Semantic Segmentation of Martian Landforms with Sparse Scribble Annotations Using a Pseudo-Labeling Strategy

Peiqi Ye[1,2], Rong Huang[1,2], Puzuo Wang[3], Yusheng Xu[*,1,2], Zhen Ye[1,2], Yongjiu Feng[1,2], Xiaohua Tong[1,2]

[1] College of Surveying and Geoinformatics, Tongji University, Shanghai, China
[2] Shanghai Key Laboratory for Planetary Mapping and Remote Sensing for Deep Space Exploration, Shanghai, China
[3] Department of Land Surveying and Geo-Informatics, The Hong Kong Polytechnic University, Hong Kong

Keywords: Mars, landforms features, semantic segmentation, weak supervision.

Abstract

Deep learning-based semantic segmentation techniques play a crucial role in the rapid mapping of Martian landforms. However, creating an annotated dataset for Martian landforms is a labor-intensive and time-consuming process. To reduce the manual effort required for labeling landform masks, this study adopts scribble annotation and develops a weakly supervised semantic segmentation framework that leverages a pseudo-labeling strategy to address the challenge of limited labeling information inherent in scribble annotations. We applied this framework to a self-constructed Martian landform semantic segmentation dataset. The experimental results demonstrate that our weakly supervised approach achieves good semantic segmentation performance in Martian environments and effectively extracts landform masks. This highlights the potential of our method to facilitate efficient and accurate mapping of Martian geomorphic features with reduced annotation effort.

1. Introduction

With geological conditions and an evolutionary history comparable to that of Earth, Mars has long been a focus of international space exploration. The surface of Mars exhibits a wide variety of geomorphologic patterns [Carr, 2007]. These landforms have formed over different periods in Mars' history due to various geological processes, including the effects of water, wind, lava, and other factors [Baker, 2001; Levy et al., 2011; Hauber et al., 2011]. With the advancement of planetary image matching and monocular image reconstruction techniques [Huang et al., 2024b; Cao et al., 2024], an increasing volume of high-quality Martian terrain data has been generated. These technological innovations enable researchers to utilize topographic data for the interpretation of Martian geomorphological events [Ye et al., 2025], which serves as a foundation for subsequent geological analyses and studies of the spatial distribution of minerals [Jiao et al., 2024]. In this context, semantic parsing based on terrain data plays a pivotal role, as it aids in accurately identifying and classifying geomorphic features, thereby enhancing the efficiency and precision of Martian geological research. Additionally, obtaining detailed geomorphological information is a critical step in selecting safe landing sites for planetary missions, significantly enhancing the probability of a successful landing [Grant et al., 2004; Arvidson et al., 2008].

Deep learning based semantic segmentation methods provide new ideas for identifying Martian landforms. In recent years, researchers have made significant progress in the semantic segmentation of Martian landforms, primarily leveraging deep learning algorithms to automatically extract geomorphic features (e.g., impact craters, dunes, and dark slope streaks) from remote sensing images, yielding promising results [Chen et al., 2023; Rubanenko et al., 2021; Wang et al., 2017]. For instance, Chen et al. [2023] introduced an advanced U-Net-based approach for detecting and segmenting impact craters at both semantic and instance levels. Using daytime infrared images from the Thermal Emission Imaging System (THEMIS), their method not only achieved accurate crater detection but also facilitated the generation of segmentation maps, contributing to geological mapping and planetary geochronology. Beyond craters, Rubanenko et al. [2021] employed a neural network to identify and delineate isolated barchan dunes from Context Camera (CTX) imagery, attaining an accuracy of 77%. These studies highlight the effectiveness of deep learning techniques in Martian landform recognition. However, two key challenges remain in the semantic segmentation of Martian landforms, indicating areas where improvements are needed. Nevertheless, the important issue is that deep learning based segmentation methods require as many mask labeling samples as possible for effective network training. During the mask labeling process, the annotation demands accurate boundaries of landforms, which is not only time-consuming, but also requires expertise of the labelers.

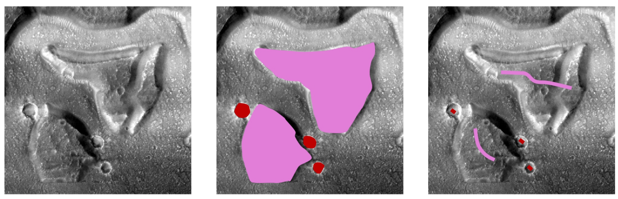

(a) (b) (c)

Figure 1. Comparison of various annotation ways. (a) Original image. (b) Mask-level and (c) Scribble-level annotation.

To this end, we use a form of scribble annotation to reduce the workload and design a weakly-supervised framework that enables the network to deal with sparsely labeled situations. Fig. 1 shows a comparison between mask annotation and scribble annotation. Unlike the mask annotation required for full supervision, scribble only requires the labelers to draw a line inside the landform to complete the annotation and does not require much expertise and knowledge. The labelers only need to draw a line where they see fit, greatly reducing the learning cost of the annotation process (up to 90%). In this case, we believe it is more effective to annotate the landforms using scribble. We proposed a new method for semantic segmentation of Martian landforms using high-resolution CTX images and CTX DEMs based on the scribble annotation.

2. Related work

Over time, an increasing variety of approaches have been applied to the semantic segmentation of Martian landforms. Recently, the rise of machine learning, particularly deep neural networks, has led to their widespread adoption as the dominant solution for segmenting Martian landform features. The following section provides a concise overview of deep learning-based segmentation methods, with a specific focus on those incorporating weakly supervised learning.

2.1 Semantic segmentation on Martian landforms

Deep learning-based methods have recently gained traction as a preferred approach for the semantic segmentation of Martian landforms. Palafox et al. [2017] pioneered the use of convolutional neural networks (CNNs) for detecting volcanic rootless cones and transverse wind ridges, demonstrating that CNNs outperform traditional support vector machine (SVM)-based classifiers by capturing a broader range of landforms with improved accuracy and recall. Expanding on this, Mulvi et al. [2022] applied CNNs to identify impact pits, transverse sand ridges, and volcanic rootless cones, confirming their effectiveness in processing images and extracting key morphological features. Advancing further, Rubanenko et al. [2021] utilized Mask-RCNN to delineate barchan dunes on Mars with high precision, while Alshehhi and Gebhardt [2022] successfully trained Mask-RCNN using Mars Dust Activity Database (MADAD) data to achieve accurate semantic segmentation of Martian dust storms. Additionally, Rubanenko et al. [2021] applied Mask-RCNN to CTX images for Martian dune segmentation, unveiling previously undetected dunes and revealing their greater prevalence in the planet's northern hemisphere.

Beyond CNN-based architectures, the transformer-based "Marsformer," introduced by Xiong et al. [2023], integrates transformer modules for Martian image analysis and demonstrates superior performance in segmenting rock masses. Furthermore, Chen et al. [2023] proposed MC-UNet, an enhanced variant of the traditional U-Net that incorporates a downsampling strategy, feature map fusion, and an attention mechanism to improve impact crater segmentation in THEMIS images, ultimately achieving higher recall and precision than conventional U-Net models.

2.2 Exploring weakly supervised learning

While deep learning-based methods have made substantial progress in the semantic segmentation of Martian landforms, the availability of large-scale, high-quality annotated training datasets remains a major challenge. Annotating Mars remote sensing images is a labor-intensive process that requires planetary geology experts to meticulously identify and classify complex geomorphic features.

To mitigate this issue, weakly supervised learning has been explored as a means to reduce the annotation burden while cite effective segmentation performance. In the field of weakly supervised semantic segmentation for Earth observation data, numerous methods have been developed and successfully applied to various Earth remote sensing tasks [Lin et al., 2022; Wang and Yao, 2022; Huang et al., 2024a]. These approaches have proven effective in reducing annotation costs while maintaining high segmentation performance. Recently, there has been a growing interest in extending weakly supervised techniques to Mars remote sensing data. Ali-Dib et al. [2020] leveraged weak supervision to successfully extract 87% of impact craters, highlighting crater depth and ellipticity from detected boundaries. Similarly, Wilhelm et al. [2020] categorized Martian landforms into 15 classes based on five major criteria: 'aeolian landforms,' 'slope-related landforms,' 'impact-related landforms,' 'topographic landforms,' and 'basic terrain types.' They also introduced the doMars16k dataset, derived from CTX images, and trained a CNN using a sliding window approach, with Markov random fields applied for post-processing to enhance segmentation results. Building upon this dataset, Zhao et al. [2024] employed super-pixel pre-generation combined with a feature extraction and fusion network to classify each super-pixel, achieving superior segmentation performance compared to existing weakly supervised models. Zhang et al. [2024] introduced the Mars sparse annotation dataset S_5Mars and utilized a weakly supervised framework to learn meaningful representations from limited labeled data, effectively segmenting landforms despite sparse annotations. Additionally, Wang et al. [2023] generated pseudo labels for unlabeled data and integrated them into a contrastive learning framework, demonstrating their method's effectiveness on the MSL dataset [Wagstaff et al., 2018].

Most of these approaches predominantly rely on optical imagery, which is well-suited for capturing landform textural features. However, the use of 3D data, such as DEMs, remains relatively underexplored in Martian landform segmentation. Given the potential benefits of incorporating 3D information, future research could focus on integrating image-based techniques with 3D semantic segmentation networks (e.g., point cloud-based models) to improve segmentation accuracy. Furthermore, adopting weakly supervised strategies for 3D datasets could help alleviate the challenges associated with manual annotation, making the segmentation of Martian landforms more efficient and scalable.

3. Method

In this work, we use scribble as a form of sparse annotation for various types of Martian landforms. We completed the annotation of 18 types of Martian landforms based on CTX DEMs and CTX images, and for comparison experiments, we also performed mask annotation for each type of landform. In this study, our input data consists of fused images and DEM data. To effectively process this fusion, we represent the data as a point cloud, which serves as the input to the neural network. In this point cloud, the first three dimensions correspond to the spatial coordinates derived from the DEM, while the fourth dimension encodes grayscale values from the image, providing essential color information. We believe that representing the data in the form of a point cloud offers a more accurate reflection of Mars' real 3D environment. This approach justifies the use of a 3D deep learning network, as it enables more effective feature extraction and spatial relationship modeling, ultimately enhancing the performance of semantic segmentation tasks. The generated Martian dataset will be transformed into point cloud format, and these point clouds will be fed into the 3D semantic segmentation network for training and testing. The landform categories include the following: "impact crater (cra)", "crater eject (eje)", "channel (cha)", "delta (del)", "mesa (mesa)", "polygons (pol)", "wrinkle ridge (wri)", "lava tube (tub)", "lava flow (flo)", "transverse sand ridges (tra)", "crescent dunes (cre)", "yardangs (yar)", "gullies (gul)", "dark slope streak (dark)", "mounds (mou)", "ridges (rid)", "cliff (cli)", and "smooth terrain (smo)". These categories represent a diverse range of Martian geomorphic features.

Towards allowing the network to fully utilize the information of unlabeled points, as shown in Fig. 2, we designed a weakly supervised framework with mean-teacher as the framework and RandlaNet [Hu et al., 2020] as the backbone network. Before feeding the points into the two models, the points will suffer perturbation, and the pseudo labels is generated by the teacher model. We carefully design three loss functions to enhance the semantic segmentation performance of this weakly supervised framework. These three loss function combinations are used for gradient back-propagation of the student model, and the teacher model is updated from an exponentially weighted average of the student model parameters.

For labeled points, the loss L_{seg} is calculated using cross-entropy loss, following typical fully supervised learning scheme.

$$L_{seg} = -\frac{1}{M}\sum_{i}^{M}\sum_{c=0}^{C-1} y_{p_{i_c}} log\tilde{p}_{p_{i_c}} \tag{1}$$

where $\tilde{p}_{p_{i_c}}$ represents c-th dimension of the softmax probability for labeled point p_{i_c} and $y_{p_{i_c}}$ denotes the ground truth.

To improve the network's robustness, data augmentation techniques were applied to the input point cloud, and consistency regularization was introduced to ensure that the student network's predictions aligned closely with those of the teacher network. This strategy helps the network generate more stable and reliable outputs. For point clouds with sparse annotations, various augmentations—such as rotation, flipping, random noise addition, and scaling—are employed to modify the data's original structure. The augmented point cloud is then passed through the student network. In contrast, the teacher network receives inputs with weaker augmentations, limited to rotation and flipping, to maintain the consistency of core geometric features.

To enforce consistency between the outputs of the student and teacher networks, a consistency loss, denoted as L_{cr}, is calculated using the mean squared error (MSE) formula:

$$L_{cr} = \frac{1}{N}\sum_{i}^{N} \|\tilde{p}_{p_i} - p_{p_i}\|^2 \tag{2}$$

where $\tilde{p}_{p_i}$ is the perturbed point p_i's softmax probability generated by the segmentation head of the student network and p_{p_i} is the softmax probability of the teacher network of point p_i that suffers another perturbation.

To mitigate the challenge posed by limited annotations, the pseudo-labeling strategy has been adopted as a simple yet effective approach in weakly supervised learning. This method involves generating pseudo labels through a network to predict labels for previously unlabeled data points.

In this study, pseudo labels are initially generated by a teacher network for unlabeled points. The network is subsequently trained using both the pseudo labels, aiming to enhance the overall quality and effectiveness of the pseudo-labeling process. A critical challenge lies in determining the uncertainty of pseudo labels. To minimize the incorporation of incorrect pseudo labels, the entropy of the probability distribution for each point is employed to ensure the selection of high-quality pseudo labels, allowing for better utilization of high-confidence predictions. Given an unlabeled point p_i that predicted by the teacher network, its entropy H_{Pi} is computed by:

$$H_{p_i} = \sum_{c=0}^{C-1} p_{ic} log p_{ic} \tag{3}$$

where p_i is the softmax probability generated by the teacher's segmentation head, p_{ic} is the value of p_i at c-th category. A larger value of H_{p_i} means that the pseudo label of the point is less reliable. Furthermore, to more effectively capture the class distribution across the Martian scene, pseudo labels are softened by assigning different weights according to the uncertainty of the class distribution. This approach helps the model better account for variations in class confidence, thereby improving segmentation performance. Following the strategy proposed in Wang and Yao [2022], the weight w_i for an unlabeled point p_i is defined as:

$$w_i = 1 - \frac{H_{p_i}}{log\, K} \tag{4}$$

The loss of reliable pseudo labels can be computed as a weighted cross-entropy loss:

$$L_{pse} = -\frac{1}{|P_r|}\sum_{p_i \in P_r} w_i \sum_{c=0}^{C-1} \hat{y}_i log p_{ic} \tag{5}$$

where P_r represents the set of reliable points in the mini-batch.

All losses proposed in previous sections participate in back-propagation simultaneously with different weights. The combined optimization problem is presented as:

$$L = L_{seg} + \lambda_1 L_{cr} + \lambda_2 L_{pse} \tag{6}$$

During training, at each epoch, the teacher network generates updated reliable pseudo labels based on the aforementioned rules. The two parameters λ_1 and λ_2 are both set to 1. These newly generated pseudo labels are then utilized to compute the pseudo-label loss, denoted as L_{pse}, which serves as a key component for back-propagation. This iterative process ensures that the model progressively refines its predictions, leveraging both labeled data and high-confidence pseudo labels to enhance overall segmentation performance.

4. Results and discussion

We present the results under a weak supervision setting. Notably, in our weakly supervised framework, the number of annotations corresponds to just 1% of the annotations used in a fully supervised setting. The baseline method in this comparison is RandlaNet, which is trained with full supervision but limited to the same 1% of sparse labels (trained by only L_{seg} without mean-teacher framework). This setup allows us to evaluate the effectiveness of our method in leveraging limited labeled data. The results demonstrate that our approach significantly outperforms the baseline, highlighting its ability to achieve superior semantic segmentation performance even with minimal supervision.

The quantitative results are presented in Tab. 1. With only 1% of labeled data, our method demonstrates significant improvements in both OA and $mIOU$ compared to the baseline. Specifically, our approach achieves an OA of 91.67% and an

https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-1641-2025 | © Author(s) 2025. CC BY 4.0 License.

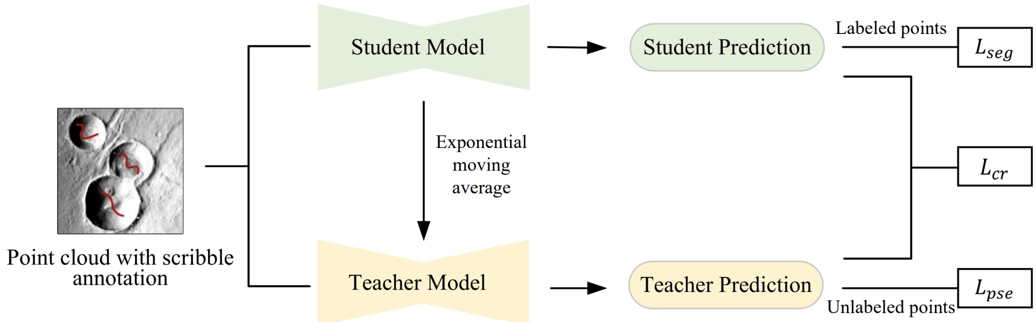

Figure 2. Processing framework of the weakly supervised semantic segmentation.

Method	OA	mIOU	IOU																	
			cra	eje	cha	del	mesa	pol	flo	tub	wri	tra	cre	yar	gul	dark	mou	rid	cli	smo
RandlaNet (100%)	93.32	80.26	76.4	47.5	90.4	69.4	83.1	87.8	75.2	88.5	86.7	88.9	91.4	80.2	91.8	82.3	81.7	68.5	63.6	91.3
RandlaNet (1%)	90.52	71.12	70.2	38.9	80.8	64.3	72.1	75.8	68.8	81.2	74.5	81.3	79.5	75.2	80.2	78.9	71.4	52.3	54.6	80.1
Ours (1%)	91.67	75.12	71.6	42.7	89.6	69.1	73.0	83.9	71.1	75.9	78.7	85.7	87.1	73.9	85.6	84.9	78.3	63.2	52.9	85.5

Table 1. Comparison of fully and weakly supervision methods.

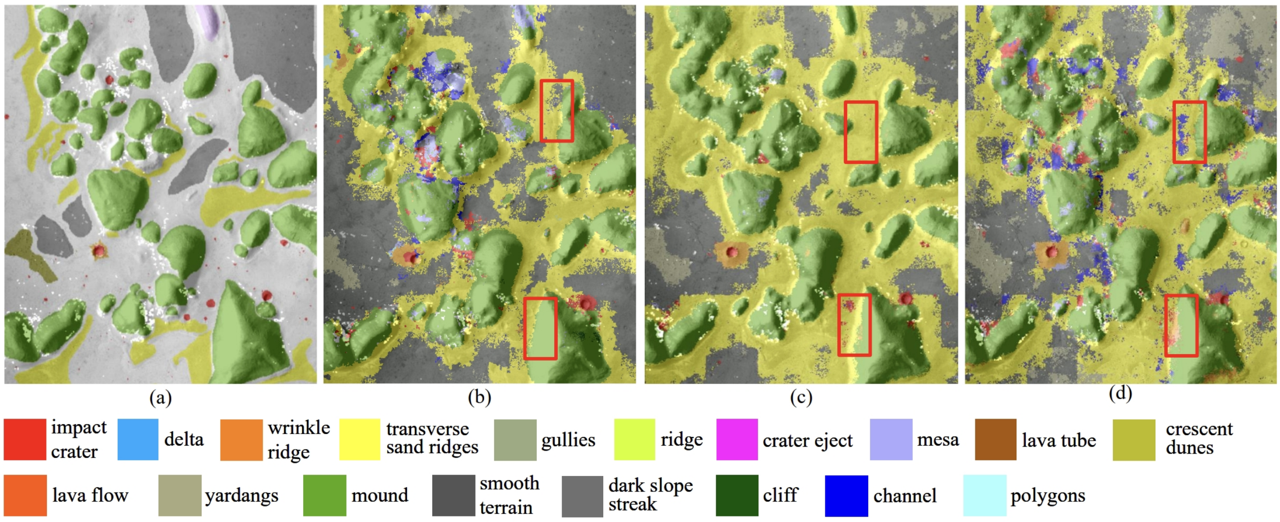

Figure 3. Semantic segmentation results by different networks in the first area. (a) is the ground truth. (b) is the result of the network that trained with full supervision. (c) is the result of the network that trained with weak supervision (1%) in our framework. (d) is the result of the network that trained only with weak supervision.

$mIOU$ of 75.12% on the test dataset. This represents an improvement of 1.15% in OA and 4.00% in $mIOU$ over the model trained solely with sparse annotations. Moreover, our results are closely aligned with those of the fully supervised RandlaNet, highlighting the effectiveness and efficiency of our method in enhancing the network's semantic segmentation performance, even with minimal labeled data. In the segmentation of deltas, mounds, and ridges categories, our method demonstrates substantial improvements compared to the baseline. These enhancements are particularly evident in the accuracy and completeness of the segmented geomorphic features. The underlying reasons for these improvements will be further elaborated in the discussion section, where we will analyze the visualization results in detail to highlight the specific advantages of our approach.

Moreover, we selected mapping results from two extensive Martian regions to assess the performance of our method. The comparison of these results clearly demonstrates that our approach closely matches the outcomes achieved through full supervision while significantly outperforming models trained solely with weak supervision. This finding highlights the robust semantic segmentation capabilities of our method, confirming its effectiveness in accurately identifying and delineating Martian landforms, even when trained with limited labeled data. Such results underscore the potential of our approach for efficient and precise Martian landform mapping in sparse annotation situation.

As illustrated by the mapping results for the two large Martian scenes, our method clearly outperforms the network trained solely with sparse annotations. In Fig. 3, our approach accurately identifies the masks for various landforms, such as

https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-1641-2025

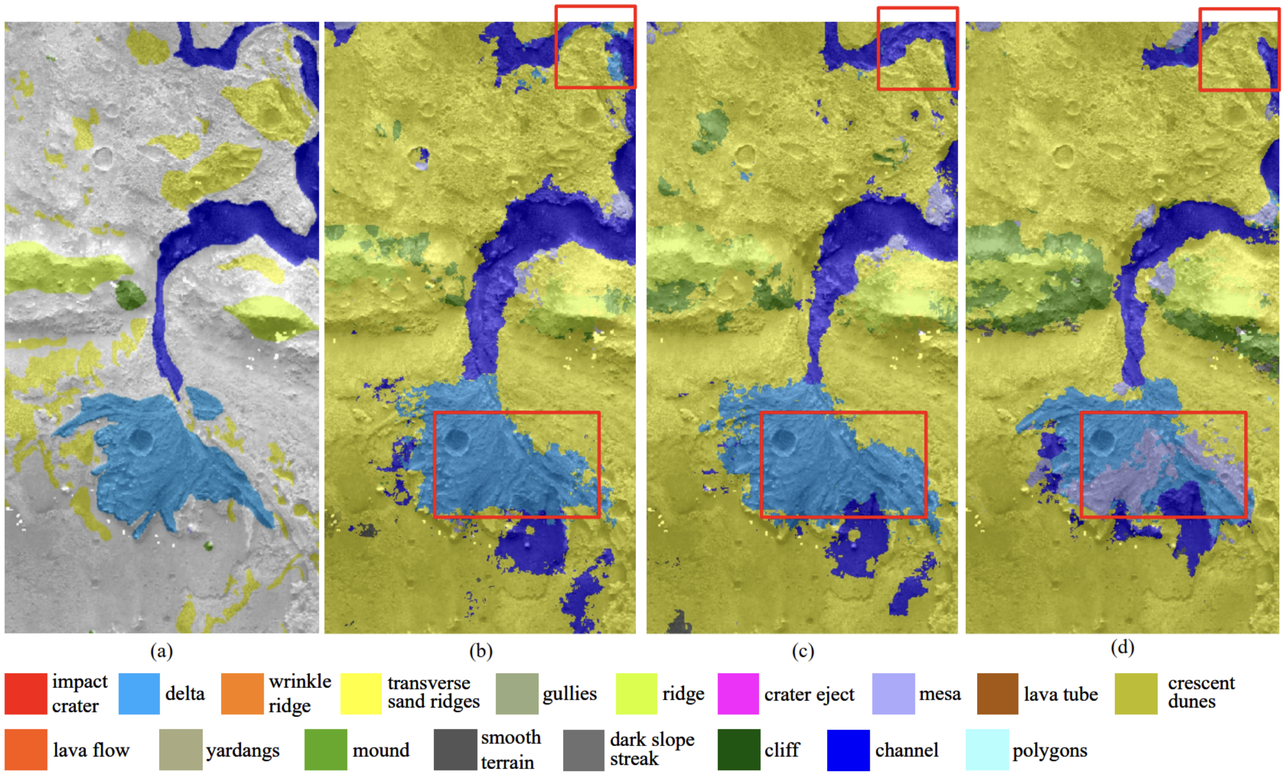

Figure 4. Semantic segmentation results by different networks in the second area. (a) is the ground truth. (b) is the result of the network that trained with full supervision. (c) is the result of the network that trained with weak supervision (1%) in our framework. (d) is the result of the network that trained only with weak supervision.

mounds and transverse sand ridges. In contrast, the results from the weakly supervised-only model exhibit more noise and less defined feature boundaries, with noticeable inconsistencies, particularly within the areas highlighted by the red boxes. This comparison underscores the strong semantic segmentation capabilities of our method. The effectiveness of our approach can be attributed to the pseudo-label learning strategy, which enhances the network's confidence when predicting unlabeled points. As a result, the segmentation outputs are smoother and less noisy, leading to more precise delineation of Martian geomorphic features. In the comparative analysis of the second region shown in Fig. 4, our method demonstrates superior semantic segmentation performance. This region contains both a delta and a channel, and while both the fully supervised network and our approach successfully segment these features with high accuracy, the model trained solely with sparse annotations exhibits significant errors. Specifically, it misclassifies a substantial portion of the delta as belonging to the mesa category. Additionally, the segmentation of the channel by the weakly-supervised-only model is incomplete, failing to capture its full extent. These results further underscore the effectiveness of our method in accurately segmenting complex geomorphic features, even with limited labeled data.

Through the analysis of both visualization and quantitative results, it is evident that the semantic segmentation outcomes of our method exhibit less noise, smoother landform boundaries, and a lower tendency for misclassification compared to the results obtained from the model trained solely with weak supervision. This improvement primarily stems from the consistency regularization applied in our approach, which enhances the model's robustness and stability during training. Additionally, the pseudo-label learning strategy effectively leverages information from unlabeled data, enriching the training process and reducing classification uncertainty. Together, these techniques contribute to the superior performance of our method in accurately segmenting Martian landforms.

5. Conclusion

In this study, we employed scribble annotations to reduce the annotation workload required for creating semantic segmentation datasets of Martian landforms. We proposed a weakly supervised semantic segmentation framework based on pseudo-label generation, and experimental results demonstrated that our approach significantly enhances the performance of the semantic segmentation network, highlighting its effectiveness. Our method not only accurately segmented geomorphic boundaries on Mars but also substantially reduced the effort needed for mask annotation during dataset generation. This contributes valuable scientific support for future Mars exploration missions by facilitating more efficient and accurate landform analysis.

Acknowledgment

The work described in this paper was supported by the National Key R&D Program of China (2022YFF0504100), Key Technology Research Project of TW-3 (TW3004), and in part by the Fundamental Research Funds for the Central Universities.

References

Ali-Dib, M., Menou, K., Jackson, A. P., Zhu, C., Hammond, N., 2020. Automated crater shape retrieval using weakly-supervised deep learning. *Icarus*, 345, 113749.

Alshehhi, R., Gebhardt, C., 2022. Detection of Martian dust storms using mask regional convolutional neural networks. *Progress in Earth and Planetary Science*, 9, 1–16.

Arvidson, R., Adams, D., Bonfiglio, G., Christensen, P., Cull, S., Golombek, M., Guinn, J., Guinness, E., Heet, T., Kirk, R. et al., 2008. Mars Exploration Program 2007 Phoenix landing site selection and characteristics. *Journal of Geophysical Research: Planets*, 113(E3).

Baker, V. R., 2001. Water and the Martian landscape. *Nature*, 412(6843), 228–236.

Cao, Z., Kang, Z., Hu, T., Yang, Z., Zhu, L., Ye, C., 2024. Cross-attention induced multilayer domain adaptation network for extraction of sub-kilometer craters from HiRIC images. *Icarus*, 407, 115776.

Carr, M. H., 2007. *The surface of Mars*. 6, Cambridge University Press.

Chen, D., Hu, F., Mathiopoulos, P. T., Zhang, Z., Peethambaran, J., 2023. MC-UNet: Martian crater segmentation at semantic and instance levels using u-net-based convolutional neural network. *Remote Sensing*, 15(1), 266.

Grant, J. A., Golombek, M. P., Parker, T. J., Crisp, J. A., Squyres, S. W., Weitz, C. M., 2004. Selecting landing sites for the 2003 Mars Exploration Rovers. *Planetary and Space Science*, 52(1-3), 11–21.

Hauber, E., Reiss, D., Ulrich, M., Preusker, F., Trauthan, F., Zanetti, M., Hiesinger, H., Jaumann, R., Johansson, L., Johnsson, A. et al., 2011. Landscape evolution in Martian mid-latitude regions: insights from analogous periglacial landforms in Svalbard. *Geological Society, London, Special Publications*, 356(1), 111–131.

Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A., 2020. Randla-net: Efficient semantic segmentation of large-scale point clouds. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11108–11117.

Huang, R., Gao, Y., Xu, Y., Hoegner, L., Tong, X., 2024a. A Simple Framework of Few-Shot Learning Using Sparse Annotations for Semantic Segmentation of 3D Point Clouds. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17, 5147–5158.

Huang, R., Wan, G., Zhou, Y., Ye, Z., Xie, H., Xu, Y., Tong, X., 2024b. Fast Double-channel Aggregated Feature Transform for Matching Planetary Remote Sensing Images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17, 9282 –9293.

Jiao, L., Xu, Y., Huang, R., Ye, Z., Liu, S., Liu, S., Tong, X., 2024. Spatial Pattern Investigation of Olivine on Mars using OMEGA and MOLA Remote Sensing Data: A Case Study at Valles Marineris and Nili Fossae. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48, 629–635.

Levy, J., Head, J., Marchant, D., 2011. Gullies, polygons and mantles in Martian permafrost environments: cold desert landforms and sedimentary processes during recent Martian geological history. 354.

Lin, Y., Vosselman, G., Yang, M. Y., 2022. Weakly supervised semantic segmentation of airborne laser scanning point clouds. *ISPRS journal of photogrammetry and remote sensing*, 187, 79–100.

Mulvi, R., Gauns, S., Gaonkar, A., Naik, S., Aswale, S., Shetgaonkar, P., 2022. Detection and segmentation of different structures from satellite images. *2022 3rd International Conference on Intelligent Engineering and Management (ICIEM)*, IEEE, 384–389.

Palafox, L. F., Hamilton, C. W., Scheidt, S. P., Alvarez, A. M., 2017. Automated detection of geological landforms on Mars using Convolutional Neural Networks. *Computers & geosciences*, 101, 48–56.

Rubanenko, L., Pérez-López, S., Schull, J., Lapôtre, M. G., 2021. Automatic detection and segmentation of barchan dunes on Mars and Earth using a convolutional neural network. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14, 9364–9371.

Wagstaff, K., Lu, Y., Stanboli, A., Grimes, K., Gowda, T., Padams, J., 2018. Deep mars: Cnn classification of mars imagery for the pds imaging atlas. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32number 1.

Wang, P., Yao, W., 2022. A new weakly supervised approach for ALS point cloud semantic segmentation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 188, 237–254.

Wang, W., Lin, L., Fan, Z., Liu, J., 2023. Semi-supervised learning for Mars imagery classification and segmentation. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(4), 1–23.

Wang, Y., Di, K., Xin, X., Wan, W., 2017. Automatic detection of Martian dark slope streaks by machine learning using HiRISE images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 129, 12–20.

Wilhelm, T., Geis, M., Püttschneider, J., Sievernich, T., Weber, T., Wohlfarth, K., Wöhler, C., 2020. Domars16k: A diverse dataset for weakly supervised geomorphologic analysis on mars. *Remote Sensing*, 12(23), 3981.

Xiong, Y., Xiao, X., Yao, M., Liu, H., Yang, H., Fu, Y., 2023. Marsformer: Martian rock semantic segmentation with transformer. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1-12.

Ye, P., Huang, R., Xu, Y., Li, W., Ye, Z., Tong, X., 2025. 3D morphometry of Martian craters from HRSC DEMs using a multi-scale semantic segmentation network and morphological analysis. *Icarus*, 426, 116358.

Zhang, J., Lin, L., Fan, Z., Wang, W., Liu, J., 2024. S^5 Mars: Semi-Supervised Learning for Mars Semantic Segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1-15.

Zhao, H., Liu, S., Tong, X., Du, Q., Bruzzone, L., Du, K., Zhang, J., Lu, X., 2024. MarsMapNet: A Novel Superpixel-Guided Multi-view Feature Fusion Network for Efficient Martian Landform Mapping. *IEEE Transactions on Geoscience and Remote Sensing*, 62, 1-16.