



Semi-supervised Learning for Mars Imagery Classification and Segmentation

WENJING WANG, LILANG LIN, ZEJIA FAN, and JIAYING LIU, Wangxuan Institute of Computer Technology, Peking University, China

With the progress of Mars exploration, numerous Mars image data are being collected and need to be analyzed. However, due to the severe train-test gap and quality distortion of Martian data, the performance of existing computer vision models is unsatisfactory. In this article, we introduce a semi-supervised framework for machine vision on Mars and try to resolve two specific tasks: classification and segmentation. Contrastive learning is a powerful representation learning technique. However, there is too much information overlap between Martian data samples, leading to a contradiction between contrastive learning and Martian data. Our key idea is to reconcile this contradiction with the help of annotations and further take advantage of unlabeled data to improve performance. For classification, we propose to ignore inner-class pairs on labeled data as well as neglect negative pairs on unlabeled data, forming supervised inter-class contrastive learning and unsupervised similarity learning. For segmentation, we extend supervised inter-class contrastive learning into an element-wise mode and use online pseudo labels for supervision on unlabeled areas. Experimental results show that our learning strategies can improve the classification and segmentation models by a large margin and outperform state-of-the-art approaches.

CCS Concepts: • **Computing methodologies** → **Semi-supervised learning settings**; *Object recognition*; *Image segmentation*;

Additional Key Words and Phrases: Mars vision tasks, image classification, image segmentation, representation learning, unsupervised learning

ACM Reference format:

Wenjing Wang, Lilang Lin, Zejia Fan, and Jiaying Liu. 2023. Semi-supervised Learning for Mars Imagery Classification and Segmentation. *ACM Trans. Multimedia Comput. Commun. Appl.* 19, 4, Article 144 (March 2023), 23 pages.

<https://doi.org/10.1145/3572916>

1 INTRODUCTION

Humanity's interest in the universe is prevalent and enduring. In recent years, machine learning has shown its great power in space exploration. For example, the first black hole image was

This work was supported by the National Key Research and Development Program of China under Grant No. 2018AAA0102702, and the National Natural Science Foundation of China under Contract No. 62172020.

Authors' address: W. Wang, L. Lin, Z. Fan, and J. Liu (Corresponding author), Wangxuan Institute of Computer Technology, Peking University, No. 128 Zhongguancun North Street, Haidian District, Beijing, China; emails: {daoshee, linlilang, zejia, liujiaying}@pku.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Association for Computing Machinery.

1551-6857/2023/03-ART144 \$15.00

<https://doi.org/10.1145/3572916>

captured by combining data from eight telescopes using a machine learning algorithm [4]. As techniques develop, machine learning will play a more and more significant role in scientific fields.

Humans have been exploring the planet Mars since the last century. Multiple rovers have been dispatched to Mars, sending an enormous amount of images to Earth. With these massive images, data-driven learning is increasingly being used in Mars research. Wagstaff et al. [48] proposed to automatically classify images from the **Mars Science Laboratory (MSL)** mission with a neural network. Targeting rover self-driving, Swan et al. [45] explored the task of Mars terrain segmentation. However, these works simply applied conventional machine learning algorithms designed for Earth object classification or street scene segmentation. They neglect the properties of Mars data and leave the features of extraterrestrial planet surfaces unexplored.

Mars data provides specific difficulties for machine learning methods. First, along with exploration progress, the rover moves to a new area and collects new data, which can cause a severe **train-test gap**. Second, the quality of Martian data suffers in many ways, such as bad weather conditions, camera equipment damage, and signal loss in Mars-to-Earth transmission, which gives rise to **limited information quality**. Detailed analysis and visualization will be given in Section 3.

Many methods have been proposed to narrow train-test gaps. Some works revise loss designs [30, 40, 50], some focus on imbalanced data distribution [5, 12, 25, 57], and some propose specific training strategies [43, 53, 54]. Although these approaches are powerful for common vision tasks, the train-test gap on Mars rover data is too challenging as we will show in Section 3, making existing methods ineffective.

Data quality improvement is a popular topic. For image quality, researchers have studied methods to deal with various kinds of distortion, including but not limited to super-resolution [16, 23], de-noising [13, 52], and illumination enhancement [24]. However, Mars rover images suffer from compound distortions, which are too complex for existing restoration methods to handle. For data diversity, image augmentation [53, 54] is not powerful enough, while adversarial-learning-based image generation [18] is not reliable. For annotation quality, researchers have explored how to train with noisy labels [32] or synthesize more labels [29], but these approaches are not effective enough on Mars rover data as we will show in Section 6.

Unlike existing methods, we solve the problem through representation learning. With a robust visual representation, the train-test gap and low data quality can be resolved simultaneously. Based on this methodology, we study two Martian vision tasks: image classification and semantic segmentation. The former is about image-level prediction, while the latter is about per-pixel prediction. Classification and segmentation are very representative tasks. Our exploration of them can also provide insights for other Martian vision tasks, such as object detection, tracking, and locating.

We adopt a widely used representation learning approach: contrastive learning [10, 21]. Contrastive learning increases the mutual information between positive pairs and decreases the similarity between negative pairs. It can improve the separability and compactness of features, providing a more suitable representation space for various downstream vision tasks. However, directly applying it to Mars rover data results in poor performance. This is because there is a severe information overlap between different Mars data samples, which negates the effect of contrastive learning.

To resolve this contradiction, we propose a semi-supervised learning strategy. On the one hand, we make use of annotations and ignore pairs within the same class, forming supervised inter-class contrastive learning. On the other hand, we train models on unlabeled images or areas to introduce more supervision. The difference between conventional contrastive learning and our learning paradigm is illustrated in Figure 1. More specifically, for Mars image classification, we abandon negative pairs and carry out unsupervised similarity learning on unlabeled images. For segmentation, we further revise contrastive learning into a pixel-wise mode with online pseudo

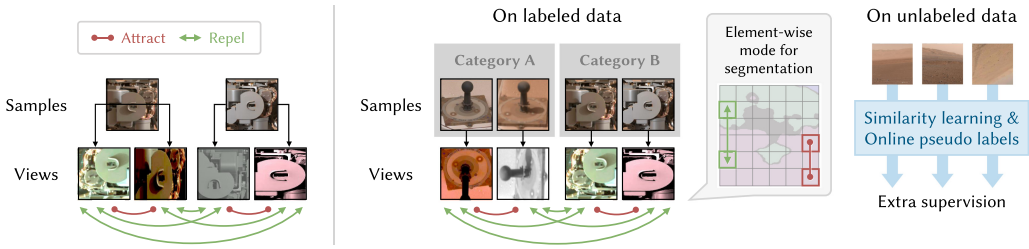


Fig. 1. The proposed semi-supervised learning paradigm for Martian machine vision tasks. Left: The framework of conventional contrastive learning. Due to the redundancy in Mars data, naive pairs can be inappropriate. Right: Our semi-supervised framework. On labeled data, we generate inner-class positive pairs and inter-class negative pairs. On unlabeled data, we use similarity learning and online pseudo labels to introduce more supervision.

labels on unlabeled areas. Experimental results demonstrate that our method achieves superior performance for Mars rover imagery classification and segmentation.

Our contributions can be summarized as follows:

- Targeting at Martian machine vision tasks, we propose a semi-supervised learning framework, which outperforms existing approaches by a large margin in terms of classification and segmentation.
- For Mars imagery classification, we propose supervised inter-class contrastive learning and unsupervised similarity learning. By abandoning inter-class pairs on labeled data as well as negative pairs on unlabeled data, we resolve the contradiction between Mars rover images and contrastive learning.
- For Mars imagery segmentation, we extend inter-class contrastive learning into an element-wise mode and introduce online pseudo labels on the unlabeled area. Our method not only suits per-pixel prediction tasks but also makes use of the unlabeled area for further supervision, improving the performance of segmentation.

The rest of the article is organized as follows. Section 2 provides a detailed review of the relevant literature. Section 3 presents an in-depth analysis of Mars rover data. Sections 4 and 5 introduce the proposed semi-supervised frameworks for classification and segmentation, respectively. Experimental results and analyses are in Sections 6 and 7. Concluding remarks are finally given in Section 8.

2 RELATED WORKS

2.1 Machine Learning in Mars Exploration

Machine learning has been utilized for a variety of planetary science tasks, such as exoplanet detection [41], comparative planetology, and exoplanet biosignatures [49]. Readers may refer to [2] for a more comprehensive review and outlook.

For Mars exploration, existing machine learning applications can be categorized into two categories: in situ (Mars edge) and ex situ (Earth edge) [33]. For in situ methods, machine learning can benefit autonomous decision-making and save bandwidth by filtering out undesired images. The **Opportunistic Rover Science (OASIS)** framework uses machine learning algorithms to identify terrain features [8, 20], dust devils, and clouds [7]. For rover navigation, Abcouwer et al. [1] presented two heuristics to rank candidate paths, where a machine learning model is applied to predict untraversable areas. For ex situ methods, machine learning can help scientists analyze data

and notify noteworthy findings. JPL scientists [17] built an impact crater classifier to analyze images captured by the Martian Reconnaissance Orbiter. Dundar et al. [19] applied machine learning algorithms to discover less common minerals and search for aqueous mineral residue. Rothrock et al. [38] designed machine learning models to identify terrain types and features in orbital and ground-based images. The analysis can alert areas of slippage for rovers and assist the determination of potential landing sites for new missions. Wagstaff et al. [48] created a dataset of the Mars surface environment and trained AlexNet [28] for content classification. Swan et al. [45] collected a terrain segmentation dataset of 35K high-resolution images and tested the performance of DeepLabv3+ [9]. To find an efficient energy distribution between different systems in the orbiter, Petkovic et al. [36] applied multi-target regression to estimate the power consumption of the thermal system. For geomorphic mapping, Wilhelm [51] built a dataset and provided an automated landform analysis strategy.

However, most of the aforementioned algorithms are non-deep, taking no advantage of powerful neural networks. Some research builds neural networks but directly applies models designed for conventional computer vision tasks [45, 48], thus having unsatisfactory performance. Moreover, advanced learning strategies such as semi-supervised and weakly supervised learning in the Martian scenario remains unexplored. The “weak supervision” in [51] refers to window sliding with Markov Random field smoothing for creating maps, which is far different from learning representation with limited data. In this article, we present an in-depth analysis of images captured by Mars rovers and introduce a more powerful semi-supervised framework, which expands the research of deep learning for Mars.

2.2 Improving Classification and Segmentation Performance

Classification and segmentation are some of the most basic and popular tasks in computer vision. There have been many techniques for improving their performance.

Several works propose loss designs to balance positive and negative samples. Triplet loss [40] minimizes the distance between positive pairs and maximizes the distance between negative pairs. Center loss [50] clusters the feature representation. Focal loss [30] aims at the imbalance between positive and negative samples. Classification problems often suffer from data imbalance across classes. To solve data imbalance, re-sampling- and re-weighting-based methods [5, 12, 25] are proposed. In work [25], researchers decouple the learning procedure into representation learning and classification, then apply class-balanced sampling for classifier retraining. CB loss [12] represents the additional benefit with a hyper-parameter related to the sample volume and re-weighting the samples based on the additional benefit. LDAM [5] aims to minimize the margin-based generalization bound along with the prior re-weighting or re-sampling strategy. CutOut [15], CutMix [53], and ClassMix [34] are powerful data augmentation mechanisms.

CutOut [15] randomly masks out regions, removing contiguous sections of images. The limitation is that CutOut only captures relationships within the samples. To introduce information across samples, CutMix [53] combines different parts of two images to generate a new image. However, the random combination may destroy the semantic structure of the original image. Thus, ClassMix [34] combines semantic classes extracted from different images to make the generated image more meaningful. ReCo [31] is a contrastive learning framework designed at a regional level to assist learning in semantic segmentation. It is computationally expensive to carry out pixel-level contrastive learning for all available pixels in high-resolution training. To reduce memory requirements, ReCo introduces an active hard sampling strategy to optimize only a few queries and keys.

However, all these methods have limited effectiveness on Mars rover data as we will demonstrate in Sections 6 and 7. In this article, we propose a more effective representation learning strategy and achieve superior performance.

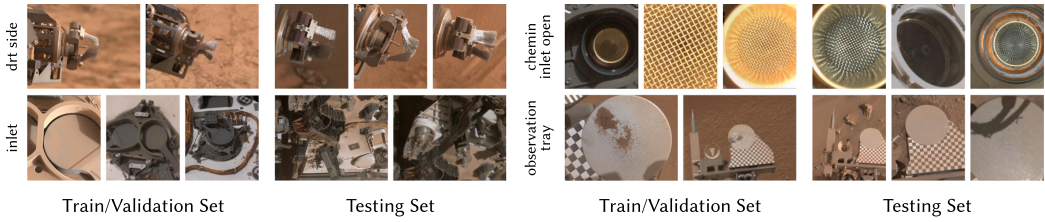


Fig. 3. The various appearances of the same objects in the MSL surface dataset [48].

Table 1. The Number of Samples for Each Category in MSL Surface Dataset [48]

	Train/Val	Test		Train/Val	Test
apxs	46	34	mahli cal target	60	57
apxs cal target	10	14	mastcam	36	32
chemcam cal target	15	21	mastcam cal target	105	48
chemin inlet open	92	84	observation tray	99	12
drill	39	20	portion box	38	48
drill holes	506	0	portion tube	128	9
drt front	6	60	portion tube opening	20	2
drt side	12	150	rem s uv sensor	32	36
ground	2,430	254	rover rear deck	57	14
horizon	299	72	scoop	190	10
inlet	261	16	turret	193	0
mahli	14	12	wheel	698	300

Table 2. The Total Area for Each Category in AI4Mars Dataset [45]

	soil	bedrock	sand	big rock
Train/Val	2.59×10^9	2.59×10^9	5.68×10^8	6.35×10^7
Test	9.57×10^8	2.10×10^9	5.89×10^8	2.28×10^7

- Limited Information Quality.** The quality of Martian data suffers in many ways. First, rover data may be affected by wrong shooting operations, camera equipment damage, and signal loss in Mars-to-Earth transmission. These errors can degrade the visual quality of images, as shown in Figure 4. Second, because of the monotonous Mars scenes and limited data sources, Mars datasets are usually less diverse than common computer vision datasets. Third, annotating Mars data requires particular expert knowledge. High labeling costs and limited budgets hamper the acquisition of high-quality annotation. Accordingly, the label quality of Mars rover data can be unsatisfactory, as shown in Figure 5.

4 SEMI-SUPERVISED MARS IMAGERY CLASSIFICATION

We first introduce our solution for the Mars imagery classification task. To make full use of annotations and unlabeled images, we design a semi-supervised contrastive learning scheme, consisting of two sub-strategies: supervised inter-class contrastive learning and unsupervised similarity learning. In the following, we first review contrastive learning, then introduce our specific method.

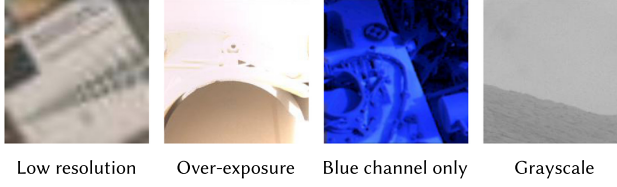


Fig. 4. Examples of low-quality images in the MSL surface dataset [48].

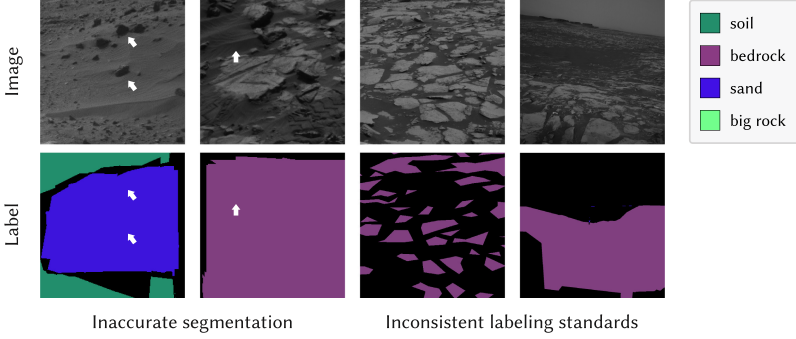


Fig. 5. Some annotations are of low quality in the AI4Mars dataset [45]. On the left, as indicated by white arrows, rocks are labeled as sand (the first column), while a patch of sand is labeled as bedrock (the second column). On the right, similar terrains are labeled inconsistently on different samples.

4.1 Review of Contrastive Learning

As stated in [10], the core of contrastive learning is to increase the mutual information between pairs of positive samples and decrease the similarity between pairs of negative samples. To generate positive and negative samples, data transformation is widely utilized. Specifically, for a sample \mathbf{x}_i from the dataset \mathbf{X} , contrastive learning performs random data transformation \mathcal{T} to obtain transformed data $\mathbf{x}_i^1 = \mathcal{T}^1(\mathbf{x}_i)$, $\mathbf{x}_i^2 = \mathcal{T}^2(\mathbf{x}_i)$, where $\mathbf{x}_i^1 \neq \mathbf{x}_i^2$ are two different views of \mathbf{x}_i , and $\mathcal{T}^1, \mathcal{T}^2$ are transformations sampled from \mathcal{T} independently. With \mathcal{T} , contrastive learning assigns positive pairs as views of the same image and negative pairs as views of different images. Then, a feature encoder f is used to extract the representation of \mathbf{x}_i^1 and \mathbf{x}_i^2 , denoted as $\mathbf{z}_i^1 = f(\mathbf{x}_i^1)$, $\mathbf{z}_i^2 = f(\mathbf{x}_i^2)$. Finally, contrastive learning normalizes the features into a spherical manifold and computes the cosine similarity between positive pairs and negative pairs. The InfoNCE loss [47] is applied for optimization:

$$\mathcal{L} = -\mathbb{E}_{\mathbf{x}_i} \left[\log \frac{\exp(K(\mathbf{z}_i^1, \mathbf{z}_i^2) / \tau)}{\sum_{\mathbf{x}_j \in \mathbf{X}} \exp(K(\mathbf{z}_i^1, \mathbf{z}_j^2) / \tau)} \right], \quad (1)$$

$$K(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \mathbf{v} / \|\mathbf{u}\| \|\mathbf{v}\|,$$

where τ is a temperature hyper-parameter, and $K(\cdot, \cdot)$ is the kernel function for computing cosine similarity.

Although contrastive learning is originally designed for unsupervised learning, it can also assist supervised learning. Compared with a solo classification loss, training with an extra contrastive loss can enrich the visual representation, improving the robustness of the neural network.

4.2 Supervised Inter-class Contrastive Learning

Current research mostly applies contrastive learning on large-scale Earth image datasets such as ImageNet [14] and JFT [44], where the difference between images is large enough. However, compared with diverse Earth scenes, Mars scenes are rather homogeneous. Moreover, Mars rovers may take photos at the same scene multiple times, e.g., when investigating the surrounding terrain or monitoring equipment degradation. These factors result in a severe information overlap between different Mars image samples. Since contrastive learning relies on the mutual information between different samples, cross-image information overlap may lead to the futility of contrastive learning on Mars data.

To address this problem, we make use of classification annotations to select more appropriate contrastive pairs. Specifically, we delete negative samples belonging to the same class and add positive pairs of different samples belonging to the same class; i.e., we increase the mutual information of all samples in the same category and decrease the similarity only for pairs of different categories. In other words, we turn the original unsupervised sample-wise contrastive learning strategy into a supervised inter-class version.

Denoting the number of categories as C , our supervised inter-class contrastive loss is

$$\mathcal{L}_{cls}^s = -\mathbb{E}_{\mathbf{x}_{c_i}} \left[\log \frac{\exp \left(K \left(\mathbf{z}_{c_i}^1, \mathbf{z}_{c_i}^2 \right) / \tau \right)}{\sum_{c_j=1}^C \exp \left(K \left(\mathbf{z}_{c_i}^1, \mathbf{z}_{c_j}^2 \right) / \tau \right)} \right], \quad (2)$$

where $\mathbf{z}_{c_i}^1 = f(\mathcal{T}^1(\mathbf{x}_{c_i}))$ and $\mathbf{z}_{c_i}^2 = f(\mathcal{T}^2(\mathbf{x}'_{c_i}))$. \mathbf{x}_{c_i} and \mathbf{x}'_{c_i} are samples belonging to category c_i .

The selection of data augmentation \mathcal{T} is one of the keys in contrastive learning. The augmentation we use can be categorized into two types: shape and pixel. Shape augmentation teaches the model to perceive objects under different camera angles and magnifications. It contains random flipping, cropping, resizing, and rotation. Pixel augmentation aims to improve the model's robustness to image quality degradations. It contains Gaussian blur, color jittering, and desaturation. With revised contrastive learning and targeted augmentation, the problems of data imbalance and low image quality can be greatly alleviated.

4.3 Unsupervised Similarity Learning

Our supervised inter-class contrastive loss relies on sufficient annotations. However, labeling data requires a lot of manpower and financial resources. For Mars data, the cost is particularly high. Identifying Martian landscapes and rover components requires expert knowledge [48]. Although annotations are expensive, pure images are relatively easy to obtain. Through past and current missions, scientists have acquired millions of images from Mars. To reduce the reliance on annotations, we explore how to use unlabeled data to further improve classification performance.

Unlabeled Mars data also has information overlap between different samples. When applying contrastive learning to these data, negative pairs may come from the same scene and should not be forced apart in the feature space. The only thing guaranteed is that different views of the same image should have similar representations. Therefore, we adopt similarity learning, where we abandon negative pairs and only consider positive ones. Denoting \mathbf{x}_u as a sample from unlabeled data, the proposed unsupervised similarity loss is

$$\mathcal{L}_{cls}^u = -\mathbb{E}_{\mathbf{x}_u} \left[K \left(\mathbf{z}_u^1, \mathbf{z}_u^2 \right) \right]. \quad (3)$$

Why similarity learning does not lead to model collapse is an interesting question. In experiments, we find that training classification with \mathcal{L}_{cls}^u does cause collapse. However, when we add \mathcal{L}_{cls}^s , model collapse is prevented. This may be because supervised inter-class contrastive learn-

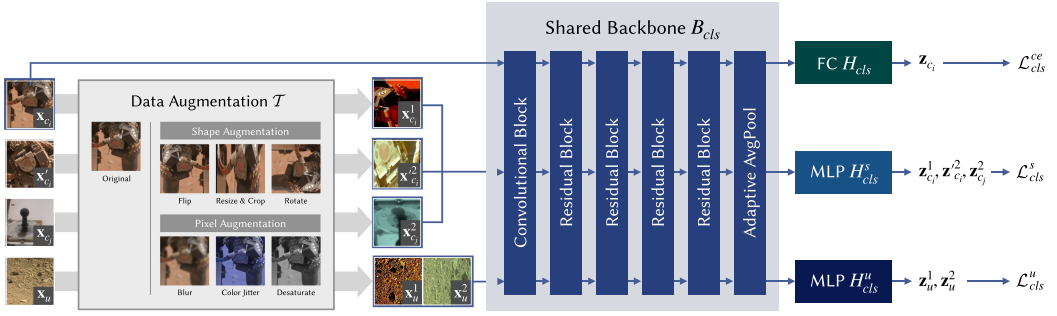


Fig. 6. The framework of our semi-supervised Mars image classification. There are three streams in the framework: (top to bottom) classification, supervised inter-class contrastive learning, and unsupervised similarity learning. Each stream shares the backbone and has independent heads.

ing provides a good restriction to the feature representation, counteracting the bad impacts of similarity learning.

4.4 Full Model

As shown in Figure 6, our framework consists of three streams: classification, supervised inter-class contrastive learning, and unsupervised similarity learning. Since we combine supervised and unsupervised learning, our approach is “semi-supervised.”

The classification objective \mathcal{L}_{cls}^{ce} is cross entropy. Given a labeled sample $\mathbf{x}_{c_i} \in \mathbf{X}$ of category c_i , \mathcal{L}_{cls}^{ce} is

$$\mathcal{L}_{cls}^{ce} = -\mathbb{E}_{\mathbf{x}_{c_i}} \left[\log \frac{\exp(\mathbf{z}_{c_i, c_i})}{\sum_{c_j=1}^C \exp(\mathbf{z}_{c_i, c_j})} \right], \quad (4)$$

where \mathbf{z}_{c_i, c_j} is the c_j th element of \mathbf{z}_{c_i} , representing the prediction of the sample belonging to the c_j category.

Our full loss function is

$$\mathcal{L}_{cls} = \mathcal{L}_{cls}^{ce} + \lambda_{cls}^s \mathcal{L}_{cls}^s + \lambda_{cls}^u \mathcal{L}_{cls}^u, \quad (5)$$

where $\lambda_{cls}^s = 1$ and $\lambda_{cls}^u = 0.2$ are hyper-parameters to balance different training objectives. The temperature hyper-parameter τ in \mathcal{L}_{cls}^s is set to 0.2.

We use ResNet-50 [22] for classification. In \mathcal{L}_{cls}^s and \mathcal{L}_{cls}^u , the feature encoder f consists of a shared ResNet-50 backbone B and a two-layer **Multi-Layer Perceptron (MLP)** head. We denote the MLP in \mathcal{L}_{cls}^s and \mathcal{L}_{cls}^u as H_{cls}^s and H_{cls}^u , respectively. The output dimension of H_{cls}^s and H_{cls}^u is 128.

5 SEMI-SUPERVISED MARS IMAGERY SEGMENTATION

In this section, we extend our semi-supervised contrastive learning scheme from Mars imagery classification to semantic segmentation. The challenge is that classification is an image-wise prediction task, while segmentation is pixel-wise. In segmentation, each pixel has its own category, which requires the model to perceive objects at a finer scale. However, conventional contrastive learning methods treat the input image as a whole. To address this problem, we propose element-wise contrastive learning.

5.1 Element-wise Inter-class Contrastive Learning

Given an input image \mathbf{x}_{seg} , we first use a feature encoder g to extract the representation $\mathbf{f} = g(\mathbf{x}_{seg})$. In classification, the extracted representation is a vector, while in segmentation, what we obtain

is a 2D feature map. The resolution of \mathbf{f} is usually lower than that of the input image \mathbf{x}_{seg} . Then, we down-sample the segmentation annotation of \mathbf{x}_{seg} to match the resolution of \mathbf{f} . In this way, each spatial element in \mathbf{f} can have its own category, which supports us to conduct element-wise inter-class contrastive learning.

Similar to our approach for classification, we want to make features of the same category closer and features of different categories more separable. Denoting \mathbf{f}_i as an element in the 2D feature map \mathbf{f} and its category as c_i , we can have a naïve form of element-wise inter-class contrastive learning:

$$\tilde{\mathcal{L}}_{seg}^s = -\mathbb{E}_{\mathbf{x}_{seg}} \mathbb{E}_{\mathbf{f}_i} \left[\log \frac{\sum_{c_j=c_i} \exp(K(\mathbf{f}_i, \mathbf{f}_j)/\tau)}{\sum_j \exp(K(\mathbf{f}_i, \mathbf{f}_j)/\tau)} \right]. \quad (6)$$

For pixels without labels, we simply ignore them.

The problem is that $\tilde{\mathcal{L}}_{seg}^s$ is impossible to calculate. If the resolution of \mathbf{f} is 128×128 , there will be 16,384 elements in \mathbf{f} . Computing the similarity among 16,384 vectors, i.e., 268 million vector multiplications, exceeds the capacity of current computation devices. Moreover, $\tilde{\mathcal{L}}_{seg}^s$ only considers elements in a single image. To make training feasible and introduce cross-image contrastive pairs, we use a memory bank to store the history average representation for each category.

The proposed memory bank $\mathbf{M} = \{\mathbf{m}_{c_i}\}_{c_i=1}^C$ consists of C queues, where C is the number of categories. Each queue \mathbf{m}_{c_i} constantly removes the oldest element and stores the average of all features labeled c_i in the current \mathbf{f} . Denoting $\bar{\mathbf{m}}_{c_i}$ as the average of \mathbf{m}_{c_i} , our element-wise inter-class contrastive learning loss is

$$\mathcal{L}_{seg}^s = -\mathbb{E}_{\mathbf{x}_{seg}} \mathbb{E}_{\mathbf{f}_i} \left[\log \frac{\exp(K(\mathbf{f}_i, \bar{\mathbf{m}}_{c_i})/\tau)}{\sum_{c_j=1}^C \exp(K(\mathbf{f}_i, \bar{\mathbf{m}}_{c_j})/\tau)} \right]. \quad (7)$$

With an assistant memory bank, \mathcal{L}_{seg}^s reduces the computation complexity and introduces cross-image information by storing previous features.

5.2 Online Pseudo Labeling for Semi-supervised Learning

Annotating segmentation is laborious and time-consuming. Accordingly, 44.83% of the area does not have labels in the AI4Mars dataset, which limits the effect of our supervised inter-class contrastive learning. Moreover, as we stated in Section 3, the quality of annotation in AI4Mars is not satisfactory. To exploit numerous unlabeled pixels and refine annotation, we expand supervision from labeled areas to more areas by semi-supervised learning.

To construct positive and negative pairs on unlabeled data, previous methods apply clustering [6] or data transformation [21]. However, these strategies rely on sufficient and diverse data, which are not suitable for Mars scenarios. We instead use a fairly good segmentation model and predict pseudo labels on unlabeled data. With the help of segmentation prediction, we can construct more accurate contrastive pairs.

We first train a segmentation network on fully supervised data. Then, we use it to predict a category for each unlabeled pixel. To ensure the quality of pseudo labels, we remove predictions with low confidence. The confidence threshold is set to 0.9. Then, pseudo labels are added to join our *Element-wise Inter-class Contrastive Learning*. Formally, for an input image \mathbf{x}_{seg} with ground-truth annotation \mathbf{y} and segmentation prediction $\hat{\mathbf{y}}$, we supplement the unlabeled part in \mathbf{y} with the high-confidence part in $\hat{\mathbf{y}}$. The merged result is denoted as $\bar{\mathbf{y}}$. For fully supervised learning, \mathcal{L}_{seg}^s refers \mathbf{y} for the category of each pixel, while in semi-supervised learning, we replace \mathbf{y} with $\bar{\mathbf{y}}$. In this way, we introduce unlabeled data into training. Semi-supervision can introduce more guidance to the network, making the feature space easier to generalize.

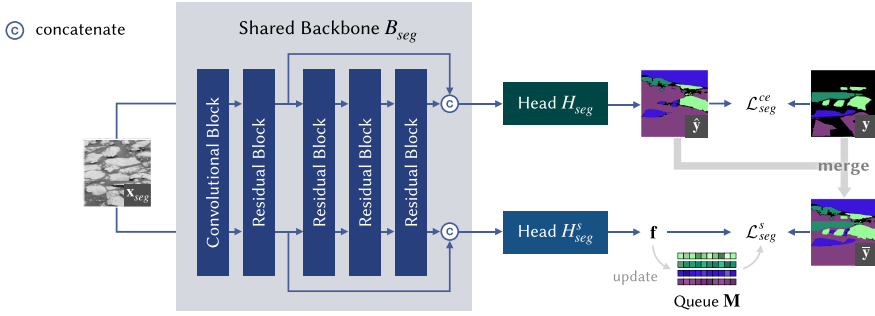


Fig. 7. The framework of our semi-supervised Mars image segmentation. There are two streams: (top to bottom) segmentation and element-wise inter-class contrastive learning. The two streams share the backbone and have independent heads.

5.3 Full Model

Similar to classification, the segmentation objective \mathcal{L}_{seg}^{ce} is also cross entropy:

$$\mathcal{L}_{seg}^{ce} = -\mathbb{E}_{x_{seg}} \mathbb{E}_{f_i} \left[\log \frac{\exp(f_{i,c_i})}{\sum_{c_j=1}^C \exp(f_{i,c_j})} \right], \quad (8)$$

where f_{i,c_j} represents the prediction of this pixel belonging to the c_j category, and c_i is the ground-truth label.

To make training stable and maximize the effectiveness of each learning design, we adopt a three-step training strategy: supervised contrastive learning pretraining, segmentation fine-tuning, and semi-supervised joint training.

We first pretrain the model with supervised contrastive learning alone, which provides a suitable feature space initialization for segmentation. We apply *Element-wise Inter-class Contrastive Learning* with ground-truth annotations.

After pretraining, the model is trained with contrastive and segmentation losses simultaneously. The training objective is

$$\mathcal{L}_{seg} = \mathcal{L}_{seg}^{ce} + \lambda_{seg}^s \mathcal{L}_{seg}^s, \quad (9)$$

where λ_{seg}^s controls the training balance between two objectives. To reduce the impact of contrastive learning, λ_{seg}^s is set to 0.001. The hyper-parameter τ in \mathcal{L}_{seg}^s is set to 0.07.

Finally, the model is trained with semi-supervised learning. We add *Online Pseudo Labeling* to element-wise inter-class contrastive learning. The previous training steps ensure the accuracy of segmentation, providing good initial pseudo labels for this step. With the training of the framework, we get better and better label estimates for unlabeled data, which promotes the semi-supervised learning process and enables the model to extract more separable feature representations.

The framework is shown in Figure 7. Our model is based on DeepLabv3+ [9]. The segmentation and contrastive streams share a same ResNet-101 [22] backbone B_{seg} . The heads in these two streams, H_{seg} and H_{seg}^s , are all DeepLabv3+ segmentation heads. The output dimension of H_{seg}^s is 128. The queue length in the memory bank M is 32.

6 EXPERIMENTS FOR MARS IMAGERY CLASSIFICATION

In this section, we evaluate our semi-supervised learning framework for classification. We first introduce the experimental settings, then show the comparison results against the state of the art, and finally provide ablation studies and more performance analysis.

Table 3. Results for MSL Rover Image Classification

Category	Method	Top-1 (%)
2*Baseline	AlexNet [28] in [48]	66.70
	ResNet-50 [22]	79.28 ± 1.76
Loss design	Focal loss [30]	82.86 ± 0.74
	Center loss [50]	82.91 ± 0.93
	Triplet loss [40]	84.87 ± 1.13
2*Training 2*design	MixUp [54]	76.19 ± 1.65
	CutMix [53]	80.61 ± 1.51
	Dropout [43]	83.37 ± 0.41
Semi- supervised learning	S4L [3], rotation	75.19 ± 1.73
	SupCon [26] + Linear Classifier	77.11 ± 2.89
	Pseudo labeling [29]	78.64 ± 0.04
	S4L [3], jigsaw	81.81 ± 2.33
	SupCon [26] + All Layers	83.22 ± 1.33
	SsCL [55]	90.81 ± 1.85
2*Re-sampling	Decoupling [25], cRT	80.94 ± 1.43
	Decoupling [25], LWS	81.30 ± 0.62
2*Re-weighting	Class-balanced loss [12]	80.02 ± 0.89
	LDAM-DRW [5]	82.12 ± 1.92
Ours		95.86 ± 1.63

6.1 Experiment Setup

For model training and evaluation, we use the MSL surface dataset [48]. For unsupervised similarity learning, we additionally collect 34K unlabeled color images from the NASA’s **Planetary Data System (PDS)**.¹

We first pretrain the model on ImageNet [14] following MoCo V2 [11]. Then, the model is fine-tuned on Mars data with Adam [27] optimizer for 30 epochs. The mini-batch size is set to 16 for \mathcal{L}_{cls}^{ce} and \mathcal{L}_{cls}^u and to 24 for \mathcal{L}_{cls}^s . Here, 24 equals the number of categories in the MSL surface dataset. The initial learning rate is set to 1e-3 for H_{cls} and 1e-6 for B_{cls} , H_{cls}^s , and H_{cls}^u , then multiplied by 0.1 at 20 and 25 epochs. The fine-tuning process takes about 1 hour with an Nvidia GeForce RTX 2080Ti.

6.2 Comparison Results

We compare the performance of our model with the other 10 related methods. For reliability, we run each experiment three times, then show the mean and standard deviation in Table 3.

We first carry out a comparison with work [48], in which the model performance is 66.70%. The low performance may be due to the limited capability of AlexNet [28]. The performance improves to 79.28% after we change the baseline to ResNet-50 [22], demonstrating the necessity of using good feature extractors.

As for the next step, we take works that aim to balance positive and negative samples through loss and training designs into account. We first consider three widely used losses: Triplet loss [40], Center loss [50], and Focal loss [30]. These loss designs enforce the embedded distances among different categories and can improve the classification performance by 3%~6%. However, they do not enrich the feature representation and therefore have limited effectiveness. Next, we consider

¹<https://pdsimage2.wr.usgs.gov/archive/MSL/>.

three widely used training strategies: MixUp [54], CutMix [53], and Dropout[43]. Their performances are below 85%, indicating that data augmentation and model regularization cannot solve the problem of Martian classification.

Since we introduce unlabeled data into our framework, the comparison also involves four semi-supervised learning methods. One is pseudo learning [29], which generates pseudo labels on unlabeled data. However, the performance degrades with pseudo learning, which is possible because the unlabeled data crawled from PDS is uncurated. Compared with the MSL surface dataset, the collected unlabeled data can be more long-tailed and unbalanced. Therefore, the pseudo labels on unlabeled data can mislead the classification model training to some extent. The second one is S4L [3], which applies self-supervised learning on unlabeled data. Among self-supervised tasks, the jigsaw pretext task improves the performance a little, while rotation brings down the classification accuracy. This could be because most unlabeled data contains Martian soil and rock, which does not offer much semantic information. In that case, unlabeled data is ambiguous for the jigsaw and rotation pretext tasks. The third one is SupCon [26]. We test two training policies: unfreezing a linear classifier with a frozen model base (originally used in [26]) and jointly fine-tuning all layers. Although SupCon also restricts positive samples to be of the same class, our Supervised Inter-class Contrastive Learning is quite different from SupCon. First, SupCon simply selects training samples randomly, while we control a training batch to equally cover all classes, which can efficiently cluster representation for classification. Second, SupCon splits the process of representation learning and classification. In comparison, our Supervised Inter-class Contrastive Learning is trained along with classification, which is more convenient to implement and better combines supervision information with contrastive learning. Because SupCon randomly selects training samples and splits the learning of representation and classification, it only achieves a top-1 accuracy of 77.11% for unfreezing a linear classifier and 83.22% for fine-tuning all layers, which is much lower than our 95.86%. The fourth one is SsCL [55], which is cross entropy and pseudo-label-based contrastive learning with pseudo label propagation according to similarity. However, the contrastive learning in SsCL is still unsupervised. Also, the unlabeled data is utilized by pseudo label co-calibration with similarity alignment, which is in direct proportion to the label number of each class and thus cannot handle the long-tailed and unbalanced unlabeled Martian data. Due to the unimproved contrastive learning scheme and the unsuitable unsupervised learning strategy, the top-1 accuracy of SsCL is only 90.81%, which is >5% worse than ours.

We also compare our methods with three state-of-the-art techniques for imbalanced data. Decoupling [25] is based on re-sampling, and Class-balanced loss [12] and LDAM-DRW [5] on re-weighting. Although these techniques can improve the performance of the ResNet-50 baseline, the effectiveness is limited compared with our strategies.

Our top-1 accuracy is 95.86%. The comparison experiments show that our training strategies facilitate the model to learn a better visual representation, which leads to better generalization and robustness.

6.3 Ablation Studies

Effect of Learning Strategies. The ablation settings of our designs are shown in Table 4. The accuracy first increases by 2.28% owing to the good feature learned by MoCo V2 [11] pretraining. The data augmentation we use in contrastive learning is stronger than normal data augmentations used for classification. If we directly apply our strong augmentation to classification, the performance degrades by 4.93%, which is because the supervision of classification is too weak to learn against strong augmentation.

Then we study the remarkable effect produced by the supervised inter-class contrastive learning \mathcal{L}_{cls}^s . The accuracy declines with the original contrastive learning formula, i.e., Equation (1),

Table 4. Ablation Studies of Our Semi-supervised Classification Framework

Method	Top-1 (%)
ResNet-50 (baseline)	79.28 ± 1.76
MoCo V2 pretraining	81.56 ± 1.36
MoCo V2 pretraining + Strong Aug.	76.63 ± 3.06
MoCo V2 pretraining + Equation(1)	75.66 ± 1.13
MoCo V2 pretraining + \mathcal{L}_{cls}^s	93.82 ± 1.57
MoCo V2 pretraining + \mathcal{L}_{cls}^u	78.65 ± 1.53
$\mathcal{L}_{cls}^s + \mathcal{L}_{cls}^u$	87.20 ± 0.93
Final (MoCo V2 pretraining + $\mathcal{L}_{cls}^s + \mathcal{L}_{cls}^u$)	95.86 ± 1.63

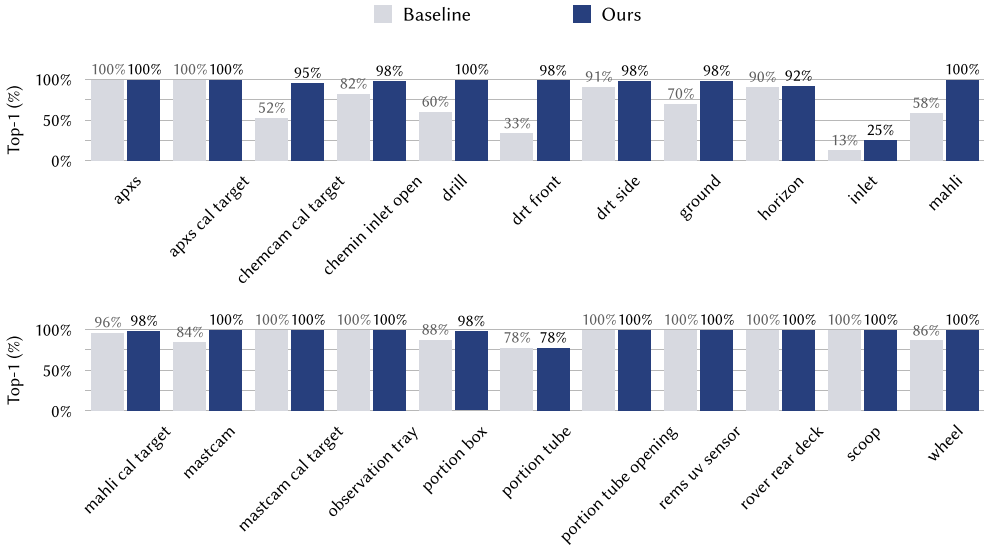


Fig. 8. Accuracy for each category of the baseline and our MSL rover image classifier.

confirming our motivation of introducing inter-class supervision. The accuracy also drops with unsupervised similarity learning \mathcal{L}_{cls}^u alone, which is in line with our analysis in Section 4. Also, removing MoCo V2 pretraining degrades performance even with \mathcal{L}_{cls}^s and \mathcal{L}_{cls}^u . The above experiments demonstrate the effectiveness of each component in our semi-supervised learning framework, and all components work together to provide the best performance, 95.86%.

A comparison of the detailed accuracy for each category is shown in Figure 8. Among all 22 categories in the testing set, our model is completely correct in 11 categories. Compared with the baseline, our model improves the performance by a large margin, especially for the *drt front*, *chemcam cal target*, *mahli*, and *drill* classes.

Some examples can be found in Figure 9. In the first row, we can see that our classifier is more robust to image quality degradations such as over-exposure, color channel error, and low resolution. There is a severe train-test object appearance gap in the *drt side* category as stated in Section 3. With the proposed semi-supervised learning scheme, this gap can be narrowed as shown in the first column of the second row in Figure 9. The other images in the second row show that our classifier is more robust to complex devices and terrains.

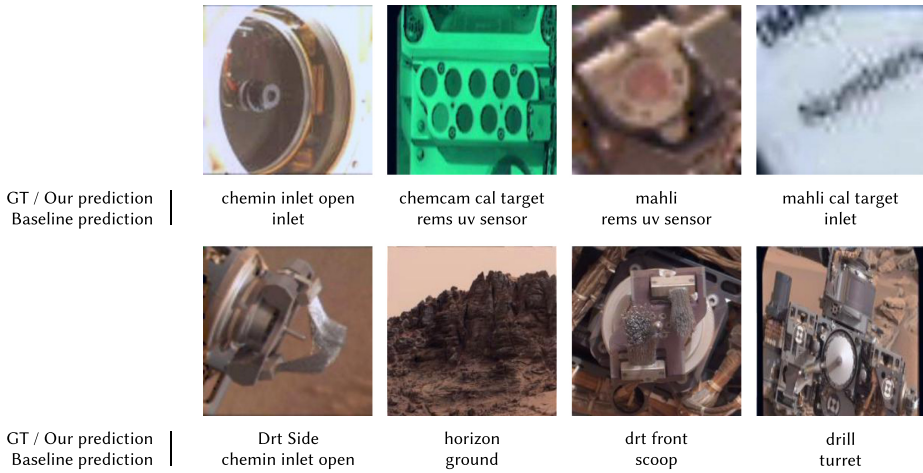


Fig. 9. Example classification results compared with the baseline.

Effect of Backbones. To evaluate the generalization of our framework, we test other classification backbones besides ResNet-50 [22]: VGG-16 [42], MobileNet V2 [39], RegNetX-1.6GF [37], ResNet-18 [22], ResNet-34 [22], ResNet-101 [22], and ResNet-152 [22] in Table 5. The baselines are first pretrained on ImageNet, then fine-tuned on the MSL dataset. In the implementation of these frameworks, we do not use the MoCo V2 [11] strategy to pretrain for convenience.

Our framework can improve the classification performance of all backbones, demonstrating the effectiveness of semi-supervised learning. It also works for lightweight models MobileNet V2 and RegNetX-1.6GF, which may be helpful for designing rover-edge deep models. The accuracy of ResNet-50 is better than VGG-16, ResNet-18, and ResNet-34, which is in line with ResNet-50 being more powerful than them on ImageNet. The performance is comparable with ResNet-101 and ResNet-152, which might mean that ResNet-50 is powerful enough for the Martian classification. Two lightweight models perform worse than ResNet-50, indicating that lightweight models are less robust when transferring from ImageNet to Mars data.

Parameter Selection. The effect of hyper-parameters λ_{cls}^s and λ_{cls}^u are shown in Tables 6 and 7. Too small values reduce the effect of contrastive learning, and too big values break the balance between different loss terms. Finally, $\lambda_{cls}^s = 1$ and $\lambda_{cls}^u = 0.2$ achieves the best performance.

6.4 Failure Case Study

Although our framework has greatly improved the classification performance, it may still make incorrect predictions. The *inlet* class suffers from a severe train-test gap. In the training set, the inlet images are mostly close-up. However, in the testing set, the inlets are only parts of the whole image. Accordingly, the accuracy of the *inlet* class is particularly low as shown in Figure 8.

More failure cases are shown in Figure 10. In the first row, although our classifier is robust to image quality degradation to some extent, it cannot cope with extremely low image quality. On the left of the second row in Figure 10, we can see that some images contain so many objects that even humans would have trouble classifying them. Our classifier may also fail when other objects occupy most areas of the image. As shown on the right of the second row in Figure 10, the ground occupies more areas than the target objects, and thus our classifier tends to recognize the images as *ground*.

Table 5. The Effect of Different Classification Backbones with Our Semi-supervised Learning

	ResNet-50	VGG-16	MobileNet V2	RegNetX-1.6GF
Baseline	79.28 ± 1.76	61.99 ± 2.86	80.00 ± 1.84	77.16 ± 6.65
Ours	87.20 ± 0.93	65.49 ± 1.68	85.70 ± 2.22	82.94 ± 1.50
	ResNet-18	ResNet-34	ResNet-101	ResNet-152
Baseline	72.69 ± 1.74	78.70 ± 2.06	81.02 ± 2.43	81.43 ± 2.29
Ours	75.43 ± 1.01	80.61 ± 0.90	87.36 ± 0.94	86.97 ± 0.67

The performance is measured by top-1 (%) accuracy on the MSL dataset. All models are based on ImageNet pretraining.

Table 6. Effect of Different λ_{cls}^s with $\lambda_{cls}^u = 0$

λ_{cls}^s	0.3	0.5	1.0	2.0	5.0
Top-1 (%)	88.40 ± 0.69	92.69 ± 0.67	93.82 ± 1.57	93.49 ± 0.27	93.44 ± 0.04

Table 7. Effect of Different λ_{cls}^u with $\lambda_{cls}^s = 1$

λ_{cls}^u	0	0.1	0.2	0.5	1.0
Top-1 (%)	93.82 ± 1.57	95.43 ± 0.69	95.86 ± 1.63	95.33 ± 1.01	94.74 ± 2.10

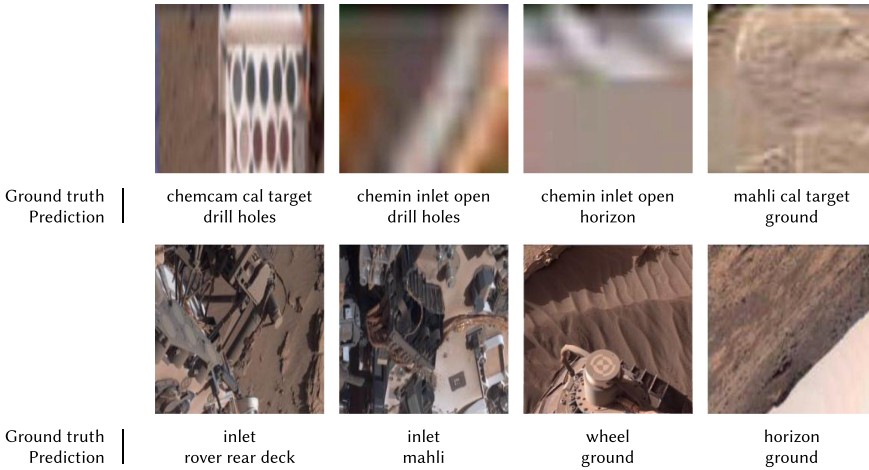


Fig. 10. Failure cases on MSL rover image classification.

7 EXPERIMENTS FOR MARS IMAGERY SEGMENTATION

In this section, we evaluate our method under semi-supervised segmentation learning settings on the AI4Mars dataset [45]. We show experimental settings, comparative results, ablation studies, and visual analysis.

7.1 Experiment Setup

We train the model with an SGD optimizer and a learning rate of 0.01. The polynomial annealing policy is applied for scheduling the learning rate. The batch size is set to 16. We first pretrain the encoder for 60 epochs with *Element-wise Inter-class Contrastive Learning*, then jointly train with

Table 8. Segmentation Performance on the AI4Mars Dataset

Method	ACC (%)	mIoU (%)
Mean Teacher [46]	72.24	55.34
ClassMix [34]	70.67	53.84
CutOut [15]	75.31	58.01
CutMix [53]	74.81	57.70
CCT [35]	77.08	48.19
LESS Within-image [56]	83.65	51.24
LESS Cross-image [56]	85.37	53.60
ReCo [31]	72.67	55.34
ReCo [31] + ClassMix [34]	68.01	51.25
ReCo [31] + CutOut [15]	75.61	58.19
ReCo [31] + CutMix [53]	73.63	56.76
ReCo [†] [31]	83.29	69.15
ReCo [†] [31] + ClassMix [34]	83.14	68.77
ReCo [†] [31] + CutOut [15]	83.23	68.73
ReCo [†] [31] + CutMix [53]	83.11	68.78
Ours	88.82	70.34

[†]indicates ImageNet pretraining.

contrastive learning loss and segmentation loss for 60 epochs, and finally with online pseudo labels for another 60 epochs.

7.2 Comparison Results

We compare our model with state-of-the-art representation learning, augmentation, and contrastive learning methods and their combinations. As shown in Table 8, our method achieves the best results on both segmentation accuracy and **mean Intersection over Union (mIoU)**.

In contrast to Mean Teacher [46], our method does not directly employ pseudo labels to train the segmentation model. Instead, pseudo labels are added to the memory bank for contrastive learning and assist segmentation by extracting representation through parameter sharing. This mechanism can mitigate the negative impact of inaccurate pseudo labels on segmentation performance.

The disadvantage of the data augmentation methods CutMix [53], CutOut [15], and ClassMix [34] is that the images of Mars are relatively simple, so the mixed data is not far from other data in the training set. As a consequence, the purpose of expanding data distribution cannot be achieved. Meanwhile, the gap between the training and the testing sets cannot be overcome by simply mixing images, but requires the network to learn better feature representations. In contrast, our method obtains a better feature space by making full use of the pixels of the training set. Also, contrastive learning can make the features of different categories more discriminative and generalized to the testing set.

We also compare semi-supervised segmentation learning methods. CCT [35] applies unlabeled data to randomly interfere with the feature, and constrains its output features to be consistent. However, consistency constraints alone cannot make features more inter-class separable. LESS [56] applies pixels within and between images for contrastive learning. It directly uses pixel features for building positive and negative samples, which leads to an excessive number of samples and thus makes training unstable. Moreover, the pseudo labels in LESS are fixed once generated. In comparison, we refine the pseudo labels with a retrained encoder for every epoch, which makes the pseudo labels more accurate. ReCo [31] also employs contrastive learning. The consumption of

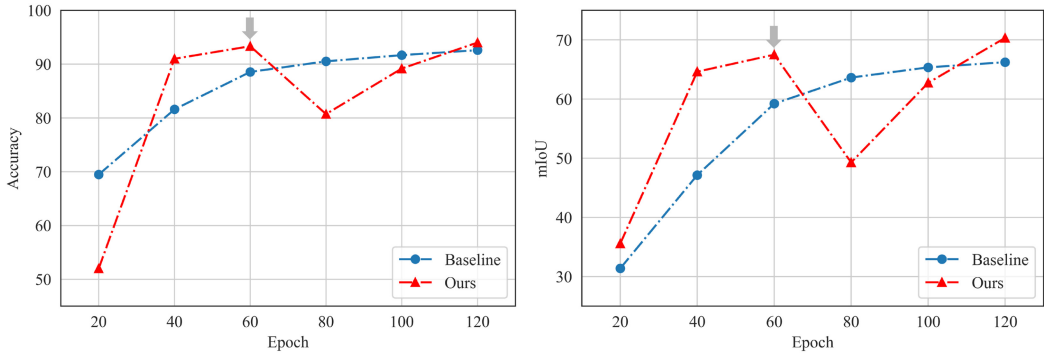


Fig. 11. Curves of segmentation accuracy (left) and mIoU (right) under different epochs. The gray arrow indicates that we change loss functions at the 60th epoch.

massive computation and memory is the key problem of pixel-level contrastive learning. To solve this problem, ReCo applies a regional contrast scheme and filters sparse queries and keys based on the confidence of segmentation prediction. We instead utilize the feature centers of each category as positive and negative samples. Compared with ReCo, our strategy requires fewer computational resources. Our features are also closer to the clustering center, leading to a better contrastive learning effect.

We also examine the effect of jointly using ReCo and augmentation. CutOut can slightly improve ReCo, while ClassMix and CutMix may even hurt the performance of ReCo. ImageNet pretraining has a significant positive effect. But when pretrained on ImageNet, none of the data augmentation methods can further improve the performance. Note that our model does not require ImageNet pretraining but still outperforms methods pretrained on the ImageNet dataset. This is because our semi-supervised method can extract more separable features, learning a representation even better than supervised learning on large-scale datasets.

7.3 Ablation Studies

Number of Epochs. Figure 11 shows the tendency of accuracy and mIoU of different epochs. For the first 60 epochs, as the training goes on, there is a significant performance gain. With contrastive learning pretraining, the network quickly achieves segmentation accuracy that exceeds the baseline algorithm. Also, for the speed of convergence, our network converges at about 60 epochs, while the baseline needs more than 100 epochs.

At the 60th epoch, we add the semi-supervised loss for co-training. The performance suddenly drops, which is because the network needs to adapt the feature space to the unlabeled data. Then, as the training progresses, the feature space is gradually refined with the help of the unlabeled data, which narrows the train-test gap and leads to a higher segmentation accuracy and mIoU score.

Effect of Modules. Table 9 shows the impact of different modules in our design. Accuracy (ACC) is the number of correct pixels divided by the total number of pixels. Mean accuracy (MAcc) is the average of the accuracy for each class. The frequency accuracy (FACC) is weighted and summed by the frequency of occurrence of each class. Jointly training the segmentation model with the contrastive loss contributes most to the overall performance gain. Contrastive learning pretraining improves the mIoU by 1%, and online pseudo labeling further promotes the comprehensive performance.

Table 9. Ablation Studies of Our Designs

Method	ACC (%)	MACC (%)	FACC (%)	mIoU (%)
Baseline	92.61	71.71	86.28	66.23
<i>w/o Joint Semi-Supervised Learning</i>	92.81	72.41	86.71	66.64
<i>w/o Pretraining</i>	93.78	74.75	88.39	69.30
<i>w/o Online Pseudo Labeling</i>	94.00	75.70	88.77	70.31
Final Method	94.02	75.68	88.82	70.34

Table 10. The Effect of Different Threshold Values

Threshold	ACC (%)	MACC (%)	FACC (%)	mIoU (%)
0.99	94.05	75.35	88.85	70.14
0.9	94.02	75.68	88.82	70.34
0.7	94.00	75.35	88.77	70.09
0.5	92.89	72.61	86.81	66.92

Table 11. The Effect of Different Loss Weight Settings

λ_{seg}^s	ACC (%)	MACC (%)	FACC (%)	mIoU (%)
0.1	93.92	75.83	88.65	70.29
0.01	93.39	75.64	88.75	70.26
0.001	94.02	75.68	88.82	70.34
0.0001	94.04	75.46	88.84	70.19

Effect of Threshold. In *Online Pseudo Labeling*, inaccurate labels will bring noise and interfere with the segmentation task. Therefore, we employ a threshold to control the pseudo label assignment of unlabeled data. Only predictions with confidence higher than the threshold can be appended to the memory bank for contrastive learning.

The effect of this threshold is shown in Table 10. We can see that as the threshold decreases, the accuracy and the mIoU of the segmentation task drop as well. Low threshold leads to inaccurate pseudo label annotations, which become training noises that interfere with segmentation learning. But if the threshold is too high, e.g., 0.99, there will be too few pseudo labels, which provides limited supervision. We finally set the threshold to 0.9.

Effect of Loss Weight. When jointly training the contrastive learning and segmentation task, the balance between the two loss functions affects the performance, as shown in Table 11. As the weight of contrastive learning λ_{seg}^s decreases, the segmentation accuracy becomes higher. When the value of λ_{seg}^s is too high, the network pays more attention to the features required for contrastive learning, while reducing the learning of the segmentation task. But when the value of λ_{seg}^s is too low, contrastive learning will have a limited impact on the training process. Finally, we find that $\lambda_{seg}^s = 0.001$ achieves the best comprehensive performance.

7.4 Visualizations and Interpretability

Feature Visualization. We visualize the output of the last layer in the encoder B_{seg} . As shown in Figure 12, features extracted by our method are more compact and separable than those extracted by the baseline algorithm. Benefiting from the assistance of contrastive learning, the distance between features of different categories increases, and the features of the same category are more concentrated.

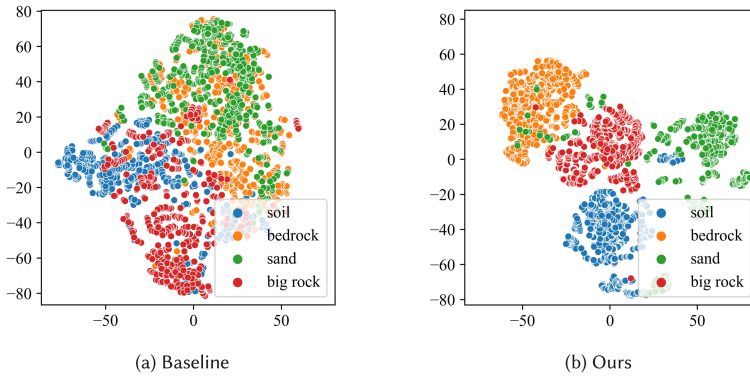


Fig. 12. Visualization of the extracted features for each pixel on the AI4Mars dataset.

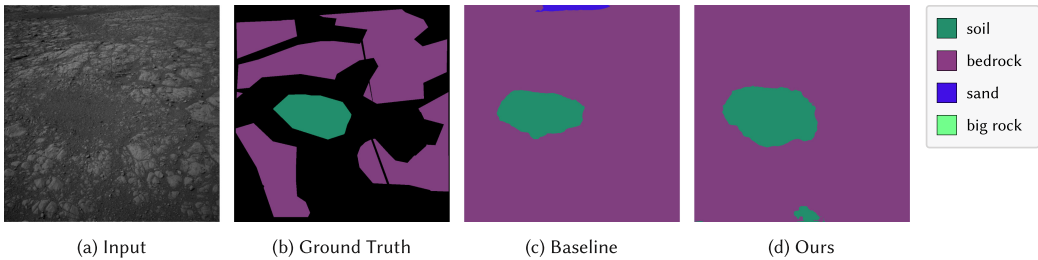


Fig. 13. Subjective segmentation results compared with the baseline.

From Figure 12, we also notice that the feature distance between soil and sand is relatively close, which is in line with soil and sand having a similar appearance. Meanwhile, big rocks have diverse appearances. Accordingly, their representations are more scattered and easier to be confused with other categories.

Subjective Segmentation Results. We show subjective segmentation results for a testing sample with many unlabeled areas in Figure 13. Compared with the baseline, our method has more reasonable predictions, such as recognizing the soil at the bottom of the image. In comparison, the baseline may misclassify rocky ground to be sand. The segmentation boundary is also more fine-grained in our prediction.

8 CONCLUSION

In this article, we propose a semi-supervised learning framework for Martian machine vision tasks. For classification, we extend contrastive learning to supervised inter-class and unsupervised similarity-only versions. For segmentation, we design element-wise contrastive learning and introduce extra supervision by online pseudo labeling. Experimental results demonstrate the superiority of our designs. In the future, we will extend our framework to more Martian vision tasks, such as object detection, tracking, and locating.

REFERENCES

- [1] Neil Abcouwer, Shreyansh Daftry, Siddarth Venkatraman, Tyler del Sesto, Olivier Toupet, Ravi Lanka, Jialin Song, Yisong Yue, and Masahiro Ono. 2021. Machine learning based path planning for improved rover navigation. In *IEEE Aerospace Conference*.

- [2] Abigail R. Azari, John B. Biersteker, Ryan M. Dewey, Gary Doran, Emily J. Forsberg, Camilla D. K. Harris, Hannah R. Kerner, Katherine A. Skinner, Andy W. Smith, Rashied Amini, Saverio Cambioni, Victoria Da Poian, Tadhg M. Garton, Michael D. Himes, Sarah Millholland, and Suranga Ruhunusiri. 2020. Integrating machine learning for planetary science: Perspectives for the next decade. *arXiv* (2020).
- [3] Lucas Beyer, Xiaohua Zhai, Avital Oliver, and Alexander Kolesnikov. 2019. S4L: Self-supervised semi-supervised learning. In *IEEE International Conference on Computer Vision*.
- [4] Katherine L. Bouman, Michael D. Johnson, Daniel Zoran, Vincent L. Fish, Sheperd S. Doeleman, and William T. Freeman. 2016. Computational imaging for VLBI image reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [5] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Aréchiga, and Tengyu Ma. 2019. Learning imbalanced datasets with label-distribution-aware margin loss. In *Conference on Neural Information Processing Systems*.
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *Conference on Neural Information Processing Systems*.
- [7] Andres Castano, Alex Fukunaga, Jeffrey J. Biesiadecki, Lynn Neakrase, Patrick L. Whelley, Ronald Greeley, Mark T. Lemmon, Rebecca Castaño, and Steve A. Chien. 2008. Automatic detection of dust devils and clouds on mars. *Machine Vision and Applications* 19, 5–6 (2008), 467–482.
- [8] Rebecca Castaño, Tara A. Estlin, Robert C. Anderson, Daniel M. Gaines, Andres Castano, Benjamin J. Bornstein, Caroline Chouinard, and Michele Judd. 2007. Oasis: Onboard autonomous science investigation system for opportunistic rover science. *Journal of Field Robotics* 24, 5 (2007), 379–397.
- [9] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *European Conference on Computer Vision*.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*.
- [11] Xinlei Chen, Haoqi Fan, Ross B. Girshick, and Kaiming He. 2020. Improved baselines with momentum contrastive learning. *arXiv* (2020).
- [12] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge J. Belongie. 2019. Class-balanced loss based on effective number of samples. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [13] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen O. Egiazarian. 2007. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* 16, 8 (2007), 2080–2095.
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. 2009. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [15] Terrance Devries and Graham W. Taylor. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv* (2017).
- [16] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2016. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 2 (2016), 295–307.
- [17] Gary Doran, Steven Lu, Maria Liukis, Lukas Mandrake, Umaa Rebbapragada, Kiri L. Wagstaff, Jimmie Young, Erik Langert, Anneliese Braunegg, Paul Horton, Daniel Jeong, and Asher Trockman. 2020. COSMIC: Content-based on-board summarization to monitor infrequent change. In *IEEE Aerospace Conference*.
- [18] Fabio Henrique Kiyoyiti dos Santos Tanaka and Claus Aranha. 2019. Data augmentation using GANs. *arXiv* (2019).
- [19] Murat Dundar, Bethany L. Ehlmann, and Ellen K. Leask. 2019. Machine-learning-driven new geologic discoveries at mars rover landing sites: Jezero and NE syrtis. *Earth and Space Science Open Archive* (2019), 23.
- [20] Tara Estlin, Rebecca Castano, Benjamin Bornstein, Daniel Gaines, Robert C. Anderson, Charles de Granville, David Thompson, Michael Burl, Michele Judd, and Steve Chien. 2009. Automated targeting for the MER rovers. In *IEEE International Conference on Space Mission Challenges for Information Technology*.
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [23] Xin Jin, Jianfeng Xu, Kazuyuki Tasaka, and Zhibo Chen. 2021. Multi-task learning-based all-in-one collaboration framework for degraded image super-resolution. *ACM Transactions on Multimedia Computing, Communications, and Applications* 17, 1 (2021), 1–21.
- [24] Daniel J. Jobson, Zia-ur Rahman, and Glenn A. Woodell. 1997. A multiscale retinex for bridging the gap between color images and the human observation of scenes. *IEEE Transactions on Image Processing* 6, 7 (1997), 965–976.
- [25] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. 2020. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*.

- [26] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. In *Conference on Neural Information Processing Systems*.
- [27] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Conference on Neural Information Processing Systems*.
- [29] Dong-Hyun Lee. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *International Conference on Machine Learning Workshops*.
- [30] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision*.
- [31] Shikun Liu, Shuaifeng Zhi, Edward Johns, and Andrew J. Davison. 2021. Bootstrapping semantic segmentation with regional contrast. *arXiv* (2021).
- [32] Zhiwu Lu, Zhenyong Fu, Tao Xiang, Peng Han, Liwei Wang, and Xin Gao. 2017. Learning from weak and noisy labels for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 3 (2017), 486–500.
- [33] Ali Momenasab. 2021. Machine learning for mars exploration. *arXiv* (2021).
- [34] Viktor Olsson, Wilhelm Tranhedén, Juliano Pinto, and Lennart Svensson. 2021. ClassMix: Segmentation-based data augmentation for semi-supervised learning. In *IEEE Winter Conference on Applications of Computer Vision*.
- [35] Yassine Ouali, Céline Hudelot, and Myriam Tami. 2020. Semi-supervised semantic segmentation with cross-consistency training. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [36] Matej Petkovic, Redouane Boumghar, Martin Breskvar, Saso Dzeroski, Dragi Kocev, Jurica Levatic, Luke Lucas, Aljaz Osojnik, Bernard Zenko, and Nikola Simidjievski. 2019. Machine learning for predicting thermal power consumption of the mars express spacecraft. *IEEE Aerospace and Electronic Systems Magazine* 34, 7 (2019), 46–60.
- [37] Ilija Radosavovic, Raj Prateek Kosaraju, Ross B. Girshick, Kaiming He, and Piotr Dollár. 2020. Designing network design spaces. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [38] Brandon Rothrock, Ryan Kennedy, Christopher T. Cunningham, Jeremie Papon, Matthew Heverly, and Masahiro Ono. 2016. SPOC: Deep learning-based terrain classification for mars rover missions. In *AIAA SPACE*. <https://arc.aiaa.org/doi/10.2514/6.2016-5539>.
- [39] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. MobileNetV2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [40] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. FaceNet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [41] Christopher J. Shallue and Andrew Vanderburg. 2018. Identifying exoplanets with deep learning: A five-planet resonant chain around kepler-80 and an eighth planet around kepler-90. *Astronomical Journal* 155, 2 (2018), 94.
- [42] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*.
- [43] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15, 56 (2014), 1929–1958.
- [44] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *IEEE International Conference on Computer Vision*.
- [45] R. Michael Swan, Deegan Atha, Henry A. Leopold, Matthew Gildner, Stephanie Oij, Cindy Chiu, and Masahiro Ono. 2021. AI4MARS: A dataset for terrain-aware autonomous driving on mars. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*.
- [46] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Conference on Neural Information Processing Systems*.
- [47] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv* (2018).
- [48] Kiri L. Wagstaff, You Lu, Alice Stanboli, Kevin Grimes, Thamme Gowda, and Jordan Padams. 2018. Deep mars: CNN classification of mars imagery for the PDS imaging atlas. In *Conference on Innovative Applications of Artificial Intelligence*.
- [49] Sara I. Walker, William Bains, Leroy Cronin, Shiladitya DasSarma, Sebastian Danielache, Shawn Domagal-Goldman, Betul Kacar, Nancy Y. Kiang, Adrian Lenardic, Christopher T. Reinhard, et al. 2018. Exoplanet biosignatures: Future directions. *Astrobiology* 18, 6 (2018), 779–824.
- [50] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. 2016. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*.
- [51] Thorsten Wilhelm, Melina Geis, Jens Püttschneider, Timo Sievernich, Tobias Weber, Kay Wohlfarth, and Christian Wöhler. 2020. DoMars16k: A diverse dataset for weakly supervised geomorphologic analysis on mars. *Remote Sensing* 12, 23 (2020), 3981.

- [52] Chenggang Yan, Zhisheng Li, Yongbing Zhang, Yutao Liu, Xiangyang Ji, and Yongdong Zhang. 2020. Depth image denoising using nuclear norm and learning graph model. *ACM Transactions on Multimedia Computing, Communications, and Applications* 16, 4 (2020), 1–17.
- [53] Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe. 2019. CutMix: Regularization strategy to train strong classifiers with localizable features. In *IEEE International Conference on Computer Vision*.
- [54] Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. 2018. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*.
- [55] Yuhang Zhang, Xiaopeng Zhang, Jie Li, Robert Qiu, Haohang Xu, and Qi Tian. 2022. Semi-supervised contrastive learning with similarity co-calibration. *IEEE Transactions on Multimedia* (2022). <https://ieeexplore.ieee.org/document/9732218>.
- [56] Xiangyun Zhao, Raviteja Vemulapalli, Philip Andrew Mansfield, Boqing Gong, Bradley Green, Lior Shapira, and Ying Wu. 2021. Contrastive learning for label efficient semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- [57] Boyan Zhou, Quan Cui, Xiu-Shen Wei, and Zhao-Min Chen. 2020. BBN: Bilateral-branch network with cumulative learning for long-tailed visual recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Received 22 August 2022; revised 25 August 2022; accepted 25 November 2022