자기-지도학습 적대적 모델과 이상 감지를 이용한 금성의 정적인 대기의 파동 감지

Detection of Stationary Atmospheric Waves in Venus with a Self-Supervised Adversarial Model Using Anomaly Detection

Abstract

In this paper, we propose an anomaly detection scheme for identifying stationary waves in Venus' atmosphere using a self-supervised model. Initially, we stack multiple images of projected maps of Venus to filter out other types of features, highlighting stationary waves. We split each image into 72×72-resolution patches and designated patches with waves as anomalies. In contrast, the rest of the grids are designated as normal data. We create a variational autoencoder-based, adversarially trained anomaly detection model. We train the model using a portion of the normal data and test it with the remaining normal and anomaly data. According to the anomaly hypothesis, the model will successfully reconstruct the trained data and fail on the "anomaly" designated untrained data. The results show that the model can differentiate between stationary waves and cloud formations, with an AUC score of 78.37%.

1. Introduction

Venus is the second planet from the Sun and is similar in size and bulk chemical composition to Earth. However, its atmosphere has evolved in a drastically different way [2]. It is mainly composed of CO_2 (carbon dioxide) and other minor compounds like SO_2, which can condense to form clouds at approximately 48-70 km of altitude [3]. The atmosphere near the cloud top flows at a speed of about 100m/s, around 60 times faster than the solid globe. The surface temperature on Venus is incredibly high, reaching up to 460°C due to a runaway greenhouse effect caused primarily by the aforementioned atmospheric composition and an atmospheric mass that can reach 100 times that of Earth's [4].

Atmospheric waves on Venus are created as disturbances that propagate through the atmosphere due to changes in temperature, pressure, and other atmospheric variables. In Venus' atmosphere, these waves can produce patterns and oscillations that could be detected. Detecting atmospheric waves in the atmosphere of Venus could provide valuable insights into the planet's atmospheric dynamics and further characterization of the phenomena itself. Moreover, the detection and characterization of stationary features on the atmosphere allows a better understanding of possible connections between the surface topography and atmospheric circulation.

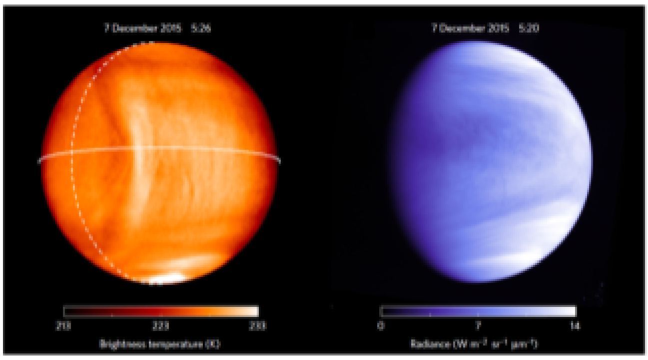

Fig. 1 Venus images taken by Akatsuki space probe on December 7th, 2015, using long-infrared wavelength (on the left), and at ultraviolet wavelength (on the right)[8].

Detection of atmospheric gravity waves on Venus has generally been done by experts. Since most wave features have a low frequency on most observed wavelengths, various processing methods are usually applied to the images to identify waves [1]. However, with increased amounts of high-quality data provided mainly by space missions such as the ongoing Akatsuki space probe [5], smaller features can now be identified in significant numbers. This leads to manual identification becoming a time-consuming process, which can prove impractical and possibly lead to increased human error, such as false positives or overlooking critical atmospheric gravity wave features.

Due to the overwhelming quantity of images with multiple cloud features, we propose a self-supervised anomaly detection scheme to detect atmospheric gravity waves as anomalies on Venus using long-range infrared (LIR) data from the Akatsuki space probe. The results show that the proposed preprocessing makes it possible to highlight the atmospheric gravity waves, and the proposed model successfully discriminates between normal features and atmospheric gravity wave features.

2. Data Preprocessing and Anomaly Detection Model

The LIR images taken by the Akatsuki probe, provided by the Japan Aerospace Exploration Agency (JAXA), are publicly available (https://darts.isas.jaxa.jp/doi/vco/vco-00012.html). The dataset consists of consecutive images with varying time separations between them, ranging from a few seconds to several hours. As many of these features are of low frequency and our target features are stationary relative to the background, we performed image stacking on consecutive images with an existing wave to highlight its features better and smooth out the remaining parts of the image.

We create 11 stacked images for the dataset creation from the days when the stationary waves were clearly identified between May 2016 and January 2017, by manual investigation, as reported in [6]. After applying the smoothing operation (stacking) to eliminate other low-frequency features, we apply a high-pass filter to highlight the stationary wave features. We split these stacked images into smaller images using a grid size of $72 \times 72\, pixels$, which we observed to fit the best for the stationary waves on our dataset.. We gathered 399 images, of which 33 contain stationary waves. The normal image data sample size reaches 2314 after applying horizontal and vertical flips, random 10-degree tilt, and random zoom operation between 80%-120%. Since the zonal wind dominates the atmospheric flow at the top of clouds, applying horizontal flips to these images would not be reasonable for the specific case of Venus. However, due to the approximate symmetry in the dynamical profile of the northern and southern hemispheres, a vertical flip operation can be performed.. This doubles the available stationary wave image number for testing. An example set of images containing "normal" cloud formations and stationary waves can be seen in Fig. 2. The "normal" terminology simply refers to all other images that don't feature the specific morphological features of stationary waves.

The anomaly detection model is an encoder-decoder-based, adversarially trained generative model, including a Discriminator. The input size of the model is 72 × 72 grayscale images, and the encoder consists of four cascading convolutional and max-pooling layers. The feature vector of length 256 becomes the input for the decoder, which has four upsampling and convolutional layers instead of transpose convolution due to observed checkerboard formations in generated images. The decoder generates the image with the exact resolution of the input image. The Discriminator in its architecture is identical to the encoder, except it also has a sigmoid layer that classifies the input image as fake and leaky rectified linear unit (l-ReLU) activation operation on each layer. During the training, the generator creates more realistic images, and the Discriminator classifies the images with better accuracy.

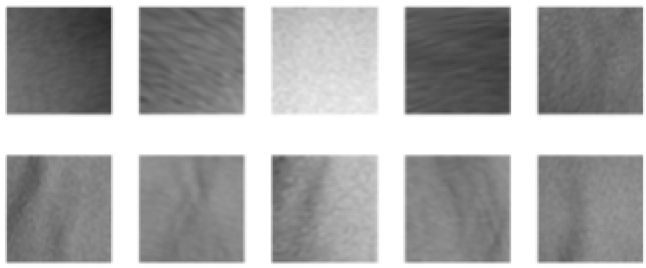

Fig. 2 Normal cloud formation examples on top, and stationary wave examples on bottom five images.

For robust training, we introduce three loss functions to the model:

1- Contextual Loss: We use the L1 loss between the input image x and the reconstruction $\hat{x}$ for high-level feature distribution learning by the model, which helps the generator create better contextual images:

$$L_{context} = E_{x\sim p_x}|x - \hat{x}|_1$$

2- Latent Loss: We apply L2 loss between the input image's feature vector and the reconstructed image's feature vector in the Discriminator's fully-connected layer. The loss is as follows:

$$L_{latent} = E_{x\sim p_x}|f(x) - f(\hat{x})|_2$$

3- Adversarial loss: We apply Wasserstein loss [7] on the Discriminator, which seeks a minimized distance between the real and the generated data distribution.

This makes the Discriminator a critic that scores the "realness" or "fakeness" of an input image rather than classifying it as authentic or fake. The loss is as follows:

$$L_{wasserstein} = - mean(D(x)) + mean(D(G(z)))$$

where $D(x)$ is the output of the Discriminator for real input images x, and $G(z)$ is the generator's output for fake images generated from noise z. The mean is taken over the batch of input images.

The total loss of the model is the weighted sum of three loss functions, as shown below:

$$L_{total} = 40 \times L_{context} + L_{latent} + L_{wasserstein}$$

The optimal weighting is found after a grid search operation.

3. Model Training and Anomaly Detection

We split the “normal” designated cloud formation data into 80% train and 20% test. The hypothesis is that the model learns the high-dimension and latent space representation of normal clouds during training. Therefore, we train the model using only "normal" designated data. We train the model in an unsupervised fashion for 100,000 iterations using Adam optimizer and a learning rate of 10^{-4} until the loss values stabilize and the reconstructed images look similar to the original data. We do not use "anomaly" designated stationary wave data during the training.

We use the area under the curve (AUC) of the receiver operating characteristics (ROC) metric to calculate the performance of the Discriminator. The model achieves a 78.37% AUC score despite the low number of data and hard-to-distinguish features. We do not employ thresholding metrics such as the F1 score since the anomaly-designated images are only available during the inference.

Anomaly detection hypothesizes that the model is expected to reconstruct the cloud formation well and is expected to fail in reconstructing stationary waves since it has not learned high-dimension or latent representation of the anomaly data. Taking advantage of this discrepancy, we calculate the anomaly score for each test image with the formula below:

$$A(\hat{x}) = (1 - \lambda) \times R(\hat{x}) + \lambda \times L(\hat{x})$$

where $R(\hat{x})$ is the $L1$ loss between the input and the generated image, $L(\hat{x})$ is the $L2$ loss between the latent vectors of input and the generated images in the Discriminator’s fully connected layer, and λ is the weighting factor, for which we designate as $\lambda = 0.2$. Following min-max normalization of the anomaly score vector Λ, the values scale within the range [0, 1]:

$$\Lambda(\hat{x}) = \frac{A(\hat{x}) - min(A)}{max(A) - min(A)}$$

We plot the anomaly score distributions of cloud and stationary waves as a qualitative metric. The resulting image can be shown in Fig. 2.

4. Conclusion

In this paper, we propose an anomaly detection approach to detect stationary atmospheric waves in Venus' atmosphere. Our model is a variational autoencoder-based, adversarially trained, self-supervised generative model. We stack long-infrared wavelength images obtained from the Akatsuki probe, where the stationary waves appear to highlight the wave features and reduce other features, apply smoothing and high-pass filters, and split each image into 72 × 72 grids. We manually select grids with stationary waves and designate them as "anomalies." We train the model with 80% of the data normal cloud data. During inference, we test the Discriminator for quantitative scoring, which results in a 78.37% AUC score, and plot the anomaly distributions between the normal and stationary wave images. For future work, we will compare the anomaly detection performance between different wavelengths (e.g., ultraviolet) and use other probe data with a more complex model for better performance. Although the techniques applied in image preprocessing were suitable, there remains another possibility for improvement with extended analysis.

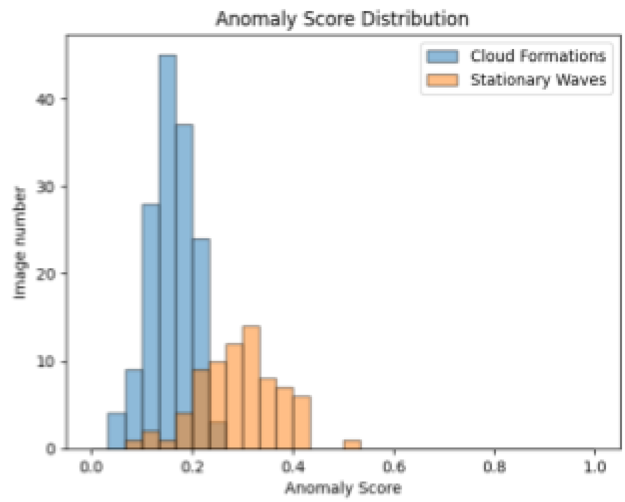

Fig. 3 Anomaly score distribution for regular cloud formations and stationary waves. Note that scores close to 0 indicate normality, and scores close to 1 indicate anomaly.

References

[1] Kitahara, T. et al. (2019), “Stationary features at the cloud top of Venus observed by Ultraviolet Imager onboard Akatsuki,” Journal of Geophysical Research,

[2] Svedhem, H., et al. (2007) – “Venus Express – The first European mission to Venus” – Planetary and Space Science 55, pp. 1636–1652.

[3] Titov, D. et al. (2018) – “Clouds and hazes of Venus” – Space Science Reviews 214:126.

[4] McFadden, L.A., Weissman, P.R., Johnson, T.V., (2007) – “Encyclopedia of the Solar System” – Second Edition, Elsevier Inc.

[5] Nakamura, M., et al. (2011), –“Overview of Venus Orbiter, Akatsuki” – Earth, Planets and Space 63, pp 443–457.

[6] Kouyama, T., et al. (2017), – “Topographical and local time dependence of large stationary gravity waves observed at the cloud top of Venus” – Geophysical Research Letters, 44, 12098–12105.

[7] Arjovsky M., Chintala S., and Bottou, L. (2017), – “Wasserstein generative adversarial networks,” – Proceedings of the 34th International Conference on Machine Learning,
vol. 70, 214–223.

[8] Fukuhara, T., et al. (2017) – “Large stationary gravity wave in the atmosphere of Venus” – Nature Geoscience, vol. 10, Iss. 2, pp. 85–88.