EPSC Abstracts
Vol. 12, EPSC2018-343, 2018
European Planetary Science Congress 2018

Atmospheric Retrieval using Deep Learning

T. Zingales (1,2) and I. Waldmann (1)
(1) University College London, London, United Kingdom (2) INAF - Osservatorio Astronomico di Palermo, Palermo, Italy
(tiziano.zingales.15@ucl.ac.uk)

Abstract

Atmospheric retrievals on exoplanets involve usually computationally intensive Bayesian methods. The choice of the fitting parameters bounds are often leaded by physical constraints and and the user experience. We introduce ExoGAN, a new generation artificial intelligence able to recognise molecular features, abundances and atmospheric physical parameters using unsupervised learning. ExoGAN will return a probability distribution for each parameters that can be used either as a final atmospheric analysis or as prior distribution for a subsequent Bayesian model.

1. Introduction

Artificial Intelligence has been used extensively in the last few years in many different fields. The use of Artificial Neural Networks in Astrophysics is a relatively young but the growing interest of the scientific community towards this tool will increase dramatically within the next decade. Independently on the research field, Neural Networks can be used to understand and describe relatively complex structures and behaviour in a wide variety of dataset. Waldmann (2016) is a pioneering work who apply a deep-belief neural network (DBN) to recognize the atmospheric features on an exoplanetary transmission spectrum. Rodriguez et al. (2018) developed a tool able learn models that can efficiently generate new, physically realistic realizations of the cosmic web using generative networks. In the exoplanetary field, and more in general, in the astrophysics field, the use of GANs is still a new method able to solve computationally intensive problems. In our work we suggest a more advanced and up-to-date unsupervised algorithm to understand and reliably reproduce the atmosphere of an exoplanets. In particular we used a Deep Convolutional Generative Adversarial Network (DCGAN) introduced for the first time by Goodfellow et al. (2014). We introduce ExoGAN, a DCGAN able to recognise spectral feature from exoplanetary spectra and return a detailed chemical and physical analysis in a completely new way, reducing the computational time from several hours to a few minutes.

2. What is a DCGAN

A GAN has two different neural networks competing each other and learning how to reproduce realistic synthetic data from an input dataset. A Deep Convolutional GAN (DCGAN) is a GAN which uses batch normalisation, it is made of two all-convolutional networks, uses the Adam Optimizer and the leaky ReLU activation function (Xu et al., 2015; Radford et al., 2015).

3. The training

A good training set is crucial to teach ExoGAN how to generate a realistic transmission spectrum. GANs are general methods, recently they have been applied to several serious problems, such as semi-supervised learning, stabilizing sequence learning methods for speech and language, and 3D modelling (Denton et al., 2015; Radford et al., 2015; Salimans et al., 2016; Lamb et al., 2016; Wu et al., 2016). However, they still remain remarkably difficult to train, with most current papers dedicated to heuristically finding stable architectures (Arjovsky and Bottou, 2017). We use a training set of 10 million of spectra varying 7 different parameters: H_2O, CO_2, CH_4 and CO abundances, the mass of the planet M_p, the radius R_p and the temperature T_p. In order to stretch as much as possible every feature we divide the input spectrum into several different bins as shown in Fig 1 We normalised each of these spectra between 0 and 1 and, to optimise the efficiency of our GAN to recognise the features and the correlations between the parameters we normalised each parameters as following (see Fig 2)

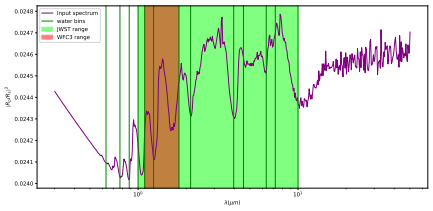

Figure 1: Spectral binning used in this work. The green vertical line are represent the bin edges according to some water features. the Red area is the wavelength range related the the Hubble WFC3 camera and the lime area is that of the JWST.

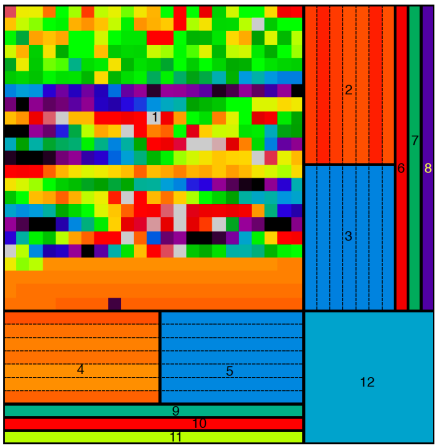

Figure 2: Normalised spectrum. Each area is dedicated to a particular atmospheric characteristic: Area **1** is the spectrum between 1μm and 50μm at resolution 100 normalised between 0 and 1 in each spectral bin. From area **2** to **5** give information about the normalisation factors used in the different section of the spectrum, clear and dark area give, respectively, information about the maximum values and the minimum values. In areas from **6** to **8** there are, respectively, CO_2, CO and CH_4 abundances. Areas **9** to **11** are, respectively M_p, R_p and T_p. Area **12** gives information on the H_2O abundance.

3.1. Method - Generative Adversarial Networks

GANs are classified as unsupervised learning algorithms. In our work we use it to recognise spectral features from exoplanetary spectra. The very same net, nevertheless, can be used to detect the characteristics of an image (Goodfellow, 2017; Creswell et al., 2018). GANs consist of two neural network facing each other and learning how to reproduce a realistic synthetic data from an input dataset. The two essential parts of a GAN are de Discriminator and the Generator nets (Fig 3).

4. The training

In order to update the weights related to the two neural nets we need to differentiate this function with respect to the discriminator and the generator. Concerning the generator training, the gradient of the first term with respect to the generator is zero (the generator does not appear there) so only the second term is relevant. Assuming that the discriminator does a really good job in discriminating the real and the fake images, it means that $D(G(z))$ is very close to zero, the slope of the cost function at point 0 is also very close to zero and the θ_g has no chance to improve and change. It means that in this case the network does not learn anything from the training set. For the generator, a possible solution to this problem could be, instead of targeting the value 0 for fake images, looking at the value 1 target (real images), it means that we are trying to minimise the negative expected value of $\log D(G(z))$ and so we can use as a cost function for the generator:

$$J^{(G)} = \mathbb{E}_{\mathbf{z}\sim p_z}[\log(D(G(z)))] \tag{1}$$

This new solution works because the term $\log D(\text{anyimage})$ corresponds to a target of value 1 in the binary cross-entropy and the other term is irrelevant because its derivative in the gradient descent is zero. For each number of training iteration and for k steps, defining as m as the batch size of or set, we update the discriminator by ascending its stochastic gradient (Goodfellow et al., 2014):

$$\nabla_{\theta_d}\frac{1}{n}\sum_{i=1}^{n}\left[\log D\left(\mathbf{x}^{(i)}\right)+\log\left(1-\log D\left(G\left(\mathbf{z}^{(i)}\right)\right)\right)\right] \tag{2}$$

At the same time we update the generator by de-

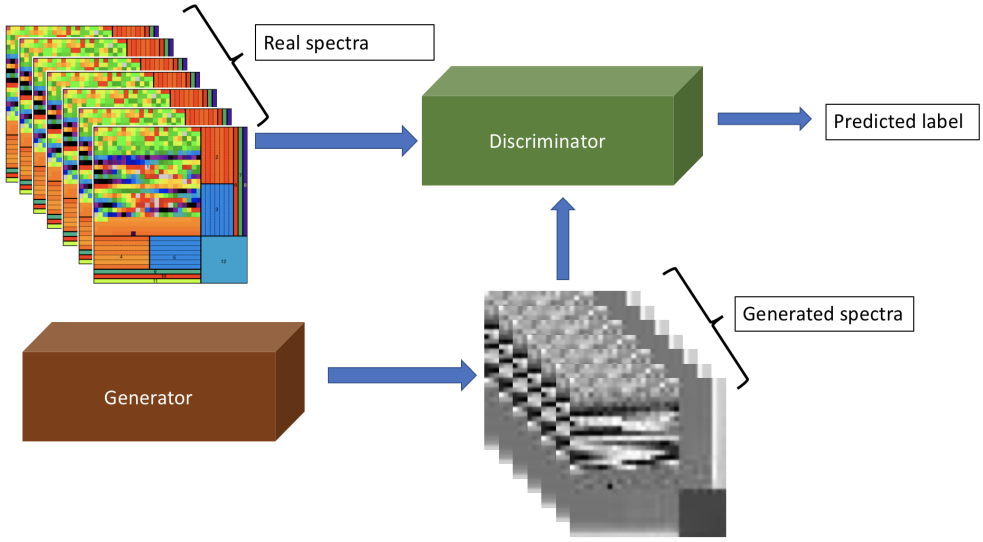

Figure 3: ExoGANscheme. From right to left, The generator reproduce a realistic spectral sample that is seen by the discriminator. The discriminator network is the only one who have access to the real dataset. By analysing both the generated sample and the real one it tries to understand whether the generator is able to reproduce a perfectly realistic spectral sample or not.

scending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{n} \sum_{i=1}^{n} \log \left(1 - \log D \left(G \left(\mathbf{z}^{(i)}\right)\right)\right) \quad (3)$$

5. Image reconstruction

Atmospheric Retrievals are often computing intensive and they need to be done by fitting many parameters. A Bayesian model on exoplanetary atmosphere parameters are done using non-informative (flat) priors within a range of values. These ranges are often fixed according to physical constraint and, once one have experience on exoplanetary atmosphere, user experience. The use of Artificial Neural network can replace the user experience with a more sofisticated and completely unsupervised tool. The use of ExoGAN, trained on a huge dataset of exoplanetary atmospheres models can help the retrieval code to have better constraint, get rid of the solutions that are likely not to give any acceptable solution and accelerate the computing analysis. In order to find the best $\hat{z}$ we define two loss functions for an arbitrary $\hat{z} \sim p_z$: the contextual loss and the perceptual loss:

$$\mathcal{L}_{contextual}(z) = \| M \odot G(z) - M \odot y \|_1 \quad (4)$$

with $\| x \|_1 = \Sigma_i |x_i|$ for some vector x.

$$\mathcal{L}_{perceptual}(z) = \log\left(1 - D(G(z))\right) \quad (5)$$

We find $\hat{z}$ defining a combination of the two losses:

$$\mathcal{L} = \mathcal{L}_{contextual}(z) + \lambda \mathcal{L}_{perceptual}(z) \quad (6)$$

where λ is a hyper-parameter which control how important is the contextual loss compared to the perceptual. At this point $\hat{z}$ is defined as:

$$\hat{z} = \arg\min_z \mathcal{L}(z). \quad (7)$$

The best reconstructed image defined in **??** is the one that uses the $\hat{z}$ defined in 7. In Fig 4 we show the three phase associated to a prediction. We see that the input parameters are masked and reproduced by a pre-trained ExoGAN.

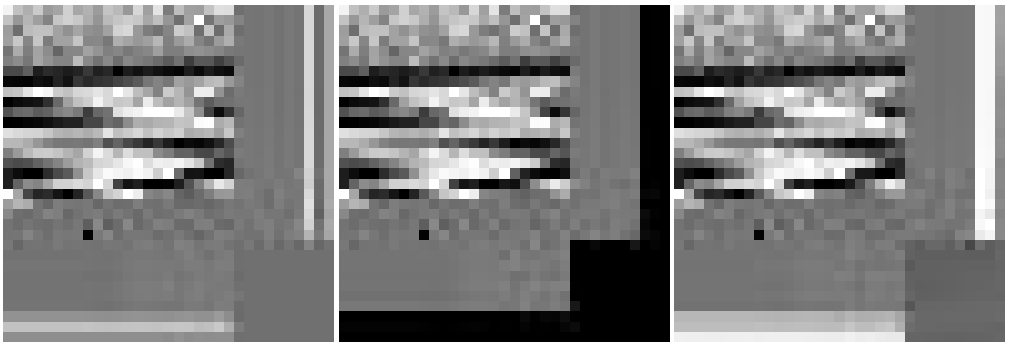

Figure 4: On the **left** we find the input spectrum together with the parameters pixels. In the **centre** there is the masked spectrum which is fed to the ExoGAN and on the **right** we find the predicted spectrum with the best pixels parameters.

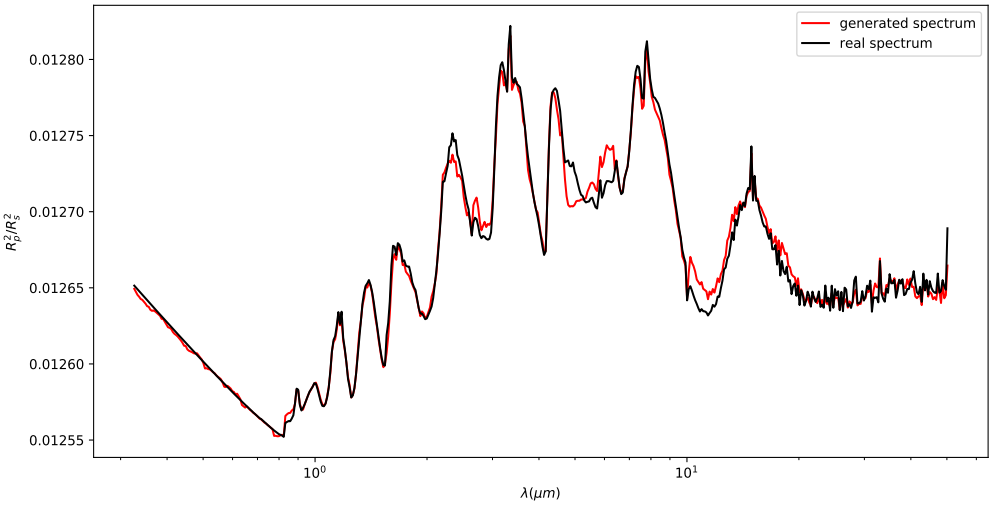

Figure 5: Feature recognition of ExoGAN. The black line is the input spectrum and the red one is the best generated one. ExoGAN recognised all the features of the input spectrum associating to it the parameters that can generate it with a classical model.

6. Results on Test set

Test set parameters			
Variable	accuracy	$\tilde{\chi}^2$ (1σ)	$\tilde{\chi}^2$ (2σ)
CO	50.2%	5.83	1.46
$CO2$	85.3%	0.85	0.21
$H2O$	79.6%	2.64	0.66
$CH4$	68.8%	0.46	0.11
Rp	98.4%	0.03	0.01
Mp	71.8%	1.59	0.40
Tp	75.5%	2.41	0.60

Table 1: Accuracies and $\tilde{\chi}^2$ associated to each parameters for 1000 test spectra. The 2 column represent the absolute accuracy of the prediction without taking into account the error bar of the prediction. The 2nd and 3rd column are the $\tilde{\chi}^2$ calculated, respectively, using one σ and 2 σ prediction errors.

7. Summary and Conclusions

We demonstrated how the use of a DCGAN can help the spectral retrieval of exoplanetary atmosphere. DCGAN in our work are 210% faster than a Bayesian analysis on the same spectrum. The output parameters distribution can be used either as a final solution or as an input prior distribution for a more efficient Bayesian modelling.

Acknowledgements

T. Z. is supported by the European Research Council ERC projects *ExoLights* (617119) and from INAF trough the "Progetti Premiali" funding scheme of the Italian Ministry of Education, University, and Research.

References

Arjovsky, M. and Bottou, L. (2017). Towards Principled Methods for Training Generative Adversarial Networks. *ArXiv e-prints*.

Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., and Bharath, A. A. (2018). Generative Adversarial Networks: An Overview. *IEEE Signal Processing Magazine*, 35:53–65.

Denton, E., Chintala, S., Szlam, A., and Fergus, R. (2015). Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks. *ArXiv e-prints*.

Goodfellow, I. (2017). NIPS 2016 Tutorial: Generative Adversarial Networks. *ArXiv e-prints*.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Networks. *ArXiv e-prints*.

Lamb, A., Dumoulin, V., and Courville, A. (2016). Discriminative Regularization for Generative Models. *ArXiv e-prints*.

Radford, A., Metz, L., and Chintala, S. (2015). Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *ArXiv e-prints*.

Rodriguez, A. C., Kacprzak, T., Lucchi, A., Amara, A., Sgier, R., Fluri, J., Hofmann, T., and Réfrégier, A. (2018). Fast Cosmic Web Simulations with Generative Adversarial Networks. *ArXiv e-prints*.

Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved Techniques for Training GANs. *ArXiv e-prints*.

Waldmann, I. P. (2016). Dreaming of Atmospheres. *ApJ*, 820:107.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, Ł., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., and Dean, J. (2016). Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *ArXiv e-prints*.

Xu, B., Wang, N., Chen, T., and Li, M. (2015). Empirical Evaluation of Rectified Activations in Convolutional Network. *ArXiv e-prints*.