


Automatic Detection and Segmentation of Barchan Dunes on Mars and Earth Using a Convolutional Neural Network

Lior Rubanenko , Sebastian Pérez-López, Joseph Schull, and Mathieu G. A. Lapôtre

Abstract—The morphology of isolated barchan dunes on Mars and Earth may shed light on the dynamic conditions that form them, their migration direction and the physical properties of the sediments composing them. Prior to this study, dune fields have been largely analyzed manually from aerial and satellite imagery, as automatic detection techniques are often not sufficiently accurate in outlining dunes. Here, we employ an instance segmentation neural network to detect and outline isolated barchan dunes on Mars and Earth. We train and test the model on martian targets using Mars reconnaissance orbiter (MRO) context camera (CTX) images, and find it sufficiently accurate (mAP=77% on the test dataset) to characterize dune field dynamics. Using our trained model, we detect and map the global distribution of barchan dunes relative to previously mapped dune fields, and find that barchan dunes are more abundant in the northern hemisphere than in the southern hemisphere. These contrasting abundances of barchans may reflect latitudinally dependent wind regimes, sediment supply, or sediment availability.

Index Terms—Machine learning, neural networks, planets: Mars, geology.

I. INTRODUCTION

SAND seas on Mars and Earth are riddled with dunes created by accumulating sand particles carried by the wind. These and other windblown features have been thoroughly studied *in situ* by rovers [1]–[3], which allow investigation of individual landforms, and remotely by orbiters [4]–[7], which provide a view of their dynamics at dune-field to global scales. In conditions of limited sand availability and approximately unidirectional wind, windblown landforms tend to take the shape of isolated and crescentic barchan dunes. Barchan dunes are characterized by horns pointing downwind, and a steep, angle of repose slipface in their lee. As the dip directions of slip-faces of barchan dunes are typically oriented in the dominant wind direction, the morphology of barchan dunes is routinely used to infer wind conditions on Mars and Earth through manual analysis of aerial or satellite imagery [6], [8]–[13]. However, performing such analyses at a global scale remained challenging and impractical until recently. To date, automatic detection

techniques based on traditional computer vision algorithms have been largely ineffective at identifying the outlines of dunes from images, due to the difficulty of separating the features of interest from the background (see Fig. 1). Digital elevation models may assist the detection process but are often not available at the required resolution. In the past decade, the advent of machine learning (ML) and deep learning (DL) has vastly improved object detection and classification from images. For example, algorithms based on support vector machines or R-Vine classifiers were reported to yield classification accuracy exceeding 80% for various types of dunes on Mars [14]–[16]. However, these techniques, which excel at detection and classification, are not sufficiently accurate for segmenting individual dunes and extracting their outlines, which is critical to determine their migration directions. Here, we employ Mask regional convolutional neural network (R-CNN), [17], a state-of-the-art instance segmentation neural network, to automatically detect and outline barchan dunes on Mars and Earth.

ML and DL techniques have great potential in the field of planetary science and planetary geomorphology in particular, as those fields often rely on analysis of remote-sensing data. The goal of this work is to demonstrate the efficacy of DL techniques in harvesting and analyzing planetary data, assess its accuracy by validating it with manually mapped dune fields on Mars and Earth, investigate the global distribution of barchan dunes on Mars relative to other landforms and discuss how transfer learning can be used to expand the applicability of an already trained model to remote sensing data on Earth and other terrestrial planetary bodies.

II. AUTOMATIC OBJECT DETECTION IN IMAGES

A. Traditional Object Detection Techniques

Traditional computer vision object detection techniques rely on sharp image gradients for edge detection [19]. As a result, these methods tend to be inaccurate when the objects have soft, less pronounced borders, and are prone to identifying false positives. In many cases, the visual similarity between dunes and their substrate make them blend in with their environment in aerial and satellite images, reducing the effectiveness of traditional edge detection techniques. To illustrate the performance of traditional non-ML computer vision algorithms in detecting dunes on Mars, we applied the Canny operator [18], a popular non-ML edge-detection algorithm, to a Mars Reconnaissance

Manuscript received August 20, 2021; accepted August 26, 2021. Date of publication September 3, 2021; date of current version September 27, 2021. (Corresponding author: Lior Rubanenko.)

The authors are with the Department of Geological Sciences, Stanford University, Stanford, CA 94305, California (e-mail: liorr@stanford.edu; sebpl01@stanford.edu; jschull@berkeley.edu; mlapotre@stanford.edu).

Digital Object Identifier 10.1109/JSTARS.2021.3109900




Fig. 1. Automatic detection of dunes on Mars in CTX images using traditional and ML methods. (a) A field of barchan dunes near 44.58°E , -41.45°N . (b) Dunes detected using the traditional, non-ML, Canny edge detection computer vision algorithm [18] with $g_{\min} = 50$ and $g_{\max} = 150$. Using different threshold values affects detection sensitivity and thus true- and false-detection statistics. (c) Dunes detected by the Mask R-CNN CNN. Throughout the figure, true positive and false positive detections are indicated in white and red, respectively.

Orbiter (MRO) Context Camera (CTX) [20] image that contains several dunes (see Fig. 1).

The Canny operator detects the edges of objects in an image by thresholding its gradient. The algorithm discards gradient values smaller than g_{\min} , and keeps gradient values greater than g_{\max} . Values between g_{\min} and g_{\max} are preserved only if they are connected to pixels whose gradient values are greater than g_{\max} . Consequently, the Canny operator is sensitive to the selection of these thresholds, which typically need to be manually optimized for each image. Although some techniques [21] may help in automating this thresholding process, reliably detecting faint objects in images without using ML remains challenging as even with optimized parameters, non-ML methods under-perform ML methods [22].

B. ML-Based Object Detection

In the past decade, artificial neural networks have revolutionized object detection in images. This family of deep supervised ML techniques achieves abstraction comparable to that of humans by stacking layers of parameters provided as inputs to nonlinear activation functions [23]. Each layer in the network extracts higher-level features from the previous layer: the first layers usually learn simple features such as edges or corners, and deeper layers can identify more abstract features such as faces, animals, or letters and numbers. Here, we employ Matterport's implementation of Mask R-CNN [17], [24], an instance segmentation neural network, to detect and outline barchan dunes on Mars and Earth. Supervised artificial neural networks are a type of ML algorithms that map problems (input) to solutions (output) by learning from labeled examples. Here, we briefly review how these models operate and refer interested readers to more detailed texts, such as [23].

Artificial neural networks are built from mathematical functions called *artificial neurons*, which are modeled after biological neurons. An artificial neuron takes one or more inputs and multiplies them by weights before feeding the outcome through a nonlinear activation function. In a neural network, neurons

are aggregated into layers organized such that the output of each layer is passed as input to the next layer. When analyzing images, it is useful to add convolutional and pooling layers to the model, which subsample or apply filters to the images and enhance the model's ability to identify localized features. In this type of CNN, each layer in the model extracts higher-level features from the previous layer.

Artificial neural networks excel at supervised learning tasks, such as finding the weights that, given a vector of inputs, best reproduce a vector of expected outcomes. To this end, the output of the last layer of a neural network is passed to a loss function that calculates the deviation of the model output from the expected output. Then, in an iterative optimization process called *training*, the loss function is minimized with respect to the weights employing methods such as back-propagation and gradient descent [23], [25]. The weight matrix that minimizes the model loss is then used to make predictions on inputs that the model has not seen as part of the training process. Neural networks which learn to recognize specific features in images are complex and may require hundreds of billions of parameters in some cases [26]. Due to their complexity, DL models typically employ techniques such as L_2 regularization, normalization layers, or dropout [23], to mitigate overfitting or gradient saturation, which affects learning efficiency. For example, in L_2 regularization (or *weight decay*), the L_2 norm of the weights matrix multiplied by a regularization parameter, ω , is added to the loss function to penalize it. As the goal of the training process is to minimize the model loss, regularization acts to decrease the magnitude of the model weights, which reduces the complexity of the model. Similarly, in the process of *dropout* neurons are randomly omitted from the network to reduce the number of degrees of freedom of the model [23].

III. METHODS

Mask R-CNN is an instance segmentation neural network that detects the locations of objects in images and the pixels contained within their outlines (masks). Building on several




Fig. 2. A simplified description of the architecture of Mask R-CNN, the neural network employed in this study. Input images are fed through a backbone neural network, chosen here to be ResNet 101 [29], which outputs a feature map that highlights objects. The feature map is passed to a RPN, trained to separate the background from objects and propose ROIs. ROIs, along with the feature map, are passed through an algorithm called *ROI align*, whose role is to fix inaccuracies due to roundoff errors. Finally, the information is passed through additional layers to different loss functions that are optimized to classify, find the bounding box, and segment detected objects.

antecedent techniques [27], [28], Mask R-CNN operates in three stages. First, a CNN called *backbone* extracts a feature map from an input image, typically employing well-studied architectures such as ResNet [29]. The number of layers in the backbone affects the model’s complexity and needs to be manually selected to avoid overfitting or underfitting. In the second stage, a region proposal network (RPN) scans the feature map and proposes positive and negative *anchors*—parts of the image that are likely to contain objects or background, respectively. The boundaries of positive anchors are determined by minimizing the ratio between the intersection and the union of the predicted bounding box with the manually labeled bounding box (also called *intersection over union*, or IoU). Then, positive anchors are joined and filtered employing nonmaximum suppression [30] to form regions of interest (ROIs). In the third and last stage, these ROIs are examined by a feature pyramid network [31] to determine their classification and their binary pixel mask employing several loss functions. Fig. 2 shows a schematic of the model architecture. When using the model to predict and segment objects in images, it returns a list of objects’ binary masks, along with labels and detection confidence. Compared with non-ML techniques such as the Canny operator, Mask R-CNN demonstrably performs better in identifying and outlining individual barchan dunes, as empirically shown in Fig. 1. A more quantitative comparison of traditional and DL based object detection techniques can be found in [22] and references therein.

1) *Training Mask R-CNN to Detect Barchan Dunes*: Our study focuses on the detection of isolated barchan dunes. To compose a training dataset, we extracted images of dune fields from the CTX global mosaic (5 m/pixel, [32]). Poleward of latitude $\pm 70^\circ$, where dune fields are more abundant [7], we extracted all images from the global mosaic. Between latitudes $\pm 70^\circ$, we focused our efforts on areas previously marked as dune fields in the Mars Global Digital Dune Database [7]. A total of 137 111 images were extracted, cylindrically projected to preserve distances, standardized, and cropped to a resolution of 832×832 pixels for memory considerations. Finally, because the input to Mask R-CNN is a three-channel (RGB) image, we replicate the single-channel CTX image three times following the procedure of [33].

TABLE I
THREE EXAMPLE SETS OF MODEL CONFIGURATIONS AND THEIR COMPUTED MEAN AVERAGE PRECISION (mAP) VALUES

Configuration	backbone, ω , P_{aug}	mAP _{val}	mAP _{test}
1	ResNet50, 10^{-5} , 0.25	0.62	0.6
2	ResNet101, 10^{-6} , 0.25	0.71	0.71
3	ResNet101, 10^{-6} , 0.5	0.78	0.77

The performance of our optimized model can be evaluated by its corresponding mAP score, which takes into account both false negatives and false positives. Higher mAP indicated better agreement between the detected bounding boxes and the labeled bounding boxes. mAP_{val} and mAP_{test} refer to the mAP score evaluated for the validation ($n = 200$) and test ($n = 50$) datasets.

We prepared a training database consisting of 1076 images that were randomly selected from the extracted CTX images. To ensure the training images were as uniformly distributed as possible across Mars, we binned the surface in equal-area bins and sampled 10 – 50 images from each bin (depending on the local density of dunes) to create a geographically balanced dataset. In addition to an imagery dataset, training Mask R-CNN requires a list of binary masks that outline the objects of interest in each training image. Using Labelbox’s online instance segmentation platform [34], we outlined 9543 instances of barchan dunes in the images, opting to ignore connected, barchanoid ridges, linear, and other complex dune types. Next, we divided the labeled data into two subsets: a training dataset (80% of the images) and a validation dataset (20% of the images) to be used to evaluate the hyperparameters of the model. We trained the model for 120 training epochs (see Appendix A for detailed training information) and used the ResNet101 backbone. At the end of the training process, the model weights were saved and used to detect dunes in extracted unlabeled images.

2) *Model Accuracy*: Model accuracy in classification tasks can be evaluated by dividing the number of correctly classified features by the total number of features. In detection, where the location of the object is as important as its classification, accuracy is typically estimated through methods such as the mean Average Precision (mAP) that consider true positives (TP), false positives (FP) and false negatives (FN). To calculate mAP_t, a threshold IoU t is selected (typically, $t = 0.5$) such that objects with $\text{IoU} > t$ are considered true positives. For a given t , the average precision (AP) for a given image is computed as the definite integral of the *precision*, given by $TP/(TP + FP)$, with respect to the *recall*, given by $TP/(TP + FN)$ from 0 to 1. The mAP is the mean of the AP for all the images in the dataset. As the model hyperparameters were tuned on the validation dataset, the mAP computed using it may be overestimated. To address that, we annotated an additional test dataset consisting of 50 randomly selected images, and compute its mAP value. In Table I, we show the mAP values for the validation dataset in addition to the test dataset.

3) *Model Initialization, Hyperparameters, and Image Augmentation*: Two algorithms were tested to optimize model weights: stochastic gradient descent (SGD, [23]) and adaptive gradient descent (Adam, [35]). We find that, as also noted in previous studies [17], SGD empirically outperforms the adaptive

TABLE II
A LIST OF HYPERPARAMETERS USED THROUGHOUT THIS WORK FOR THREE
EXAMPLE MODEL CONFIGURATIONS

Configuration	1	2	3
Learning Rate	10^{-3} (20 epochs), 10^{-5} (40 ep.), 10^{-7} (60 ep.)		
Learning Momentum	0.9	0.9	0.9
Backbone	ResNet50	ResNet101	ResNet101
Batch size	3 images		
Weight Decay Parameter ω	10^{-5}	10^{-6}	10^{-6}
P_{aug}	0.25	0.25	0.5
RPN anchor scale	(32, 64, 128, 256, 512)		
RPN classification loss weight	1	1	2
RPN bounding box loss weight	1	1	1
MRCNN bounding box loss weight	1	1	1
MRCNN classification loss weight	1	1	1
RPN Non-Max Suppression Threshold	0.9	0.9	0.9

For relevant mAP values, please see Table I.

methods, which tend to overshoot (miss) Mask R-CNN’s loss minima.

Weights of ML models can either be initialized with random values or with values that were previously optimized for a different dataset. The latter method, referred to as *transfer learning* [23], [36], [37], accelerates model convergence but may lead to inaccuracies if the original model is under- or overfitted. We tested both random weight initialization and transfer learning, using weights from the common objects in Context (COCO) dataset [38] and weights derived from a crater detection model that was trained using digital terrain model (DTM) data [33]. As expected, transfer learning outperformed random weight initialization but no significant differences were found between the COCO initialization and the crater model initialization. We eventually elected to use COCO, as it yielded a slightly higher model mAP (see Table I). We note that the small difference between the crater weights [33] and COCO initializations may result from the fact that [33] also trained their model on the COCO dataset, and their additional training on DTM does not significantly contribute to the convergence of our model.

In contrast to model weights, which are optimized during training, model hyperparameters, which control the training process, are prescribed at the onset. Mask R-CNN has a wide range of hyperparameters that can be adjusted to improve model optimization. For example, the user may tune the depth (number of layers) of the backbone, the typical scale of anchors used by the RPN, or adjust parameters related to the training itself such as the learning rate and learning momentum, which affect model convergence rate and accuracy [17]. The model we adopted for dune detection achieved an mAP_{50} of 0.77 on the test dataset, which we found to be sufficient to gather representative morphometric statistics from dune outlines (see Table I).

Our training strategy was to first overfit the model, and then gradually add L_2 regularization and image augmentation [39] to reduce validation loss. Several popular image augmentation techniques were tested. Through an iterative process, we found that rotating the image by a random angle effectively reduced false identification of impact craters as dunes due to the circular shape of craters. Additionally, varying the image contrast helped improve detection of dunes in CTX with very low contrast. The choice of the regularization parameter, ω , and the probability of augmenting an image during training, P_{aug} , are also hyperparameters of our model. We tested a wide range of ω and P_{aug} , and empirically found a relatively small ω and high P_{aug} performed best (see Table I).

The total loss outputted by Mask R-CNN is the weighted sum of the individual loss functions of the RPN and the classifier. By adjusting the weights of this sum (not the model weights optimized during training), the user can tune different aspects of model performance, such as the losses associated with detecting objects and their bounding boxes, and the loss associated with classifying detected objects and their masks. Consequently, these loss weights are important hyperparameters. We empirically found that doubling the loss weight controlling the performance of the RPN in distinguishing objects from background (RPN classification loss) improved the model mAP by approximately 10%. Adjusting other loss weights did not have a measurable effect on model loss and accuracy.

IV. RESULTS

A. Detection of Barchan Dunes on Mars

Of the 137 111 images in our dataset, 55 674 were found by the model to contain objects suspected to be barchan dunes. Because barchan dunes rarely appear as solitary landforms but tend to occur in fields of isolated dunes, all downselected images containing less than two objects were discarded to avoid the inclusion of spurious detections and therefore, increase the robustness of our results. Our final compilation contains 1 005 701 objects, distributed across 36 656 images. Fig. 3 illustrates the global distribution of detected barchan dunes, displayed as the number density of detected objects per 100 km^2 . Our model was not optimized to detect all instances of barchan dunes (see Table I, which lists models mAP values). However, since accuracy is not expected to vary systematically across the surface of Mars, Fig. 3 provides a reliable estimate of the relative spatial number density of martian barchan dunes on Mars.

To characterize the latitudinal distribution of barchan dunes on Mars, we calculate the ratio between the surface area of barchan dune fields and the surface area of all dune fields between latitudes $\pm 70^\circ \text{N}$ [Fig. 3(b)]. First, we compute the surface area of a barchan dune field contained within a single CTX image as the area of the convex hull defined by the centers of mass of dunes found within the image, as illustrated in Fig. 3(c). Then, we bin and sum the areas of these convex hulls in bins of 10° latitude. To calculate the ratio between the area of barchan dune fields and all other types of dune fields, we similarly bin and sum the area of all types of dune fields as reported by [7]. Although many dunes are located in high latitudes, we choose to focus

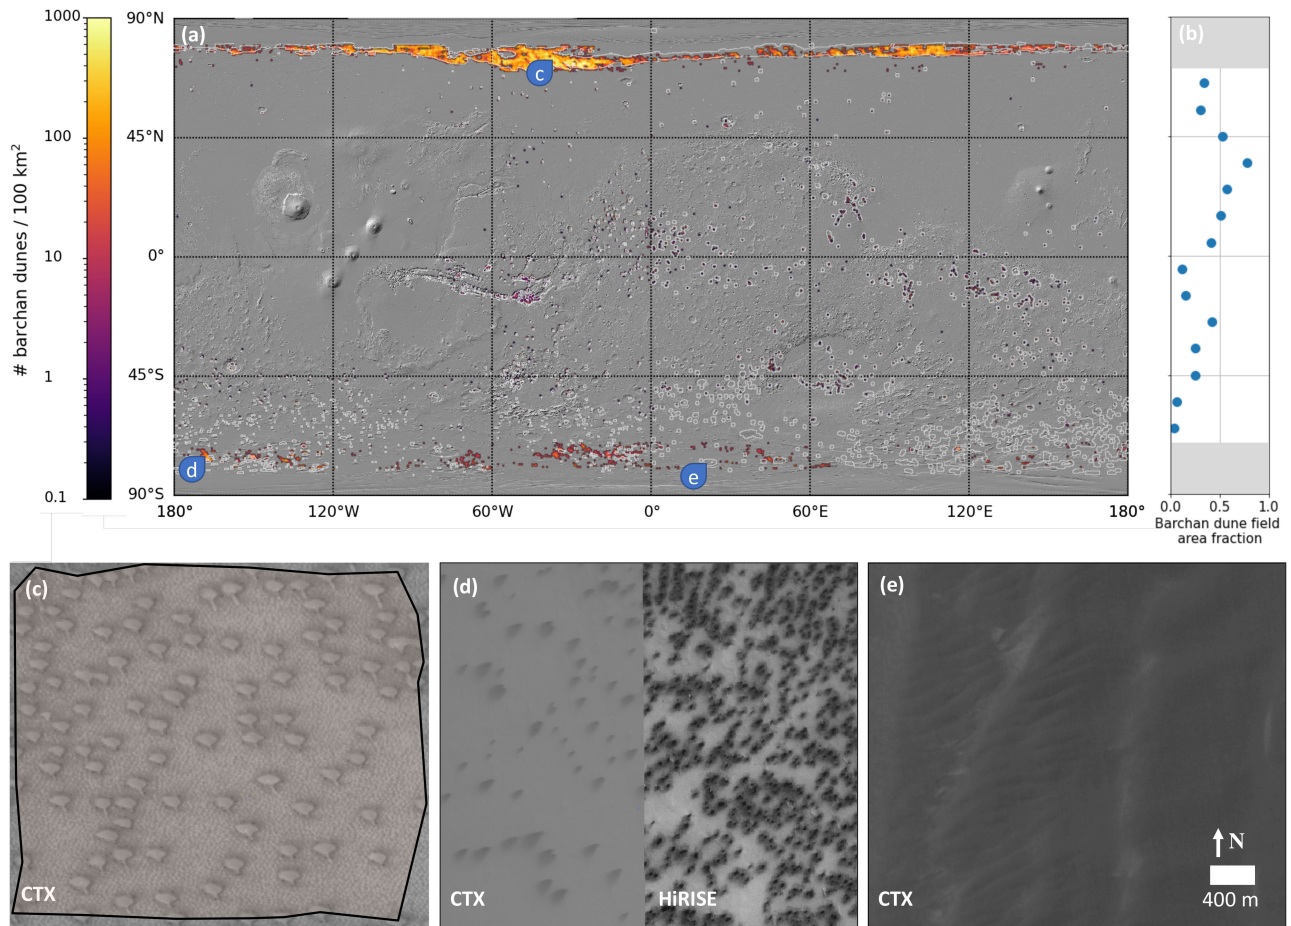


Fig. 3. (a) The global number density of barchan dunes on Mars drawn over the MRO CTX global mosaic [32]. White contours are dune fields of all types, as previously mapped by [6], [40], [41], and [7]. In the northern polar region, the majority of the dune fields on Mars are barchan dunes, in contrast to the southern polar region. In some cases, the barchan dune density in the north polar region exceeds 1000 dunes / 100 km². (b) The area fraction of barchan dune fields normalized by the area fraction of all dune fields for a given latitude, plotted in bins of 10° latitude. (c)–(e) Three examples of objects identified as barchan dunes by the neural network. (c) A true positive detection of barchan dunes near the north pole (−47.7°E, 75.3°N). Many of the isolated barchan dunes in the northern polar region have an atypical extended horn or a second slip-face, which may be caused by bimodal-wind regime or collisions between dunes [8]. The black outlines indicate the boundaries of the convex hull used to estimate the surface area coverage of the dune field in the image. (d) Objects classified as barchan dunes in a CTX image (right) were discovered to potentially be a likely false positive in a higher resolution HiRISE image (−167.4°E, −76.2°N). (e) A true negative example of dunes in South Crater (24.1°E, −76.5°N).

this analysis to $\pm 70^\circ\text{N}$ to avoid complications linked to suspected noneolian features [such as those illustrated in Fig. 3(d)]. In the northern hemisphere, we find barchan dunes constitute $\sim 40 - 80\%$ of the dune fields in areal extent. In the southern hemisphere, this ratio is significantly lower, $\sim 0 - 45\%$. This difference may reflect latitudinally dependent wind regimes or differences in sediment supply and availability. Notably, several studies reported evidence for the presence of near-surface ice poleward of latitude -45°N [42], [43]. The relative paucity of barchan dunes south of -45°N could thus reflect a decreased sand availability due to stabilization by near-surface volatiles.

In the north-polar region ($> 70^\circ\text{N}$), we find $\sim 92\%$ overlap between dunes previously mapped by [6] and the dunes detected by our model, which indicates that the majority of dune fields near the north pole consist of isolated barchan dunes. The overlap is computed as the surface area of dunes in our dataset divided by the surface area of both datasets. In some cases, as shown in Fig. 3(c), isolated barchan dunes are asymmetric, with an elongated horn or a second slip-face, which may form as a

result of wind bimodality or collisions between dunes [8], [9]. In the south polar region ($< -70^\circ\text{N}$), the average areal overlap between our dataset and [6] is $\sim 14\%$, reflecting the propensity of south-polar dunes to occur as connected rather than as isolated dunes [40], which could reflect a greater availability of loose sand at high southern latitudes [44]. Potentially contributing to this discrepancy is the abundance of dark, possibly sublimation-driven features, which at CTX resolution have a similar appearance to some of the fainter dunes in our training set [see Fig. 3 (d)]. The noneolian nature of these features only becomes more evident upon inspection of higher-resolution imagery from the high resolution imaging science experiment (HiRISE).

B. Detecting Terrestrial Barchan Dunes Employing Our Model Trained on Mars

The weights produced by the neural network are not only suitable for detecting barchan dunes on Mars, but may also be used to detect barchan dunes on other planetary bodies.




Fig. 4. Detecting terrestrial barchan dunes using a model that was trained on martian barchan dunes. (a) Dunes in the northern Sahara, (-12.98°E , 27.69°N ; source: Google Earth, CNES/Airbus). (b) Barchan dunes in the Bodélé Depression, Chad (17.79°E , 16.68°N ; source: Google Earth, Maxar Technologies).

Here, we present a preliminary analysis, which demonstrates the performance of our model, trained only using satellite imagery of the martian surface, when applied directly to terrestrial dune fields without any further training. Fig. 4 illustrates the model’s performance on two terrestrial barchan dune fields. The dunes in the first site we studied, located in the northern Sahara, are composed of beach sand and migrate toward the south-southwest. The second dune field is located in the Bodélé Depression, a $\sim 133500\text{ km}^2$ paleolake in Chad [45]. The dunes in the depression are composed of low-density diatomite sand, mobilized by strong winds coming from the northeast. Unlike most barchan dunes on Mars, which are composed of basaltic sand [46]–[48] and appear darker than their substrate, dunes in the Bodélé Depression are lighter than the surface over which they migrate due to the light-toned source material [45].

We find that even without further tuning of the model weights using terrestrial examples, it produces useful results, detecting $\sim 50 - 70\%$ of all barchan dunes in the scenes we reviewed noting that a more comprehensive survey is warranted to accurately determine the performance the model in analyzing terrestrial images. In some cases, such as in the Bodélé depression, Chad, the dune-background contrast is opposite to that of typical martian examples. This caused the model to initially underperform and detect only a small fraction of the dunes in the field. We significantly improved the model performance up to a $\sim 70\%$ detection rate by inverting image colors, to better mimic martian dune-substrate contrasts. Therefore, we recommend users

interested in employing our model weights to detect terrestrial barchans to preprocess the images by enhancing their contrast or inverting their color, where the dunes are brighter than their substrate. To further increase the model’s accuracy, our weights could be further optimized using labeled terrestrial examples in a process of transfer learning [36].

V. CONCLUSION

Research in Earth and planetary sciences increasingly relies on global, high-resolution remote sensing data acquired by orbiting satellites. Performing systematic analyses of surface features at these scales, whether manually or even automatically through non-ML techniques, has remained a challenge. In this work, we employ an instance segmentation neural network, Mask R-CNN, to detect and outline barchan dunes on Mars and Earth. We trained the model on 9543 isolated barchan dunes, achieving a mAP score of 77%, which we found sufficient to investigate morphological trends in barchan dune fields.

Using the trained model weights we detected and outlined over 10^6 features, demonstrating the efficacy of DL techniques in harvesting planetary remote sensing imagery data. We compare the global surface area coverage of barchan dune fields detected by our model to that of previously mapped dune fields, and find a latitudinal trend in the relative abundance of barchan dunes. In the northern hemisphere, on average $\sim 60\%$ of the total area covered by dune fields is covered by barchan dunes—much higher than the relative area fraction of barchan dune fields in the southern hemisphere, which is on average $\sim 30\%$. Poleward of latitude -45°N , we find that much fewer dune fields contain barchan dunes.

The results presented in this study could be used to further explore barchan dunes on Mars and Earth. As demonstrated above, our trained model weights could be used to detect barchan dunes on Earth with an accuracy of $\sim 50 - 70\%$. If a higher accuracy is desired, our trained weights could be used for transfer learning of a separated terrestrial instance segmentation DL model by adding a small number of terrestrial examples. Finally, the global catalog of barchan dunes we compiled in this study could be used to infer the migration direction of barchan dunes globally on Mars, which could prove useful for studying surface winds and constraining global climate models.

APPENDIX A

COMPLETE LIST OF THE MODEL HYPERPARAMETERS

For the purpose of reproducing our results, we list below the hyperparameters of the three models discussed in the manuscript. Hyperparameters not appearing in Table II were set to their default values in the Matterport implementation of Mask R-CNN. The model was trained on an NVIDIA GeForce RTX 3090, with 24 GB VRAM for 120 epoches. The walltime of each training session was approximately 24 h.

ACKNOWLEDGMENT

The authors would like to thank Lori Fenton, Ryan Ewing, Andrew Gunn, and Michael Hasson for helpful comments and discussion. The authors would also like to thank Adrian Brown

and two anonymous reviewers for constructive reviews. Model weights and link to code that can be used to reproduce the results are archived here [49].

REFERENCES

- [1] R. Greeley *et al.*, "Wind-related processes detected by the spirit rover at Gusev Crater, Mars," *Science*, vol. 305, no. 5685, pp. 810–813, 2004.
- [2] R. Sullivan *et al.*, "Wind-driven particle mobility on Mars: Insights from Mars Exploration Rover observations at "El Dorado" and surroundings at Gusev Crater," *J. Geophys. Res., Planets*, vol. 113, no. E6, 2008.
- [3] M. Lapôtre and E. Rampe, "Curiosity's investigation of the bagnold dunes, gale crater: Overview of the two-phase scientific campaign and introduction to the special collection," *Geophys. Res. Lett.*, vol. 45, no. 19, pp. 10–200, 2018.
- [4] N. Bridges, F. Ayoub, J. Avouac, S. Leprince, A. Lucas, and S. Mattson, "Earth-like sand fluxes on Mars," *Nature*, vol. 485, no. 7398, pp. 339–342, 2012.
- [5] S. Silvestro, D. Vaz, L. Fenton, and P. E. Geissler, "Active aeolian processes on Mars: A regional study in Arabia and Meridiani Terrae," *Geophys. Res. Lett.*, vol. 38, no. 20, 2011.
- [6] R. K. Hayward, L. Fenton, and T. N. Titus, "Mars global digital dune database (MGD3): Global dune distribution and wind pattern observations," *Icarus*, vol. 230, pp. 38–46, 2014.
- [7] L. K. Fenton, "Updating the global inventory of dune fields on Mars and identification of many small dune fields," *Icarus*, vol. 352, 2020, Art. no. 114018.
- [8] E. J. Parteli, O. Durán, M. C. Bourke, H. Tsoar, T. Pöschel, and H. Herrmann, "Origins of barchan dune asymmetry: Insights from numerical simulations," *Aeolian Res.*, vol. 12, pp. 121–133, 2014.
- [9] H. Tsoar and E. J. Parteli, "Bidirectional winds, barchan dune asymmetry and formation of seif dunes from barchans: A discussion," *Environ. Earth Sci.*, vol. 75, no. 18, pp. 1–10, 2016.
- [10] C. S. Breed, M. J. Grolier, and J. F. McCauley, "Morphology and distribution of common 'sand' dunes on Mars: Comparison with the Earth," *J. Geophys. Res., Solid Earth*, vol. 84, no. B14, pp. 8183–8204, 1979.
- [11] H. Tsoar, R. Greeley, and A. R. Peterfreund, "Mars: The north polar sand sea and related wind patterns," *J. Geophys. Res., Solid Earth*, vol. 84, no. B14, pp. 8167–8180, 1979.
- [12] P. Lee and P. C. Thomas, "Longitudinal dunes on Mars: Relation to current wind regimes," *J. Geophys. Res., Planets*, vol. 100, no. E3, pp. 5381–5395, 1995.
- [13] Z. Zhang, Z. Dong, G. Hu, and E. J. Parteli, "Migration and morphology of asymmetric barchans in the central hexi corridor of northwest China," *Geosciences*, vol. 8, no. 6, 2018, Art. no. 204.
- [14] L. Bandeira, J. S. Marques, J. Saraiva, and P. Pina, "Advances in automated detection of sand dunes on Mars," *Earth Surface Processes Landforms*, vol. 38, no. 3, pp. 275–283, 2013.
- [15] M. Azzaoui, M. Adnani, H. El Belrhiti, I. Chaouki, and C. Masmoudi, "Detection of barchan dunes in high resolution satellite images," international archives of the photogrammetry," *Remote Sens. Spatial Inf. Sci.*, vol. 41, 2016.
- [16] D. Carrera, L. Bandeira, R. Santana, and J. A. Lozano, "Detection of sand dunes on Mars using a regular vine-based classification approach," *Knowl.-Based Syst.*, vol. 163, pp. 858–874, 2019.
- [17] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2961–2969.
- [18] J. Canny, "A Computational approach to edge detection," *IEEE Trans. Pattern Ana. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
- [19] D. A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*. Pearson Education, Inc., Upper Saddle River, NJ 07458, 2012.
- [20] M. C. Malin *et al.*, "Context camera investigation on board the mars reconnaissance orbiter," *J. Geophys. Res., Planets*, vol. 112, no. E5, 2007.
- [21] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. 9, no. 1, pp. 62–66, Jan. 1979.
- [22] A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, "Deep learning for computer vision: A brief review," *Comput. Intell. Neurosci.*, pp. 404–411, 2018.
- [23] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, vol. 1, no. 2. Cambridge, MA, USA: MIT Press, 2016.
- [24] W. Abdulla, "Mask R-CNN for object detection and instance segmentation on keras and tensorflow," 2017. [Online]. Available: https://github.com/matterport/Mask_RCNN
- [25] S. Ruder, "An overview of gradient descent optimization algorithms," 2016. *arXiv:1609.04747*
- [26] J. Rasley, S. Rajbhandari, O. Ruwase, and Y. He, "Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters," in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3505–3506.
- [27] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 580–587.
- [28] R. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1440–1448.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [30] A. Neubeck and L. Van Gool, "Efficient non-maximum suppression," in *Proc. IEEE 18th Int. Conf. Pattern Recognit.*, 2006, vol. 3, pp. 850–855.
- [31] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2117–2125.
- [32] J. Dickson, L. Kerber, C. Fassett, and B. Ehlmann, "A global, blended CTX mosaic of Mars with vectorized seam mapping: A new mosaicking pipeline using principles of non-destructive image editing," in *Proc. Lunar Planetary Sci. Conf.*, 2018, vol. 49, pp. 1–2.
- [33] M. Ali-Dib, K. Menou, A. P. Jackson, C. Zhu, and N. Hammond, "Automated crater shape retrieval using weakly-supervised deep learning," *Icarus*, vol. 345, 2020, Art. no. 113749.
- [34] Labelbox Image Labeling Platform. [Online]. Available: <https://labelbox.com/education>
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015, *arXiv:1412.6980*.
- [36] Z. U. Rehman *et al.*, "Recognizing apple leaf diseases using a novel parallel real-time processing framework based on mask rcnn and transfer learning: An application for smart agriculture," *IET Image Process.*, vol. 15, no. 10, pp. 2157–2168, 2021.
- [37] S. L. Ullo *et al.*, "A new mask R-CNN-based method for improved landslide detection," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 3799–3810, 2021, doi: [10.1109/JSTARS.2021.3064981](https://doi.org/10.1109/JSTARS.2021.3064981).
- [38] T.-Y. Lin *et al.*, "Microsoft coco: Common objects in context," *Comput. Vis. ECCV*. Fleet D., Pajdla T., Schiele B., Tuytelaars T. Eds., Lecture Notes in Computer Science, Springer, Cham, vol. 8693. [Online]. Available: https://doi.org/10.1007/978-3-319-10602-1_48
- [39] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *J. Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [40] R. K. Hayward *et al.*, "Mars global digital dune database and initial science results," *J. Geophys. Res., Planets*, vol. 112, no. E11, 2007.
- [41] R. K. Hayward, L. Fenton, T. Titus, A. Colaprete, and P. Christensen, "Mars global digital dune database: Mc-30," *US Geological Survey Open-File Report*, vol. 1259, 2012.
- [42] S. J. Robbins and B. M. Hynek, "A new global database of Mars impact craters > 1 km: 2. global crater properties and regional variations of the simple-to-complex transition diameter," *J. Geophys. Res., Planets*, vol. 117, no. E6, 2012.
- [43] S. Piqueux *et al.*, "Widespread shallow water ice on Mars at high latitude sand mid latitudes," *Geophys. Res. Lett.*, vol. 46, no. 24, pp. 14290–14298, 2019.
- [44] R. Wasson and R. Hyde, "Factors determining desert dune type," *Nature*, vol. 304, no. 5924, pp. 337–339, 1983.
- [45] T. Baird, C. S. Bristow, and P. Vermeesch, "Measuring sand dune migration rates with COSI-Corr and landsat: Opportunities and challenges," *Remote Sens.*, vol. 11, no. 20, 2019, Art. no. 2423.
- [46] D. Tirsch, R. Jaumann, A. Pacifici, and F. Poulet, "Dark aeolian sediments in Martian craters: Composition and sources," *J. Geophys. Res., Planets*, vol. 116, no. E3, 2011.
- [47] B. Ehlmann *et al.*, "Chemistry, mineralogy, and grain properties at Namib and High, danes, Bagnold dune field, Gale crater, Mars: A synthesis of Curiosity rover observations," *J. Geophys. Res., Planets*, vol. 122, no. 12, pp. 2510–2543, 2017.
- [48] M. G. Lapotre *et al.*, "Compositional variations in sands of the Bagnold Dunes, Gale Crater, Mars, from visible-shortwave infrared spectroscopy and comparison with ground truth from the curiosity rover," *J. Geophys. Res., Planets*, vol. 122, no. 12, pp. 2489–2509, 2017.
- [49] L. Rubanenko, "Trained weights for mask rcnn to detect barchan dunes on Mars," 2021. [Online]. Available: <https://dx.doi.org/10.21227/mbrx-hp61>




Lior Rubanenko received the B.Sc. degree in geophysics and planetary sciences from Tel-Aviv University, Tel Aviv, Israel, in 2012 and the M.Sc. degree in planetary sciences from the Weizmann Institute of Science, Rehovot, Israel, in 2015, and the Ph.D. degree in geophysics and space physics from the University of California, Los Angeles, CA, USA, in 2020.

He is currently a Postdoctoral Scholar with the Department of Geological Sciences in Stanford University, focusing on applying machine learning and computer vision to study the surfaces of Earth, Mars, the Moon and Mercury. His research focuses on understanding how planetary surfaces form and evolve through mechanical, cryogenic, and atmospheric interactions.




Sebastian Pérez-López is currently working toward the undergraduate degree (fourth-year) in majoring in geological sciences in geological sciences with the Stanford University, Stanford, CA, USA.

His current research interests lie in the fields of planetary geology, surface processes, and astrobiology.



Joseph Schull is currently working toward the undergraduate degree (first year) in majoring in cognitive science and data science with the University of California, Berkeley, CA, USA.

He intends to double major in Cognitive Science and Data Science. His interests include machine learning, neurobiology, evolutionary biology, and astronomy.



Mathieu G. A. Lapôte received the B.S. degree in geophysics and the M.S. degree in environmental sciences and engineering from the Université de Strasbourg, Strasbourg, France, in 2009 and 2011, respectively, the M.S. degree in geophysical engineering from the Ecole et Observatoire des Sciences de la Terre, Strasbourg, France, in 2011, and the M.S. degree in planetary science and the Ph.D. degree in geological sciences from the California Institute of Technology, Pasadena, CA, USA, in 2014 and 2017, respectively.

From 2017 to 2019, he was a John Harvard Distinguished Science Fellow with Harvard University. He is currently an Assistant Professor of geological sciences and, by courtesy, of geophysics with Stanford University, Stanford, CA, USA. His research focuses on the mechanics of sedimentary and geomorphic processes that shape planetary surfaces (including Earth's) and aims to untangle what landforms and rocks tell us about past hydrology, climate, and habitability of planets.