Expert Guided Rule Based Prioritization of Scientifically Relevant Images for Downlinking over Limited Bandwidth from Planetary Orbiters

Srija Chakraborty,[1] **Subhasish Das,**[1] **Ayan Banerjee,**[1]
Sandeep K. S. Gupta,[1] **Philip R. Christensen**[2]

[1]School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281, USA
[2]School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA
{schakr34, sdas48, abanerj3, sandeep.gupta, phil.christensen}@asu.edu

Abstract

Instruments onboard spacecraft acquire large amounts of data which is to be transmitted over a very low bandwidth. Consequently for some missions, the volume of data collected greatly exceeds the volume that can be downlinked before the next orbit. This necessitates the introduction of an intelligent autonomous decision making module that maximizes the return of the most scientifically relevant dataset over the low bandwidth for experts to analyze further. We propose an iterative rule based approach, guided by expert knowledge, to represent scientifically interesting geological landforms with respect to expert selected attributes. The rules are utilized to assign a priority based on how novel a test instance is with respect to its rule. High priority instances from the test set are used to iteratively update the learned rules. We then determine the effectiveness of the proposed approach on images acquired by a Mars orbiter and observe an expert-acceptable prioritization order generated by the rules that can potentially increase the return of scientifically relevant observations.

Introduction

Vast amounts of data acquired by instruments onboard spacecraft offer rich insights that improve our understanding of the observed planetary surfaces. However, as the data from deep space missions are to be transmitted over a limited bandwidth before the next orbit, the available downlink volume is significantly lower than acquired data volume. Under the current setting, only a subset of observations is targeted such that it can be transmitted over the available bandwidth to Earth. This subset selection does not leverage any intelligent decision making strategy aimed at identifying the datasets that are scientifically more significant for analysis by domain experts. This minimizes the possibility of discovering novel instances to expand the existing knowledge base and may result in the loss of scientifically valuable observations for experts to analyze. Thus, it is crucial to introduce an intelligent autonomous decision making module operating remotely onboard the spacecraft, to prioritize the transmission of datasets that have higher significance for further analysis. Efforts in this domain have focused on detecting application specific classes of interest such as Martian polar ice caps (Castano et al. 2007), dust devil tracks and

clouds (Castano et al. 2008) to determine scientific value. We introduce a methodology that also considers contextual attributes, such as relevance of a class or landform to experts, co-occurrence of landforms and its spatial and temporal properties to determine its relevance. This resembles expert like interpretation more closely instead of determining relevance from the presence or absence of interesting features alone. We then determine the relative importance of the instances based on these attributes, where it is crucial for the module to *understand* the content in the image like an expert. Prioritization has been studied with planetary rover data to select rocks with a specific target signature (Estlin et al. 2012), for autonomous exploration by integrating past orbital datasets with a science hypothesis (Thompson et al. 2018). We focus on formulating a generalized representation of the already observed data and allow the module to determine if a new instance is unusual. We combine strategies such as rule based knowledge repository construction and novelty detection to form a meaningful representation of interesting classes and identify outliers with respect to these representations to favor exploration. Additionally, what makes an instance *interesting* is subjective to expert interests. Our strategy therefore allows expert guidance to refine the rules with domain knowledge.

The objective of the proposed method is to emulate expert-like identification of *relevant* datasets to maximize the return of high value data over limited bandwidth. We only focus on geological features in planetary images *interesting* to experts. All observations with these expert identified interesting landforms (ILFs) have a high downlink priority. Additionally, downlink priority also depends on how novel or anomalous the instance is with respect to expert selected attributes. The more anomalous the instance is with respect to expert defined attributes, the higher is its value for analysis and is hence assigned a higher priority.

We introduce an expert guided rule formation methodology where the rules represent the expected characteristics of each ILF with respect to the attributes. Deviation of an ILF bearing instance from its rule causes it to be more significant for further investigation. We quantify this deviation as *relevance score* to characterize the degree of novelty of an ILF bearing instance. In order for the proposed approach to generate an expert acceptable priority order, the rules have to be aware of domain knowledge. We elicit this knowledge by

seeking expert rating on observed relevant instances from which the rules are formed. The rating allows incorporation of prior domain knowledge to guide the rules in the decision making process. Along with deciding on the priority of an instance, the prioritization step also lists the attributes with respect to which the instance is novel, explaining *why* the decision was made to increase the expert's trust in the system. We evaluate our approach on images from the early operation stage of the Thermal Emission Imaging System (THEMIS) instrument onboard the Mars Odyssey spacecraft. We repeat our evaluation over successive orbits where the rules are updated to account for changes and observe that the acceptability of the rule generated ranking to experts increases when we incorporate domain knowledge.

The contribution of our work is as follows. Firstly, we introduce a rule based representation of interesting landmarks on a planetary surface with respect to expert selected attributes. Secondly, we allow expert feedback to guide the rules with domain specific preferences to determine the relevance score of an instance. Finally, we evaluate the acceptability of our assigned priority scores to experts and provide them with explanations of the decision. We show examples of successful detection of novel instances and demonstrate that the proposed approach is capable of assigning acceptable priorities to relevant images.

Related Work

Relevant instance retrieval and *learning to rank* class of problems have been studied with applications in both document (Silva, Gonçalves, and Veloso 2011) and image retrieval (Zhou, Li, and Tian 2017), where conventionally features such as BM25 and Page Rank have been used to symbolize documents, while SIFT features, Fisher vectors have been utilized to represent images. Recent studies have utilized autoencoder features for document ranking given a query (Albuquerque et al. 2018) and patch-based (Paulin et al. 2015) or holistic neural codes extracted from the intermediate layers of neural networks for image retrieval. However, in the proposed work, we emphasize on detecting novelty (by considering association, spatial and temporal attributes to interpret the scene content) instead of similar instances while allowing expert preference to guide the prioritization. Requirement based prioritization of test cases using k-means clustering ranks test cases that meet pre-specified requirements (Arafeen and Do 2013). In contrast, this work has to effectively identify unseen ILF associations autonomously instead of assigning it to an already known cluster.

We adopt a rule based approach to represent the classes of interest as it can extract meaningful relations, if any, even from a small amount of data. The proposed approach differs from existing rule based approaches (Lee et al. 2007) and semantics-enabled information mining (Kurte et al. 2017) as the extracted rules favor discovering novel instances and incorporate domain specific preferences.

System Model

The proposed framework of rule based determination of scientific value of relevant instances is shown in Figure 1. In
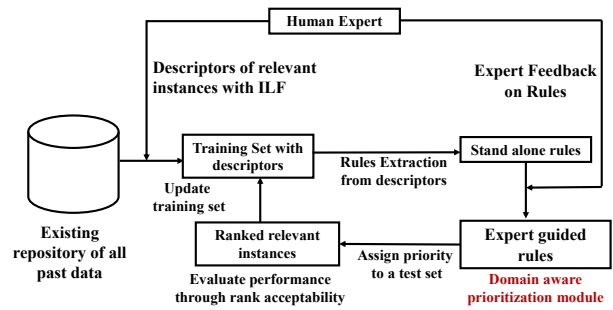

Figure 1: Proposed expert guided rule extraction framework

this study, the expert selected ILFs are dunes, dust devil tracks, channels and wind streaks and the rules are to be formed with respect to each ILF. The instances with ILFs are selected by experts from the existing observation repository. The expert also describes the selected instances with respect to the attributes to form the training set. The proposed approach then forms a rule based representation of each class from the training set. The rules are then updated using expert feedback to form the prioritization module, which is then tested to assign priority to a test set containing relevant instances from subsequent orbits. The instances from the test set are then used to update the rules, which are used to prioritize the next batch of instances. In this work, we focus on determining the relevance of expert annotated instances. Thus, both training and test sets contain relevant instances only as the focus here is on the prioritization module to assign priorities to instances based on its scientific value.

Proposed Approach: Prioritization by Identifying Anomalous Instances

In this section, we propose a rule based representation that defines the expected attributes of each ILF. Under the prioritization criterion chosen by experts, any instance containing one or more ILFs has a high priority for downlinking, expressed in terms of *relevance score*. However, the more unusual the instance is with respect to the rules, the higher its novelty and relevance score is for further investigation. Such unusual instances have higher downlink priority than an ILF bearing instance that meets the expected rules. The novelty of a relevant instance with respect to the rules is determined based on expert selected attributes which are:

- **class relevance**: This denotes the relative relevance of the classes present in the instance for further analysis. Although all relevant instances contain one or more ILF, which increases its downlink priority, the relevance of each class may vary from expert to expert. Presence of ILFs that are more relevant to the expert increases the relevance score of the instance.

- **spatial context or inter-class association**: The context of the images determined by the association of ILFs within an instance. For example, all past observations of a given ILF i may appear with another ILF j in the same instance. However, in a test image, if ILF i occurs by itself or is

present with another ILF l, it increases the relevance score of the instance as it implies aberration in the data generating process that is significant for further investigation.

- **spatial location**: Unusual spatial location of any ILF with respect to its previously observed spatial distribution. This increases the relevance score of the instance and assists experts to explore unresearched regions thereby enhancing existing surface maps.

- **seasonal information**: An ILF may be prevalent during a particular season. Any occurrence of that ILF in a different season is considered anomalous and increases the relevance score of the instance it appears in.

Rule Formation of Relevant Classes

Rule based characterization of ILFs defines the expected attributes for each of the C ILFs from its observed instances. Presence of an ILF renders it important for downlinking however, the rules are used to identify the most unique instances and hence impact prioritization. The rules are learned iteratively from the observations acquired in each orbit. In a given orbit y, from the set $D_y = \{d_1, d_2, \ldots, d_N\}$ of all N observations acquired, the relevant subset $\Theta_y = \{\theta_1, \theta_2, \ldots, \theta_{N_C}\}$ containing N_C instances with ILFs are identified by experts. For all N_C instances, experts also provide the instance descriptor $\theta \equiv < labels, latitude, longitude, season >$ for ground based rule generation with respect to the attributes. These instance descriptors form the training dataset for rule generation.

Strength or frequency of any combination of inter-class association between ILFs $(i, j, \ldots, n)$ is the joint probability of its occurrence together within an instance and is computed empirically from the training dataset. Strength indicates the likelihood of co-occurrence of two or more ILFs in one instance and does not necessarily imply causality. If the ILFs $(i, j, \ldots, n)$ co-occur in θ_s instances in the set Θ_y and $|\theta_s|$ denotes the number of such instances, then the strength of association $s_{i,j,\ldots,n}$, between the given combinations of ILFs is computed using

$$s_{i,j,\ldots,n} = \frac{|\theta_s|}{N_c}, \tag{1}$$

where $\theta_s \subseteq \Theta_y$. While a high s indicates that the association between the ILFs is frequently observed, low s implies that the association is rarely observed, which increases its relevance for further analysis. For unseen associations, $s = 0$.

In addition to inter-class associations, the rules also incorporate the set of expected latitude and longitude center for each ILF i to determine spatial anomalies. If a given ILF is detected in a spatial location that is unusual with respect to all previously observed instances of that ILF, such a test instance has a high downlink priority as it prompts experts towards new and previously unknown locations of that ILF, thereby aiding in exploration and enhancing existing surface maps. If X_i and Y_i are the vector of latitudes and longitudes of all instances of ILF i in the training data, its joint density function p_{XY_i} is determined using the kernel density estimate (KDE) to identify spatial clusters of the given class. The latitude and longitude centers are those (X_i, Y_i), which

causes density function p_{XY_i} at X_i, Y_i to be a local maxima. All instances of ILF$_i$ are then assigned to the closest cluster center $(X_{\omega_i}, Y_{\omega_i})$, where ω is the number of clusters of the ILF i determined from the KDE. Following cluster assignment, the rules for each ILF is augmented with the cluster centers and spatial variances $(\sigma_{lat_{\omega_i}}, \sigma_{long_{\omega_i}})$ of each cluster of that ILF. The predominant spatial locations of each ILF is the set $\Omega_i = \{(X_{\omega_i} \pm \sigma_{lat_{\omega_i}}, Y_{\omega_i} \pm \sigma_{long_{\omega_i}})\}$.

Season of prevalence learned from past data of each ILF is used to characterize its seasonality. Season of occurrence of each observed instance of every ILF i is expressed as a binary vector $\mathbf{t}$, whose length corresponds to the number of seasons ϕ on Mars (Cantor, James, and Calvin 2010). For every past instance θ of ILF i, if the associated season is φ, then $t_\theta[k] = 1$, for $k = \varphi$ and $t_\theta[k] = 0$ otherwise, where $k = 1, \ldots, \phi$. The expected season of prevalence μ_{k_i}, of each ILF i can be expressed as $\mu_{k_i} = E[k_i] = \sum_{k=1}^{\phi} k \cdot f(k)_i$, while the seasonal variance is given by $\sigma(k_i) = \sum_{k=1}^{\phi} (k - E[k_i])^2 \cdot f(k)_i$. Here $f(k)_i = \frac{1}{N_{c_i}} \sum_{I \in \text{ILF}_{c_i}} t_\theta[k]$, represents the prevalence of ILF i in season k learned from N_{c_i} instances of N where i appears.

Incorporating Expert Knowledge

Rules extracted for every relevant class is uninformed of expert preferences and domain knowledge. This may hamper the onboard module's effectiveness in generating expert acceptable ranking. In this section, we describe a technique to enrich the extracted rules by incorporating it with expert knowledge. This is essential to determine expert preference of relative importance of each class. Additionally, expert preferences are crucial for refining association strength. This is important when the rules are formulated from local knowledge on a small training dataset. For example, in the early stages of operation when only a small fraction of the surface has been observed, expert knowledge is useful to incorporate any known prior global context, which reduces any bias that may arise due to limited imaging of the surface. Moreover, this knowledge is utilized to override any data driven rule and guide it with expert knowledge when the relevance of that rule cannot be directly inferred from the data. For example, a frequent association may still bear high significance due to a domain specific reason. Seeking expert feedback ensures that the rules are aware of this context.

However, expecting accurate quantification of this knowledge even by experts is unrealistic and remains a challenge (Castano et al. 2005). We elicit this knowledge from experts by recording their rating of images from the training data of orbit y. For every expert identified relevant instance, we determine their preference by recording the *factor* (ILF relevance, interesting inter-class association) that caused the instance to be relevant on a scale of 1 to 5, indicating the *strength of preference*.

These ratings are used to guide the rules with expert opinion in the form of weights governing the class importance w_i and weights for all combinations of observed inter-class associations $w_{i,j,\ldots,n}$. Here w_i represents the relative importance of ILF i to the experts, whereas $w_{i,j,\ldots,n}$ represents the expert rating on the novelty or importance of associa-

tion between ILFs $i, j, \ldots, n$. These weights are initially set to unity and then updated to the average normalized expert rating of every instance bearing the ILF i or observed inter-class association $i, j, \ldots, n$ respectively. The rules formed for each ILF after expert feedback can be expressed as $< w_i, s_i, \Omega_i, (\mu_{k_i}, \sigma_{k_i}) >$, where s_i is a vector indicating association strength of i with all other class combinations. We utilized a feedback form that displays an image of each ILF category and each observed ILF association to record the corresponding expert weight. These weights are used to determine the class relevance and update the inter-class association strength, while the expected spatial location and season is learned in a data driven manner.

Inference Rules for Prioritization

The rules extracted define the expected attributes with respect to each relevant class. In this section, we propose an inference methodology to determine the relevance score of the test instances from orbit $y+1$ with respect to each attribute of the rules. A higher relevance score is indicative of greater significance of the given instance for downlinking because it is more likely to be relevant to experts. The novelty of each instance with the ILFs are determined using the attributes defined above, namely class weight, association strengths of multiple classes, spatial and temporal anomaly. For an ILF bearing instance τ in the test set, that can be described as $< \{i, \ldots, n\}, (X_\tau, Y_\tau), t_\tau >$, where $\{i, \ldots, n\}$ denotes ILFs present in τ, (X_τ, Y_τ), t_τ indicates the latitude/longitude and season of τ respectively, the inference rules to compute the relevance scores of τ with respect to each attribute is determined by:

- overall relevance of each ILF i in the instance, computed from expert weight $\sum_{\forall i \in \tau} w_i$,

- novelty of association $a_{i,j,\ldots,n}$ between all ILFs in the given instance, which is given as

$$a_{i,j,\ldots,n} = \frac{w_{(i,j,\ldots,n)}}{(s_{(i,j,\ldots,n)} + \epsilon)}, \forall (i, j, \ldots, n) \in \tau \quad (2)$$

where s is the association strength of all ILFs in τ computed using equation (1) and ϵ is a constant ($\epsilon << s$), representing the strength of previously unseen associations. For unobserved associations, $w_{(i,j,\ldots,n)} = 1$. Here, the inverse of strength indicates how rare the association is.

- occurrence of ILFs in a previously unknown location. This is indicated when (X_τ, Y_τ) is further away than the expected variance from its nearest cluster center. Spatial novelty is computed using

$$\sum_{\forall i \in \tau} \delta_i - \sigma_{lat_{\omega_i}}, \text{if} \quad \delta_i > \sigma_{lat_{\omega_i}}, \quad (3)$$

where $\delta_i = min(|X_\tau - X_{\omega_i}|)$ for all spatial clusters of ILF i. Anomalous longitude occurrence also contributes to relevance score similarly. This step is similar to assigning an instance to a previously known cluster.

- occurrence of i in its non-prevalent season, indicating temporal anomaly. For a given instance τ, the temporal novelty score is computed using

$$\sum_{\forall i \in \tau} \delta_t - \sigma_{k_i}, \text{if} \quad \delta_t > \sigma_{k_i}, \text{where}(\delta_t = |t_\tau - \mu_{k_i}|). \quad (4)$$

Thus, the overall relevance score of an instance is the sum of the ILF relevance score, inter-class association, spatial and temporal relevance scores. However, when two or more instances have equal relevance scores, higher priority is assigned to the instance that has ILFs with higher w_i and $a_{i,j,\ldots,n}$. A second tie-breaker is added that assigns higher priority to the instances with higher types of novelties. For example, an instance with inter-class association, spatial and temporal novelties is assigned higher priority than an instance with inter-class association and spatial novelties only.

Rule Update

The relevant instances from orbit $y+1$ are added to the training dataset to update the previously known rules using the proposed rule formation approach. Although the expert knowledge elicited from the previous training dataset is retained to guide the updated rules, an expert can update their ratings at this phase to appropriately direct the rules to reflect any modifications to the geological classes and its attributes of interest. However, feedback during update phases are expected to be infrequent. These updated rules are in turn utilized for prioritizing the test data in the following orbit.

Evaluation Metrics

Our proposed approach ranks all instances in a test set batch based on how interesting its contents are for experts to analyze. The effectiveness of our approach is determined by its ability to generate expert acceptable ranking in a given test batch. Hence, for every test batch, we evaluate the similarity in order of ranking $\mathbf{r} = [r_{\tau_1}, \ldots, r_{\tau_B}]$ generated from the rules for each test set batch with that of the expert order of ranking $\mathbf{e} = [e_{\tau_1}, \ldots, e_{\tau_B}]$ of the same set, where B denotes the number of instances in each batch and r_{τ_b} and e_{τ_b} denotes the rule generated rank and expert rank of instance τ_b. The metrics used to quantify this similarity in ranking are

- Mean Average Precision (MAP): For a given test set batch q, we compute the average precision (AP) at every position b where the rule generated rank r_{τ_b} of the instance τ_b matches its expert rank e_{τ_b}. This is expressed as

$$AP(b)_q = \frac{1}{B} \sum_{b=1}^{B} \gamma_b, \quad (5)$$

where

$$\gamma_b = \begin{cases} p_b, & \text{if } r_{\tau_b} = e_{\tau_b} \\ 0, & \text{otherwise.} \end{cases}$$

Here p_b denotes the precision in detecting the correct ranking of the relevant instances up to position b. MAP is then computed as the mean AP over all Q test set batches and is given by $MAP = \frac{1}{Q} \sum_{q=1}^{Q} AP(b)_q$.

- Spearman Rank Correlation (SRC): This computes the correlation in the rule generated rank with the expert ranking of each dataset.

These metrics quantify the goodness of ranking and hence the prioritization order. A high value (≈ 1) denotes acceptability of the rule generated ranking to experts and implies effectiveness of the rule based methodology in capturing expert preferences by identifying the most relevant instances.
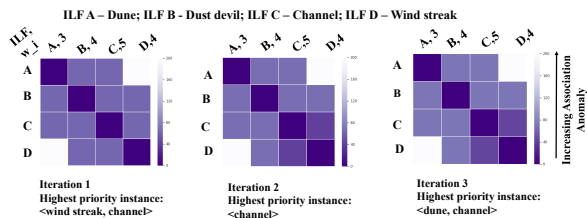
Figure 2: Inter-class association anomaly score (shown between two classes only) update over iterations and identification of novel associations.
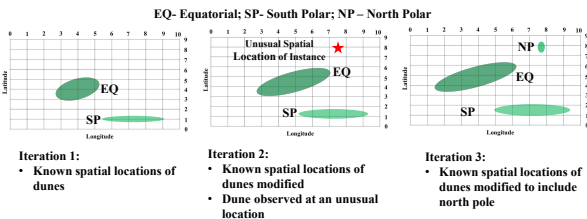

Figure 3: Identifying novel instance at a new spatial location of an ILF (dune) and updating existing spatial maps.

Dataset and Results

Dataset Description

The visible imager datasets from THEMIS has been used in this study (http://viewer.mars.asu.edu). Our dataset consists of all 1600 (single channel band 3) datasets collected from orbits 81 to 239, out of which 300 instances were identified by experts for its relevance due to presence of ILFs. Along with the spatial and seasonal attributes of these instances, experts annotated each instance to identify the interesting classes that appeared in it. We divided these annotated images into training and test sets. The training set consists of 99 relevant images from the first 71 orbits. This training dataset is used in extracting the rules which are then guided with expert feedback. The remaining 201 relevant images from the 87 subsequent orbits form the test set batches on which the rules are applied to rank it according to its scientific value. Each test set batch consists of eight images which are ranked and validated with expert acceptance of the ranking to evaluate the effectiveness of our approach. All instances from the test batch are used to update the rules, which is then used to rank the next batch.

Results

In this section we discuss the results obtained to evaluate the effectiveness of our proposed rule based methodology in identifying unusual instances and generating an expert acceptable ranking based on its relevance. Here, the test sets contain only expert identified instances with the ILFs from each orbit. We compare ranking from both stand-alone rules that are solely extracted from the data and the expert guided rules with that of expert ranking (from one expert) of the test set using the evaluation metrics discussed before. The expert ranking of the test set is only utilized to validate our
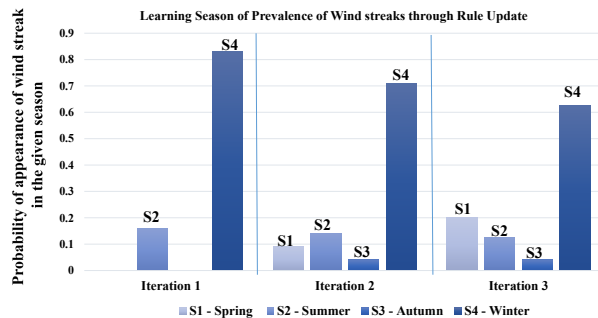

Figure 4: Learning Season of Prevalence of a given ILF (wind streak) through rule update.

Table 1: Acceptability of rule generated priority to experts using MAP and SRC over all test set batches.

Rule Forming Methodology	MAP	SRC
Standalone Rules	0.1864	0.5482
Expert Guided Rules	**0.765**	**0.984**

approach. Table 1 shows the average acceptability of the prioritization order to an expert over all test set batches using all attributes of the inference rules. The stand-alone rules score poorly using both metrics as these are uninformed of expert preferences. These rules give equal weight to all ILFs and inter-class associations. However, both MAP and SRC demonstrate the improvement in generating expert acceptable rank after incorporating expert knowledge. This is expected as the rules are aware of expert preferences. We also explain our decisions to increase interpretability by listing all attributes with respect to which an instance was novel.

We demonstrate the detection of novel instances by highlighting some examples of high priority instances identified in the test set batches. In Figure 2, the grids denote the association anomaly score of each ILF with every other ILF (only two-way associations are shown here). The corresponding expert weights of each ILF w_i, are shown as (ILF$_i$, w_i). Based on these association anomaly scores in iteration 1, the highest priority instance identified contains wind streak and channel, that have high class relevance and association anomaly scores. In the subsequent iterations, the relevance of this association therefore decreases, while the association relevance scores of unobserved associations such as dune and wind streak remains high. The proposed method also favors instances with channels as it has highest relevance to the expert. In Figure 3, we show the discovery of dunes at a location that was unknown based on previous observations. This discovery of dunes at the north pole then updates the existing planetary map for subsequent iterations. In Figure 4, we show an example of the seasonal rule update for wind streak. In iteration 1, the knowledge of predominant season is winter, which makes wind streak instances during autumn and summer in the next orbit/batch very unusual and hence is assigned high downlink priority. This is used to update the knowledge of predominant season which results in

the updated probabilities in iteration 2. By discovering different types of novel instances, the rules are updated such that highest preference is always assigned to instances that are most novel with respect to the attributes or are previously unobserved. This refinement of the existing knowledge base allows learning the ILF characteristics over time while informing experts about such updates.

Conclusion and Future Work

In this work, we presented an expert guided rule based methodology for identifying the most relevant images with interesting landforms and rank them in order of their scientific context. The rules thus formed represent the expected attributes of the interesting landforms and creates a knowledge repository from planetary images. We propose a priority inference approach where the relevance of an image is determined by its deviation from the formulated rules. We evaluate the effectiveness of our approach by determining how acceptable the prioritization order is to experts. As expected, the prioritization acceptability increases when the rules are guided by expert feedback as rules formulated solely from the data is uninformed of domain specific preferences. The proposed approach has the potential to operate as an independent remote module on planetary orbiters that assists domain experts by prioritizing downlinking of interesting instances allowed by the limited bandwidth along with adaptive region of interest aware compression techniques (Chakraborty et al. 2018). Another area where a similar approach can be effective is ground based retrieval of the most relevant instances for a phenomenon of interest to an expert, which are often rare in the repository.

In the future, we will focus on improving the knowledge base and consequently the prioritization approach by including additional attributes such as overlapping landforms and fraction of the instance occupied by ILFs. Additionally, for this module to function remotely, a landform classifier that summarizes the instance statistics is to precede the prioritization module. We will also extend it to subsequent years, evaluate our approach using feedback from multiple experts, which presents a foreseeable challenge as experts usually have varied research interests and create an interface to visualize the ranking order of novel instances in each orbit as well as explain the ranking.

Acknowledgements

References

Albuquerque, A.; Amador, T.; Ferreira, R.; Veloso, A.; and Ziviani, N. 2018. Learning to Rank with Deep Autoencoder Features. In *2018 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.

Arafeen, M. J., and Do, H. 2013. Test Case Prioritization Using Requirements-Based Clustering. In *Software Testing, Verification and Validation (ICST), 2013 IEEE Sixth International Conference on*, 312–321. IEEE.

Cantor, B. A.; James, P. B.; and Calvin, W. M. 2010. MARCI and MOC Observations of the Atmosphere and Surface Cap in the North Polar Region of Mars. *Icarus* 208(1):61–81.

Castano, R.; Wagstaff, K.; Song, L.; and Anderson, R. 2005. Validating Rover Image Prioritizations. *The Interplanetary Network Progress Report* 42:160.

Castano, R.; Wagstaff, K. L.; Chien, S.; Stough, T. M.; and Tang, B. 2007. On-board Analysis of Uncalibrated Data for a Spacecraft at Mars. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 922–930. ACM.

Castano, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic Detection of Dust Devils and Clouds on Mars. *Machine Vision and Applications* 19(5-6):467–482.

Chakraborty, S.; Banerjee, A.; Gupta, S. K.; and Christensen, P. R. 2018. Region of interest aware compressive sensing of themis images and its reconstruction quality. In *2018 IEEE Aerospace Conference*, 1–11. IEEE.

Estlin, T. A.; Bornstein, B. J.; Gaines, D. M.; Anderson, R. C.; Thompson, D. R.; Burl, M.; Castaño, R.; and Judd, M. 2012. AEGIS Automated Science Targeting for the MER Opportunity Rover. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3(3):50.

Kurte, K. R.; Durbha, S. S.; King, R. L.; Younan, N. H.; and Vatsavai, R. 2017. Semantics-enabled framework for spatial image information mining of linked earth observation data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 10(1):29–44.

Lee, A. J.; Hong, R.-W.; Ko, W.-M.; Tsao, W.-K.; and Lin, H.-H. 2007. Mining Spatial Association Rules in Image Databases. *Information Sciences* 177(7):1593–1608.

Paulin, M.; Douze, M.; Harchaoui, Z.; Mairal, J.; Perronin, F.; and Schmid, C. 2015. Local Convolutional Features with Unsupervised Training for Image Retrieval. In *Proceedings of the IEEE International Conference on Computer Vision*, 91–99.

Silva, R.; Gonçalves, M. A.; and Veloso, A. 2011. Rule-based Active Sampling for Learning to Rank. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 240–255. Springer.

Thompson, D. R.; Candela, A.; Wettergreen, D. S.; Dobrea, E. N.; Swayze, G. A.; Clark, R. N.; and Greenberger, R. 2018. Spatial Spectroscopic Models for Remote Exploration. *Astrobiology* 18(7):934–954.

Zhou, W.; Li, H.; and Tian, Q. 2017. Recent Advance in Content-based Image Retrieval: A Literature Survey. *arXiv preprint arXiv:1706.06064*.