



Contents lists available at ScienceDirect

Planetary and Space Science

journal homepage: www.elsevier.com/locate/pss

SCOTI: Science Captioning of Terrain Images for data prioritization and local image search[☆]



Dicong Qiu^{a,*}, Brandon Rothrock^b, Tanvir Islam^b, Annie K. Didier^b, Vivian Z. Sun^b,
Chris A. Mattmann^b, Masahiro Ono^{b,**}

^a Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213, USA

^b Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109, USA

ARTICLE INFO

Keywords:

Planetary exploration
Data prioritization
Image search
Image captioning
Machine learning

ABSTRACT

Planetary exploration is full of challenges. Data bandwidth is very limited between planetary rovers and ground-based data system. What's worse, even though NASA has accumulated over 34 million images from various missions, it requires significant effort and is hardly possible for any scientist to go through all of them. In order to improve the degree of automation and the efficiency of these processes, we propose a system leveraging machine learning for planetary rovers to actively look for scientifically interesting and valuable features according to text instructions from scientists and prioritize the images captured onboard with those features for downlink. Such an image prioritization mechanism can also be naturally applied to content-based image search through text description in any local planetary image data server, allowing scientists to search for images with desired features without going through them one by one. Besides theoretical and engineering details of our proposed approach, we also present both quantitative and qualitative evaluation of the system along with some concrete examples.

1. Introduction

Over the last few decades, NASA has acquired an enormous amount of image data. More than 34 million images are made available to the public through NASA's Planetary Data System (PDS), of which over 25 million are from various Mars missions, including Spirit (MER-A), Opportunity (MER-B) and Curiosity (MSL). The number continues to grow due to ongoing missions, where more than 200 K images have been accumulated from mission InSight alone in less than a year. And much more image data are expected to be produced from incoming and future missions, such as Mars 2020. Two crucial questions then arise: (1) How can we prioritize the images collected onboard to downlink due to the limited bandwidth? (2) How can we search our local database efficiently for images with desired geologic and/or non-geologic features?

These two questions are nontrivial to answer, since there is no straightforward rule-based logic that can deal with images conveniently and efficiently. With the latest advance in deep learning (Deng et al., 2009; Krizhevsky et al., 2012; LeCun et al., 2015), promising methods were proposed for problems involving planetary image data processing

(Thompson et al., 2012; Rothrock et al., 2016; Wagstaff et al., 2018). However, to the best of our knowledge, research has seldom been done on helping machines "understand" multiple objects and their relationships in visual data (e.g. images) in planetary exploration domain.

In this paper, we propose an approach to tackle these problems by leveraging an LSTM-based image captioning neural network architecture with visual attention mechanism, which pays attention on different parts of an image at different times to encode them into embeddings and then translates these embeddings into a meaningful sequence of words (a caption) for the image. Relations of different parts in the image are captured by the ordering of attentions, and are recovered into relation/connection words in its caption. Remotely, images captured by a planetary rover can be captioned in such a fashion and prioritized in accord to the text similarity between their captions and scientist inputs (text descriptions of desired features to look for). Locally, present image data (e.g. the ones in NASA PDS) can also be captioned, after which users can search these image data through text and results are returned in decreasing order of text similarity. An implementation of text similarity metric is also proposed in our work. Moreover, besides validating the

[☆] This work was carried out at NASA's Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109, U.S.A.

* Corresponding author.

** Corresponding author.

E-mail addresses: david@davidqiu.com (D. Qiu), masahiro.ono@jpl.nasa.gov (M. Ono).

<https://doi.org/10.1016/j.pss.2020.104943>

Received 4 April 2019; Received in revised form 13 April 2020; Accepted 15 April 2020

Available online 30 April 2020

0032-0633/© 2020 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

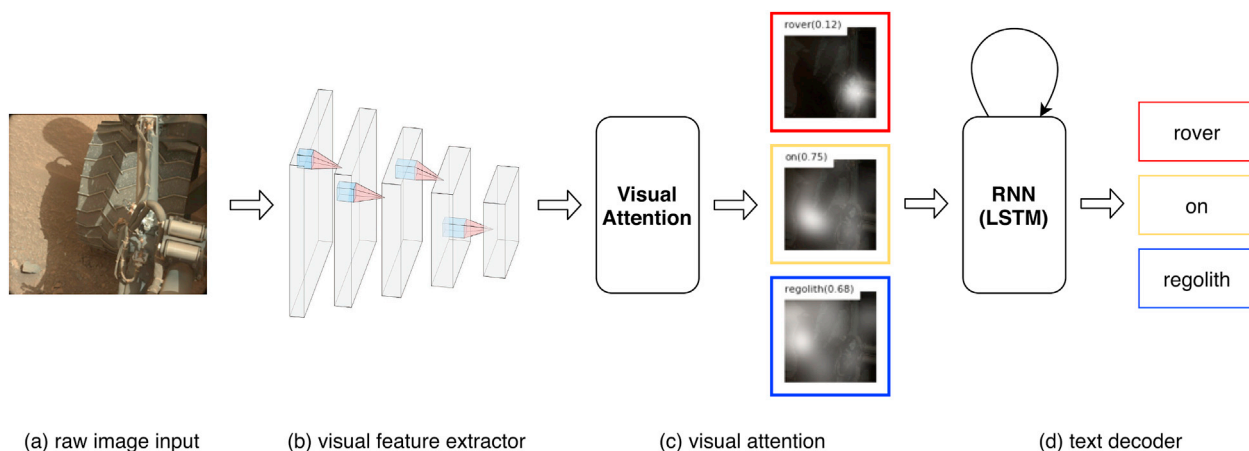


Fig. 1. An illustration on the architecture of SCOTI.

image captioning performance of the architecture, we also developed an internal web tool to simulate remote image downlink tasks with text descriptions for desired features from scientists, and integrated a beta version of text-based search tool for local images onto NASA PDS to encourage public test (Ono et al., 2019).

2. Related work

Researchers have been working for many years on automating the data analysis process for planetary data.

Autonomous image data processing through feature engineering and rule-based methods (Castano et al., 2003, 2007a) have been studied and applied on prior missions through the integration into the Onboard Autonomous Science Investigation System (OASIS) (Castano et al., 2007b). These methods, however, depend heavily on the bank of carefully selected features and hand-crafted rules, which limit their scalability beyond images of rocks, clouds and dust devils. A method that combines both feature engineering and learning for image classification (Woods et al., 2011) has also been proposed. But since it only learns how to combine predefined descriptors (e.g. SURF and SIFT) rather than learning those low-level feature descriptors (filters) from scratch, the bank of descriptors constrain its capability of identifying other meaningful features from an image.

Recent advance in machine learning and deep learning introduced a variety of new approaches in processing planetary data, especially visual data. Deep Mars (Wagstaff et al., 2018) takes advantage of recent development in convolutional neural networks (CNNs) to classify engineering-focused rover images (e.g. those of rover wheels, drill holes, etc.) and orbital images. However, since Deep Mars adopts the architecture of AlexNet (Krizhevsky et al., 2012), it can only recognize one single object in an image. Pixel-wise segmentation approaches were explored by researchers. TextureCam (Thompson et al. Wagstaff) leverages random forests to detect and classify rocks onboard. The Soil Property and Object Classification (SPOC) (Rothrock et al., 2016) segments Mars terrains in an image by utilizing a fully-convolutional neural network (FCNN). Although these methods are capable of segmenting and classifying multiple objects (regions) in an image, they require training datasets with pixel-wise segmentation annotations that are expensive to collect, and fall short of understanding the relationships among different objects in a scene.

Extracted features and predicted labels from the methods mentioned above can be used in both data prioritization and local image search (Castano et al., 2003, 2007a). prioritizes images to downlink through key target signature identification, and (Woods et al., 2011) determines data priority by feature-based image classification results. Again, these approaches rely heavily on manual feature engineering, require users to input unintuitive feature vectors to specify high-priority targets and are

constrained to limited object categories (Wagstaff et al., 2018). searches engineering-focused images through predicted labels and was integrated into NASA PDS. But it is incapable of identifying multiple objects or their relationships in an image.

In this paper, we describe our work that leverages recently developed end-to-end image captioning methods (Kiros et al., 2014; Vinyals et al., 2015) for a machine to “understand” planetary images. These methods are advantageous because they can identify multiple objects along with their relationships in an image, and translate the visual information into human understandable text descriptions. The problems of data prioritization and local image search are then converted to the problem of finding similar text descriptions, which has been researched on for a long time (Gomaa and Fahmy, 2013; Han, 2018). More concretely, the image captioning architecture we use is based on Long Short-term Memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) along with attention mechanisms (Bahdanau et al., 2015; Xu et al., 2015), and we adapt the Bilingual Evaluation Understudy (BLEU) metric to evaluate similarity between generated image captions and user inputs.

3. Science Captioning of Terrain Images

This section describes the Science Captioning of Terrain Images (SCOTI) network for tackling the problem of “understanding” planetary images (primarily terrain images) for a machine. SCOTI extracts visual features from a raw image input into a feature map (a multidimensional vector), repeatedly pays attention to different parts of the image (represented by the feature map) and generates a caption (a text description) for the image word by word.

3.1. Image captioning networks

The objective of an image captioning network is to learn an optimal parameters θ^* from a given training dataset D that maximizes the (logarithmic) probability of occurrence of all the image-caption pairs in the dataset.

$$\theta^* = \operatorname{argmax}_{\theta} \sum_{(I,y) \in D} \log p(y|I; \theta) = \operatorname{argmax}_{\theta} \sum_{(I,y) \in D} \sum_{t=1}^K \log p(y_t|I, y_1, \dots, y_{t-1}; \theta) \quad (1)$$

where I is an input image, $y = \{y_1, \dots, y_K\}$ with $K < K_{\max}$ is the corresponding caption with length K , and K_{\max} is the maximum number of words in any caption. We also expand $p(y|I; \theta)$ by the fact that a word y_t in a caption y depends on both the input image I the previously generated word sequence $\{y_1, \dots, y_{t-1}\}$. The basic structure of an end-to-end image captioning network is straightforward (Vinyals et al., 2015), consisting of

two components: a visual feature extractor and a text decoder, as shown in Fig. 1 (b) and (d).

A raw input image can be very large in size, which is computationally expensive to process directly. So we need a visual feature extractor to extract the most informative features from the input image for the processing steps that follow. The visual feature extractor f extracts visual features from an image I and encodes them in to a fixed and lower dimensional feature vector x representing the image. Practically, to make the feature extractor differentiable and trainable, we implement it with convolutional layers following the VGG-19 architecture (Simonyan and Zisserman, 2015) and parameterize it with θ_f .

$$x = f(I; \theta_f) \quad (2)$$

A text decoder T then follows to translate the feature vector x into a caption y . One challenge is that the length of y may vary. Thus we implement the text decoder with an LSTM (see Appendix A for details about LSTMs), which is capable of generating outputs with variable lengths. And it is natural as well to use an LSTM to model the dependency $p(y_t | I, y_1, \dots, y_{t-1})$ of a word y_t on the prior word sequence. Similarly, we parameterize the text decoder with θ_T .

$$y = T(x; \theta_T) \quad (3)$$

Putting together, an end-to-end image captioning network can be simply written as $y = T(f(I; \theta_f); \theta_T)$ and we denote $\theta = (\theta_f, \theta_T)$ as the complete set of parameters for the network. Given the image-caption pairs from a training dataset D , we can train the network by rolling out the LSTM units in according to the target caption length and optimizing θ with any gradient method (Kingma and BaAdam, 2015) in a supervised manner with target images as input and target captions as output.

3.2. Visual attention mechanisms

At each stage $1 \leq t \leq K$, the text decoder LSTM generates a word y_t and requires an external input x_t (see Appendix A). With only one single feature vector $x = f(I; \theta_f)$ extracted from the input image I , there are two workarounds to fulfill the requirements. (1) We may set the first external input $x_1 = x$ and the rest $x_t = 0, \forall t > 1$, but this approach may lead to the lack of sufficient input information for the LSTM units at $t > 1$ to generate the proper words. Or (2) we may set $x_t = x, \forall 1 \leq t \leq K$, which however may also cause confusion easily for the network, since the external input at any stages is identical.

The above two issues can be resolved by the introduction of visual attention, as shown in Fig. 1 (c), which constrains the attention of the network to particular parts of the input image at each stage t . Then the corresponding external input x_t becomes the attention-weighted visual information from the input image, and thus varies at different stages. In this work, we adopt the Bahdanau (soft) attention (Bahdanau et al., 2015; Xu et al., 2015) among other approaches (Luong et al., 2015; Karpathy and Fei-Fei, 2015). A attention generator g parameterized by θ_g , considering both the current step feature x_t and the previous step output y_{t-1} , produces an attention weight vector α_t .

$$e_{t,i} = g(x_{t,i}, h_{t-1}; \theta_g) \alpha_{t,i} = \frac{\exp(e_{t,i})}{\sum_{j=1}^L \exp(e_{t,j})} \quad (4)$$

In Eq. (4) above, the generator g generates a pre-normalized weight $e_{t,i}$ for each dimension $x_{t,i}$ in x_t individually, where $1 \leq i \leq L$ and L is the size of x_t . The weights are then normalized with a softmax function to produce an attention mask $\alpha_t = [\alpha_{t,1} \quad \dots \quad \alpha_{t,L}]^T$ such that $\alpha_{t,i} > 0, \forall t, i$ and $\sum_{i=1}^L \alpha_{t,i} = 1, \forall t$. The raw feature vector input x_t to the LSTM is then augmented by the attention weighted average feature to indicate the attention region.

$$x_t^{(LSTM)} = \left(x_t, \sum_{i=1}^L \alpha_{t,i} x_{t,i} \right) \quad (5)$$

Here $x_t^{(LSTM)}$ denotes the external input to the LSTM at stage t . The training procedure remains the same as mentioned in the previous section.

4. Text similarity for image prioritization and local search

While onboard image prioritization requires a remote device to prioritize image data in accord to scientist inputs to downlink, local image search requires a local server to prioritize image search results in accord to user inputs to return to the user. These two tasks are similar in the sense that they both prioritize images according to some inputs. We propose that natural language description (in text) is an intuitive and convenient form of such inputs, since it does not require comprehensive understanding of the system to design a set of rules and target feature vectors to look for as in (Castano et al., 2003, 2007a; Woods et al., 2011). With text descriptions being inputs and the SCOTI network generating captions for images captured onboard and/or stored on a local server, we essentially converted the remote/local image prioritization problem into a text similarity evaluation task.

A large amount of prior work has been done on text similarity evaluation in the machine translation community (Han, 2018). Among those proposed approaches, we follow the BLEU metric (Papineni et al., 2002) to evaluate the similarity between auto-generated image captions and scientist/user inputs. BLEU can evaluate the similarity between a candidate text (an auto-generated image caption) and a set of reference texts (scientist/user inputs) jointly, which is advantageous since it is common to have multiple inputs from scientists to describe a set of target features to look for. And BLEU also accounts for both the word-wise and the phase-wise match rate between two texts through n -gram precision. An n -gram is a consecutive sequence of n word(s). Given a candidate c and a set of reference texts R , the n -gram precision $p_n(c, R)$ defines a metric that tells how much the candidate c matches the reference set R .

$$p_n(c, R) = \frac{\sum_{g \in G_{n,c}} \min\{C_{g|c}, \max_{r \in R} C_{g|r}\}}{\sum_{g \in G_{n,c}} C_{g|c}} \quad (6)$$

where $G_{n,t}$ denotes the set of all unique n -grams in text t , while $C_{g|t}$ counts the occurrences of an n -gram g in t . And the minimization term in the numerator essentially gives the number of non-repeated occurrences of an n -gram g in both the candidate c and any of the reference $r \in R$. However, p_n tends to favor shorter candidates. To reduce such an artifact, a penalty term

$$\eta(c, r) = \begin{cases} 1 & \text{if } L(c) > L(r) \\ \exp\left(1 - \frac{L(r)}{L(c)}\right) & \text{if } L(c) \leq L(r) \end{cases}$$

is introduced to penalize candidates with short length, where $L(c)$ and $L(r)$ denote the lengths of candidate c and reference r respectively. The BLEU score $s_N(c, R)$, which accounts both the length of candidate c and all n -grams with n from 1 up to N , can then be defined.

$$s_N(c, R) = \eta(c, r_{ML}) \exp\left(\sum_{n=1}^N w_n \log p_n(c, R)\right)$$

where r_{ML} is the reference text with the maximum length in the reference set R , and the hyperparameter w_n weights the importance of n -grams with different length n . The higher s_N a caption is associated with, the high similarity it has to the set of user/scientist inputs. Practically, the logarithmic BLEU score $\log s_N$ is more efficient to compute, and we may further improve onboard image prioritization by combining other common strategies, such as novelty detection and representative sampling (Castaño et al., 2003).

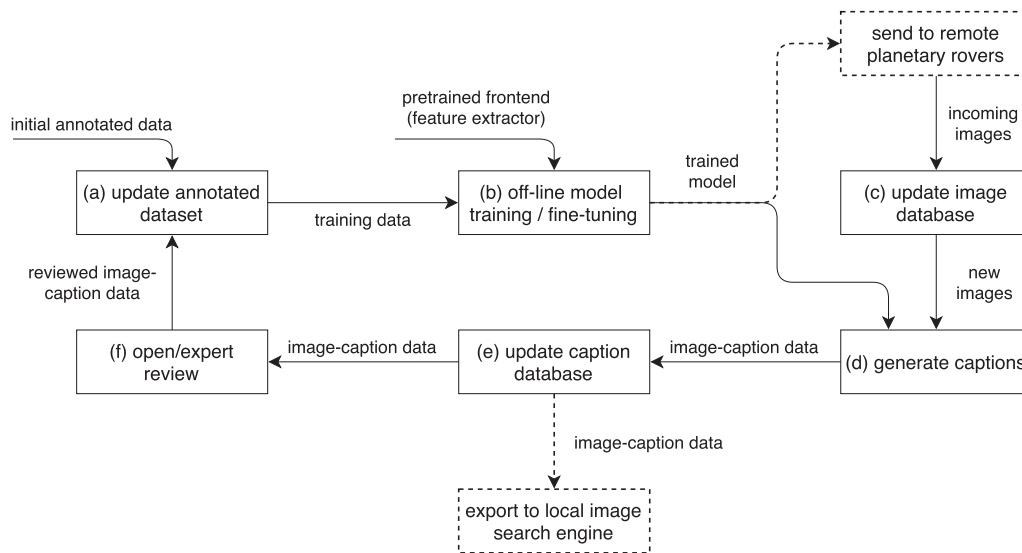
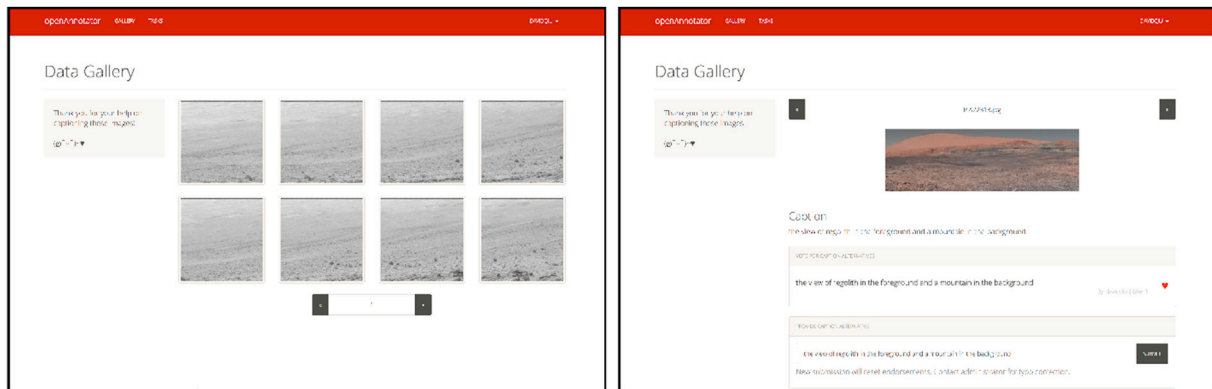


Fig. 2. Workflow of the data aggregation and model retraining pipeline.



(a) Gallery Page

(b) Annotation Page

Fig. 3. The user interface of *OpenAnnotator*.

5. Data pipeline

An important assumption made by supervised learning is data from the training dataset and the testing dataset should be drawn from the same distribution. It implies besides sufficient quantity and variance of the training data, it is also vital to have the training dataset update-to-date so that the SCOTI network model can be retrained on it and hence is applicable to the acquired images from latest missions. For this reason, we build a data pipeline to import and caption incoming images, allow expert reviews, and retrain the SCOTI model to keep it update-to-date.

As shown in Fig. 2, the data pipeline starts with initial annotated data. After (a) updating the annotated dataset, training data are used to (b) train the SCOTI model. The trained model can be sent to a remote rover to support onboard image prioritization, and used to (d) generate captions from new images from the image database (c) updated with incoming images. The auto-generated captions for all images will be (e) updated to the caption database, where the image-caption data can be exported to a local server to support local image search. New images with captions generated by SCOTI are sent to *OpenAnnotator* for open/expert review, after which reviewed and annotated data are (a) updated to the annotated dataset for the next round of model training. A data aggregation rate δ is also defined to indicate the proportion of new annotated data compared to those used in the last round of model training. Only

after $\delta > 0.5$ will the next round of model training begin.

5.1. *OpenAnnotator*: A web-based multiuser annotation tool

OpenAnnotator is a web-based multi-user tool we developed for image-caption data review and annotation. With this tool, users/experts can vote for a caption or propose a different version for an image from the database.

OpenAnnotator helps minimize the labor required for data annotation, since an auto-generated caption is provided in the first place for each image, which users can directly vote for. On the other hand, users may propose a different caption by referring to existing ones so that they do not have to conceive a new one from scratch. Moreover, it enables data annotation in a distributed manner. Its potential integration to public data servers such as NASA's PDS prefigures the access to the help and the collective intelligence from all users on these servers for data review and annotation, which may also bring about positive social effect by encouraging users' participation (Fig. 3).

6. The Martian Image Caption Dataset

We created the *Martian Image Caption Dataset (MICD)* using *OpenAnnotator* (see section 5.1) to evaluate the SCOTI network as well as the




Fig. 4. Some data samples from MICD, where 4 image-caption pairs annotated by experts are presented on the left, and 16 more image samples are presented on the right to demonstrate the data variety of MICD.

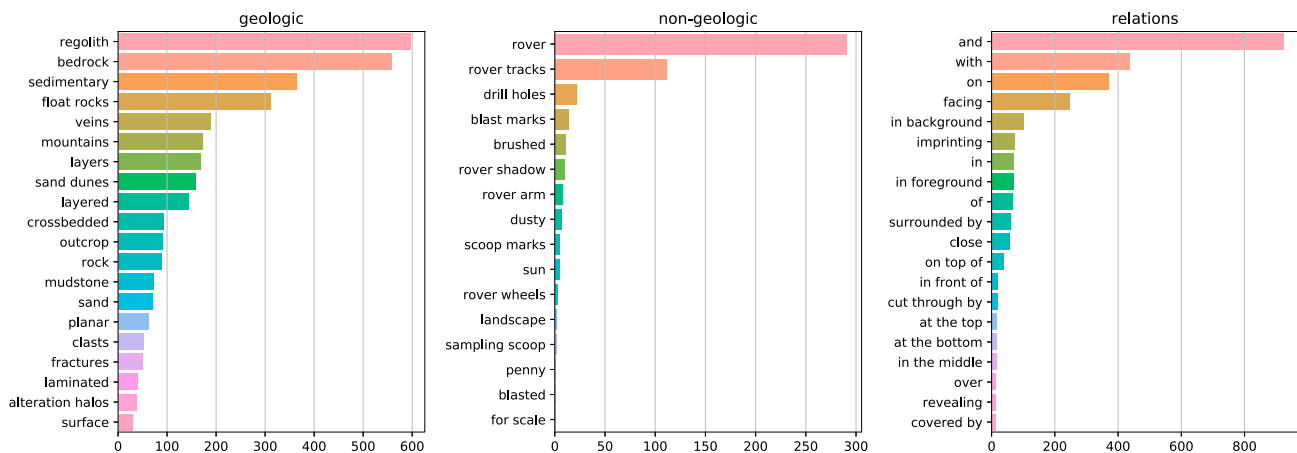


Fig. 5. Top-20 features in different categories identified from the MICD dataset.

proposed image data prioritization and local image search approaches. The dataset contains more than 12,500 pre-processed images captured by the Mars Science Laboratory (MSL) rover, among which 1,250 images are annotated with expert captions. And the number of annotated images continues to grow.

The images from MICD primarily capture Martian geologic features, especially terrain features, where more than one objects or feature categories may exist in the same image. These images can be grayscale or colorful, and non-geologic features may also be found in them, such as the ones in the last two rows on the right in Fig. 4.

Expert-annotated captions in MICD consists of a rich set of words and phrases describing geologic features, non-geologic features, colors and relations. The top-20 features in several different categories are shown in Fig. 5, and a more comprehensive summary of the features identified from MICD is also provided in Appendix B.

7. Experiment and results

We conducted both quantitative and qualitative experiment to evaluate the our proposed approach and system. In section 7.1, results of quantitative experiment on evaluating the performance of the SCOTI

network are presented, including details about its training process and metric scores on validation data. In sections 7.2 and 7.3, qualitative experiment results of onboard image prioritization simulation and local image search are discussed.

7.1. Training and evaluating SCOTI on MICD

Regardless of the extensive potential applications where the SCOTI system can be used in planetary and astronomical domains, we evaluated it primarily on Martian data from MICD in this work. In this section, we will discuss the training details and systematical evaluation of the SCOTI network. We also present quantitative results through the learning curve and metric scores, along with some concrete examples.

In the experiment, we split the annotated data from MICD into the training and validation datasets in a ratio of 0.9 : 0.1. Unlabelled data from MICD constitutes the test dataset. And we adopt the architecture of VGG-19 convolutional layers (Simonyan and Zisserman, 2015) as the visual feature extractor, as shown in 1 (b). Due to the limited size of annotated data from MICD, the weights θ_f of the visual feature extractor are pretrained on ImageNet (Deng et al., 2009) to prevent overfitting. After pretrained, θ_f is fixed while training (fine-tuning) the joint

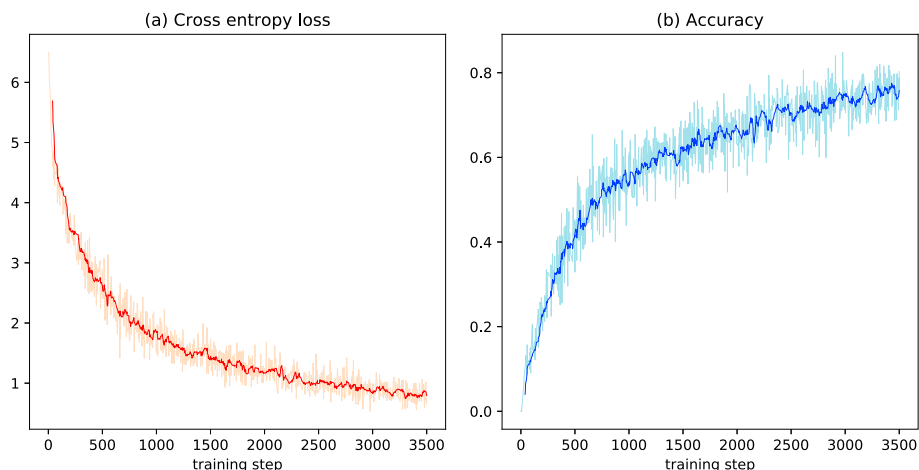


Fig. 6. Learning curves of SCOTI network on *MICD* training dataset throughout the 3500 training steps (100 training epochs). The highlighted curves are smoothed by averaging 10 nearby sample points from the actual learning curves shown on the background.




Fig. 7. BLEU-1, -2, -3, -4, METEOR, ROUGE-L, and CIDEr scores on the *MICD* validation dataset along 3500 training steps (100 training epochs).

Table 1

Evaluation of the SCOTI model trained on *MICD* with respect to BLEU-1, -2, -3, -4, METEOR, ROUGE-L and CIDEr scores, in comparison to those reported by Xu et al. (2015) on *Flickr8k*, *Flickr30k* and *MS COCO* datasets.

Model/Dataset	BLEU (B)				METEOR	ROUGE-L	CIDEr
	B-1	B-2	B-3	B-4			
<i>SCOTI/MICD</i>	0.547	0.482	0.431	0.388	0.309	0.606	3.939
Xu et al./ <i>Flickr8k</i>	0.670	0.448	0.299	0.195	0.189	–	–
Xu et al./ <i>Flickr30k</i>	0.667	0.434	0.288	0.191	0.185	–	–
Xu et al./ <i>COCO</i>	0.707	0.492	0.344	0.243	0.239	–	–

parameter θ_T of the visual attention layer and the text decoder, as shown in Fig. 1 (c) and (d), on the training dataset. The training process is set up as a supervised learning problem, where the inputs are raw images and the outputs are the corresponding expert annotated captions. Adam optimizer (Kingma and BaAdam, 2015) is used to optimize the cross-entropy loss between the predicted and the ground-truth captions. A training process for 3500 steps (100 epochs with 35 batches in each one of them) is shown with the learning curves, including the cross entropy loss curve and the accuracy curve, in Fig. 6 below.

Along the training process, the SCOTI model performance is evaluated with metric scores, including BLEU-1, -2, -3, -4 (Papineni et al., 2002), the Metric for Evaluation of Translation with Explicit Ordering (METEOR) (Banerjee and Lavie, 2005), the Recall-Oriented Understudy for Gisting Evaluation for Longest Common Subsequence (ROUGE-L) (Lin, 2004) and the Consensus-based Image Description Evaluation (CIDEr) (Vedantam et al., 2015). Fig. 7 visualizes these scores on the

MICD validation dataset along the 3500 training steps.

After trained for 3500 steps, the SCOTI network performs reasonably well with respect to the validation metric scores in comparison to the state-of-the-art reported in (Xu et al., 2015), where the visual feature extractor is also pretrained but the image captioning model is fine-tuned in more complex domains implicitly defined by *Flickr8k*, *Flickr30k* and *MS COCO* datasets. Table 1 compares the metric scores between SCOTI on *MICD* and the state-of-the-art.

An interesting observation is that though SCOTI trained on *MICD* almost outperforms the state-of-the-art trained on more datasets in more complex domains, but has a lower BLEU-1 score. More training steps until overfitting can further improve the performance of SCOTI, but it is important to note that the amount and variety of training data can also affect the model performance significantly. Compared to merely over a thousand annotated images in *MICD*, most public datasets such as *Flickr8k*, *Flickr30k* and *MS COCO* have more than tens of thousands of

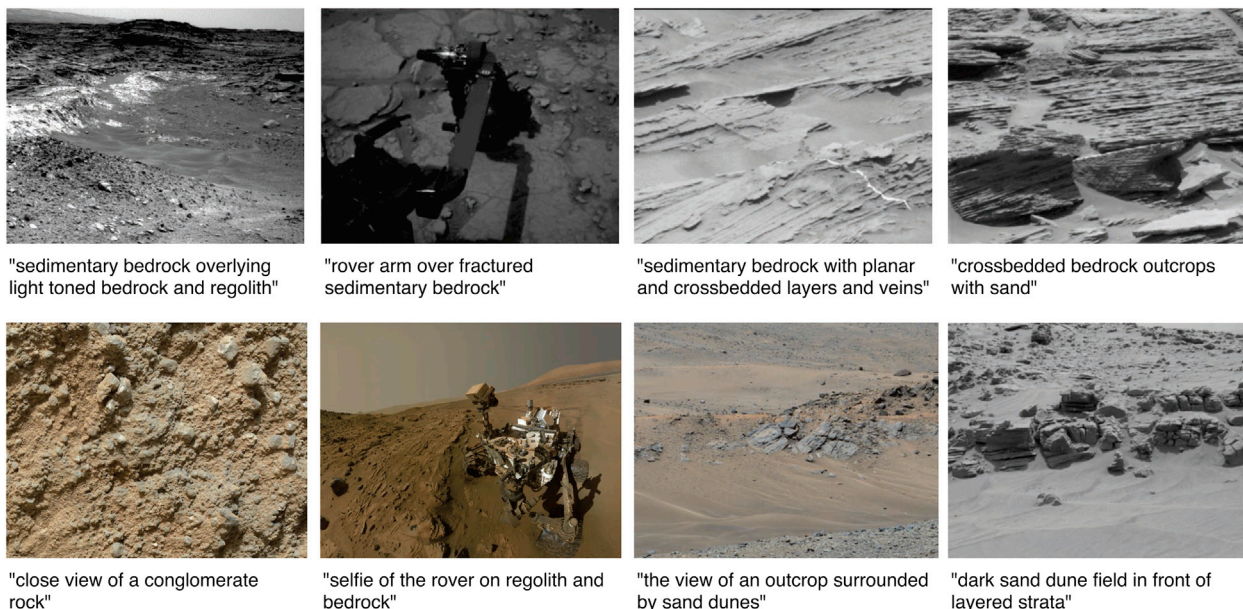
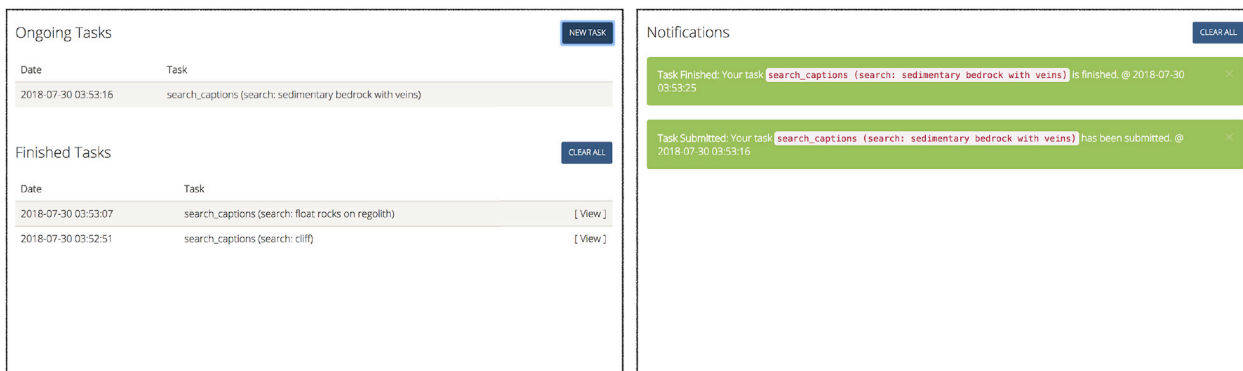


Fig. 8. Some captioning results of unlabelled images from MICD.



(a) Prioritized Downlink Tasks Page.

(b) Result Notification Page.

Fig. 9. User interface of the onboard image prioritization and data downlink simulator.

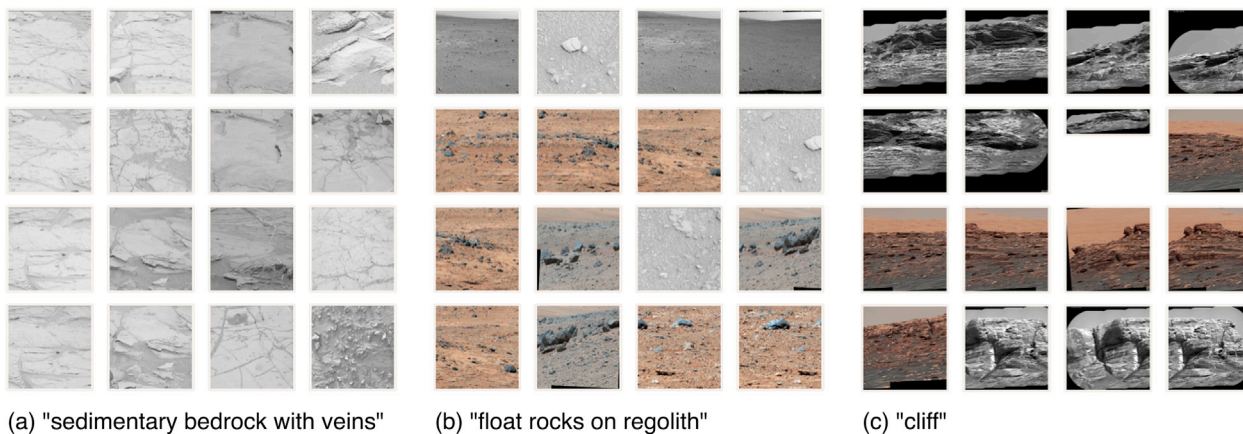


Fig. 10. Simulated prioritized downlink results given user (scientist) inputs.

annotated image-caption data. It is anticipated to see better performance of the SCOTI network after more annotated planetary image-caption data become available.

We also present in Fig. 8 some concrete image captioning results generated by the SCOTI model, showing that our model is capable of captioning images in various conditions including both greyscale and colored ones, taken from the planet Mars with a rich set of landscape features and objects in them.

7.2. Onboard image prioritization and downlink simulation

Integration of the proposed system, including the SCOTI network and the image prioritization pipeline, to a planetary rover or any other onboard system is nontrivial and requires much further research, optimization and validation. In this paper, we only explore the proof-of-concept level potential of leveraging machine learning to facilitate future missions in planetary exploration. To that end, we build a simulation system to simulate the onboard image prioritization and data downlink processes.

The simulation system consists of a simulated onboard server and a simulated ground-based server. A trained SCOTI network model and an image data prioritization pipeline are deployed to the simulated onboard server, which simulates a planetary rover or any other onboard device. A downlink task commander and an image data collector are implemented on the simulated ground-based server. A user can access the ground-based server through its user interface, as shown in Fig. 9, to submit a new prioritized data downlink task request, check task status and view finished tasks.

We load the unlabelled images from MICD to the simulated onboard server as the images captured by the onboard device. Fig. 10 shows the first 16 images downlinked after onboard image prioritization in three different prioritized downlink task simulations. In the above experiment, we use weights $\{w_n\}_{n=1}^4 = \{0.8, 0.15, 0.045, 0.005\}$ for BLEU in image prioritization, which is a set of hyperparameters we found to perform reasonably well in preliminary experiment.

7.3. Local image search through text descriptions

Local image search is essentially equivalent to onboard image prioritization using the same text similarity metric to return the images with the most similar auto-generated captions to the text description inputs from users. The onboard image prioritization and downlink test from the previous section directly reflects the capability of our proposed method on local image search. Furthermore, a beta version of text-based local image search tool was also integrated onto NASA PDS for public test (Ono et al., 2019).

8. Conclusion and discussion

We propose an approach that generalizes the tasks of both onboard image prioritization for data downlink and local image search for planetary image data servers into the problem of image captioning and text similarity evaluation. Though it is still a proof-of-concept level system, SCOTI demonstrates its potential in facilitating data transmission efficiency from planetary rovers to ground-based data system by prioritizing the image data with highest scientific values according to requests from scientist through text-based instructions. Such an approach avoids the need for scientists to have a deep understanding of the in-flight software engineering details in order to come up with a set of unintuitive feature vectors in order to instruct what data to prioritize for downlink. Furthermore, the similar essence behind image search and image prioritization allows us to naturally apply the same image captioning and text similarity based data prioritization mechanism to local image search. And a beta version of the local image search tool was deployed to NASA's PDS being available for the public (Ono et al., 2019). With the data

pipeline and the *OpenAnnotator* tool we developed, the number of annotated images in the *Martian Image Caption Dataset* continues to grow, and hence the performance of the SCOTI network and the overall system can keep improving.

Though in this paper we mainly focus on terrain images captured by Martian rovers, but it does not preclude the applications of the SCOTI system in any other planetary and astronomical context. In fact, the same framework as proposed can be applied on different kinds of visual data captured by any type of space device besides planetary rovers, simply by retraining SCOTI on the corresponding data. One typical example of such an extended application is images captured by satellites, which can be an interesting future work. And as aforementioned, we focus this work on the proof-of-concept level validation of leveraging machine learning, more specifically image captioning, on planetary exploration, further optimization, such as its implementation on a High Performance Space-flight Computing (HPSC) platform, and hardware-in-the-loop validation will be required in order to actually deploy the SCOTI system onto a planetary rover in any future flight mission. And these will be our following work to carry out.

CRedit authorship contribution statement

Dicong Qiu: Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization. **Brandon Rothrock:** Conceptualization, Methodology, Investigation, Resources, Writing - review & editing, Supervision, Project administration. **Tanvir Islam:** Software, Validation, Investigation, Visualization. **Annie K. Didier:** Resources, Writing - review & editing. **Vivian Z. Sun:** Resources, Data curation, Writing - review & editing. **Chris A. Mattmann:** Resources, Writing - review & editing. **Masahiro Ono:** Conceptualization, Methodology, Resources, Writing - review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgement

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. U.S. Government sponsorship is acknowledged. The authors would also like to thank Bhavin T. Shah and Adrian Stoica for supporting this work.

Appendix. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.pss.2020.104943>.

References

- Bahdanau, D., Cho, K., Bengio, Y., 2015. Neural machine translation by jointly learning to align and translate. In: The 3rd International Conference on Learning Representations.
- Banerjee, S., Lavie, A., 2005. Meteor: an automatic metric for mt evaluation with improved correlation with human judgments. In: Proceedings of the Acl Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation And/or Summarization, pp. 65–72.
- Castano, R., Anderson, R., Estlin, T., DeCoste, D., Fisher, F., Gaines, D., Mazzoni, D., Judd, M., 2003. Techniques for onboard prioritization of science data for transmission. IPN Progress Report, pp. 42–153.
- Castano, R., Judd, M., Anderson, R.C., Estlin, T., 2003. Machine learning challenges in mars rover traverse science. In: ICML Workshop on Machine Learning Technologies for Autonomous Space.
- Castano, R., Estlin, T., Gaines, D., Chouinard, C., Bornstein, B., Anderson, R.C., Burl, M., Thompson, D., Castano, A., Judd, M., 2007. Onboard autonomous rover science. In: Aerospace Conference. IEEE, pp. 1–13.
- Castano, R., Estlin, T., Anderson, R.C., Gaines, D.M., Castano, A., Bornstein, B., Chouinard, C., Judd, M., 2007. Oasis: onboard autonomous science investigation system for opportunistic rover science. J. Field Robot. 24 (5), 379–397.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: a large-scale hierarchical image database. In: Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on, IEEE, pp. 248–255.

- Gomaa, W.H., Fahmy, A.A., 2013. A survey of text similarity approaches. *Int. J. Comput. Appl.* 68 (13), 13–18.
- Han, L., 2018. Machine translation evaluation resources and methods: a survey. In: *Proceedings of the Ireland Postgraduate Research Conference*.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Karpathy, A., Fei-Fei, L., 2015. Deep visual-semantic alignments for generating image descriptions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3128–3137.
- Kingma, D.P., Ba, J., Adam, 2015. A method for stochastic optimization. In: *The 3rd International Conference on Learning Representations. ICLR*.
- Kiros, R., Salakhutdinov, R., Zemel, R., 2014. Multimodal neural language models. In: *International Conference on Machine Learning*, pp. 595–603.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning, *nature* 521 (7553), 436–444.
- Thompson, D.R., Abbey, W., Allwood, A., Bekker, D., Bornstein, B., Cabrol, N.A., Castaño, R., Estlin, T., Fuchs, T., Wagstaff, K.L., 2012. Smart Cameras for Remote Science Survey. JPL TRS.
- Lin, C.-Y., 2004. Rouge: a package for automatic evaluation of summaries. *Text Summarization Branches Out* 74–81.
- Ono, M., Rothrock, B., Mattmann, C., Islam, T., Didier, A., Sun, V., Qiu, D., Ramirez, P., Grimes, K., Hedrick, G., 2019. Make planetary images searchable: content-based search for pds and on-board datasets. In: *Lunar and Planetary Science Conference*, vol. 50.
- Papineni, K., Roukos, S., Ward, T., Zhu, W.-J., 2002. Bleu: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics. Association for Computational Linguistics*, pp. 311–318.
- Rothrock, B., Kennedy, R., Cunningham, C., Papon, J., Heverly, M., Ono, M., 2016. Spoc: deep learning-based terrain classification for mars rover missions. In: *AIAA SPACE. AIAA*, p. 5539, 2016.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. In: *The 3rd International Conference on Learning Representations*.
- Luong, M.-T., Pham, H., Manning, C.D., 2015. Effective Approaches to Attention-Based Neural Machine Translation *arXiv preprint arXiv. 1508.04025*.
- Vedantam, R., Lawrence Zitnick, C., Parikh, D., 2015. Cider: Consensus-based image description evaluation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4566–4575.
- Vinyals, O., Toshev, A., Bengio, S., Erhan, D., 2015. Show and tell: a neural image caption generator. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3156–3164.
- Wagstaff, K.L., Lu, Y., Stanboli, A., Grimes, K., Gowda, T., Padams, J., 2018. Deep mars: cnn classification of mars imagery for the pds imaging atlas. In: *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Woods, M., Shaw, A., Rendell, P., 2011. High-level autonomy and image prioritisation for long distance mars rovers. In: *Proc. 11th Symposium on Advanced Space Technologies in Robotics and Automation. ESA/ESTEC, The Netherlands*, pp. 12–14.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y., 2015. Show, attend and tell: neural image caption generation with visual attention. In: *International Conference on Machine Learning*, pp. 2048–2057.