

Patch-GAN Transfer Learning with Reconstructive Models for Cloud Removal

Wanli Ma[✉], Oktay Karakuş[✉], Paul L. Rosin[✉]

School of Computer Science and Informatics

Cardiff University, Cardiff, CF24 4AG, UK.

{maw13, karakuso, rosinpl}@cardiff.ac.uk

Abstract—Cloud removal plays a crucial role in enhancing remote sensing image analysis, yet accurately reconstructing cloud-obscured regions remains a significant challenge. Recent advancements in generative models have made the generation of realistic images increasingly accessible, offering new opportunities for this task. Given the conceptual alignment between image generation and cloud removal tasks, generative models present a promising approach for addressing cloud removal in remote sensing. In this work, we propose a deep transfer learning approach built on a generative adversarial network (GAN) framework to explore the potential of the novel masked autoencoder (MAE) image reconstruction model in cloud removal. Due to the complexity of remote sensing imagery, we further propose using a patch-wise discriminator to determine whether each patch of the image is real or not. The proposed reconstructive transfer learning approach demonstrates significant improvements in cloud removal performance compared to other GAN-based methods. Additionally, whilst direct comparisons with some of the state-of-the-art cloud removal techniques are limited due to unclear details regarding their train/test data splits, the proposed model achieves competitive results based on available benchmarks.

Index Terms—Cloud Removal, Transfer Learning, Masked Auto-encoder.

I. INTRODUCTION

Cloud removal is an essential topic in remote sensing and earth observation, as clouds frequently obscure the Earth’s surface in optical satellite images, hindering the effective analysis of remote sensing data. For example, images from satellite platforms such as Landsat, Sentinel-2, and MODIS are highly susceptible to cloud interference. Accurate cloud removal methods are crucial for obtaining reliable data used in environmental monitoring, land-use analysis [1, 2], agriculture [3], and other applications [4]. Therefore, it is vital to develop efficient cloud removal methods for optical remote sensing images.

With the rapid development of generative and reconstructive models, these models, which exhibit exceptional capabilities, are now capable of creating images that are virtually indistinguishable from real ones. The Masked Autoencoder (MAE) [5] is one such model, designed to reconstruct the original image from its partial observations. The obstruction caused by any obstacle (such as cloud cover in remote sensing images) restricts the observable information to only partial details. MAE effectively reconstructs the hidden information, making it particularly suitable for tasks like cloud removal, where the goal is to eliminate the clouds and generate the content of

the corresponding area. In this sense, cloud removal is akin to image reconstruction, but it presents additional challenges, as the input cloudy image may not provide sufficient information for accurate reconstruction. Considering that the MAE model has the capability to reconstruct an entire image using only partial image information as input, the motivation of this paper is to fine-tune a MAE-based pre-trained reconstructive model (originally trained on a large natural image dataset) and transfer its knowledge to the cloud removal task. This approach aims to extract information from cloudy multi-spectral remote sensing images and generates cloud-free counterparts, ultimately supporting remote sensing image analysis.

Remote sensing images are more complex than natural images due to their inherent diversity and variability. Unlike natural images, which often depict a limited number of static objects, remote sensing images typically contain multiple objects—such as buildings, vegetation, roads, and bodies of water—that vary in location and appearance. These objects are not fixed, and their spatial arrangement can change significantly between images. This variability makes it challenging to apply traditional image generation models directly. To address this complexity, we opted for a patch-based GAN structure, rather than an image-wise discriminator. In a patch-based GAN, the discriminator evaluates whether each individual patch (a small, localized portion of the image) is real or fake, rather than assessing the entire image at once. This approach allows the model to focus on local details and effectively handle the diverse structures and spatial relationships present in remote sensing imagery. By processing smaller patches independently, the model can better capture the intricate features of different objects and their varying locations, leading to improved accuracy in reconstructing missing or obscured information.

Specifically, the contributions of this paper are as follows: (1) We propose a novel deep transfer learning approach integrated with a patch-based GAN framework, which effectively handles the complexity of remote sensing images. This approach allows for localized evaluation of image patches, improving the accuracy of cloud removal by capturing the intricate features and spatial variability typical of remote sensing data. (2) We identify the inherent similarity between the masked autoencoder-based reconstruction task and cloud removal. By recognizing this connection, we successfully explore and transfer the knowledge of reconstructive models,

pre-trained on large natural image datasets, to the specific challenge of cloud removal in remote sensing imagery. This transfer not only enhances the cloud removal process but also leverages the power of pre-trained models to address the unique challenges posed by remote sensing data.

II. RELATED WORK

Clouds in multi-spectral remote sensing imagery can be classified into two categories: *thick* and *thin* clouds. When an area is covered by thick clouds, the information available for that region in optical remote sensing images becomes very limited. Widely used approaches for the thick cloud removal task are to use multi-modal or multi-temporal remote sensing data to collect additional information [1, 6–8]. On the other hand, image enhancement and filter techniques are used for thin cloud removal [9, 4]. Before applying deep learning in cloud removal, traditional image processing techniques addressed the task through pixel correction. Gafurov *et al.* [10] proposed a method involving six successive steps to estimate pixel coverage by utilizing various temporal and spatial information sources. Wang *et al.* [11] proposes using grayscale morphological operations and contrast expansion based on traditional homomorphic filtering to preserve surface information, enhance image clarity, and remove cloud effects.

With the rapid development of deep learning, the application of neural networks in cloud removal tasks has become more widespread where GANs are commonly used for this purpose in the last decade [12, 13]. GANs, in particular, comprise two components — a *generator* and a *discriminator* — that are trained concurrently through a competitive process. The generator is tasked with producing cloud-free images, while the discriminator differentiates between real data (from the original dataset) and synthetic data (created by the generator).

Masked Autoencoder (MAE) [5] is an autoencoder (encoder-decoder) technique that reconstructs the original image based on its partial observation. MAE’s encoder is a standard vision transformer (ViT) [14] but it operates solely on visible, unmasked patches. Similar to a standard ViT, the encoder embeds patches through a linear projection combined with positional embeddings and processes the resulting set using a series of Transformer blocks. During the MAE training phase, masked patches are excluded, meaning only a small subset of the full set is utilized. This approach enables the training of very large encoders while requiring only a fraction of the computational and memory resources. The MAE’s decoder is also a ViT, designed to process the full set of tokens, which includes both the encoded visible patches and the mask tokens. The decoder generates a reconstructed image, filling in the areas that were masked in the input image. Although the MAE was primarily designed for image representation, with the encoder typically used for downstream tasks, its reconstruction capability is also highly effective, particularly when combined with GAN loss.

Deep transfer learning is a machine learning technique that utilizes the knowledge learned from one task to improve the performance of another related task [15]. Transfer learning

has been utilized in the literature for tasks such as cloud detection [16] and cloud removal [17], but its application to cloud removal remains limited and is still in its early stages.

III. METHODOLOGY

The proposed deep transfer learning approach uses a GAN architecture, which comprises both a generator and a discriminator. Figure 1 illustrates the overall structure of the proposed method.

A. Generator

The generator is a cloud removal model featuring a vision transformer (ViT-large) encoder and a ViT decoder. It is initially pre-trained for an image reconstruction task using MAE on the ImageNet dataset. As the model is built on a vision transformer, the input image is randomly cropped to 224×224 during the training phase to match the required input format. The cropped image is then divided into 196 patches, each measuring 16×16 pixels, before putting it into the encoder. The decoder generates predicted cloud-free patches based on the input patches, which are then combined to form a complete image.

B. Discriminator

The Discriminator network is a fully connected neural network as shown in Figure 1. The generated cloud-free image is partitioned into 196 patches, each measuring 16×16 pixels. Each patch has RGB channels, which are flattened into a one-dimensional signal with a length of 768. The output of the discriminator is a one-dimensional signal of length 196, where each value represents whether the corresponding patch is predicted to be real or fake.

C. Loss

The generator is guided by both MSE and GAN losses. The ground truth is used to supervise the generator by minimizing the MSE loss $\mathcal{L}_{MSE}$ between the predictions p and the ground truth g .

$$\mathcal{L}_{MSE} = \frac{1}{C \times W \times H} \sum_{k=1}^{C \times W \times H} |p_k - g_k|^2, \quad (1)$$

where C , W and H refer to the channel, weight and height of features, respectively.

Another component of the losses is the GAN loss L_{GAN} , which consists of the discriminator loss L_{GAN}^D and generator loss L_{GAN}^G . Thus GAN loss is defined as

$$L_{GAN} = L_{GAN}^D + L_{GAN}^G, \quad (2)$$

where the discriminator loss is

$$L_{GAN}^D = -\frac{1}{P} \sum_{k=1}^P [\log x_k + \log (1 - \hat{x}_k)], \quad (3)$$

with P represents the number of patches in the image. x represents the Discriminator’s prediction for real input, and $\hat{x}$ represents its prediction for generated input. The generator




Fig. 1. The overall framework of the proposed deep transfer learning.

aims to deceive the discriminator, and the GAN loss for the generator is expressed as follows:

$$L_{GAN}^G = -\frac{1}{P} \sum_{k=1}^P \log(\hat{x}_k) \quad (4)$$

where P represents the number of patches in the image and $\hat{x}$ represents its prediction for generated input.

We used a layer-wise learning rate decay strategy in the fine-tuning where different layers of the neural network are assigned different learning rates during the training process. Earlier layers (closer to the input) are given smaller learning rates, while later layers (closer to the output) receive higher learning rates. This is because the weights in the earlier layers usually contain more general features (such as edges, textures, etc.) learned from the pre-training stage, which may not need to be adjusted as much. In contrast, the later layers, which tend to capture more task-specific features, may require more significant updates to adapt to the fine-tuning task.

IV. EXPERIMENTS AND RESULTS

A. Dataset

We used the RICE dataset [18] to evaluate the proposed method. It is divided into two subsets: RICE1 comprises 500 image pairs, each consisting of a cloud-covered image and its corresponding cloud-free image, both with dimensions of 512×512. A total of 400 image pairs are used for training, and 100 image pairs are reserved for testing. The dataset is collected from Google Earth. RICE2 includes 736 image sets, with each set containing three 512×512 images: a reference



Fig. 2. Dataset samples of RICE dataset. The top row displays cloudy images, while the bottom row shows cloud-free images. The first two samples on the left are from RICE1, while the two samples on the right are from RICE2.

cloud-free image, a cloud-covered image, and the corresponding cloud mask. A total of 588 image pairs are used for training, and 148 image pairs are reserved for testing. RICE1 and RICE2 are obtained from Google Earth and Landsat 8OLI/TIRS data, respectively. Following [19], we used 1/5 training data for validation during the training phase. Examples of data from the RICE dataset are displayed in Figure 2.

B. Experimental Setup

Our experiments were conducted using PyTorch with a mini-batch ADAM optimizer. The base model¹ we fine-tuned is the MAE-large model, enhanced with GAN loss. The experiments ran on an NVIDIA 3060 GPU, and the models

¹Model link: https://dl.fbaipublicfiles.com/mae/visualize/mae_visualize_vit_large_ganloss.pth

TABLE I

PERFORMANCE COMPARISON BETWEEN THE PROPOSED METHOD AND OTHER GAN-BASED TECHNIQUES ON THE RICE1 AND RICE2 DATASETS.

Model	RICE1		RICE2	
	PSNR	SSIM	PSNR	SSIM
Conditional GAN	26.547	0.903	25.386	0.811
Cycle GAN	25.880	0.893	23.910	0.793
SpA GAN	30.232	0.954	28.368	0.906
The Proposed Model	33.659	0.976	34.056	0.955

were evaluated using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) metrics.

$$\text{PSNR} = 10 \log_{10} \left(\frac{(MAX_I)^2}{MSE} \right) \quad (5)$$

where MAX_I is the maximum value that a pixel can take, and MSE is

$$MSE = \frac{1}{H \times W} \sum_{i=1}^H \sum_{j=1}^W (X(i, j) - Y(i, j))^2 \quad (6)$$

where X represents the predicted, and Y the ground truth images. SSIM is also defined as

$$\text{SSIM} = \frac{2\mu_X\mu_Y + C_1}{\mu_X^2 + \mu_Y^2 + C_1} \cdot \frac{2\sigma_X\sigma_Y + C_2}{\sigma_X^2 + \sigma_Y^2 + C_2} \cdot \frac{\sigma_{XY} + C_3}{\sigma_X\sigma_Y + C_3} \quad (7)$$

where, C_1, C_2, C_3 are constants used to avoid system errors. μ and σ represent the mean and variance of two image windows, respectively, while the overall score corresponds to the mean.

C. Experimental Results

We compared the performance of the proposed method on the RICE1 dataset with other GAN-based cloud removal approaches, as shown in Table I. The results of the reference methods are collected from [19]. The proposed method demonstrates a clear improvement over other GAN-based approaches in terms of both PSNR and SSIM. The improvement in SSIM (0.976) highlights its ability to maintain structural integrity, while the high PSNR (33.659) reflects enhanced image quality. These results suggest that the proposed method is more effective at addressing the complexities of cloud removal in remote sensing images.

Apart from RICE1, we also evaluate the proposed method on the RICE2 dataset, and results are also shown in Table I. Results of reference method are collected from [19]. The proposed achieves the best results in both PSNR and SSIM, showcasing its superiority in generating cloud-free images with excellent quality and strong structural preservation.

Furthermore, we evaluated our proposed method against state-of-the-art cloud removal approaches using SSIM as the performance metric for two benchmark datasets, as presented in Table II. However, it is important to note that the authors [20] of the compared method – *CMNet* – and remaining methods did not disclose critical details about their dataset configurations, such as the specific splits used for training, testing, and validation. This lack of transparency makes a

TABLE II

SSIM OF STATE-OF-THE-ART APPROACHES.

Method	DCP	YUV-GAN	MS-GAN	RCAN	CVAE	CMNet	Ours
RICE1	0.769	0.888	0.932	0.935	0.886	0.970	0.976
RICE2	0.645	0.819	0.881	0.840	0.919	0.926	0.955




Fig. 3. Visual results on the RICE1 (top) and RICE2 (bottom) datasets. From left to right are the cloud-covered image, the generated cloud-free image, and the ground truth cloud-free image.

direct, fair comparison challenging. Despite these limitations, the SSIM values indicate that our proposed method achieves competitive performance, demonstrating its effectiveness in addressing the cloud removal task even under these less-than-ideal comparative conditions.

Some visual examples are displayed in Figure 3. The generated cloud-free images efficiently remove clouds and closely resemble the ground truth cloud-free images in both the RICE1 and RICE2 datasets.

V. CONCLUSION

In this study, we introduced a novel deep transfer learning approach that adapts a model pre-trained on a large natural image dataset for image reconstruction to the task of cloud removal in remote sensing. The proposed method leverages a GAN framework with a patch-wise discriminator, which simplifies the discriminator's task by focusing on localized regions of the image. This design enhances the discriminator's accuracy and, consequently, the overall performance of the model.

Our approach significantly outperforms existing GAN-based methods for cloud removal and demonstrates competitive results compared to state-of-the-art techniques, as evaluated using SSIM. This highlights the effectiveness of integrating deep transfer learning and patch-based evaluation for addressing the complexities of remote sensing imagery.

In future work, we aim to investigate the potential of larger vision and language models for remote sensing tasks. Given their remarkable performance across various domains, these models hold promise for advancing cloud removal and other remote sensing applications.

REFERENCES

- [1] H. Liu, B. Huang, and J. Cai, "Thick cloud removal under land cover changes using multisource satellite imagery and a spatiotemporal attention network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–18, 2023.
- [2] H. Shen, J. Wu, Q. Cheng, M. Aihemaiti, C. Zhang, and Z. Li, "A spatiotemporal fusion based cloud removal method for remote sensing images with land cover changes," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 3, pp. 862–874, 2019.
- [3] S. B. N. Sarè, R. Gaetano, R. Interdonato, Y. Hountondji, D. Ienco, and C. Dantas, "Joint cloud removal and classification of Sentinel-2 image time series for agricultural land cover mapping in northern Benin," in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 4824–4827.
- [4] H. Shen, H. Li, Y. Qian, L. Zhang, and Q. Yuan, "An effective thin cloud removal procedure for visible remote sensing images," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 96, pp. 224–235, 2014.
- [5] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 16 000–16 009.
- [6] A. Meraner, P. Ebel, X. X. Zhu, and M. Schmitt, "Cloud removal in Sentinel-2 imagery using a deep residual neural network and SAR-optical data fusion," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 166, pp. 333–346, 2020.
- [7] Y. Chen, W. He, N. Yokoya, and T.-Z. Huang, "Blind cloud and cloud shadow removal of multitemporal images based on total variation regularized low-rank sparsity decomposition," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 157, pp. 93–107, 2019.
- [8] Y. Chen, Q. Weng, L. Tang, X. Zhang, M. Bilal, and Q. Li, "Thick clouds removing from multitemporal Landsat images using spatiotemporal neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–14, 2020.
- [9] J. Liu, X. Wang, M. Chen, S. Liu, X. Zhou, Z. Shao, and P. Liu, "Thin cloud removal from single satellite images," *Optics Express*, vol. 22, no. 1, pp. 618–632, 2014.
- [10] A. Gafurov and A. Bárdossy, "Cloud removal methodology from modis snow cover product," *Hydrology and Earth System Sciences*, vol. 13, no. 7, pp. 1361–1373, 2009.
- [11] X. Wang, M. Li, and H. Tang, "A modified homomorphism filtering algorithm for cloud removal," in *2010 International Conference on Computational Intelligence and Software Engineering*, 2010, pp. 1–4.
- [12] P. Singh and N. Komodakis, "Cloud-GAN: Cloud removal for Sentinel-2 imagery using a cyclic consistent generative adversarial networks," in *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2018, pp. 1772–1775.
- [13] J. Li, Z. Wu, Z. Hu, J. Zhang, M. Li, L. Mo, and M. Molinier, "Thin cloud removal in optical remote sensing images based on generative adversarial networks and physical model of cloud distortion," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 166, pp. 373–389, 2020.
- [14] A. Dosovitskiy, "An image is worth 16×16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [15] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4-7, 2018, Proceedings, Part III 27*. Springer, 2018, pp. 270–279.
- [16] S. Pang, L. Sun, Y. Tian, Y. Ma, and J. Wei, "Convolutional neural network-driven improvements in global cloud detection for Landsat 8 and transfer learning on Sentinel-2 imagery," *Remote Sensing*, vol. 15, no. 6, p. 1706, 2023.
- [17] S. Ahn, S. Kim, J. Do, J. Park, and J. Kang, "Cloud removal on satellite image using transfer learning based generative adversarial network," in *2020 International Conference on Information and Communication Technology Convergence (ICTC)*. IEEE, 2020, pp. 203–205.
- [18] D. Lin, G. Xu, X. Wang, Y. Wang, X. Sun, and K. Fu, "A remote sensing image dataset for cloud removal," *arXiv preprint arXiv:1901.00600*, 2019.
- [19] H. Pan, "Cloud removal for remote sensing imagery via spatial attention generative adversarial network," *arXiv preprint arXiv:2009.13015*, 2020.
- [20] J. Liu, B. Pan, and Z. Shi, "Cascaded memory network for optical remote sensing imagery cloud removal," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–11, 2024.