Patch-Level Unsupervised Planetary Change Detection

Sudipan Saha, *Member, IEEE*, and Xiao Xiang Zhu, *Fellow, IEEE*

Abstract—Change detection (CD) is critical for analyzing data collected by planetary exploration missions, e.g., for identification of new impact craters. However, CD is still a relatively new topic in the context of planetary exploration. Sheer variation of planetary data makes CD much more challenging than in the case of Earth observation (EO). Unlike CD for EO, patch-level decision is preferred in planetary exploration as it is difficult to obtain perfect pixelwise alignment/coregistration between the bi-temporal planetary images. Lack of labeled bi-temporal data impedes supervised CD. To overcome these challenges, we propose an unsupervised CD method that exploits a pretrained feature extractor to obtain bi-temporal deep features that are further processed using global max-pooling to obtain patch-level feature description. Bi-temporal patch-level features are further analyzed based on difference to determine whether a patch is changed. Additionally, a self-supervised method is proposed to estimate the decision boundary between the changed and unchanged patches. Experimental results on three planetary CD datasets from two different planetary bodies (Mars and Moon) demonstrate that the proposed method often outperforms supervised planetary CD methods. Code is available at https://gitlab.lrz.de/ai4eo/cd/-/tree/main/planetaryCDUnsup.

Index Terms—Change detection (CD), planetary exploration, pooling, transfer learning, unsupervised learning.

I. INTRODUCTION

INTEREST in planetary exploration missions has increased significantly in the last decade [1] as such missions enrich our knowledge about the solar system [2]. A crucial role in most such missions is played by the scientific imaging instruments that are used for various purposes, including planetary surface characterization and spectral mapping for mineralogy.

Large volume of images are currently being captured by ongoing planetary imaging investigations, e.g., high-resolution imaging science experiment (HiRISE) and ConTeXtCamera (CTX) on the MarsReconnaissance Orbiter [3], [4].

Change detection (CD) is one of the most studied topics in Earth observation (EO). CD plays crucial role in several EO tasks, e.g., disaster management [5], urban monitoring [6], and military applications. Despite its established significance in EO, CD has not been explored much in context of planetary exploration. However, just like EO, CD may play a significant role in planetary explorations. As detailed by Kerner *et al.* [3], one such application of CD in planetary exploration is to monitor the changes induced by meteorite impact. Such impacts strongly alter the landscape of the planets. Another such application is monitoring of recurring slope lineae (RSL) that appear/disappear on surface of Mars on timescales close to a year.

Kerner *et al.* [3] studied several supervised methods in context of planetary CD. However, when talking about CD for EO, unsupervised methods are more popular than the supervised ones [6]. This is because collection of labeled multitemporal training data is difficult in context of CD. Moreover, even if training data are collected for one particular application or geography, supervised methods do not generalize well for other applications or geography. While variation of geography limits applicability of supervised CD methods for EO (i.e., one planet), it certainly limits their applicability when considering many planets and their hundreds of moons. This is confirmed by the work of Kerner *et al.* [3], where accuracy significantly drops when a model trained on HiRISE RSL dataset is applied on CTX meteorite impacts dataset. This shows the necessity of moving beyond supervised methods for planetary CD. Another difference between planetary CD and CD for EO is that the latter assumes near-perfect pixelwise alignment between bi-temporal input while such alignment is difficult to obtain for the former. Due to this, previous work on planetary CD [3] focuses on patch-level CD instead of pixelwise CD.

Deep transfer learning methods that exploit a pretrained model for bi-temporal feature extraction and comparison have shown excellent performance in different CD applications [5], [6]. Inspired by this we propose a deep transfer learning based CD method that ingests bi-temporal patches and processes it through a set of convolution layers (pretrained network). The bi-temporal feature maps are processed using a global max-pooling to summarize the content of both patches and account for possible presence of misalignment. Finally, the difference of features after global max-pooling are taken,

thresholded using a decision boundary obtained with a self-supervised method, to determine whether the considered patch is changed/unchanged.

The key contributions of this letter are as follows.

1) In context of planetary CD, this letter proposes an unsupervised deep transfer learning based method that can determine whether a pair of bi-temporal patches are changed, even if they are not perfectly coregistered.
2) This letter further proposes a method using pseudo unchanged pairs to determine threshold for distinguishing changed and unchanged patches.
3) The letter validates the proposed method on three diverse planetary CD datasets, two of which show misalignment between pretemporal and post-temporal images.

We organize the rest of the letter as follows. Related works are discussed in Section II. Section III outlines the proposed method. Datasets and experimental results are detailed in Section IV. Finally, we conclude the letter in Section V.

II. RELATED WORK

Considering relevance to our work, in this Section, we briefly discuss unsupervised CD and planetary CD.

A. Unsupervised CD in EO

Prior to the emergence of deep learning, most unsupervised CD methods for EO used the concept of pixelwise image differencing, i.e., change vector analysis (CVA) [7]. Many variants of CVA, e.g., parcel CVA (PCVA) [8] and robust CVA (RCVA) [9], incorporated the notion of spatial context in CVA. Deep learning-based unsupervised CD methods are generally based on transfer learning [6]. Deep CVA (DCVA) [6] is one such framework that incorporates CVA with pretrained deep network-based feature extraction. DCVA has shown excellent performance in many tasks, e.g., building CD [5] and agricultural monitoring [10]. Another class of unsupervised CD methods preclassifies some samples with high confidence as changed/unchanged using some traditional approach and further uses those confident samples for training a CD model [11]. Such methods have limited applicability in planetary CD as error introduced by sensor or unseen geographic characteristics may significantly impact the choice of confident samples.

B. CD in Planetary Exploration

Kerner *et al.* [3] presented a detailed study of the supervised methods for planetary CD. To the best of our knowledge, this is the only work on planetary CD. They specifically promoted the use of convolutional autoencoder along with different supervised classifiers for planetary CD. For the scenario where training data is not available, they reused the supervised CD model trained on another dataset. However, their results clearly show that such straightforward reuse of supervised model does not perform satisfactorily when the source and target domains are significantly different, e.g., when source domain is HiRISE RSL dataset and the target domain is CTX meteorite impacts dataset.
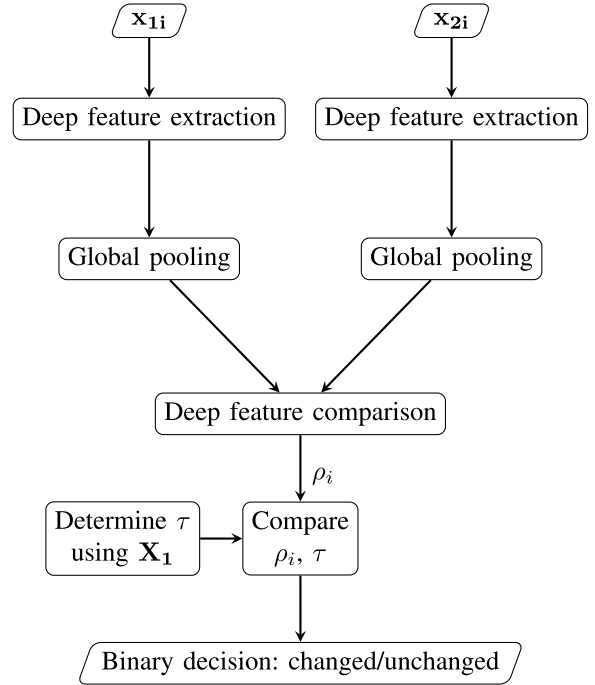

Fig. 1. Proposed patch-level planetary unsupervised CD technique.

III. PROPOSED METHOD

Let us consider two sets of I unlabeled patches $\mathbf{X_1} = \{\mathbf{x_{1i}}, \forall i = 1, \ldots, I\}$ and $\mathbf{X_2} = \{\mathbf{x_{2i}}, \forall i = 1, \ldots, I\}$, captured over same planetary surface at time t_1 and t_2. Spatial dimension of the patches is $R \times C$. We assume that $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ represent same area/object, however, may not be properly aligned, which is in stark contrast to near-perfect alignment assumption made in most CD methods [6], [9]. Instead of pixel-level CD, we are interested to assign each patch pair a label: changed or unchanged.

Both $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are separately processed through a multilayered pretrained convolutional neural network (CNN) feature extractor, originally trained for some other task (see Section III-A). Pixelwise features are summarized using global max-pooling to obtain $\mathbf{f_{1i}}$ and $\mathbf{f_{2i}}$, corresponding to $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$, respectively, (see Section III-B). They are compared to obtain a deep feature difference ρ_i, larger value of which indicates the patch-pair $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ is more likely to be changed. We can obtain a binary changed/unchanged decision by comparing ρ_i to a threshold τ (see Section III-C). We determine the optimum value of τ using a self-supervised mechanism that does not require any label information and uses only the pretemporal sets of patches $\mathbf{X_1}$ (see Section III-D). The proposed patch-level CD method is shown in Fig. 1.

While the proposed method is intended for patch-level CD, pixelwise CD map can also be obtained, which is briefly discussed in Section III-E.

A. Deep Features Extraction

Similar to the DCVA framework [6], bi-temporal patches $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are separately processed using a pretrained CNN to obtain features corresponding to both $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$. In this

work, we use VGG-16 model [12] trained on natural image dataset, however, any other suitable pretrained model can also be used Sumbul et al. [13]. By applying this network on our target planetary images, we reuse the trained CNN model to transfer the visual descriptors learned by the CNN for its original training task to solve the planetary CD problem. As previous works have shown [6], the intermediate layers are more suitable for transfer learning on targets that are semantically different from original training data, as the lower layers of the CNN capture primitive features like edges and the very last layers capture features specific to the original training dataset. Following this, we use the sixth convolution layer of VGG-16 [12] for feature extraction. The dimension of the feature maps are $R' \times C' \times D$ where D is the number of features. R' may or may not be equal to R depending on whether there are downsampling operations in the pretrained feature extractor. Unlike [6], due to possible misalignment between $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$, straightforward pixelwise comparison of $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ may lead to errors.

Please note that the above-mentioned process of feature extraction does not require prior availability of any planetary data, labeled or unlabeled.

B. Patch Summarization

The output feature maps obtained in the previous step are sensitive to the location of the features in the input patches. While this property is useful to obtain pixelwise CD map if $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are perfectly aligned, this may cause error in our case, as $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are possibly misaligned. For patch-level CD, we need to summarize/downsample the feature maps in a way such that summarization process is robust to the shift in the position of the feature in the image.

We use global pooling operation as a global image descriptor, i.e., to summarize the patch. In contrast to popularly used average pooling, we use max pooling. Average pooling assumes that the feature descriptors in a patch are independent and identically distributed. Thus, average pooling is sensitive to more frequently occurring descriptors, a phenomenon known as visual burstiness [14] that hinders average pooling from capturing feature relevant to distinguish different samples.

The use of spatial max pooling was introduced by LeCun et al. [15] and was later extended for global max pooling [16]. Remarkably, global max pooling achieves partial invariance to small translations because the max of a patch depends only on the single largest element in the given patch. If a small translation does not bring in a new largest element at the edge of the patch or does not remove the largest element by taking it outside of the patch, then the outcome of global max pooling does not change.

Global max pooling summarizes $R' \times C' \times D$ feature maps obtained from $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ to D-dimensional $\mathbf{f_{1i}}$ and $\mathbf{f_{2i}}$, respectively.

C. Determining Whether Patch Is Changed

$\mathbf{f_{1i}}$ is subtracted from $\mathbf{f_{2i}}$ and ℓ_1 norm is applied on difference to obtain a change indicator ρ_i. Denoting the steps from deep feature extraction to obtaining change indicator as g

$$\rho_i = g(\mathbf{x_{1i}}, \mathbf{x_{2i}}). \tag{1}$$

Larger value of ρ_i indicates the patch-pair $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are more likely to be changed. We can determine whether the patch-pair $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ are changed or not by comparing ρ_i to a threshold τ.

D. Threshold Determination

We further propose an automatic self-supervised method of determining τ used in Section III-C. Taking an unlabeled patch $\mathbf{x_{1i}}$ from $\mathbf{X_1}$, we obtain a noisy version $\mathbf{x'_{1i}}$ as

$$\mathbf{x'_{1i}} = h(\mathbf{x_{1i}}) \tag{2}$$

where $h(.)$ is equivalent to applying Gaussian noise and shifting the patch by few pixels. In practice, we applied shift of up to 10 pixels. $\mathbf{x_{1i}}$ and $\mathbf{x'_{1i}}$ can be treated as pseudo unchanged pair as they represent the same scene however, with slight differences as commonly observed in multitemporal planetary images. Thus, $\mathbf{x_{1i}}$ and $\mathbf{x'_{1i}}$ are processed following the steps in Sections III-B–III-D to obtain ρ'_i:

$$\rho'_i = g\left(\mathbf{x_{1i}}, \mathbf{x'_{1i}}\right) \tag{3}$$

ρ'_i is a sample ρ value for an unchanged pair. Similarly, I such values can be generated for $i = 1, \ldots, I$, which provides us a distribution of ρ for the unchanged pairs. An upper bound for ρ for the unchanged pairs can be used as threshold τ to distinguish the unchanged patch pairs from the changed ones. τ can be obtained as

$$\tau = \max\left(\rho'_1, \rho'_2, \ldots, \rho'_I\right). \tag{4}$$

In practice, accounting for possible anomaly we exclude the top 5 percentile of ρ'_i while calculating τ.

E. Pixelwise CD

If the proposed method determines a patch-pair $\mathbf{x_{1i}}$ and $\mathbf{x_{2i}}$ to be changed, pixelwise CD map can also be obtained if the patches are reasonably aligned, by pixelwise comparison of feature maps obtained in Section III-A using the DCVA framework [6].

IV. Experimental Results

A. Datasets

Following planetary datasets are used for evaluation of the proposed method.

1) *HiRISE RSL:* RSL are dark and narrow features that are thought to be formed due to the shallow subsurface water flow. They incrementally fade and recur throughout the year. RSL dataset collected by Kerner et al. [3] focuses on Garni crater on Mars. The images in the dataset are collected using the HiRISE camera which is onboard the Mars Reconnaissance Orbiter. The images are of approximately 30 cm/pixel. The dataset is formed using red channel [3]. A total of 254 pairs of 100×100 pixels size (examples in Fig. 2) are used in the test set [3]. The dataset also has training and validation set that are not
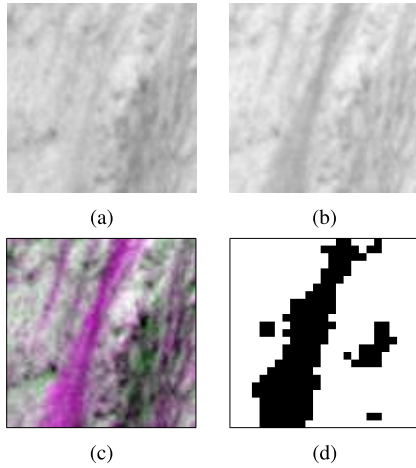
Fig. 2. Example from HiRISE RSL dataset: (a) prechange patch; (b) postchange patch; (c) false color composition created using R: prechange, G: postchange, B: prechange patch; and (d) pixelwise CD map, as detected by the proposed method.


Fig. 3. Example from CTX meteorite impact dataset: (a) prechange patch; (b) postchange patch; and (c) pixelwise CD map, as detected by the proposed method.

used by us since our proposed method is unsupervised and thus do not require training.

2) *CTX Meteorite Impacts:* Many planets including Mars are continuously impacted by meteorites. It is important to know the occurrence and location of such impact as this data helps scientists to estimate the current cratering rate in our solar system [17], [18]. The meteorite impact dataset collected by Kerner *et al.* [3] is composed of 96 images pairs (example in Fig. 3) that are collected by CTX onboard the Mars Reconnaissance Orbiter with a spatial resolution of 6 m/pixel and size of 150 × 150 pixels. Compared to HiRISE RSL, the dataset size is much smaller, 254 test cases versus only 96 test cases.

3) *Lunar Reconnaissance Orbiter Camera (LROC) Moon Dataset:* In this dataset, surface changes are the result of meteorite impacts and a spacecraft landing. The bitemporal images in this dataset are misregistered by as many as 40 pixels, making it a challenging dataset. The test set consists of five changed pairs and five unchanged pairs of 100 × 100 pixels each (example in Fig. 4).

B. Experiment Objectives

The only existing work on planetary CD [3] experimented on a set of supervised learning paradigms (combination of classifier and input representation) to find out suitable supervised learning paradigm for planetary CD. This experiment


Fig. 4. Example from LROC (Moon) dataset: (a) prechange patch; (b) postchange patch; and (c) pixelwise CD map as detected by the proposed method.

TABLE I
COMPARISON OF THE PROPOSED UNSUPERVISED METHOD WITH THE SUPERVISED METHODS IN [3] FOR HIRISE RSL

Method	Accuracy %
Proposed	94.1
Proposed (average pooling)	92.1
Proposed (k-Means clustering)	91.7
Proposed (Layer 7)	92.1
Proposed (Layer 5)	90.9
Inception-v3 absolute difference	**94.5**
Naive Bayes pixel difference	**94.5**
Inception-v3 signed difference	91.7
Inception-v3 composite grayscale	89.8
SVM absolute difference	87.8
Inception-v3 bottleneck representation	70.1
RCVA	64.6

was performed on HiRISE RSL dataset. They further experimented on CTX meteorite impacts dataset and LROC dataset to investigate if models trained on HiRISE RSL dataset can be transferred for CD in those datasets.

We compare our proposed unsupervised CD method to the following methods.

1) All supervised and transfer learning paradigms in [3].
2) Unsupervised RCVA [9]. While RCVA is proposed for pixel-based prediction, we designed a patch-based version of it by following similar strategies used for the proposed method.
3) A variant of the proposed method using global average pooling instead of max pooling.
4) A variant of the proposed method where instead of the proposed thresholding scheme, the ρ_i values corresponding to all patch pairs are clustered using k-means clustering with $k = 2$ to obtain two clusters, corresponding to the changed and unchanged patches, respectively.

C. Result Analysis

1) HiRISE RSL Dataset: Two supervised methods slightly outperform (by 0.4%) the proposed unsupervised method. However, surprisingly the proposed method outperforms as many as four supervised methods (Table I). The proposed method outperforms Inception-v3 based three different methods, Inception-v3 with signed difference, Inception v-3 with bottleneck representation, and Inception-v3 with composite grayscale. For more details of these supervised methods, refer to [3]. This result shows that proposed method's performance is almost comparable to the best performing supervised model.

TABLE II

COMPARISON OF THE PROPOSED UNSUPERVISED METHOD WITH THE
TRANSFER LEARNING CAPABILITY OF THE SUPERVISED METHODS
IN [3] FOR CTX METEORITE IMPACT DATASET

Method	Accuracy %
Proposed	**81.3**
Proposed (average pooling)	76.0
Proposed (k-Means clustering)	75.0
Inception-v3 bottleneck representation	75.0
Inception-v3 absolute difference	71.9
Inception-v3 signed difference	54.2
Inception-v3 composite grayscale	55.2
Naive Bayes pixel difference	52.1
SVM absolute difference	67.7
RCVA	59.4

The proposed method also outperforms its average pooling based version, k-means clustering based version, and unsupervised RCVA. Moreover, Table I shows result of the proposed method using fifth and seventh convolution layer of VGG-16 for feature extraction. One example from HiRISE dataset is shown in Fig. 2 that shows the proposed method is capable to indicate the location of change.

2) CTX Meteorite Impacts Dataset: Here we investigate the proposed method to the transfer learning capability of the supervised methods in [3]. As tabulated in Table II, the proposed method outperforms all supervised methods on CTX meteorite impacts dataset. Notably this dataset has some pairs with significant misalignment. Superior result of the proposed method indicates that it is able to handle misalignment by effectively summarizing patch with global max-pooling. Moreover, this proves the superiority of the proposed unsupervised method in comparison to merely transferring supervised CD model from another dataset. An example from the CTX meteorite impacts dataset is shown in Fig. 3.

3) LROC Moon Dataset: In spite of strong misalignment error in this dataset, the proposed method is able to distinguish the changed patches from the unchanged ones. The proposed method successfully labels 9 out of 10 test patches, in comparison to the method in [3] that can only correctly label 8 patches. Proposed method also outperforms its average pooling variant, RCVA (both correctly label 7 patches), and k-means clustering based variant (correctly labels 8 patches). This shows that global max-pooling successfully summarizes the content of bi-temporal feature and subsequent comparison effectively identifies the changed patches. Example of a prediction on this dataset is shown in Fig. 4 that shows the proposed method is capable to indicate the location of impact despite misalignment in dataset.

V. CONCLUSION

To the best of our knowledge, this letter introduces deep transfer learning based unsupervised CD for first time in the context of planetary exploration. Humankind has reached far beyond Earth and missions to new destinations are launched periodically. This poses us with challenge of processing multitemporal planetary data with huge variations. The proposed unsupervised method enables us to process varieties of unlabeled multitemporal planetary data without using any label. Toward this, the proposed method cleverly exploits deep transfer learning along with automatic threshold determination.

Though unsupervised, proposed method outperforms most of the existing supervised methods. This shows the proposed method as a suitable option when labeled multitemporal planetary data is not available. Considering the varieties of planetary mission and applications, it is impossible to always have abundant labeled multitemporal data. Though proposed for multispectral input, the proposed method can be easily modified for hyperspectral input by choosing a different pretrained network. Though proposed in the context of planetary CD, proposed method can be extended for any CD applications that require patch-level decision. Our future work will experiment on more planetary data. To conclude, our work is one step further in better understanding the temporal evolution of space and other planets.

REFERENCES

[1] H. R. Kerner *et al.*, "Comparison of novelty detection methods for multispectral images in rover-based planetary exploration missions," *Data Mining Knowl. Discovery*, vol. 34, no. 6, pp. 1642–1675, Nov. 2020.

[2] J. Vago *et al.*, "ESA ExoMars program: The next step in exploring Mars," *Sol. Syst. Res.*, vol. 49, no. 7, pp. 518–528, Dec. 2015.

[3] H. R. Kerner, K. L. Wagstaff, B. D. Bue, P. C. Gray, J. F. Bell, and H. B. Amor, "Toward generalized change detection on planetary surfaces with convolutional autoencoders and transfer learning," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 10, pp. 3900–3918, Oct. 2019.

[4] T. Wilhelm *et al.*, "DoMars16k: A diverse dataset for weakly supervised geomorphologic analysis on Mars," *Remote Sens.*, vol. 12, no. 23, p. 3981, Dec. 2020.

[5] S. Saha, F. Bovolo, and L. Bruzzone, "Building change detection in VHR SAR images via unsupervised deep transcoding," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 3, pp. 1917–1929, Mar. 2021.

[6] S. Saha, F. Bovolo, and L. Bruzzone, "Unsupervised deep change vector analysis for multiple-change detection in VHR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 6, pp. 3677–3693, Jun. 2019.

[7] W. A. Malila, "Change vector analysis: An approach for detecting forest changes with Landsat," in *Proc. LARS Symp.*, 1980, p. 385.

[8] F. Bovolo, "A multilevel parcel-based approach to change detection in very high resolution multitemporal images," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 1, pp. 33–37, Jan. 2009.

[9] F. Thonfeld, H. Feilhauer, M. Braun, and G. Menz, "Robust change vector analysis (RCVA) for multi-sensor very high resolution optical satellite data," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 50, pp. 131–140, Aug. 2016.

[10] S. Saha, Y. T. Solano-Correa, F. Bovolo, and L. Bruzzone, "Unsupervised deep transfer learning-based change detection for HR multispectral images," *IEEE Geosci. Remote Sens. Lett.*, vol. 18, no. 5, pp. 856–860, May 2021.

[11] H. M. Keshk and X.-C. Yin, "Change detection in SAR images based on deep learning," *Int. J. Aeronaut. Space Sci.*, vol. 21, no. 2, pp. 1–11, 2019.

[12] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*.

[13] G. Sumbul *et al.*, "BigEarthNet dataset with a new class-nomenclature for remote sensing image understanding," 2020, *arXiv:2001.06372*.

[14] V. Christlein, L. Spranger, M. Seuret, A. Nicolaou, P. Kral, and A. Maier, "Deep generalized max pooling," in *Proc. Int. Conf. Document Anal. Recognit. (ICDAR)*, Sep. 2019, pp. 1090–1096.

[15] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.

[16] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, "Learning and transferring mid-level image representations using convolutional neural networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 1717–1724.

[17] I. Daubar, A. S. McEwen, S. Byrne, M. R. Kennedy, and B. Ivanov, "The current Martian cratering rate," *Icarus*, vol. 225, no. 1, pp. 506–516, Jul. 2013.

[18] W. K. Hartmann and G. Neukum, "Cratering chronology and the evolution of Mars," in *Chronology and Evolution of Mars*. Bern, Switzerland: Int. Space Sci. Inst., 2001, pp. 165–194. [Online]. Available: https://link.springer.com/content/pdf/bfm%3A978-94-017-1035-0%2F1.pdf