


Getting a feel for the terrain

June 28 2012, by Michael Schirber



A TextureCam analysis of a Mars image is able to distinguish rocks from soil.
Credit: NASA/JPL/Caltech/Cornell

A team is developing a rover-based computer system that can identify geologically relevant elements in a camera scene.

It's a hot summer day, and your eyes spot an ice cream cart up ahead. Without even really thinking, you start walking that direction. Planetary scientists would like to give robots that kind of visual recognition---not for getting ice cream, but for finding scientifically interesting targets.


Currently, rovers and other [space vehicles](#) are still largely dependent on commands from their human controllers back on Earth. But to decide what commands to send, operators must wait to receive images and other

pertinent information from the spacecraft. Because rovers don't have powerful antennas, this so-called downlink usually takes a lot of time.

The data bottleneck means rovers often "twiddle their thumbs" between subsequent commands.

"Our goal is to make smart instruments that can do more within each command cycle," says David Thompson of the Jet Propulsion Laboratory in Pasadena, Calif.

Thompson is heading a project called TextureCam, which involves creating a computer vision package that can map a surface by identifying [geological features](#). It is primarily envisioned for a rover, but it could also benefit a spacecraft visiting an asteroid or an aerobot hovering in the atmosphere of a distant world.



A photo of a stromatolite (left) from Western Australia analyzed by TextureCam (right). The program assigns a color to each patch in the image according to how it matches the criteria for stromatolite rocks (red means good match, or high probability). Credit: NASA/JPL

With funds from NASA's [Astrobiology](#) Science and Technology for Exploring Planets (ASTEP), Thompson's team is currently refining their [computer algorithm](#), with an eventual plan to build a prototype instrument that can map an astrobiologically-relevant field site.

Roam rover, roam rover

Rovers have already made great advances in autonomy. Current prototypes can travel as much as a kilometer on their own using on-board navigation software. This allows these vehicles to cover a much larger territory.

But one concern is that a rover may literally drive over a potentially valuable piece of scientific real estate and not even realize it. Giving a rover some rudimentary visual identification capabilities could help avoid missing "the needle in the haystack," as Thompson refers to the hidden clues that astrobiologists hope to uncover on other planets.

"If the rover can make simple distinctions, we can speed up the reconnaissance," he says. As it drives along, the rover could snap several images and use on-board software to prioritize which images to downlink to Earth.

And while waiting for its next set of commands, it could pick a potentially interesting geological feature and then drive up close to take a detailed picture or even perform some simple chemical analysis.

"You could start the next day with the instrument sitting in front of a prime location," Thompson says.

Instead of spending time trying to get the rover from point A to point B, mission controllers could concentrate on doing the higher level scientific investigation that the rover can't do. At least, not yet.

"The field being investigated by David Thomson is vital to cope with the flood of remote sensing data returned from spacecraft," says Anthony Cook of Aberystwyth University in the UK, who is not involved with TextureCam.



The Zoë rover driving through the Atacama Desert in Chile in 2005. Credit: Carnegie Mellon University

There are a other projects working on computer vision for rovers. In 2010, the Mars rover Opportunity received a software upgrade called AEGIS that can identify scientifically interesting rocks. A project in the Atacama desert in Chile used a similar rock detector system on its rover called Zoë. And ESA's ExoMars mission is developing computer vision that can detect objects in the rover's vicinity.

TextureCam is unique from these other efforts in that it is mapping the surface, rather than trying to isolate particular objects. It's a more general strategy that can identify terrain characteristics, such as weathering or fracturing.

Recognizing a rock face

The new approach by Thompson's group focuses on the "texture" of an image, which is computer vision terminology for the statistical patterns that exist in an array of pixels. The same kind of image analysis is being used in more common day-to-day applications.

For example, the web is inundated with huge photo archives that haven't been sorted in any systematic way. Several companies are developing "search engines" that can identify objects in digital images. If you were looking for, say, an image with a "blue dog" or a "telephone booth," these programs could sift through a collection of photos to find those that match the particular criteria.

Additionally, many digital cameras detect faces in the camera frame and automatically adjust the focus depending on how far away the faces are. And some new video game consoles have sensors to detect the bodily pose of a game player.

What all these technologies have in common is a sophisticated analysis of image pixels. The relevant software programs typically look for signals in the variations of brightness or the shades of color that are characteristic of a telephone or a face or a rock.



A concept image for the ExoMars rover that is being developed for a 2018 mission to Mars. Credit: ESA

These signals often have little to do with the way we might describe these objects.

"The software identifies statistical properties that might not be obvious to the human eye," Thompson says.

Let the computer do the guesswork

In the case of TextureCam, the computer program takes a small patch, or thumbnail, inside the image and performs a number of different pixel-to-pixel comparisons. Which comparisons? Actually, the computer decides.

"We train the system from examples," Thompson explains. They take

images that were previously analyzed by a geologist as having an outcrop or a sediment or a rock of a particular variety. The computer program compares its pixel analysis to these labels and builds a decision tree (or a more elaborate "decision forest") that best discriminates between the different possibilities.

"These decision trees can be quite efficient even after just a few branches," Thompson says.

This so-called "machine learning" has advantages over other techniques that construct a visual model of what the computer should be looking for.

"The disadvantage with visual models is that you have to build a new rule for every new thing you want to identify," Thompson says. It can be hard for humans to find reliable distinctions that can help a computer. It makes more sense to let the computer go out and explore the possibilities with trial and error.

"The system trains itself, so we don't have to anticipate," Thompson says.

The "training regimen" for TextureCam began with a set of images from Mars and is now moving onto photos from the Mojave Desert.

The team plans to integrate their algorithm into a field programmable gate array (FPGA), which is basically a special purpose computer that would connect directly to a rover camera. This would allow TextureCam to work faster, without relying on the rover's main computer.

"Computers and software are not ready to take over the interpretation tasks of human geologists, but they will help to pre-sort and pre-identify regions of interest, thus reducing the amount of remote sensing data that geologists must examine," Cook says.

Source: Astrobio.net

Citation: Getting a feel for the terrain (2012, June 28) retrieved 20 March 2024 from <https://phys.org/news/2012-06-terrain.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.