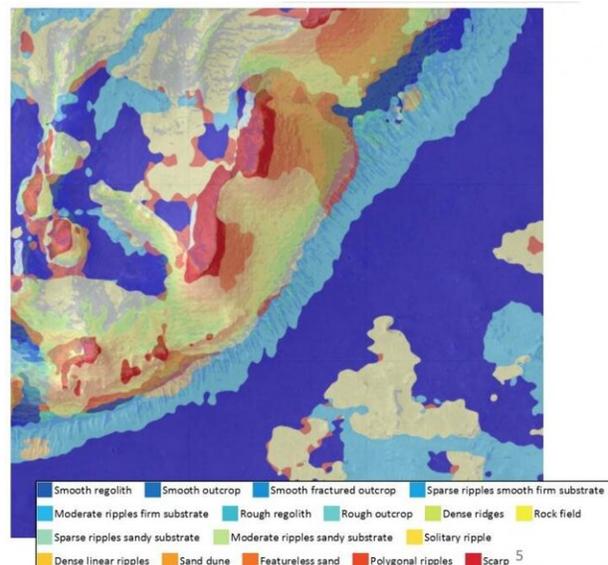
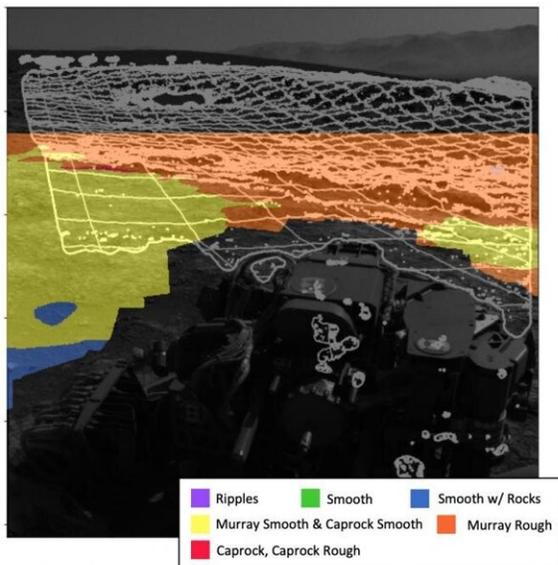


Deep learning will help future Mars rovers go farther, faster, and do more science

August 19 2020



The Machine Learning-based Analytics for Autonomous Rover Systems (MAARS) program encompasses a range of areas where artificial intelligence could be useful. The team presented results of the MAARS project at IEEE Aerospace Conference in March 2020. The project was a finalist for the NASA Software Award. Credit: NASA JPL

NASA's Mars rovers have been one of the great scientific and space successes of the past two decades.

Four generations of rovers have traversed the red planet gathering

[scientific data](#), sending back evocative photographs, and surviving incredibly harsh conditions—all using on-board computers less powerful than an iPhone 1. The latest [rover](#), Perseverance, was launched on July 30, 2020, and engineers are already dreaming of a future generation of rovers.

While a major achievement, these missions have only scratched the surface (literally and figuratively) of the planet and its geology, geography, and atmosphere.

"The surface area of Mars is approximately the same as the total area of the land on Earth," said Masahiro (Hiro) Ono, group lead of the Robotic Surface Mobility Group at the NASA Jet Propulsion Laboratory (JPL)—which has led all the Mars rover missions—and one of the researchers who developed the software that allows the current rover to operate.

"Imagine, you're an alien and you know almost nothing about Earth, and you land on seven or eight points on Earth and drive a few hundred kilometers. Does that alien species know enough about Earth?" Ono asked. "No. If we want to represent the huge diversity of Mars we'll need more measurements on the ground, and the key is substantially extended distance, hopefully covering thousands of miles."

Traveling across Mars' diverse, treacherous terrain with limited computing power and a restricted energy diet—only as much sun as the rover can capture and convert to power in a single Martian day, or sol—is a huge challenge.

The first rover, Sojourner, covered 330 feet over 91 sols; the second, Spirit, traveled 4.8 miles in about five years; Opportunity, traveled 28 miles over 15 years; and Curiosity has traveled more than 12 miles since it landed in 2012.

"Our team is working on Mars robot autonomy to make future rovers more intelligent, to enhance safety, to improve productivity, and in particular to drive faster and farther," Ono said.

New Hardware, New Possibilities

The Perseverance rover, which launched this summer, computes using RAD 750s—radiation-hardened single board computers manufactured by BAE Systems Electronics.

Future missions, however, would potentially use new high-performance, multi-core radiation hardened processors designed through the High Performance Spaceflight Computing (HPSC) project. (Qualcomm's Snapdragon processor is also being tested for missions.) These chips will provide about one hundred times the computational capacity of current flight processors using the same amount of power.

"All of the autonomy that you see on our latest Mars rover is largely human-in-the-loop"—meaning it requires human interaction to operate, according to Chris Mattmann, the deputy chief technology and innovation officer at JPL. "Part of the reason for that is the limits of the processors that are running on them. One of the core missions for these new chips is to do [deep learning](#) and machine learning, like we do terrestrially, on board. What are the killer apps given that new computing environment?"

The Machine Learning-based Analytics for Autonomous Rover Systems (MAARS) program—which started three years ago and will conclude this year—encompasses a range of areas where artificial intelligence could be useful. The team presented results of the MAARS project at IEEE Aerospace Conference in March 2020. The project was a finalist for the NASA Software Award.

"Terrestrial high performance computing has enabled incredible breakthroughs in autonomous vehicle navigation, machine learning, and data analysis for Earth-based applications," the team wrote in their IEEE paper. "The main roadblock to a Mars exploration rollout of such advances is that the best computers are on Earth, while the most valuable data is located on Mars."

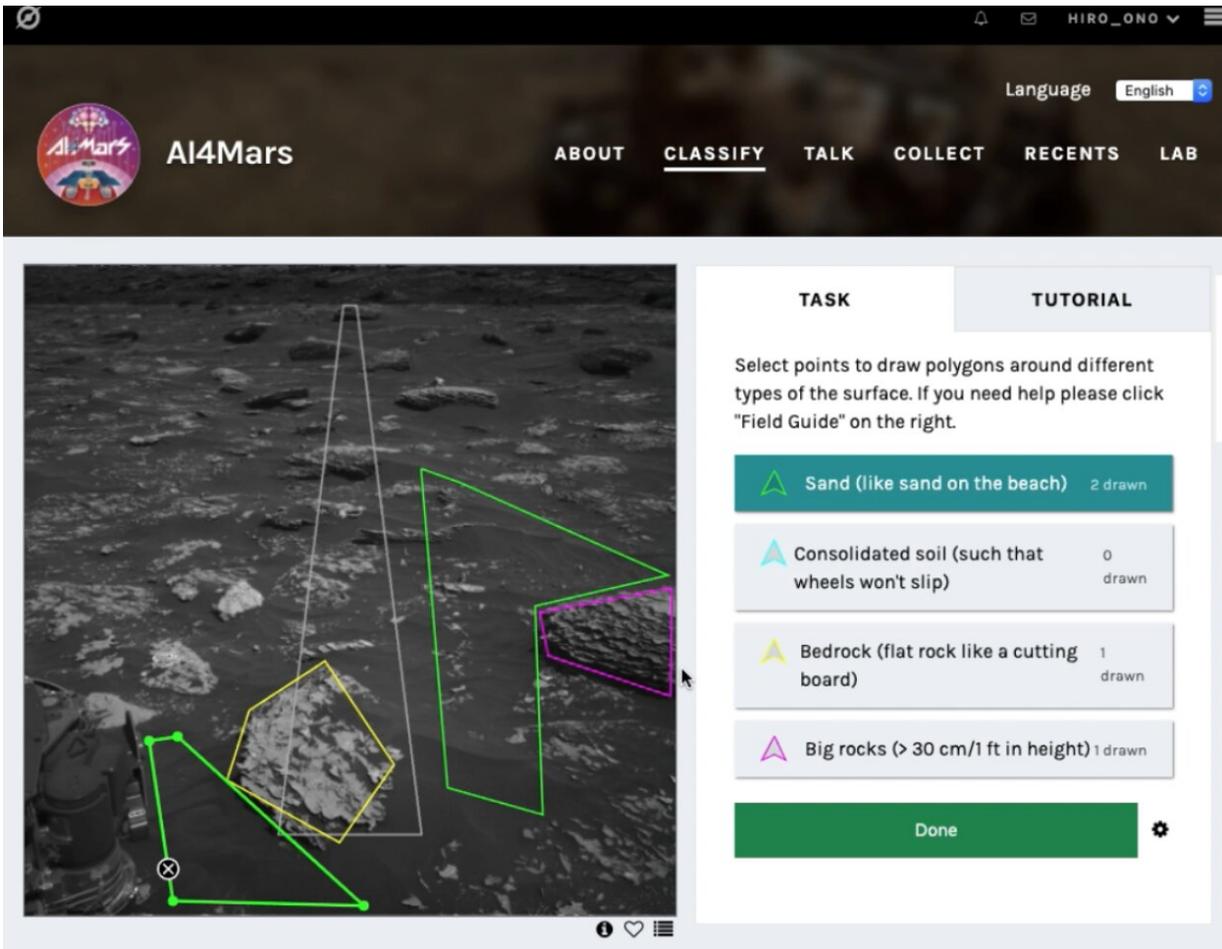
Training machine learning models on the Maverick2 supercomputer at the Texas Advanced Computing Center (TACC), as well as on Amazon Web Services and JPL clusters, Ono, Mattmann and their team have been developing two novel capabilities for future Mars rovers, which they call Drive-By Science and Energy-Optimal Autonomous Navigation.

Energy-Optimal Autonomous Navigation

Ono was part of the team that wrote the on-board pathfinding software for Perseverance. Perseverance's software includes some machine learning abilities, but the way it does pathfinding is still fairly naïve.

"We'd like future rovers to have a human-like ability to see and understand terrain," Ono said. "For rovers, energy is very important. There's no paved highway on Mars. The drivability varies substantially based on the terrain—for instance beach versus bedrock. That is not currently considered. Coming up with a path with all of these constraints is complicated, but that's the level of computation that we can handle with the HPSC or Snapdragon chips. But to do so we're going to need to change the paradigm a little bit."

Ono explains that new paradigm as commanding by policy, a middle ground between the human-dictated: "Go from A to B and do C," and the purely autonomous: "Go do science."



The public tool AI4Mars shows different kinds of Martian terrain as seen by NASA's Curiosity rover. By drawing borders around terrain features and assigning one of four labels to them, users can help train an algorithm that will automatically identify favorable and hazardous terrain for Curiosity's rover planners. Credit: NASA/JPL-Caltech

Commanding by policy involves pre-planning for a range of scenarios, and then allowing the rover to determine what conditions it is encountering and what it should do.

"We use a supercomputer on the ground, where we have infinite computational resources like those at TACC, to develop a plan where a policy is: if X, then do this; if y, then do that," Ono explained. "We'll basically make a huge to-do list and send gigabytes of data to the rover, compressing it in huge tables. Then we'll use the increased power of the rover to de-compress the policy and execute it."

The pre-planned list is generated using machine learning-derived optimizations. The on-board chip can then use those plans to perform inference: taking the inputs from its environment and plugging them into the pre-trained model. The inference tasks are computationally much easier and can be computed on a chip like those that may accompany future rovers to Mars.

"The rover has the flexibility of changing the plan on board instead of just sticking to a sequence of pre-planned options," Ono said. "This is important in case something bad happens or it finds something interesting."

Drive-By Science

Current Mars missions typically use tens of images a Sol from the rover to decide what to do the next day, according to Mattmann. "But what if in the future we could use one million image captions instead? That's the core tenet of Drive-By Science," he said. "If the rover can return text labels and captions that were scientifically validated, our mission team would have a lot more to go on."

Mattmann and the team adapted Google's Show and Tell software—a neural image caption generator first launched in 2014—for the rover missions, the first non-Google application of the technology.

The algorithm takes in images and spits out human-readable captions.

These include basic, but critical information, like cardinality—how many rocks, how far away?—and properties like the vein structure in outcrops near bedrock. "The types of science knowledge that we currently use images for to decide what's interesting," Mattmann said.

Over the past few years, planetary geologists have labeled and curated Mars-specific image annotations to train the model.

"We use the one million captions to find 100 more important things," Mattmann said. "Using search and information retrieval capabilities, we can prioritize targets. Humans are still in the loop, but they're getting much more information and are able to search it a lot faster."

Results of the team's work appear in the September 2020 issue of *Planetary and Space Science*.

TACC's supercomputers proved instrumental in helping the JPL team test the system. On Maverick 2, the team trained, validated, and improved their model using 6,700 labels created by experts.

The ability to travel much farther would be a necessity for future Mars rovers. An example is the Sample Fetch Rover, proposed to be developed by the European Space Association and launched in late 2020s, whose main task will be to pick up samples dug up by the Mars 2020 rover and collect them.

"Those rovers in a period of years would have to drive 10 times further than previous rovers to collect all the samples and to get them to a rendezvous site," Mattmann said. "We'll need to be smarter about the way we drive and use energy."

Before the new models and algorithms are loaded onto a rover destined for space, they are tested on a dirt training ground next to JPL that serves

as an Earth-based analog for the surface of Mars.

The team developed a demonstration that shows an overhead map, streaming images collected by the rover, and the algorithms running live on the rover, and then exposes the rover doing terrain classification and captioning on board. They had hoped to finish testing the new system this spring, but COVID-19 shuttered the lab and delayed testing.

In the meantime, Ono and his team developed a citizen science app, AI4Mars, that allows the public to annotate more than 20,000 images taken by the Curiosity rover. These will be used to further train [machine learning](#) algorithms to identify and avoid hazardous terrains.

The public have generated 170,000 labels so far in less than three months. "People are excited. It's an opportunity for people to help," Ono said. "The labels that people create will help us make the rover safer."

The efforts to develop a new AI-based paradigm for future autonomous missions can be applied not just to rovers but to any autonomous space mission, from orbiters to fly-bys to interstellar probes, Ono says.

"The combination of more powerful on-board computing power, pre-planned commands computed on high performance computers like those at TACC, and new algorithms has the potential to allow future rovers to travel much further and do more science."

More information: Dicong Qiu et al, SCOTI: Science Captioning of Terrain Images for data prioritization and local image search, *Planetary and Space Science* (2020). [DOI: 10.1016/j.pss.2020.104943](https://doi.org/10.1016/j.pss.2020.104943)

Provided by University of Texas at Austin

Citation: Deep learning will help future Mars rovers go farther, faster, and do more science (2020, August 19) retrieved 26 April 2024 from <https://phys.org/news/2020-08-deep-future-mars-rovers-faster.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.