ARTIFICIAL INTELLIGENCE

# EELS: Autonomous snake-like robot with task and motion planning capabilities for ice world exploration

T. S. Vaquero[1]*†, G. Daddi[2]†, R. Thakker[1]†, M. Paton[1]†, A. Jasour[1], M. P. Strub[1], R. M. Swan[1], R. Royce[1], M. Gildner[1], P. Tosi[1], M. Veismann[1], P. Gavrilov[1], E. Marteau[1], J. Bowkett[1], D. Loret de Mola Lemus[1], Y. Nakka[1], B. Hockman[1], A. Orekhov[3], T. D. Hasseler[1], C. Leake[1], B. Nuernberger[1], P. Proença[1], W. Reid[1], W. Talbot[1], N. Georgiev[1], T. Pailevanian[1], A. Archanian[1], E. Ambrose[1], J. Jasper[1], R. Etheredge[1], C. Roman[1], D. Levine[1], K. Otsu[1], S. Yearicks[1], H. Melikyan[1], R. R. Rieber[1], K. Carpenter[1], J. Nash[1], A. Jain[1], L. Shiraishi[1], M. Robinson[1], M. Travers[1], H. Choset[3], J. Burdick[4], A. Gardner[1], M. Cable[1], M. Ingham[1], M. Ono[1]



**Ice worlds are at the forefront of astrobiological interest because of the evidence of subsurface oceans. Enceladus in particular is unique among the icy moons because there are known vent systems that are likely connected to a subsurface ocean, through which the ocean water is ejected to space. An existing study has shown that sending small robots into the vents and directly sampling the ocean water is likely possible. To enable such a mission, NASA's Jet Propulsion Laboratory is developing a snake-like robot called Exobiology Extant Life Surveyor (EELS) that can navigate Enceladus' extreme surface and descend an erupting vent to capture unaltered liquid samples and potentially reach the ocean. However, navigating to and through Enceladus' environment is challenging: Because of the limitations of existing orbital reconnaissance, there is substantial uncertainty with respect to its geometry and the physical properties of the surface/vents; communication is limited, which requires highly autonomous robots to execute the mission with limited human supervision. Here, we provide an overview of the EELS project and its development effort to create a risk-aware autonomous robot to navigate these extreme ice terrains/environments. We describe the robot's architecture and the technical challenges to navigate and sense the icy environment safely and effectively. We focus on the challenges related to surface mobility, task and motion planning under uncertainty, and risk quantification. We provide initial results on mobility and risk-aware task and motion planning from field tests and simulated scenarios.**

## INTRODUCTION

Icy moons of gas giant planets, such as Europa, Titan, and Enceladus, are at the forefront of astrobiological interest because of the evidence of subsurface oceans. Saturn's icy moon Enceladus in particular is a compelling moon that is believed to have liquid water oceans beneath its frozen surface (*1*–*5*). Enceladus is unique not only because of its subsurface ocean but also because it has known geysers (*1*, *3*), its ocean is likely in direct contact with rocky core that supplies minerals (*6*), and Cassini's flythrough data indicated the potential existence of hydrothermal vents (*7*). The detection of biosignatures might be possible from samples captured via plume flythrough or captured on the surface while landing on it. Nevertheless, existing studies indicated that some of the key astrobiological questions cannot be answered without directly sampling the ocean water in the vent (*8*, *9*). The search for and characterization of these biosignatures, and consequently life, on Enceladus would ideally require samples collected from the vent, its surroundings, and the ocean.

Unfortunately, we only have very limited knowledge about the internal environment of the vents and the surface around them, most of which comes from a single mission, Cassini. The best resolution of the surface image is ~6 m per pixel for a very limited area. Most of the close-up images are monocular; hence, the surface topography at the scale of our robots is largely unknown. Furthermore, there are multiple conflicting hypotheses for the eruption mechanisms (*10*–*13*), which cannot be resolved by the existing data from Cassini. As a result, the vent geometry, depth to the liquid interface (if it exists), velocity of the flow, and the physical property of the ice walls are largely unconstrained.

Exploring Enceladus' uncertain ice environments without a precursor mission requires the development of resilient autonomous robotic systems that can operate, navigate, and adapt under challenging conditions. For example, a priori knowledge of the environment and terrain properties/dynamic is severely limited, leading to high uncertainty and risk during navigation and operation. Sampling from the underground ocean itself requires both surface and subsurface navigation and exploration where environmental conditions and physical properties of the ice/environment might change markedly over the surface-subsurface trajectories (in the subsurface case, there are added elements of limited visibility while navigating narrow passages and potentially turbulent flows in the vents). Communication with Earth is constrained because of long distances (the round-trip latency will be 130 to 155 min, whereas the downlink rate would be highly limited compared with Mars missions) and/or no line of sight when the system is in the subsurface environment. The system's lifetime might be limited because of harsh environmental conditions due to radiation or temperature extremes and limited resources, such as depletable batteries. These challenges preclude us from manually commanding a robot from the ground in such an environment, and, therefore, a robot explorer needs to be highly autonomous, adaptive, and resilient to execute the mission with limited to no human supervision.

[1]Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA. [2]Department of Mechanical and Aerospace Engineering, Politecnico di Torino, Torino TO 10129, Italy. [3]Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. [4]Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125, USA.
*Corresponding author. Email: tiago.stegun.vaquero@jpl.nasa.gov
†These authors contributed equally to this work.

Here, we present an autonomous, adaptive, and resilient robotic platform, the Exobiology Extant Life Surveyor (EELS) robot, developed at NASA's Jet Propulsion Laboratory (JPL), that aims to address the aforementioned challenges/needs. Figure 1 illustrates the EELS concept for an Enceladus mission. We introduce its resilient software architecture and focus on how we address the specific problem of safe surface navigation, exploration, and sensing through onboard autonomous task and motion planning under uncertainty. Our contributions refer specifically to the software architecture and the task and motion planning system for highly uncertain environments. We also provide a summary of the surface mobility and motion planning approach and experiments introduced in (*14*) while elaborating on the autonomy requirements that are needed to operate on Enceladus, the software architecture that can fulfill those requirements, and the high-level task-and-motion planning component of the software stack.

## The EELS project

The EELS robot is a serpent-like mobile instrument platform conceived to explore both surface and underground structures, assess habitability of the environment, and ultimately search for evidence of life. It is designed to be adaptable to traverse ocean world–inspired terrain, fluidized media, enclosed labyrinthian environments, and liquids. Enceladus is the main driver for the design of EELS hardware and software architecture, as well as its mobility and autonomous capabilities. We have been using glaciers as Earth analog ice environments to develop and test its architecture as a stepping stone toward Enceladus.

Planetary subglacial access missions aimed at achieving similar science goals to EELS have been proposed throughout the years. Thermal ice drilling via melt probes is the primary approach to overcoming planetary ice sheets (*15*–*18*). The commonality between these mission concepts is that mobility occurs through bulk ice, posing challenges because of the kilometer-scale cryogenic ice sheets of Enceladus or Europa. Thermal drilling in cryogenic ice is inefficient because of the heat lost through the probe's sidewalls and sediment accumulation, which can decrease thermal conductivity. A future challenge is presented by the potential presence of rocks or voids within the ice sheet. Another approach is via mobility systems capable of traversing the rough glacial terrain and climbing down the vent system. Noteworthy is the Enceladus vent explorer concept (*19*) that investigated the feasibility of using ice screws and articulating arms as a descent technology. As in (*19*), EELS leverages the existence of an open pathway to Enceladus' subglacial ocean and sidesteps the problems of thermal ice drilling in cryogenic ice using ice as terrain over which to move rather than a medium through which to move. Differently from (*19*), EELS is better suited for an Enceladus vent mobility mission because it is capable of mobility in a wider spectrum of terrains.

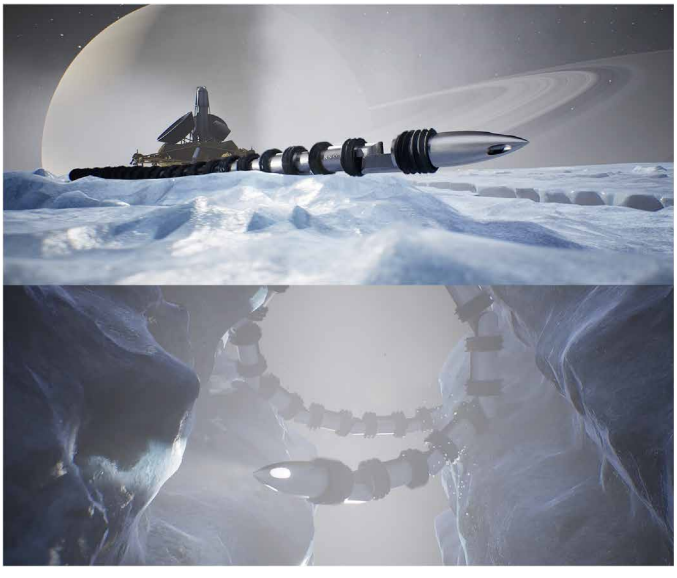

CREDIT: NASA/JPL-CALTECH

**Fig. 1. Artist renditions of the EELS concept.** (**Top**) The system moving with the help of its active skin propulsion toward a vent after being deployed by a lander on Enceladus. (**Bottom**) The EELS platform climbing down a vent, resisting the upward fluid dynamic pressure by pushing out against the ice walls.

The development of EELS encompasses nine specific capabilities essential for enabling an Enceladus vent exploration mission. These include proprioceptive control for reliable interaction with the ice surface and holding the robot in the high-speed jet; gait/motion control for fully using the available degrees of freedom and traveling through various types of environments; exteroceptive perception for acquiring geometric and semantic information of the local environment; motion planning for continuously finding a feasible robot path toward position goals; global localization and mapping to provide situational awareness in a global coordinate; global locomotion planning to cooptimize the path and gait/motion pattern; scientific sensing and actuation for observation, sampling, and sensor deployment; planning and robust execution of high-level mobility, science, and engineering activities and effective failure responses; and risk management to ensure that the mission plan is executed within user-specified risk thresholds. Given the high uncertainty aspect of the target environment, reasoning under uncertainty across these capabilities plays a central role in the EELS project and drives system-level autonomy requirements. In this work, we focus exclusively on addressing uncertainty in the processes of motion planning and task planning while managing risk of violating safety constraints. In what follows, we elaborate on the requirements for these particular three items.

A number of works have outlined future autonomy needs and challenges for planetary exploration (*20*, *21*). The common themes that always emerge for operations in remote regions of the solar system are the needs for system-level autonomy and decision-making under uncertainty. High uncertainty, coupled with long communication latency, low frequency of uplink-downlink cycles, and short mission time, exclude the classical operator-/ground-in-the-loop planning paradigm that has long been adopted by space missions. Successful surface (and subsurface) operations on Enceladus will thus require a system-level task and motion planner aware of the shifting environmental uncertainty and risk as more of the environment is perceived through active observations.

Another desirable characteristic of a system-level planner for an extreme environment is a degree of awareness of the robot's motion planning problem. When moving in an uncertain environment, there is an inherent trade-off among risk-taking, information gain, and use of the mission's finite resources. A system-level planning capability is the ideal coordinator for this trade-off because it has the broadest understanding of the system's state. Thus, the task planning capability should be cognizant of the lower-level motion planning problem at a high enough resolution to generate a sparse sequence of navigation tasks that break down the scientific intent of reaching the vent into waypoints for the path planning process.

An effective system-level planner should also be capable of inferring the likely consequences of making a decision given a prior belief and extracting an understanding of the risk that is associated with each action. It should be capable of inferring the most likely consequence as it navigates the environment and taking the appropriate information-gaining action as needed when necessary. For Enceladus, the necessity for risk awareness is visible, for example, in the surface mobility phase. When moving in a glacial environment, there are hazards of varying degrees of lethality ranging from minor inconveniences, such as small dips, to mission-ending crevasses. A planner capable of modulating its uncertainty tolerance on the basis of risk will produce plans that can achieve a higher science return with a fixed cost and without increasing the probability of catastrophic events. This planner capability will allow the system to adapt its behaviors on the basis of the risks that actions pose to the mission's outcome, given the system's state and external environment conditions. This is similar to the ground-in-the-loop, risk matrix–based operations paradigm that is currently in use, but in the Enceladus case, that risk matrix would be guiding the onboard decision-making process as illustrated in Fig. 2. Note that, although in Fig. 2, we have focused on a mobility example, the same concept of onboard risk assessment can be extended to a broader action space to include actions such as communication, scanning the environment, and more. Understanding the likely consequences of an action, however, requires deep knowledge of the system's state. Among the system state information that is needed for accurate consequence assessment, we include the probability of failure of the robot's functions. A simple example of why this is the case is a robot that is deciding whether to enter a crevasse. The likelihood of having operational low-light exteroception, for example, should certainly be a central component of the robot's decision-making process.

In addition to risk-aware deliberative planning, the system-level planning process also requires optimization because of limited mission onboard resources and time. Thus, it should be capable of finding optimized solutions to the relevant planning problem within these constraints. Moreover, the system requires the capability of rapidly reacting to off-nominal events, as well as a robust mission executive capable of meshing reaction and deliberation (which is usually computationally more expensive). Fulfilling these system-level planning and execution capability needs is the goal of the EELS project.

## RESULTS

### System architecture overview

The EELS robot combines versatile hardware capabilities (Fig. 3) and a software architecture that supports the proposed autonomous risk-aware task and motion planning functions.

#### *Hardware*

The EELS platform is a large-scale snake robot with active skin propulsion as shown in Fig. 3. Snake robots have long been studied for Earth-based applications for their high reconfigurability and capacity to move through unconsolidated, complex terrains (*22*, *23*). Previous work has traditionally focused on shape-based locomotion enabled by anisotropic skin friction, which limits the system's size because of shape actuators operating outside of a quasi-static regime when used for propulsion. The EELS platform, on the other hand, proposes an active skin propulsion. Only a handful of studies have looked into active skin propulsion using screws and tracks, but it has been demonstrated as an effective approach (*24*–*27*). Screw-based propulsion in particular allows the EELS platform to have a larger scale than shape-based snake robots, which enables, for example, larger science payload, more energy efficiency than shape-based locomotion, pushing against the surface of a vent with larger diameter and navigating over rough terrain, and potentially more efficient surface mobility and climbing on ice than traditional active skin mechanisms.

The system is composed of 10 identical segments that give EELS the self-repeating structure that is typical of snake robots. All segments sum up to a length of ~4 m and a mass of ~100 kg. Each segment is equipped with three actuators, two of which enable the robot to change shape, whereas the third actuates the skin (screw). Each shape actuator has a peak torque of 400 N·m. The large number of degrees of freedom allows for reconfigurability and a wide range of possible gaits. Alongside adaptable mobility performance, the robot's self-repeating

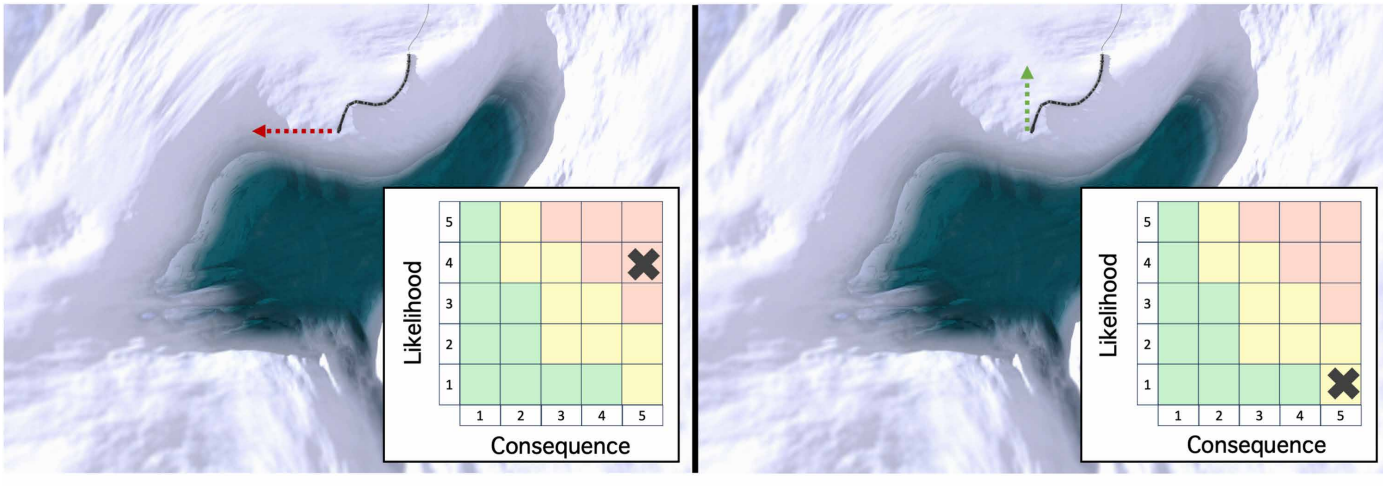


**Fig. 2. Onboard risk assessment.** A movement action is evaluated under the lens of risk. A matrix can be used to compute risk by multiplying the severity of the consequences of a failure times the likelihood of that failure. (**A**) High risk due to a high likelihood of falling into the crevasse, with potentially severe consequences for the mission. (**B**) An action that leads to lower probability of falling and hence a lower risk.

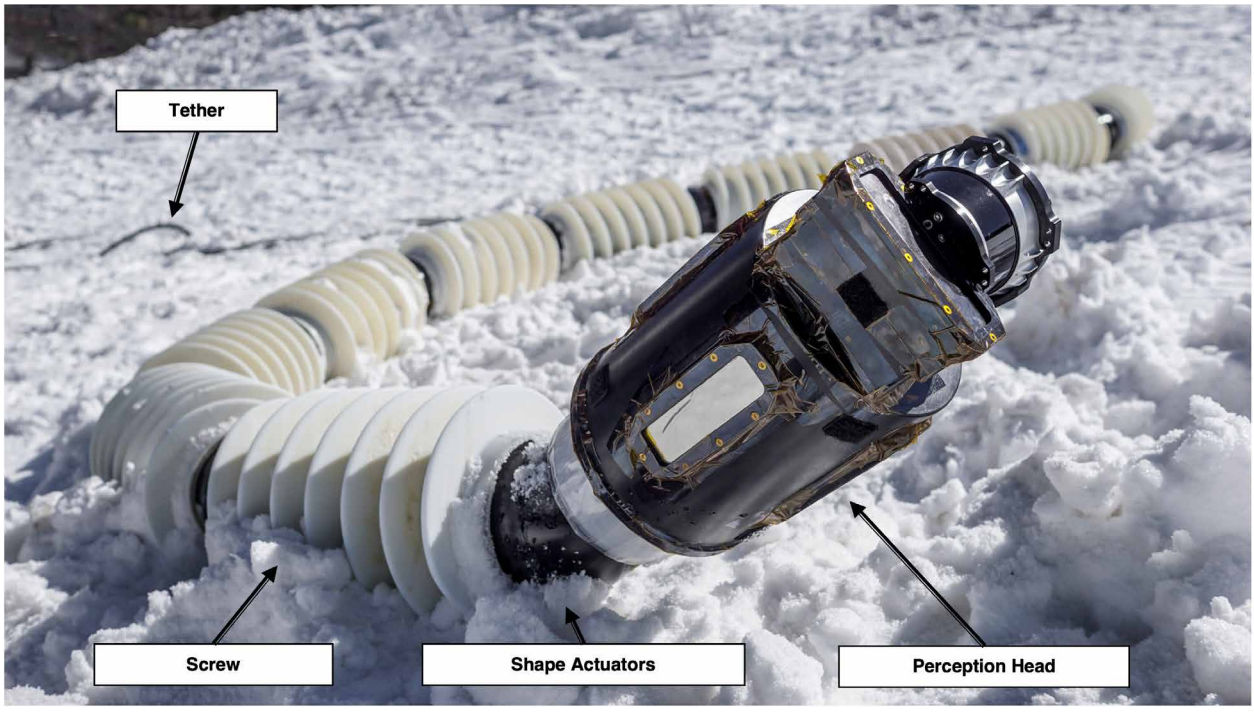


**Fig. 3. EELS hardware.** Core components of EELS robot, including perception head, shape actuators that connect each body segment, screws attached to each module, and the tether for data and electrical power transfer.

structure allows for graceful degradation and redundancy. Active skin propulsion is achieved through the use of counterrotating screws attached to each module. This characteristic enables movement while keeping the shape actuators in a quasi-static regime of operations.

EELS draws electrical power and communicates via a tether that is attached to the robot's terminal segment. Long surface and subsurface traversals will be enabled by high voltage, high communication bandwidth, lightweight tethers, such as (*28*). Future work will explore approaches for power management, communication, and tethering.

On the opposite end of the tail segment, there is a perception head composed of an Ouster OS0 light detection and ranging (LiDAR), four stereo cameras, an inertial measurement unit (IMU), and a barometer. To facilitate vision in enclosed environments such as glacial moulins, the perception head is equipped with light-emitting diode lights. The system is also equipped with proprioceptive sensors. Each segment's actuators have encoders that provide feedback about the robot's shape, whereas IMUs are located at the head and tail to estimate the pose without relying on the perception head. Moreover, current sensing allows characterization of robot-environment interactions.

### *Software*

Our software architecture takes inspiration from existing autonomy frameworks [such as JPL's Mission Data System (*29*), CLARAty (*30*), FRESCO (*31*), and ERGO (*32*)] by following a goal-oriented autonomy paradigm and introducing both uncertainty reasoning and risk awareness. At a high level, EELS' software stack is divided into estimator and controller modules as shown in Fig. 4. Estimators are tasked with inferring information about the system under control and the environment's state. In increasing levels of abstraction, there is a proprioceptive estimation module that is used to obtain information about the robot's internal state. The main outputs of this module are the shape and screw contact-state estimates that are consumed by controllers. Proprioceptive estimates are crucial both to achieve resilience to exteroceptive failures by estimating the robot's pose through proprioceptive sensors only and to provide controllers with robot-environment contact information, which is necessary to achieve screw propulsion and vertical/horizontal ice mobility. The exteroceptive estimation module concurrently infers the robot's pose and builds a map of the surrounding environment. It ingests raw information from LiDAR, cameras, IMU, and barometer and transforms the input into a map and a sequence of poses that are used by path planning and task and motion planning modules. A fault/risk management module is tasked with using state information and progress toward achieving goals from each module to estimate the likelihood of functionality failure and the uncertainty in that estimate. This module can be seen as a system-level fault and risk monitor that reasons at a time resolution and abstraction level sufficiently large to capture system-level trends that are a necessary input for the mission planning module.

In EELS, controllers are hierarchically ordered in a way that facilitates problem formulation, software development, component testing, and code maintainability. Each controller takes as input a goal from a higher level controller and translates it into decomposed goals for a lower-level controller. At the bottom of that hierarchy, we have low-level hardware commands, whereas at the top, we have mission-level goals expressed as operator intent. The actuator (joint-level) controller translates desired screw velocities and bend and twist actuator positions into hardware-understandable commands. Shape and screw

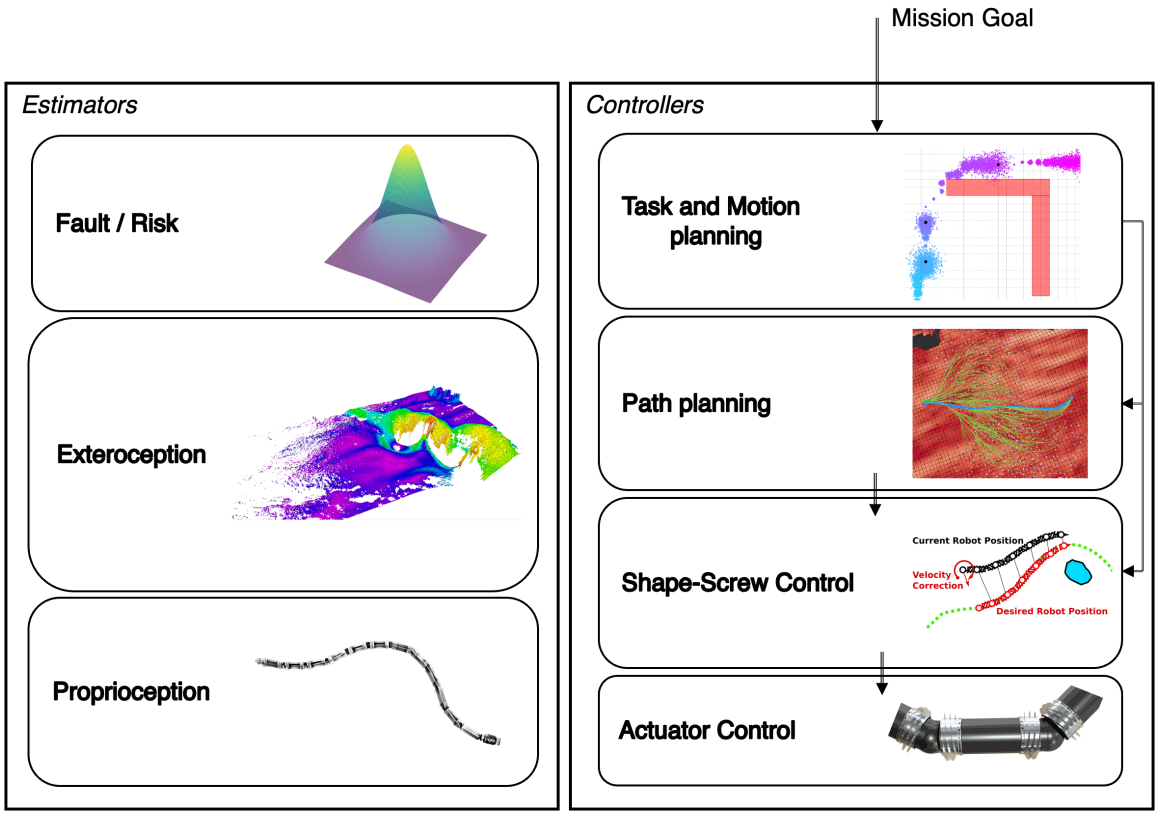


**Fig. 4. High-level view of EELS software architecture.** Modules are hierarchically organized and fall into two categories: estimators and controllers. The level of abstraction decreases going from top to bottom in the image. The lowest level estimators infer state from proprioceptive information. The exteroceptive estimators generate a map and localize the system. The fault and risk component reasons probabilistically about the system's health state. On the controllers side, the lowest layer is actuator control that sends commands to the motor controllers. Further up is shape and screw control that generates desired screw velocities actuator positions. Higher up, there are path planning and task and motion planning that respectively generate detailed local paths and high-level schedules of waypoints and actions.

control includes a set of controllers that receive desired path, desired controller, and deviation information and output desired joint angle and screw velocity. Different gaits require different control schemes. The path planning module consumes goals in the form of poses in the environment and generates an appropriate path toward these goals. The interested reader can find further details about the control and path planning modules, referred to here as the surface locomotion modules, in (*14*)—we summarize the main strategies used for locomotion in the next section. The task and motion planning module (one of the main contributions of this work) receives a high-level goal from operators in the form of constraints over the system's state and plans a sequence of tasks/behaviors (including motion tasks) that satisfies these constraints. Behaviors are highly abstracted representations of the robot's capabilities, such as "moving toward a waypoint," "scanning the environment," and more. The behaviors produce goals for various lower level controllers. In the examples above, the moving behavior produces goals for the path planning module, whereas the scanning behavior interacts directly with the shape controller.

## Surface locomotion strategies

In this section we describe how the flexibility of EELS' hardware, coupled with controller and estimators, was leveraged to achieve surface mobility in extreme terrains. To demonstrate surface mobility in such terrains, we deployed and tested two gaits alongside two different state estimation schemes. The main locomotion strategy is inspired by multiagent autonomy and decouples shape- and screw-based locomotion. This gait is known as leader follower and consists of generating a path that is followed by the robot's perception head. Propulsive forces are achieved by turning the screw actuators, and each hardware module is controlled in a way that follows the path traced by the perception head. When using exteroceptive state estimation, an error term between the robot's actual and the desired position can be generated and used to control the robot toward the planned path. By commanding a separate velocity to each screw, the robot can move holonomically toward a desired position, using the screws to correct any drift. When exteroceptive state estimation is unavailable, the robot can follow a path using the same screw propulsion and shape adaptation method, but errors and progress along the path will not be estimated. The second gait that was tested was shape-based sidewinding. With this locomotion strategy, the robot's time-varying shape in combination with the screw's anisotropic friction properties is leveraged to generate thrust. The capability of switching between screw-based and shape-based gaits provides flexibility and adaptability to different environments.

## Task and motion planning module

At the task and motion planning module, we have focused on facilitating surface mobility through the combination of task and motion planning under uncertainty. During operations, exteroceptive state estimation can fail both because of hardware faults or fundamental unobservability present in the environment. Observability issues are especially marked for snake robots because the perception head is very low to the environment, and the number of visible

features is often limited. When exteroceptive state estimation is operating nominally (see Fig. 5A), an almost optimal path can be followed, with a small amount of uncertainty introduced by the natural covariance growth typical simultaneous localization and mapping (SLAM) systems in absence of loop closures. On the other hand, when exteroception is not working (Fig. 5B), the motion and path must be conservatively planned by projecting uncertainty growth and ensuring that the likelihood of violating safety constraints is never exceeded. The key intuition that justifies a mixed control-task action space is that raising the perception head to scan the environment can improve the robot's pose knowledge by observing more features and potentially directly observing the goal. Being able to relocalize allows the module to plan a path with scan activities that is less costly than the conservative path without scans (Fig. 5C). Furthermore, this path will not have a greater risk associated with it because relocalizing allows the module to guarantee that the safety constraints are not exceeded when moving closer to the obstacles. Another advantage of this mixed approach is that high cumulative fractional area (*33*) environments, where narrow passages must be navigated to reach the goal, could be unsolvable for a conservative path planner. Conversely, the task and motion planner would be able to progress toward the goal under these environmental conditions until the observability issues are left behind and the system can switch back to exteroceptive navigation.

We developed a computationally efficient risk-aware task and motion planner capable of autonomously performing the trade-off between information-gaining actions and trajectory selection (handling the scenario in Fig. 5C). The challenges of combining continuous motion with discrete task planning is a growing research area (*26*, *34*), where the key limiting factor is scalability because of the problem's intractable nature. Our task and motion planner differs from previous work because it formulates the problem using mathematical programming rather than a sequential combination of task planning and motion planning and explicitly reasons over belief space. We also differ from previously flown system-level autonomy stacks, such as EO-1 with ASE (*35*) or perseverance's mission planning component (*36*), by explicitly reasoning about risk and uncertainty and by combining task planning with motion reasoning. Previous missions operated in environments where planning assumptions could be grounded in vastly more environmental knowledge than what is available for Enceladus. When presented with large environmental uncertainty, the planning process needs to account for this uncertainty, and tighter integration between tasks and motion planning enables more robust solutions when compared with a classical

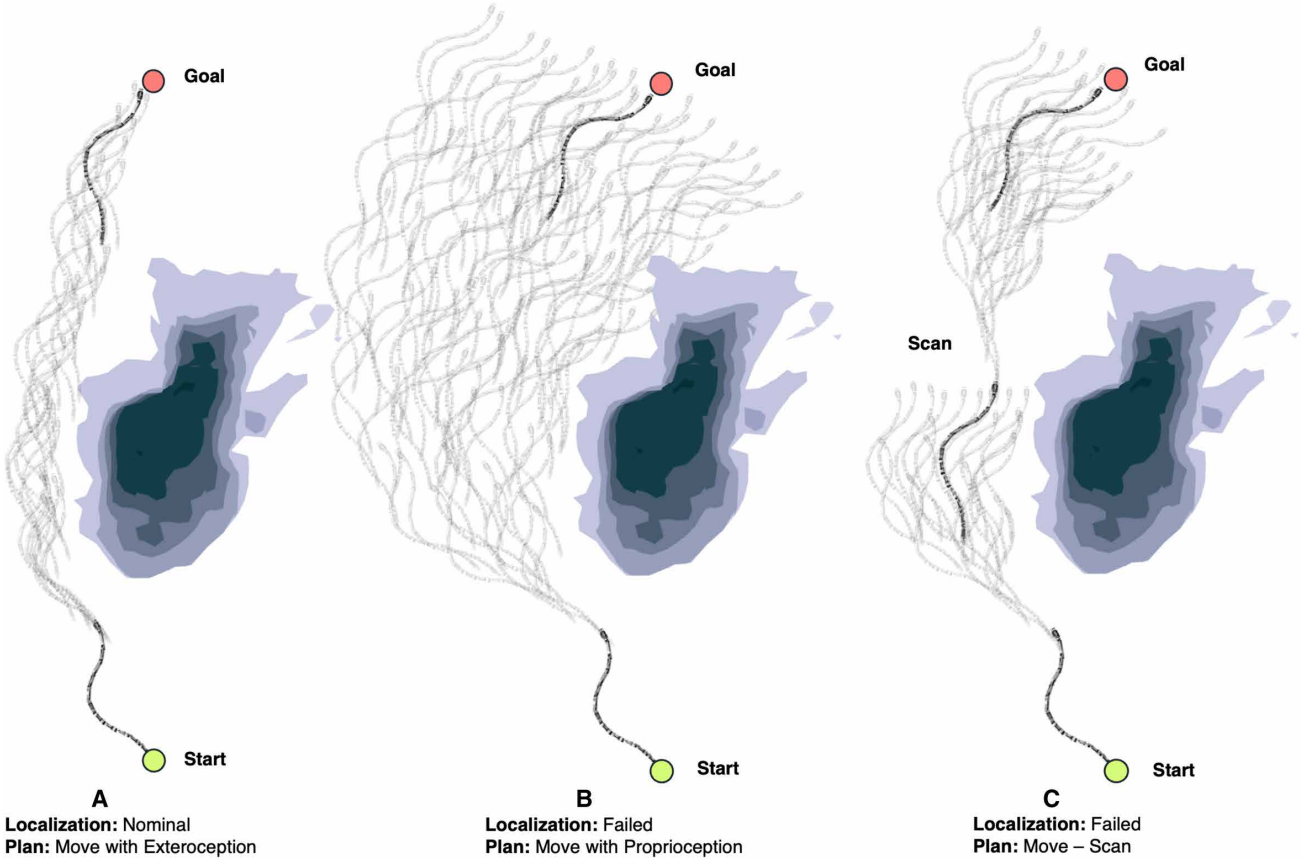


**Fig. 5. Integrated task and motion planning at the mission planning system.** Three surface mobility scenarios are shown, each with same start and goal location and a crevasse interposed between start and goal. State uncertainty is depicted through duplicating the mobility system and decreasing its opacity. (**A**) Nominal mobility with nominal exteroception, where uncertainty growth is low and the trajectory passes close to the crevasse. (**B**) Mobility under exteroception failures relying solely on proprioception. Under these conditions, uncertainty grows rapidly. (**C**) Mobility under exteroception failure with the task and motion planner enabled. Uncertainty grows with the same rate as for (B), but relocalizing actions allow the robot to take a less conservative and safer trajectory.

decoupled approach because that allows a trade-off between trajectory safety and information-gaining action cost. We framed this planning problem as a mixed integer linear programming (MILP) and tested it on hardware in a laboratory environment on a plastic-ice simulant as an analog to a hard ice surface.

### Summary of surface locomotion performance

Surface mobility was assessed through a campaign of laboratory and field tests in sand terrains at the JPL's Mars Yard and in ice and snow terrains in Big Bear, California. The surface mobility tests are described in detail in (*14*); here, we summarize the results. Tests were performed on flat ground, on inclined slopes, and in the presence of obstacles and holes. In addition, sand and snow were tested under both consolidated and unconsolidated conditions as shown in Fig. 6. The large variety of test terrain configurations highlighted the robot's adaptability enabled by a combination of screw- and shape-based locomotion. The primary gaits that were the focus of the test campaign were leader-follower and shape-based gaits. Both exteroceptive and proprioceptive control were demonstrated as effective locomotion methods. Screw-based locomotion allowed closed-loop path tracking for the leader-follower gait, and shape-based locomotion was shown to be a useful strategy to get the robot unstuck from situations that would have spelled the end of a traditional mobility system such as a rover. In addition, shape-based locomotion was shown as an effective way to navigate unconsolidated terrains, such as powder snow or fine sand, when screw locomotion proved to be less reliable.

It is worth highlighting that exteroceptive localization and mapping failed on multiple occasions in the field. These failures stemmed from a mix of hardware faults and interactions between exteroception algorithms with the environment. These failures emphasized the need for a proprioceptive control strategy and system-level autonomy capable of switching between control schemas. Exteroception failures were commonly caused by degeneracy in the environment and can be overcome by moving proprioceptively until more features become visible, allowing relocalization. In addition, when exteroception fails because of a hardware fault, there is still value in being able to continue with the mission, albeit with a more conservative path relying on proprioception.

### Summary of task and motion planning performance

In the task and motion planning system, we showed how reasoning over uncertainty, risk, and explicitly planning information-gaining actions allow operating in the presence of exteroceptive failures. In Fig. 7, we can see a representation of a plan in belief space in a 2.5-dimensional simulation environment. The planner is tasked with reaching a goal from a starting position and generates a sequence of control inputs, mixed with scan actions, accounting for uncertainty growth (represented in the figure as shaded circles). In the hardware system's tests, we compared an open-loop behavior, where the robot moves without exteroceptive feedback (no scan tasks are performed, and SLAM is not running), against the proposed risk-aware planner and observed that planning tasks jointly with the robot's motion yields considerably safer behavior while not engaging in excessively




**Fig. 6. Field test overview for surface scenarios.** Robot tests in different environments: (**A**) consolidated and unconsolidated snow in Big Bear, California; (**B**) ice in JPL's Table Mountain facility; and (**C**) consolidated and unconsolidated sand at JPL's Mars Yard.

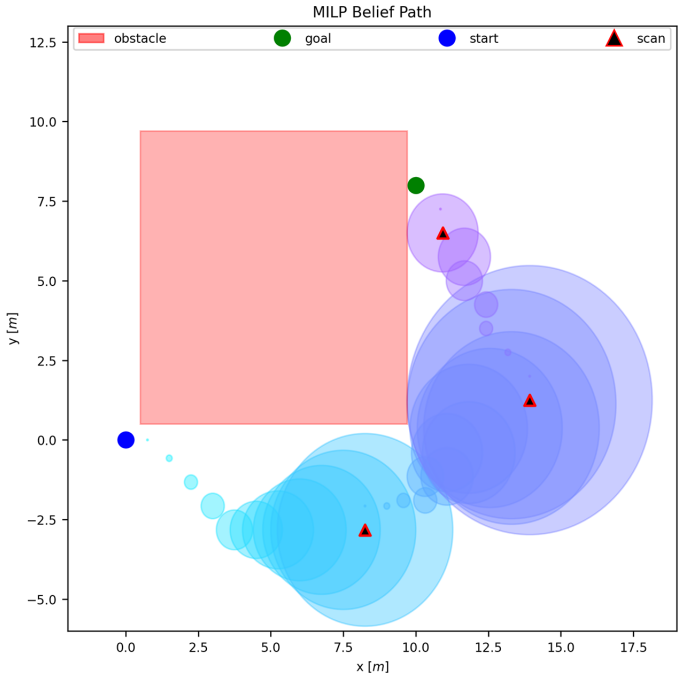


**Fig. 7. Risk-aware task and motion planner based on MILP formulation.** Belief space plan visualization. The agent is tasked with reaching a goal (green dot) from a starting point (blue dot) and plans a path represented by the moving circles with color gradient representing time. The diameter of each circle of the path represents predicted state uncertainty, and the points in space with black triangles are where the agent scheduled an information-seeking activity.

conservative behavior. Figure 8 highlights the difference between open-loop behavior and a risk-aware task and motion planner. We assume that the robot's exteroceptive localization is not functioning close to ground and that an information-gaining action can be performed by reconfiguring the sensor's position. For the EELS robot, these information-seeking behaviors consist of raising the robot's head and looking for features. We can see how the open-loop policy violates safety constraints, whereas the risk-aware policy reaches the goal by operating in open loop but gaining new information and replanning when state uncertainty has grown too much. In these laboratory settings, the task and motion planner was able to consistently find a plan within 10 s of being invoked on a base station (Intel i7 NUC at 2.65 GHz, 32-gigabyte memory). This task and motion planning problem is a subset of the full mission-planning problem that needs to be solved to operate on Enceladus' surface because the action space that we have considered is limited to movement and information-gaining actions. The full-mission planning problem will include tasks such as science collection, communication, and other failure recovery behaviors.

## DISCUSSION

This work has outlined the vision for the EELS project and the science questions it is designed to answer. Alongside these, we have described the system-level autonomy capabilities needed to enable an in situ, subglacial access mission to Enceladus and developed a hardware and software architecture intended to fulfill these requirements. The system we developed was field-deployed in relevant environments, where it showed the effectiveness and robustness of a combination of leader-follower and shape-based gaits and the benefits of decoupling the proprioceptive and exteroceptive control layers in the presence of exteroceptive failures. Furthermore, we have presented preliminary results for the task and motion planning system that has been built to aid the system in balancing information-gaining actions with trajectory selection.

### Limitations

Although the robot was capable of surface traversal over snow and ice, the autonomy/control challenge of matching gait selection with terrain conditions and executing a desired trajectory requires further work. In addition, subsurface mobility and the transition between surface and subsurface remain open research areas, and more work is needed to integrate surface and subsurface components into a single robotic system capable of safely executing a full mission from lander to subglacial ocean.

### Lessons learned from surface locomotion testing

A locomotion technology, such as EELS, tested in such a wide variety of terrains is bound to lead to valuable lessons learned about low-level locomotion assumptions ranging from the way that contact estimation should be performed to what makes for an effective compliance strategy. Assumptions about the way different screw geometries interact with various terrains have also been challenged and have led to insightful discoveries. Further details can be found in (*14*). The focus of this work is the lessons relevant to the system-level autonomy problem. Exteroception failures and the effect of degraded SLAM performance over closed-loop path tracking capability have guided the task and motion planner problem formulation. The possibility of getting stuck in unconsolidated material and the possibility to recover from these mobility failure conditions by changing locomotion strategy point to a need for a system-level coordinator capable of deciding how to recover from mobility failures. This has also been observed in previous robotic research and development efforts for underground exploration, specifically in the Defense Advanced Research Projects Agency (DARPA) Subterranean Challenge (*37*). A project-level lesson learned is that developing high-level autonomy features for field robotics in close contact with the lower-level problems leads to a greater understanding of the real mission's needs and thus positively affects capability development. In addition, fast-paced research projects are often subject to varying system and software requirements, and thus being intimately familiar with the project's evolving needs is another way to improve the high-level autonomy relevance and improve the chances of deploying autonomy features sooner. This is an argument in favor of adopting a primarily inductive research process, where experience from the field is used to narrow down the scope of the problems being worked on and then generalized into architectures and frameworks.

### Lessons learned from the task and motion planning testing

From a system-level autonomy algorithm design perspective, there is an intrinsic trade-off between scalability, model accuracy, and solution quality. The most accurate way of formulating sequential decision-making processes under imperfect-state information is as a partially observable Markov decision process (POMDP), which suffers from severe scalability issues—especially in continuous action or state spaces. POMDPs can be solved with approximate methods that introduce solution suboptimality and lack of guarantees, which are needed for highly complex, safety-critical flight missions. On the other side of the trade, excessively simple planners can be useful under certain conditions and are easy to study and verify, but, in most

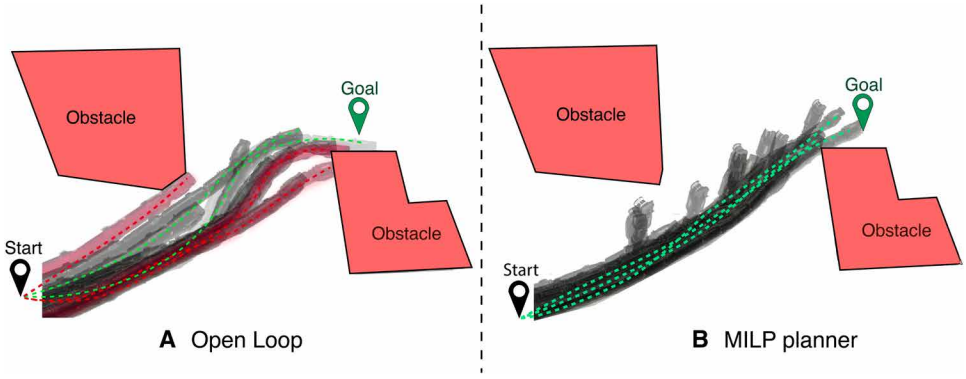


**Fig. 8. Mission planning system performance.** Risk-aware task and motion planner, compared against open-loop behavior. The figure shows multiple hardware runs laid over each other. The background has been removed, obstacles have been substituted with red geometric shapes, start and goal poses are marked with pins, and the robot's trajectories are represented as dashed lines. Trajectories that led to a collision are marked in red, whereas successful trajectories are in green. (**A**) How open-loop behavior in the presence of state uncertainty growth leads to several plan execution failures. (**B**) How the MILP risk-aware planner achieves its goal safely by scheduling information-gaining actions along the path.

cases, they do not fulfill the autonomy capability that would enable a subglacial access mission to be successfully undertaken. Moreover, finding a balance among model accuracy, performance, and solver guarantees is key in autonomous missions targeted for extreme, uncertain terrains. The task and motion planner designed for surface mobility under exteroception failures achieves that balance.

Although an EELS robot moves with a speed on the order of centimeters per second, fast reaction times in the control layer are required during vertical mobility and during surface mobility because of potential slips or other dynamic events. Strategic mission and path planning can occur at a much lower rate without added risk to the mission. The hierarchical organization of EELS' software architecture enabled reasoning and reactions at different timescales. This is a critical design decision that makes tractable and solvable a problem that would otherwise be too complex to formulate and solve. The control layer was designed to iterate hundreds of times per second, whereas the task and motion planner was invoked every 1 to 10 s. Another aspect that is worth considering is reaction to capability failures. These are system-level adaptations that are often time sensitive and need to occur with a latency in the order of 1 s. This requires an approach to high-level planning that meshes deliberation and reaction to failures. A disadvantage of hierarchical architectures is that system models at different levels might interact with each other unexpectedly if not carefully designed. An example that appeared in the formulation of the task and motion planning component is the treatment of nonholonomic constraints. We chose to neglect these constraints in the highest layers of reasoning but were justified in doing so only under conditions where the high-level planner is reasoning over much larger spatial scales than the range in which these constraints matter. For example, if the nonholonomic constraints have influence over a scale of meters, then the task and motion planner should reason over a scale of hundreds of meters or kilometers.

## MATERIALS AND METHODS

### Surface locomotion strategies

Specific details on implementation, approaches, and algorithms used in surface mobility can be found in (*14*). In this section, we provide a summary of methods used to verify the locomotion systems' performance. Field tests have proven to be the best way to validate control assumption and assess the true performance characteristics of the system. Although simulations are a critical component of software integration testing and preliminary control development, only field tests can reveal unknown unknowns, validate modeling assumptions, and help with requirement discovery.

In our tests, we used the following software components that make up the ingredients for surface mobility: proprioceptive and exteroceptive estimators, path planner, shape and screw control, and joint-level control. The robot was exposed to a wide variety of terrains through three separate field locations and one laboratory setting. The field tests were performed in JPL's Table Mountain facility in California, with slopes up to 35° and a mixture of ice, consolidated snow, and microterrain features. More snow tests were performed at Big Bear, California, during a snowstorm that tested the robotic platform's capacity to navigate unconsolidated snow and operate under harsh environmental conditions. Sand, both consolidated and unconsolidated, was tested at JPL's Mars Yard, California. Field test results were qualitatively analyzed, taking into account distance traveled and average movement speed as a way to determine the robot's capability to traverse a diverse set of microterrain features. Further mobility tests were performed in a laboratory environment on a planar surface made of synthetic ice sheets. Start, goal, and obstacles were kept consistent throughout laboratory tests. Path tracking error, distance traveled, average speed, and the number of manual interventions were used as qualitative performance metrics. A detailed description of mobility results and testing methodology can be found in (*14*). Although Enceladus' surface might be much less cohesive than sand or snow (*38*), the field environments in which we tested are a crucial step toward demonstrating mobility in environments where prior knowledge is scarce, by demonstrating the robustness on different terrains.

### Task and motion planning module

In this work, we consider the following inputs for the task and motion planning system: (i) operator-specified goals in the form of a set of surface waypoints for the robot to visit in the given sequence; (ii) a map of the environment and the uncertainty model that represents how the uncertainty grows as we navigate the environment; and (iii) a threshold on the probability of violating safety constraint (hitting a obstacle). On the basis of this set of inputs and the action space available (move and scan), our planning system will generate a plan of tasks and execute that plan by dispatching tasks and monitoring progress. There are many ways to formulate such a task and motion planning problem. Commonly, this type of decision-under-uncertainty problem is framed as a POMDP (*39*). Recent search algorithm advancements in tree search (*40*) and reinforcement learning (*41*) have allowed for solving very large problems. Because of the importance of solver

guarantees, we chose to formulate this problem using an MILP approach. This mathematical programming paradigm, commonly used in operations research (*42*), has been adapted for trajectory planning under uncertainty (*43*, *44*). Differently from other approaches, we combined task and motion planning in a continuous domain while accounting explicitly for risk and uncertainty and finding the optimal balance between cost of moving (path length) and the cost of gathering information (scanning the environment). In our MILP formulation, we modeled constraints that represent action-dependent state and uncertainty propagation, risk bounds on the probability of collision with polygonal obstacles, start and end conditions, and a set of supporting constraints that avoid solutions/plans that go over obstacle corners. The resulting mathematical program was solved with the Gurobi solver, chosen because of its efficient MILP-solving heuristics. Solution times for problems with ~50 time steps and three obstacles were on the order of 10 s of compute time on a gigahertz class, 12-core central processing unit.

We tested the MILP planner in a 2.5-dimensional simulation environment for quantitative insight, comparing it against a state-of-the-art POMDP formulation solved through Monte Carlo tree search and a two-stage approach that decouples task from motion planning (see Fig. 7 for an example of a simulated scenario), and then deployed it on hardware in a laboratory environment over an ice simulant material (synthetic ice sheets). We ran multiple tests with consistent start and goal positions and observed trajectories and success rates qualitatively, i.e., whether the robot achieved the goal position. Three distinct maps were tested in simulation, each with a different number of obstacles. Start and goal poses were placed ~14 m from each other. The maps have, respectively, a single rectangular obstacle, two obstacles arranged into an L shape, and three rectangular obstacles arranged into an irregular pattern forming narrow corridors.

The task and motion planner's laboratory experiments were conducted with constant start and goal positions. The goal was located ~1 body length in front of the head module's starting position because of laboratory space constraints. A single artificial obstacle was placed between start and end goal, and the building's walls and operations station delimited the robot's workspace, funneling the available paths. See Fig. 8 for an overview of the laboratory environment.

The EELS project is poised to focus next on mission planning and execution for vertical mobility. Further iterations of risk estimation and failure detection, isolation, and recovery are also under development and will be integrated with the mission planning layer. Further work needs to be done to achieve multiple-time-scale, high-level planning capable of meshing reactions to failures with deliberative task and motion planning. Longer term goals include the unification of surface with subsurface mobility into a single high-level task and motion planning formulation and the demonstration of fully autonomous operations in a realistic mock mission.

## REFERENCES AND NOTES


1. C. J. Hansen, L. Esposito, A. I. F. Stewart, J. Colwell, A. Hendrix, W. Pryor, D. Shemansky, R. West, Enceladus' water vapor plume, *Science* **311**, 1422–1425 (2006).
2. L. less, D. J. Stevenson, M. Parisi, D. Hemingway, R. A. Jacobson, J. I. Lunine, F. Nimmo, J. W. Armstrong, S. W. Asmar, M. Ducci, P. Tortora, The gravity field and interior structure of Enceladus. *Science* **344**, 78–80 (2014).
3. C. C. Porco, P. Helfenstein, P. C. Thomas, A. P. Ingersoll, J. Wisdom, R. West, G. Neukum, T. Denk, R. Wagner, T. Roatsch, S. Kieffer, E. Turtle, A. McEwen, T. V. Johnson, J. Rathbun, J. Veverka, D. Wilson, J. Perry, J. Spitale, A. Brahic, J. A. Burns, A. D. Del Genio, L. Dones, C. D. Murray, S. Squyres, Cassini observes the active south pole of Enceladus. *Science* **311**, 1393–1401 (2006).
4. P. C. Thomas, R. Tajeddine, M. S. Tiscareno, J. A. Burns, J. Joseph, T. J. Loredo, P. Helfenstein, C. Porco, Enceladus's measured physical libration requires a global subsurface ocean. *Icarus* **264**, 37–47 (2016).
5. W. B. McKinnon, Effect of Enceladus's rapid synchronous spin on interpretation of Cassini gravity. *Geophys. Res. Lett.* **42**, 2137–2143 (2015).
6. F. Postberg, S. Kempf, J. Schmidt, N. Brilliantov, A. Beinsen, B. Abel, U. Buck, R. Srama, Sodium salts in E-ring ice grains from an ocean below the surface of Enceladus. *Nature* **459**, 1098–1101 (2009).
7. J. H. Waite, C. R. Glein, R. S. Perryman, B. D. Teolis, B. A. Magee, G. Miller, J. Grimes, M. E. Perry, K. E. Miller, A. Bouquet, J. I. Lunine, T. Brockwell, S. J. Bolton, Cassini finds molecular hydrogen in the Enceladus plume: Evidence for hydrothermal processes. *Science* **356**, 155–159 (2017).
8. M. Chodas, M. Ono, J. Weber, L. Rodriguez, M. Ingham, B. Hockman, K. Mitchell, M. Cable, J. Rabinovitch, Enceladus vent explorer mission architecture trade study, in *Proceedings of IEEE Aerospace Conference* (IEEE, 2023), pp. 1–16.
9. J. M. Weber, L. E. Rodriguez, M. Ono, M. Cable, M. Chodas, M. Ingham, A general STM for the Enceladus vent explorer mission concept and a case for vent exploration, paper presented at the 54th Lunar and Planetary Science Conference, Woodlands, TX, 13 to 17 March 2023.
10. E. S. Kite, A. M. Rubin, Sustained eruptions on Enceladus explained by turbulent dissipation in tiger stripes. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 3972–3975 (2016).
11. K. L. Mitchell, M. Ono, C. Parcheta, S. Iacoponi, Dynamic pressure at Enceladus' vents and implications for vent and conduit in-situ studies, paper presented at the 48th Lunar and Planetary Science Conference, Woodlands, TX, 20 to 24 March 2017.
12. M. Nakajima, A. P. Ingersoll, Controlled boiling on Enceladus. 1. Model of the vapor-driven jets. *Icarus* **272**, 309–318 (2016).
13. A. P. Ingersoll, M. Nakajima, Controlled boiling on Enceladus. *Icarus* **272**, 319–326 (2016).
14. R. Thakker, M. Paton, M. P. Strub, M. Swan, G. Daddi, R. Royce, P. Tosi, M. Gildner, T. Vaquero, M. Veismann, P. Gavrilov, E. Marteau, J. Bowkett, D. Loret, Y. Nakka, B. Hockman, A. Orekhov, T. Hasseler, C. Leake, B. Nuernberger, P. Proença, W. Reid, W. Talbot, N. Georgiev, T. Pailevanian, A. Archanian, E. Ambrose, J. Jasper, R. Etheredge, C. Roman, D. Levine, K. Otsu, H. Melikyan, J. Nash, R. Rieber, K. Carpenter, A. Jain, L. Shiraishi, D. Pastor, S. Yearicks, M. Ingham, M. Robinson, A. Agha, M. Travers, H. Choset, J. Burdick, M. Ono, EELS: Towards autonomous mobility in extreme environments with a novel large-scale screw driven snake robot, in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2023), pp. 9886–9893.
15. S. Ulamec, J. Biele, O. Funke, M. Engelhardt, Access to glacial and subglacial environments in the Solar System by melting probe technology. *Rev. Environ. Sci. Biotechnol.* **6**, 71–94 (2007).
16. K. Konstantinidis, C. Martinez, B. Dachwald, A. Ohndorf, P. Dykta, P. Bowitz, M. Rudolph, I. Digel, J. Kowalski, K. Voigt, R. Förstner, A lander mission to probe subglacial water on Saturn's moon Enceladus for life. *Acta Astronaut.* **106**, 63–89 (2015).
17. W. C. Stone, B. Hogan, V. Siegel, S. Lelievre, C. Flesher, Progress towards an optically powered cryobot. *Ann. Glaciol.* **55**, 2–13 (2014).
18. K. Zacny, J. Mueller, T. Costa, T. Cwik, A. Gray, W. Zimmerman, P. Chow, F. Rehnmark, G. Adams, SLUSH: Europa hybrid deep drill, in *IEEE Aerospace Conference* (IEEE, 2018), pp. 1–14.
19. M. Ono, K. Mitchel, A. Parness, K. Carpenter, S. Iacoponi, E. Simonson, A. Curtis, M. Ingham, C. Budney, T. Estlin, C. Parcheta, R. Detry, J. Nash, J. P. de la Croix, J. Kawata, K. Hand, Enceladus vent explorer concept, in *Outer Solar System: Prospective Energy and Material Resources* (Springer, 2018), pp. 665–717.
20. I. Nesnas, L. M. Fesq, R. A. Volpe, Autonomy for space robots: Past, present, and future. *Curr. Robot. Rep.* **2**, 251–263 (2021).
21. J. A. Starek, B. Açıkmeşe, I. A. Nesnas, M. Pavone, Spacecraft autonomy challenges for next-generation space missions, in *Advances in Control System Technology for Aerospace Applications* (Springer, 2016), pp. 1–48.
22. P. Liljebäck, K. Y. Pettersen, Ø. Stavdahl, J. T. Gravdahl, A review on modelling, implementation, and control of snake robots. *Robot. Auton. Syst.* **60**, 29–40 (2012).
23. K. Y. Pettersen, Snake robots. *Annu. Rev. Control* **44**, 19–44 (2017).
24. F. Richter, P. V. Gavrilov, H. M. Lam, A. Degani, M. C. Yip, ARCSnake: Reconfigurable snakelike robot with Archimedean screw propulsion for multi-domain mobility. *IEEE Trans. Robot.* **38**, 797–809 (2022).
25. T. Dachlika, D. Zarrouk, Mechanics of locomotion of a double screw crawling robot. *Mech. Mach. Theory* **153**, 104010 (2020).
26. K. Nagaoka, M. Otsuki, T. Kubota, S. Tanaka, Terramechanics-based propulsive characteristics of mobile robot driven by Archimedean screw mechanism on soft soil, in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2010), pp. 4946–4951.
27. M. Rashid, M. Yakub, S. Salim, N. Mamat, S. Putra, S. Roslan, Modeling of the in-pipe inspection robot: A comprehensive review. *Ocean Eng.* **203**, 107206 (2020).
28. A. Barchowsky, A. Amirahmadi, K. Botteon, G. Carr, C. Jin, P. McGarey, S. Sposato, S. Yang, A high voltage tethered power system for planetary surface applications, in *IEEE Aerospace Conference* (*AERO*) (IEEE, 2022), pp. 1–8.

29. M. Ingham, R. Rasmussen, M. Bennett, A. Moncada, Engineering complex embedded systems with state analysis and the mission data system. *J. Aerosp. Comput. Inf. Commun.* **2**, 507–536 (2004).
30. I. A. Nesnas, R. Simmons, D. Gaines, C. Kunz, A. Diaz-Calderon, T. Estlin, R. Madison, J. Guineau, M. McHenry, I.-H. Shu, D. Apfelbaum, CLARAty: Challenges and steps toward reusable robotic software. *Int. J. Adv. Robot. Syst.* **3**, 23–30 (2006).
31. R. Amini, L. Fesq, R. Mackey, F. Mirza, R. Rasmussen, M. Troesch, K. Kolcio, FRESCO: A framework for spacecraft systems autonomy, in *IEEE Aerospace Conference (50100)* (IEEE, 2021), pp. 1–18.
32. J. Ocón, F. J. Colmenero, K. Buckley, M. Alonso, E. Heredia, J. Garcia, A. I. Coles, A. J. Coles, M. Martínez, E. Savaş, F. Pommerening, T. Keller, S. Karachalios, M. Woods, I. Dragomir, S. Bensalem, P. Dissaux, A. Schach, The ERGO framework and its use in planetary/orbital scenarios, in *Proceedings of the 69th International Astronautical Congress* (*IAC*) (International Astronautical Federation, 2018), pp. 1–13.
33. M. Golombek, D. Rapp, Size-frequency distributions of rocks on Mars and Earth analog sites: Implications for future landed missions. *J. Geophys. Res. Planets* **102**, 4117–4129 (1997).
34. C. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, T. Lozano-Pérez, Integrated task and motion planning. *Annu. Rev. Control Robot. Auton. Syst.* **4**, 265–293 (2021).
35. S. Chien, R. Sherwood, D. Tran, B. Cichy, G. Rabideau, R. Castano, A. Davies, D. Mandl, S. Frye, B. Trout, S. Shulman, D. Boyer, Using autonomy flight software to improve science return on earth observing one. *J. Aerosp. Comput. Inf. Commun.* **2**, 196–216 (2005).
36. D. Gaines, S. Chien, G. Rabideau, S. Kuhn, V. Wong, A. Yelamanchili, S. Towey, J. Agrawal, W. Chi, A. Connell, E. Davis, C. Lohr, Onboard planning for Mars 2020 Perseverance rover, in *Proceedings of the 16th Symposium on Advanced Space Technologies in Robotics and Automation* (European Space Agency, 2022), pp. 1–5.
37. A. Agha, K. Otsu, B. Morrell, D. D. Fan, R. Thakker, A. Santamaria-Navarro, S.-K. Kim, A. Bouman, X. Lei, J. Edlund, M. F. Ginting, K. Ebadi, M. Anderson, T. Pailevanian, E. Terry, M. Wolf, A. Tagliabue, T. S. Vaquero, M. Palieri, S. Tepsuporn, NeBula: Quest for robotic autonomy in challenging environments; TEAM CoSTAR at the DARPA Subterranean challenge. arXiv:2103.11470 [cs.RO] (18 October 2021).
38. E. S. Martin, J. L. Whitten, S. A. Kattenhorn, G. C. Collins, B. S. Southworth, L. S. Wiser, S. Prindle, Measurements of regolith thicknesses on Enceladus: Uncovering the record of plume activity. *Icarus* **392**, 115369 (2023).
39. H. Kurniawati, Partially observable Markov decision processes and robotics. *Annu. Rev. Control Robot. Auton. Syst.* **5**, 253–277 (2022).
40. D. Silver, J. Veness, Monte-Carlo planning in large POMDPs. *Adv. Neural Inf. Process. Syst.* **23**, 2164–2172 (2010).
41. S. Mousavi, M. Schukat, E. Howley, Deep reinforcement learning: An overview, in *Proceedings of SAI Intelligent Systems Conference* (*IntelliSys*) (Springer International Publishing, 2018), vol. 2, pp.426–440.
42. C. Floudas, L. Xiaoxia, Mixed integer linear programming in process scheduling: Modeling, algorithms, and applications. *Ann. Oper. Res.* **139**, 131–162 (2005).
43. D. Ioan, I. Prodan, S. Olaru, F. Stoican, S. I. Niculescu, Mixed-integer programming in motion planning. *Annu. Rev. Control.* **51**, 65–87 (2021).
44. M. da Silva Arantes, C. F. M. Toledo, B. C. Williams, M. Ono, Collision-free encoding for chance-constrained nonconvex path planning. *IEEE Trans. Robot.* **35**, 433–448 (2019).



**Acknowledgments**

**Funding:** The research was carried out at the JPL, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). **Author contributions:** The EELS project's success is attributed to a collaborative effort across various teams. T.S.V., G.D., R.T., M.P., A.J., M.P.S., R.M.S., R.R., J.B., D.L.d.M.L., Y.N., B.H., A.O., T.D.H., C.L., B.N., P.P., W.R., W.T., D.L., K.O., J.N., and C.R. contributed to the development of EELS' software stack. The system's active skin was studied by P.T., M.V., and E.M. EELS' hardware platform was developed by N.G., M.G., P.G., T.P., A.A., E.A., and H.M. M.O. and K.C. played a pivotal role in the project as the principal investigators. Project management was performed by M.R. and R.E. The project's scientists were A.G. and M.C. Logistics and simulants were managed by S.Y. Systems engineering was performed by B.H., R.R.R., and J.J. A.J., L.S., M.T., H.C., J.B., and M.I. advised the team. The task and motion planning team was led by T.S.V., G.D., R.T., A.J., and M.I. This manuscript was prepared and edited jointly by T.S.V., G.D., and M.O. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper.

Submitted 22 March 2023
Accepted 14 February 2024
Published 13 March 2024
10.1126/scirobotics.adh8332

Science Robotics

# EELS: Autonomous snake-like robot with task and motion planning capabilities for ice world exploration

T. S. Vaquero, G. Daddi, R. Thakker, M. Paton, A. Jasour, M. P. Strub, R. M. Swan, R. Royce, M. Gildner, P. Tosi, M. Veismann, P. Gavrilov, E. Marteau, J. Bowkett, D. Loret de Mola Lemus, Y. Nakka, B. Hockman, A. Orekhov, T. D. Hasseler, C. Leake, B. Nuernberger, P. Proença, W. Reid, W. Talbot, N. Georgiev, T. Pailevanian, A. Archanian, E. Ambrose, J. Jasper, R. Etheredge, C. Roman, D. Levine, K. Otsu, S. Yearicks, H. Melikyan, R. R. Rieber, K. Carpenter, J. Nash, A. Jain, L. Shiraishi, M. Robinson, M. Travers, H. Choset, J. Burdick, A. Gardner, M. Cable, M. Ingham, and M. Ono

*Sci. Robot.* **9** (88), eadh8332. DOI: 10.1126/scirobotics.adh8332

**Editor's summary**

There is growing interest in the exploration of icy moons, such as Enceladus, which may have astrobiological implications. However, obtaining samples is challenging because of the environmental extremities on the surface or within ice vents. Vaquero *et al.* developed a snake-like robot named Exobiology Extant Life Surveyor (EELS) that was capable of autonomously navigating on icy surfaces. EELS has a perception head that contains a series of sensors and cameras to observe its environment, whereas the body has articulated segments for shape-changing and a screw-like outer surface to enable motility. EELS shows potential for risk-aware autonomous exploration of complex icy terrain. —Amos Matsiko







*Science Robotics* (ISSN 2470-9476) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.