NOTE: This document was prepared by a team that participated in the 2018 Workshop on Autonomy for Future NASA Science Missions. It is for informational purposes to inform discussions regarding the use of autonomy in notional science missions and does not specify Agency plans or directives.

2018 Workshop on Autonomy for Future NASA Science Missions: Mars Design Reference Mission Reports

Table of Contents

Introduction 2
The Mars Design Reference Mission Report 3
Mars Design Reference Mission Report Summary 17

NOTE: This document was prepared by a team that participated in the 2018 Workshop on Autonomy for Future NASA Science Missions. It is for informational purposes to inform discussions regarding the use of autonomy in notional science missions and does not specify Agency plans or directives.

Introduction

Autonomy is changing our world; commercial enterprises and academic institutions are developing and deploying drones, robots, self-driving vehicles and other autonomous capabilities to great effect here on Earth. Autonomous technologies will also play a critical and enabling role in future NASA science missions, and the Agency requires a specific strategy to leverage these advances and infuse them into its missions. To address this need, NASA sponsored the 2018 Workshop on Autonomy for NASA Science Missions, held at Carnegie Mellon University, October 10-11, 2018.

The Workshop goals included:

- Identifying emerging autonomy technologies (10-15 years) that will:
    - Enable or enhance mission capabilities
    - Reduce risk
    - Reduce cost
- Identifying potential collaborations, partnerships, or linkages involving government, industry, and/or academia to enable these technologies

Capturing crosscutting autonomy technology requirements for future NASA missions
Over 90 individuals from industry, academia, and NASA participated in the workshop, which included presentations by keynote speakers, panel discussions, and small group discussions.

To provide structure for workshop discussions and post-workshop analysis, NASA established eight teams to examine the following Design Reference Mission (DRM) areas: Astrophysics, Earth Science, Heliophysics, Mars, Moon, Ocean Worlds, Small Bodies and Venus. Each DRM team was led by a scientist and a technologist, and team members consisted of workshop participants with relevant experience and interest. NASA asked each team to develop one or more mission scenarios that would be enabled by infusion of autonomous technology. The Agency provided guidance to support these team discussions; in particular, NASA urged the DRM teams to “think out of the box” and to consider bold missions that would be enabled by autonomous technology to provide valuable science results. Each DRM team developed mission scenarios that included defined science objectives, capability and technology needs, system requirements, and a concept of operations. Teams also identified gaps where autonomy technologies and other supporting technologies need to be developed and/or infused to enable each mission.

The DRM teams conducted small group discussions at the workshop and then presented a summary of their findings to all workshop attendees. Each DRM team continued to refine its mission scenarios after the workshop, creating both a full report and a summary report to document team findings. DRM teams also reported results at the December 2019 meeting of the American Geophysical Union.

This document contains the full report and summary report generated by the Mars DRM team. Full and summary reports generated by all eight DRM teams, plus a summary of workshop results are available online.

NOTE: This document was prepared by a team that participated in the 2018 Workshop on Autonomy for Future NASA Science Missions. It is for informational purposes to inform discussions regarding the use of autonomy in notional science missions and does not specify Agency plans or directives.

The Mars Design Reference Mission Report

Part I: Abstract

Mars is special. It is our closest planetary neighbor and shares commonalities with Earth. NASA has studied Mars more than any other solar system object outside the Earth and Moon. The scientific exploration of Earth's planetary neighbor has largely focused on addressing the presence and persistence of water, geochemistry, geology, and atmospheric evolution. Prior, current, and near-term missions are filling in fundamental Mars knowledge gaps and in doing so, support models of how the Mars planetary system functions and has evolved. These missions also take the first steps necessary for addressing whether or not Mars ever hosted microbial life. However, in situ data collections are limited to singular spacecraft in singular localities. All but one (the European Space Agency's ExoMars mission) are largely limited to a surface investigation. Past, current, and near-term mission architectures, while critical for exploration on a broad scale via multiple missions, do not support the system-level understanding of processes and conditions at regional scales.

The time has come for a paradigm shift. Sustained, wide-area study is needed to take the next step: to explore Mars as a system. This document describes not a single mission, but a practical, scalable and sustainable **Mars Exploration Campaign** that establishes an exploration framework on Mars. In this framework, new spacecraft, new rovers, and missions themselves become new elements within the campaign's framework.

Mars is expected to be the first destination for humans beyond the Moon. The human exploration zone will be regional in scale (~100-km radius). It is expected that humans will investigate, utilize in situ resources, and change the environment at this scale. Establishing in-depth knowledge of the surrounding environment, from subsurface to atmosphere, may be critical to the success of human missions at Mars. This Design Reference Mission (DRM) describes a practical mission that precedes human exploration and provides a detailed reconnaissance survey that will support initial human activities and provide an informational, infrastructural, and operational foundation for sustained human-robotic activities. The infrastructure is scalable (spatial), mission-extendable (time), and extensible to other missions (integration and growth).

As the foundational mission in the Mars Exploration Campaign, this Mars DRM aims to study the ground-water ice in the context of climate and regional geology, local weather, and possible biology while also providing detailed insight on the location and potential exploitation of subsurface water on Mars. These aims address ***NASA's 2018 strategic plan*** [1] by specifically addressing: Objective 1.1 to understand the Solar System, in particular with respect to searching for life elsewhere; preparing for Objective 2.2 to "conduct human exploration in Deep

Space…"; and Object 4.6 paving a path forward to establishing sustainable infrastructure capabilities and operations on Mars. The DRM also addresses the three high-priority science goals for the exploration of Mars as described in the current ***Planetary Decadal Survey [2]***: "Understand the processes and history of climate," "Determine if life ever arose on Mars," and "Determine the evolution of the surface and interior." The crosscutting nature of this DRM effectively addresses all four goals of the ***Mars Exploration Program Analysis Group's 2015 goals document [3]***.

The investigation is not possible without substantial developments in autonomy. The sheer area involved requires many surface assets, including rovers, helicopters, and fixed landers. Each asset cannot wait for an Earth-based team to provide daily instructions on where to move, which targets to select, and whether or not the target is of interest. In particular, this investigation requires surface navigation, individual-agent planning, multi-agent planning, and automated science analysis.

Comparison to State of the Art

Mars rovers to date have used onboard stereo vision to detect and avoid obstacles and to do visual dead reckoning of their position relative to the start of each drive. At the end of each Mars day ("sol"), the rover position relative to orbiter imagery has been estimated by human operators, who manually register downlinked images from the rovers to orbiter images. The 2020 Mars rover is expected to be able to drive up to about 300 meters per sol, using a new computer vision coprocessor to accelerate obstacle detection and visual dead reckoning. The total rover traverse objective for the 2020 mission, including time spent on science operations, is to cover about 20 km in 1.5 Mars years (about 2.8 Earth years). For comparison, the Opportunity rover, which landed on Mars in 2004, drove a total of about 45 km in about 14 Earth years. The 2020 Mars mission plans to carry a 2 kg helicopter to conduct the first ever technology demonstration of a heavier-than-air aircraft on another planet. If successful, this helicopter will execute about 5 flights, up to on the order of 100 m long.

Driving and flight distances are constrained by the power required for mobility and by the amount of energy available per sol from onboard solar arrays or radioisotope power systems. For future missions, energy-limited traverse distances on the order of 1 km/sol or more may be possible. The Curiosity rover, which landed in 2012, on average has driven on approximately one third of the sols in the mission; non-driving sols were spent on a variety of functions, including science operations.

Autonomous vehicles on Earth can operate much faster than vehicles on Mars, but have access to much more energy, such as hydrocarbon fuels that are manually replenished, and use non-space-qualified onboard computers that have much higher performance than is available now for spacecraft. This and other factors make direct performance comparisons of Earth and Mars vehicles of limited value.

In the area of intelligent science instruments or "autonomous science," only relatively limited demonstrations of onboard autonomy have been done, such as automatic detection of clouds and dust devils [10]. Some instruments contain simple optimization algorithms. The Sample Analysis at Mars instrument on the Mars Science Laboratory contains such an algorithm. However, these simple algorithms do not constitute autonomy. The value to NASA of science autonomy will become enormous over time. Current science analysis on all missions to Mars, including Mars 2020 and ExoMars 2020, relies on relaying complete science data to Earth for analysis where a large team of scientists manually evaluates the data and makes decisions about the next steps for the mission. This approach creates a data volume limitation. In 2021 three rovers may be operating on Mars with as many as four (Mars Reconnaissance Orbiter [MRO], MarsExpress, Mars Atmosphere and Volatile EvolutioN [MAVEN] and Trace Gas Orbiter [TGO]) relay satellites transmitting data to Earth and yet each mission is bandwidth limited.

Findings

The Mars DRM team finds that the following actions and activities would facilitate implementation of the DRM scenario described above.

- Embrace the paradigm of **Exploration Campaigns** with a scalable network of cooperating, independent assets.
- Continue to develop **autonomous navigation and operation skills**, such as the ability to drill and handle samples. This technology cuts across almost any robotic planetary mission.
- Develop artificial intelligence techniques for **in situ science data analysis** for each type of instrument expected to be deployed on Mars or other planetary missions.
- Immediately start developing very small, low powered, **peer-to-peer interface standards** for multiple agents.
- Develop much **more powerful spaceflight compatible computing** platforms. Make base ship platform capable of performing the equivalent of "cloud" computing services for surface assets.
- Develop artificial intelligence **techniques to monitor health of surface assets** to identify and work around faults for reduced risk and increased operational efficiency.

Part II: The Case for Mars

Introduction

This Mars DRM aims to study the ground-water ice in the context of climate and regional geology, local weather, and biology while also providing detailed insight on the location and potential exploitation of subsurface water on Mars.

Why Mars?

Mars is considered a possible abode for past, modern, and future (human) life. As such, it is a key planetary target for exploration. ***From an astrobiological perspective***, Mars may have hosted ancient microbial life when the planet was warmer and wetter than today and it is possible that microbial life persists on modern (last 5 million years to present) Mars in the subsurface, away from the intense ionizing radiation and dryness of the surface. Models indicate that the obliquity cycle of Mars has a significant influence on the climate and geohydrology of the planet, such that mid- to high-latitude near-subsurface ice (several meters) may have been flowing ground water during times of high obliquity [4, 5, 6]. Furthermore, between wet periods, the ground ice can be lost to sublimation or can be mixed with other materials by periglacial freeze/thaw churning of near surface sediments [7, 8, 9]. ***From a Mars system perspective***, piecing together the reservoirs and dynamics of the Mars climate and its hydrologic cycles is critical to understanding planetary evolution, atmospheric composition, where water resources are most likely concentrated, and even the modern-day surface conditions (e.g., frost formation, near-surface moisture mobility, salt distribution, static charge). ***From a human exploration perspective***, water resources may fulfill a critical resource need for humans and their habitat, and present potential hazards such as biology or high salt concentrations. What is more, determining the physical and chemical properties of subsurface water, its distribution and mobility, and its biological potential may influence human activities. Humans will change the Mars environment at least on a local scale if not a regional scale, and they will need to monitor this change for the sake of science and to also mitigate risks to human safety and equipment longevity. This requires a fundamental understanding of the Mars surface and near-subsurface prior to direct human influence. Robotic missions may fill this knowledge gap, but an array of mobile platforms is necessary to cover regions on the scale of a human exploration zone (100-km radius).

DRM Science Objectives

There are three overarching objectives of the DRM. Addressing these objectives will enable scientists to answer key science questions.

Objective 1: Determine the distribution and physical context of subsurface (0-5 m) water at regional scale (approximately a 100-km radius).

a. Does the in situ map corroborate remote sensing water maps?
b. Is it primarily pore ice, layered-ice, icy regolith, or mineral hydration?
c. Is the presence or nature of water related to geomorphic and other geological features within the study region? What is the nature of the water reservoir?
d. What processes and sources are responsible for water detected? Do they reveal anything about changes in climate with respect to obliquity?

Objective 2: Determine subsurface physical, chemical, and biological water qualities

a. What is the water activity, Eh, and pH?
b. What is the composition of impurities? Do they support habitability?

c. Is there any indication of recent biology in the water? Recent biology includes extant life and dead organisms that may be recorded in ice since the last thawing, as these two groups will have the greatest impact on future missions.

Objective 3: Monitor weather conditions at regional scale.

a. How do surface environmental conditions (temperature, humidity, wind, radiation) affect the physical state of the subsurface water?

These objectives might also support human exploration interest in knowing where the best places are for accessing subsurface water in the actual exploration zone (if humans go to the same region) or in an analogous site; what to expect in terms of water qualities that pose advantages and disadvantages to human activities; meteorology data that might be relevant to human missions; and an understanding of effects of meteorological conditions on subsurface water or water brought to the surface for use.

Part III: Design Reference Mission Scenario

A Mars Subsurface Geohydrology Investigation

As the first stage of the Mars Exploration Campaign, the science-motivated Mars Subsurface Geohydrology Investigation will consist of multiple missions to Mars in order to survey on the scale required. Each mission consists of several surface assets. We conceive the first mission to use a small number of assets with a target zone of tens of square kilometers. The number of assets will be scaled up at each mission until sufficient assets are in place to meet the objectives and complete a detailed geohydrology map on the scale of the human exploration zone.

The Concept of Operations

The concept of operations for the Mars Subsurface Geohydrology Investigation consists of a fleet of small rovers, helicopters and a fixed lander.

Each rover contains instruments capable of providing: ice and hydrated mineral measurements; subsurface sounding measurements, such as ground penetrating radar; ice solute composition measurements, such as Phoenix Ion Selective Electrodes (ISE); drilling and sample acquisition; weather measurements; imagers for surface feature detection and navigation; and, communication with other surface assets and orbiters. Each rover is also capable of caching samples and delivering them to the fixed lander, or eventually to a human base station.

Small, independent helicopters provide aerial atmospheric measurements and surface imagery. Weather measurements at altitude complement surface measurements and enhance the understanding of Martian water system and weather patterns. Note that the helicopters may not be used if the selected exploration zone is at relatively high altitude.

The fixed lander, or "base ship," contains a laboratory of instruments to perform a detailed analysis of samples delivered by the rover fleet. It has robust communication with orbiting assets as well as direct communication with Earth. Instruments onboard the base ship are capable of biosignature detection. The base ship also contains a powerful computer capable of supporting neural networks.

The rovers use the base ship's computing ability for detailed analysis as they perform field sample collections. Rovers transmit instrument science data to the base ship where the computer's neural networks analyze the data to identify the fundamental composition of the sample. This high-level science information is useful for three purposes:

- The rovers' instruments use this information to determine how they should tune themselves and whether the sample requires further analysis.
- The fundamental composition results are automatically integrated into the geohydrology map.
- Fundamental composition results are transmitted to Earth rather than the complete science data set from each rover, dramatically reducing data volume. Science team can selectively request supporting data for the most interesting results.

Rovers traverse outward from the landing site in a cooperative search pattern. Samples are drilled at intervals and at likely places based on geology, surface features, and information from remote sensing water maps. As the mission progresses, science teams on Earth use the growing subsurface geohydrology map along with weather data to refine the description of likely places for drilling and gain a better understanding of the global Mars system.

Assumption(s)

There will be a few assumptions for this mission:

- Orbiters will be in place to support surface communications with Earth.
- Computing power for surface assets will be powerful enough to perform neural network algorithms.
- Tactical planning will be performed in situ on Mars, strategic planning will be done from Earth.
- Hardware to support relatively high-bandwidth peer-to-peer communications on the surface at rates on the order of 5Mbits/second.
- More energy will be available to rovers either through reduced power needs for mobility or improved solar or other energy production methods.
- Lightweight drilling systems capable of delivering samples from 1-5 meters below the surface.
- Advances in ground-penetrating radar and magnetic induction spectrometry to identify subsurface water and quantify the state of the water as liquid, ice, or within a clay mineral.

Autonomy is needed for this DRM for the following purposes:

A. Individual Agent Task Planning
B. Collaborative Multi-agent Task Planning
C. Sample Acquisition and Delivery
D. Surface Navigation
E. Scientific Autonomy

Each of these items is described in detail below. The autonomous technologies needed for this DRM are summarized in the following table, using NASA's Autonomous Systems Capability Leadership Team (AS-CLT) Taxonomy document as a guide.

Capability	Functionality	Autonomous System Taxonomy
Individual Agent Task Planning	Collection and processing of information internal and external to the system from sensors and instruments.	*Sensing and perception*
	Selection of goals, objectives, and activities to achieve a mission, subject to the situation and constraints.	*Mission planning and scheduling*
	Selection and ordering of activities to be performed while managing system resources to achieve mission goals.	*Activity and resource planning*
	Agreement on current and future activities, their priorities, and their disposition among elements or systems.	*Goal and task negotiation*
	Change of system state to meet mission goals and objectives according to a plan or schedule, subject to control authority and permission and based on mission phase, environment, or system state.	*Execution and control*
Collaborative Multi-agent Task Planning	Agreement on current and future activities, their priorities, and their disposition among elements or systems	*Goal and task negotiation*
	Collection, assembly, sharing, and interpretation of information and intent among elements to solve problems and plan actions/responses.	*Joint knowledge and understanding*
	Estimation of internal and external states from raw or processed inputs generated by multiple sensors/instruments, ascertainment, and continual comparison to expected states.	*State estimation and monitoring*
	Selection of goals, objectives, and activities to achieve a mission, subject to the situation and constraints.	*Mission planning and scheduling*
	Selection and ordering of activities to be performed while managing system resources to achieve mission goals.	*Activity and resource planning and scheduling*
	Change of system state to meet mission goals and objectives according to a plan or schedule, subject to	*Execution and control*

	control authority and permission and based on mission phase, environment, or system state.	
	Assurance that the system is operating in a manner consistent with expectations of all elements.	*Operational trust building*
Sample Acquisition and Delivery	Collection and processing of information internal and external to the system from sensors and instruments.	*Sensing and Perception*
	Creation of information sources about the environment or the system from sensing, perception, and human interaction that can be queried.	*Knowledge and Model Building*
	Evaluation of whether the state of the environment, the state of the system, and/or their interaction pose a threat to the safety of actions (or inactions) that are contemplated, which could compromise the system or mission.	*Hazard Assessment*
	Analyses of data (about environment or system) to identify events and trends that may affect future state, operations, or decision-making.	*Event and Trend Identification*
	Determination that the environment or system does not exhibit expected characteristics.	*Anomaly Detection*
	Selection of goals, objectives, and activities to achieve a mission, subject to the situation and constraints.	*Mission Planning and Scheduling*
	Selection and ordering of activities to be performed while managing system resources to achieve mission goals.	*Activity and Resource Planning and Scheduling*
	Generation or modification of a path or trajectory to reach a desired target physical location or configuration, subject to system and environment constraints.	*Motion Planning*
	Change of system state to meet mission goals and objectives according to a plan or schedule, subject to control authority and permission and based on mission phase, environment, or system state.	*Execution and Control*
	Identification of faults, prediction of future faults, and assessment of system capability as a consequence of those faults.	*Fault Diagnosis and Prognosis*
	Restoration of nominal or best-possible system configuration and operations after a fault.	*Fault Response*
	Adapting to changing environments and conditions without explicit re-programming, using knowledge collected from the past or from other systems' experiences.	*Learning and Adapting*
Surface Navigation	Generation or modification of a path or trajectory to reach a desired target physical location or configuration, subject to system and environment constraints.	*Motion Planning*
	Change of system state to meet mission goals and objectives according to a plan or schedule, subject to	*Execution and Control*

	control authority and permission and based on mission phase, environment or system state.	
	Adapting to changing environments and conditions without explicit re-programming, using knowledge collected from the past or from other systems' experiences.	*Learning and Adapting*
	Creation of information sources about the environment or the system from sensing, perception, and human interaction that can be queried.	*Knowledge and Model Building*
	Estimation of internal and external states from raw or processed inputs generated by multiple sensors/instruments, ascertainment, and continual comparison to expected states.	*State Estimation and Monitoring*
	Collection and processing of information internal and external to the system from sensors and instruments.	*Sensing and Perception*
Scientific Autonomy	In situ calibration and parameter-setting for instrumentation.	*Learning and adapting.*
	Assessment of measurement quality.	*State estimation and monitoring*
	Automated target selection for sampling.	*Reasoning and Acting.*

A. Individual Agent Task Planning

This individual rover should be able to inspect its surroundings, identify a target location to study, and determine if the science data is sufficient or if another target should be identified and analyzed. This would an enabling technology.

A first requirement is a framework for specifying the rover's high-level mission for the duration of its autonomous operation, as determined by a combination of the base ship, orbiter, and Earth. From this high-level specification (e.g., map certain area), the rover should be able to autonomously select its lower-level objectives and activities (including both those necessary for the mission and its own continued operation). The basis of such planning will be the rover's model of its current state and interpretation of its scientific measurements, the latter of which requires new techniques for processing raw data into intelligible and actionable observations. The rover must periodically re-evaluate its plan and schedule in the face of new information, and be able to respond immediately to urgent situations such as system failures and transient events of interest.

B. Collaborative Multi-Agent Task Planning

The individual agents need to cooperate to efficiently implement a larger plan and automatically adjust the plan based on new data (e.g., maintaining an overall map and selecting targets for each agent based on minimum movement or based on expectation of findings). This would be an enabling technology.

With a heterogeneous team of rovers and helicopters, one of the first collaborative tasks to be performed will be high-resolution mapping by the helicopters. This mapping activity will determine terrain trafficability for rovers with a spatial resolution at least an order of magnitude better than is possible from orbit. Cameras on the helicopters will be able to obtain millimeter-scale imagery, which can be analyzed by neural network algorithms on the lander or even onboard the helicopters to identify stratigraphic formations of scientific interest. It may also be possible for helicopters to carry spectral instruments to do some mineralogical characterization, or miniature neutron spectrometers to measure shallow subsurface bulk hydrogen content. Helicopters will also perform basic meteorological measurements.

The initial helicopter mission will be planned using regional map information from orbiters. Higher-resolution map information collected by helicopters will be integrated on the lander. The integrated map will be used to refine and extend helicopter mission plans and to create rover mission plans. As further mapping and science information is integrated from the rover(s), that will also affect subsequent rover mission planning. The rate of progress of individual rovers will depend on science opportunities and results that are discovered on the way, so plans for each rover may be affected by progress and discoveries made by the others.

Other non-autonomous technologies that are needed include delay-tolerant networking (DTN), mesh networking, peer-to-peer interface standards for multiple interacting agents, and high-performance, remote computing.

C. Sample Acquisition and Delivery

Section E “Scientific Autonomy” describes instruments capable of subsurface water detection that provide the rover with a likely location and depth to drill for a sample. Section D “Surface Navigation” describes how the rover approaches the drill location. This section describes the technology to safely operate the drill, manipulate samples returned by the drill, and deliver the samples to the instruments within the same agent or on another agent.

Automated sample collection and manipulation require hazard assessment, anomaly detection, sensing and perception, and self-awareness.

Subsurface obstacles such as a hard rock could damage or permanently disable a drill. Onboard analysis of the subsurface instrument data allows the agent to assess the hazard to the drill. During drilling operations, anomaly detection is required to reduce damage and prevent jamming the drill into unexpectedly hard rock. A machine-learning algorithm resident on the base ship combines the subsurface data with past drill performance (of all mobile agents) to improve identification of hazardous subsurface materials as the mission progresses.

The rovers must know their location with respect to the base ship for the sample handoff. Section D, “Surface Navigation,” describes the means to navigate to the base ship. The handoff of the sample will be done using image analysis of the base ship’s sample receptacle, specifically designed for visual identification.

Other non-autonomous technologies that are needed include a lightweight drill capable of 5 m (TBR); sample collection capability, the handoff of potentially wet samples to the base ship, ground-penetrating radar and magnetic induction spectroscopy tuned for water detection, and sample mass or volume verification.

D. Surface Navigation

Each individual agent traverses an area to a target specified by the plan. The agent determines the best route and avoids obstacles to reach the target with the optimum route based on risk, time and energy. This would be an enabling technology.

Navigation functions include state estimation, terrain perception, and path planning. State variables to be estimated include the position, velocity, heading, and tilt of rovers, plus the altitude of helicopters. Most of these variables will be estimated using a combination of visual and inertial measurements plus wheel odometry for rovers and altimeters for helicopters. The position of all vehicles relative to regional maps created from orbiter images will be measured by corresponding features seen in images onboard the vehicles and in orbital imagery. Tilt and heading measurements may also be obtained by imaging the Sun or by recognizing landmark features on the horizon. The lander will maintain knowledge of the position of all vehicles and landmark features. The lander may detect when the same landmarks are visible to more than one platform and perform a joint optimization of the landmark and multiple vehicle positions.

Terrain perception includes perceiving the geometry of the terrain, as in creating digital elevation maps, and estimating other physical properties relevant to trafficability, such as parameters like soil cohesion that affect rover slip. Trafficability parameters that are determined by direct contact with the terrain can be associated with the geometry (e.g., slope) and appearance (e.g., texture) of the terrain, so that it will be possible to predict soil parameters ahead of rovers based on the geometry and appearance of the terrain. This form of learning and adaptation may be generalized; for example, if it is possible to associate learned soil parameters with terrain appearance in orbital imagery, and thereby to propagate locally-learned trafficability inferences to the entire region covered by the orbital imagery.

Motion planning will start with the map available from orbiter knowledge and will be revised as better map knowledge is accumulated from helicopters and rovers. Inferences about trafficability, such as parameters affecting slip, will be uncertain, so both the terrain representation and the motion planning algorithms will need to model and reason about such uncertainty.

E. Scientific Autonomy

Instruments on this mission require onboard intelligence. Subsurface instruments such as ground penetrating radar need to identify likely locations for subsurface water, identify rocks that might damage the drill, and know when water is not likely in the area being studied. In addition, instruments on each rover or the base ship need to analyze samples drilled from several meters below the Martian surface. The instruments will characterize any ice found in the sample, as well as identify minerology and signs of recent life (see the DRM Science Objectives outlined in Part II).

Time and bandwidth requirements require that the individual rovers' science instruments be intelligent, thus this is an enabling technology. Scientific autonomy, or the ability to analyze the science data in situ, will be required for three purposes:

1. The science instruments need to be able to adjust and tune themselves based on data. If the instruments see something of interest in the data, they should be able to adjust themselves without a human in the loop to further analyze the target.
2. The high rate of target acquisition and analysis on multiple-surface assets will result in data volumes too high to return to Earth. Science instruments need to reduce data volume by identifying interesting data and culling uninteresting data.
3. The instruments should provide decisional information to the local rover and the larger network of assets to determine future targets. This information may influence the decision to move, search out a new location, or to drill deeper for another sample.

Instruments capable of the detailed analysis of samples required by this mission will have numerous tunable parameters. A typical analysis experiment would start with a survey experiment where the contents of the sample are entirely unknown, and the instrument's parameters are configured for a wide range. Follow-on experiments may then be performed to provide more detail or confirm autonomously derived hypotheses.

In this DRM, the instruments send the results of the survey experiment data to the base ship for analysis. The base ship analyzes the data using its knowledge of other samples, potentially from other rovers, and responds to the rover with a set of further experiments to be performed on the sample. The rover tunes its own parameters to implement the experiments suggested by the base ship. It may need to verify the existence of particular constituents, or more accurately measure a quantity, or possibly discard the sample and either drill to a different depth or move to a new location. The base ship's analysis may illicit more than one detailed analysis experiment.

Instruments can be expected to generate large amounts of data. Several rovers working independently and at several times the speed of current rovers make it impossible to transmit all the science data back to Earth. This scenario requires the basic decision-making ability to understand what data is worth sending back to Earth when bandwidth is limited.

Commercial activity in the realm of automated science data analysis is too focused to be applicable to the discovery-driven science necessary for this mission. Also, approaches to autonomy will be unique to each class of instrument. For instance, a completely different learning algorithm will have to be applied to mass spectrometer data than to a Laser Induced Breakdown Spectrometer.

Other non-autonomous Technologies

Other non-autonomous technologies that are needed include surface-imaging computing into the Digital Terrain and Geology Map (DTGM), high-performance computing power, in situ subsurface structure remote sensing at rover scale for integration with DTGM for 3-D models, an onboard spectral analysis to mineralogical content, and an onboard interest operator to analyze, prioritize, and decide next activity especially for transient events.

Relevant Research and Development projects for this Mars Subsurface Geohydrology Investigation DRM
Develop an integration and test approach for each system of autonomy above, including independent safety management at a “do-no-harm” level.
Develop calibration plans for science instruments centered around creating large data sets explicitly designed to train machine learning algorithms.

Part IV: Findings

The cost of developing the autonomy technologies described in this DRM are enormous. Yet the cost of not developing them is even larger. The increased science return on any planetary mission, not just missions to Mars, vastly outweighs the cost of developing these technologies. Autonomy increases the rate of science collection, improves the quality of science data, and ensures the data returned to Earth includes the most interesting science information. Once the autonomy technology is developed on the ground, Mars is the place to prove it out and then it can be applied to many planetary mission scenarios, such as missions to hostile environments like Venus or Europa.

While vast resources are being committed commercially to similar problems, commercial developments in autonomy assume powerful computers and high-bandwidth connections to an essentially limitless Internet of support. These assumptions do not apply to NASA planetary missions, including this DRM. Investments should be made to fill in the gap between what the commercial companies are doing and what is possible on planetary missions.

The Mars DRM team finds that the following actions and activities would facilitate implementation of the DRM scenario described above.

- Embrace the paradigm of **Exploration Campaigns** with a scalable network of cooperating, independent assets.
- Continue to **develop autonomous navigation and operation skills**, such as the ability to drill and handle samples. This technology cuts across almost any robotic planetary mission.
- Develop artificial intelligence techniques for **in situ science data analysis** for each type of instrument expected to be deployed on Mars or other planetary missions.

- Immediately start developing very small, low powered, **peer-to-peer interface standards** for multiple agents.
- Develop much more **powerful spaceflight-compatible computing platforms**. Make the base ship platform capable of performing the equivalent of "cloud" computing services for surface assets.
- Develop artificial intelligence techniques to **monitor health of surface assets** to identify and work around faults to reduce risk and increase operational efficiency.

Why is this DRM important?
The **Mars Exploration Campaign** paradigm defined in this DRM is a blueprint not only for Mars exploration, but exploration of most planetary targets. Once developed, the technologies of autonomous navigation, cooperation among a team of independent assets, and science autonomy will be enabling to any planetary mission.

Mars represents the best place to establish these technologies. NASA has a generation of experience in robotic operations on Mars. The environment and terrain are well known. Yet each mission raises more questions than it answers. Humans may someday help answer these questions, but they will need enormous support to do so. This DRM provides crucial data to support human life on Mars, such as the location and nature of in situ resource. It also provides a framework for human-robotic interaction once humans do arrive.

Part V: Mars DRM Team

The Mars Design Reference Mission team is comprised of:
Eric Lyness (Co-Lead), NASA GSFC
Jennifer Eigenbrode (Co-Lead), NASA GSFC
Larry Matthies, NASA JPL
Rich Doyle, NASA JPL
Jay Falker, NASA STMD
Eugene Fang, Carnegie Mellon University
Philip Koopman, Carnegie Mellon University
Rob Manning, NASA JPL
Bryan O'Gorman, UC-Berkeley
Florence Tan, NASA SMD

Part VI: References

1. National Aeronautics and Space Administration, Developed by the Strategic Investments Division (2018) NASA_2018_Strategic_Plan. 1–64.
2. National Research Council. 2011. *Vision and Voyages for Planetary Science in the Decade 2013-2022*. Washington, DC: The National Academies Press. https://doi.org/10.17226/13117.
3. MEPAG (2018), Mars Scientific Goals, Objectives, Investigations, and Priorities: 2018. D. Banfield, ed., 81 p. white paper posted October, 2018 by the Mars Exploration Program Analysis Group (MEPAG) at https://mepag.jpl.nasa.gov/reports.cfm.
4. Laskar, J., Levrard, B., and Mustard, J.F. (2002) Orbital Forcing of the Martian Polar Layered Deposits. Nature 419: 375–77.
5. Richardson, M.I., and Mischna, M.A. (2005) Long-Term Evolution of Transient Liquid Water on Mars. Journal of Geophysical Research 110 (E3): E03003.
6. Zent, A. (2008) A Historical Search for Habitable Ice at the PHX Landing Site. Icarus 196 (2): 385–408.
7. Mellon M. T., Malin M. C., Arvidson R. E., Searls M. L., Sizemore H. G., Heet T. L., Lemmon M. T., Keller H. U. and Marshall J. (2009) The periglacial landscape at the Phoenix landing site. *J. Geophys. Res.* **114**, E02S02.
8. Jawin E. R., Head J. W. and Marchant D. R. (2018) Transient post-glacial processes on Mars: Geomorphologic evidence for a paraglacial period. *Icarus* **309**, 187–206.
9. Soare R. J., Conway S. J. and Dohm J. M. (2014) Possible ice-wedge polygons and recent landscape modification by "wet" periglacial processes in and around the Argyre impact basin, Mars. *Icarus* **233**, 214–228.
10. Castano, A., Fukunaga, A., et al. (2006) Autonomous Detection of Dust Devils and Clouds on Mars. IEEE International Conference on Image Processing

Mars Design Reference Mission Report Summary

NASA has studied Mars more than any other solar system object outside Earth and the Moon. The scientific exploration of Earth's planetary neighbor has largely focused on addressing the presence and persistence of water, geochemistry, geology, and atmospheric evolution.

Prior, current, and near-term missions are filling in fundamental knowledge gaps regarding Mars and in doing so, support models of how the Mars system functions and has evolved. But these missions involve singular spacecraft in singular localities. A sustained, wide-area study is needed to enable astrobiological research concerning potential past, modern, and future (human) life on Mars, to support system-level understanding of Mars processes and conditions on a regional scale, and to support future human exploration.

Because of this need, the Mars team's suggested DRM is not just a single mission but a practical, scalable, and sustainable Mars exploration campaign that establishes an exploration framework on Mars. In this framework, new spacecraft, new rovers, and missions themselves become new elements within the campaign's framework. This campaign will study the groundwater ice in the context of climate and regional geology, local weather, and possible biology, while also providing detailed insight into the location and potential exploitation of subsurface water on Mars.

The Mars DRM team suggests the following DRM scenario.

DRM Scenario: A Mars Subsurface Geohydrology Investigation

This science-motivated investigation will consist of multiple missions to Mars to survey the planet on the scale required. Each mission will consist of several surface assets. The first mission will use a small number of assets with a target zone of tens of square kilometers. The number of assets will be scaled up for each mission in this scenario until sufficient assets are in place to meet the objectives and complete a detailed geohydrology map on the scale of the expected human exploration zone (~100-km radius).

The investigation scenario is not possible without substantial developments in autonomy. The sheer area to be investigated requires many agents—including a fleet of rovers, helicopters, a fixed lander, and an orbiter. Each asset cannot wait for an Earth-based team to provide daily instructions on where to move, which targets to select, and whether the target is of interest.

This DRM scenario requires a level of autonomy that is not currently available. Advancements in autonomy technology are required for this mission scenario to perform the following:

Individual Agent Task Planning: Autonomy will allow an individual rover to inspect its surroundings, identify a target location to study, and determine if the science data is sufficient or if another target should be identified and analyzed.

Collaborative Multi-agent Task Planning: Autonomy will allow the individual agents to cooperate and efficiently implement a larger plan and automatically adjust the plan based on new data. For instance, the system must be capable of maintaining an overall map and selecting targets for each agent based on minimum movement or based on expectation of findings.

Sample Acquisition and Delivery: Autonomy will allow for automated sample collection and manipulation, including activities such as safely operating a drill, manipulating samples returned by the drill, and delivering the samples to the instruments on the same agent or on another agent.

Surface Navigation: Autonomy will allow each individual agent to traverse an area to a target specified by the mission plan. For example, the agent will determine the best route and avoid obstacles to reach the target using the optimum route based on risk, time, and energy.

Scientific Autonomy: Autonomy will provide the ability to analyze the science data in situ. The science instruments will need to adjust and tune themselves based on data obtained. Science instruments will also need to reduce data volume by identifying interesting data and culling uninteresting data. The instruments should also provide decisional information to the local rover and the larger network of assets to determine future targets.

To enable autonomy in this DRM scenario, advancements in the following supporting technology areas are required:

- Surface imaging computing into the Digital Terrain and Geology Map (DTGM)
- High-performance remote computing power to support machine learning, including neural networks
- In situ, remote sensing of subsurface structure at rover scale for integration with DTGM for 3-D models
- An onboard interest operator to analyze, prioritize, and decide the next activity, especially for transient events
- Delay-tolerant networking (DTN) and mesh networking.
- Peer-to-peer interface standards for multiple interacting agents
- High bandwidth (on the order of 5Mbits/second), surface-to-surface, over-the-horizon data communications
- A lightweight drill capable of delivering potentially wet samples from minimum depth of 1-5m
- Ground-penetrating radar and magnetic induction spectroscopy tuned for water detection

Findings

The Mars DRM team finds that the following actions and activities would facilitate implementation of the DRM scenario described above:

1. Embrace the paradigm of the Mars Exploration Campaign with a scalable network of cooperating, independent assets.
2. Continue to develop autonomous navigation and operation skills, such as the ability to drill and handle samples. This technology cuts across almost any robotic planetary mission.
3. Develop artificial intelligence techniques for in situ science data analysis for each type of instrument expected to be deployed on Mars or other planetary missions.
4. Immediately start to develop very small, low-powered, peer-to-peer interface standards for multiple agents.

5. Develop high-bandwidth, peer-to-peer data communication devices.
6. Develop much more powerful spaceflight-compatible computing platforms. The base ship platform should be capable of performing the equivalent of “cloud computing” services for surface assets.
7. Develop artificial intelligence techniques to monitor health of surface assets and identify and work around faults to reduce risk and increase operational efficiency.