AUTONOMOUS SOIL ASSESSMENT SYSTEM: A DATA-DRIVEN APPROACH TO PLANETARY MOBILITY HAZARD DETECTION

K. Raimalwala[1], M. Faragalli[1] and E. Reid[1], [1]Mission Control Space Services Inc. (1125 Colonel By Drive, 311 St. Patrick's Building, Ottawa, ON K1S 5B6; kaizad@missioncontrolspaceservices.com, michele@missioncontrolspaceservices.com, ewan@missioncontrolspaceservices.com).

Motivation: Planetary rover driving operations are conducted by strategic path planning using orbital assets while tactically accounting for hazards and science targets within the field of view of the rover. Current strategies include a variety of automated and manual processes to plan and adapt rover driving paths while leveraging a science and operations backroom to support tactical decision making. Tactically, science team members are engaged to provide daily advice to rover planners to determine safest routes using imagery downlinked such that new commands can be uplinked to the rover for the next sol [Arvidson et al. 2017]. The rover's actual drive path deviates from the planned route to avoid hazards and obstacles not evident in the map derived from orbital data. Pre-set thresholds for risks on distinct terrain classes, based on terrestrial testing and experience, are used to stop the rover's traverse and await further commands from the ground [1]. For example, slip thresholds are set where traverses are automatically stopped if slip exceeds a defined threshold. For Curiosity's traverse from Yellowknife Bay to Mount Sharp, a route was planned using HiRISE orbiter imagery to avoid bedrock-dominated slopes greater than 25°, sandy slopes steeper than 12.5°, aeolian dunes and ripples and rocks taller than one wheel diameter [2]. Figure 1 depicts the geomorphic map along Curiosity's traverse between Yellowknife Bay and the base of Mount Sharp.

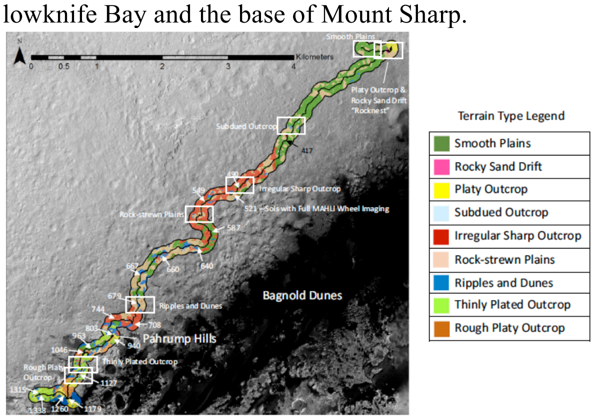

Figure 1: Geomorphic map of Curiosity traverse from Yellowknife Bay to Mount Sharp [1]

However, a limitation in this approach is that it is difficult to rapidly and accurately classify terrain types and hazards for tactical operations – particularly when encountering new terrain types. For example, Curiosity sustained significant wheel damage upon traversing rock-strewn terrain. As a result of punctures and tears, the remaining wheel lifetime is estimated at 10km motivating more careful path planning. To avoid further damage, Curiosity was directed towards megaripples (windblown, sand-sized deposit covered by coarser grains) to cushion wheel loads. However, the megaripples led to unexpected mobility difficulties, with high sinkage (approximately 30% of the wheel diameter) and high slip (up to 77%). Neither the rocky plateaus nor the megaripples were initially thought to pose mobility hazards to the rover by tactical operations teams and several sols were lost due to tactical planning and analysis.

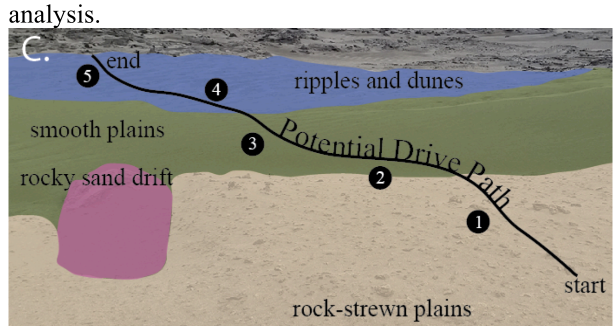

Figure 2: Image of classified terrain from Curiosity Navigation Camera (Navcam) on Sol 664 based on classes by [1].

The gap between rover navigation by human input and automated hazard detection and avoidance is filled increasingly by machine learning algorithms, which seek to make human-like geologic terrain observations [3,4]. These algorithms associate proprioceptive trafficability measurements and exteroceptive terrain sensing to make basic geologic terrain interpretations (e.g., sand vs. gravel) and predict terrain traversability ahead of the rover automatically. Because automated geologic terrain maps have not been extensively used in planetary rover operations, the efficiency and capability of operations workflows that incorporate such data have not been tested. We have tested this emerging protocol using the machine-learning terrain analysis software described below.

Autonomous Soil Assessment System Overview: The Autonomous Soil Assessment System (ASAS) is a software tool developed to predict non-geometric mobility hazards in the rover's field of view. ASAS uses a rover's navigation sensors to measure its mobility performance as it drives, and relates it to visual terrain features of the terrain. By fusing exteroceptive and proprioceptive data through machine learning algorithms, i.e. associating "feeling" with "seeing", ASAS learns to predict non-geometric hazards such as soft sand along the rover's path. In doing so, ASAS offers increased autonomy and efficiency in tactical workflows for rover navigation. Its three key components are:

Data-driven terramechanics modeling. Classical physics-based models have been used in the past to relate terrain mechanical properties such as cohesion and internal angle of friction to rover trafficability properties such as slip and sinkage. Such models rely on the accuracy of empirically derived soil parameters and are restrictive in structure, unable to capture unmodeled phenomena.

Our data-driven approach stores pairs of terrain slope and rover slip in finite queues and assigns heuristic hazard levels. This allows inference of trafficability from empirical data without strict model enforcement, offering a more robust solution over estimating soil properties and modeling wheel-soil mechanics. Additionally, this approach allows for online adaptability whereby the model is updated in real-time. Separate terramechanics models can be developed for discrete terrain classes; sand, gravel, and bedrock were chosen for our testing. Another advantage of the data-driven approach is that a new terrain class can be inferred from sufficient and consistent deviation of trafficability from existing models.

Automated terrain classification. Separating terramechanics models by terrain type requires the ability to classify that type. This is done by a supervised classification algorithm that uses images from the rover's stereo camera. Prior to running ASAS operationally, this algorithm is trained on images captured during a 'training phase', of terrain similar to what would be traversed in operation.

Real-time proprioceptive and exteroceptive data fusion and prediction. ASAS works by correlating what the rover sees to what it experiences, and then leveraging that correlation to predict what it would experience given what it sees ahead. Given a trained terrain classifier, ASAS identifies the terrain type within the field of view ahead of the rover. Once the rover reaches the area classified, with pose being estimated by localization techniques, ASAS stores the rover's slope angle (from its IMU) and instantaneous slip in the terramechanics model for that terrain class identified previously. This fusion of exteroceptive and proprioceptive information builds the data-driven terramechanics model over time. In operation, ASAS then identifies the terrain class and estimates the slope of an area ahead of the rover, and uses that information to predict rover slip and hazard level.

Operations, Field Tests and Demonstrations. ASAS has been developed to Technology Readiness Level 4 and field-tested at several sites in Ontario and Quebec, Canada, including the Canadian Space Agency's Mars Emulation Terrain, a.k.a. *Mars Yard*. Our platform is the Argo J5 built by Ontario, Drive & Gear and employs TRL 6 flexible metallic wheels. Additional field tests will take place at the dunes of White Sands, NM, in the winter of 2018, to test ASAS in a high-fidelity analogue environment consisting of aeolian landforms and duricrust-like features. Results from these tests will be presented at PSIDA-2018.

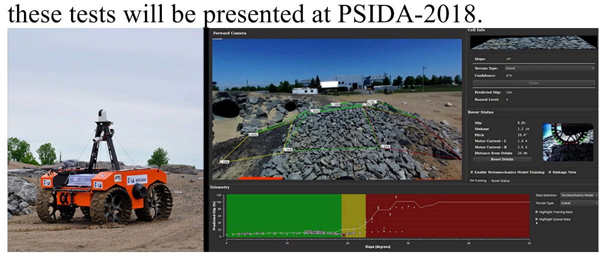

Figure 3: The Argo J5 rover and a screenshot of the ASAS user interface during a demonstration at the CSA Mars Yard.

Concept of Operations. Demonstrations of ASAS have involved a user interface through which an operator can monitor live information and train the terrain classifier and terramechanics model as the rover is driven manually in a low-latency teleoperation (LLT) environment (see Figure 3). The operator can use ASAS to inform their driving decisions based on hazards predicted ahead of the rover.

ASAS is designed to 'learn' to predict mobility hazards, thus improving the efficiency of tactical workflows in rover navigation and increasing the expected rover lifetime. It can be configured to be used as a path planning aid to suit an operation characterized by varying degrees of rover autonomy and command latencies. In the high-latency teleoperations (HLT) environment that characterizes Mars rover operations, ASAS incorporates non-geometric obstacle avoidance in autonomous navigation. The current strategy of stopping the rover when slip thresholds are reached can be augmented with ASAS' capability to predict that slip before the rover even reaches that hazardous terrain. For LLT of a lunar rover, an operator can use ASAS to aid their driving decisions in real-time.

References: [1] Arvidson R.E. et al. (2017) *J. Terramech., 73-93*. [2] Arvidson et. Al., (2014) *J. Geophysical Reaserch: Planets, 1322-1344*. [3] Gonzalez R. et al. (2017) *J. Field Robotics, 1-17*. [4] Wong C. et al. (2017) *IEEE Conference on Adaptive Hardware and Systems (AHS), 237-244*.