# **Conditional Denoising of Remote Sensing Imagery Using Cycle-Consistent Deep Generative Models**

Michael Zotov\* Jevgenij Gamper<sup>†</sup> Cervest Ltd. London, UK

#### Abstract

The potential of using remote sensing imagery for environmental modelling and for providing real time support to humanitarian operations such as hurricane relief efforts is well established. These applications are substantially affected by missing data due to non-structural noise such as clouds, shadows and other atmospheric effects. In this work we probe the potential of applying a cycle-consistent latent variable deep generative model (DGM) for denoising cloudy Sentinel-2 observations conditioned on the information in cloud penetrating bands. We adapt the recently proposed Fréchet Distance metric to remote sensing images for evaluating performance of the generator, demonstrate the potential of DGMs for conditional denoising, and discuss future directions as well as the limitations of DGMs in Earth science and humanitarian applications.

## 1 Introduction

Climate variability induced impacts are becoming increasingly important for land resource management for commercial aims and social good. One mechanism for greater insights is employment of remote sensing imagery [15, 5, 18, 6, 22, 12]. However, Earth observation data in high temporal and spatial resolution is not widely available. The effectiveness and accuracy of remote sensing based solutions (for instance, machine learning solutions for crop yield prediction) is hindered by the frequent appearance of atmospheric noise in such data, in the form of outliers induced by small shadows or missing data due to clouds. Besides the social benefit of efficient use of land resources through remote sensing, multiple humanitarian missions directly benefit from the support of satellite Earth observations [16]. Thus publicly available missions like Sentinel-2, MODIS, Landsat are of crucial importance, as these are accessible to non-profit or research organisations. However, it has been shown that even with partially clouded skies there is just a 25% chance that an arbitrary active field will appear in daily optical satellite imagery [21], which highlights the influence of cloud cover on the availability of optical imagery.

One approach that has been widely adopted is cloud masking. However, obtaining pixel-based segmentation ground truth in remote sensing is not as straightforward as for natural images. Current methods mostly resort to computationally intensive models that use temporal information or combine multiple physics-based models to create supervised learning data, analogous to the idea of data programming [7, 1, 24]. Furthermore, recent work has demonstrated that for the purpose of crop classification, deep learning models could deal with cloud induced outliers and still obtain sufficient performance [17]. While these approaches prevent only some outliers in the downstream tasks, it

 $<sup>^{*}\</sup>mbox{Work}$  completed as part of the Cervest Research Residency program. Correspondence to  $\mbox{jev@cervest.earth}$ 

<sup>&</sup>lt;sup>†</sup>Mentor

AI for Social Good Workshop (Track 1), 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

does not solve the problem of missing data due to clouds, which fundamentally undermines the benefit of higher temporal resolution.

In this work, we therefore propose a novel treatment of surface reflectance data denoising from atmospheric artifacts as an unpaired translation from the domain of cloudy/shaded images X to the domain of clear images Y. We contribute by: (i) demonstrating preliminary results of applying cycle-consistent deep generative models to this problem, (ii) building a deep convolutional neural network to obtain state of the art performance in land cover classification on the existing public dataset, and (iii) developing a method to assess the fidelity of the denoised remote sensing images by calculating the Fréchet Distance between the distributions of generated and real images in the classifier feature space.

## 2 Methods and Results

#### 2.1 Remote sensing data

Sentinel-2 (S-2) is a European Space Agency program that provides surface reflectance data. Compared to other existing programs such as Landsat and MODIS, it provides 6-8 day temporal resolution at higher spatial resolution per pixel of 60m to 10m, depending on the band. S-2 samples data in 13 channels in the visible, near-infrared and short wave infrared spectral range. For this study, whilst the raw S-2 data is publicly available<sup>3</sup>, we chose a curated dataset created from S-2 observations, that has been established to standardise land cover classification performance in remote sensing - BigEarthNet [20]. The dataset contains 590,326 images with multi-class labels of 43 ground cover categories and demonstrates the challenges of supervised learning for remote sensing classification tasks: some labels are present in over 200,000 examples while others appear in less than 1,000 examples.

While there are other standardised public surface reflectance datasets, BigEarthNet also provides a breakdown of the dataset into images containing clouds and cloud shadows, and clear images<sup>4</sup> with respective cardinality 9,280 and 519,339 (images with seasonal snow are excluded) [9]. This breakdown is beneficial to this problem, and represents the general challenge of cloud and shadow removal in remote sensing: it is possible to identify images that contain clouds or cloud shadows, but it is not easy to obtain pixel-wise ground truth to train segmentation models.

#### 2.2 Land cover classification for denoising evaluation

The performance of DGMs is difficult to measure [19], particularly in the context of remote sensing imagery, as an image that "appears" realistic to the human eye may still contain outliers for downstream tasks. Furthermore, given the social and economic importance of the task at hand, these models require a robust testing approach. As such we adopt the Fréchet Distance (FD) based score, commonly used to evaluate DGMs, to remote sensing [10]. FD captures similarity between the two curves. To evaluate DGMs using FD, one needs to train a sufficiently powerful image classifier, whose features would parametrise multivariate normal distributions of generated and real samples in the classifier feature space, and then compare the two distributions using FD. We therefore train an 18 layer ResNet [8] on the BigEarthNet multi-label classification task, and achieve a state of the art performance. Instead of sigmoid activations in the final layer commonly used in multilabel classification, we use softmax activations [11]. To regularise against dataset imbalance during training we randomly set images to only contain a single label with a probability of 0.5, selected uniformly from the original list of labels for that image. Our results in land cover classification, compared to the model used by the creators of the dataset, are presented in Table 1.

Model	Precision(%)	Recall(%)	$F_1$	$F_2$
S-CNN-All [20]	70	77	0.70	0.74
Regularized ResNet (our model)	85	77	0.79	0.77

Table 1: Experimental results of training the landcover classifier

<sup>&</sup>lt;sup>3</sup>https://registry.opendata.aws/sentinel-2/

<sup>&</sup>lt;sup>4</sup>http://bigearth.net/

We then proceed to establish a baseline validation FD score for our task. We computed a FD score of 228 for cloudy versus clear validation sets.



Figure 1: a) The first column corresponds to training images in domain X, middle column is the same images mapped to domain Y, and in the third column these images are restored back to domain X for the evaluation of cycle-consistency loss. Images were captured at the 100th epoch, with FD of 106. b) Fréchet Distance over the course of training of an Aug CycleGAN.

#### 2.3 DGM for cloud and shadow removal

We approach the problem from a novel perspective of unpaired image-to-image translation between the domain of cloudy/shaded images, X, and clear images Y. We acknowledge the randomness in spatial distribution of clouds or shadows by introducing a latent variable  $\mathbf{z}_{\mathbf{x}}$ , and the distribution over possible land cover categories underneath the cloud as  $\mathbf{z}_{\mathbf{y}}$ . As such, we have two joint distributions describing the respective pairs  $(\mathbf{x}, \mathbf{z}_{\mathbf{y}}) \sim p_d(\mathbf{x})p(\mathbf{z}_{\mathbf{y}})$  of cloudy/shaded images and the possible land classes underneath, and  $(\mathbf{y}, \mathbf{z}_{\mathbf{x}}) \sim p_d(\mathbf{y})p(\mathbf{z}_{\mathbf{x}})$  of clear images and possible locations of clouds and shadows respectively. Our task is to learn the implicit, parameterised generative models  $p_{\theta}(\mathbf{x}|\mathbf{y}, \mathbf{z}_{\mathbf{x}})$ and  $p_{\phi}(\mathbf{y}|\mathbf{x}, \mathbf{z}_{\mathbf{y}})$ . By construction, every cloudy image in domain X, when mapped into the clear domain Y, is conditioned on the observed cloud penetrating bands. While in this preliminary work the model construction is not ideal, in future work we will include radar based bands that in fact do contain significant information about the underlying land cover and are not affected by clouds [2].

To infer the parameters of the implicit models we adopt the inference scheme of Augmented Cycle-Consistent Generative Adversarial Networks (Aug CycleGAN) [3]. Namely, we utilise 8 deep neural networks: two Generator - Discriminator pairs, where generators learn maps  $X \times Z_y \to Y$ and  $Y \times Z_x \to X$  respectively, and two Latent Encoder - Discriminator pairs. We add Spectral Normalization to the convolutional layers in the architectures [13], and optimise the model using the loss functions described by the authors of the original paper. The model is optimised to map a sample between domains and then recover it consistently (see Figure 2a), where cycle-consistency is enforced by the L1 norm. During training, we iterate over a subset of cloudy images, and select a batch of clear images at random (from a training subset) for each batch of cloudy images. The validation set consists of fixed pairs, for better performance comparison between epochs. We compute FD for validation set after every epoch (see Figure 1b).

In Figure 2b we show that the generator is able to capture information from cloud penetrating bands to inpaint the land segment covered by the cloud (e.g. images [1,3] (row 1 column 3) and [1,1]) as well as address shadows (image [3,3]), and preserve the landscape, but struggles with the thicker parts of the cloud, as might be expected (seen in [2,1] and [1,2]). The light brown spots in such areas are indicative of some mode collapse. We also provide an example of a model failure to translate the image to the domain Y ([3,2]).





Figure 2: a) Red green and blue bands of validation images originally in the domain X (cloudy and shaded). b) These same images mapped to domain Y (clear).

Additionally, we observed that lower FD does not always correspond to high fidelity samples in this particular problem. A model that has suffered mode collapse (characterised by inpainting the clouds by light brown) may have lower FD on validation than one that produces a diverse set of samples. Inception score, which measures the entropy of predicted classes over the generated images would be a better metric for this case [19]. In future work we are aiming to extend analogous score to multi-label cases.

## **3** Discussion & Future Directions

While we provide the evidence of success that the proposed approach addresses the problem of missing values and outliers in remote sensing data, it is important to point out the lack of precision in available tools to evaluate DGMs, particularly for critical applications of land modeling or humanitarian missions, or even for more sophisticated modelling applications in Earth sciences. Furthermore, the proposed approach is rooted in the assumptions of generative modeling, yet the capacity of DGMs to learn the underlying density has been put under question recently [14, 23]. Furthermore, due to the imbalance of land-cover classes in remotely sensed images and larger number of channels per image, issues of mode-collapse are only emphasized during adversarial training and addressing these simply by bigger models and larger scale training is not practical for remote sensing or Earth science applications [4]. Resolving these issues will be fundamental for successfully applying adversarial inference to climate, remote sensing, or Earth science tasks. Furthermore, cycle consistency loss defined by the L1 norm does not work well in practice for latent variable translation models - an image mapped to the cloudy domain should be considered cycle-consistent regardless of the spatial position of the cloud, given one or more were generated, as highlighted by the setting of the problem in Section 2.3.

In this work we have demonstrated the potential of the proposed approach for producing cleaner data for remote sensing applications, and identified some of the weaknesses and shortcomings of DGMs, which are to be addressed in future.

### References

- [1] Sen2cor | STEP. URL http://step.esa.int/main/third-party-plugins-2/sen2cor/.
- [2] S. Abdikan, F. B. Sanli, M. Ustuner, and F. Calò. Land Cover Mapping Using Sentinel-1 Data. ISPRS -International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLI-B7: 757–761, June 2016. ISSN 2194-9034.

- [3] Amjad Almahairi, Sai Rajeswar, Alessandro Sordoni, Philip Bachman, and Aaron Courville. Augmented CycleGAN: Learning Many-to-Many Mappings from Unpaired Data. arXiv:1802.10151 [cs], February 2018. arXiv: 1802.10151.
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity Natural Image Synthesis. *arXiv:1809.11096 [cs, stat]*, September 2018. arXiv: 1809.11096.
- [5] R. Delécolle, S. J. Maas, M. Guérif, and F. Baret. Remote sensing and crop production models: present trends. *ISPRS Journal of Photogrammetry and Remote Sensing*, 47(2):145–161, April 1992. ISSN 0924-2716. doi: 10.1016/0924-2716(92)90030-D.
- [6] Paul C. Doraiswamy, Sophie Moulin, Paul W. Cook, and Alan Stern. Crop Yield Assessment from Remote Sensing, June 2003. URL https://www.ingentaconnect.com/content/asprs/pers/2003/ 00000069/00000006/art00005.
- [7] Olivier Hagolle, Mireille Huc, Villa Pascual David, and Gérard Dedieu. A multi-temporal method for cloud detection, applied to FORMOSAT-2, VENUS, LANDSAT and SENTINEL-2 images. *Remote Sensing of Environment*, 114(8):1747–1755, August 2010.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, December 2015. arXiv: 1512.03385.
- [9] Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *arXiv:1709.00029 [cs]*, August 2017. arXiv: 1709.00029.
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. arXiv:1706.08500 [cs, stat], June 2017. arXiv: 1706.08500.
- [11] Armand Joulin, Laurens van der Maaten, Allan Jabri, and Nicolas Vasilache. Learning Visual Features from Large Weakly Supervised Data. *arXiv:1511.02251 [cs]*, November 2015. arXiv: 1511.02251.
- [12] Zbyněk Malenovský, Helmut Rott, Josef Cihlar, Michael E. Schaepman, Glenda García-Santos, Richard Fernandes, and Michael Berger. Sentinels for science: Potential of Sentinel-1, -2, and -3 missions for scientific observations of ocean, cryosphere, and land. *Remote Sensing of Environment*, 120:91–101, May 2012. ISSN 0034-4257.
- [13] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral Normalization for Generative Adversarial Networks. arXiv:1802.05957 [cs, stat], February 2018. arXiv: 1802.05957.
- [14] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do Deep Generative Models Know What They Don't Know? September 2018. URL https://openreview.net/ forum?id=H1xwNhCcYm.
- [15] N. A. Quarmby, M. Milnes, T. L. Hindle, and N. Silleos. The use of multi-temporal NDVI measurements from AVHRR data for crop yield estimation and prediction. *International Journal of Remote Sensing*, 14 (2):199–210, January 1993. ISSN 0143-1161.
- [16] Tim G. J. Rudner, Marc Rußwurm, Jakub Fil, Ramona Pelich, Benjamin Bischke, Veronika Kopackova, and Piotr Bilinski. Multi3net: Segmenting Flooded Buildings via Fusion of Multiresolution, Multisensor, and Multitemporal Satellite Imagery. arXiv:1812.01756 [cs], December 2018. arXiv: 1812.01756.
- [17] Marc Rußwurm and Marco Körner. Multi-Temporal Land Cover Classification with Sequential Recurrent Encoders. ISPRS International Journal of Geo-Information, 7(4):129, April 2018. URL https://www. mdpi.com/2220-9964/7/4/129.
- [18] Marc Rußwurm, Romain Tavenard, Sébastien Lefèvre, and Marco Körner. Early Classification for Agricultural Monitoring from Satellite Time Series. page 5.
- [19] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved Techniques for Training GANs. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 29, pages 2234–2242. Curran Associates, Inc., 2016.
- [20] Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. BigEarthNet: A Large-Scale Benchmark Archive For Remote Sensing Image Understanding. arXiv:1902.06148 [cs], February 2019. arXiv: 1902.06148.

- [21] T. van der Wal, B. Abma, A. Viguria, E. Prévinaire, P. J. Zarco-Tejada, P. Serruys, E. van Valkengoed, and P. van der Voet. Fieldcopter: unmanned aerial systems for crop monitoring services. *Precision agriculture* '13, pages 169–175, 2013.
- [22] Sherrie Wang, George Azzari, and David B. Lobell. Crop type mapping without field-level labels: Random forest transfer and unsupervised clustering techniques. *Remote Sensing of Environment*, 222:303–317, March 2019. ISSN 0034-4257.
- [23] Shengjia Zhao, Hongyu Ren, Arianna Yuan, Jiaming Song, Noah Goodman, and Stefano Ermon. Bias and Generalization in Deep Generative Models: An Empirical Study. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 31, pages 10792–10801. Curran Associates, Inc., 2018.
- [24] Anze Zupanc. Improving Cloud Detection with Machine Learning, May 2019. URL https://medium. com/sentinel-hub/improving-cloud-detection-with-machine-learning-c09dc5d7cf13.