

Mars Image Content Classification: Three Years of NASA Deployment and Recent Advances

Kiri Wagstaff¹, Steven Lu¹, Emily Dunkel¹, Kevin Grimes¹, Brandon Zhao², Jesse Cai³,
Shoshanna B. Cole⁴, Gary Doran¹, Raymond Francis¹, Jake Lee¹, and Lukas Mandrake¹

¹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109-8099, USA

² Duke University, 2138 Campus Drive, Durham, NC 27708, USA

³ California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA

⁴ Space Science Institute, 4765 Walnut St, Suite B, Boulder, CO 80301, USA

{kiri.l.wagstaff,you.lu,emily.dunkel,kevin.m.grimes}@jpl.nasa.gov, brandon.zhao@duke.edu, jycai@caltech.edu,
scole@spacescience.org, {gary.b.doran.jr,raymond.francis,jake.h.lee,lukas.mandrake}@jpl.nasa.gov

Abstract

The NASA Planetary Data System hosts millions of images acquired from the planet Mars. To help users quickly find images of interest, we have developed and deployed content-based classification and search capabilities for Mars orbital and surface images. The deployed systems are publicly accessible using the PDS Image Atlas. We describe the process of training, evaluating, calibrating, and deploying updates to two CNN classifiers for images collected by Mars missions. We also report on three years of deployment including usage statistics, lessons learned, and plans for the future.

Introduction

The NASA Planetary Data System (PDS) maintains archives of data collected by NASA missions that explore our solar system. The PDS Cartography and Imaging Sciences Node (Imaging Node) provides access to millions of images of planets, moons, comets, and other bodies. Given the large and continually growing volume of data, there is a need for tools that enable users to quickly search for images of interest. Each image product is described by a rich set of searchable metadata properties such as the time it was collected, the instrument used, the image target, local season, etc.

However, users often wish to search on the *content* of the image to zero in on those images most relevant to a scientific investigation or individual curiosity. Manually searching through millions of images is infeasible. In previous work, we trained image classifiers to detect classes of interest in Mars orbital and surface images (Wagstaff et al. 2018). Using the predictions made by these classifiers, users can interactively search for classes of interest using the PDS Image Atlas¹. Since the deployment of these classifiers in late 2016 and through August 2020, their predictions have been used to satisfy 62,613 searches on the Atlas website.

In this paper, we report on several new advances within this domain. First, we expanded the set of classes known to each classifier to broaden their coverage of different content types. Second, we employed classifier calibration to

produce more reliable posterior probabilities, which is vital since only classifications with a posterior probability of at least 0.9 are displayed to users. Finally, we now report on three years of deployment including usage statistics, lessons learned, and plans for the future.

Related Work

Machine learning image classification has achieved high levels of performance since the adoption of convolutional neural networks (CNNs) trained on millions of images (Krizhevsky, Sutskever, and Hinton 2012). In addition to demonstrated improvements in accuracy, the use of a CNN removes the need for manual feature engineering. The ability to adapt or “fine-tune” large networks enables the re-use of learned lower levels of the network on new image collections while customizing the output nodes to the classes of interest. Palafox et al. (2017) showed that a CNN out-performed a support vector machine classifier on finding Mars landforms of interest. In previous work, we demonstrated the ability to fine-tune the AlexNet classifier for application to images collected by instruments in Mars orbit and on the Mars surface (Wagstaff et al. 2018). Other approaches with relevance for planetary exploration are terrain classification of regions within an image to inform navigation (Rothrock et al. 2016) and generating text captions for planetary images and enable a larger search vocabulary (Qiu et al. 2020), as opposed to a fixed set of image classes.

New Mars Classifier Data Sets

We created two new labeled data sets to train and evaluate the latest versions of our Mars image classifiers. The HiRISE images were collected by the High Resolution Imaging Experiment (HiRISE) instrument onboard the Mars Reconnaissance Orbiter (MRO) (McEwen et al. 2007), while the MSL images were collected by the Mast Camera (Mastcam) and Mars Hand Lens Imager (MAHLI) instruments mounted on Mars Science Laboratory (MSL) Curiosity rover (Grotzinger et al. 2012). To ensure high quality, the labels for both data sets were acquired using crowdsourcing with local volunteers who received specific training for each data set.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://pds-imaging.jpl.nasa.gov/search/>

Class Name	Count	Percent
Bright dune	250	2.31%
Crater	794	7.34%
Dark dune	166	1.53%
Impact ejecta	74	0.68%
Other	8,802	81.39%
Slope streak	267	2.47%
Spider	164	1.52%
Swiss cheese	298	2.76%
Total	10,815	100%

Table 1: HiRISE (Mars orbital) data set class distribution.

HiRISE Orbital Data Set (v3)

In previous work (Wagstaff et al. 2018), we compiled 3,820 images of Mars surface features that covered five classes of interest. The new HiRISE data set (v3) increases the number of labeled images to 10,815 (before augmentation), with eight classes (Doran et al. 2020)².

HiRISE images consist of long strips that cover up to 60 km with a 6-km wide swath at a resolution of 30 centimeters/pixel. To identify surface features of interest, we employed a focus of attention mechanism known as *dynamic landmarking* (Wagstaff et al. 2012). This process scans through a large image to identify visually salient regions, which are termed *landmarks*. The salience of each pixel is defined as a linear combination of the response to a Canny filter and the Earth mover’s distance (Rubner, Tomasi, and Guibas 1998) between the distribution of pixel intensity values within a window around the pixel compared to values within a larger enclosing window. We employed a genetic algorithm to optimize the parameters (analysis window sizes, weighting of individual filters, and salience threshold) based on fourteen HiRISE images with hand-labeled salient regions. The salient landmarks within an image were obtained by identifying connected components of regions that exceed the salience threshold. We cropped the salient landmarks from the “browse” (reduced resolution) version of each HiRISE image using a square bounding box plus a 30-pixel border, then resized each image to 227×227 pixels.

The resulting HiRISE image data set contains 10,815 landmark images derived from 232 separate HiRISE source images. The class distribution is shown in Table 1 in alphabetical order. The classes are highly imbalanced, with the majority of images classified as “Other” and “Impact ejecta” the least common class.

Examples of each class are shown in Figure 1. “Bright dune”, “Crater”, “Dark dune”, “Other”, and “Slope streak” classes were included in the v1 HiRISE data set. “Bright dune” and “Dark dune” are two sand dune classes found on Mars. Dark dunes are completely defrosted, whereas bright dunes are not. The “Crater” class consists of crater images in which the diameter of the crater is greater than or equal to 1/5 the width of the image and the circular rim is visible for at least half the crater’s circumference. The “Slope streak” class consists of images of dark flow-like features on

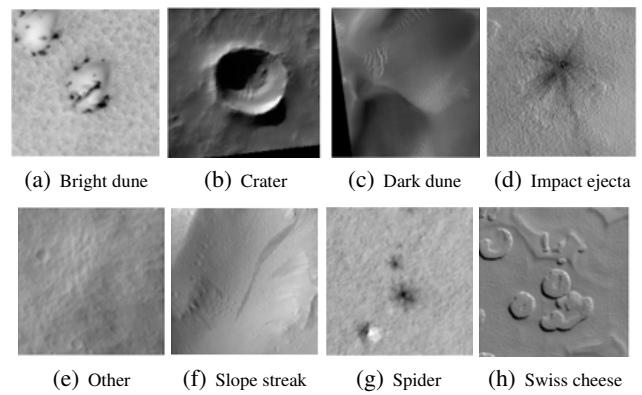


Figure 1: Examples of each class in the HiRISE v3 data set.

slopes. These features are believed to be formed by a dry process in which overlying (bright) dust slides down a slope and reveals a darker sub-surface. “Other” is a catch-all class that contains images that fit none of the defined classes of interest (e.g., Figure 1(e)).

We introduce three new classes of interest, which are “Impact ejecta”, “Spider”, and “Swiss cheese”. “Impact ejecta” refers to evidence of a meteorite impact on the surface. “Spiders” and “Swiss cheese” are phenomena that occur in the south polar region of Mars. Spiders have a central pit with radial troughs, and they are believed to form as a result of seasonal jets expelling carbon dioxide gas through an overlying ice layer (Aye et al. 2019). “Swiss cheese” is terrain that consists of pits that are formed when the sun heats the ice making it sublime (change solid to gas).

We used a combination of labeling platforms to label the HiRISE landmark images. Early images were labeled by in-house volunteers using the Zooniverse.org platform. We conducted a second labeling campaign that targeted three minority classes: Impact ejecta, Spiders, and Swiss cheese. Landmark images from this campaign were labeled using the Interactive Data Analyzer and Reviewer (IDAR) browser-based image labeling tool³. We obtained labels for each image from three volunteers. Images for which the three labels did not agree (for the second campaign, this amounted to approximately 30% of the images) were manually reviewed to select the final label. To guide labeling when more than one class was present in the image, we instructed volunteers to prioritize classes as Impact ejecta, Slope streak, Spider, Dark dune, Bright dune, Swiss cheese, Crater, or Other.

MSL Surface Data Set (v2)

We created a new data set of Mars surface images collected by the Mastcam and MAHLI instruments on the MSL Curiosity rover. Mastcam is a two-instrument suite with left- and right-eye cameras. MAHLI is a single focusable camera located on the turret at the end of the rover’s robotic arm. In our previous work (Wagstaff et al. 2018), we created a

²<https://doi.org/10.5281/zenodo.4002935>

³<https://github.com/stevenlujpl/IDAR>

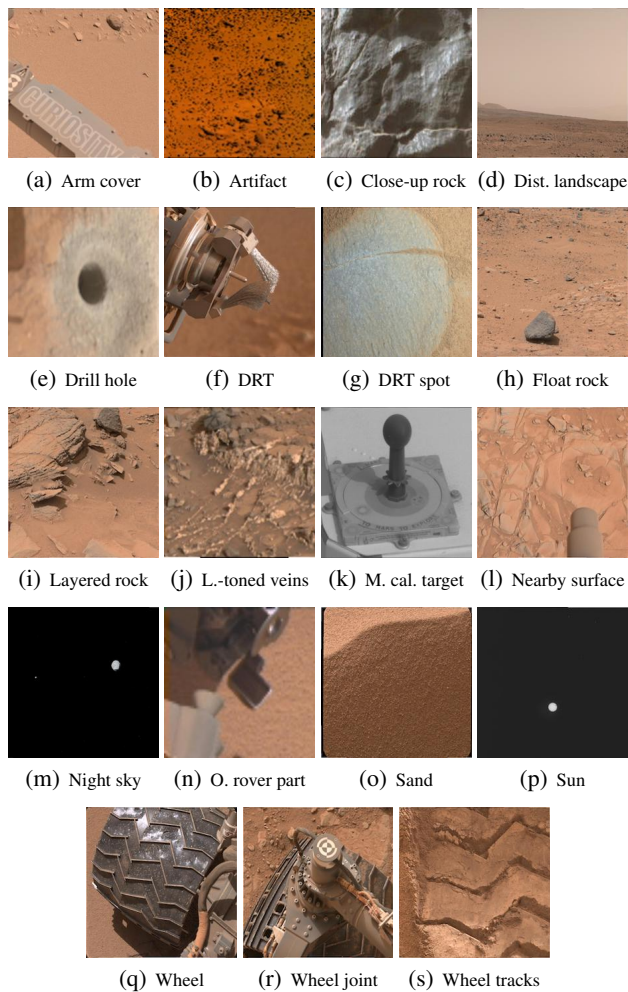


Figure 2: Examples of each class in the MSL surface data set. Subfigure (d) contains the circular shape of the sun, while subfigure (e) is an irregularly shaped moon of Mars.

data set of 6,691 images spanning 24 classes that primarily focused on rover hardware parts. The new data set (v2) includes 2,900 images spanning 19 classes that primarily focus on objects of scientific interest (Lu and Wagstaff 2020)⁴.

The MSL data set consists of RGB and grayscale images that are 8-bit, decompressed, radiometrically calibrated, color corrected, and geometrically linearized browse images from the MSL mission archive hosted by the PDS Imaging Node. We resized the smallest side to 227 pixels while preserving aspect ratio and center-cropped the other side to 227 pixels. We randomly sampled a total of 2,900 images from sol (MSL mission day) 1 to 2,224 composed of 1,172 Mastcam left eye camera, 1,000 Mastcam right eye camera, and 728 MAHLI images.

Our first task was to define the set of MSL classes of interest, which was not known in advance. We analyzed a subset of 1,600 images covering the full sol range with the browser-

⁴<https://doi.org/10.5281/zenodo.4033453>

Class Name	Count	Percent
Arm cover	10	0.34%
Artifact	408	14.07%
Close-up rock	373	12.86%
Distant landscape	197	6.79%
Drill hole	65	2.24%
DRT	14	0.48%
DRT spot	47	1.62%
Float rock	80	2.76%
Layered rock	105	3.62%
Light-toned veins	42	1.45%
Mastcam calibration target	100	3.45%
Nearby surface	1,008	34.76%
Night sky	23	0.79%
Other rover part	86	2.97%
Sand	123	4.24%
Sun	115	3.97%
Wheel	56	1.93%
Wheel joint	33	1.14%
Wheel tracks	15	0.52%
Total	2,900	100%

Table 2: MSL (Mars surface) data set class distribution.

based Class Discovery Tool from the IDAR software suite. This tool allows users to interactively associate images with dynamically created categories as they are discovered. We started with an initial list of classes from a domain expert on the MSL science team and pre-sorted the images using the DEMUD algorithm (Wagstaff et al. 2013) so that the most “interesting” or unusual images were displayed first. The DEMUD algorithm is efficient in terms of class discovery (Wagstaff and Lu 2020). The class discovery process yielded 19 classes of interest.

Examples of each class are shown in Figure 2 in alphabetical order. They include three classes that describe rover-created features (“Drill hole”, “DRT (Dust Removal Tool) spot”, and “Wheel tracks”), “Sun” and “Night sky”, seven Mars surface feature classes (e.g., “Light-toned veins”, “Layered rock”, “Float rock”), five rover part classes (e.g., “DRT”, “Mastcam calibration target”, “Wheel” and a generic “Other rover part” class), and “Artifact” used for images that are low in quality or contain missing data. The pixel resolutions and lighting conditions in these images vary a lot as they were imaged at different distances and different times.

We used the IDAR labeling tool to label MSL surface data set images. We divided the 2,900 images into 29 batches, and each batch of 100 images were distributed to three volunteers for labeling. We provided detailed labeling instructions with definitions of each class and prioritization guidance when multiple classes appeared in a single image. As with the HiRISE data set, we resolved disagreement using expert review. The MSL surface data set class distribution is shown in Table 2.

Instrument	Train sol 1 - 948		Val. sol 949 - 1920		Test sol 1921 - 2224		Full Archive sol 1 - 2224	
	Count	Percent	Count	Percent	Count	Percent	Count	Percent
Mastcam Left	842	42.1%	108	36.0%	222	37.0%	50,480	40.0%
Mastcam Right	678	33.9%	107	35.7%	215	35.8%	43,041	34.1%
MAHLI	480	24.0%	85	28.3%	163	27.2%	32,612	25.9%
Total	2,000	100%	300	100%	600	100%	126,133	100%

Table 3: MSL surface data set training, validation, and test data sets by instrument and sol range.

CNN Classification for Mars Images

We trained and deployed two Convolutional Neural Network (CNN) classifiers, denoted as HiRISENet and MSLNet, for MRO HiRISE images and MSL Mastcam and MAHLI images. We employed transfer learning to adapt the weights of networks pre-trained on Earth images for use with Mars orbital and surface images.

HiRISENet: CNN Classifier for Mars Orbital Images

We adapted the AlexNet image classifier (Krizhevsky, Sutskever, and Hinton 2012) for use with HiRISE classes. AlexNet was trained on 1.2 million Earth images from 1000 classes in the ImageNet data set. We started with Caffe’s BVLC reference model (Jia et al. 2014), which is a copy of AlexNet that was trained for 310,000 iterations, provided by Jeff Donahue⁵. We removed the final fully connected layer, added a new layer with eight output classes, and re-trained the network with Caffe (Jia et al. 2014). We followed Caffe’s recommendations for fine-tuning, including using a small base learning rate and small step size and a larger learning rate multiplier for the final layer only. We used a learning rate of 0.0001, weight decay of 0.0005, and relatively small step size of 20,000. The initial layers were almost fixed; they used learning rate multipliers of 1 (weight) and 2 (bias). The final layer was allowed greater adaptation with multipliers of 10 (weight) and 20 (bias). We trained the model for 78,900 iterations. Caffe computes the per-band mean pixel values from the training set and uses these values to normalize all images during training and prediction.

We split the HiRISE dataset into train, validation, and test sets using each landmark’s HiRISE source image identifier to ensure no overlap in source images between the sets. We used approximately 65% of the data for training, 19% for validation, and 17% for testing. Images obtained from our second labeling campaign (to target minority classes) appear in the training and validation sets only so that we could assess improvements against an unchanged test set.

We applied data augmentation to the training and validation sets. The augmentation includes three rotations: 90, 180, and 270 degrees, horizontal and vertical flips, and a random brightness adjustment. In addition, we up-sampled data obtained in the second labeling campaign by a factor of two.

⁵https://github.com/BVLC/caffe/tree/master/models/bvlc_reference_caffenet

MSLNet: CNN Classifier for Mars Surface Images

MSLNet is a hybrid of two CNN classifiers. The version 1 (v1) classifier focused on rover hardware classes (Wagstaff et al. 2018; Lu et al. 2019). The primary objectives of the Mastcam and MAHLI instruments are to enable science analysis of rover investigation sites, which motivated the creation of version 2 (v2) classifier to expand the set of classes to include science targets (e.g. “Float rock”, “Layered rock”) and activities (e.g. “DRT spot”, “Drill hole”). The v1 classifier initially focused on engineering considerations and rover hardware classes due to requests by the MSL rover planning team as well as pre-existing availability of labels for those items. The “Wheel” class was of particular interest due to growing awareness in 2017 that the rover’s wheels were experiencing a higher than expected level of degradation due to driving on the rough surface. The success of the v1 classifier led to new requests to also accommodate science-related classes in support of MSL mission to explore and understand Mars. Observations that contain classes such as “Layered rock” and “Light-toned veins” are very high science priorities to help determine the history and evolution of water activity, which can also have implications for habitability. The deployed MSLNet classifier covers both areas of interest (engineering and science) to meet the needs of diverse users with different priorities.

MSLNet first classifies images with the v2 classifier, then reclassifies any images classified as “Other rover part” with the v1 classifier to get a fine-grained classification of rover parts. The creation and evaluation of the v1 classifier were reported in previous work (Wagstaff et al. 2018). The v2 classifier was trained and evaluated using the MSL surface data set described above. We divided this data set into training, validation, and test data sets according to their sol of acquisition to enable the evaluation of generalization performance on newly acquired images. The sol splitting boundaries, as shown in Table 3, were chosen to enable per-camera distributions that roughly match the full archive.

To improve the generalization performance of the classifier, we augmented the images in the training data set (but not the validation and test data sets). The MAHLI images (which come from a rotatable platform) were augmented using rotation (90°, 180°, and 270°) and flipping (horizontal and vertical); the Mastcam images (which come from a fixed platform) were augmented using only horizontal and vertical flipping methods.

As with HiRISENet, for the MSLNet v2 classifier we fine-tuned AlexNet for 2,050 iterations with a fixed base learning rate of 0.0001. We set the learning rate multipliers of the first four convolution layers, the fifth convolution layers,

and the final fully connected layers to 0, 1, 20 respectively. We set the dropout rate to 0.5 for the first and second fully connected layers. The final hybrid MSLNet classifier combines v1 and v2 and classifies images into 35 classes of both science and engineering relevance.

Classifier Calibration

The deployed classifiers use a confidence threshold to determine which results are shown to users, so it is vital that the models are well calibrated. Modern neural networks have achieved higher accuracies but in many cases have suffered an increase in calibration error, which means that the predicted class probabilities deviate from the true empirical probabilities. In many cases, the networks are consistently over-confident in their predictions. This effect appears to be tied to an increase in model capacity and lack of regularization (Guo et al. 2017). For a quantitative measure of model calibration, we calculate the Expected Calibration Error (ECE), which is the expected difference between posterior probability (confidence) and accuracy. We partition n predictions into M equally spaced bins and computed a population-weighted average of the difference between accuracy and confidence within each bin:

$$ECE = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|.$$

We evaluated four post-hoc calibration methods that extend Platt Scaling (Platt et al. 1999) to multiclass problems. Temperature scaling (Guo et al. 2017) uses a single parameter T to rescale model output. Given the model output for item \mathbf{x} , which is a logit vector $\mathbf{z} \in \mathcal{R}^K$, the calibrated probability of class k is: $p(y = k|x) = \frac{e^{z_k/T}}{\sum_{j=1}^K e^{z_j/T}}$. The parameter T is optimized with respect to the log likelihood on the validation set. Since the parameter T does not change the maximum of the softmax function, the accuracy of the model is unchanged. Bias-corrected temperature scaling (BCTS) (Alexandari, Kundaje, and Shrikumar 2020) adds a bias term b_k for each class: $p(y = k|x) = \frac{e^{z_k/T+b_k}}{\sum_{j=1}^K e^{z_j/T+b_j}}$. Vector and matrix scaling (Guo et al. 2017) add additional flexibility with per-class scaling using a $K \times K$ linear transformation matrix \mathbf{W} by computing $\mathbf{W}\mathbf{z} + \mathbf{b}$ and then normalizing across classes to get $p(y = k|x)$. Vector scaling constrains \mathbf{W} to be a diagonal matrix whose entries function as class-specific temperature values.

CNN Classification Evaluation

To evaluate HiRISENet and MSLNet, we used the overall (post-threshold) accuracy score and abstention rate as the primary performance metrics and ECE to measure the calibration level of the classifiers. We also analyzed the precision and recall scores to understand per-class performance.

HiRISENet Evaluation

HiRISENet classification accuracy results are shown in Table 4. Random class prediction on this data set achieves 11.1% accuracy (given eight classes). Compared to a simple baseline that predicts the most common class from the training set (“Other”), HiRISENet exhibits a strong improvement from 81.1% to 92.8% on the test set.

Classifier	Train ($n = 6997$) ($n_{aug} = 51058$)	Val ($n = 2025$) ($n_{aug} = 14959$)	Test ($n = 1793$)
Most common	78.4%	75.0%	81.1%
HiRISENet	99.6%	88.6%	92.8%

Table 4: Classification accuracy on HiRISE (Mars orbital) images. The best performance on each data set is in bold.

	ECE	Acc	0.9 confidence	
			Acc	Abst Rate
Uncalibrated	0.073	88.6	94.2	13%
Temperature scaling	0.013	88.6	97.3	31%
BCTS	0.014	89.2	97.3	29%
Vector scaling	0.010	89.3	97.2	27%
Matrix scaling	0.013	90.3	97.7	24%

Table 5: HiRISENet calibration results on validation set; best performance in bold.

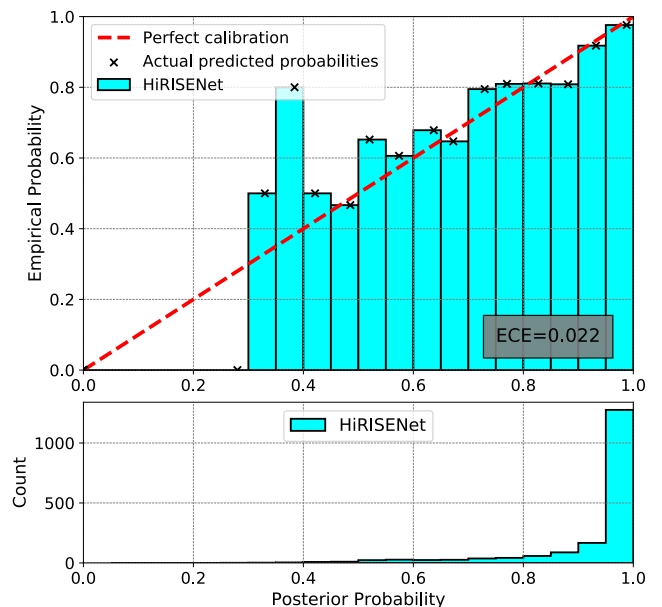


Figure 3: Calibrated HiRISENet reliability (test set).

For this application, reliable posterior probabilities are essential, since predictions are thresholded so that only those of at least 0.9 probability are shown to users. We compared four calibration methods in terms of their impact on accuracy, ECE, and abstention (Table 5). We found that matrix scaling achieved the highest validation accuracy (90.3%) as well as the lowest abstention rate (24%). Vector scaling achieved the lowest ECE but with higher abstention and lower accuracy. Therefore, we adopted matrix scaling for deployment. On the test set, the calibrated HiRISENet model achieved 96.7% accuracy with an abstention rate of 20%.

Reliability diagrams (DeGroot and Fienberg 1983; Niculescu-Mizil and Caruana 2005) provide a visual repre-

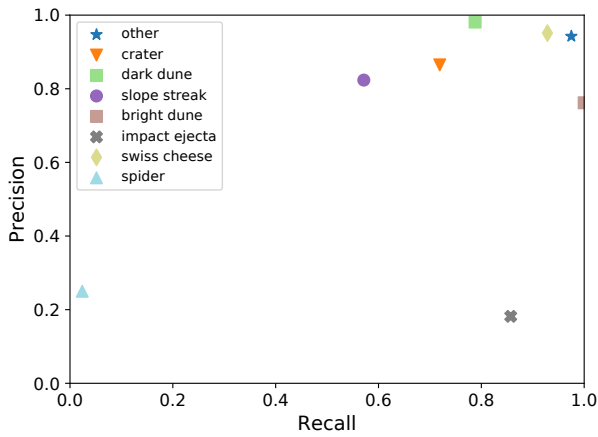


Figure 4: Calibrated HiRISENet per-class precision and recall (test set).

sensation of model calibration. The empirical per-bin accuracy is plotted as a function of model posterior probability. For a perfectly calibrated model, these values are equal, following the diagonal line. Figure 3 shows the reliability diagram for HiRISENet. The bottom panel shows the data set distribution in terms of predicted probability. We find that HiRISENet is well calibrated with an ECE of just 0.022 and the majority of predictions in the most-confident bin.

Figure 4 shows per-class precision and recall on the test set after matrix scaling calibration. Most classes achieve precision above 0.70 with recall above 0.50 (even after thresholding). The “Spider” class has the lowest recall (0.02, out of only 42 items), while the “Impact ejecta” class has the lowest precision (0.18). Figure 5 shows the confusion matrix on the test set after matrix scaling calibration. Diagonal (correct) entries have a dark background. A comparison to the confusion matrix before calibration (not shown) indicates that two images that were previously incorrectly classified into the “Spider” class are now correctly classified as the “Impact ejecta” class; however, the confusion between the “Crater” and “Impact ejecta” classes has increased. In addition, the “Spider” class suffered from significant domain shift, which is evident in Figure 6. The “Spider” images in the validation set as shown in Figure 6(a) are extremely visually different compared to the “Spider” images in the test set as shown in Figure 6(b). We found that even human labelers had trouble recognizing them as the same phenomena. Future updates to this data set will target the “Spider” class.

MSLNet Evaluation

The performance of the MSLNet v2 classifier is shown in Table 6 in comparison to the most-common-class (“Nearby surface”) baseline. The MSLNet classifier significantly outperforms the baseline method, achieving 74.5% accuracy, or 87.2% with 36% abstention using a confidence threshold of 0.9, on the test set. Recall that images in the training, validation, and test sets were divided according to their sol (date)




Figure 5: Calibrated HiRISENet confusion matrix (test set).




Figure 6: Domain shift in HiRISE “Spider” landmarks.

of acquisition. The performance of the classifier gradually decreases over time as the rover traversed to new locations, possibly due to label shift, in which the prior class probabilities change over space or time (Lipton, Wang, and Smola 2018). We plan to investigate label shift adaptation to enable the classifier to accommodate such change.

MSLNet achieves lower accuracy and higher abstention than HiRISENet on its corresponding test set. Given the larger number of classes and smaller number of labeled images, we believe that this classifier is likely even more data-limited and would benefit from additional data collection.

Reliable posterior probabilities are likewise essential for MSLNet. We calibrated the MSLNet classifier using temperature scaling, the most computationally efficient method among the four calibration methods discussed in this paper (e.g., matrix scaling scales quadratically with the number of classes, which is problematic for MSLNet). After calibration, test set accuracy using the confidence threshold improved to 90.3%, at the cost of increased abstention. For this application, we are willing to sacrifice coverage to ensure that the classifications provided to users are highly reliable. MSLNet’s ECE improved from 0.142 to 0.08 with temperature scaling. The reliability diagram of MSLNet after cali-

	Train (n=5920)			Validation (n=300)			Test (n=600)		
	Acc	Acc (0.9)	Abst Rate	Acc	Acc (0.9)	Abst Rate	Acc	Acc (0.9)	Abst Rate
Most Common	26.3%	-	-	24.7%	-	-	31.2%	-	-
MSLNet	99.6%	100%	6.5%	78.3%	89.8%	38.0%	74.5%	87.2%	36.2%
MSLNet-calibrated	99.6%	100%	18.8%	78.3%	96.5%	52.7%	74.5%	90.3%	51.8%

Table 6: Performance results for MSLNet classifiers (best for each data set is in bold). Note that Acc (0.9) in the title means accuracy score computed with a 0.9 confidence threshold.




Figure 7: Calibrated MSLNet reliability diagram (test set).

bration is shown in Figure 7.

Per-class precision and recall scores are shown in Figure 8. The green and red F1-score curves separate the classes into three groups. The first group, those above the green curve, includes ten classes (e.g. “Nearby surface”, “Mastcam calibration target”) whose F1 scores are greater than 0.6; the second (intermediate) group includes five classes (e.g. “Layered rock”, “Drill hole”) whose F1 scores are between 0.2 and 0.6; and the third group, those below the red curve, includes four classes (e.g. “Float rock”, “Wheel tracks”) whose F1 scores are less than 0.2. We note that the classes in the third group were evaluated on very few images, so their performances are not statistically robust. These classes require further improvement, and we plan to investigate up-sampling or additional data acquisition to increase the number of images of these classes.

PDS Image Atlas Deployment

HiRISENet and MSLNet generate classifications that enable PDS users to quickly find images of interest via content-based search. The public interface to the PDS image archives is the PDS Image Atlas⁶. Users can select an instrument

⁶<https://pds-imaging.jpl.nasa.gov/search/>




Figure 8: Calibrated MSLNet per-class precision and recall (test set).




Figure 9: View from the Atlas of craters found in HiRISE image PSP_002912_2075_RED.

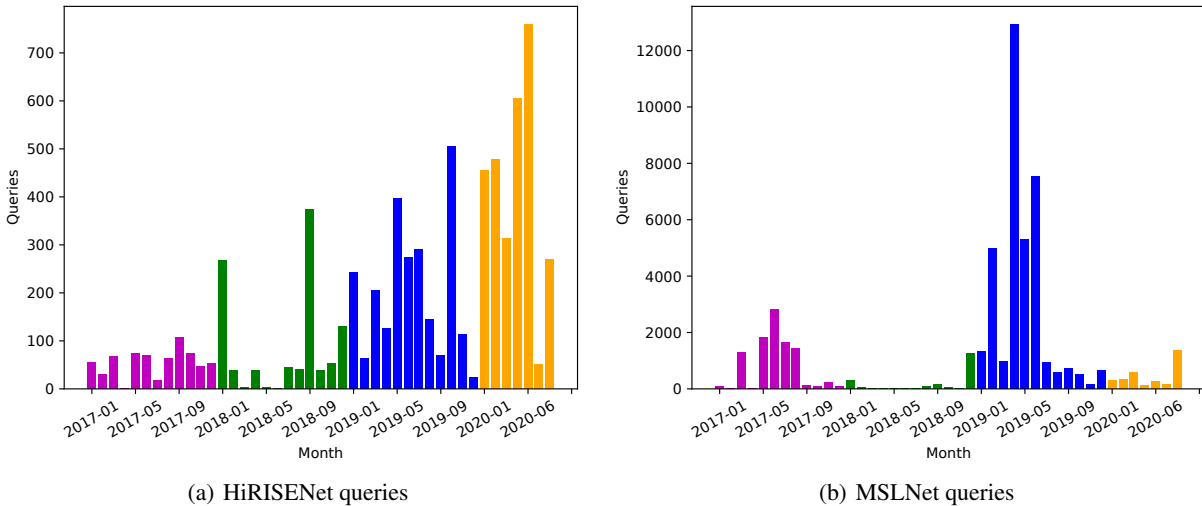


Figure 10: Number of monthly queries for HiRISENet and MSLNet classifications over 3.5 years (colors distinguish years).

(e.g., HiRISE) and filter the images to only contain a particular class of interest. To enable this kind of search, we applied HiRISENet and MSLNet to the full archive of images collected by the relevant instruments on Mars. These archives total 65,091 HiRISE and 1,057,912 MSL images, far more than the labeled subsets used for training and evaluation. Figure 9 shows the Atlas user view of a HiRISE image with all confidently classified craters enclosed in red bounding boxes. Craters that are small, faint, degraded, or distorted are less likely to pass the confidence threshold, but those identified as craters are highly reliable. In response to user requests, we added the ability to download a file that contains the latitude and longitude coordinates of each detected landmark, using PDSC⁷ to convert from pixel to geographic coordinates.

Classifying all HiRISE images took about five days on a GPU system and yielded 29,608 landmarks with a posterior probability of at least 0.9, from classes that were not “Other”. We also filtered out predictions for “Spider” or “Swiss cheese” at latitudes outside of the south polar region (Aye et al. 2019). This total represents an 81% increase over the number of classified landmarks available in the first classifier release (Wagstaff et al. 2018). MSLNet classified 136,967 images with a posterior probability of at least 0.9. Both classifiers are integrated into the data ingestion pipeline for the Atlas. As new data is delivered from HiRISE or MSL, the images are automatically processed and tagged with their predicted classes.

We track the number of Atlas queries that make use of HiRISENet and MSLNet classifications. As shown in Figure 10(a), HiRISENet exhibits regular and increasing usage over time. The most popular HiRISE class to be queried is “Crater”. MSLNet shows more varied activity (Figure 10(b)), dominated by heavy usage in early 2019 when the number of queries for “Wheel” increased by sev-

eral orders of magnitude (note the difference in y axis range). Given the small separation in query timestamps, most likely it was the result of a large number of scripted queries to the Atlas, which provides a public API. It is possible that this classifier output is serving to help train other investigators’ models.

We also analyzed the top 20 domain names from which the queries came. From January to July of 2020, we found that 40% of these queries came from hosts through an ISP, including Spectrum and Comcast as well as ISPs in the U.K., the Netherlands, Romania, and Taiwan. Another 33% of these queries came from JPL domains, which is not surprising given the relevance of the content to JPL projects. A minority of the top 20 domain queries came from the Remote Sensing Technology Center of Japan (2%), the University of Wyoming (1%), and SoftBank (Japan) (1%), while 23% of hostnames did not resolve to a domain.

Finally, we used Google Analytics to examine the global distribution of visitors who made classification queries. Between July 2017 (the oldest data available) and August 2020, there were 62,613 visitors. The top ten countries and share of visitors were: United States (51%), India (6%), United Kingdom (4%), Germany (3%), Canada (2%), France (2%), Italy (2%), Spain (2%), Australia (2%), and Russia (2%). In all, visitors came from 180 different countries.

Lessons learned. The deployment of Mars image classifiers at scale has yielded several lessons learned. First, it is worth highlighting the challenge of defining meaningful and relevant classes up front. Unlike a fixed benchmark data set, new Mars images are continually collected and new classes could arise at any time. Collaboration with domain experts is vital for ensuring the relevance of the class definitions.

Second, domain shift is an important factor in both data sets. Figure 11 compares the class distribution (excluding “Other”) for the labeled HiRISE data set (brown) to the predictions made across the full HiRISE archive (orange). While the “Crater” class is dominant in both, its share of the

⁷<https://github.com/JPLMLIA/pdsc>

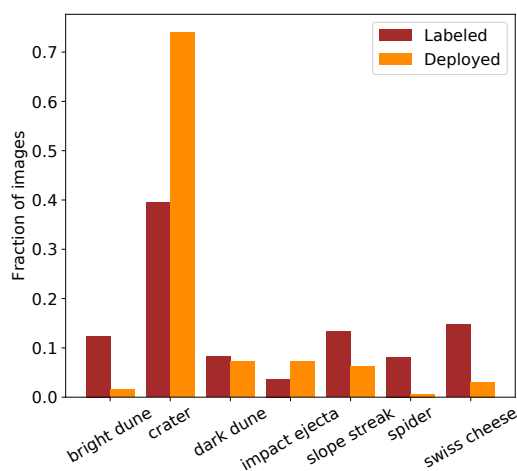


Figure 11: Class distribution for HiRISE labeled data set versus predictions on the full archive.

images classified nearly doubles when deployed. There are likely two factors involved: our labeled data set is not fully representative of all of Mars, and “Crater” predictions may in general be more confident and thus more likely to pass the 0.9 threshold and be retained here. Similar effects are seen in the MSL data set. We are currently investigating the use of label shift adaptation increase the robustness of both classifiers.

Finally, we found that our initial simplifying assumption of one class per image is sometimes inadequate (a crater might contain a dark slope streak; an MSL image might contain rover parts and the horizon). Even with guidance on class priorities, human labelers sometimes found it difficult to select a single class label. We plan to adopt a multi-label approach in future versions to allow more flexibility and reduce label noise.

Conclusions and Future Work

This paper presents the latest updates to the deployment of machine learning image classifiers in support of planetary science. Two classifiers, for orbital and surface images of Mars, have been operating since late 2016 to enable the first content-based search of large NASA image archives. Usage has increased over time, and feedback from users as well as internal assessments have guided recent improvements. These include acquiring additional training data to improve minority class performance for the HiRISE classifier, defining new classes of scientific interest for the MSL classifier, employing calibration to increase classifier reliability, and making landmark coordinates downloadable. To increase our understanding of how and why users employ the machine learning classifications of Mars images, we are investigating minimally intrusive ways to learn more about user motivations and use cases. Meanwhile, the performance and scope of these classifiers continues to grow. Each time new images are delivered by the instruments at Mars, they are automatically classified and added to the archive.

We are currently developing a new classifier, MERNet, that will operate on images collected by the two Panoramic Camera (Pancam) instruments on the Opportunity and Spirit Mars Exploration Rovers (MER) rovers. Based on our lessons learned, we are employing a multi-label approach so multiple labels can be assigned to a single image. MERNet will classify all Pancam images using classes of both science and engineering interest that were identified in a survey conducted by the MER Data Catalog project (Cole et al. 2020). MERNet will provide users with the first content-based search capability for MER images.

Finally, we plan to incorporate label shift adaptation (Alexandari, Kundaje, and Shrikumar 2020) into future upgrades of the Mars image classifiers. It is evident that the data collected by Mars instruments is not i.i.d.; class frequencies change as spacecraft study different locations on Mars, globally from orbit or locally via rover traverse. By enabling the classifiers to adapt to class probability changes, we expect to obtain more reliable classifications.

Acknowledgments

We thank Michael McAuley from the PDS Imaging Node for the continuing support of this work and Anil Natha for assistance with the Google Analytics results. We also thank the numerous volunteers who helped label the Mars images. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. This publication uses data generated via the Zooniverse.org platform, development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation.

References

- Alexandari, A.; Kundaje, A.; and Shrikumar, A. 2020. Maximum likelihood with bias-corrected calibration is hard-to-beat at label shift adaptation. In *Proceedings of the 2020 International Conference on Machine Learning*, 222–232.
- Aye, K.-M.; Schwamb, M. E.; Portyankina, G.; Hansen, C. J.; McMaster, A.; Miller, G. R. M.; Carstensen, B.; Snyder, C.; Parrish, M.; Lynn, S.; Maic, C.; Miller, D.; Simpson, R. J.; and Smith, A. M. 2019. Planet Four: Probing spring-time winds on Mars by mapping the southern polar CO₂ jet deposits. *Icarus* 319: 558–598.
- Cole, S. B.; Aubele, J. C.; Cohen, B. A.; Milkovich, S. M.; and Shields, S. R. 2020. Identifying community needs for a Mars Exploration Rovers (MER) Data Catalog. In *Proceedings of the 51st Lunar and Planetary Science Conference*, Abstract #1709.
- DeGroot, M. H.; and Fienberg, S. E. 1983. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)* 32(1-2): 12–22.
- Doran, G.; Dunkel, E.; Lu, S.; and Wagstaff, K. 2020. Mars orbital image (HiRISE) labeled data set, version 3.2.0. URL <http://doi.org/10.5281/zenodo.4002935>.
- Grotzinger, J. P.; Crisp, J.; Vasavada, A. R.; Anderson, R. C.; Baker, C. J.; Barry, R.; Blake, D. F.; Conrad, P.; Edgett,

- K. S.; Ferdowski, B.; Gellert, R.; Gilbert, J. B.; Golombek, M.; Gómez-Elvira, J.; Hassler, D. M.; Jandura, L.; Litvak, M.; Mahaffy, P.; Maki, J.; Meyer, M.; Malin, M. C.; Mitrofanov, I.; Simmonds, J. J.; Vaniman, D.; Welch, R. V.; and Wiens, R. C. 2012. Mars Science Laboratory mission and science investigation. *Space Science Reviews* 170: 5–56.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, 1321–1330.
- Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional Architecture for Fast Feature Embedding. *arXiv preprint arXiv:1408.5093*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* 25, 1097–1105.
- Lipton, Z. C.; Wang, Y.-X.; and Smola, A. 2018. Detecting and correcting for label shift with black box predictors. In *Proceedings of the 2018 International Conference on Machine Learning*, 3128–3136.
- Lu, S.; and Wagstaff, K. L. 2020. MSL Curiosity rover images with science and engineering classes, version 2.1.0. URL <http://doi.org/10.5281/zenodo.4033453>.
- Lu, S.; Wagstaff, K. L.; Cai, J.; Doran, G.; Grimes, K.; Lee, J.; Mandrake, L.; and Yue., Y. 2019. Improved content-based image classifiers for the PDS Image Atlas. In *4th Planetary Data Workshop*.
- McEwen, A. S.; Eliason, E. M.; Bergstrom, J. W.; Bridges, N. T.; Hansen, C. J.; Delamere, W. A.; Grant, J. A.; Gulick, V. A.; Herkenhoff, K. E.; Keszthelyi, L.; Kirkand, R. L.; Mellon, M. T.; Squyres, S. W.; Thomas, N.; and Weitz, C. M. 2007. Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment (HiRISE). *Journal of Geophysical Research (Planets)* 112(E5). doi:10.1029/2005JE002605.
- Niculescu-Mizil, A.; and Caruana, R. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine learning*, 625–632.
- Palafox, L. F.; Hamilton, C. W.; Scheidt, S. P.; and Alvarez, A. M. 2017. Automated detection of geological landforms on Mars using Convolutional Neural Networks. *Computers & Geosciences* 101: 48–56.
- Platt, J.; et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers* 10(3): 61–74.
- Qiu, D.; Rothrock, B.; Islam, T.; Didier, A. K.; Sun, V. Z.; A.Mattmann, C.; and Ono, M. 2020. SCOTI: Science Captioning of Terrain Images for data prioritization and local image search. *Planetary and Space Science* 188.
- Rothrock, B.; Kennedy, R.; Cunningham, C.; Papon, J.; Heverly, M.; and Ono, M. 2016. SPOC: Deep learning-based terrain classification for Mars rover missions. In *Proceedings of the AIAA SPACE Forum*.
- Rubner, Y.; Tomasi, C.; and Guibas, L. J. 1998. A metric for distributions with applications to image databases. In *Proceedings of the Sixth International Conference on Computer Vision*, 59–66. doi:10.1109/ICCV.1998.710701.
- Wagstaff, K. L.; Lanza, N. L.; Thompson, D. R.; Dietterich, T. G.; and Gilmore, M. S. 2013. Guiding scientific discovery with explanations using DEMUD. In *Proceedings of the Twenty-Seventh Conference on Artificial Intelligence*, 905–911.
- Wagstaff, K. L.; and Lu, S. 2020. Efficient active learning for new domains. In *Proceedings of the 2020 ICML on Real World Experiment Design and Active Learning*.
- Wagstaff, K. L.; Lu, Y.; Stanboli, A.; Grimes, K.; Gowda, T.; and Padams, J. 2018. Deep Mars: CNN classification of Mars imagery for the PDS Imaging Atlas. In *Proceedings of the Thirtieth Annual Conference on Innovative Applications of Artificial Intelligence*, 7867–7872.
- Wagstaff, K. L.; Panetta, J.; Ansar, A.; Greeley, R.; Hoffer, M. P.; Bunte, M.; and Schorghofer, N. 2012. Dynamic land-marking for surface feature identification and change detection. *ACM Transactions on Intelligent Systems and Technology* 3(3). Article number 49.