



Universiteit
Leiden
The Netherlands

Using a neural network approach to accelerate disequilibrium chemistry calculations in exoplanet atmospheres

Hendrix, J.L.A.M.; Louca, A.J.; Miguel, Y.

Citation

Hendrix, J. L. A. M., Louca, A. J., & Miguel, Y. (2023). Using a neural network approach to accelerate disequilibrium chemistry calculations in exoplanet atmospheres. *Monthly Notices Of The Royal Astronomical Society*, 524(1), 643-655. doi:10.1093/mnras/stad1763

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3719172>

Note: To cite this publication please use the final published version (if applicable).

Using a neural network approach to accelerate disequilibrium chemistry calculations in exoplanet atmospheres

Julius L. A. M. Hendrix,¹ Amy J. Louca¹   and Yamila Miguel^{1,2}

¹*Leiden Observatory, Leiden University, Postbus 9513, NL-2300 RA Leiden, the Netherlands*

²*SRON Netherlands Institute for Space Research, Niels Bohrweg 4, NL-2333 CA Leiden, the Netherlands*

Accepted 2023 June 7. Received 2023 June 5; in original form 2022 December 1

ABSTRACT

In this era of exoplanet characterization with *JWST*, the need for a fast implementation of classical forward models to understand the chemical and physical processes in exoplanet atmospheres is more important than ever. Notably, the time-dependent ordinary differential equations to be solved by chemical kinetics codes are very time-consuming to compute. In this study, we focus on the implementation of neural networks to replace mathematical frameworks in one-dimensional chemical kinetics codes. Using the gravity gradient, temperature-pressure profiles, initial mixing ratios, and stellar flux of a sample of hot-Jupiter's atmospheres as free parameters, the neural network is built to predict the mixing ratio outputs in steady state. The architecture of the network is composed of individual autoencoders for each input variable to reduce the input dimensionality, which is then used as the input training data for an LSTM-like neural network. Results show that the autoencoders for the mixing ratios, stellar spectra, and pressure gradients are exceedingly successful in encoding and decoding the data. Our results show that in 90 per cent of the cases, the fully trained model is able to predict the evolved mixing ratios of the species in the hot-Jupiter atmosphere simulations. The fully trained model is $\sim 10^3$ times faster than the simulations done with the forward, chemical kinetics model while making accurate predictions.

Key words: exoplanets – planets and satellites: atmospheres – planets and satellites: gaseous planets.

1 INTRODUCTION

There are two methods commonly used for calculating the abundance of different species in an atmosphere: thermochemical equilibrium and chemical kinetics (Bahn & Zuckoski 1960; Zeleznik & Gordon 1960). Thermochemical equilibrium calculations treat each species independently and do not require an extensive list of reactions between different species. Consequently, this method is fast for estimating the abundance of different species in an exoplanet atmosphere and has been widely used in the community (e.g. Stock et al. 2018a; Woitke et al. 2018). However, the atmospheres of exoplanets are dynamic environments. Both physical and chemical processes can alter the compositions and thermal structures of the atmosphere. In particular, atmospheric processes like photochemistry, mixing and condensation of different species can affect atmospheric abundances, deviating the concentrations observed from what would be found by chemical equilibrium calculations (Cooper & Showman 2006; Swain, Vasisht & Tinetti 2008; Moses et al. 2011; Baxter et al. 2021; Kawashima & Min 2021; Roudier et al. 2021). For example, the recent detection of SO₂ (Ahrer et al. 2022; Alderson et al. 2022; Rustamkulov et al. 2022; Feinstein et al. 2023) and the determination of this species as direct evidence of photochemical processes shaping the atmosphere of WASP 39b (Tsai et al. 2022), suggest that certain exoplanet atmospheres are in disequilibrium, and we need chemical

disequilibrium models using chemical kinetics to correctly interpret the observations.

Chemical kinetics codes consider the effects that lead to a non-equilibrium state in the atmosphere. These codes incorporate a wide range of atmospheric processes, such as the radiation from the host star that can drive the dissociation of molecules or photochemistry, the mixing of species at different pressures due to the planet's winds, or the diffusion of species, and calculate the one-dimensional abundances of species in exoplanetary atmospheres (e.g. Moses et al. 2011; Venot et al. 2012; Miguel & Kaltenecker 2014; Tsai et al. 2017; Hobbs et al. 2019). However, to calculate the abundance of different species using chemical kinetics, a system of coupled differential equations involving all the species must be solved, and prior knowledge of reaction rates and a reaction list is necessary to estimate the production and loss of each species. Therefore, as more species and reactions are incorporated into the chemical networks, the complexity of these simulations increases, and so does the computational cost these simulations require. The result is that chemical kinetics codes have long computational times, and cannot be used by more detailed calculations (e.g. circulation models) or as a fast way of interpreting observations (by retrieval codes), which are usually subject to simplifications.

For the past few decades, the use of machine learning techniques, specifically neural networks (NNs), has become more prevalent in research fields outside of computer science. Within astronomy, NNs have been used for applications like image processing (Dattilo et al. 2019), adaptive optics (Landman et al. 2021), exoplanet detection

* E-mail: louca@strw.leidenuniv.nl

(Shallue & Vanderburg 2018), exoplanetary atmospheric retrieval (Cobb et al. 2019), chemical modelling (Holdship et al. 2021), and more traditional machine learning techniques have been used for applications like exoplanetary atmospheric retrieval (Nixon & Madhusudhan 2020) and chemistry modelling of protoplanetary discs (Smirnov-Pinchukov et al. 2022). Trained NNs are fast to use, so a NN trained to accurately reproduce the outcomes of chemical kinetics codes could greatly reduce computational times. Such a NN could simulate a large amount of atmospheric conditions in a short period of time, which is for example useful for atmospheric retrievals from observational constraints. It could also be incorporated into a multidimensional atmospheric simulation that connects a multitude of individual one-dimensional simulations by the implementation of atmospheric mixing and other global processes.

In this study, we investigate the feasibility of machine learning techniques for speeding up a one-dimensional chemical kinetics code. To this end, we perform calculations on a fiducial giant planet as an example to show how this technique can be used to bring the best of these two worlds: the detailed information of chemical kinetics calculations and the speed of NNs techniques. In the next section, we explain in more detail how we obtain the data set and the specifics of the architectures used. The results of our networks are presented in the following section (Section 3), which are discussed afterwards in Section 4. Finally, we summarize and conclude our findings in Section 5.

2 METHODS

2.1 Chemical kinetics

Chemical kinetics is the most realistic way of calculating abundances and is necessary, particularly at low temperatures ($T < 2000$ K) and pressures ($P < 10$ – 100 bars), where time-scales of processes such as mixing in the atmosphere are shorter than chemical equilibrium and dominate the chemistry and abundances in the atmosphere.

We make use of the one-dimensional chemical kinetics code VULCAN (Tsai et al. 2017, 2021) to create a large data set on the atmospheres of gaseous exoplanets. The code is validated for hot-Jupiter atmospheres from 500 to 2500 K. VULCAN calculates a set of mass differential equations:

$$\frac{\partial n_i}{\partial t} = \mathcal{P}_i - \mathcal{L}_i - \frac{\partial \Phi_i}{\partial z}, \quad (1)$$

where n_i is the number density of the species i , t is the time, \mathcal{P}_i and \mathcal{L}_i are the production and loss rates of the i -th species, and Φ_i its transport flux that includes the effects of dynamics caused by convection and turbulence in the atmosphere. For a more complete derivation of this equation from the general diffusion equation, we refer the reader to Hu, Seager & Bains (2012). VULCAN starts from initial atmospheric abundances calculated using the chemical equilibrium chemistry code *FastChem* (Stock et al. 2018b), although we note that the final disequilibrium abundances are not affected by the choice of initial values adopted by Tsai et al. (2017), and further evolves these abundances by solving a set of Eulerian continuity equations that includes various physical processes (e.g. vertical mixing and photochemistry). To solve these partial differential equations, VULCAN numerically transforms them into a set of stiff ordinary differential equations (ODEs). These ODEs are solved using the *Rosenbrock* method, which is described in detail in the appendix of Tsai et al. (2017). In this study, we make use of machine learning techniques to solve these equations and hence speed up the process.

2.2 Building the data set

2.2.1 Parameter space

To construct the data set, we vary the following parameters:

- (i) **Planet mass, M [M_J]**: within the range [0.5, 20] M_J .
- (ii) **Orbit radius, r [au]**: within the range [0.01, 0.5] au.
- (iii) **Stellar radius, R_\star [R_\odot]**: within the range [1, 1.5] R_\odot .

Other parameters such as surface gravity, irradiation temperature, and stellar effective temperature are derived from these free parameters.

(i) **Planet radius [R_{Jup}]**: This is derived from the planet mass using the relation from Chen & Kipping (2017), shown in equation (2), where R is the planet radius and M is the planet mass:

$$\frac{R}{R_\oplus} = 17.78 \left(\frac{M}{M_\oplus} \right)^{-0.044}. \quad (2)$$

We note that our aim is to present the results for a simple general case, and the mass–radius relation we use is suitable for this purpose. However, we must emphasize that the relation between mass and radius for giant exoplanets is not unique and depends on various factors, such as the mass of metals, core mass, irradiation received by the planet, and their effect on the inflation of radius. All of these factors can impact the evolution path of giant planets and their final radius, leading to a dispersion in the mass–radius relation.

(i) **Temperature–pressure profile**: As our aim is to demonstrate the use of NNs for calculating non-equilibrium chemical abundances in a general case, we have utilized an analytical, non-inverted temperature–pressure profile from Heng, Mendonça & Lee (2014). While these analytical profiles are simplistic, they are widely used in the literature to explore general cases and are suitable for our purposes. However, for calculating the chemistry of a real planet, more detailed calculations that take into account the opacities of different species and their abundances in the atmosphere should be included. The assumptions for this calculation are $T_{\text{int}} = 120$ K, $\kappa_L = 0.1$, $\kappa_S = 0.02$, $\beta_S = 1$, and $\beta_L = 1$, based on the default values included in the VULCAN code Tsai et al. (2017). The pressure profile is constructed within the range [10^{-2} , 10^9] dyne cm^{-2} . This calculation is an important step, as it determines whether the set of parameters is valid for the data set. If any part of the temperature profile falls outside of the range [500, 2500] K, the temperature range for which VULCAN is validated, the example is rejected from the data set.

(ii) **Stellar flux**: The stellar spectra used for the data set have two sources: the Measurements of the Ultraviolet Spectra Characteristics of Low-mass Exoplanetary Systems (MUSCLES) collaboration (France et al. 2016; Loyd et al. 2016; Youngblood et al. 2016) and the PHOENIX Stellar and Planetary Atmosphere Code (Baron, Chen & Hauschildt 2010). The MUSCLES data base contains observations from M- and K-dwarf exoplanet host stars in the optical, UV, and X-ray regime, and is used for stars with an effective temperature lower than 6000 K. For effective temperatures of 6000 K and above, stellar spectra are generated by the PHOENIX model. Flux values below 10^{-14} erg nm^{-1} cm^{-2} s^{-1} are cut-off.

The remaining parameters of the VULCAN configuration files are kept constant throughout the data set. Eddy and molecular diffusion are both taken into account as fixed parameters. For the eddy diffusion constant, we use a constant of $K_{zz} = 10^{10}$ cm^2 s^{-1} . The molecular diffusion constant is taken for a hydrogen-dominated gas as described

in Banks & Kockarts (1973). As a standard chemical network, we make use of VULCAN's reduced default N–C–H–O network that includes photochemistry. We assume solar elemental abundances for the hot-Jupiters, and we make use of 150 pressure levels for the height layers. The output of VULCAN is saved every 10 steps. In total, 13291 valid configurations are generated within the parameter space.

2.2.2 Formatting

In order to limit the computation times when training the network, the input–output pairs do not contain all of the information supplied by VULCAN.

For the inputs, a selection of six properties is made. These properties are extracted from VULCAN before time integration starts, so they can be interpreted as the initial conditions of the simulation. The six properties are:

- (i) **Initial mixing ratios:** the initial mixing ratios of the species in the simulation. These mixing ratios are calculated by VULCAN using FastChem (Stock et al. 2018b). The shape of the array containing the mixing ratios is (69, 150), as the mixing ratios are defined for 69 species for 150 height layers each.
- (ii) **Temperature profile:** the temperature profile as calculated by the analytical expression from Heng et al. (2014). The temperature is defined for every height layer, so it has a shape of (150).
- (iii) **Pressure profile:** the pressure profile that is calculated as part of the temperature profile calculation. It is of the same shape, (150).
- (iv) **Gravitational profile:** the gravitational acceleration per height layer. It has the shape (150).
- (v) **Stellar flux component:** one of the two components that make up the stellar spectrum contains the flux values. This is generated from either the MUSCLES data base and/or the PHOENIX model and interpolated to a shape of (2500).
- (vi) **Stellar wavelength component:** the second component of the stellar spectrum contains the wavelengths corresponding to the flux values. It has the same shape, (2500).

For the outputs, we make use of the time-dependent mixing ratios.

Because not every simulation takes the same amount of time to converge to a solution, the number of saved abundances differs per VULCAN simulations. To include the information contained in the evolution of the abundances through time, 10 sets of abundances, including the steady-state abundances, are saved in each output. This set of abundances is evenly spaced through time, so the simulation time between abundances will vary for different VULCAN simulation runs. Before the abundances are saved, they are converted to mixing ratios. The shape of the outputs is (10, 69, 150).

2.2.3 Data standardization

The inputs and outputs of the various components differ by several orders in magnitude. To ensure that the NN trained on the data set is not biased towards higher-valued parameters, the data has to be standardized. First, the distributions of the properties are standardized according to equation (3):

$$p_s = \frac{\log_{10}(p) - \mu}{\sigma}, \quad (3)$$

with

$$\mu = \frac{1}{n} \sum_{i=0}^n \log_{10}(p_i), \quad (4)$$

and

$$\sigma = \sqrt{\frac{1}{n} \sum_{i=0}^n (\log_{10}(p_i) - \mu)^2}, \quad (5)$$

where p is the property to be scaled, n is the size of the data set, and p_s is the standardized property. After standardization, the properties are normalized in the range [0, 1]:

$$p_{s,n} = \frac{p_s - \min(p_s)}{\max(p_s) - \min(p_s)}, \quad (6)$$

where $p_{s,n}$ is the final normalized property.

Once the input properties are normalized, the output mixing ratios are normalized with the same scaling parameters as were used for the input mixing ratios. When the trained NN is presented with an input for which to predict the mixing ratios, it only has information about the scaling parameters of the inputs. To be able to unnormalize the outputs, they need to be scaled with the same scaling parameters.

2.3 Model architecture

2.3.1 Autoencoder structure

The input of each configuration within the data set consists of roughly 15 800 values. To speed up the training process and complexity of the NN, we make use of an *autoencoder* (AE) for reducing the dimensionality of the examples in the data set. In previous studies, this approach has been shown to be an effective way to reduce dimensionality within chemical kinetics (e.g. Grassi et al. 2022). An AE consists of two collaborating NNs: an *encoder* and a *decoder*. The task of the encoder is to reduce the dimensionality of the input data by extracting characterizing features from the example and encoding them in a lower dimensionality representation called the *latent representation*. The task of the decoder is to take the latent representation and use it to reconstruct the original input data with as little loss of information as possible. The encoder and decoder are trained simultaneously, and no restraints are placed on the way the AE uses its *latent space*, apart from the size of the latent representations.

As discussed in Section 2.2.2, the inputs of the model consist of six properties. Because these properties do not share the same shape, we cannot encode and decode them using a single AE. Instead, we construct six unique AE, one for each property of the model inputs. Fig. 1 shows an overview of the process of encoding the initial conditions. The decoding process is not shown but is symmetrical to the encoding process.

Each encoder conceals a specific property into a corresponding latent representation. To get the latent presentation of the entire input example, l_i , the property latent representations $\{l_{MR}, l_F, l_W, l_T, l_P, l_G\}$ are concatenated. When decoding the latent representation of the input, the latent vector l_i is split back into the different property latent representations, and given to that property's decoder. Every encoder–decoder pair is trained separately.

The hyperparameters of each AE are optimized by trial and error. A summary of each set of hyperparameters is shown in Table 1.

2.3.1.1 Mixing ratios autoencoder The mixing ratios are the largest contributor to the size of the model inputs. Compressing each of the species' mixing ratios efficiently reduces the size of the input latent representation, l_i , by a substantial amount. Because of the limited size of the training data set, a compromise has to be made to successfully train this AE. Rather than concurrently encoding all 69 species for each example, each species is encoded individually.

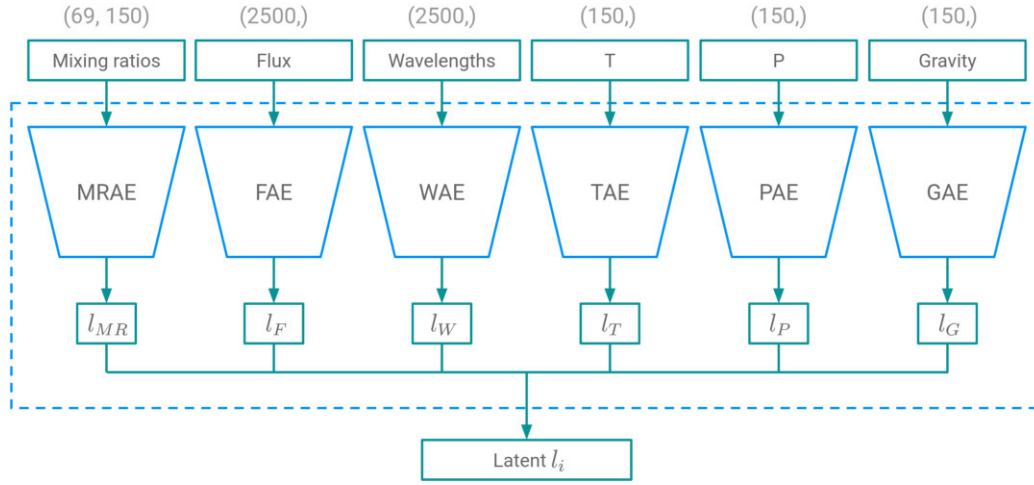


Figure 1. The different properties and their corresponding encoders that together encode the input data. The decoding process is symmetrical to this encoding process, where every property has a corresponding property decoder.

Table 1. Hyperparameters for the property AEs. Each AE has an encoder and decoder NN. These are MLPs, consisting of seven fully connected layers with hyperbolic tangent activation functions.

Model	Hidden size	Latent size	Optimizer	Learning rate	Batch size	Epochs
MRAE	256	30	Adam	10^{-5}	32	200
PAE	256	2	Adam	10^{-6}	4	100
TAE	256	30	Adam	10^{-5}	4	100
GAE	256	30	Adam	10^{-5}	4	100
FAE	1024	256	Adam	10^{-5}	4	200
WAE	1024	2	Adam	10^{-7}	4	200

As a result, the training data set expands by a factor of 69 while disregarding any potential correlations in species abundances during the encoding procedure.

Fig. 2 shows the application of such an AE: for a given input, each of the 69 species' mixing ratios is encoded into corresponding latent vectors $\{l_1, l_2, l_3, \dots, l_{69}\}$. The concatenation of these 69 latent vectors then makes up the latent representation of the mixing ratios, l_{MR} .

All encoders and decoders are multilayer perceptron (MLP) NNs. For the mixing ratio autoencoder (MRAE), the encoder and decoder both consist of seven fully connected layers, followed by hyperbolic tangent activation functions. The encoder input layer has a size of 150 and the output layer has a size of 30. The hidden layers have a size of 256. Adversely, the decoder has an input layer size of 30 and an output layer size of 150. The compression factor of the MRAE is therefore $150/30 = 5$.

To train the MRAE, the data set is split into a train data set (70 per cent), a validation data set (20 per cent), and a test data set (10 per cent). To increase the size of the data set, the MRAE is trained on a shuffled set¹ of the mixing ratios of both the inputs and the output of the chemical kinetics simulations. The performance of the AE is measured using the loss function in equation (7):

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left(\frac{p_i - a_i}{a_i} \right)^2. \quad (7)$$

where \mathcal{L} is the loss, N is the number of elements in the actual/predicted vector, and p_i and a_i are the i -th elements of the

¹Using a random sampler function within the PYTORCH package.

predicted and actual vectors, respectively. The MRAE is optimized using the Adam optimizer (Kingma & Ba 2014), with a learning rate of 10^{-5} . A batch size of 32 is used, and the model is trained for 200 epochs. These hyperparameters can also be found in Table 1.

2.3.1.2 Atmospheric profile autoencoders The temperature, pressure, and gravity profiles all have the same shape of (150), so their AE can use the same architecture. Moreover, the atmospheric profiles share their shape with the mixing ratios of individual species (i.e. the height layers in the atmosphere). Therefore, a very similar NN structure to that of the MRAE is used for the atmospheric profile AE. The encoder input layer shape, the decoder output layer shape, and the hidden layer shapes are taken directly from the MRAE for all atmospheric profile AE. An important parameter to tune for each atmospheric profile AE separately is the size of the latent representations.

The pressure profile is set logarithmically for all examples in the data set. By taking the logarithm of the pressures, the spacing becomes linear, $\log(P_i) - \log(P_{i+1})$. Theoretically, we then only need two values to fully describe the pressure profile: the pressure at the first and last height layers. To encode these values, no AE is needed. One could take this one step further, and provide input parameters like mass and radius, from which the pressure and gravity profiles are dependent, directly to the core model as inputs. While this more specialized approach is suitable for these input parameters, it is not generalizable to other input parameters. To keep the model architecture more general and adaptable to different input parameters, an AE is used none the less. The size of the pressure profile autoencoder (PAE) latent representations is set to 2. This corresponds to a compression factor of $150/2 = 75$.

The temperature and gravity profiles are not linear. For both the temperature autoencoder (TAE) and gravity autoencoder (GAE), a latent representation size of 30 is used. This corresponds to a compression factor of $150/30 = 5$, the same as for the MRAE.

All profile AEs are evaluated using the loss function previously defined in equation (7) and are optimized using the Adam optimizer. The TAE and GAE use a learning rate of 10^{-5} , and the PAE uses a learning rate of 10^{-6} . All profile AEs are trained with a batch size of 4, for 100 epochs (see also Table 1).

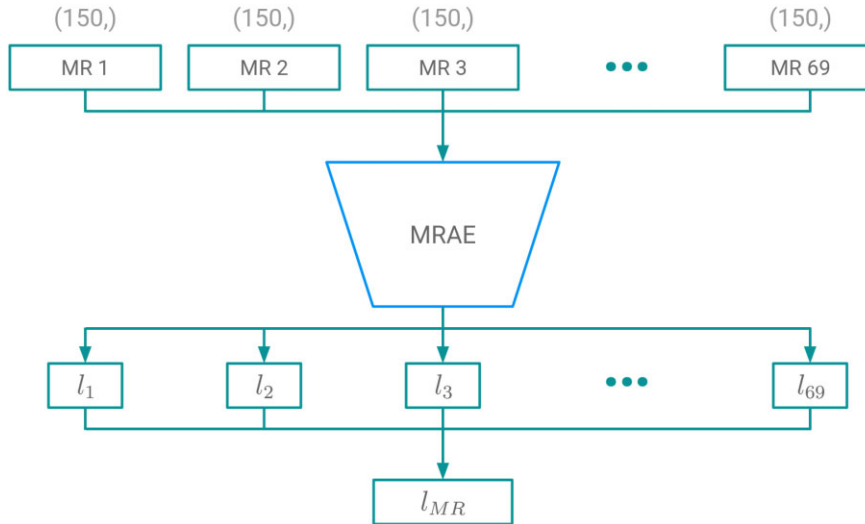


Figure 2. A more detailed sketch of the architecture of the mixing ratio AE. MR i denotes the mixing ratio of a certain species, i , for all height layers, and l_i denotes the encoded mixing ratios for species i .

2.3.1.3 Stellar spectrum autoencoders After the mixing ratios, the stellar spectrum components contribute predominantly to the input data size. The stellar spectrum is comprised of a flux and a wavelength component. These components share the same shape, so one NN structure can be used for both AEs. The structure of the encoder and decoder is similar to that of the MRAE: a 7-layer, fully connected MLP with hyperbolic tangent activation functions after each layer. The encoder input layer and decoder output layer have a size of 2500, and the hidden layers have a size of 1024.

Similarly to the PAE, the wavelength bins are spaced logarithmically. Again only two values are needed to fully describe the wavelength range. The latent representation size for the wavelength autoencoder (WAE) is therefore also 2. The compression factor for this network is $2500/2 = 1250$. The flux autoencoder (FAE) has a latent representation size of 256, which gives it a compression factor of $2500/256 \approx 10$.

Both AEs are evaluated using the loss function from equation (7). They are optimized using the Adam optimizer, the WAE with a learning rate of 10^{-7} , and the FAE with a learning rate of 10^{-5} . They are both trained for 200 epochs with batches of four examples (see also Table 1).

2.3.2 Core network

As mentioned before, the outputs are also large in dimensionality. Because the outputs contain mixing ratios for all species, for 10-time steps (Section 2.2), they can be encoded using the MRAE. Fig. 3 shows how the AE would encode both the inputs and the last time step of the outputs to their latent representations l_i and l_o , respectively. Note that even though the AE is shown twice in this figure, the same AE is used to encode both the inputs and the outputs. In the middle of the figure, connecting the two latent spaces, a second NN called the *core network* is located.

The function of the core network is to learn a mapping between the latent representations of the inputs and the evolved outputs. The design of the core network takes advantage of some of the characteristics of VULCAN. From Section 2.1, we know that VULCAN solves ODEs for specific atmospheric configurations over a simulated period of time. To impart this sense of time in the core

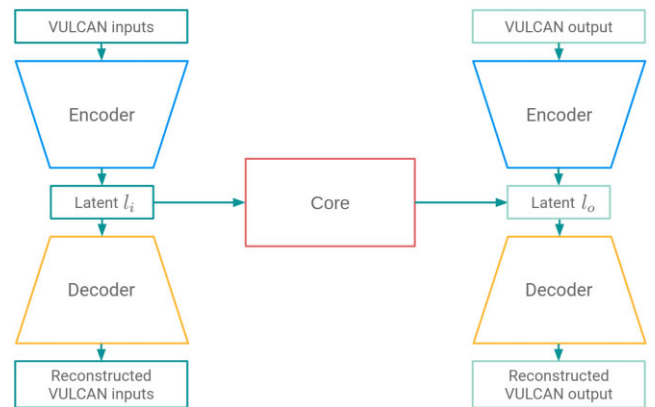


Figure 3. An overview of the model architecture. The core NN maps between the latent representations of the VULCAN inputs l_i and outputs l_o .

NN, a *Long–Short Term Memory* (LSTM) is used as the base of the design. The LSTM was chosen for its proven performance in numerous applications, from stellar variability (e.g. Jamal & Bloom 2020) to Core-collapse supernovae search (e.g. Less et al. 2023) and solar radio spectrum classification (e.g. Xu et al. 2019), as well as the ease of implementation. The LSTM has known shortcomings like the vanishing gradient problem and long training times when dealing with long sequences. However, with the short sequence length used with our model (i.e. 10 time-steps), these shortcomings are not considered problematic for this proof of concept.

The input of the core network is not sequential in nature. With some changes, we can use the LSTM in a ‘one-to-many’ configuration. In this configuration, the initial output of the LSTM h_0 is given to an MLP. This MLP produces a vector with the same shape as the initial input x_0 , which can be interpreted as the ‘evolved’ input x_1 . This evolved input is fed back into the LSTM to produce h_1 , from which the MLP produces x_2 , and so forth. This can be repeated for an arbitrary number of steps.

The design of the core network is visualized in Fig. 4. We interpret the latent representation of the inputs l_i as the initial value x_0 . The LSTM and MLP configuration produces nine intermediary ‘evolved’ latent representations $\{x_1, \dots, x_9\}$ before arriving at the final evolved

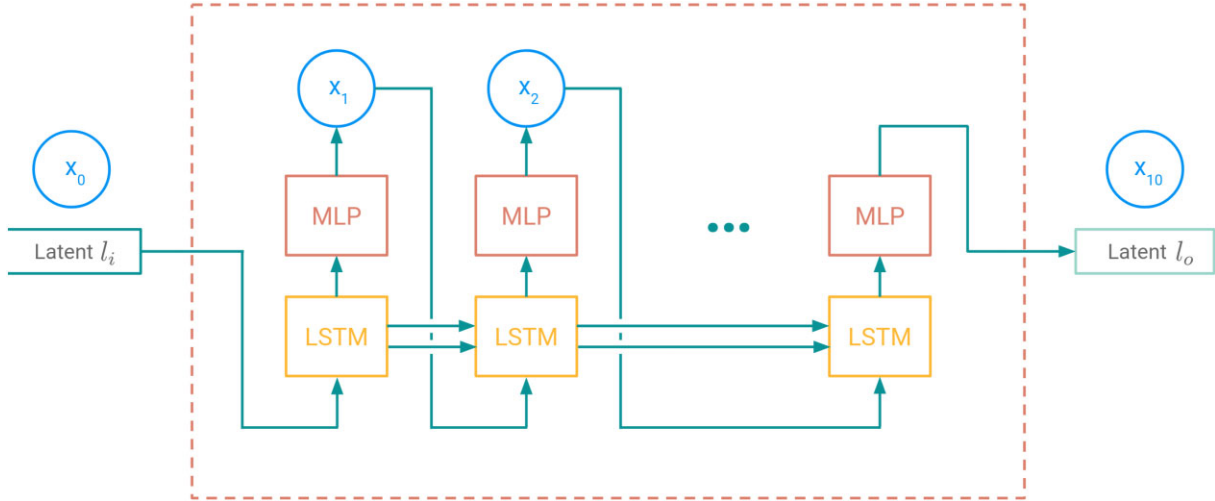


Figure 4. The design of the core network. It consists of a one-to-many LSTM + MLP configuration that is run for 10 steps.

latent representation x_{10} . We interpret this latent representation as the prediction of the latent representation of the evolved output l_o .

2.3.2.1 Training When the core model predicts a sequence of 10 latent representations, it is essentially traversing the latent space. We can guide the network to learn to traverse the latent space similarly to how VULCAN simulations evolve by using the sequence of outputs saved in the data set (Section 2.2). We do this in two ways: first, we construct a loss function that not only depends on the accuracy of the prediction of the latent representation of the final output l_o , but also on the accuracy of the intermittent latent representation predictions:

$$\mathcal{L} = \sum_{t=1}^{10} \left(\frac{1}{N} \sum_{i=1}^N (p_{t,i} - a_{t,i})^2 \right), \quad (8)$$

where \mathcal{L} is the loss, N is the number of elements in the actual/predicted vector, $p_{t,i}$ is the i -th element of the latent representation prediction vector at time step t , and $a_{t,i}$ is the i -th element of the latent representation vector of the output at time step t . With this notation, $a_{10} = l_o$. By training a network with this loss function, we force the core network to evolve the latent mixing ratios similarly to how VULCAN evolves mixing ratios. It should be noted that the latent representation of the inputs l_i is larger than the latent representation of the outputs l_o , as it contains more properties than just mixing ratios. The predicted latent representations x_t are therefore also larger than l_o . To account for this, we only look at the elements corresponding to the encoded mixing ratios in l_i when comparing the predicted latent representations x_t and the output mixing ratios a_t/l_o .

To further incentivize the core network to adhere to VULCAN's evolution patterns, we can intercept the predicted latent representations x_t before they get fed back into the LSTM, and replace them with the latent representation of the actual output of the corresponding time step. This way, the core network is always learning from latent representations that follow VULCAN's evolution, even if the network is predicting poorly. This is only done during the training of the network when the true VULCAN outputs are known. During validation and testing, the predicted latent representations x_t are not altered.

The core model LSTM has a hidden and cell size of 4096. The MLP has only two layers: an input layer of size 4096, and an output

layer of the same size as the latent representation of the inputs l_i , followed by a hyperbolic tangent function. It is optimized with the Adam optimizer, with a learning rate of 10^{-4} and a batch size of 8. It is trained for 100 epochs.

2.3.3 Deployment

When the trained model is deployed on the validation and test data set, the first step is encoding the inputs into their latent representation l_i using the encoder part of the AE (top left section in Fig. 3). The core network then predicts the latent representation of the evolved VULCAN output l_o by traversing the latent space in 10 steps (centre section in Fig. 3). The prediction of the latent representation of the VULCAN output is then decoded by the decoder part of the AE to obtain the predicted mixing ratios (bottom right section in Fig. 3).

3 RESULTS

3.1 Autoencoders

3.1.1 Mixing ratio autoencoder

The top row of Fig. 5 shows the reconstructed mixing ratio values against their actual values (left plot) for all examples from the test data set. For the entire range of mixing ratios, the majority of the reconstructions lie within an order of magnitude of the diagonal line that marks perfect reconstructions, with an R -squared value of $R^2 = 0.9997$.

The right plot of the top row of Fig. 5 shows the reconstruction error of the mixing ratios in logarithmic space. This scale is chosen because the AEs are trained in log space (Section 2.2). The solid line shows the median and the dashed lines show the 5th and 95th percentiles. From the right figure, we can see that 90 percent of the reconstructions have an error between -0.39 and 0.40 orders of magnitude.

3.1.2 Flux autoencoder

The middle row of Fig. 5 (left) shows the reconstructed flux values against their actual value. All reconstructed flux values are within

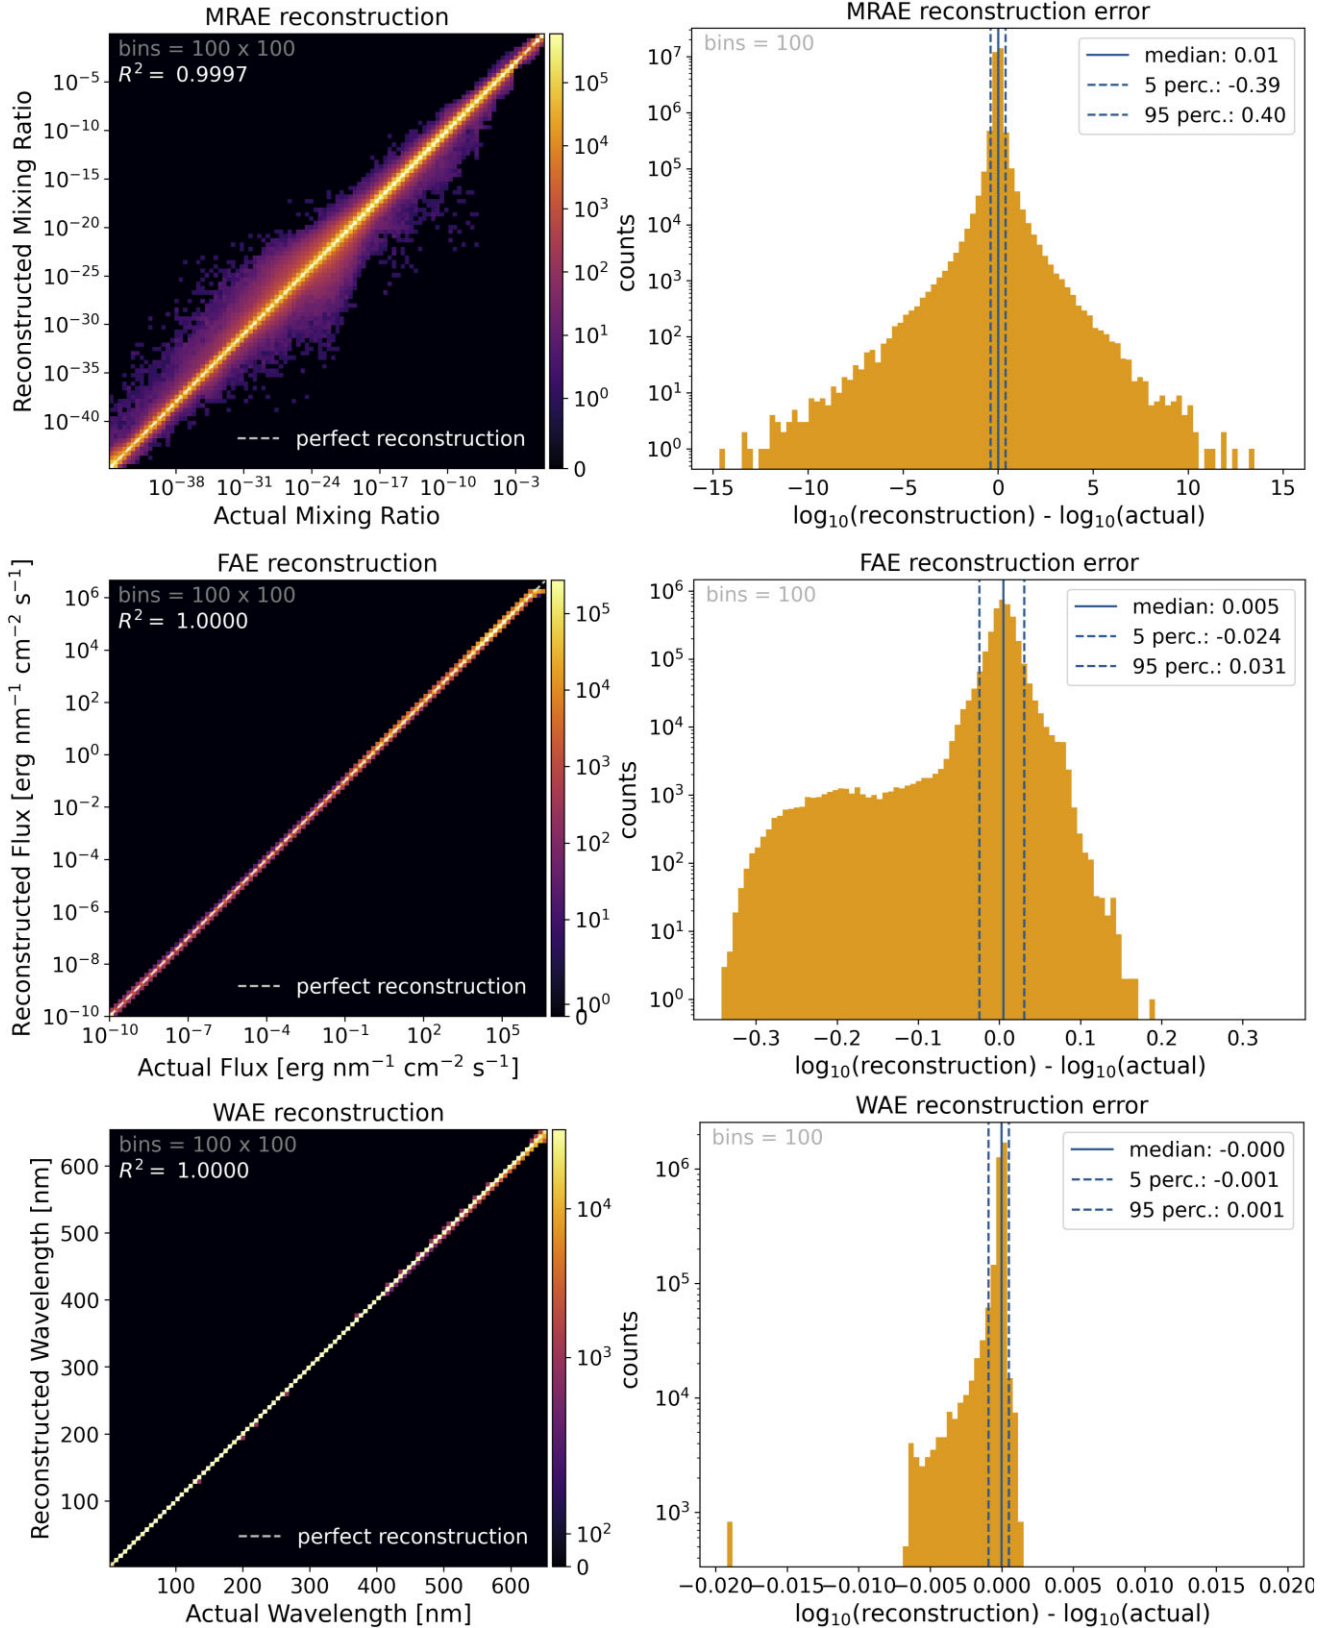


Figure 5. The reconstructed against the actual input values (left column) and the reconstruction error in log space (right column) for the mixing ratios (top row), stellar flux (middle row), and the wavelengths (bottom row). The diagonal dashed line in the reconstructed versus actual mixing ratios plot shows the performance of a perfectly reconstructing model. Here, the colour represents the number of examples within each bin. In the reconstruction error figure, the solid line shows the median value, and the dashed lines show the 5th and 95th percentiles. The R^2 values of each reconstruction plot are shown in the left column.

0.5 order magnitude of the graph diagonal. At fluxes with values around $10^6 \text{ erg nm}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$, the FAE is slightly underpredicting the actual flux values.

From the reconstruction error plot (right), we can see that 90 per cent of the reconstructions have an error between -0.024 and 0.031 orders of magnitude. In this figure, we see a distinct underprediction of a small number of examples, which are the high flux values we see being underpredicted.

3.1.3 Wavelength autoencoder

The bottom row of Fig. 5 (left) shows the reconstructed wavelength values against their actual value. All reconstructed wavelength values are close to the actual values, deviating less than $\sim 10 \text{ nm}$ from the graph diagonal. We can see that for wavelengths with values around 650 nm , the FAE has a tendency to underpredict.

From the reconstruction error of the wavelength values plot (right), we can see that 90 per cent of the reconstructions have an error between -0.001 and 0.001 orders of magnitude. The slight underprediction of higher wavelength values is also visible in this figure.

3.1.4 Pressure profile autoencoder

The top row of Fig. 6 (left) shows the reconstructed pressure values against their actual value. All reconstructed pressure values are well within 0.1 order magnitude of the graph diagonal.

From the reconstruction error plot of the pressure values (right), we can see that 90 per cent of the reconstructions have an error between -0.0003 and 0.0003 orders of magnitude. This figure also shows a very minor underprediction of some pressure values, which correspond to pressure values of $\sim 10^3 \text{ bar}$.

3.1.5 Temperature profile autoencoder

The middle row of Fig. 6 (left) shows the reconstructed temperature values against their actual value for samples from the test data set. It is immediately obvious that this AE cannot accurately reconstruct the temperature profiles. A fraction of temperatures in the range $\sim 750 \text{ K} < T < \sim 1400 \text{ K}$ are reconstructed close to the graph diagonal, but the TAE is largely overpredicting temperatures below $\sim 750 \text{ K}$ and underpredicting temperatures above $\sim 750 \text{ K}$.

The histogram of reconstruction errors of the temperature values (right) shows under- and overprediction. 90 per cent of the reconstructions have an error between 15.8 and 754.97 K . Most predictions outside this range are underpredictions.

3.1.6 Gravity profile autoencoder

The GAE shows similar behaviour to the TAE. The bottom row of Fig. 6 (left) shows the reconstructed gravity values against their actual value, for samples from the test data set. Gravity values below $\sim 4000 \text{ cm s}^{-2}$ are underpredicted by the AE, while values above $\sim 4000 \text{ cm s}^{-2}$ are overpredicted. Only values around $\sim 4000 \text{ cm s}^{-2}$ are predicted accurately by the GAE. In the reconstruction errors plot of the GAE of Fig. 6 (right), we can see that gravity values are consistently being over- and underpredicted. 90 per cent of the reconstructions have an error between 197.46 and 4939.6 cm s^{-2} . Most predictions outside this range are overpredictions.

3.2 Core network

Because the TAE and GAE do not accurately reconstruct the temperature and gravity profiles, these profiles were not encoded for the final model. Instead, they were put directly in the latent representations of the inputs. This way, no information contained in these profiles is lost.

The left plot in Fig. 7 shows the mixing ratios predicted by the trained NN model against the actual mixing ratios for the test data set. The histogram shows that most of the model predictions lie within ~ 1 order magnitude of the diagonal of the graph. A notable exception are the predictions for the few mixing ratios with values lower than $\sim 10^{-44}$, for which the model overpredicts. These species can be neglected since they are not abundant enough to play a big role in the chemistry or to show features in the observable spectra.

Fig. 7 (right) shows the mixing ratio prediction error of the NN model in log-space. The solid line shows the median and the dashed lines show the 5th and 95th percentiles. From the figure, we can see that 90 per cent of the model predictions have an error between -0.66 and 0.65 orders magnitude. Outside of this range, the model does not show a clear tendency to either over- or underpredict.

Fig. 8 shows selected examples (best, typical, and worst cases) of predictions by the NN compared with the output of the VULCAN model for a selection of seven species. The best case (top panel) shows a prediction that is almost indistinguishable from the actual mixing ratios. The examples in the typical case (middle panel) and worst case (lower panel) show larger prediction errors. In the typical case, CO_2 , CO , and HCN have the largest prediction errors in the lower atmosphere, though still negligible. The worst case shows the largest prediction errors, with H having prediction errors of up to almost 1 order of magnitude in the upper atmosphere. Notable is that this case has very strong photochemistry in the upper atmosphere as it's positioned nearby its host star.

To compare the computational efficiency of VULCAN and the NN model, the computational time to calculate or predict every example in the full data set was recorded. The results are presented in Table 2. It should be noted that the VULCAN simulations were run on similar, but older hardware than the NN model. The median computational times show a $\sim 7.5 \cdot 10^3$ decrease in computational time for the NN model. The longest computational time required by the NN model still shows a $\sim 10^3$ decrease in computational time compared to the fastest VULCAN simulation.

4 DISCUSSION

In this study, we successfully used AEs to extract most of the characterizing input features and encode them into latent representations for the mixing ratios, stellar flux, wavelengths, and pressure profiles. Within these four groups, the largest prediction errors stem from the MRAE due to the high variability in input values, as opposed to the other input sources. We included the initial and evolved mixing ratios of 69 species over 150 height layers. Additionally, the mixing ratio profiles among species differed significantly from one another (e.g. CH_4 and CO in Fig. 8). This made the complexity of extracting and encoding the fundamental input features, highest for this particular AE. In opposition, the variety in the flux from stellar spectra was much less. We obtained the stellar spectra from either the MUSCLES data base or generated them using the PHOENIX model. The spectra from these different sources were quite distinct from each other in the EUV ($0.5\text{--}200 \text{ nm}$). The PHOENIX models assume the spectra to follow blackbodies, while, in reality, M- and

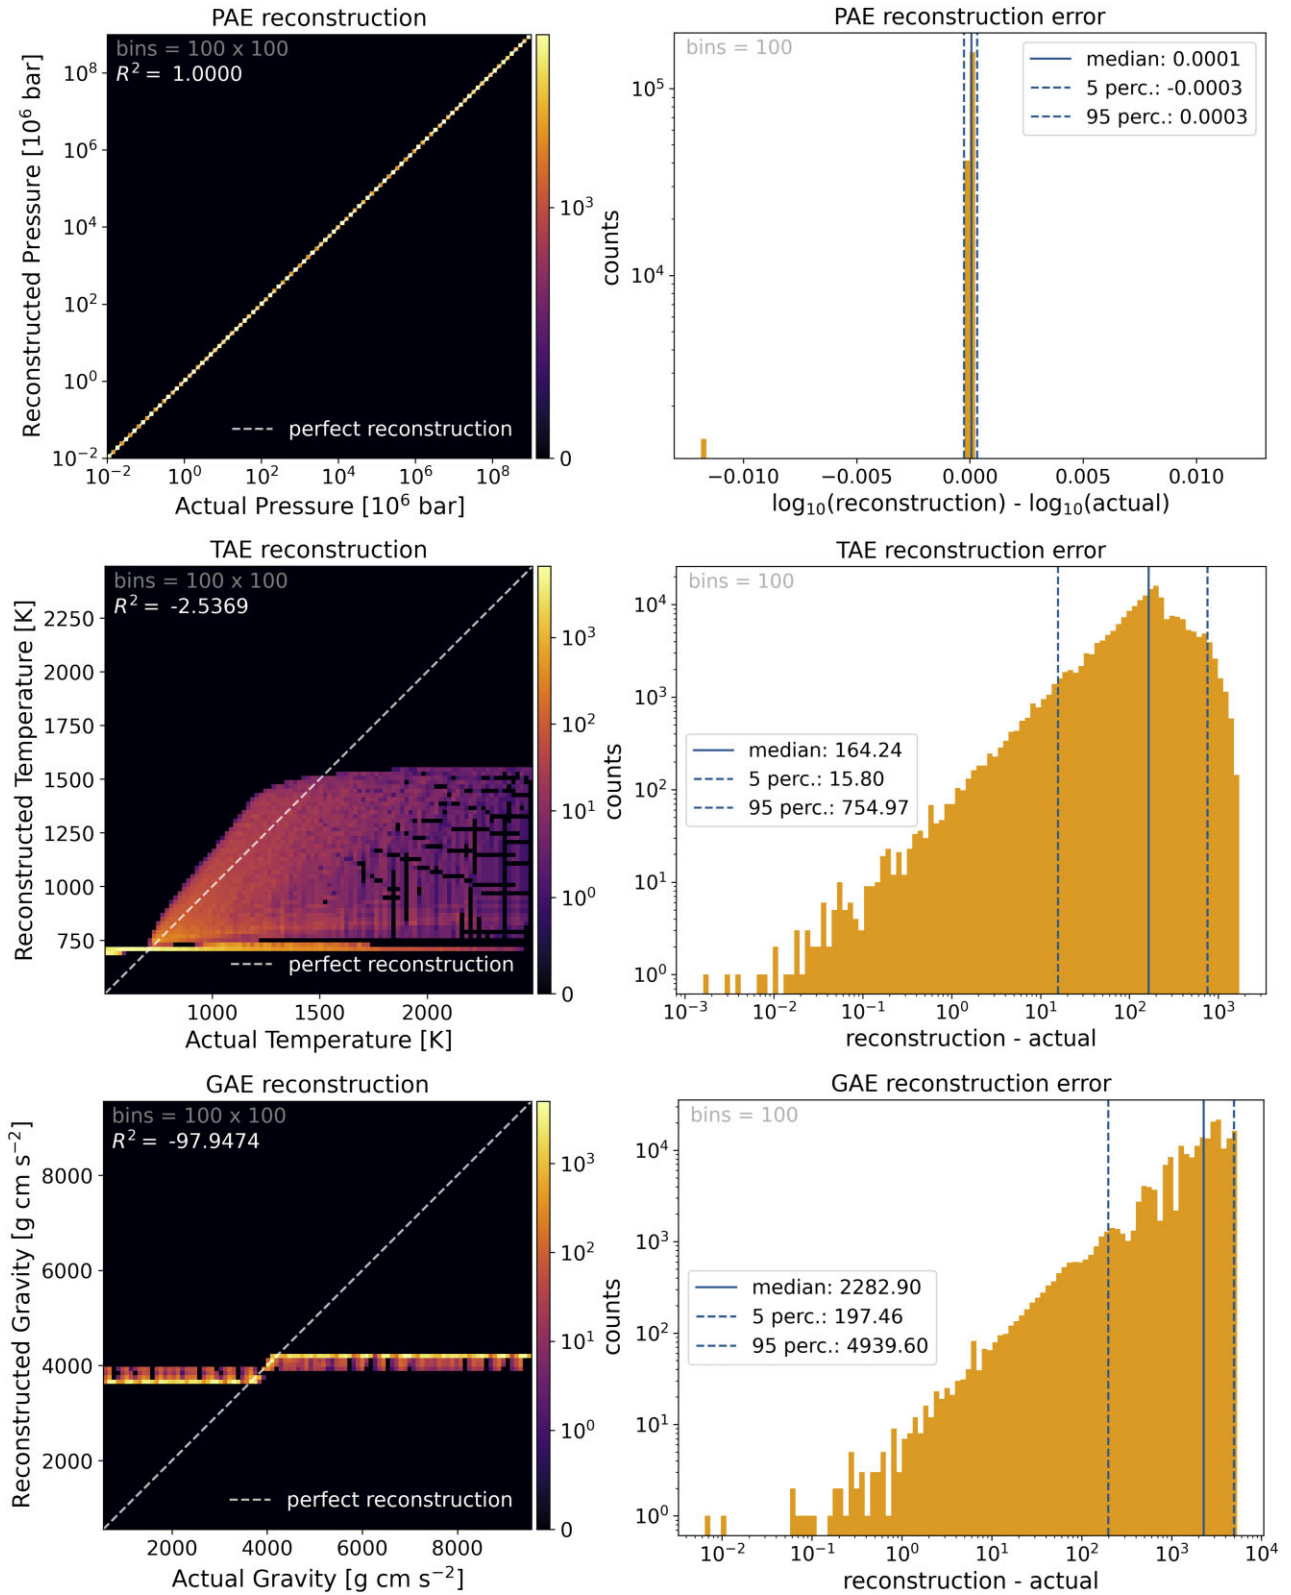


Figure 6. The reconstructed against the actual input values (left column), and the reconstruction error in log space (right column) for the pressure profile (top row), temperature profile (middle row), and the gravity profile (bottom row). Note that the reconstruction errors plot of the temperature and gravity values is calculated in linear space.

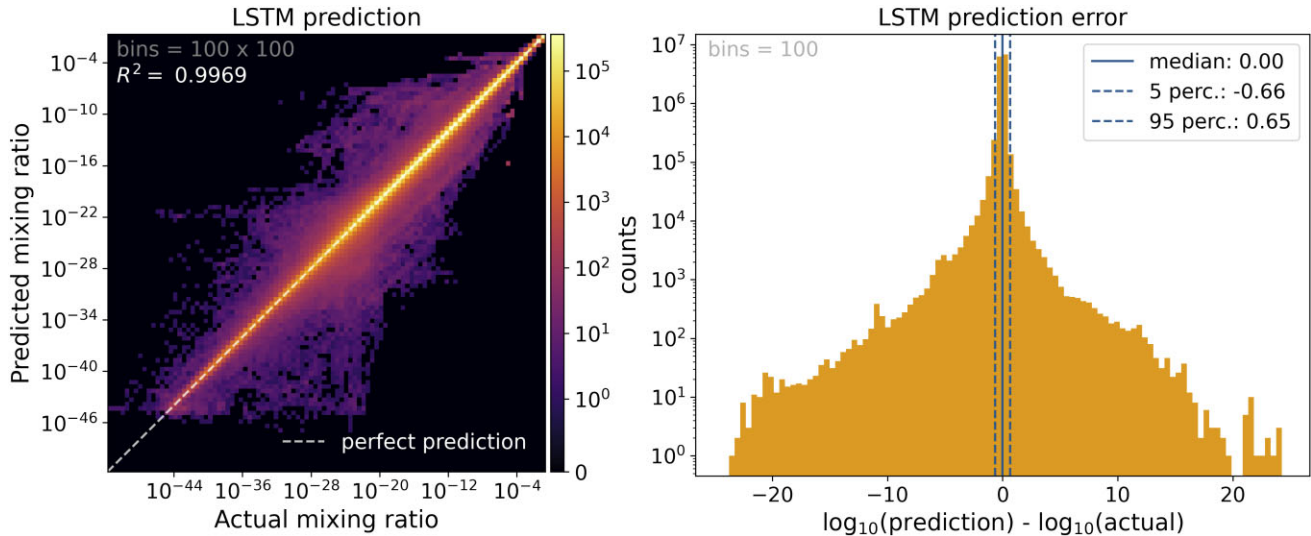


Figure 7. The LSTM predicted mixing ratios plotted against the actual mixing ratios (left) and the LSTM mixing ratio prediction error in log space (right). The dashed diagonal line in the left plot shows the performance of a perfectly predicting model, and the colour of each bin represents the number of predictions. The solid line in the right plot shows the median value, and the dashed lines show the 5th and 95th percentiles.

K-stars have shown to be highly active in the EUV (Reiners & Basri (2008)), as was observed by the MUSCLES collaboration. None the less, the spectra within each method seemed largely similar, which made it more straightforward for the FAE to learn how to accurately reconstruct them. For both the PAE and the WAE, the profiles were linearly spaced in logarithmic space, which made it easy for the AEs to learn how to encode these parameters. It is remarkable, however, that the WAE was not able to perfectly reproduce the wavelengths. A solution would be to make use of a handcrafted algorithm that encodes merely the first and last elements in the array. It is recommended to make use of such an algorithm for future use. Finally, the temperature and gravity profile AEs were not successful at encoding and reconstructing their inputs. Both AEs produced the same solutions for each input example. The limited data set size and large variations in the temperature and gravity example cases could explain the AEs to be prone to errors. Future studies could focus on improving these specific AEs by performing root cause analysis. However, a more specialized approach to encoding the pressure and gravity profiles would be to provide hyperparameters, such as the planet mass and radius, directly to the core model. Such an approach negates the need to train AEs for these input parameters.

The prediction of the core network, LSTM, is within one order of magnitude for the majority (>90 per cent) of the predictions. These errors are comparable to the discrepancies between different chemical kinetic codes (Venot et al. 2012). However, the accuracy of predictions of different examples varies. This inconsistency can arise due to some bias within the data set. Example cases similar to the best-case scenario (see Fig. 8) were more prevalent in the data set, causing the core network to produce better predictions of this type of hot-Jupiters. Additionally, by plotting the loss of each validation case against input parameters (see Fig. 9), it becomes apparent that some specific system parameters perform better than others. From Fig. 9, we see that planets with smaller orbit radii seem to have worse predictions. One explanation could be that these planets endure more irradiation from their host star, ensuring photochemistry is the dominant process in the upper atmosphere. The abrupt and severe changes in abundances for some species due to photodissociation in the upper atmosphere could be difficult for the core network to

learn from with a limited data set as provided in this study. Also noteworthy is the correlation between the planetary mass and the performance of the core network. Higher-mass planets tend to have lower losses as compared to lower-mass planets. Future work could focus on improving the prediction losses of the chemistry profiles for lower-mass planets and planets that orbit their host star close in.

We also showed that the trained model consistently overpredicts mixing ratios that have a value lower than 10^{-44} . This can again be explained by the lack of examples that have such lower values. Species with mixing ratios this low are small contributors to the atmospheric composition nevertheless and are not expected to affect forward models.

Finally, we want to note that the hyperparameters used in this study have been found by trial and error and have not been proven to be the most optimal values. Future studies could focus on a hyperparameter search for each individual AE and core network to find the most optimal parameters.

Due to all mentioned caveats, there is room for improvement in future work. Here, we detail some of the aspects that are out of the scope of this paper, but we will be looking into them in future publications. Evidently, a larger size of the data set, with which the network is trained, is expected to improve the results significantly. In order to train a NN to be more generalized and less biased, a more diverse and extensive data set should be created. Free parameters that can be taken into account, which were not explored in this study, are variables such as the eddy diffusion coefficient,² condensation, and composition of the atmosphere (e.g. varying metallicity and the C/O ratio).

Another approach to improve possibly the results is to change the model itself. The traditional AEs can be replaced with *variational AEs*, VAEs (Kingma & Welling 2013). These types of AEs are based on Bayesian statistics. It is possible to regulate the latent space such that similar input examples have similar latent representations that lie close to each other within the latent space. The core network

²Note that vertical mixing is taken into account in every simulation, but is kept constant throughout the data set.




Figure 8. The mixing ratios per height layer for the best (top), typical (middle), and the worst (bottom) case of the validation set. The planet parameters for each case are given at the top of the plot. The solid lines show the actual mixing ratios as calculated by VULCAN, and the dashed lines show the NN model predictions.

Table 2. Median, minimum, and maximum running times of VULCAN and the NN model for all configurations in the data set. VULCAN runs were performed on a single CPU core using an Intel (R) Xeon (R) CPU E5-4620 0 @ 2.20 GHz. The NN model was run on a single core using an Intel (R) Xeon (R) W-1250 CPU @ 3.30 GHz.

Code	Median	Minimum	Maximum
VULCAN	5994.3 s	1236.7 s	102223.0 s
NN model	0.77 s	0.73 s	0.93 s




Figure 9. The loss as a function of the semimajor axis of each validation case. The colour represents the planet mass in M_J and the size of each scatter point represents the size of the host star which ranges between 1 and $1.5 M_\odot$. The loss is calculated by making use of equation (7).

might then be able to learn how to traverse a regulated latent space and predict more accurately. The core network itself can be improved by, for example, including more time steps within the LSTM. Adding more time steps will ensure that the network predicts the solutions in a more similar way as VULCAN integrates towards the solution. A disadvantage of this is that the training time for the network will increase. Lastly, the recurrent NN architecture could be changed to a *transformer* design. Recently, the transformer NN architecture, proposed by Vaswani et al. (2017), revolutionized the field of sequence transactions within machine learning. By using a so-called *attention* mechanism, the transformer NN outperforms recurrent NNs in accuracy and efficiency. Because of the similarity between transformer and recurrent NN applications, the core model may perform better when changing to this new type of architecture. However, because transformers require tokenized inputs, the AEs will also have to be changed to produce the expected outputs. The implementation of a transformer would therefore increase the complexity of the entire model and should be done carefully.

A different direction in approach would be, for example, to use interpolation methods in the already existing data set. The limitation of such methods is the ease of distribution. The size of the data set used in this study is 600 GB, as opposed to a size of 3GB for the weights of the NN used in this study.

5 SUMMARY AND CONCLUSIONS

In this study, we investigated the ability of a NN to replace the time-dependent ordinary differential equations in the chemical kinetics code VULCAN (Tsai et al. 2017, 2021). The aim of this research was to explore the LSTM architecture for solving ordinary differential equations that include vertical mixing and photochemistry. We

first created a data set that contains the inputs and outputs of VULCAN simulations of hot-Jupiter atmospheres. We made use of the planetary mass ($0.5\text{--}20 M_J$), the semimajor axis ($0.01\text{--}0.5$ au), and the stellar radius ($1\text{--}1.5 R_\odot$) as free parameters. Other parameters for the VULCAN configurations were derived either from analytical relations or kept constant throughout the data set. The input of the data set comprises the initial mixing ratios, the stellar spectrum, the temperature and pressure profiles, and the gravity profiles. Note that the NN trained in this study is limited to the chosen free parameters and cannot be used for atmospheric models that include: e.g. condensation. The outputs of the data set contain the mixing ratios of the species in the atmosphere, taken from 10-time steps (including the steady state) during the VULCAN simulation. This data was used to train a NN that consists of two parts: the *AE* network and the *core* network. The AE was used to reduce the dimensionality of the input and output data from the data set by encoding them into lower dimensional *latent representations*. The AE network consisted of six smaller AEs, designed and trained to encode and decode the mixing ratios, flux, wavelengths, and temperature-, pressure-, and gravity profiles to and from their respective latent representations. The total input latent representation was the concatenation of these six smaller ones. The core network was designed to have an LSTM-based architecture, and it mapped between the latent representation of the inputs to the encoded evolved output by traversing the *latent space* in ten steps. During the training, the latent representations at these ten steps were compared to the ten sets of mixing ratios saved in the outputs of the data set to ensure that the core network is evolving the latent representation in a similar fashion as the VULCAN simulation evolves the mixing ratios. To summarize, we found that:

- (i) the mixing ratios, flux, wavelengths, and pressure profile AEs were able to efficiently encode and accurately reconstruct their respective input properties
- (ii) the AEs were not able to encode and decode the temperature- and gravity profiles successfully. These AEs were, therefore, not used and instead, these profiles were put directly into the latent representation of the inputs
- (iii) the fully trained model (i.e. including the core network) was able to predict the mixing ratios of the species with the errors in the range $[-0.66, 0.65]$ orders magnitude for 90 per cent of the cases. Due to imbalances in the data set, the model is biased to more accurately solve for some examples as compared to others
- (iv) the fully trained model is $\sim 10^3$ times faster than the VULCAN simulations

Overall, this study has shown that machine learning is a suitable approach to accelerate chemical kinetics codes for modelling exoplanet atmospheres.

DATA AVAILABILITY

All simulated data created in this study will be shared upon reasonable request to the corresponding author. The code and results are publicly available on github.com/JuliusHendrix/MRP.

REFERENCES

- Ahrer E.-M. et al., 2022, *Nature*, 614, 653
 Alderson L. et al., 2022, *Nature*, 614, 664
 Bahn G. S., Zukoski E. E., 1960, *Kinetics, Equilibria and Performance of High Temperature Systems: Proceedings of the First Conference*. Butterworths, Los Angeles

- Banks P., Kockarts G., 1973, *Effects of Diffusion in the Heterosphere*, 1st edn., Elsevier, Amsterdam, NL
- Baron E., Chen B., Hauschildt P. H., 2010, PHOENIX: A General-purpose State-of-the-art Stellar and Planetary Atmosphere Code, *Astrophysics Source Code Library*, record ascl:1010.056
- Baxter C. et al., 2021, *A&A*, 648, A127
- Chen J., Kipping D., 2017, *ApJ*, 834, 17
- Cobb A. D. et al., 2019, *AJ*, 158, 33
- Cooper C., Showman A., 2006, *ApJ*, 649, 1048
- Dattilo A. et al., 2019, *AJ*, 157, 169
- Feinstein A. D. et al., 2023, *Nature*, 614, 670
- France K. et al., 2016, *ApJ*, 820, 89
- Grassi T., Nauman F., Ramsey J. P., Bovino S., Picogna G., Ercolano B., 2022, *A&A*, 668, A139
- Heng K., Mendonça J. M., Lee J.-M., 2014, *ApJS*, 215, 4
- Hobbs R., Shorttle O., Madhusudhan N., Rimmer P., 2019, *MNRAS*, 487, 2242
- Holdship J., Viti S., Haworth T. J., Ilee J. D., 2021, *A&A*, 653, A76
- Hu R., Seager S., Bains W., 2012, *ApJ*, 761, 166
- Iess A., Cuoco E., Morawski F., Nicolaou C., Lahav O., 2023, *A&A*, 669, A42
- Jamal S., Bloom J. S., 2020, *ApJS*, 250, 30
- Kawashima Y., Min M., 2021, *A&A*, 656, A90
- Kingma D. P., Ba J., 2014, preprint ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980))
- Kingma D. P., Welling M., 2013, preprint ([arXiv:1312.6114](https://arxiv.org/abs/1312.6114))
- Landman R., Haffert S. Y., Radhakrishnan V. M., Keller C. U., 2021, *J. Astron. Tel. Instr. Syst.*, 7, 039002
- Lloyd R. O. P. et al., 2016, *ApJ*, 824, 102
- Miguel Y., Kaltenegger L., 2014, *ApJ*, 780, 166
- Moses J. I. et al., 2011, *ApJ*, 737, 15
- Nixon M. C., Madhusudhan N., 2020, *MNRAS*, 496, 269
- Reiners A., Basri G., 2008, *ApJ*, 684, 1390
- Roudier G. M., Swain M. R., Gudipati M. S., West R. A., Estrela R., Zellem R. T., 2021, *AJ*, 162, 37
- Rustamkulov Z. et al., 2022, *Nature*, 614, 659
- Shallue C. J., Vanderburg A., 2018, *AJ*, 155, 94
- Smirnov-Pinchukov G. V., Molyarova T., Semenov D. A., Akimkin V. V., van Terwisga S., Francheschi R., Henning T., 2022, *A&A*, 666, L8
- Stock J. W., Kitzmann D., Patzer A. B. C., Sedlmayr E., 2018a, *MNRAS*, 479, 865
- Stock J. W., Kitzmann D., Patzer A. B. C., Sedlmayr E., 2018b, *MNRAS*, 479, 865
- Swain M. R., Vasisht G., Tinetti G., 2008, *Nature*, 452, 329
- Tsai S.-M., Lyons J. R., Grosheintz L., Rimmer P. B., Kitzmann D., Heng K., 2017, *ApJS*, 228, 20
- Tsai S.-M., Malik M., Kitzmann D., Lyons J. R., Fateev A., Lee E., Heng K., 2021, *ApJ*, 923, 264
- Tsai S.-M. et al., 2022, *Nature*, 617, 483
- Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A. N., Kaiser Ł., Polosukhin I., 2017, preprint ([arXiv:1706.03762](https://arxiv.org/abs/1706.03762))
- Venot O., Hébrard E., Agúndez M., Dobrijevic M., Selsis F., Hersant F., Iro N., Bounaceur R., 2012, *A&A*, 546, A43
- Woitke P., Helling C., Hunter G. H., Millard J. D., Turner G. E., Worters M., Blečić J., Stock J. W., 2018, *A&A*, 614, A1
- Xu L., Yan Y.-H., Yu X.-X., Zhang W.-Q., Chen J., Duan L.-Y., 2019, *Res. Astron. Astrophys.*, 19, 135
- Youngblood A. et al., 2016, *ApJ*, 824, 101
- Zelevnik F. J., Gordon S., 1960, *An Analytical Investigation of Three General Methods for of Calculating Chemical Equilibrium Compositions*. NASA Technical Note

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.