Computers & Geosciences 101 (2017) 48–56

Contents lists available at ScienceDirect

Computers & Geosciences

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

Research paper

Automated detection of geological landforms on Mars using Convolutional Neural Networks

Leon F. Palafox[a,*], Christopher W. Hamilton[a], Stephen P. Scheidt[a], Alexander M. Alvarez[b]

[a] *Lunar and Planetary Laboratory, University of Arizona, Tucson, AZ, USA*
[b] *College of Medicine, University of Arizona, Tucson, AZ, USA*

ARTICLE INFO

Keywords:
Convolutional neural networks
Support vector machines
Volcanic rootless cones
Transverse aeolian ridges
Mars

ABSTRACT

The large volume of high-resolution images acquired by the Mars Reconnaissance Orbiter has opened a new frontier for developing automated approaches to detecting landforms on the surface of Mars. However, most landform classifiers focus on crater detection, which represents only one of many geological landforms of scientific interest. In this work, we use Convolutional Neural Networks (ConvNets) to detect both volcanic rootless cones and transverse aeolian ridges. Our system, named MarsNet, consists of five networks, each of which is trained to detect landforms of different sizes. We compare our detection algorithm with a widely used method for image recognition, Support Vector Machines (SVMs) using Histogram of Oriented Gradients (HOG) features. We show that ConvNets can detect a wide range of landforms and has better accuracy and recall in testing data than traditional classifiers based on SVMs.

1. Introduction

During the past ten years, the Mars Reconnaissance Orbiter (MRO) has collected over 30 Terabytes of data. Two of the cameras onboard MRO that are routinely used to study geological landforms include the High Resolution Imaging Science Experiment (HiRISE; 0.3 m/pixel resolution; McEwen et al., 2007) and the Context Camera (CTX; 6 m/pixel resolution; Malin et al., 2007). However, the total data volume of these images poses new challenges for the planetary remote-sensing community. For instance, each image includes limited metadata about its content, and it is time consuming to manually analyze each image to search for non-indexed information. Therefore, there is a need for computational techniques to search the HiRISE and CTX image databases and discover new content.

Many algorithms can classify image content, such as Support Vector Machines (SVMs) and logistic regression. Yet, most of these algorithms require pre-processing steps, like smoothing filters or Histogram of Oriented Gradients (HOG) methods (Dalal et al., 2005), which are typically tailored to address a specific classification problem. These pre-processing steps extract characteristics of the data, like edges in a picture, or patterns of illumination in a remote sensing scene. The signal processing and computer science communities refer to these characteristics as features. Convolutional Neural Networks (ConvNets) have become an increasingly popular alternative for image classification (LeCun, 2016), and compared with other classifiers, ConvNets have the best performance for recognition of both characters (Ciresan et al., 2012) and images (Graham, 2015). ConvNet architectures are the best performing algorithms in both the Mixed National Institute of Standards and Technology (MNIST) and Canadian Institute for Advanced Research (CIFAR) data sets, which are the standard classification data sets within the computer vision community. ConvNets learn their own input features, which alleviates the need to test different pre-processing algorithms. Furthermore, Graphical Processing Units (GPUs) can significantly increase the speed of training and classification steps in ConvNets. Using GPUs is not unique of ConvNets, and other Deep Learning architectures can also benefit from GPU acceleration.

In this paper, we address the problem of automated landform detection using ConvNets to identify Volcanic Rootless Cones (VRCs) and Transverse Aeolian Ridges (TARs) in two types of Mars satellite imagery by:

1. Training a ConvNet to detect landforms of varying size and shape, using VRCs as an example;
2. Showing that, for VRCs, a ConvNet performs better than optimized SVMs with HOG features; and
3. Showing that ConvNets also have the ability to detect a variety of other landforms, such as TARs.

Although our classifier is designed to detect many geologic features, the scope of the current study focuses on identifying VRCs and TARs as two examples of morphological distinct landforms, which are intended

* Corresponding author.

http://dx.doi.org/10.1016/j.cageo.2016.12.015
Received 18 October 2016; Accepted 28 December 2016
Available online 16 January 2017
0098-3004/

to highlight the broad applicability of our classifier to a wide range of geological classification problems.

2. Background information

2.1. Automated landform detection

Previous applications of machine learning in planetary sciences have typically focused on the automated detection of impact craters (Urbach and Stepinski, 2009; Bandeira et al., 2012; Stepinski et al., 2012; Emami et al., 2015; Cohen et al., 2016). Such Crater Detection Algorithms (CDAs) diminish the need for an operator to delimit manually all craters within a region, which is useful for generating impact crater inventories over large areas; however, manual inspection is still required to validate the results. The most popular CDAs first extract features from the data (e.g., shapes and patterns of light and shadow) and then apply a classifier (Stepinski et al., 2012). For instance, Urbach and Stepinski (2009) proposed a popular and efficient CDA, which applies a series of filters to remove the background noise and then creates a set of features that look for the characteristic crescent-shaped shadow of a crater. Bandeira et al. (2012) used the same approach, but added texture recognition to improve the precision of the algorithm. Cohen et al. (2016) showed preliminary results using a ConvNet for crater detection and demonstrated that they outperformed previously tested methods in the same dataset.

Aside from detecting impact craters, machine learning methods have only been used to identify a few other landforms in a planetary science context. These efforts include using Self Organizing Feature Maps (SOFMs) to identify VRCs in Mars Global Surveyor (MGS) Mars Orbiter Camera (MOC) imagery (Hamilton and Plug, 2004), applications of SVMs to detect dunes in MOC images (Bandeira et al., 2011), and object-based approaches to estimating the orientation of TARs with HiRISE data (Vaz and Silvestro, 2014). More recently, Palafox et al. (2015) and Scheidt et al. (2015) have also demonstrated the utility of ConvNets for detecting VRCs and TARs in HiRISE images. However, in general, little work has been done to develop generalized classifiers to detect other geological landforms using planetary remote sensing data—with the exception of the hazard navigation and automated rock analysis by robotic rovers on Mars. For instance, Gor and Castano (2001) designed an automated classifier to detect and analyze rocks for both of NASAs Mars Exploration Rovers (MERs) Spirit and Opportunity (Gor et al., 2001). Biesiadecki and Maimone (2006) also designed a self-navigation system using stereo matching and Random Sample Consensus (RANSAC) algorithms, and used these algorithms to estimate the position of the rover by identifying landmarks in the image data (Biesiadecki and Maimone, 2006).

2.2. The characteristics and geological significance of VRCs and TARs

Volcanic Rootless Cones (VRCs) are generated by explosive interactions between lava and external sources of water (Thorarinsson, 1951, 1953), and are commonly associated with the flow of lava into marshes, lacustrian basins, littoral environments, glacial outwash plains, snow, and ice. Terrestrial VRCs cover areas of up to ~150 km^2 and generally include numerous cratered cones ranging from 1 to 35 m in height and ~2–500 m in diameter (Fagents and Thordarson, 2007). VRCs on Mars (Fig. 1) are generally larger, typically ranging from tens of meters to ~1 km in diameter, and can form groups covering thousands of square kilometers (Hamilton et al., 2010a, 2010b, 2011). Rootless cone morphologies and spatial organization strongly depend upon lava emplacement processes (Hamilton et al., 2010a, 2010c) and a balance between the availability and utilization of lava (fuel) and groundwater (coolant) in molten fuel–coolant interactions (MFCIs; Sheridan and Wohletz, 1981, 1983; Wohletz, 1983, 1986, 2002; Zimanowski et al., 1991; Zimanowski, 1998). However, in the presence of excess lava (e.g., in regions inundated by large sheet-like flows of molten lava), it may be assumed that the location of VRC groups will strongly depend on the distribution of near-surface H_2O and that VRCs may be used a proxy for former H_2O deposits (Frey et al., 1979; Frey and Jarosewich, 1982; Greeley and Fagents, 2001; Fagents and Thordarson, 2007; Head and Wilson, 2002; Fagents et al., 2002; Jaeger et al., 2007; Hamilton et al., 2010a, 2010c, 2011). Cratered cones, resembling terrestrial VRCs, have been identified in many regions on Mars (Fagents and Thordarson, 2007) and their widespread occurrence makes them important as a paleoenvironmental indicator that can be used to infer the locations of near-surface H_2O at the time of lava flow emplacement.

Wind plays a significant role in shaping the surface of Earth and Mars by moving small particles to generate a variety of depositional and erosional features. Aeolian bedforms include ripples and dunes, as well as a distinct class of bedforms termed Transverse Aeolian Ridges (TARs) (Bourke et al., 2003). TARs occur in the equatorial and mid-latitude regions of Mars (Balme et al., 2008; Berman et al., 2011), but it is uncertain whether or not they form by ripple- or dune-forming processes. It is clear that many martian TARs are constructional landforms, resulting from the transport and deposition of granular material, alternative hypotheses have been proposed for some examples. For instance, Montgomery et al. (2012) explain several TAR-like features on Mars as periodic bedrock ridges, which are erosional landforms with crests that are transverse to the prevailing wind direction (Greeley et al., 1992; Hugenholtz et al., 2015). These contrasting interpretations carry different implications for surface–atmospheric interactions on Mars and the deposition, or erosion, of sedimentary units through time. Mapping the spatial distribution of TARs over regional and global scales could provide important new constraints for their formation processes, but their small size and widespread distribution makes automated approaches to TAR identification preferable to manual mapping efforts.

3. Methods

3.1. Support Vector Machines (SVMs)

In planetary remote sensing, SVMs have been used to detect impact craters on the Moon (Burl, 2000) and to study volcanic landforms on Venus (Burl, 2001; Decoste and Schölkopf, 2002). SVM algorithms use a function, known as a kernel, to create a decision boundary that separates data into distinguishable classes (Boser et al., 1992; Hastie et al., 2009). In remote sensing, these kernels become especially important as objects from different classes may have overlapping characteristics.

Our SVM classifier uses Histogram of Oriented Gradients (HOG) features to accentuate landforms in HiRISE and CTX images. In the HOG transformation, a series of oriented gradients—discrete angles between 0 and 360°—are drawn in small, adjacent non-overlapping units. A histogram representing the number of elements in line with these oriented gradients is created for each unit and depicted as an intensity vector in that unit. An array of HOG features representing the linear landforms of an image can provide additional information beyond the original data set. HOG is very robust to changes in illumination and shadowing, which is a desirable characteristic in a landform detection algorithm (Dalal et al., 2005).

3.2. Convolutional Neural Networks (ConvNets)

Artificial Neural Networks (ANNs) are composed of connected set of linear classifiers, each of which is trained to generate a specific decision boundary and classify simple spaces. Layers within an ANN are connected in sequential order, such that the input of a layer is the output of the previous one. Traditionally, ANNs have an input layer, which receives the input data; a set of hidden layers, which serve as the classifier; and an output layer that provides the result of the classifica-