = 100$, obtained using TauREx ([Al-Refaie et al., 2021](#)) varying 7 planetary parameters: CH_4 , CO_2 , CO , H_2O mixing ratios together with the mass of the planet (M_p), its radius (R_p) and its isothermal temperature (T_p). In table 4.1 I present the upper and lower limits of the 7 parameters used in the dataset.

Parameter	Lower Bound	Upper Bound
H_2O	10^{-8}	10^{-1}
CO_2	10^{-8}	10^{-1}
CO	10^{-8}	10^{-1}
CH_4	10^{-8}	10^{-1}
M_p	$0.8 M_J$	$2.0 M_J$
R_p	$0.8 R_J$	$1.5 R_J$
T_p	1000 K	2000 K

Table 4.1: Parameters boundaries used in the dataset.

I normalized the parameters and the spectral value between 0 and 1 following the procedure used in [Zingales, Waldmann \(2018\)](#). This normalization helps the model identify relationships between parameters and spectral features which may have values differing by several orders of magnitude (for example, mixing ratios and spectral features). Moreover, it keeps the weight values small, which facilitates the optimization process.

In particular, mixing ratios are expressed in logarithmic form and divided by the logarithmic lower bound (-8),

while all other parameters are normalized by their respective upper bounds.

Spectral values are normalized in two steps: band subdivision and scaling. Band subdivision divides the spectrum into segments within which the variation between spectral values is minimal. Each spectral value x in the i -th segment is rescaled by applying the following relation:

$$\text{normalized_spectral_value} = \frac{x - \min_i}{\max_i - \min_i} \quad (4.1)$$

where $\min_i$ and $\max_i$ denote the minimum and maximum spectral values of the i -th segment. If $\min_i = \max_i$ the normalized spectral value is forced to be 0 to avoid division by 0 in Equation 4.1. The $\min_i$ and $\max_i$ values are also normalized, as they are included as inputs to the model. In particular:

$$\text{normalized_min}_i = \frac{\text{global_min}}{\min_i}, \quad \text{normalized_max}_i = \frac{\max_i}{\text{global_max}} \quad (4.2)$$

where $\text{global_min} = 0.006461$ and $\text{global_max} = 0.030191$ represent the minimum and maximum spectral values across the entire dataset.

The spectrum is divided into 14 segments. The choice of these segments is based on the filter pass bands of the main instruments of the James Webb telescope and on the Hubble Wide Field Camera 3 passband (see Section 2.1.4). Figure 4.1 provides a schematic representation of the spectral segment together with the instrumental passbands.




Figure 4.1: Normalization segment scheme and associated instrumental passbands. Instrumental passbands are obtained from [PandExo](#) online calculator.

This segmentation is designed to facilitate the use of these models for retrieval and inpainting tasks with real

observational data from JWST, which are often obtained through one or more filters rather than a continuous spectrum.

This normalization process ensures that all input values lie between 0 and 1, thus simplifying the training process and enhancing the model's ability to recognize patterns.

After normalization, all normalized input data (spectral values, normalization values, and parameters) are stored in a 33×33 pixel array called the Atmospheric Spectra and Parameters Array (ASPA) similar to [Zingales, Waldmann \(2018\)](#) (see Figure 3.4). Organizing the input data into a two-dimensional image-like structure enables us to leverage the power of Convolutional Neural Networks in efficiently identify patterns and correlations between spatially separated data (see Section 3.1.2). This approach is ideal for our goal of retrieving atmospheric and planetary parameters from spectral features, as well as for completing missing spectral regions.

As the final step of preprocessing, I added a third dimension to each ASPA matrix to simulate an image color channel. This modification is necessary to ensure compatibility with TensorFlow's built-in training functions.

4.2 Training

95% of the dataset was used as training set while the remaining 5% was used as the test set to study the generalization of the model. I adopted a learning rate of 10^{-4} , as the optimal balance between training speed and accuracy in gradient descent.

To train all models, I used the [ADAM optimizer](#) provided by the TensorFlow Python package ([Abadi et al., 2015](#)) with default hyperparameters. The ADAM optimizer, short for Adaptive Moment Estimation, is a widely used algorithm in deep learning due to its efficiency and effectiveness in training neural networks. It works by adaptively adjusting the learning rate for each parameter based on estimates of both the first moment (the mean) and the second moment (the variance) of the gradients. Essentially, ADAM keeps track of the average of recent gradients and the average of their squares, which helps it to understand both the direction and variability of the parameter updates. This adaptive approach allows ADAM to lead to faster convergence and better performance. For a more comprehensive explanation of the ADAM optimizer refer to [Ruder \(2017\)](#). Each step of gradient descent was performed by analyzing a batch of 950 data points (mini-batch size of 100).

I used two loss functions: the Mean Squared Error (MSE) and a combination of MSE with Binary Cross Entropy¹ (BCE). In both cases, MSE is used to minimize errors in the reconstruction of input data. In the second case, I add to the MSE the average BCE, computed between each element of the batch and each element of a representative subset corresponding to 5% of the training set. This additional term acts as a penalty for reconstructions that deviate too much from the subset used as a reference, which is intended to represent the entire dataset. This strategy helps organize the latent space of the autoencoder, limiting non-physical reconstructions, and improving the physical plausibility of the output. Both the MSE and the BCE functions are implemented in TensorFlow.

The training was performed on a virtual machine with a NVIDIA TESLA T4 GPU on the [CloudVeneto](#) server. The GPU allowed to keep the training time for each model under 10 hours (roughly 1000 epochs) even for the

¹The Binary Cross Entropy (BCE) loss for a single prediction $\hat{y}$ and true label y is defined as:

$$\text{BCE}(y, \hat{y}) = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})].$$

BCE measures the dissimilarity between predicted values and true labels. It is commonly used for classification tasks and here used to penalize large deviations in reconstruction similarity.

more complex models.

4.3 Best models

In this section, I present the three best-performing models, describing their architectures, training procedures, and reconstruction performance. To differentiate between them, I named the models **AURA**, **AURA-Small**, and **AURA-BCE-Loss**. AURA stands for Autoencoder Utility for Retrieval and Reconstruction of exoAtmospheres. The operations employed in these models include:

- **Bi-dimensional convolution and its inverse**, referred to as `Conv2D` and `Conv2DTranspose`, respectively. These operations are used to combine data that are spatially separated within the ASPA matrix.
- **Flatten** and **Reshape** operations, which convert bi-dimensional arrays into vectors and vice versa.
- **Dense layers**, corresponding to standard fully connected feed-forward operations (see Section 3.1.1). They are typically applied after a `Flatten` operation to obtain the latent vector, or before a `Reshape` operation to reconstruct the ASPA matrix.

To assess the reconstruction performance of the models, I employed the Absolute Relative Error (ARE), which, for a parameter x , is defined as:

$$\text{ARE} = \frac{|x_{\text{reconstructed}} - x_{\text{true}}|}{|x_{\text{true}}|} \quad (4.3)$$

where x_{true} is the original value and $x_{\text{reconstructed}}$ is the reconstructed value.

4.3.1 AURA

The AURA model is the most complex model developed. Its architecture is summarized in Table 4.2.

Layer type	Output Shape	Number of Parameters
Encoder		
Conv2D	(33, 33, 16)	80
Conv2D	(33, 33, 32)	2,080
Flatten	(34848)	0
Dense	(512)	17,842,688
Dense	(100)	51,300
Total params:		17,896,148
Decoder		
Dense	(512)	51,712
Dense	(34848)	17,877,024
Reshape	(33, 33, 32)	0
Conv2DTranspose	(33, 33, 16)	2,064
Conv2DTranspose	(33, 33, 1)	65
Total params:		17,930,865

Table 4.2: Summary of Encoder and Decoder of the AURA Model architecture. Layer type indicates the operation performed to obtain that layer. Conv2D stands for a 2 dimensional convolution, and its inverse operation is Conv2DTranspose. Flatten operation transforms an array into a vector, and its inverse operation is Reshape. Dense layers represent the standard feed-forward fully connected operations (see Section 3.1.1). The amount of storage space required for this model is 410 MB.

For the two convolutional layers, I used kernels of size $2 \times 2 \times 16$ and $2 \times 2 \times 32$, respectively. These kernels

extended color channels to effectively extract relevant features. No pooling operations were applied, preserving the original 33×33 ASPA dimensions. The model was trained for 800 epochs using the Mean Squared Error loss function. Table 4.3 summarizes the model's reconstruction performance on the test set.

Metric	Mean absolute relative error (%)	Std (%)
Entire ASPA	0.81	4.23
Spectrum region	1.57	8.86
Normalization regions	0.08	0.03
Parameters regions	0.15	0.07

Table 4.3: Mean absolute relative error and its standard deviation values for the entire reconstructed ASPA matrix and its main sub region obtained using ASPA model on the test set.

The AURA model has an extremely high overall reconstruction accuracy. One of the main reasons for this result is probably the large latent vector with a dimension of 100, which allows the model to encode the input data without requiring significant compression.

In figure 4.2 I present the original data against the reconstructed data for one test example. We can clearly see the extremely high reconstruction accuracy in both the spectrum data points (with an average absolute relative error of 0.03 %) and the parameters. The CO mixing ratio exhibits the lowest accuracy, probably due to the lack of relevant spectral features in this region of the electromagnetic spectrum.




Figure 4.2: Original and reconstructed data for one test example (id: 3992) using AURA model. The errors are given as absolute relative errors.

4.3.2 AURA-Small

AURA-Small is a simpler model than AURA. Its architecture is presented in Table 4.4.

For the two convolutional layers, I used kernels of size $2 \times 2 \times 32$ and $2 \times 2 \times 64$, respectively. In the first layer, the kernel was moved across the ASPA matrix in steps of 3 pixels at the time ($strides = 3$). This resulted in reducing the dimension of the ASPA matrix from 33×33 to 11×11 . The dimensionality reduction combined with a latent vector with a dimension of 50 forced the model to find stronger and more interconnected patterns,

Layer type	Output Shape	Number of parameters
Encoder		
Conv2D	(11, 11, 32)	832
Conv2D	(11, 11, 64)	51,264
Flatten	(7,744)	0
Dense	(512)	3,965,440
Dense	(100)	51,300
Dense	(50)	5,050
Total params:		4,073,886
Decoder		
Dense	(100)	5,100
Dense	(512)	51,712
Dense	(7,744)	3,972,672
Reshape	(11, 11, 64)	0
Conv2DTranspose	(11, 11, 32)	51,232
Conv2DTranspose	(33, 33, 1)	801
Total params:		4,081,517

Table 4.4: Summary of Encoder and Decoder of the AURA-Small Model architecture. Layer type indicates the operation performed to obtain that layer. Conv2D stands for a 2 dimensional convolution, and its inverse operation is Conv2DTranspose. Flatten operation transforms an array into a vector, and its inverse operation is Reshape. Dense layers represent the standard feed-forward fully connected operations (see Section 3.1.1). The amount of storage space required for this model is 94 MB.

leveraging the properties of Convolutional Neural Networks. The model was trained for 2500 epochs using the Mean Squared Error loss function. Table 4.5 summarizes the model reconstruction performance in the test set.

Metric	Mean absolute relative error (%)	Std (%)
Entire ASPA	1.00	3.70
Spectrum region	1.95	7.75
Normalization regions	0.14	0.05
Parameters regions	0.14	0.06

Table 4.5: Mean absolute relative error and its standard deviation values for the entire reconstructed ASPA matrix and its main subregions obtained using ASPA-Small model on the test set.

From Table 4.5 we observe that the overall reconstruction accuracy is high, though slightly lower than that of the AURA model. This decrease in precision is also evident in Figure 4.3, where the reconstructed spectrum shows a high overall accuracy but appears less precise, and the parameters exhibit higher absolute relative errors.

4.3.3 AURA-BCE-Loss

The AURA-BCE-Loss model has the same architecture as the AURA model (see Table 4.2). The difference between these two models lies in the training phase. AURA-BCE-Loss has been trained for 1000 epochs with a personalized loss function that combines the Mean Squared Error (MSE) function and the Binary Cross Entropy (BCE) function. The personalized loss function is the sum of the MSE and the average BCE computed between each element of the training batch and each element of a representative subset corresponding to 5% of the training set. As previously described, the MSE is used to minimize errors in the reconstruction of the input data. The additional BCE term acts as a penalty for reconstructions that deviate too much from the subset used as a reference, which is intended to represent the entire dataset. This strategy helps organize the latent space of the autoencoder, thereby limiting nonphysical reconstructions. The balance between the two loss terms is




Figure 4.3: Original and reconstructed data for one test example (id: 3992) using AURA-Small model. The errors are given as absolute relative errors.

controlled by the α factor:

$$LOSS = MSE + \alpha \cdot BCE \quad (4.4)$$

I used $\alpha = 10^{-4}$ to ensure that the BCE term becomes significant only after the model has learned to reconstruct the ASPA matrix. This approach allows the model to first focus on accurate reconstruction and, in a second phase, to organize its latent space to prevent nonphysical reconstructions.

Metric	Mean absolute relative error (%)	Std (%)
Entire ASPA	2.38	14.89
Spectrum region	4.83	31.21
Normalization regions	0.14	0.06
Parameters regions	0.15	0.09

Table 4.6: Mean absolute relative error and its standard deviation values for the entire reconstructed ASPA matrix and its main subregions obtained using ASPA-BCE-Loss model on the test set.

As we can observe from Table 4.6 the performance of this model on the test set is significantly worse even though, in some cases, the reconstruction accuracy is comparable to the one of the previous models (see Figure 4.4). For this reason, caution must be taken when using it.

4.4 Models Summary and Comparison

In Table 4.7, I summarize the performance of the three models. The AURA model demonstrates superior accuracy across all metrics when compared to the other models. While the AURA-Small model delivers slightly less precise reconstructions, its low storage requirements make it more practical for use in publicly available tools. The AURA-BCE-Loss model, due to its lower accuracy, should be used with caution. Considering these factors, the next chapter will primarily focus on the AURA model, although some discussion of the AURA-




Figure 4.4: Original and reconstructed data for one test example (id: 3992) using AURA-BCE-Loss model. The errors are given as absolute relative errors.

BCE-Loss model will also be included.

Model	Metric	Mean absolute relative error (%)	Std (%)
AURA	Entire ASPA	0.81	4.23
	Spectrum region	1.57	8.86
	Normalization regions	0.08	0.03
	Parameters regions	0.15	0.07
AURA-Small	Entire ASPA	1.00	3.70
	Spectrum region	1.95	7.75
	Normalization regions	0.14	0.05
	Parameters regions	0.14	0.06
AURA-BCE-Loss	Entire ASPA	2.38	14.89
	Spectrum region	4.83	31.21
	Normalization regions	0.14	0.06
	Parameters regions	0.15	0.09

Table 4.7: Mean absolute relative error and standard deviation values for the entire reconstructed ASPA matrix and its main subregions for the three models: AURA, AURA-Small, and AURA-BCE-Loss.

Chapter 5

Retrieval and Inpainting algorithm

In this chapter, I describe the retrieval and inpainting algorithm, detailing its strengths and limitations. The results of its application to the test set are also presented to verify its reliability with previously unseen data. Then I present the case study of WASP-39 b to test the algorithm in a real-case scenario.

5.1 Description of the Algorithm

The main goal of the retrieval and inpainting algorithms is to obtain missing data from an incomplete ASPA matrix. In this thesis, I define the distinction between retrieval and inpainting based on the portion of the incomplete ASPA matrix that is targeted for reconstruction. Specifically, when the reconstruction is limited only to the seven parameter values, the process is referred to as retrieval. When the goal extends beyond the parameters to include also the recovery of spectral regions, the process is referred to as inpainting.

The algorithm leverages the reconstruction ability of the models described in Chapter 4. The central idea is to find a latent vector that, when decoded by the decoder of the model, produces a reconstruction that closely matches the known parts of the input data while plausibly filling in the missing areas.

The algorithm takes as input an ASPA matrix containing missing values. A boolean mask is generated to distinguish between observed unmasked values (identified as True in the mask) and missing masked pixels (identified as False in the mask). The mask is constructed based on the spectral passbands described in Section 4.1 to facilitate normalization of the observed data to make them machine readable. Therefore, the algorithm can be used only when one or more passbands are completely available (refer to Figures 4.1 for the passband boundaries). Parameters may be unmasked individually if their values are derived from alternative methods or obtained from the literature.

The reconstruction process begins by initializing a latent vector with values drawn from a uniform distribution between 0 and 1. In the case of the AURA model this vector has a dimension of 100. The core of the algorithm is an iterative optimization process that updates the latent vector to minimize the mean squared error between the reconstructed ASPA and the original data, but only over the known (unmasked) regions. The optimization is performed over a maximum of 2×10^4 iterations using the Adam optimizer. It employs a learning rate of 0.1 for the first 1000 iterations, a learning rate of 0.01 for the following 2000 iterations, and a learning rate of 0.001 for the remaining iterations. This learning rate schedule is chosen to speed up convergence while maintaining a high accuracy in gradient descent. Early stopping is applied based on the change in the loss function between

consecutive iterations: if the absolute difference between the current and previous loss values remains below 2×10^{-8} for 500 consecutive iterations (the number of consecutive iteration is referred to as patience), the optimization process is terminated to prevent unnecessary computations. The use of the absolute difference ensures that the optimization is halted even if the loss value fluctuates near a minimum, avoiding unnecessary iterations when no further improvements are obtained.

The algorithm is extremely flexible, allowing us to easily change most of the hyperparameters such as the learning rate schedule, the number of maximum iterations, and the loss minimum variation. The final output of the algorithm consists of the ASPA with reconstructed missing values and the optimized latent vector that produced this reconstruction.

This approach effectively exploits the learnt latent space captured by the autoencoder, allowing a plausible and data-consistent reconstruction of missing regions based on available information and the model’s internal representation.

5.2 Retrieval and Inpainting using AURA model

To test the reliability of the model when used for retrieval or inpainting tasks, I applied the algorithm to 100 ASPA matrices randomly selected from the test set. Then I compared these 100 reconstructions with the original ASPAs to derive statistical measures of the reconstruction errors. I computed the absolute relative error on the seven parameters and the average absolute relative error of the spectral data points. For the mixing ratios, the absolute relative error was computed after converting the reconstructed values back from their logarithmic form to their original scale. In the following sections I present the main results of the retrieval and inpainting algorithm.

5.2.1 Retrieval

As mentioned above, in a retrieval task we aim to obtain the 7 parameters starting from the entire spectrum between 0.3 and 50 μm . Table 5.1 presents the results.

Parameter	Median Absolute Relative Error (%)
CH ₄ mixing ratio	2.670071
CO ₂ mixing ratio	1.941011
CO mixing ratio	29.428775
H ₂ O mixing ratio	1.726851
M_p	3.724719
R_p	0.032856
T_p	3.352776

Table 5.1: Median absolute relative error for 7 parameters computed from 100 retrievals using AURA model.

The results show a high degree of reliability, with all the parameters presenting a relative absolute retrieval error below 4%. The only exception is the CO mixing ratio with an error of roughly 30%. This could be explained by the lack of strong spectral features, which limits the ability of the model to find correlation between the spectrum and the CO mixing ratio value.

This test is particularly useful for evaluating the reliability of the retrieved parameters and to analyze how well the model has learned to associate spectral features with the parameters. However, it is important to note that this approach is not typically applicable in real-world scenarios, as such wide spectral range is rarely available.

5.2.2 Inpainting

NIRSpec Prism passband

I applied the inpainting algorithm to 100 ASPAs leaving the NIRSpec Prism passband (from $0.61\mu m$ to $5.14\mu m$) unmasked to recover the missing spectral regions and parameters.

Parameter	Median Absolute Relative Error (%)
CH ₄ mixing ratio	3.376949
CO ₂ mixing ratio	2.822296
CO mixing ratio	33.007576
H ₂ O mixing ratio	2.771587
M_p	10.535164
R_p	0.039996
T_p	10.070746

Table 5.2: Median Absolute Relative Error for 7 parameters computed from 100 inpainted ASPAs using AURA model, leaving the NIRSpec Prism passband unmasked.

The results presented in Table 5.2 show high reliability in recovering the mixing ratios of CH₄, CO₂ and H₂O and the radius of the planet (R_p). In most cases, the mass of the planet and its temperature are recovered with an error lower than 10%. The CO mixing ratio obtained with this method presents a median absolute relative error greater than 30%, indicating that the results for this parameter are unreliable.




Figure 5.1: Example of the results obtained by the inpainting algorithm leaving the NIRSpec Prism passband unmasked.

The average absolute relative error of the spectral data is 0.03%, indicating an extremely high reliability in reconstructing missing spectral regions. This reconstruction accuracy can be observed in Figure 5.1, which presents an inpainted spectrum example.

The high reliability of the inpainting results obtained from the NIRSpec Prism passband can be attributed to its relatively wide wavelength coverage and the abundance of spectral features it contains. This enables the algorithm to better exploit correlations between spectral data and parameters.

NIRSpec G235 f170lp filter

When applied to single filters, such as the NIRSpec G235 f170lp (passband covering the 1.69 - 3.06 μm wavelength range) the inpainting results are extremely poor, especially for the retrieved parameters that present median absolute relative errors exceeding 10%. The radius of the planet is the only parameter reliably retrieved, with a median absolute relative error of 0.19%.

Leaving the mixing ratio of H_2O unmasked, which is a common atmospheric trace gas marker in Hot Jupiters, does not improve the recovery of spectral values or parameters.

These poor results can probably be explained by the small spectral range covered by this filter, which limits the correlation that the AURA model can use to fill the missing values.

To test a better signal extraction algorithm, I performed the retrieval using the AURA-BCE-Loss model, which incorporates a penalty term to discourage nonphysical reconstructions. While the model still performs poorly in retrieving the 7 parameters, it shows promising results in recovering the missing spectral region (see Figure 5.2).




Figure 5.2: Comparison of the recovered spectrum using the AURA model (upper panel) and the AURA-BCE-Loss model (bottom panel) starting from the same test example (id:3992) in the NIRSpec G235 f170lp filter passband.

The difference in the accuracy of recovering spectral values can be attributed to the different loss functions employed during the training of the two models. The penalization term introduced by the BCE in the AURA-BCE-Loss model forces the reconstruction to resemble the spectra of the reference set (see Section 4.3.3), thereby limiting non-physical reconstructions and improving reconstruction accuracy. However, the still imper-

fect reconstruction and the poorly retrieved parameters indicate that further investigation and improvement are required.

NIRISS and MIRI passbands

I tested the inpainting algorithm by combining two passbands: NIRISS passband (from $0.82\mu m$ to $2.72\mu m$) and MIRI passband (from $5.14\mu m$ to $13.76\mu m$). This test aims to simulate the combination of multiple instrument observations of the same target. The results are summarized in Table 5.3.

Parameter	Median Absolute Relative Error (%)
CH ₄ mixing ratio	7.674695
CO ₂ mixing ratio	7.733877
CO mixing ratio	68.043462
H ₂ O mixing ratio	6.557396
M_p	17.437857
R_p	0.091591
T_p	16.745207

Table 5.3: Median absolute retrieval error for 7 parameters computed from 100 retrievals using AURA model, leaving the NIRISS and MIRI passbands unmasked.

The only parameter recovered below 1% precision is the radius of the planet R_p with an absolute relative error of 0.09%. The mixing ratios of CH₄, CO₂ and H₂O are recovered with a slightly larger median error compared to the inpainting using the NIRSpec Prism passband. The reconstruction of spectral values, with an average absolute relative error of 0.05%, is comparable to the result obtained with the NIRSpec Prism passband. In Figure 5.3, I present an example of inpainted spectrum.




Figure 5.3: Example of inpainting leaving the NIRISS and MIRI passbands unmasked using AURA model.

5.3 Case study

In a realistic scenario, spectral data are affected by both astrophysical and instrumental noise. To handle this within the inpainting procedure, the central value of each data point is perturbed by drawing it from a Gaussian

distribution. This distribution has a mean equal to the central value and a standard deviation corresponding to the measurement's error bars.

Specifically, for each spectral data point (x_i), multiple perturbed instances ($\tilde{x}_i^{(j)}$) are generated as follows:

$$\tilde{x}_i^{(j)} \sim \mathcal{N}(x_i, \sigma_i), \quad (5.1)$$

where x_i is the observed value and σ_i is the associated uncertainty.

The inpainting algorithm is then applied separately to each perturbed instance, resulting in a set of reconstructed spectra. The final estimate for each spectral and parameter value is given by the median of these reconstructions. To quantify the uncertainty in this estimate, the 16th and 84th percentiles of the distribution of reconstructed values are considered to correspond to a one-sigma confidence interval for a Gaussian distribution. This approach accounts for the uncertainties in the original data, propagating them through the inpainting process.

5.3.1 Inpainting test with synthetic noise

To simulate a real case application of the algorithm, I perturbed one of the test set examples (id=3992) with a constant noise of 30 ppm, generating 1000 different instances. I then applied the retrieval and inpainting algorithm leaving the NIRSpec Prism passband unmasked. I chose this passband because it was the one that showed the highest reconstruction accuracy.

Parameter	Retrieved value	True value
$\log(\text{CH}_4)$	$-3.7^{+1.0}_{-1.2}$	-4.1
$\log(\text{CO}_2)$	$-2.8^{+1.0}_{-1.2}$	-3.3
$\log(\text{CO})$	$-5.2^{+2.3}_{-2.2}$	-4.9
$\log(\text{H}_2\text{O})$	$-1.7^{+0.7}_{-0.9}$	-2.6
M_p	$1.36^{+0.46}_{-0.50}$	1.07
R_p	$1.37^{+0.04}_{-0.05}$	1.42
T_p	1100^{+440}_{-390}	1100

Table 5.4: Results of the inpainting algorithm (leaving the NIRSpec Prism passband unmasked) applied to the test example 3992 compared with the true values.

Results are presented in Table 5.4. The retrieved values correspond to the median of the distributions of the recovered parameters. Uncertainties are computed as the difference between the 16th and 84th percentiles of the distribution and the median value. The distribution of the seven parameters are presented in Appendix C.

As shown in Table 5.4, all seven parameters are consistent with the true values within their respective error bars. However, it is worth noting that the uncertainties are relatively large.

Figure 5.4 presents the retrieved spectrum and its error bars compared to the original synthetic spectrum. The spectral data are compatible with the original spectrum within the error bars; however, the uncertainties are large.

These results show that the algorithm is able to converge towards the true data but the uncertainties prevent giving accurate results.




Figure 5.4: Retrieved spectrum and its error bars compared with the original denoised synthetic spectrum for test example 3992.

5.3.2 WASP-39 b case study

WASP-39 b is a hot Saturn mass exoplanet located in the Virgo constellation, approximately 230 ± 80 pc from Earth. This exoplanet orbits its host star, WASP-39, with orbital period of 4.0552842 ± 0.0000034 days at a remarkably close distance of 0.04828 ± 0.00082 astronomical units (AU). Due to this proximity, WASP-39b experiences intense stellar irradiation, resulting in a high equilibrium temperature of approximately 1166 ± 14 K. The planet’s physical properties include a radius of $1.279 \pm 0.040 R_{Jup}$ and a mass of $0.281 \pm 0.032 M_{Jup}$, placing it among the class of inflated hot Jupiters. It was discovered in 2011 (Faedi et al., 2011) via transit method during the SuperWASP survey (Pollacco et al., 2006).

I report in Tables 5.5 and 5.6 the stellar and planetary parameters of the WASP-39 system (Faedi et al., 2011; Mancini et al., 2018).

Parameter	Value
Mass (M_*) [$M_\odot$]	0.913 ± 0.047
T_{eff} [K]	5485 ± 50
Radius (R_*) [$R_\odot$]	0.939 ± 0.22
Density (ρ_*) [g cm^{-3}]	1.555 ± 0.080
Distance [pc]	230 ± 80
Apparent Magnitude (V)	12.095 ± 0.252
Spectral type	G8
RA	14h29m18.40s
Dec	-03°26'40.20"

Table 5.5: Stellar parameters of WASP-39 obtained from NASA Exoplanet Archive website, in particular from Mancini et al. (2018). The spectral type and distance were found in Faedi et al. (2011).

Following its detection, further studies of WASP-39 b have focused on its atmospheric composition and structure. The observations from JWST NIRSpec PRISM allowed spectral analysis of WASP-39 b, detecting key molecular species such as H_2O , CO_2 , CO , and SO_2 (Khorshid et al., 2024; Rustamkulov et al., 2023). The detected CO_2 absorption at $4.3 \mu\text{m}$, along with the absence of CH_4 , strongly suggests that WASP-39 b has a super-solar metallicity ($10\times\text{solar}$)(Khorshid et al., 2024).

Another important characteristic of this exoplanet is the evidence for atmospheric asymmetry across the morning and evening terminators. As suggested by Espinoza et al. (2024) there is a temperature difference of about

Parameter	Value
Mass (M_P) [M_{Jup}]	0.281 ± 0.032
T_{eq} [K]	1166 ± 14
Radius (R_P) [R_{Jup}]	1.279 ± 0.040
Density (ρ_P) [g cm^{-3}]	0.167 ± 0.023
Orbital period (P) [days]	4.0552842 ± 0.0000034
Semi-major axis (a) [AU]	0.04828 ± 0.00082
Inclination of the orbit (i) [$^\circ$]	87.32 ± 0.17

Table 5.6: Planetary parameters of WASP-39b obtained from Nasa Exoplanet Archive website, in particular from [Mancini et al. \(2018\)](#).

177 ± 57 K between these regions, with the evening terminator being significantly hotter. This asymmetry is likely driven by strong equatorial jet streams redistributing heat from the dayside to the nightside. Additionally, the morning terminator exhibits greater cloud coverage compared to the evening terminator ([Carone et al., 2023b](#)). Cloud coverage is an important aspect to consider when studying transmission spectra, since clouds properties play a crucial role in shaping the observed data (see Section 2.1.4). In particular, [Carone et al. \(2023b\)](#), proposes that WASP-39b is characterized by a patchy cloud structure composed of silicates (MgSiO_3 , Mg_2SiO_4) and metal oxides (Fe_2O_3 , TiO_2). The complex nature and distribution of WASP-39b's clouds ([Carone et al., 2023b](#); [Espinoza et al., 2024](#)) require careful modeling in order to obtain valuable atmospheric retrievals.

In addition to this, the observed atmospheric abundances of WASP-39b are closely linked to its formation history ([Khorshid et al., 2024](#)). As already mentioned, the planet's atmosphere exhibits a super-solar metallicity, in particular it presents a carbon-to-oxygen (C/O) ratio of 0.7 ([Rustamkulov et al., 2023](#)). The significant presence of CO_2 implies an origin beyond the CO_2 snow line, where CO_2 and other volatile species could condense into solid icy particles and be efficiently accreted into the forming planet ([Khorshid et al., 2024](#)). The presence of sulfur-rich compounds, like SO_2 , suggests that the planet also accreted material from regions containing refractory elements. This may have influenced the abundance of SO_2 and other trace species. This accretion of refractory elements probably happened during WASP-39b's disk driven inward migration ([Khorshid et al., 2024](#)).

Inpainting retrieval

I analyzed the transmission spectrum of WASP-39b obtained from a transit observation on July 10, 2022, using JWST's Near InfraRed Spectrograph (NIRSpec) in PRISM mode. This observation was part of the JWST Transiting Exoplanet Community Early Release Science Program (ERS Program) ([Rustamkulov et al., 2023](#)). Figure 5.5 presents the observed transmission spectrum along with its associated uncertainties. The transit depth is expressed as a percentage.

These data were pre-processed using FIREFLY (Fast InfraRed Exoplanet Fitting for Lightcurves), and calibrated applying the JWST Python pipeline. The resulting spectrophotometry was then binned in wavelength to create 207 variable-width spectral channels with roughly equal counts in each. Finally, to extract the atmosphere's transmission spectrum, the planet's transit depth was modeled in each wavelength bin using a limb darkened transit light curve. This data reduction was performed by [Rustamkulov et al. \(2023\)](#) and the resulting transmission spectrum is publicly available in the dedicated section of the [NASA Exoplanet Archive website](#).

To match the wavelength grid required to build the ASPA, I interpolated the observed spectral data. Because




Figure 5.5: JWST NIRSpec PRISM transmission spectrum of WASP-39 b obtained by [Rustamkulov et al. \(2023\)](#).

uncertainties for the interpolated points are not available, a constant error equal to the maximum observed error is used for all spectral values. I generated 1000 perturbed spectral instances (as described in previous sections) and performed the inpainting retrieval. It is worth noting that the mass of the planet, the CO mixing ratio, and the CH₄ mixing ratio (imposed equal to 0 due to its non-detection) are outside the training limits of AURA models. This could limit the accuracy and reliability of the reconstruction.

Parameter	Retrieved value
$\log(\text{CH}_4)$	$-8.1^{+2.4}_{-2.5}$
$\log(\text{CO}_2)$	$-5.4^{+1.6}_{-2.3}$
$\log(\text{CO})$	$-6.7^{+1.7}_{-2.0}$
$\log(\text{H}_2\text{O})$	$-3.6^{+0.9}_{-1.1}$
M_p	$1.60^{+0.52}_{-0.55}$
R_p	$1.30^{+0.07}_{-0.07}$
T_p	1370^{+430}_{-390}

Table 5.7: Retrieved parameters from the inpainting of WASP-39 b NIRSpec Prism observational data using the AURA model.

In Table 5.7, I present the retrieved values. All the parameters present large uncertainties. The only exception is the radius, which shows high precision and it is compatible with the value found in the literature (see Table 5.6). Figure 5.6 presents the retrieved spectrum compared to the observed spectral values.

We observe that almost every retrieved spectral value in the NIRSpec Prism passband is compatible with the original observations. Despite this, both regions, inside and outside the observational passband, are poorly recovered and affected by large uncertainties.

Comparison with TauREx retrieval

To compare the result of the inpainting process I performed a classical Bayesian atmospheric retrieval using TauREx. I set up the retrieval's parameters assuming an isothermal temperature profile with an initial temperature guess of 1000 K, allowing it to vary within a reasonable range based on prior expectations from the literature ([Carone et al., 2023b](#); [Espinoza et al., 2024](#)). The pressure structure was defined by a profile spanning from 10^0 Pa to 10^6 Pa, with 100 vertical layers. The molecular opacities were precomputed using cross-section




Figure 5.6: Retrieved spectrum and its error bars compared with the original data observed by [Rustamkulov et al. \(2023\)](#).

data. The model used accounts for molecular absorption, Rayleigh scattering, and, collision-induced absorption (CIA) arising from $\text{H}_2\text{-H}_2$ and $\text{H}_2\text{-He}$ interactions, ensuring a comprehensive treatment of continuum sources alongside the molecular absorption features.

The input transmission spectrum, as already mentioned, was obtained from the JWST NIRSpec PRISM mode observation ([Rustamkulov et al., 2023](#)).

The chemistry of the atmosphere was modeled using a simplified atmospheric structure. The primary gases (H_2 and He) were included with a ratio (H_2/He) of 0.3. Other molecular species such as H_2O , CO_2 , CO, and CH_4 were added with initial estimates based on the expected relative importance as found in the literature ([Rustamkulov et al., 2023](#)). Precise values for the initial guesses of these parameters are presented in Table 5.8. No explicit cloud or haze parametrization was included in this simplified model.

The stellar flux was modeled as a blackbody spectrum, with the host star’s effective temperature fixed at 5485 K. The planetary radius was initially set to $1.279 R_J$, allowing it to vary within the boundaries. Given the low sensitivity of this analysis to planetary mass, it was fixed at $0.281 M_J$. The retrieval was carried out using the Nested Sampling optimizer in TauREx, to ensure a robust exploration of the posterior distributions of the retrieved atmospheric parameters.

Table 5.8 summarizes the retrieved parameters with the associated uncertainties, their initial guess values, and the accepted values boundaries.

Parameter	Initial Guess	Boundaries	Retrieved Value
Planet Radius (R_p) [R_J]	1.279	[1, 1.5]	1.30 ± 0.01
Temperature (T) [K]	1000	[800, 1200]	883^{+58}_{-42}
$\log(\text{H}_2\text{O})$	-2	[-5, -1]	$-3.27^{+0.11}_{-0.13}$
$\log(\text{CO}_2)$	-4	[-5, -3]	$-4.14^{+0.11}_{-0.10}$
$\log(\text{CO})$	-2	[-3, -0.3]	$-0.87^{+0.06}_{-0.05}$
$\log(\text{CH}_4)$	-6	[-12, -4]	$-9.5^{+1.6}_{-1.7}$

Table 5.8: Retrieved atmospheric parameters of WASP-39b with their initial guesses and boundaries. the retrieval was performed using TauREx.

The corner plot showing the posterior distributions with the median values of the fitted parameters is shown in Appendix D. The results obtained for the radius and the chemical abundances are compatible with the literature. The planetary temperature is significantly lower than expected, probably because of the simple temperature profile used to model the data.

In Table 5.9, I present a summary of the values recovered using the two methods.

Parameter	TauREx	Inpainting
Planet Radius (R_p) [R_J]	1.30 ± 0.01	$1.30^{+0.07}_{-0.07}$
Temperature (T_p) [K]	883^{+58}_{-42}	1370^{+430}_{-390}
$\log(\text{H}_2\text{O})$	$-3.27^{+0.11}_{-0.13}$	$-3.6^{+0.9}_{-1.1}$
$\log(\text{CO}_2)$	$-4.14^{+0.11}_{-0.10}$	$-5.4^{+1.6}_{-2.3}$
$\log(\text{CO})$	$-0.87^{+0.06}_{-0.05}$	$-6.7^{+1.7}_{-2.0}$
$\log(\text{CH}_4)$	$-9.5^{+1.6}_{-1.7}$	$-8.1^{+2.4}_{-2.5}$
M_p [M_J]	—	$1.60^{+0.52}_{-0.55}$

Table 5.9: Comparison of retrieved parameters of WASP-39b from two retrieval methods. Dashes indicate parameters not retrieved by that method.

Figure 5.7 shows the distribution of retrieved parameters obtained through the retrieval and inpainting algorithm along with their median values, 16th and 84th percentiles, and expected values (from TauREx retrieval and literature). The CO mixing ratios and the planetary mass, are poorly constrained. This outcome is expected since these parameters have expected values outside the training range of the AURA model. Interestingly, the CH_4 mixing ratio retrieved by the inpainting algorithm is not zero (as expected from the literature) but lies close to the lower boundary of the training range, indicating that no significant CH_4 traces were detected. Moreover, the TauREx retrieved value is compatible with the inpainted value within the associated error bars. All other parameters are consistent with their expected values within the associated uncertainties. In particular, the planetary radius and the H_2O were retrieved with high accuracy.

When interpreting these results, it is important to consider that the presence of parameters outside the model’s training boundaries, combined with the complex, cloud-rich atmospheric structure, strongly affecting the observational data, likely contributes to the poorer quality of some retrieved parameters and spectral values. Further investigations are necessary to assess the robustness and reliability of the retrieval and inpainting algorithm when applied to real observational data. Moreover, future model improvements should consider expanding the training parameter’s boundaries and incorporating a simplified cloud model in the synthetic training data to better resemble real observations.

5.4 Discussion of the results

Although the training phase demonstrates very high reconstruction accuracy for the autoencoder models, this performance does not always directly translate when applying the retrieval and inpainting algorithm to incomplete or partial spectral data. Specifically, the algorithm gives excellent results when the complete spectrum is available (as in the retrieval case) or when a large portion of the spectrum is accessible (such as with the NIRSpec Prism passband, or with the combination of two filters). Under these conditions, the algorithm can effectively exploit the strong correlations between spectral features and planetary parameters learned during training. These results show that the Auto Encoder managed to generalize the atmospheric transmission spectra coherently, despite using only 1% of the ExoGAN dataset.

However, the performance degrades significantly when the algorithm is applied to more limited spectral regions




Figure 5.7: Retrieved parameter distributions obtained from the retrieval and inpainting algorithm are presented along with their median values, 16th and 84th percentiles, and expected values. The expected values for the four mixing ratios were retrieved using TauREx, while the three planetary parameters were taken from the literature.

(for example, the NIRSpec G235 f170lp filter) or to passbands that lack prominent spectral features (such as the MIRI bandpass). The low performance is particularly evident in the retrieval accuracy of the seven planetary and atmospheric parameters. These results suggest that the convolutional autoencoder model, despite its ability to reconstruct full spectra with high fidelity, struggles to extract sufficiently robust correlations between spectral features and physical parameters when limited spectral information is available.

Interestingly, this limitation is somewhat mitigated when using the AURA-BCE-Loss model, which incorporates a penalty term in its loss function designed to discourage non-physical reconstructions. This indicates that adding a constraint to organize the latent space more effectively can improve the reconstruction accuracy of the spectral values. Such a penalty term might guide the model in creating more physically meaningful latent representations. Future improvements could therefore focus on refining this latent space regularization and exploring network architectures that enforce stronger, more independent interactions across different input regions, thereby enhancing generalization when spectral data are sparse or incomplete. A potential improvement in the regularization of the latent space could be achieved by using variational autoencoders (see Appendix B).

Furthermore, when the retrieval and inpainting algorithm is applied to data affected by noise, the results remain consistent with the true values within error bars, but the uncertainties, quantified by the 16th and 84th percentiles, are relatively large. This highlights a critical limitation of the current model when reconstructing

parameters from noisy inputs. To address this, incorporating noise directly to training data, effectively training the autoencoder as a denoising autoencoder, could significantly improve robustness and reduce uncertainties. Given that autoencoders have been successfully employed for denoising tasks in many domains, this approach shows great potential for enhancing retrieval reliability under realistic observational conditions.

Additionally, to better align the training data with real-world scenarios, introducing noise profiles generated by dedicated astrophysical noise simulators, such as the PandExo tool ([Batalha et al., 2017](#)), which models the expected noise from JWST observations, would make the training data more representative of real measurements. This would prepare the models to handle observational noise, improving performance when applied to real telescope data. Moreover, as shown by the WASP-39 b case study, future model improvements should consider expanding the training parameters' boundaries and incorporating a simplified cloud model in the synthetic training data to better resemble real observations.

In conclusion, the retrieval and inpainting algorithm demonstrates promising reconstruction and parameter retrieval performance on synthetic data when a broad spectral window is available. However, its accuracy and reliability diminish significantly with restricted spectral coverage or in the presence of noise. Potential pathways to overcome these limitations include enhancing the network architecture to enforce stronger connections between input regions, integrating latent space regularization through penalization terms to avoid nonphysical reconstructions, and incorporating realistic noise and clouds models during training to improve denoising capabilities. These improvements would facilitate the application of the developed models and algorithms to real exoplanetary spectral observations.

Chapter 6

Conclusions

This thesis has explored the development and application of machine learning techniques, specifically convolutional autoencoders, to the modeling and retrieval of exoplanetary atmospheres from transmission spectroscopy data. In Chapter 2, I described the fundamentals of transmission spectroscopy, detailing the physical processes such as atomic and molecular absorption, scattering mechanisms, and radiative transfer, which govern the formation of transmission spectra in exoplanet atmospheres. The focus was placed on hot Jupiters, their atmospheric properties, and the current observational instruments and data analysis techniques that enable atmospheric characterization.

Chapter 3 provided a comprehensive overview of machine learning fundamentals, including neural networks and convolutional architectures, with a particular emphasis on autoencoders as powerful unsupervised models for feature extraction and dimensionality reduction. I discussed prior machine learning applications in astrophysics, in particular for exoplanet atmospheric retrieval.

In Chapter 4, I presented the construction and training of several convolutional autoencoder architectures designed to efficiently encode transmission spectra and retrieve key planetary and atmospheric parameters. The preprocessing of the dataset, normalization strategies, and the Atmospheric Spectra and Parameters Array (ASPA) representation were described in detail. Among the developed models, AURA demonstrated the best reconstruction accuracy and robustness, while alternative architectures provided useful comparisons.

Chapter 5 introduced the retrieval and inpainting algorithm that leverages the trained autoencoder to reconstruct missing spectral data and infer atmospheric parameters from incomplete observations. This algorithm was tested on synthetic data and applied to WASP-39 b real case study, demonstrating its capabilities and limitations.

The results demonstrate that the AURA convolutional autoencoder model achieves excellent reconstruction accuracy on synthetic transmission spectra and reliably retrieves key atmospheric and planetary parameters when a broad spectral coverage is available. This confirms the model's capacity to learn complex correlations between spectral features and physical parameters. Moreover, the AURA model demonstrated high generalization skills, since it provided coherent results even by using only 1% of the full ExoGAN training set. However, the performance notably decreases when the available spectral data are limited to narrow passbands or regions with few distinctive features, as seen with certain JWST filters. This limitation highlights the challenges in extracting robust parameter correlations under partial spectral conditions.

The alternative AURA-BCE-Loss model, which incorporates a penalty term to discourage nonphysical recon-

structions and promote a more organized latent space, partially alleviates these issues. While this approach improves the physical plausibility of reconstructions, it results in lower parameter accuracy, indicating a trade-off that requires further investigation.

Additionally, the application of the retrieval and inpainting algorithm to noisy synthetic data shows that although the reconstructed parameters remain consistent within uncertainty boundaries, the resulting error intervals are relatively large. This underscores the sensitivity of the current models to noise and the importance of incorporating noise during training.

Overall, these findings suggest that further advancements are needed to improve robustness and reliability for real-world observational scenarios. Promising directions include training the model on the complete ExoGAN training set, refining network architectures to enforce stronger and more independent feature interactions across spectral regions, integrating latent space regularization terms to reduce nonphysical outputs, and incorporating realistic noise models into the training process. Furthermore, as the WASP-39 b case study showed, future model improvements should consider expanding the limits of the training parameters and incorporating a simplified cloud model in synthetic training data to better resemble real observations. Implementing these improvements will be essential to effectively apply convolutional autoencoders for exoplanet atmospheric retrieval in incoming observational missions.

Despite some limitations under restricted spectral coverage and noisy conditions, the results demonstrate the feasibility of using convolutional autoencoders for atmospheric retrieval and spectral inpainting of exoplanet transmission data. These promising outcomes provide a solid foundation for further model refinement and development, which are necessary steps toward applying this methodology to real observational datasets in future exoplanet characterization efforts.

Appendix A

Derivation of the Atmospheric Scale Height and Its Impact on Transit Depth

Atmospheric Scale Height

We begin with the condition of hydrostatic equilibrium, which states that the vertical pressure gradient in a planetary atmosphere balances the downward gravitational force:

$$\frac{dP}{dz} = -\rho g \quad (\text{A.1})$$

where P is pressure, z is altitude, ρ is the mass density, and g is the gravitational acceleration (assumed constant with altitude).

From the ideal gas law:

$$PV = nRT \quad \Rightarrow \quad \frac{n}{V} = \frac{P}{RT} \quad (\text{A.2})$$

where R is the universal gas constant, T is temperature, and μ_{mol} is the mean molecular weight (in kg/mol). This gives the density in terms of pressure:

$$\rho = \frac{n\mu_{mol}}{V} = \frac{P\mu_{mol}}{RT} \quad (\text{A.3})$$

Substituting into Equation [A.1](#):

$$\frac{dP}{dz} = -\frac{P\mu_{mol}g}{RT} \quad (\text{A.4})$$

Rearranging:

$$\frac{dP}{P} = -\frac{\mu_{mol}g}{RT} dz \quad (\text{A.5})$$

Assuming an isothermal atmosphere (constant T), integrating both sides:

$$\ln P = -\frac{\mu_{mol}g}{RT}z + C \quad \Rightarrow \quad P(z) = P_0 e^{-z/H} \quad (\text{A.6})$$

where P_0 is the pressure at the reference level and we define the scale height as:

$$H = \frac{RT}{\mu_{mol}g} \quad (\text{A.7})$$

Alternatively, using Boltzmann's constant k ($k = R/N_A$, where N_A is the Avogadro's number) and expressing the mean molecular weight per particle μ_m ($\mu_{mol} = N_A \mu_m$):

$$H = \frac{kT}{\mu_m g} \quad (\text{A.8})$$

This is the atmospheric scale height, which represents the characteristic vertical distance over which pressure or density decreases by a factor of e .

Transit Depth Contribution from the Atmosphere

The transit depth is defined as the fraction of stellar flux blocked by the planet:

$$\delta = \left(\frac{R_{p,\lambda}}{R_s} \right)^2 \quad (\text{A.9})$$

where $R_{p,\lambda}$ is the apparent planetary radius at wavelength λ . In transmission spectroscopy, the planet appears larger at wavelengths where the atmosphere becomes opaque. The observed radius includes an atmospheric annulus of width approximately $z_\lambda \sim NH$, where $N \sim 5$ is an empirical factor for the altitude where the atmosphere becomes optically thick.

$$R_{p,\lambda} \approx R_p + NH \quad (\text{A.10})$$

Substituting into the transit depth expression:

$$\delta_\lambda = \left(\frac{R_p + NH}{R_s} \right)^2 \quad (\text{A.11})$$

Assuming $NH \ll R_s$, we apply the binomial approximation:

$$\delta_\lambda \approx \left(\frac{R_p}{R_s} \right)^2 + \frac{2R_p NH}{R_s^2} \quad (\text{A.12})$$

The second term represents the incremental contribution from the atmosphere to the total transit depth. Therefore, the approximate contribution of the atmosphere to the transmission signal is:

$$\delta_\lambda \approx 5 \times \frac{2R_p H}{R_s^2} \quad (\text{A.13})$$

This expression highlights how the detectability of atmospheric features in transmission spectroscopy scales with the planetary radius, atmospheric scale height, and stellar radius.

Appendix B

Variational Autoencoders

Variational Autoencoders (VAEs) are a subtype of autoencoders specifically designed for generative purposes. Unlike traditional autoencoders, which encode each input into a single point in the latent space, VAEs encode inputs into a probability distribution over the latent space, typically modeled as a multivariate Gaussian distribution. This means that, instead of producing a deterministic latent vector $\mathbf{h}$, the encoder outputs two vectors: a mean vector $\boldsymbol{\mu}$ and a standard deviation vector $\boldsymbol{\sigma}$, which define a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$ in the latent space.

During training, a latent vector $\mathbf{z}$ is sampled from this distribution to be passed to the decoder. However, since direct sampling is a stochastic operation and does not allow gradients to propagate, VAEs employ the reparameterization trick. This technique expresses the sampling process as a deterministic and differentiable function by rewriting the sample as:

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \cdot \epsilon \quad \text{with} \quad \epsilon \sim \mathcal{N}(0, 1) \quad (\text{B.1})$$

In this way, the randomness is isolated in ϵ , allowing backpropagation through $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ during optimization. The decoder then uses the sampled $\mathbf{z}$ to reconstruct the input data. This stochastic approach encourages the latent space to be continuous and well-structured, which is essential for generating new, meaningful data samples by interpolating or sampling within the latent space.

The loss function of a VAE is the sum of two terms:

- **Reconstruction Loss:** Typically measured using Mean Squared Error (MSE), this term ensures that the decoded output is as close as possible to the original input.
- **Kullback-Leibler (KL) Divergence:** This regularization term measures how much the learned latent distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$ deviates from a prior distribution, usually a standard normal distribution $\mathcal{N}(0, \mathbf{I})$. The KL divergence enforces a smooth and continuous latent space by pushing the encoded distributions towards the prior.

Thanks to this structured latent space, VAEs are powerful generative models capable of producing new data samples similar to those in the training set by sampling from the prior distribution and passing these samples through the decoder. This makes VAEs widely used in tasks such as image generation.

Appendix C

Distribution of the retrieved parameters of the test example 3992




Figure C.1: Parameter distribution obtained from the inpainting of the test example 3992 perturbed with noise, leaving the NIRSpec Prism passband unmasked.

Appendix D

Posterior distributions of WASP-39 b

Figure D.1 present the corner plot showing the posterior distributions with the median values of the fitted parameters. The 0.16, 0.50 (median value), and 0.84 percentile of the posterior distribution are indicated with dashed-lines.




Figure D.1: Posterior distributions of the retrieved parameters.

Bibliography

- Abadi Martín, Agarwal Ashish, Barham Paul, Brevdo Eugene, Chen Zhifeng, Citro Craig, Corrado Greg S., Davis Andy, Dean Jeffrey, Devin Matthieu, Ghemawat Sanjay, Goodfellow Ian, Harp Andrew, Irving Geoffrey, Isard Michael, Jia Yangqing, Jozefowicz Rafal, Kaiser Lukasz, Kudlur Manjunath, Levenberg Josh, Mané Dandelion, Monga Rajat, Moore Sherry, Murray Derek, Olah Chris, Schuster Mike, Shlens Jonathon, Steiner Benoit, Sutskever Ilya, Talwar Kunal, Tucker Paul, Vanhoucke Vincent, Vasudevan Vijay, Viégas Fernanda, Vinyals Oriol, Warden Pete, Wattenberg Martin, Wicke Martin, Yu Yuan, Zheng Xiaoqiang. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015. Software available from tensorflow.org.
- Al-Refaie A. F., Changeat Q., Waldmann I. P., Tinetti G. TauREx 3: A Fast, Dynamic, and Extendable Framework for Retrievals // *The Astrophysical Journal*. VIII 2021. 917, 1. 37.
- Bakos Gáspár Á. The HATNet and HATSouth Exoplanet Surveys // *Handbook of Exoplanets*. Cham: Springer International Publishing, 2018. 957–967.
- Batalha Natasha E., Mandell Avi, Pontoppidan Klaus, Stevenson Kevin B., Lewis Nikole K., Kalirai Jason, Earl Nick, Greene Thomas, Albert Loïc, Nielsen Louise D. PandExo: A Community Tool for Transiting Exoplanet Science with JWST & HST // *Publications of the Astronomical Society of the Pacific*. IV 2017. 129, 976. 064501.
- Beleznyay Maya, Kunitomo M. Exploring the dependence of hot Jupiter occurrence rates on stellar mass with TESS // *Monthly Notices of the Royal Astronomical Society*. VIII 2022. 516, 1. 75–83.
- Benz W., Broeg C., Fortier A., Rando N., Beck T., Beck M., Queloz D., Ehrenreich D., Maxted P. F. L., Isaak K. G., Billot N., Alibert Y., Alonso R., António C., Asquier J., Bandy T., Bárczy T., Barrado D., Barros S. C. C., Baumjohann W., Bekkelien A., Bergomi M., Biondi F., Bonfils X., Borsato L., Brandeker A., Busch M.-D., Cabrera J., Cessa V., Charnoz S., Chazelas B., Collier Cameron A., Corral Van Damme C., Cortes D., Davies M. B., Deleuil M., Deline A., Delrez L., Demangeon O., Demory B. O., Erikson A., Farinato J., Fossati L., Fridlund M., Futyan D., Gandolfi D., Garcia Munoz A., Gillon M., Guterman P., Gutierrez A., Hasiba J., Heng K., Hernandez E., Hoyer S., Kiss L. L., Kovacs Z., Kuntzer T., Laskar J., Etangs A. Lecavelier des, Lendl M., López A., Lora I., Lovis C., Lüftinger T., Magrin D., Malvasio L., Marafatto L., Michaelis H., Miguel D. de, Modrego D., Munari M., Nascimbeni V., Olofsson G., Ottacher H., Ottensamer R., Pagano I., Palacios R., Pallé E., Peter G., Piazza D., Piotto G., Pizarro A., Pollaco D., Ragazzoni R., Ratti F., Rauer H., Ribas I., Rieder M., Rohlfs R., Safa F., Salatti M., Santos N. C., Scandariato G., Ségransan D., Simon A. E., Smith A. M. S., Sordet M., Sousa S. G., Steller M., Szabó G. M., Szoke J., Thomas N., Tschentscher M., Udry S., Van Grootel V., Viotto V., Walter I., Walton N. A., Wildi F., Wolter D. The CHEOPS mission // *Experimental Astronomy*. XI 2020. 51, 1. 109–151.

- Birkmann, S. M. , Ferruit, P. , Giardino, G. , Nielsen, L. D. , García Muñoz, A. , Kendrew, S. , Rauscher, B. J. , Beck, T. L. , Keyes, C. , Valenti, J. A. , Jakobsen, P. , Dorner, B. , Alves de Oliveira, C. , Arribas, S. , Böker, T. , Bunker, A. J. , Charlot, S. , de Marchi, G. , Kumari, N. , López-Caniego, M. , Lützgendorf, N. , Maiolino, R. , Manjavacas, E. , Marston, A. , Moseley, S. H. , Prizkal, N. , Proffitt, C. , Rawle, T. , Rix, H.-W. , te Plate, M. , Sabbi, E. , Sirianni, M. , Willott, C. J. , Zeidler, P. . The Near-Infrared Spectrograph (NIRSpec) on the James Webb Space Telescope - IV. Capabilities and predicted performance for exoplanet characterization // A&A. 2022. 661. A83.*
- Bishop Christopher M.* Pattern Recognition and Machine Learning (Information Science and Statistics). Berlin, Heidelberg: Springer-Verlag, 2006. 100.
- Borucki William J., Koch David, Basri Gibor, Batalha Natalie, Brown Timothy, Caldwell Douglas, Caldwell John, Christensen-Dalsgaard Jørgen, Cochran William D., DeVore Edna, Dunham Edward W., Dupree Andrea K., Gautier Thomas N., Geary John C., Gilliland Ronald, Gould Alan, Howell Steve B., Jenkins Jon M., Kondo Yoji, Latham . Kepler Planet-Detection Mission: Introduction and First Results // Science. II 2010. 327, 5968. 977.*
- Carone Ludmila, Lewis David A., Samra Dominic, Schneider Aaron D., Helling Christiane.* WASP-39b: exo-Saturn with patchy cloud composition, moderate metallicity, and underdepleted S/O. 2023a.
- Carone Ludmila, Lewis David A., Samra Dominic, Schneider Aaron D., Helling Christiane.* WASP-39b: exo-Saturn with patchy cloud composition, moderate metallicity, and underdepleted S/O. 2023b.
- Chintarungruangchai Pattana, Jiang Ing-Guey.* Detecting Exoplanet Transits through Machine-learning Techniques with Convolutional Neural Networks // Publications of the Astronomical Society of the Pacific. may 2019. 131, 1000. 064502.
- Cosentino Rosario, Lovis Christophe, Pepe Francesco, Collier Cameron Andrew, Latham David W., Molinari Emilio, Udry Stephane, Bezawada Naidu, Black Martin, Born Andy, Buchschacher Nicolas, Charbonneau Dave, Figueira Pedro, Fleury Michel, Galli Alberto, Gallie Angus, Gao Xiaofeng, Ghedina Adriano, Gonzalez Carlos, Gonzalez Manuel, Guerra Jose, Henry David, Horne Keith, Hughes Ian, Kelly Dennis, Lodi Marcello, Lunney David, Maire Charles, Mayor Michel, Micela Giusi, Ordway Mark P., Peacock John, Phillips David, Piotto Giampaolo, Pollacco Don, Queloz Didier, Rice Ken, Riverol Carlos, Riverol Luis, San Juan Jose, Sasselov Dimitar, Segransan Damien, Sozzetti Alessandro, Sosnowska Danuta, Stobie Brian, Szentgyorgyi Andrew, Vick Andy, Weber Luc.* Harps-N: the new planet hunter at TNG // Ground-based and Airborne Instrumentation for Astronomy IV. 8446. IX 2012. 84461V. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series).
- Espinoza Néstor, Steinrueck Maria E., Kirk James, MacDonald Ryan J., Savel Arjun B., Arnold Kenneth, Kempton Eliza M. R., Murphy Matthew M., Carone Ludmila, Zamyatina Maria, Lewis David A., Samra Dominic, Kiefer Sven, Rauscher Emily, Christie Duncan, Mayne Nathan, Helling Christiane, Rustamkulov Zafar, Parmentier Vivien, May Erin M., Carter Aarynn L., Zhang Xi, López-Morales Mercedes, Allen Natalie, Blecic Jasmina, Decin Leen, Mancini Luigi, Molaverdikhani Karan, Rackham Benjamin V., Palle Enric, Tsai Shang-Min, Ahrer Eva-Maria, Bean Jacob L., Crossfield Ian J. M., Haegele David, Hébrard Eric, Kreidberg Laura, Powell Diana, Schneider Aaron D., Welbanks Luis, Wheatley Peter, Brahm Rafael, Crouzet Nicolas.* Inhomogeneous terminators on the exoplanet WASP-39 b // Nat. VIII 2024. 632, 8027. 1017–1020.

- Etsebeth V, Lochner M, Walmsley M, Grespan M.* Astronomaly at scale: searching for anomalies amongst 4 million galaxies // *Monthly Notices of the Royal Astronomical Society*. 02 2024. 529, 1. 732–747.
- Faedi F., Barros S. C. C., Anderson D. R., Brown D. J. A., Collier Cameron A., Pollacco D., Boisse I., Hébrard G., Lendl M., Lister T. A., Smalley B., Street R. A., Triaud A. H. M. J., Bento J., Bouchy F., Butters O. W., Enoch B., Haswell C. A., Hellier C., Keenan F. P., Miller G. R. M., Moulds V., Moutou C., Norton A. J., Queloz D., Santerne A., Simpson E. K., Skillen I., Smith A. M. S., Udry S., Watson C. A., West R. G., Wheatley P. J.* WASP-39b: a highly inflated Saturn-mass planet orbiting a late G-type star // *A&A*. VI 2011. 531. A40.
- Foreman-Mackey Daniel, Conley Alex, Meierjurgen Farr Will, Hogg David W., Lang Dustin, Marshall Phil, Price-Whelan Adrian, Sanders Jeremy, Zuntz Joe.* emcee: The MCMC Hammer. III 2013.
- Gardner Jonathan P., Mather John C., Abbott Randy, Abell James S., Abernathy Mark, Abney Faith E., Abraham John G., Abraham Roberto, Abul-Huda Yasin M., Acton Scott, Adams Cynthia K., Adams Evan, Adler David S., Adriaensen Maarten, Aguilar Jonathan Albert, Ahmed .* The James Webb Space Telescope Mission // *Publications of the Astronomical Society of the Pacific*. VI 2023. 135, 1048. 068001.
- George Daniel, Huerta E. A.* Deep Learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data // *Physics Letters B*. III 2018. 778. 64–70.
- Gibson N. P., Aigrain S., Pont F., Sing D. K., Désert J. M., Evans T. M., Henry G., Husnoo N., Knutson H.* Probing the haze in the atmosphere of HD 189733b with Hubble Space Telescope/WFC3 transmission spectroscopy // *MNRAS*. V 2012. 422, 1. 753–760.
- Goodfellow Ian, Bengio Yoshua, Courville Aaron.* Deep Learning. 2016. 100. <http://www.deeplearningbook.org>.
- Gratton Raffaele, Mesa Dino, Bonavita Mariangela, Zurlo Alice, Estay Sebastian, Kervella Pierre, Desidera Silvano, D’Orazi Valentina, Rigliaco Elisabetta.* Jupiter-like planets are common in a low density environment // *Nat*. 02 2023. 14. 14.
- Gressier, A. , Mori, M. , Changeat, Q. , Edwards, B. , Beaulieu, J. P. , Marcq, E. , Charnay, B. .* Near-infrared transmission spectrum of TRAPPIST-1 h using Hubble WFC3 G141 observations // *A&A*. 2022. 658. A133.
- Hassija Vikas, Chamola Vinay, Mahapatra Atmesh, Singal Abhinandan, Goel Divyansh, Huang Kaizhu, Scardapane Simone, Spinelli Indro, Mahmud Mufti, Hussain Amir.* Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence // *Cognitive Computation*. 08 2023. 16.
- Howell Steve B., Sobeck Charlie, Haas Michael, Still Martin, Barclay Thomas, Mullally Fergal, Troeltzsch John, Aigrain Suzanne, Bryson Stephen T., Caldwell Doug, Chaplin William J., Cochran William D., Huber Daniel, Marcy Geoffrey W., Miglio Andrea, Najita Joan R., Smith Marcie, Twicken J. D., Fortney Jonathan J.* The K2 Mission: Characterization and Early Results // *Publications of the Astronomical Society of the Pacific*. IV 2014. 126, 938. 398–408.
- Irwin P.G.J., Teanby N.A., de Kok R., Fletcher L.N., Howett C.J.A., Tsang C.C.C., Wilson C.F., Calcutt S.B., Nixon C.A., Parrish P.D.* The NEMESIS planetary atmosphere radiative transfer and retrieval tool // *Journal of Quantitative Spectroscopy and Radiative Transfer*. 2008. 109, 6. 1136–1150. Spectroscopy and Radiative Transfer in Planetary Atmospheres.

- Khorshid N., Min M., Polman J., Waters L. B. F. M.* Constraining the formation of WASP-39b using JWST transit spectroscopy // *A&A*. V 2024. 685. A64.
- Kreidberg Laura.* batman: BASic Transit Model cAlculation in Python // *PASP*. XI 2015. 127, 957. 1161.
- LeCun Yann.* Generalization and network design strategies // *Connectionism in perspective*. 1989. 19, 143-155. 18.
- Line Michael R., Wolf Aaron S., Zhang Xi, Knutson Heather, Kammer Joshua A., Ellison Elias, Deroo Pieter, Crisp Dave, Yung Yuk L.* A SYSTEMATIC RETRIEVAL ANALYSIS OF SECONDARY ECLIPSE SPECTRA. I. A COMPARISON OF ATMOSPHERIC RETRIEVAL TECHNIQUES // *The Astrophysical Journal*. IX 2013. 775, 2. 137.
- Liou K.N.* An Introduction to Atmospheric Radiation. 2002. 123. (International Geophysics).
- Lothringer Joshua D., Barman Travis, Koskinen Tommi.* Extremely Irradiated Hot Jupiters: Non-oxide Inversions, H— Opacity, and Thermal Dissociation of Molecules // *The Astrophysical Journal*. oct 2018. 866, 1. 27.
- Mancini L., Esposito M., Covino E., Southworth J., Biazzo K., Bruni I., Ciceri S., Evans D., Lanza A. F., Poretti E., Sarkis P., Smith A. M. S., Brogi M., Affer L., Benatti S., Bignamini A., Boccato C., Bonomo A. S., Borsa F., Carleo I., Claudi R., Cosentino R., Damasso M., Desidera S., Giacobbe P., González-Álvarez E., Gratton R., Harutyunyan A., Leto G., Maggio A., Malavolta L., Maldonado J., Martinez-Fiorenciano A., Masiero S., Micela G., Molinari E., Nascimbeni V., Pagano I., Pedani M., Piotto G., Rainer M., Scandariato G., Smareglia R., Sozzetti A., Andreuzzi G., Henning Th.* The GAPS programme with HARPS-N at TNG. XVI. Measurement of the Rossiter-McLaughlin effect of transiting planetary systems HAT-P-3, HAT-P-12, HAT-P-22, WASP-39, and WASP-60 // *A&A*. V 2018. 613. A41.
- Marconi Alessandro, Abreu Manuel, Adibekyan Vardan, Aliverti Matteo, Allende Prieto Carlos, Amado Pedro, Amate Manuel, Artigau Etienne, Augusto Sergio, Barros Susana, Becerril Santiago, Benneke Björn, Bergin Edwin, Berio Philippe, Bezawada Naidu, Boisse Isabelle, Bonfils Xavier, Bouchy François, Broeg Christopher, Cabral Alexandre, Calvo-Ortega Rocio, Canto Martins Bruno Leonardo, Chazelas Bruno, Chiavassa Andrea, Christensen Lise, Cirami Roberto, Coretti Igor, Covino Stefano, Cresci Giovanni, Cristiani Stefano, Cunha Parro Vanderlei, Cupani Guido, De Castro Leão Izan, Renan De Medeiros José, Furlande Souza Marco Antonio, Di Marcantonio Paolo, Di Varano Igor, D’Odorico Valentina, Doyon René, Drass Holger, Figueira Pedro, Belen Fragoso Ana, Uldall Fynbo Johan Peter, Gallo Elena, Genoni Matteo, González Hernández Jonay, Haehnelt Martin, Hlavacek-Larrondo Julie, Hughes Ian, Huke Philipp, Humphrey Andrew, Kjeldsen Hans, Korn Andreas, Kouach Driss, Landoni Marco, Liske Jochen, Lovis Christophe, Lunney David, Maiolino Roberto, Malo Lison, Marquart Thomas, Martins Carlos, Mason Elena, Molaro Paolo, Monnier John, Monteiro Manuel, Mordasini Christoph, Morris Tim, Mucciarelli Alessio, Murray Graham, Niedzielski Andrzej, Nunes Nelson, Oliva Ernesto, Origlia Livia, Pallé Enric, Pariani Giorgio, Parr-Burman Phil, Peñate José, Pepe Francesco, Pinna Enrico, Piskunov Nikolai, Rasilla Piñeiro José Luis, Rebolo Rafael, Rees Phil, Reiners Ansgar, Riva Marco, Romano Donatella, Rousseau Sylvain, Sanna Nicoletta, Santos Nuno, Sarajlic Mirsad, Shen Tzu-Chiang, Sortino Francesca, Sosnowska Danuta, Sousa Sérgio, Stempels Eric, Strassmeier Klaus, Tenegi Fabio, Tozzi Andrea, Udry Stéphane, Valenziano Luca, Vanzi Leonardo, Weber Michael, Woche Manfred, Xompero Marco, Zackrisson Erik, Zapatero Osorio María Rosa.* HIRES, the High-resolution Spectrograph for the ELT // *Published in The Messenger* vol. 182. 2021. pp. 27-32. March 2021.

- Min, Michiel , Ormel, Chris W. , Chubb, Katy , Helling, Christiane , Kawashima, Yui .* The ARCIS framework for exoplanet atmospheres - Modelling philosophy and retrieval // *A&A*. 2020. 642. A28.
- Mitchell Tom M.* Machine learning. 1, 9. 1997. 100.
- Mollière, P. , Wardenier, J. P. , van Boekel, R. , Henning, Th. , Molaverdikhani, K. , Snellen, I. A. G. .* petitRADTRANS - A Python radiative transfer package for exoplanet characterization and retrieval // *A&A*. 2019. 627. A67.
- Murray-Clay Ruth A., Chiang Eugene I., Murray Norman.* ATMOSPHERIC ESCAPE FROM HOT JUPITERS // *The Astrophysical Journal*. II 2009. 693, 1. 23–42.
- Padovani Paolo, Cirasuolo Michele.* The Extremely Large Telescope // *Contemporary Physics*. I 2023. 64, 1. 47–64.
- Pepe F., Cristiani S., Rebolo R., Santos N. C., Dekker H., Cabral A., Di Marcantonio P., Figueira P., Lo Curto G., Lovis C., Mayor M., Mégevand D., Molaro P., Riva M., Zapatero Osorio M. R., Amate M., Manescau A., Pasquini L., Zerbi F. M., Adibekyan V., Abreu M., Affolter M., Alibert Y., Aliverti M., Allart R., Allende Prieto C., Álvarez D., Alves D., Avila G., Baldini V., Bandy T., Barros S. C. C., Benz W., Bianco A., Borsa F., Bourrier V., Bouchy F., Broeg C., Calderone G., Cirami R., Coelho J., Conconi P., Coretti I., Cumani C., Cupani G., D’Odorico V., Damasso M., Deiries S., Delabre B., Demangeon O. D. S., Dumusque X., Ehrenreich D., Faria J. P., Fragoso A., Genolet L., Genoni M., Génova Santos R., González Hernández J. I., Hughes I., Iwert O., Kerber F., Knudstrup J., Landoni M., Lavie B., Lillo-Box J., Lizon J. L., Maire C., Martins C. J. A. P., Mehner A., Micela G., Modigliani A., Monteiro M. A., Monteiro M. J. P. F. G., Moschetti M., Murphy M. T., Nunes N., Oggioni L., Oliveira A., Oshagh M., Pallé E., Pariani G., Poretti E., Rasilla J. L., Rebordão J., Redaelli E. M., Santana Tschudi S., Santin P., Santos P., Ségransan D., Schmidt T. M., Segovia A., Sosnowska D., Sozzetti A., Sousa S. G., Spanò P., Suárez Mascareño A., Tabernero H., Tenegi F., Udry S., Zanutta A.* ESPRESSO at VLT. On-sky performance and first results // *A&A*. I 2021. 645. A96.
- Pepe Francesco, Mayor Michel, Delabre Bernard, Kohler Dominique, Lacroix Daniel, Queloz Didier, Udry Stephane, Benz Willy, Bertaux Jean-Loup, Sivan Jean-Pierre.* HARPS: a new high-resolution spectrograph for the search of extrasolar planets // *Optical and IR Telescope Instrumentation and Detectors*. 4008. VIII 2000. 582–592. (Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series).
- Perryman Michael.* The Exoplanet Handbook. 2018. 123.
- Pollacco D. L., Skillen I., Collier Cameron A., Christian D. J., Hellier C., Irwin J., Lister T. A., Street R. A., West R. G., Anderson D. R., Clarkson W. I., Deeg H., Enoch B., Evans A., Fitzsimmons A., Haswell C. A., Hodgkin S., Horne K., Kane S. R., Keenan F. P., Maxted P. F. L., Norton A. J., Osborne J., Parley N. R., Ryans R. S. I., Smalley B., Wheatley P. J., Wilson D. M.* The WASP Project and the SuperWASP Cameras // *PASP*. X 2006. 118, 848. 1407–1418.
- Rajpurohit, A. S. , Allard, F. , Homeier, D. , Mousis, O. , Rajpurohit, S. .* Dayside thermal inversion in the atmosphere of WASP-19b // *A&A*. 2020. 642. A39.
- Rauer Heike, Aerts Conny, Cabrera Juan, Deleuil Magali, Erikson Anders, Laurent , Piotto Giampaolo, Pollacco Don, Ragazzoni Roberto.* The PLATO Mission. 2024.
- Raymond Sean N., Morbidelli Alessandro.* Planet Formation: Key Mechanisms and Global Models // *Demographics of Exoplanetary Systems*. 2022. 3–82.

Ruder Sebastian. An overview of gradient descent optimization algorithms. 2017.

Rustamkulov Z., Sing D. K., Mukherjee S., May E. M., Kirk J., Schlawin E., Line M. R., Piaulet C., Carter A. L., Batalha N. E., Goyal J. M., López-Morales M., Lothringer J. D., MacDonald R. J., Moran S. E., Stevenson K. B., Wakeford H. R., Espinoza N., Bean J. L., Batalha N. M., Benneke B., Berta-Thompson Z. K., Crossfield I. J. M., Gao P., Kreidberg L., Powell D. K., Cubillos P. E., Gibson N. P., Leconte J., Molaverdikhani K., Nikolov N. K., Parmentier V., Roy P., Taylor J., Turner J. D., Wheatley P. J., Aggarwal K., Ahrer E., Alam M. K., Alderson L., Allen N. H., Banerjee A., Barat S., Barrado D., Barstow J. K., Bell T. J., Blecic J., Brande J., Casewell S., Changeat Q., Chubb K. L., Crouzet N., Daylan T., Decin L., Désert J., Mikal-Evans T., Feinstein A. D., Flagg L., Fortney J. J., Harrington J., Heng K., Hong Y., Hu R., Iro N., Kataria T., Kempton E. M. R., Krick J., Lendl M., Lillo-Box J., Louca A., Lustig-Yaeger J., Mancini L., Mansfield M., Mayne N. J., Miguel Y., Morello G., Ohno K., Palle E., Petit dit de la Roche D. J. M., Rackham B. V., Radica M., Ramos-Rosado L., Redfield S., Rogers L. K., Shkolnik E. L., Southworth J., Teske J., Tremblin P., Tucker G. S., Venot O., Waalkes W. C., Welbanks L., Zhang X., Zieba S. Early Release Science of the exoplanet WASP-39b with JWST NIRSpec PRISM // *Nat. II* 2023. 614, 7949. 659–663.

Schwieterman Edward W., Leung Michaela. An Overview of Exoplanet Biosignatures. 2024.

Seager S., Deming D., Valenti J. A. Transiting Exoplanets with JWST // *Astrophysics in the Next Decade*. 10. I 2009. 123. (Astrophysics and Space Science Proceedings).

Selvaraju Ramprasaath R., Cogswell Michael, Das Abhishek, Vedantam Ramakrishna, Parikh Devi, Batra Dhruv. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization // *International Journal of Computer Vision*. X 2019. 128, 2. 336–359.

Skilling John. Nested Sampling // *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. 735. XI 2004. 395–405. (American Institute of Physics Conference Series).

Sudarsky David, Burrows Adam, Hubeny Ivan. Theoretical Spectra and Atmospheres of Extrasolar Giant Planets // *ApJ*. V 2003. 588, 2. 1121–1148.

Tadaki Ken-ichi, Iye Masanori, Fukumoto Hideya, Hayashi Masao, Rusu Cristian E, Shimakawa Rhythm, Tosaki Tomoka. Spin parity of spiral galaxies II: a catalogue of 80k spiral galaxies using big data from the Subaru Hyper Suprime-Cam survey and deep learning // *Monthly Notices of the Royal Astronomical Society*. 07 2020. 496, 4. 4276–4286.

Teyssandier Jean, Lai Dong, Vick Michelle. Formation of hot Jupiters through secular chaos and dynamical tides // *Monthly Notices of the Royal Astronomical Society*. IV 2019. 486, 2. 2265–2280.

Tinetti Giovanna, Eccleston Paul, Haswell Carole, Lagage Pierre-Olivier, Leconte Jérémy, Lüftinger Theresa, Micela Giusi, Min Michel, Pilbratt Göran, Puig Ludovic, Swain Mark, Testi Leonardo, Turrini Diego. Ariel: Enabling planetary science across light-years. 2021.

Villar V. Ashley, Hosseinzadeh Griffin, Berger Edo, Ntampaka Michelle, Jones David O., Challis Peter, Chornock Ryan, Drout Maria R., Foley Ryan J., Kirshner Robert P., Lunnan Ragnhild, Margutti Raffaella, Milisavljevic Dan, Sanders Nathan, Pan Yen-Chen, Rest Armin, Scolnic Daniel M., Magnier Eugene, Metcalfe Nigel, Wainscoat Richard, Waters Christopher. SuperRAENN: A Semisupervised Supernova Photo-

metric Classification Pipeline Trained on Pan-STARRS1 Medium-Deep Survey Supernovae // The Astrophysical Journal. dec 2020. 905, 2. 94.

Winn Joshua N. Transits and Occultations. 2014.

Zeraatgari Fatemeh Zahra, Hafezianzadeh Fatemeh, Zhang Yanxia, Mei Liqun, Ayubinia Ashraf, Mosallanezhad Amin, Zhang Jingyi. Machine learning-based photometric classification of galaxies, quasars, emission-line galaxies, and stars // Monthly Notices of the Royal Astronomical Society. 11 2023. 527, 3. 4677–4689.

Zhou Li, Ma Bo, Wang Yonghao, Zhu Yinan. Hubble WFC3 Spectroscopy of the Rocky Planet L 98–59 b: No Evidence for a Cloud-free Primordial Atmosphere // The Astronomical Journal. X 2022. 164, 5. 203.

Zingales Tiziano, Waldmann Ingo P. ExoGAN: Retrieving Exoplanetary Atmospheres Using Deep Convolutional Generative Adversarial Networks // The Astronomical Journal. XI 2018. 156, 6. 268.