



Exoplanetary atmospheres retrieval via a quantum extreme learning machine

Marco Vetrano¹ · Tiziano Zingales^{2,3} · G. Massimo Palma¹ · Salvatore Lorenzo¹

Received: 22 October 2025 / Accepted: 14 June 2026
© The Author(s) 2026

Abstract

The study of exoplanetary atmospheres traditionally relies on forward models to analytically compute the spectrum of an exoplanet by fine-tuning numerous chemical and physical parameters. However, the high-dimensionality of parameter space often results in a significant computational overhead. In this work, we introduce a novel approach to atmospheric retrieval leveraging on quantum extreme learning machines (QELMs), a quantum version of the classical Extreme Learning Machine (ELM) – a fast machine learning model typically used for regression and classification. Our method combines classical spectral patching and PCA-based dimensionality reduction with a factorized quantum reservoir that provides nonlinear features for the final linear retrieval map. We distinguish the classical PCA filtering used to mitigate noise in the input spectra from the readout-level robustness observed when the reservoir is executed on noisy quantum hardware. We demonstrate the robustness of our approach through a direct implementation on IBM Fez. The proposed QELM architecture highlights the potential of quantum computing in the analysis of astrophysical datasets, retrieving successfully the concentration of CH_4 , CO_2 , H_2O and the radius of over the 90% of the dataset in the infinite statistics limit, while remaining robust under realistic noise conditions on IBM Fez, paving the way, in the near future, to faster, more efficient, and more accurate models for the study of exoplanetary atmospheres.

Keywords Quantum Machine Learning · Exoplanets · Atmospheric Retrieval · Experimental Quantum Computation

Tiziano Zingales, G. Massimo Palma and Salvatore Lorenzo contributed equally to this work.

✉ Marco Vetrano
marco.vetrano02@unipa.it
Tiziano Zingales
tiziano.zingales@unipd.it
G. Massimo Palma
massimo.palma@unipa.it
Salvatore Lorenzo
salvatore.lorenzo@unipa.it

- ¹ Dipartimento di Fisica e Chimica, Università degli Studi di Palermo, Via Archirafi, 36, Palermo I-9012, Italy
- ² Dipartimento di Fisica e Astronomia “Galileo Galilei”, Università degli Studi di Padova, Vicolo dell’Osservatorio, 3, Padova 35122, Italy
- ³ INAF Osservatorio Astronomico di Padova, Vicolo dell’Osservatorio, 5, Padova 35122, Italy

1 Introduction

The study of exoplanetary atmospheres represents one of the most fascinating astrophysical challenges. The interpretation of atmospheric features and the analysis of absorption and emission spectra (Kreidberg et al. 2018; Tsiaras et al. 2018; Bruno et al. 2018; Mansfield et al. 2018; Spake et al. 2018; Sheppard et al. 2017; Barstow et al. 2016; Rocchetto et al. 2016) can be performed thanks to the mathematical modeling of the exoplanet atmosphere. Generation of a forward model is crucial to simulate expected spectra from a given atmospheric composition, temperature profiles, and physical conditions. In atmospheric retrievals, forward models are compared with observed spectra to retrieve atmospheric parameters (Irwin et al. 2008; Madhusudhan and Seager 2009; Line et al. 2013; Benneke and Seager 2013; Cubillos et al. 2020; Gandhi and Madhusudhan 2018; Lavie et al. 2017). Models may be very different in their complexity, ranging from simply parameterised descriptions to complex physical simulations, usually at the price of a large increase in computing resources due to the increased precision.

Finding the right balance between computational cost and model accuracy in exoplanetary science has been a challenge. Typically, simple models enable a fast analysis at the cost of oversimplifying the atmospheric processes involved and provide biased results or miss subtle features that could be present in the data (Pluriel et al. 2020, 2022). More complex models can capture intricate atmospheric phenomena at a higher computational cost as they usually require advanced statistical tools to estimate atmospheric parameters (Feroz and Hobson 2008; Skilling 2004; Feroz et al. 2009; Gregory 2011).

In the context of atmospheric retrieval, several machine learning techniques have shown the ability to make both qualitative and quantitative identifications of the spectral type (Waldmann 2016; Zingales and Waldmann 2018; Márquez-Neila et al. 2018; Cobb et al. 2019). These algorithms allow for fast and optimal parameter estimation, significantly improving computational time over classic retrieval methods. Furthermore, machine learning tools can be designed to return, depending on the particular problem, different output types such as posterior distributions for statistical inference, generated spectra for validation against observation, and straightforward prediction of atmospheric parameters such as temperature profiles, molecular abundance, and clouds.

One crucial aspect of reliable machine learning algorithms, are the high-quality training datasets. Such datasets require either extensive simulations or extensive observation campaigns, both options being very resource intensive. In addition to this, most machine learning models cannot make good extrapolations outside the parameter space spanned by their training data and may become inaccurate for new or complex atmospheric features. For this reason, it is important to generate robust datasets to train future algorithms, reproducing all those features that could be observed with space-based and ground-based instruments.

State-of-the-art retrieval methods are promising in exploiting low-resolution spectroscopy, given the capabilities of missions like JWST (Gardner 2007) and the upcoming Ariel space mission (Tinetti et al. 2016). The broad wavelength coverage and high sensitivity of JWST combined with the extensive survey of various exoplanets by Ariel yield a rich set of low-resolution spectra. In the same respect, the increased complexity of the model needed to interpret these observations underscores that the refinement of machine learning tools must go hand in hand with new challenges related to precision and scalability.

In the context of deep learning, Extreme Learning Machines (ELMs) (Konkoli 2017; Huang et al. 2004, 2011; Wang et al. 2022; Markowska-Kaczmar and Kosturek 2021) have emerged as a promising alternative to traditional approaches due to their highly efficient training protocol.

An ELM is a neural network model in which only the final layer is trained, while all other layers, generally referred to as “reservoir”, are randomly initialized and kept fixed. Despite this randomness, its inherent complexity and non-linearity enable it to extract meaningful, non-trivial features from data. ELMs have been successfully implemented both with randomly generated neural networks (Huang et al. 2006) and with a variety of physical systems (Soriano et al. 2015; Bhovad and Li 2021; Nakajima et al. 2015; Coulombe et al. 2017; Goudarzi et al. 2013; Tanaka et al. 2019; Nokkala et al. 2022).

With the advent of Quantum Machine Learning (QML) (Biamonte et al. 2017; Schuld et al. 2015; Cerezo et al. 2022; Banchi et al. 2021) ELMs have been extended to quantum reservoirs, leading to the development of Quantum Extreme Learning Machines (QELMs). While deep learning retrieval models such as (Waldmann 2016; Zingales and Waldmann 2018; Márquez-Neila et al. 2018; Cobb et al. 2019) achieve fast parameter estimation, they still rely on large classical datasets and are computationally demanding to train. QELMs, in contrast, leverage quantum reservoirs to perform the same mapping more compactly and with potentially lower resource requirements. Such models have demonstrated their effectiveness in solving quantum tasks, with applications on state estimations, entanglement witnessing and quantum state preparation (Vetrano et al. 2025; Suprano et al. 2024; Ghosh et al. 2019a, 2021a, b; Krisnanda et al. 2021; Lo Monaco et al. 2024; Zia et al. 2025), but also on classical tasks, for time-series analysis and image classifications (De Lorenzis et al. 2024, 2025; Mujal et al. 2021, 2023; Govia et al. 2021; Martínez-Peña and Ortega 2023; Ghosh et al. 2019b; Martínez-Peña et al. 2021; García-Beni et al. 2023; Martínez-Peña et al. 2023), outperforming their classical counterpart in terms of computational resources (Mujal et al. 2021; Fujii and Nakajima 2017; Xiong et al. 2025; Nakajima et al. 2019).

However, processing high-dimensional data with QML remains a challenge from both the simulation and the experimental point of view. In fact, while classical simulations of many-body systems are very expensive in terms of computational resources, real quantum devices still suffer from decoherence noise caused by unwanted interactions with the environment (Preskill 2018; Lau et al. 2022). In some cases, the challenges posed by high-dimensionality can be addressed by patching data arrays to reduce the dimensions of the task, effectively breaking it into subtasks (Huang et al. 2021a; Tsang et al. 2023).

Inspired by the architecture used in Huang et al. (2021a), in this work, we construct a QELM to extract information from synthetic spectra generated by TauREx (Al-Refaie et al. 2021; Waldmann et al. 2015), focusing on minimizing quantum resources while maintaining resilience to noise in

near-term implementations. In this article, we describe the first QELM applied to an astrophysical analysis problem. To test QELM's robustness, the algorithm is evaluated on datasets within the JWST spectral range, both with and without added artificial shot noise, to assess its response to realistic data imperfections. Additionally, robustness is proved by running the JWST dataset on the IBM Fez quantum hardware, confirming the performances and so the potential of QELM for NISQ applications.

The article is organized as follows: in Section 2 we briefly review ELM and QELM models; in Section 3 we give a detailed explanation of the structure of the datasets employed in this work, its pre-processing and the implementation of our QELM framework. In Section 4 we report the results obtained by implementing the algorithm on IBM Fez showing its robustness against noise. Finally in Section 5 we draw our conclusions.

2 QELMs in a nutshell

ELM - An ELM is a fast, efficient machine learning algorithm (Goodfellow et al. 2016) that uses single-hidden-layer feedforward neural networks with randomly assigned weights and biases, called “**reservoir**”, requiring only the output weights to be learned through a final linear post-processing step.

The reservoir can be described as a complex function $\mathcal{R} : \mathbb{R}^{D_{\text{in}}} \rightarrow \mathbb{R}^{D_{\text{out}}}$ mapping the input data in a higher dimensional space. Thus the complete function applied by an ELM can be written as the composition $f_W^{\text{ELM}} = W \circ \mathcal{R}$, where W is the trained linear map applied by the output layer.

Given the fixed and untrained nature of the reservoir function $\mathcal{R}$, this approach simplifies the learning process to solving a linear regression problem, eliminating the need for classical optimization algorithms to fine-tune the entire network's parameters. As a result, it achieves generalization much more efficiently.

ELMs belong to the class of supervised learning models. Given a training dataset $\{S^{\text{train}}, Y^{\text{train}}\}$, the training process consists in finding the best W which approximates

$$f_W^{\text{ELM}}(S^{\text{train}}) \equiv W\mathcal{R}(S^{\text{train}}) = Y^{\text{train}}, \quad (1)$$

where $S^{\text{train}} \equiv \{s_k^{\text{train}}\}$ is the ensemble of training input data, $\mathcal{R}(S^{\text{train}})$ is the matrix whose k th-column corresponds to the outcomes of the reservoir relative to the input s_k^{train} and Y^{train} is the matrix containing the target associated to each input data. More specifically, $\mathcal{R}(S^{\text{train}})$ results to be a $D_{\text{out}} \times D_{\text{train}}$ matrix, while Y^{train} is $D_{\text{feat}} \times D_{\text{train}}$. We

indicate with D_{out} the number of reservoir outcomes, with D_{train} the number of training samples and with D_{feat} the number of target features. Having in mind an implementation with a physical reservoir, the input data S^{train} can be encoded in a perturbation of the reservoir state, and the response of the reservoir is described by $\mathcal{R}(S^{\text{train}})$.

A common technique to find the map W in Equation 1 is to use the Moore-Penrose pseudo-inverse:

$$W = Y^{\text{train}}\mathcal{R}(S^{\text{train}})^+, \quad (2)$$

where $\mathcal{R}(S^{\text{train}})^+$ is computed via the Singular Value Decomposition (SVD) $\mathcal{R}(S^{\text{train}})^+ = U\Sigma^+V^\dagger$, with U , V isometries and Σ^+ a positive diagonal matrix containing the inverse of the singular values of $\mathcal{R}(S^{\text{train}})$ (Serre 2010).

QELM - The transition from classical to quantum extreme learning machines is accomplished by replacing the classical reservoir with a quantum system. In this framework, the reservoir function $\mathcal{R}$ is replaced by a quantum dynamical map Φ followed by a measurement step (Vetrano et al. 2025; Innocenti et al. 2023; Mujal et al. 2021).

Furthermore, whenever the input data are not purely quantum, they need to be encoded in a quantum state in order to be processed by the reservoir (Mujal et al. 2021, 2023, 2021; De Lorenzis et al. 2024; Mujal et al. 2023), giving us a dataset of training states ($s_k^{\text{train}} \rightarrow \rho_k^{\text{train}}$). If we assume that the measurement step is performed evaluating the expectation values of a Positive Operator Valued Measure (POVM) Π_i , we have the following correspondence:


$$\mathcal{R}(S^{\text{train}}) \rightarrow P_{ik}^{\text{train}} = \text{Tr}[\Pi_i \Phi[\rho_k^{\text{train}}]] \quad (3)$$

where P^{train} is the probability matrix relative to the training states. The dimension of P entirely depends on the number of measured operators on the reservoir and on the dimension of the training dataset, i.e. $D_{\text{out}} \times D_{\text{train}}$. In an ideal noiseless scenario, the linearity of the quantum map Φ together with the post-processing layer makes QELMs naturally resistant to overfitting, since their limited capacity prevents them from memorizing noise in the training data. A pictorial representation of QELM is shown in Fig. 1.

Finally, to evaluate the system's generalization capabilities, we use a testing dataset along with its associated probability matrix P^{test} . The performance is assessed using a task-specific metric; for this work, we employ the mean squared relative error (MSRE) between the reconstructed and testing features (respectively y^{pred} , y^{test}):

$$\epsilon = \frac{1}{D_{\text{test}}} \sum_k \frac{(y_k^{\text{test}} - y_k^{\text{pred}})^2}{(y_k^{\text{test}})^2} = \frac{1}{D_{\text{test}}} \sum_k \epsilon_k. \quad (4)$$

Fig. 1 Pictorial representation of how classical data analysis with QELM- Classical data are encoded in a physical system from which we collect a set of measurements. The outcomes undergo a linear post-processing which maps the outcomes of the reservoir to the output of task



Furthermore, in order to have a more precise idea of how many retrievals we are doing with a reasonable accuracy, we define a score considering the function $A(\epsilon^*)$, counting all the parameters retrieved with an MSRE below a threshold ϵ^*

$$A(\epsilon^*) = \frac{100}{D_{\text{test}}} \sum_k^{D_{\text{test}}} \delta(\epsilon_k; \epsilon^*) \quad (5)$$

$$\delta(\epsilon_k, \epsilon^*) = \begin{cases} 1, & \text{if } \epsilon_k \leq \epsilon^*, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

We observed that the accuracy has a weak linear dependence on ϵ^* , thus from now on it will be set to $\epsilon^* = 5\%$ which we argue to be reasonably low.

3 Methodology

In this section, we review the structure of the training dataset, its pre-processing and the working of the QELM.

3.1 Dataset

As mentioned the aim is to build a quantum extreme learning framework able to make atmospheric retrieval of exoplanets. In particular, we are going to process datasets of spectra generated by TauREx (Al-Refaie et al. 2021; Waldmann et al. 2015). TauREx is an open-source Bayesian code for exoplanetary atmospheric modeling and retrieval which allows for the generation and retrieval of transmission, emission and reflection spectra over a wide range of compositions and physical parameters. TauREx computes radiative transfer on exoplanetary atmospheres taking into account different chemical and physical conditions to produce the resulting spectrum. Its object-oriented structure allows the

inclusion of many different atmospheric scenarios and it is widely used in the exoplanetary community.

The dataset we used in this context is the same used in Zingales and Waldmann (2018). It depends in particular on 4 parameters for the molecular volume mixing ratios (CH_4 , CO_2 , CO , H_2O) and 3 physical parameters M , R , T , respectively, planetary mass, radius and equilibrium temperature. We assume a constant vertical profile for each volume mixing ratio and isothermal temperature-pressure profile. Each parameter has its lower and upper bounds and the intervals are divided in 10 discrete values. Each spectrum is generated in the spectral range $[0.3, 50]\mu m$, divided in 515 spectral bins, considering a combination of randomly selected values for each parameter. Thus we have 10^7 spectra available for training and testing. In this work we process a sampled subset of 10^4 spectra for both the training and the testing stage. All the details regarding the lower and upper bound for each of the atmospheric parameters are summarized in Table 1. This dataset will be labeled as **TauREx**.

In order to test QELM in a more realistic scenario, we considered a spectrum of HAT-P-18b (Fu et al. 2022), measured by JWST and processed the **TauREx** spectra by interpolating the same spectral binning of JWST in the spectral

Table 1 Spectral parameters bounds - The spectra considered in this work have been produced considering atmospheric parameters in the upper intervals. Each interval has been divided in 10 values. The spectra are produced considering all the 10^7 possible combination of these values

Variable	Lower Bound	Upper Bound
CH_4	10^{-8}	10^{-1}
CO_2	10^{-8}	10^{-1}
CO	10^{-8}	10^{-1}
H_2O	10^{-8}	10^{-1}
M	$0.8M_J$	$2.0M_J$
R	$0.8R_J$	$1.5R_J$
T	$1000K$	$2000K$

range $[0.6, 2.8]\mu\text{m}$. This second dataset will be labeled as **JWST**.

The **JWST** spectra have also been processed considering shot noise obtained by a source of photons at $T_\star=6460\text{K}$. The shot noise has been computed as shown in Cowan et al. (2015). This dataset will be labeled as Noisy JWST (**NJWST**).

Finally, we considered one last dataset by filtering out the shot noise with a pre-processing shown in the next subsection and we are going to label it as Filtered JWST (**FJWST**). In Fig. 2 we show an example of all the previously described data.

3.2 Data pre-processing

The primary challenge in quantum machine learning lies in efficiently encoding high-dimensional data into a Hilbert space of limited dimension. The choice of encoding will impact on the quantum resources needed to solve a given task. Furthermore in the context of Near Intermediate Scale Quantum (NISQ) devices, solving a task with too many quantum resources, for instance deep quantum circuits, results in very noisy outputs which do not allow for the implementation on real near-term quantum devices. In our context, by “quantum resources” we mean the number of qubits and the number of quantum operations involved. Therefore, when handling high dimensional data, the common approach consists in reducing their dimension with either convolutional neural networks or principal component analysis (PCA) (De Lorenzis et al. 2024; Huang et al.

2021b; Rebentrost et al. 2014; Senokosov et al. 2024) to extract the most relevant and discernible features of a dataset. For instance, PCA is a linear transformation which consists in a change of coordinates, with the aim to find the frame in which the variance of the data is maximized.

On the other hand, in the context of generative quantum machine learning such as Quantum Generative Adversarial Networks (qGANs) (Huang et al. 2021a), the number of qubits determines also the resolution (number of pixels) of the images which the qGAN is able to produce. Thus in this context it is not possible to minimize the number of qubits; nonetheless, it is still possible to limit the depth of the quantum circuits (number of gates) by factoring the task into a set of sub-tasks. This reduces the long range interactions between qubits and thus the number of quantum operations applied in a circuit.

To handle the spectra, in this work we decided to use a hybrid approach which combines both the dimensionality reduction and the organization of the task into sub-tasks. In particular, we divided the spectral range according to the major water bands in the IR while also taking into account the instrument pass-bands of JWST/NIRISS, NIRcam, MIRI and Hubble/WFC3, for a total of $N_{\text{TauREx}}=14$ spectral bands, which from now on we are going to address as “patches”.

The number of patches has been reduced to $N_{\text{JWST}}=8$ when considering only the JWST/NIRISS spectral range. Beyond the need to factorize the problem, this operation has also been performed in order to normalize the patches one by one and amplify their spectral features.




Fig. 2 Spectra - Example of a simulated spectrum produced by TauREx (a), interpolated on JWST spectral range (b), with the addition of shot noise (c) and with the addition of a principal component filter

taking 10 components in account (d). As can be seen in the spectra, adding the filter the presence of noise still changes a bit the spectrum with respect to its clean counterpart in b

In the **JWST** case, we also test the model considering the artificial noise added to the dataset. The shot noise is computed as in Cowan et al. (2015).

We assumed that measurement noise is dominated by quantum detection noise following Poisson statistics with equal mean and variance between ~ 0.6 and $2.8 \mu\text{m}$. We estimated the mean number of collected photoelectrons to be:

$$N_{\text{ph}} = \frac{\pi\tau\Delta t}{hc} \left(\frac{R_* D}{2d} \right)^2 \int_{\lambda_1}^{\lambda_2} B(\lambda, T_*) \lambda d\lambda, \quad (7)$$

where λ_1 and λ_2 are the limiting wavelengths of the spectral bin, d is the distance of the star (we took here 270 pc by default), $R_* = 1.46 R_{\text{Sun}}$, $T_* = 6460\text{K}$, $d = 270 \text{ pc}$ and $\Delta t = 21340\text{s}$, respectively. The parameters $D = 16$, $\tau = 0.4$, and Δt are the diameter of the telescope, the performance of the system, and the integration time, respectively, whose values were fixed for JWST according to Cowan et al. (2015). Since systematics may prevent us from reaching a 10 ppm precision with the JWST, we assumed a floor noise of 30 ppm throughout the whole spectral domain with a normal distribution (Greene et al. 2016) from 0.6 to $\sim 2.8 \mu\text{m}$, wherever the noise was lower than 30 ppm.

In this last case, before proceeding with the extraction of the principal components, the spectra are filtered using a PCA considering the first 10 components to reconstruct the filtered data. The number of components used to filter out the noise was chosen by examining the cumulative explained variance of the principal components, which gives an idea on how much information we are capturing by retaining a certain number of features. It is important to stress the fact that the filtering procedure in this case is necessary since applying the pre-processing as we do in the other two cases amplifies both the spectral features and the noise, leading to worse reconstructions.

The final step of the pre-processing pipeline consists in extracting the M principal components from each spectral

band to reduce the dimensionality of the dataset. Applying the PCA to the normalized patches enables the extraction of features which are comparable in magnitude. Thus, from each patch we obtain an M -dimensional vector of principal components, $\mathbf{x}^i$, yielding a total of N_p samples. To this sample matrix we further append an additional vector containing the global maximum and minimum of each spectrum and its average value. The resulting matrix X , collecting the classical data of each spectrum, is thus organized as follow:

$$X = \begin{bmatrix} \mathbf{x}^1 \\ \mathbf{x}^2 \\ \vdots \\ \mathbf{x}^{N_p+1} \end{bmatrix} \quad \text{with } N_p = \begin{cases} N_{\text{TauREx}}=14 \\ N_{\text{JWST}}=8 \end{cases} \quad (8)$$


3.3 Reservoir dynamics

To limit the number of qubits and the depth of the reservoir, with a similar reasoning of Huang et al. (2021a); Tsang et al. (2023) we consider a factorized reservoir composed in the following way: the first sub-reservoir encodes the average of the spectra and their global maxima and minima; while, the remaining N_p sub-reservoirs, are dedicated to processing the information from each individual spectral patch.

Each of the sub-reservoir has to process M input data so we initialize a quantum register, with a number of qubits $Q \geq M$, in the $|0\rangle$ state and build a quantum circuit with the following structure, (see Fig. 3):

- **encoding layer:** the input data $\mathbf{x}^i$ encoded in a RX rotation acting on the first M qubits;
- **reservoir layers:** a simple multilayer reservoir composed of an initial layer of RY rotations applied to all qubits, followed by a layer of CNOT gates, and a final layer of RY rotations. The rotation angles are randomly initialized but kept fixed throughout the processing of all spectra;

Fig. 3 Pictorial representation of a single reservoir - The **encoding** of the classical data is performed by means of RX gates whose rotation angles depend on the principal component extracted by each patch; the **reservoir layer** consists of random RY rotations, a set of CNOT and other random RY rotations. Finally the outcomes are collected performing a projective measurement in the computational basis



- **measurement layer:** this layer performs projective measurements on all qubits in the computational basis, yielding the probability distribution over the 2^Q possible measurement outcomes.

Encoding layer - Let O be a single qubit operator, from now on we will denote with O_k the tensor product:

$$O_k \equiv \mathbb{I} \otimes \mathbb{I} \otimes \dots \otimes \overbrace{O}^{k\text{-th}} \otimes \dots \otimes \mathbb{I} \otimes \mathbb{I} \quad (9)$$

The RX gate performs local phase rotation on a qubit along the X axis, thus RX is

$$RX(\theta) = e^{-i\theta\sigma_k^x} \quad (10)$$

Therefore, the encoding layer applies a unitary evolution

$$U_{Enc}^i = \prod_{k=1}^M RX(x_k^i), \quad (11)$$

where we have encoded the k -th principal components of the i -th patch in the rotation angle of RX.

Reservoir layer - While in the encoding layer we only acted on a subset of qubits, in the reservoir we work with all of them. In particular, we perform random local rotations along the Y axis by applying the RY gate on each qubit, with RY

$$RY(\theta) = e^{-i\theta\sigma_k^y} \quad (12)$$

The CNOT gates in the reservoir defined as

$$CNOT_k^{k+1} = |0\rangle\langle 0|_k + |1\rangle\langle 1|_k \sigma_{k+1}^x \quad (13)$$

act on consecutive pairs of qubits $k, k+1$ with a conditional dynamics where the first qubit of the pair acts as the control and the second qubit as the target.

Then we apply a second layer of random RY rotations. Thus, each reservoir in the reservoir layer consists of a unitary evolution

$$U_{Res}^i = \prod_{k=1}^Q RY(\alpha_k) \prod_{k=1}^{Q-1} CNOT_k^{k+1} \prod_{k=1}^Q RY(\beta_k).$$

The rotation angles α_k and β_k are the same for each sub-reservoir.

Measurement Layer - In the last layer we perform a projective measurement on each individual qubit, obtaining a set of outcome probabilities p^i for each reservoir:

$$p^i \equiv \{p_m^i\} = \langle 0 | U_{Enc}^{i\dagger} U_{Res}^{i\dagger} \Pi_m U_{Res}^i U_{Enc}^i | 0 \rangle \quad (14)$$

where Π_m stands for the projector in the m -th vector of the computational basis. The output from each sub-reservoir p^i , corresponding to the processed outcome of the respective patch x^i in Equation 8, are aggregated into a single output matrix P structured as:

$$P = \begin{bmatrix} p^1 \\ p^2 \\ \vdots \\ p^{N_p+1} \end{bmatrix} \quad (15)$$

The aggregated output vector P is then processed by a linear output layer, which extracts the atmospheric parameters of the spectra. A pictorial representation of the QELM model used in this work is visualized in Fig. 4.

The expressivity of the present QELM is constrained by the data encoding. Following the Fourier analysis of QELMs (Xiong et al. 2025), the encoding determines the set of achievable Fourier frequencies, while the reservoir dynamics and the measurement determine the corresponding coefficients and accessible feature functions. In our case, each spectral patch is represented by M PCA coordinates and encoded through local Pauli rotations Equation 11. Therefore, for each coordinate x_k^i , the accessible frequencies are restricted to the differences of the eigenvalues of the Pauli generator. For the hardware experiment, $M = 5$ the Fourier support of each block is small and implies that the demonstrated model should not be interpreted as a generic highly expressive quantum model or as evidence of quantum advantage. The factorized architecture further implies that the complete feature map is a direct sum of patchwise feature maps,

$$\Phi_{QELM}(x) = \bigoplus_{i=0}^{N_p} \Phi_i(x^i),$$

rather than a single global feature map over all spectral coordinates. We emphasize that the novelty of our approach is the structured use of this primitive for high-dimensional continuous inverse regression. This architecture provides a concrete route for applying QELMs to high-dimensional scientific data while respecting near-term hardware constraints. Increasing the input resolution does not require a single reservoir acting on all spectral bins; it can instead be handled by increasing the number of structured spectral blocks. The present implementation does not create global quantum correlations between different patches. Cross-patch correlations are learned only through the final classical linear regression layer. Thus the framework should




Fig. 4 Resuming scheme of the whole processing - Spectra are patched according to Zingales and Waldmann (2018) and pre-processed by the processing layer consisting of an interpolator (for JWST spectra), a normalization layer and a principal component analysis layer acting

on each of the patches. The components are encoded in a reservoir as shown in Fig. 3 and the outcome probabilities are then post-processed via the output layer to obtain the atmospheric parameters

be understood as a modular, hardware-compatible QELM architecture for structured inverse problems, rather than as a claim of a new universal QELM model or of quantum advantage.

4 Results

With the purpose of testing our method on a real quantum device, we performed some benchmarks to find the optimal conditions in terms of extracted principal components and number of training spectra. The tests were carried out with the idea of finally processing the **JWST** data, which we argue to be a good compromise between the more complete data generated by TauREx and the more realistic **NJWST** data, see Fig. 2(a-c).

The results shown in Fig. 5 let us determine the number of principal components and the number of training spectra

needed to perform the retrieval. These preliminary simulations are performed in the ideal case of infinite resources, which means that the processing is purely mathematical and does not take into account statistical noise in the measuring process. In Fig. 5(a), we report the Mean Squared Relative Error (MSRE) [see Equation 4] for each parameter as a function of the number of principal components extracted per patch, while in Fig. 5(b) we show its dependence on the number of training spectra for the JWST dataset. From this exploratory analysis we conclude that the minimal number of principal components to extract is $M=5$, although some subtle improvement in retrieving H_2O and CO_2 concentrations can be obtained considering one more feature. Furthermore, at least 4000 spectra are needed in the training data set to allow the algorithm to converge. The simulations were performed considering reservoirs of 8 qubits, a training dataset of 8000 spectra for Fig. a and $M=6$ for Fig. b.




Fig. 5 Benchmark simulations using JWST dataset - Benchmark simulations at infinite statistics to determine the minimal number of principal components to extract from each spectral patch (a) and the minimal number of training data. Both benchmarks were performed using reservoirs of $N=8$ qubits. The results in the upper plot are

obtained considering a training dataset of 8000 spectra, while those in the lower plot considering a $M=6$ principal components. For each point in the plot we made the training 10 times with different training datasets, while testing always on the same set

From the accuracy achieved in retrieving the concentration of CO , we observe that, regardless of the number of training samples and the number of principal components retained, the MSRE remains on the order of 10^{-1} . This behavior is consistent with that observed in spectral reconstruction using DCGANs (Zingales and Waldmann 2018). The CO molecule does not exhibit prominent spectral features in the training set; consequently, its identification is more challenging compared to that of the other three molecules.

The algorithm was ultimately deployed on IBM Fez, a superconducting quantum device featuring 156 qubits. As outlined in Section 3, processing the JWST dataset requires 8 distinct sub-reservoirs, each dedicated to the principal components of a given spectral patch. An additional sub-reservoir is employed to encode the mean, maximum, and minimum spectral values, bringing the total number of circuits to 9.

Building on the insights gained from the benchmark simulations, we employed a dataset of 4080 spectra, divided into 75% for training and 25% for testing, and extracted a total of $M=5$ principal components from each spectral patch. Under the present experimental conditions, the model is classically simulable because it is composed of independent five-qubit reservoirs. Nonetheless, while performing the real experiment on QPU requires a computational time which only depends on the depth of the quantum circuit, from a simulational point of view the computational time strongly depends on the considered number of qubits per reservoir. However, from both perspectives this in general strongly depends on the number of qubits on the QPU for the quantum case and on the specifics of the classical computer from the simulation standpoint. We therefore do not interpret the hardware experiment as evidence of quantum advantage. The relevant scalability property of the proposed architecture is instead modularity: spectra depending on a higher number of parameters could generally require a finer patching, increasing the number of associated sub-reservoirs, while keeping each reservoir below the regime in which concentration effects make the measured observables input-insensitive. This patchwise scaling differs from the scaling of a single unstructured many-qubit reservoir. In a monolithic reservoir, concentration effects may cause observable readouts to approach an input-independent value exponentially with the number of qubits, requiring an exponentially increasing number of shots to resolve the input dependence. See Appendix B for the simulation results obtained under optimal conditions.

To parallelize the running of the nine sub-reservoirs, we assigned nine subsets of five qubits each on the QPU. Each subset processes the input data corresponding to a specific spectral patch, as illustrated in Fig. 4. In the reservoirs

handling the principal components, each qubit was encoded with one component, as shown in Fig. 3.

In Fig. 6, we show the accuracies (Equation 5) obtained from the retrievals performed with the QELM implemented on IBM Fez. The blue bars correspond to the results of an infinite-statistics simulation. The orange bars indicate the results of a simulation with a sampling statistics of 20000 shots, meaning that each data point was processed 20000 times by the reservoir layer, yielding the same number of measurement outcomes from which the qubit-state probabilities were extracted. The green bars show the results obtained by running the same task on the IBM Fez hardware, using the same sampling statistics as in the simulation. As expected, the accuracies decrease slightly when moving from the ideal simulation to finite sampling, and further when executing on the real hardware. These deviations can be attributed to a combination of shot noise, readout errors, and gate imperfections inherent to the NISQ regime. Nonetheless, the overall consistency among the three cases confirms the robustness of the proposed QELM architecture and demonstrates that the implemented reservoir computing scheme remains effective even under realistic experimental conditions.

The numerical results in Fig. 6 show that, with the current dataset and set of principal components, the model is able—at infinite statistics—to retrieve the planetary radius for all samples with a relative error below 5%. Under the same conditions, the concentrations of CH_4 and H_2O are successfully reconstructed for nearly the entire dataset within the same 5% threshold. The concentration of CO_2 is accurately retrieved for about 86% of the dataset, while the remaining atmospheric components are reconstructed with progressively lower accuracy.

At finite statistics, one can observe a degradation of about 10% in the accuracy of the best-retrieved parameters, and even larger deviations for the less well-predicted quantities. This loss of precision arises primarily from the statistical noise associated with the limited number of measurement shots, which affects the estimation of output probabilities. Nevertheless, this effect can be mitigated by increasing the sampling statistics, thereby reducing the variance in the probability estimation.

Since the QELM receives features already extracted by classical patching, normalization and PCA, the retrieval accuracy cannot be attributed solely to the quantum reservoir. We therefore compare against linear regression on the same PCA features and against a classical ELM with a matched random-feature dimension. These comparisons are necessary to identify the contribution of the quantum feature map relative to the information already present in the preprocessed classical representation. In Fig. 6 (black and purple bars) we show the accuracies obtained performing




Fig. 6 Hardware Results with JWST dataset - Results obtained by processing the dataset using IBM Fez. The results have been obtained using reservoirs of 5 qubits each and $M=5$. In this context we considered a dataset of $D=4080$ spectra, with the 75% used for training and the rest for testing. Both the results obtained on hardware and the

finite statistics simulation were obtained with 20000 shots. The algorithm was also tested considering a classical ELM (brown bars), and considering the QELM without any reservoir - using only the encoding layers with $M=5$ (purple bars) and $M=8$ (black bars) - as a baseline to show how the reservoir is actually improving the predictions

the linear regression directly on the preprocessed data. The purple bars for $M = 5$, corresponds to the same number of PCA components used in the hardware QELM experiment, and the black bars for $M = 8$, which tests whether simply retaining more PCA components can reproduce the QELM performance. The results show that the classical preprocessing is indeed informative. In particular, the radius is already well captured by the preprocessed representation, consistently with the fact that its information is largely contained in global spectral quantities. However, the no-reservoir baselines do not reproduce the QELM accuracy for most retrieved parameters. For example, at $M = 5$ the direct linear baseline is substantially below the IBM Fez QELM results for CH_4 , CO_2 , CO , H_2O , M , R , and T . Increasing the number of PCA components to $M = 8$ does not close the gap and, in several cases, performs worse. Thus the observed performance cannot be explained simply as a consequence of applying linear regression to a low-dimensional PCA representation.

For the implementation of the classical ELM we considered a similar setup as in the quantum case, considering 9

random neural networks of 32 neurons each. In this context, the input data are multiplied by a random matrix of dimension 32×5 , followed by a sigmoid function. The choice of the number of neurons for each reservoir has been chosen accordingly to the dimension of the Hilbert space in the quantum case for a consistent comparison. The output of each reservoirs has been collected in a single output vector used for the training of the output layer as in the quantum case. Figure 6 show that with this setup the results are much lower than in the quantum counterpart. Further experiments were done considering a single reservoir 288 neurons, leading to similar results. Furthermore, the figure shows that the PCA without the reservoir is not powerful enough, with the model achieving only the 90% of accuracy for the radius which should be the easiest parameter to be retrieved, and decreasing accuracy also for the other parameters. Interestingly, increasing the number of principal components is detrimental for the model. In Appendix A, indeed, we will show that the components of $M = 5 - 6$ are pretty much uninformative also using the reservoir.

What is particularly remarkable is that the differences between the simulation results and the experimental outcomes obtained on the IBM Fez processor are negligible compared to the effects of sampling noise. As shown in Fig. 7, the predictions of all molecular concentrations obtained with the simulator closely match those produced by the quantum hardware. The agreement between the finite-shot simulation and the IBM Fez results should be understood as readout-level adaptation to the effective noisy feature map. In the QELM protocol, the output weights are trained directly on the probability matrix produced by the reservoir. On hardware, this matrix is not the ideal noiseless matrix P_{ideal} , but an experimentally realized matrix $P_{\text{hw}} = P_{\text{ideal}} + \Delta P_{\text{hw}}$, where ΔP_{hw} includes the effects of finite sampling, readout errors, gate imperfections, crosstalk, and other device-specific noise sources. Since training and testing are performed using the same hardware-generated feature map, the final linear readout learns the regression map associated with the experimentally realized features. This mechanism has clear limitations. It can only compensate distortions that are sufficiently stable between training and testing and that do not erase the input-dependent information contained in the feature matrix. In the simple case where the dominant effect of hardware noise is an approximate attenuation, $P_{\text{hw}} \simeq \lambda P_{\text{ideal}}$, the Moore-Penrose pseudo-inverse can absorb this factor through $(\lambda P_{\text{ideal}})^+ = \lambda^{-1} P_{\text{ideal}}^+$. This compensation is algebraic and occurs in the classical readout layer; it should not be interpreted as a quantum error-mitigation protocol. Moreover, it is not cost-free: the corresponding increase in the norm of the learned weights can amplify finite-shot fluctuations in the measured probabilities. Therefore, the hardware results demonstrate robustness under the stable noise conditions and shallow circuits used in the IBM Fez experiment.


5 Conclusion and future works

In this work, we have introduced and experimentally validated a Quantum Extreme Learning Machine for the retrieval of exoplanetary atmospheric parameters. Thanks to intrinsic noise resilience of this approach, we demonstrated a viable route for applying near-term quantum devices to complex astrophysical data analysis.

The QELM successfully processed synthetic spectra generated with TauREx, accurately retrieving key parameters such as molecular abundances, planetary mass, radius, and equilibrium temperature using only 1% of the total dataset. Under ideal conditions, the radius, CH_4 , and H_2O concentrations were retrieved with accuracies above 90%. When implemented on the IBM Fez processor the performance degraded only marginally compared to simulations. The model naturally adapts to the hardware's noise, learning the device's effective measurement process and thus self-calibrating during training.

The hybrid classical-quantum pipeline, based on spectral patching, PCA, and a factorized quantum reservoir, reduces qubit use and circuit depth while keeping key spectral features. It also remains robust against realistic shot noise, with PCA pre-filtering helping to recover most of the lost accuracy. From a broader perspective, the QELM can be integrated with classical retrieval frameworks, where its rapid training can provide intelligent priors for more computationally intensive algorithms, improving convergence and reducing overall cost. Future research will focus on extending the method to more complex atmospheric models including more complex chemical models, clouds and temperature pressure profiles, and on leveraging advances in quantum hardware to handle higher-resolution data with less pre-processing. Furthermore, the QELM framework

Fig. 7 Compared predictions using simulator and hardware with JWST dataset - Parameters estimated by IBM Fez and by finite statistics simulations for each parameter



serves as a versatile platform for investigating alternative reservoir dynamics and encoding schemes to further enhance efficiency and scalability.

In conclusion, our results demonstrate that quantum machine learning, even with today's hardware, can contribute meaningfully to exoplanet science. The robustness and adaptability of the proposed QELM mark a significant step toward the practical application of quantum computing in astrophysics and open new possibilities for analyzing the rich datasets from current and future missions such as JWST and Ariel.

Appendix

In this section, we review all the results obtained from processing the 4 datasets presented in Fig. 2. In particular, Appendix A presents the benchmark results, previously presented for the JWST data, now applied to the other datasets. Conversely, Appendix B reports the results obtained by processing the four datasets using the optimal setup identified in Appendix A.

Appendix A Training and Feature Tests

As done previously for the JWST dataset, in this section we perform a series of benchmarks to determine the optimal conditions for the atmospheric retrieval with QELM, processing all the datasets presented in Fig. 2. As in Fig. 5, for the feature test we fixed the dimension of the training dataset to $D_{\text{train}} = 8000$ and evaluated the model on a test set of $D_{\text{test}} = 2000$. The results of the feature test are shown in Fig. 8. In the **TauREx** and **JWST** cases (Fig. 8 (a) and (b)) we observe that the MSRE in the retrieval of all the parameters, reaches a minimum $M = 5-6$ components. It is also evident that, in all the cases, the relative error in the retrieval of CO does not fall below 0.1.

Furthermore, while in the **TauREx** and **JWST** cases the relative error reaches its minimum with $M = 6$ components, in the other two datasets increasing the number of principal components does not appear to improve accuracy. This suggests that shot noise and the filtering procedure likely destroy most of the information captured through PCA. On the other hand, the accuracy in the retrieval of radius, mass and temperature is largely unaffected by the




Fig. 8 Benchmark simulations on all datasets with variation of principal components - Mean squared relative error of the algorithm as a function of the number of encoded principal components. The simulations were performed on a dataset of 10^4 spectra, with the 80% used for training and the remaining 20% for testing. Panel(a) shows

the results for the **TauREx** dataset, (b) for the interpolated dataset within JWST spectral range, in (c) for the NJWST dataset and (d) for the FJWST dataset. Each point has been obtained by doing 10 different training instances, considering different training sets, while keeping fixed the testing set

number of principal components. For mass and temperature this indicates that their information is fully encoded in the first principal component. On the other hand, the radius could only be reconstructed by processing the information contained in the average, maximum, and minimum values of the spectra, implying that its information is encoded there rather than in the principal components.

All the results discussed here are obtained at infinite statistics.

The tests aimed at determining the minimal size of the training set were thus performed using $M=6$ principal components for all the datasets.

The evaluation was again carried out on $D_{test}=2000$ spectra at infinite statistics. The results are shown in Fig. 9, where we observe that, while the **TauREx** dataset requires approximately $D_{train} = 5000$ samples for the QELM to achieve full convergence, in the other cases convergence is reached already with $D_{train} = 3000-4000$.

Appendix B Simulations Results

The results obtained [A](#) allow us to test the algorithm at its full potential, at least in simulation.





Fig. 9 Benchmark simulation with all datasets varying the number of training data - Accuracy of the algorithm varying the number of training states and considering a test of 2000 spectra. In the panel (a) there are the performances using the **TauREx** dataset, in (b) there are those using the interpolated dataset in the spectral range of **JWST**, in

In Fig. 10, we show the results of the atmospheric retrieval performed on the four types of datasets. As illustrated by the histograms, in the **TauREx** case nearly all the parameters can be retrieved at infinite statistics with an accuracy of 99 – 100% except for the mass and temperature which still reach a success rate of 96%. Also in this case the retrieval accuracy of CO remains the lowest even at infinite statistics, achieving only about 80%. In the lower histogram of Fig. 10 we note that, while CH_4 , CO_2 , H_2O and R are retrieved with success rate exceeding 90%, the same does not hold for mass and temperature which require a higher statistics to be estimated accurately. For the **JWST** dataset, the QELM exhibits a comparable level of accuracy to that observed in the **TauREx** case, with only a slight degradation in the retrievability of CO_2 and CO . This suggest that part of their spectral features lies within the spectral range that was excluded.

Adding shot noise to this dataset and processing it without first filtering the noise using the principal component analysis results in a retrieval accuracy below the 70% for all parameters, except for the radius. The latter shows robustness against noise, as previously seen in Fig. 8, since its information is encoded in the average values of

(c) those with the dataset of **NJWST** and in (d) those with the dataset of **FJWST**. The number of components used for each patch is $M=6$. Each point has been obtained by doing 10 different training instances, considering different training sets, while keeping fixed the testing set

Fig. 10 Simulation results in the best conditions - Accuracy of the QELM in retrieving each atmospheric parameter using the datasets shown in Fig. 2. In this context we used a dataset of 10000 spectra split in 80 – 20% respectively for training and testing and used $M = 6$ principal components for each patch. In (a) are reported the results at infinite statistics, while in (b) the results with a statistics of 20000 shots



the spectra and is therefore less affected by the presence of shot noise. These performances can be enhanced by applying a filtering procedure during pre-processing. Indeed, without such filtering, the normalization of each spectral patch amplifies not only the spectral features but also the shot noise, invalidating the principal components extracted before the encoding. As shown in Fig. 10, the principal component filtering significantly improves retrieval performance. Although the success rate remains lower than those obtained for the JWST dataset, they still outperform the NJWST, reaching 70 – 80% of accuracy at finite statistics for CH_4 , CO_2 and H_2O . Naturally, all the results at finite statistics can be improved with a higher number of shots which should let us approach the infinite statistics results. It is important to distinguish between two different sources of noise in our study. The first is noise in the input spectra, as represented by the NJWST dataset. The second is hardware noise arising from the execution of the quantum circuits on a real device. For the NJWST dataset, the improvement obtained after constructing the FJWST dataset is primarily due to the classical PCA pre-filtering step. Without this filtering, the patchwise normalization amplifies both the relevant spectral features and the spectral shot noise, degrading the principal components that are subsequently encoded into the quantum reservoirs. The hardware results address a different question. There, the input to the QELM is already the preprocessed JWST representation, and the relevant noise source is the imperfect implementation of the reservoir probabilities on the quantum processor. The agreement between finite-shot simulation and IBM Fez execution shows that, for the shallow circuits and stable device conditions considered here, the final linear readout can adapt to the effective noisy feature map produced by the hardware.

Acknowledgements TZi acknowledges support from CHEOPS ASI-INAF agreement n. 2019-29-HH.0, NVIDIA Academic Hardware Grant Program for the use of the Titan V GPU card. We acknowledge the use of IBM Quantum Credits for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

Author Contributions All the authors have contributed equally to this work

Funding Open access funding provided by Università degli Studi di Palermo within the CRUI-CARE Agreement. TZ acknowledges funding from the Italian MUR Departments of Excellence grant 2023-2027 “Quantum Frontiers”. MV and GMP acknowledge funding under project PNRR - Research infrastructures: Strengthening the Italian leadership in ELT and SKA (STILES). SL and GMP acknowledge support by MUR under PRIN Project No. 2022FEXLYB, Quantum Reservoir Computing (QuReCo) and by the “National Centre for HPC, Big Data and Quantum Computing (HPC)” Project CN00000013 HyQELM – SPOKE 10.

Data Availability The dataset used for training and testing of this work are available for download on GitHub at <https://github.com/MarcoVet-rano02/Data-ExoQELM>

Declarations

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Al-Refaie AF, Changeat Q, Waldmann IP, Tinetti G (2021) TauREx 3: A Fast, Dynamic, and Extendable Framework for Retrievals. *Astrophys J* 917(1):37
- Banchi L, Pereira J, Pirandola S (2021) Generalization in Quantum Machine Learning: A Quantum Information Standpoint. *PRX Quantum* 2(4):40321. <https://doi.org/10.1103/PRXQuantum.2.040321>
- Barstow JK, Aigrain S, Irwin PGJ, Sing DK (2016) A consistent retrieval analysis of 10 hot jupiters observed in transmission. *Astrophys J* 834(1):50
- Benneke B, Seager S (2013) How to distinguish between cloudy mini-neptunes and water/volatile-dominated super-earths. *Astrophys J* 778(2):153
- Bhavad P, Li S (2021) Physical reservoir computing with origami and its application to robotic crawling. *Sci Rep* 11(1):13002
- Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S (2017) Quantum machine learning. *Nature* 549(7671):195–202
- Bruno G, Lewis NK, Stevenson KB, Filippazzo J, Hill M, Fraine JD, Wakeford HR, Deming D, Kilpatrick B, Line MR, Morley CV, Collins KA, Conti DM, Garlitz J, Rodriguez JE (2018) A Comparative Study of WASP-67 b and HAT-P-38 b from WFC3 Data. *Astron J* 155(2):55
- Cerezo M, Verdon G, Huang HY, Cincio L, Coles PJ (2022) Challenges and opportunities in quantum machine learning. *Nat Comput Sci* 2(9):567–576
- Cobb AD, Himes MD, Soboczenski F, Zorzan S, O’Beirne MD, Güneş Baydin A, Gal Y, Domagal-Goldman SD, Arney GN, Angerhausen D (2019) An ensemble of bayesian neural networks for exoplanetary atmospheric retrieval. *Astron J* 158(1):33
- Coulombe JC, York MCA, Sylvestre J (2017) Computing with networks of nonlinear mechanical oscillators. *PLoS ONE* 12(6):0178663
- Cowan NB, Greene T, Angerhausen D, Batalha NE, Clampin M, Colón K, Crossfield IJM, Fortney JJ, Gaudi BS, Harrington J (2015) Others: Characterizing transiting planet atmospheres through 2025. *Publ Astron Soc Pac* 127(949):311
- Cubillos P, Bleic J, Harrington J, Rojo P, Lust N, Bowman O, Stemm M, Foster A, Loredó TJ, Fortney J et al (2020) BART: Bayesian Atmospheric Radiative Transfer code. *Astrophysics Source Code Library* 1608
- De Lorenzis A, Casado M, Estarellas M, Gullo NL, Lux T, Plastina F, Riera A, Settino J (2024) Harnessing quantum extreme learning machines for image classification. *arXiv preprint arXiv:2409.00998*
- De Lorenzis A, Casado M, Gullo NL, Lux T, Plastina F, Riera A (2025) Behind the scenes of the quantum extreme learning machines. *arXiv preprint arXiv:2509.06873*
- Feroz F, Hobson MP (2008) Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses. *Mon Not R Astron Soc* 384(2):449–463. <https://doi.org/10.1111/j.1365-2966.2007.12353.x>
- Feroz F, Gair JR, Hobson MP, Porter EK (2009) Use of the MULTINEST algorithm for gravitational wave data analysis. *Class Quantum Gravity* 26(21):215003
- Fu G, Espinoza N, Sing DK, Lothringer JD, Dos Santos LA, Rustamkulov Z, Deming D, Kempton EM-R, Komacek TD, Knutson HA, Albert L, Pontoppidan K, Volk K, Filippazzo J (2022) Water and an Escaping Helium Tail Detected in the Hazy and Methane-depleted Atmosphere of HAT-P-18b from JWST NIRISS/SOSS. *Astrophys J Lett* 940(2):35
- Fujii K, Nakajima K (2017) Harnessing Disordered-Ensemble Quantum Dynamics for Machine Learning. *Phys Rev Appl* 8(2):24030. <https://doi.org/10.1103/PhysRevApplied.8.024030>
- Gandhi S, Madhusudhan N (2018) Retrieval of exoplanet emission spectra with HyDRA. *Mon Not R Astron Soc* 474(1):271–278
- García-Bení J, Giorgi GL, Soriano MC, Zambrini R (2023) Scalable Photonic Platform for Real-Time Quantum Reservoir Computing. *Phys Rev Appl* 20(1):14051
- Gardner JP (2007) The james webb space telescope. In: *Proceedings of the 18th International Symposium on Space Terahertz Technology 2007, ISSTT 2007* 123:237. https://doi.org/10.1007/978-1-4020-9457-6_1
- Ghosh S, Paterek T, Liew TCH (2019a) Quantum Neuromorphic Platform for Quantum State Preparation. *Phys Rev Lett* 123(26):260404. <https://doi.org/10.1103/PhysRevLett.123.260404>
- Ghosh S, Opala A, Matuszewski M, Paterek T, Liew TCH (2019b) Quantum reservoir processing. *Quantum Information* 5(1):35
- Ghosh S, Krisnanda T, Paterek T, Liew TCH (2021a) Realising and compressing quantum circuits with quantum reservoir computing. *Communications Physics* 4(1):1–7
- Ghosh S, Opala A, Matuszewski M, Paterek T, Liew TCH (2021b) Reconstructing quantum states with quantum reservoir networks. *IEEE Trans Neural Netw Learn Syst* 32(7):3148–3155
- Goodfellow I, Bengio Y, Courville A (2016) Deep Learning. MIT Press. <http://www.deeplearningbook.org>
- Goudarzi A, Lakin MR, Stefanovic D (2013) DNA reservoir computing: A novel molecular computing approach. In: Soloveichik D, Yurke B (eds) *Lecture Notes in Computer Science (including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Lecture Notes in Computer Science, vol 8141 LNCS, pp 76–89. Springer, Cham. https://doi.org/10.1007/978-3-319-01928-4_6
- Govia LCG, Ribeill GJ, Rowlands GE, Krovi HK, Ohki TA (2021) Quantum reservoir computing with a single nonlinear oscillator. *Phys Rev Res* 3(1):13077
- Greene TP, Line MR, Montero C, Fortney JJ, Lustig-Yaeger J, Luther K (2016) Characterizing transiting exoplanet atmospheres with jwst. *Astrophys J* 817(1):17. <https://doi.org/10.3847/0004-637x/817/1/17>
- Gregory PC (2011) Bayesian exoplanet tests of a new method for MCMC sampling in highly correlated model parameter spaces. *Mon Not R Astron Soc* 410(1):94–110. <https://doi.org/10.1111/j.1365-2966.2010.17428.x>
- Huang GB, Zhu QY, Siew CK (2004) Extreme learning machine: A new learning scheme of feedforward neural networks. In: *IEEE International Conference on Neural Networks - Conference Proceedings*, vol 2, pp 985–990. <https://doi.org/10.1109/IJCNN.2004.1380068>
- Huang GB, Chen L, Siew CK (2006) Universal approximation using incremental constructive feedforward networks with random hidden nodes. *IEEE Trans Neural Netw* 17(4):879–892. <https://doi.org/10.1109/TNN.2006.875977>
- Huang GB, Wang DH, Lan Y (2011) Extreme learning machines: A survey. *Int J Mach Learn Cybern* 2(2):107–122. <https://doi.org/10.1007/s13042-011-0019-y>
- Huang HL, Du Y, Gong M, Zhao Y, Wu Y, Wang C, Li S, Liang F, Lin J, Xu Y, Yang R, Liu T, Hsieh MH, Deng H, Rong H, Peng CZ, Lu CY, Chen YA, Tao D, Zhu X, Pan JW (2021a) Experimental Quantum Generative Adversarial Networks for Image Generation. *Phys Rev Appl* 16(2):24051
- Huang HY, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, McClean JR (2021b) Power of data in quantum machine learning. *Nat Commun* 12(1):2631
- Innocenti L, Lorenzo S, Palmisano I, Ferraro A, Paternostro M, Palma GM (2023) Potential and limitations of quantum extreme learning machines. <https://doi.org/10.1038/s42005-023-01233-w>, *arXiv:2210.00780*
- Irwin PGJ, Teanby NA, Kok R, Fletcher LN, Howett CJA, Tsang CCC, Wilson CF, Calcutt SB, Nixon CA, Parrish PD (2008) The

- NEMESIS planetary atmosphere radiative transfer and retrieval tool. *J Quant Spectrosc Radiat Transfer* 109(6):1136–1150. <https://doi.org/10.1016/j.jqsrt.2007.11.006>
- Konkoli Z (2017) On Reservoir Computing: From Mathematical Foundations to Unconventional Applications. In: Adamatzky A (ed) *Advances in Unconventional Computing: Volume 1: Theory. Emergence, Complexity and Computation*, pp 573–607. Springer, Cham. DOI:10.1007/978-3-319-33924-5_23
- Kreidberg L, Line MR, Parmentier V, Stevenson KB, Louden T, Bonnefoy M, Faherty JK, Henry GW, Williamson MH, Stassun K, Beatty TG, Bean JL, Fortney JJ, Showman AP, Désert J-M, Arcangeli J (2018) Global Climate and Atmospheric Composition of the Ultra-hot Jupiter WASP-103b from HST and Spitzer Phase Curve Observations. *Astron J* 156(1):17
- Krisnanda T, Ghosh S, Paterek T, Liew TCH (2021) Creating and concentrating quantum resource states in noisy environments using a quantum neural network. *Neural Networks* 136:141–151. <https://doi.org/10.1016/j.neunet.2021.01.003>, arXiv:2101.05994
- Lau JWZ, Lim KH, Shrotriya H, Kwek LC (2022) NISQ computing: where are we and where do we go? *AAPPS Bull* 32(1):27
- Lavie B, Mendonça JM, Mordasini C, Malik M, Bonnefoy M, Demory B-O, Oreshenko M, Grimm SL, Ehrenreich D, Heng K (2017) HELIOS-RETRIEVAL: An Open-source, Nested Sampling Atmospheric Retrieval Code; Application to the HR 8799 Exoplanets and Inferred Constraints for Planet Formation. *Astron J* 154(3):91
- Line MR, Wolf AS, Zhang X, Knutson H, Kammer JA, Ellison E, Deroo P, Crisp D, Yung YL (2013) A systematic retrieval analysis of secondary eclipse spectra. I. A comparison of atmospheric retrieval techniques. *Astrophys J* 775(2):137. <https://doi.org/10.1088/0004-637X/775/2/137>, arXiv:1304.5561
- Lo Monaco G, Bertini M, Lorenzo S, Massimo Palma G (2024) Quantum extreme learning of molecular potential energy surfaces and force fields. *Mach Learn Sci Technol* 5(3):35014
- Madhusudhan N, Seager S (2009) A temperature and abundance retrieval method for exoplanet atmospheres. *Astrophys J* 707(1):24–39
- Mansfield M, Bean JL, Line MR, Parmentier V, Kreidberg L, Désert J-M, Fortney JJ, Stevenson KB, Arcangeli J, Dragomir D (2018) An HST/WFC3 Thermal Emission Spectrum of the Hot Jupiter HAT-P-7b. *Astron J* 156(1):10
- Markowska-Kaczmar U, Kosturek M (2021) Extreme learning machine versus classical feedforward network: Comparison from the usability perspective. *Neural Comput Appl* 33(22):15121–15144
- Márquez-Neila P, Fisher C, Sznitman R, Heng K (2018) Supervised machine learning for analysing spectra of exoplanetary atmospheres. *Nature Astronomy* 2(9):719–724
- Martínez-Peña R, Ortega JP (2023) Quantum reservoir computing in finite dimensions. *Phys Rev E* 107(3):35306
- Martínez-Peña R, Giorgi GL, Nokkala J, Soriano MC, Zambrini R (2021) Dynamical phase transitions in quantum reservoir computing. *Phys Rev Lett* 127(10):100502
- Martínez-Peña R, Nokkala J, Giorgi GL, Zambrini R, Soriano MC (2023) Information processing capacity of spin-based quantum reservoir computing systems. *Cogn Comput* 15(5):1440–1451
- Mujal P, Martínez-Peña R, Nokkala J, García-Bení J, Giorgi GL, Soriano MC, Zambrini R (2021) Opportunities in quantum reservoir computing and extreme learning machines. *Adv Quantum Technol* 4(8):2100027
- Mujal P, Martínez-Peña R, Giorgi GL, Soriano MC, Zambrini R (2023) Time-series quantum reservoir computing with weak and projective measurements. *npj Quantum Information* 9(1):1–10. <https://doi.org/10.1038/s41534-023-00682-z>, arXiv:2205.06809
- Nakajima K, Hauser H, Li T, Pfeifer R (2015) Information processing via physical soft body. *Sci Rep* 5(1):10487. <https://doi.org/10.1038/srep10487>
- Nakajima K, Fujii K, Negoro M, Mitarai K, Kitagawa M (2019) Boosting computational power through spatial multiplexing in quantum reservoir computing. *Phys Rev Appl* 11(3):34021
- Nokkala J, Martínez-Peña R, Zambrini R, Soriano MC (2022) High-performance reservoir computing with fluctuations in linear networks. *IEEE Trans Neural Netw Learn Syst* 33(6):2664–2675. <https://doi.org/10.1109/TNNLS.2021.3105695>
- Pluriel W, Zingales T, Leconte J, Parmentier V (2020) Strong biases in retrieved atmospheric composition caused by day-night chemical heterogeneities. *Astronomy and Astrophysics* 636:66. <https://doi.org/10.1051/0004-6361/202037678>, arXiv:2003.05943
- Pluriel W, Leconte J, Parmentier V, Zingales T, Falco A, Selsis F, Bordé P (2022) Toward a multidimensional analysis of transmission spectroscopy: II. Day-night-induced biases in retrievals from hot to ultrahot Jupiters. *Astronomy and Astrophysics* 658:42. <https://doi.org/10.1051/0004-6361/202141943>, arXiv:2110.09080
- Preskill J (2018) Quantum computing in the NISQ era and beyond. *Quantum* 2:79. <https://doi.org/10.22331/q-2018-08-06-79>, arXiv:1801.00862
- Rebentrost P, Mohseni M, Lloyd S (2014) Quantum support vector machine for big data classification. *Phys Rev Lett* 113(3):130503
- Rocchetto M, Waldmann IP, Venot O, Lagage P-O, Tinetti G (2016) Exploring biases of atmospheric retrievals in simulated jwst transmission spectra of hot jupiters. *Astrophys J* 833(1):120
- Schuld M, Sinayskiy I, Petruccione F (2015) An introduction to quantum machine learning. *Contemp Phys* 56(2):172–185
- Senokosov A, Sedykh A, Sagingalieva A, Kyriacou B, Melnikov A (2024) Quantum machine learning for image classification. *Mach Learn Sci Technol* 5(1):15040. <https://doi.org/10.1088/2632-2153/ad2aef>
- Serre D (2010) Matrices with Entries in a Principal Ideal Domain; Jordan Reduction. In: *Matrices: Theory and Applications*, pp 163–181. Springer. https://doi.org/10.1007/978-1-4419-7683-3_9
- Sheppard KB, Mandell AM, Tamburo P, Gandhi S, Pinhas A, Madhusudhan N, Deming D (2017) Evidence for a Dayside Thermal Inversion and High Metallicity for the Hot Jupiter WASP-18b. *Astrophys J Lett* 850(2):32
- Skilling J (2004) Nested sampling. Bayesian inference and maximum entropy methods in science and engineering 735:395–405
- Soriano MC, Ortín S, Keuninckx L, Appeltant L, Danckaert J, Pesquera L, Van Der Sande G (2015) Delay-based reservoir computing: Noise effects in a combined analog and digital implementation. *IEEE Trans Neural Netw Learn Syst* 26(2):388–393. <https://doi.org/10.1109/TNNLS.2014.2311855>
- Spake JJ, Sing DK, Evans TM, Oklopić A, Bourrier V, Kreidberg L, Rackham BV, Irwin J, Ehrenreich D, Wyttenbach A, Wakeford HR, Zhou Y, Chubb KL, Nikolov N, Goyal JM, Henry GW, Williamson MH, Blumenthal S, Anderson DR, Hellier C, Charbonneau D, Udry S, Madhusudhan N (2018) Helium in the eroding atmosphere of an exoplanet. *Nature* 557(7703):68–70. <https://doi.org/10.1038/s41586-018-0067-5>
- Suprano A, Zia D, Innocenti L, Lorenzo S, Cimini V, Giordani T, Palmisano I, Polino E, Spagnolo N, Sciarrino F, Palma GM, Ferraro A, Paternostro M (2024) Experimental Property Reconstruction in a Photonic Quantum Extreme Learning Machine. <https://doi.org/10.1103/PhysRevLett.132.160802>, arXiv:2308.04543
- Tanaka G, Yamane T, Hérault JB, Nakane R, Kanazawa N, Takeda S, Numata H, Nakano D, Hirose A (2019) Recent advances in physical reservoir computing: A review. *Neural Networks* 115:100–123. <https://doi.org/10.1016/j.neunet.2019.03.005>, arXiv:1808.04962
- Tinetti G, Drossart P, Eccleston P, Hartogh P, Heske A, Leconte J, Micela G, Ollivier M, Pilbratt G, Puig L, Turrini D, Vandenbussche B, Wolkenberg P, Pascale E, Beaulieu J-P, Güdel M, Min M, Rataj M, Ray T, Ribas I, Barstow J, Bowles N, Coustenis A, Coudé du Foresto V, Decin L, Encrenaz T, Forget F, Friswell M,

- Griffin M, Lagage PO, Malaguti P, Moneti A, Morales JC, Pace E, Rocchetto M, Sarkar S, Selsis F, Taylor W, Tennyson J, Venot O, Waldmann IP, Wright G, Zingales T, Zapatero-Osorio MR (2016) The science of ARIEL (Atmospheric Remote-sensing Infrared Exoplanet Large-survey). In: *Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, vol 9904, p 99041. <https://doi.org/10.1117/12.2232370>. SPIE
- Tsang SL, West MT, Erfani SM, Usman M (2023) Hybrid Quantum-Classical Generative Adversarial Network for High-Resolution Image Generation. *IEEE Trans Quant Eng* 4. <https://doi.org/10.1109/TQE.2023.3319319>, [arXiv:2212.11614](https://arxiv.org/abs/2212.11614)
- Tsiaras A, Waldmann IP, Zingales T, Rocchetto M, Morello G, Damiano M, Karpouzas K, Tinetti G, McKemmish LK, Tennyson J, Yurchenko SN (2018) A Population Study of Gaseous Exoplanets. *Astron J* 155(4):156
- Vetrano M, Lo Monaco G, Innocenti L, Lorenzo S, Palma GM (2025) State estimation with quantum extreme learning machines beyond the scrambling time. *Quantum Information* 11(1). <https://doi.org/10.1038/s41534-024-00927-5>
- Waldmann IP (2016) Dreaming of atmospheres. *Astrophys J* 820(2):107
- Waldmann IP, Tinetti G, Rocchetto M, Barton EJ, Yurchenko SN, Tennyson J (2015) Tau-REx I: A next generation retrieval code for exoplanetary atmospheres. *Astrophys J* 802(2):107
- Wang J, Lu S, Wang SH, Zhang YD (2022) A review on extreme learning machine. *Multimedia Tools and Applications* 81(29):41611–41660. <https://doi.org/10.1007/s11042-021-11007-7>
- Xiong W, Facelli G, Sahebi M, Agnel O, Chotibut T, Thanasilp S, Holmes Z (2025) On fundamental aspects of quantum extreme learning machines. <https://doi.org/10.1007/s42484-025-00239-7>, [arXiv:2312.15124](https://arxiv.org/abs/2312.15124)
- Zia D, Innocenti L, Minati G, Lorenzo S, Suprano A, Di Bartolo R, Spagnolo N, Giordani T, Cimini V, Palma GM, Ferraro A, Sciarino F, Paternostro M (2025) Quantum extreme learning machines for photonic entanglement witnessing. <https://doi.org/10.48550/arXiv.2502.18361>, [arXiv:2502.18361](https://arxiv.org/abs/2502.18361)
- Zingales T, Waldmann IP (2018) ExoGAN: Retrieving Exoplanetary Atmospheres Using Deep Convolutional Generative Adversarial Networks. *Astron J* 156(6):268

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.