Expert-Informed Autonomous Science Planning for In-situ Observations and Discoveries

Jay McMahon
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
jay.mcmahon@colorado.edu

Nisar Ahmed
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
nisar.ahmed@colorado.edu

Morteza Lahijanian
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
morteza.lahijania@colorado.edu

Peter Amorese
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
peter.amorese@colorado.edu

Taralicin Deka
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
taralicin.deka@colorado.edu

Karan Muvvala
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
karan.muvvala@colorado.edu

Trevor Slack
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
trevor.slack@colorado.edu

Shohei Wakayama
University of Colorado Boulder
3775 Discovery Drive
Boulder, CO 80303
shohei.wakayama@colorado.edu

Abstract— **Future planetary exploration missions on the surface of distant bodies such as Europa or Enceladus can't rely on human-in-the-loop operations due to time delays, dynamic environments, limited mission lifetimes, as well as the many unknown unknowns inherent in the exploration of such environments. Thus our robotic explorers must be capable of autonomous operations to ensure continued operations and to try to maximize the amount and quality of the scientific data gathered from each mission. To advance our technology toward this goal, we are developing a system to maximize the science obtained by a robotic lander and delivered to scientists on Earth with minimal asynchronous human interaction. The autonomy architecture consists of three main components: Shared Science Value Maps (SSVMs), which function as an interface between REASON (Robust Exploration with Autonomous Science on-board) and RECOURSE (Ranked Evaluation of Contingent Opportunities for Uninterrupted Remote Science Exploration) for efficient and useful scientific communication between scientists and robot. The key advantage to this design is in its ability to continuously operate and adapt despite the constraints of high-latency, low-bandwith communications and an uncertain environment which today would require ground-in-the-loop operations. This paper presents the overview of our architecture and initial results on the development of such a system. These results will focus on progress made in developing the details of the SSVM interface between human scientists and robotic explorer and the ability of REASON to act on the SSVM to develop plans on-board that attempt to maximize science obtained while being guaranteed to respect any relevant system and safety constraints.**

Table of Contents

978-1-6654-3760-8/22/$31.00 ©2022 IEEE

1. Introduction

Science operations with remote autonomy currently focus on managing highly limited vehicle resources as well as operational impacts from direct and indirect task couplings. This has led to an increased dependency on ground-based human analysis and decision-making in such missions. Automated software tools and human-machine interfaces for operational planning and decision-making support have cut down on the overall effort required. Examples include autonomous science data collection software suites, e.g. AEGIS for MER [1], [2], smart geological feature detectors [3], [4], [5], AI-based activity planners and scheduling systems like ASPEN, CASPER, OASIS, [6], [7], [8], [9] and for Perseverance [11], [10], [9], [12], [13].

There have been various levels of autonomy implemented in spacecraft systems over the years, ranging from guidance navigation and control applications (e.g. [16], [14], [19], [15], [18], [17]) to autonomous operations and science (e.g. [21], [20], [22]). On deep space scientific missions, like OSIRIS-REx [28], Rosetta [30], [29], or New Horizons [31], [33], [32], [34] considerable expense and effort is put into planning observations on the ground and uploading them to the spacecraft which are then executed on-board in an open-loop framework with minimal autonomy. If something abnormal occurs, observations are missed or the spacecraft may enter safe mode which causes significant operation delays.

Overall, capable systems for particular autonomous tasks already exist and they let ground operations manage the traditional 'single-pathway' science plans in relatively well-understood environments. However, such systems are not extendable to missions in remote unexplored environments like ocean worlds and icy moons, where opportunistic multi-contingency autonomous on-board decision-making is needed in the face of evolving uncertainties and science. Furthermore, a fundamental issue missing in the state-of-the-art is a way for scientists to be able to naturally "talk" to their robotic counterparts without the barrier of complex sequencing and command interfaces.

Our proposed solution, in this paper, seeks to address the main issues that prevent current work from being put into use on an ocean worlds exploration lander. We propose to

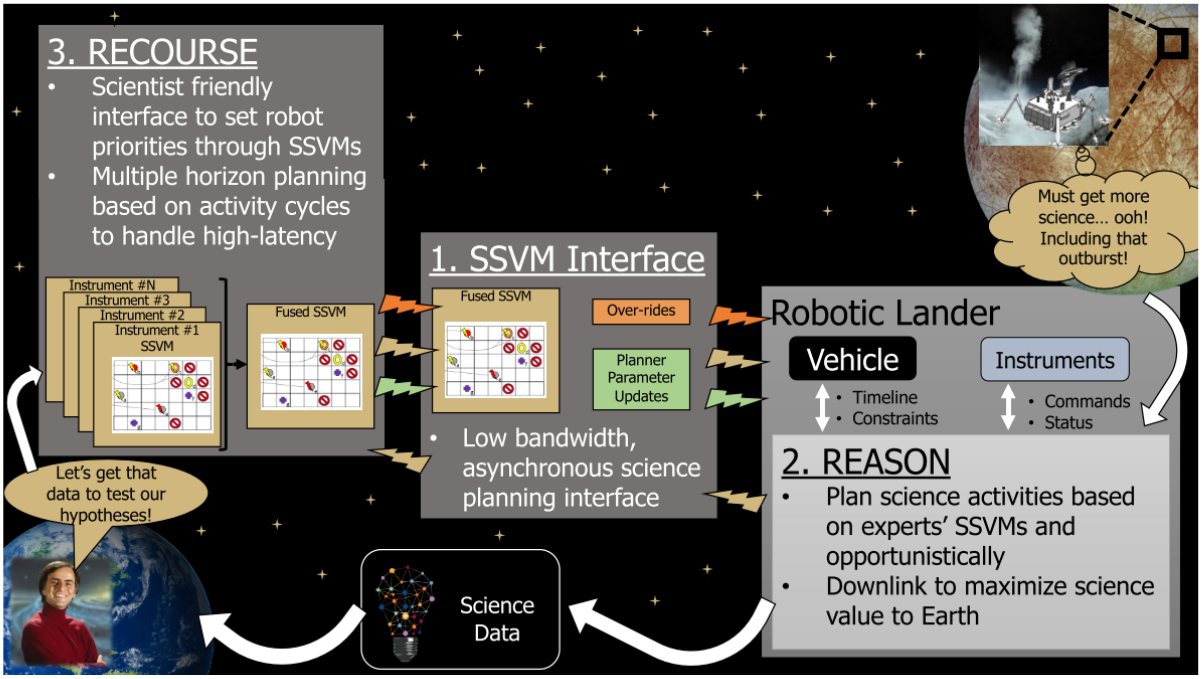

Figure 1: The proposed effort will design and implement (1) the SSVM interface between (2) the REASON on-board planner and (3) the RECOURSE scientist ground system

design, develop, and demonstrate an autonomy architecture illustrated in Fig. 1 that consists of three main components: **SSVM** (Shared Science Value Maps), which function as an interface between **REASON** (Robust Exploration with Autonomous Science on-board) and **RECOURSE** (Ranked Evaluation of Contingent Opportunities for Uninterrupted Remote Science Exploration) using for efficient and useful scientific communication between scientists and robot. In particular, the combination of REASON and RECOURSE using SSVMs as an interface has been explicitly designed to successfully handle high-latency, low-bandwidth communications while incorporating scientists' inputs and continuously gathering and downlinking valuable science data. Our solution focuses on creating a near-future flight system implementation that provides confidence and transparency about how the autonomous system is performing and the ability for operators to update the performance if desired. Our system effectively does not allow safe-modes to stop science, since it always has further actions prioritized and ready to execute.

This paper outlines our initial design decisions and investigations into implementation of our proposed architecture. As such, we discuss in turn each of the main components - REASON, SSVMs, and RECOURSE - and finally show some initial implementation results in NASA's OceanWATERS virtual testbed.

2. REASON

REASON (Robust Exploration with Autonomous Science on-board) is the on-board component of our proposed autonomy architecture. This system is responsible for all the critical tasks taking place aboard the lander, all of which can be broadly categorized into two main modules, namely: 1) Science Activity Planner, and 2) Downlink planner, as shown in Fig.2.

Science Activity Planner

The Science Activity Planner is responsible for autonomously determining both a high-level discrete sequence of actions (task plan) as well as the low-level continuous trajectory (motion plan) of moving parts of the lander to allow the system to efficiently interact with the environment and collect more scientific data per activity segment. The sequence of actions satisfy the formally specified tasks given to the system, while also accounting for the current world-configuration and environment of the lander [42]. To obtain an abstraction of the lander, we first discretize the actions of the system to a set of motion primitives specific to each instrument. The plan from the high-level planner is then used as a guide for low-level continuous planner for each motion primitive [43]. For some specific action primitives, such as moving the robotic manipulator from a start configuration to a goal configuration, it is critical to generate a continuous trajectory that also avoids collision with any obstacles in the environment. Using the Science Activity Planner we can readily interpret the high-level, verify the safety of the low-level plan, and guarantee task completion at the high-level. The Science Activity Planner is comprised of two main planning components: the task planner and the motion planner as shown in Fig. 3.

Task Planner—A scientific expert can formally define a task specification that will define a goal. The task planner (High-Level Planner) takes in a task goal (specification) and return a discrete sequence of actions that the real system can then execute. The task specification is generated by the scientists on ground in the form of a temporal logic formula such as LTLf (Linear Temporal Logic over Finite Traces) [44]. Using a task in the form of a logical formula is beneficial for unambiguous interpretation, which allows the scientific experts to mathematically decide if a task has been completed or violated. A scientist may have a conceptual idea of what the system should accomplish, however it is necessary convert this idea into a logical task specification that can only

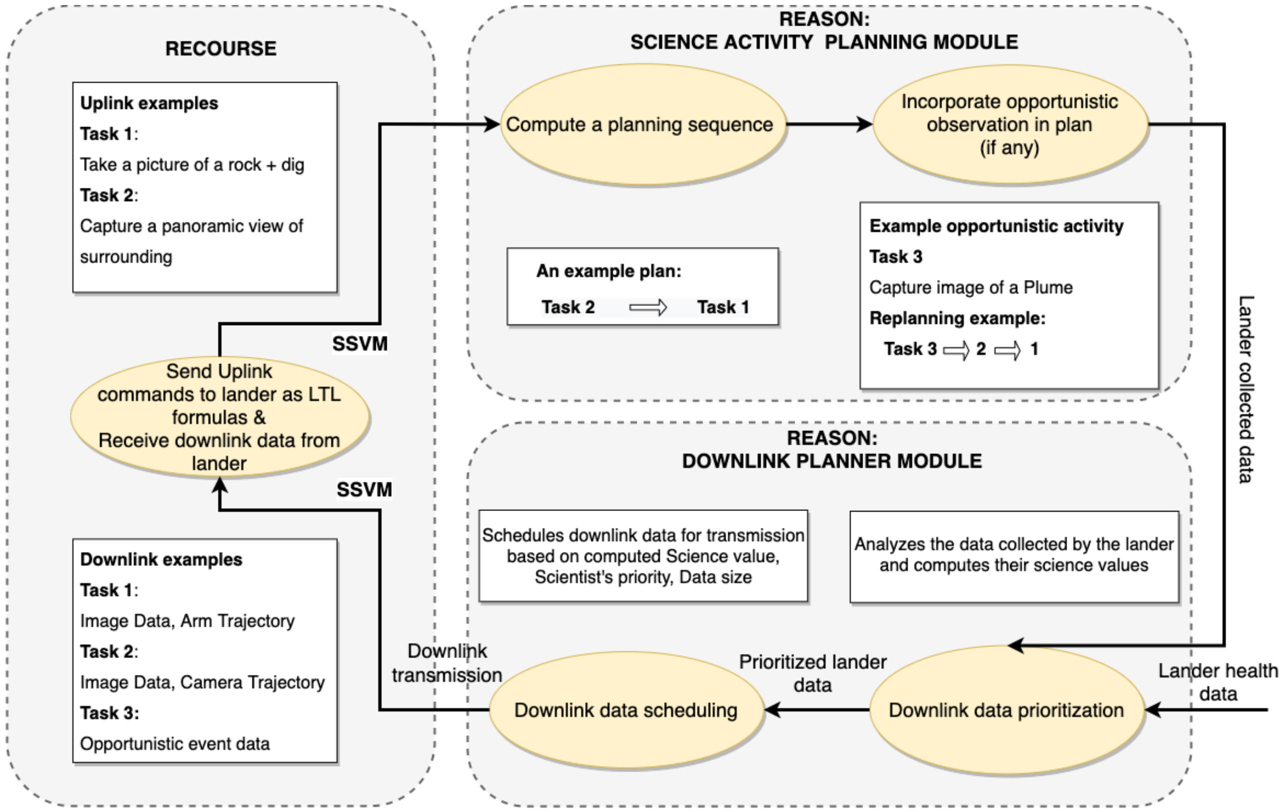

Figure 2: Data flow between different components of the proposed autonomous system

be interpreted one way. Temporal logic refers to logic that can reason over time. Specifically, LTLf refers to tasks that can be satisfied in finite time. Using a temporal logical specification allows the scientist to be much more expressive with what the system should accomplish and how the system should go about accomplishing the task [49] [50] [51] .

In order for the task planner to combine a task specification with the physical constraints of the autonomous system and environment, a discrete system model must be provided. This model, referred to as the *system abstraction* [45] [46], should reason over discrete states that include information about the status of the lander, as well as parts of the environment that the lander interacts with. Transitions between states in this model determine the physical capability and constraints of the system. These transitions can take the form of discrete actions, referred to as an *motion primitive*.

Using both the temporal logic task specification as well as the system abstraction, the task planner determines a sequence of high-level actions that will take the system from the current (initial) state to a final state in a manner that satisfies the given task specification without violating any task-specific and physical constraints of the system.

Motion Planner—The motion planner (Low-Level Planner) is responsible for low-level continuous path planning. The lander may have tools equipped that have moving parts and require low-level controls. The purpose of the motion planner is to determine a continuous trajectory for a tool to follow that will prevent collision with other parts of the lander itself, or obstacles in the environment. Using a motion planner can provide safety guarantees for the determined trajectory. The lander concept being used in the OceanWATERS (Ocean

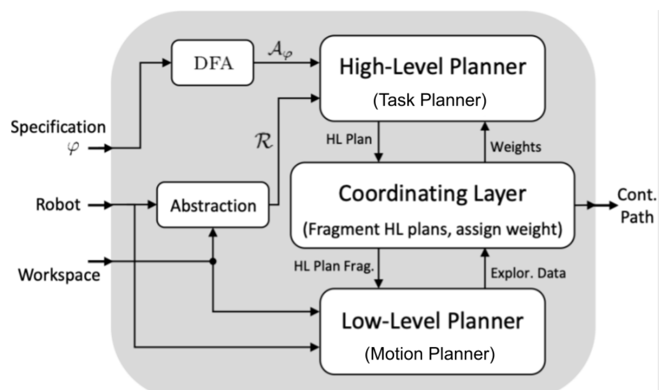

Figure 3: Synergistic Planning Diagram from [43].

Worlds Autonomy Testbed for Exploration Research & Simulation, [41]) comes equipped with a robotic manipulator, and a movable camera and communication module. For safely controlling the robotic manipulator, it is assumed that the motion planner is provided with information of obstacle boundaries, including collision boundaries of the lander. The motion planner will look for a collision free solution between the start configuration and a goal configuration. A sampling-based motion planning algorithm, such as RRT, is used to determine the trajectory for the manipulator [47] [48]. For the camera and communication module, a simple linear trajectory can be implemented that can pan and tilt the apparatus to the desired pose, while staying within the safety limits. For the purposes of this architecture, it is assumed that the movable tools on the lander also have stable low-level controllers that can safely follow a desired trajectory.

Synergistic Planning Framework—The purpose of the Synergistic Planning Framework [43] is to allow the task planner to better reason about the physical capabilities of the robotic manipulator and to accomplish tasks more smartly. For example, the system model might include an action to move the robotic manipulator between two states, however the transition might be physically infeasible due to the placement of obstacles. A naive approach is to replan for this trajectory until we find a valid motion plan to execute this motion primitive. This approach suffers from computation time and fails to account for difficulty is planning this particular continuous trajectory.

The synergistic framework introduces a *Coordinating Layer* which allows for the low-level motion planner to inform the high-level task planner and reasons quantitatively about each motion primitive. In this framework, the high-level planner computes and feeds an initial action sequence to the motion planner. If the motion planner fails when trying to plan for a certain action, the cost of executing that action is increased in the task planner. Since the task planner searches for an optimal sequence of actions to complete the task, increasing the cost of executing a certain action will encourage the high level planner to avoid planning for that action. This process is repeated until a sequence of actions is found where all actions in the sequence have a low-level motion planning solution. The Synergistic Planning Framework visualized in Fig. 3 displays the interaction between the task planner and the motion planner [43].

Downlink planner

The goal of the downlink planner is to maximize the science information sent back to Earth, while successfully handling the high-latency and low-bandwidth communication limitations between the lander and the Earth-based ground stations. This module is responsible for two main tasks on-board the lander: Downlink data prioritization and Downlink data scheduling.

Downlink data prioritization—After the lander has collected data according to the plan generated by the science activity planner (as described earlier), the downlink planning module will compute the science value of the collected data onboard. This process is called downlink data prioritization. Downlink prioritization can be achieved by methods such as *Target Signature* and *Novelty detection* methods [52], [53], [54]. Methods for downlinking selective data or processed data (such as *Image Masking*) are also being considered to reduce downlink bandwidth requirements [53], [55].

Downlink data scheduling— After science value has been assigned to the data collected by the lander (using a suitable downlink prioritization scheme), this data will be scheduled for downlinking on the basis of a weighted measure between science value, scientist's preference and data size.
A list of high-level types of data expected from the lander have been identified (Table 1) and a set of ground rules to govern the downlink data scheduling has been constructed as shown below:

- Data with high science value (computed on the basis of the weighted measure as described above), will have high downlink scheduling priority.
- If two types of data are found to have same science value, any of the data can be downlinked first.
- Fault data and information/notification about incomplete crucial tasks needs to transmitted with High priority.
- Data asked by the scientists have downlink scheduling priority of 'Medium to High', depending on the data's science value as compared to other data.
- Downlink scheduling priority of regular health check data as well as heavy data such as videos or high resolution photos can vary from 'Low to Medium', depending on other high priority transmission requirements.

These rules assume that the science value of the lander's observed data and scientist's preference are already known and the data size has not been considered here.

Based on the above rules and the types of data given in Table 1, an example downlink scheduling scheme is constructed that, given the science value of the lander data and the scientist's priority, decides the sequence in which the data should be downlinked for maximum science return, as shown in Table 2.

It is to be noted that, the constructed rules as well as the downlink data scheduling metric will be continually updated as more clarity is attained on the downlink data desired by the scientist and the data size.

3. SSVMs

In most exploration missions, detailed decisions are made on the ground about which tasks were to be executed when, where, and with which instrument, taking into account low-level operational/resource constraints and inter-task dependencies, etc. On the contrary, by effectively utilizing an abstracted symbolic representation of the science activities, we can delegate the handling of operational/resource constraints to on-board autonomy and easily describe spatial/temporal dependencies between tasks. This allows the REASON module to determine actual activity sequences to be performed by considering the information uplinked from the ground and ambient (e.g. the Sun location and surrounding temperature) and internal (e.g. voltage and current of each instrument) information of robotic explorers.

In Figs. 1 - 2 we see that the SSVMs (Shared Science Value Maps) are situated as this interface between the ground and robot. So what are they exactly? It turns out that much of our early work has been in trying to identify appropriate representations that are useful as such an interface. Regardless of the particular format that is ultimately chosen, the SSVMs are defined as symbolic expressions with scientists' science preferences transmitted from the Earth to the lander, an example is depicted in Fig.4.

In essence, SSVMs are common operating pictures for coordination shared by scientists and autonomy. As pictured, SSVMs can be imagined as a map generated and updated for each instrument/science investigation after each activity segment that is easy to interpret and interactive. Such an abstraction of the exploration space for the various instruments - whether that is surface location or pointing direction or time - allows a framework for science activities to be planned, and the results to be returned to scientists on the ground.

In short SSVMs capture scientist's preferences and rules to be communicated to the robot to inform REASON's planning and execution, as well as to inform what acquired science data is most valuable to downlink. Then a report about the robot's executed activities can be captured by updating the "status" of scientist's requested activities. Additionally, modification of

Lander's internal data	Lander's external data
• Instrument status. • Instrument fault notifications. • Current, Voltage, Pressure, Temperature of each instrument.	• Observation data. • Opportunistic data. • Photos/Videos of surroundings. • Information on the level of task completion by the planner.

Table 1: High level types of downlink data expected from lander. Includes data asked by scientists and regular health check data.

Lander's internal data	Scheduling Priority	Overall scheduling priority
Instrument status (**Regular health check**)	2	Needs to be sent regularly. But if no instrument fault is detected, this can be skipped depending on other high priority transmission requirements.
Current, Voltage, Pressure, Temperature of internal equipments. (**Regular health check**)	2	Same as above.
Instrument status (**Asked by scientists**)	3	Depends on scientist's priority as compared to other data.
Current, Voltage, Pressure, Temperature of internal equipments. (**Asked by scientists**)	3	Depends on scientist's priority as compared to other data.
Instrument fault status	5	Instrument faults needs to be reported urgently.
Lander's external data	**Scheduling Priority**	**Overall scheduling priority**
Photos/Videos of data (**Opportunistic science**)	2	Depends on science value of observed data.
Photos/Videos of data (**Asked by scientists**)	3	Send updates to past photos/videos, if the data is already observed. Else send new (low resolution) photos/videos.
Observed interesting data (**Opportunistic science**).	4	Depends on Science Value of the data.
Observation data (**Asked by scientists**)	4	Depends of scientist's priority as compared to other data.
Highly anomalous observed data (**Opportunistic science**).	5	Depends on science values of both asked data and observed data. If they are equal, then send asked data first then immediately send the observed opportunistic data.
Information indicating level of completeness of a task by the planner.	5	Lander needs to downlink details of any incomplete portion of task so that ground can re-work its next uplink sequence accordingly.

Table 2: An example downlink data scheduling metric on a scale of 1-5, 1 meaning lowest downlink priority and 5 meaning highest downlink priority (assuming data science value and scientist's preference are given)

activity segments performed by the REASON module can be confirmed in the RECOURSE module as described in Sec.4 along with the reasoning for the particular executed plan, which improves the reliability and transparency of remote autonomy.

4. RECOURSE

In the presence of low-bandwidth and high-latency communication constraints, there are limited opportunities for end-users (scientists, i.e. exploration domain experts who are not robotics/autonomy experts) and robot autonomy (REASON module) to interact. In addition, the user interfaces (UIs) developed in previous space exploration missions [6], [35] are too complex for mission scientists to operate and demand too much mental workload. Taking these issues

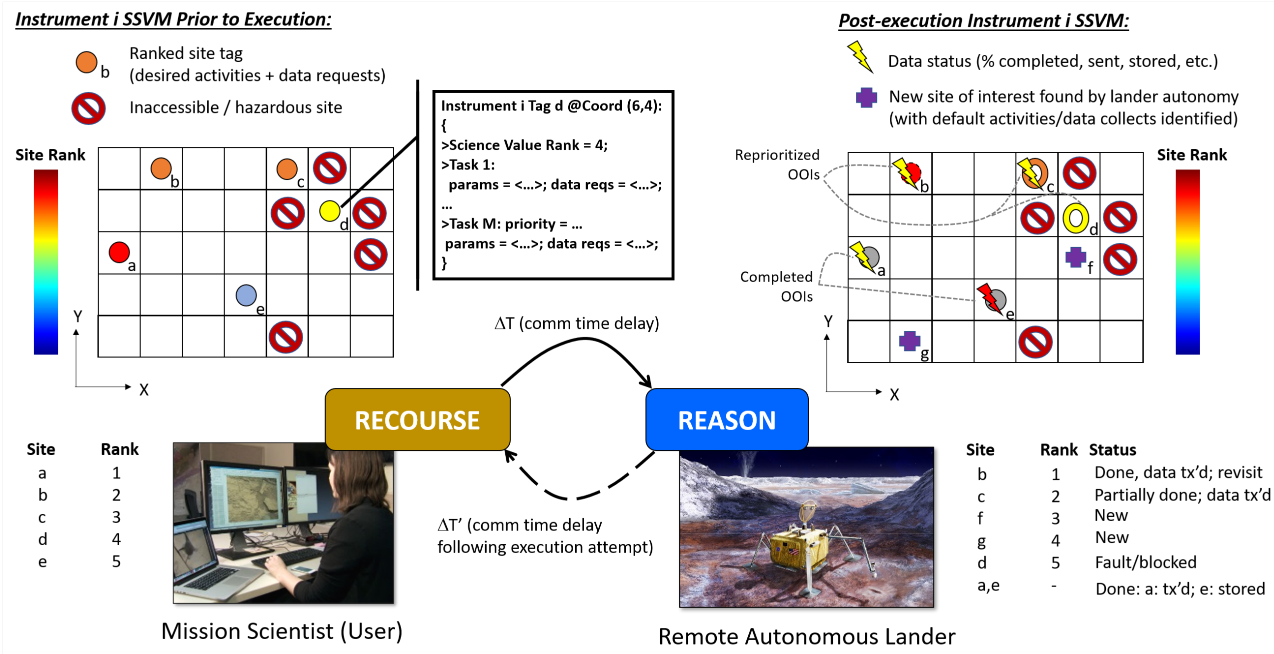

Figure 4: Illustration of SSVM concept definition and updating process.

into consideration, a novel ground-based system tool called RECOURSE (Ranked Evaluation of Contingent Opportunities for Uninterrupted Remote Science Exploration) has been developed. The primary aims of RECOURSE are to design a schedule (i.e. activity segments processed in the REASON module) that is expected to maximize the science return with intuitive operations and to bridge a gap of situational awareness of unknown environments between human and autonomy. RECOURSE is designed with an asynchronous communication framework in mind, as depicted in Fig. 5. Uploading scientist preferences on a staggered timeline with time gaps between when related data is downlinked allows scientists to constantly provide input to the robot without time pressure influencing the execution. In the following, the main two components of RECOURSE, the Uplink UI and the Downlink UI, are described.

Uplink UI

User interfaces (UIs) for uplinking are ground-side tools for transmitting signals from Earth to remote robotic explorers. As mentioned in [35], since these UIs developed in previous space exploration missions are too complicated to manipulate, there is a need to develop user-friendly ones for non-robotic experts, and in order to perform more science-driven operation, the desired UI may need to be able to allow scientists to specify search targets in a simple and intuitive way. And, as explained in Sec.2, it may be desirable for activity segments that are transmitted to robotic explorers to guarantee the completion of tasks. Furthermore, since the survey targets are the surfaces of ocean worlds and icy moons such as Europa or Enceladus, each communication takes a long time and the mission lifetime is short, thus it is not possible to modify the schedule while checking the status of the robot autonomy and surrounding environments sequentially as in the case of near-Earth exploration. Hence, it is desirable to have an (even rough) idea of how the robot's state will transition after executing an activity segment. This will help scientists avoid generating schedules that clearly fail tasks. Based on these motivations, the Uplink UI in RECOURSE has been prototyped as illustrated in Fig.6. This UI allows mission scientists to specify high-level science targets via semantic sketches, and these targets are registered as new atomic propositions (APs) as shown in top-left of Fig.6. Then, by using a formal language (linear temporal logic, *LTL* [36]), it is possible to generate activity segments that can investigate objects of interest while guaranteeing the constraints of missions and instruments (Fig.6, top-right). The REASON module prevents robot autonomy from entering safe mode for a long period of time in case of task failure/unprecedented accidents, however we prefer to avoid entering such mode as much as possible for efficient exploration. Thus, a high-fidelity physical and visual simulator (in this case, the OceanWATERS simulator described in Sec.5), which reproduces realistic environment on Europa, can be directly accessible on the Uplink UI so that mission scientists can immediately check the possible behavior of the robot autonomy (Fig.6, bottom-left) and the simulation log (Fig.6, bottom-right) when it executes LTL-based tasks.

Downlink UI

The Downlink UI is intended to summarize what actions happened and why they were taken by the lander. The UI has three main uses cases to achieve these goals: visualizing received data, explaining differences in the uplinked and executed plan, and reproducing actions taken by the lander. After data from an activity segment is received, it will be separated into two categories. Internal state data such as instrument temperature, current, and voltage will be displayed in a manner intuitive to an engineering audience while external data such as instrument collected information will be presented in a manner conducive to a scientific audience. The executed plan may differ from the uplinked plan due to external and internal factors and therefore it is important to denote and explain these differences. A display directly contrasting the two plans illustrates the differences in the execution order. This visual is useful to a trained logician but does not provide detailed insight to the scientists of the motivation behind why REASON chose to perform one

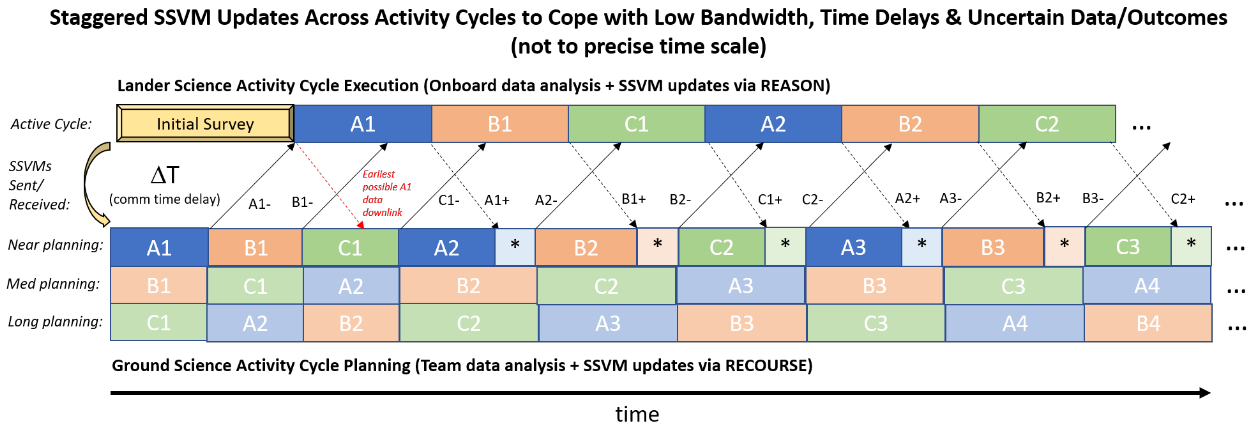

Figure 5: Coordinated SSVM updates between lander platform and ground science teams.

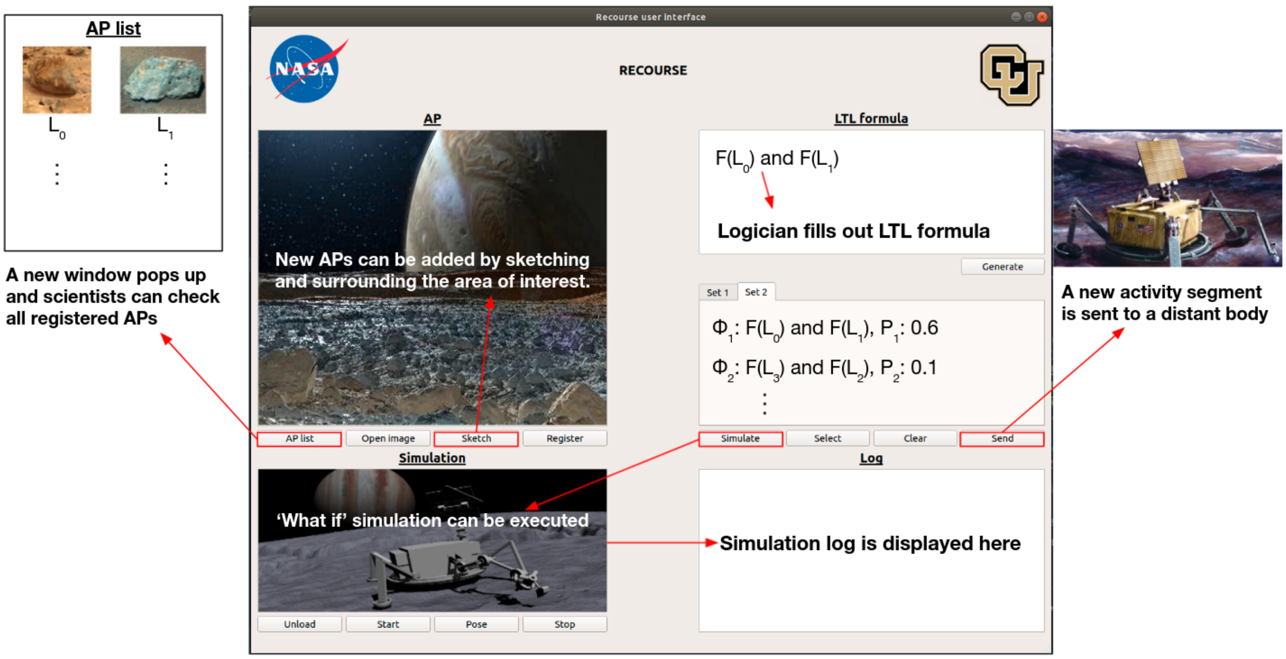

Figure 6: Prototype of the Uplink UI.

task over another. Another section of the UI will focus on explaining the contrasting plans using information provided by REASON. The exact manner of achieving explainability is left for future research building upon previous works exploring formal plan explanation [37] [38]. Additionally task information such as a name, description, and its success or failure status will be shown to the user. An option to simulate the executed plan on the OceanWATERS simulator will allow scientists to view an approximation of the actions the lander took. Downlinked, discritized trajectory data from the lander will be used to create the simulation. The simulation intends to help scientists better understand the underlying autonomy.

5. OceanWATERS Examples

OceanWATERS (Ocean Worlds Autonomy Testbed for Exploration Research & Simulation) is a simulation test-bed for a lander concept on an icy moon [41]. The simulation environment uses ROS Melodic with Gazebo and MoveIt RViz [39] [40] simulation environments. A visualization of the Gazebo environment can be seen in Fig.7.

Instrument Tools

The simulation test-bed comes with many built in tool functionality for simulating different robotic actions. The robotic components of the lander include a robotic manipulator and a revolving camera/communication module. Each of these tools can be controlled by simulating trajectory information. A drill and a digger tool are attached to the end-effector of the robotic manipulator. The test-bed has built in robotic actions that can be used to move the drill and the digger.

Examples

The Synergistic Planning Framework was implemented in the OceanWATERS simulation test-bed, using the robotic arm. The purpose of this case study is to both demonstrate the eft

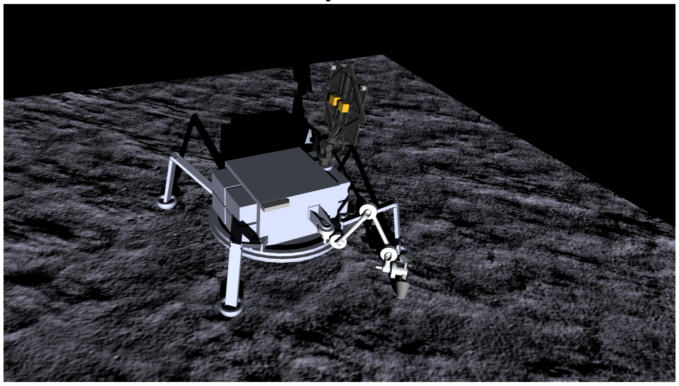

Figure 7: OceanWATERS Gazebo

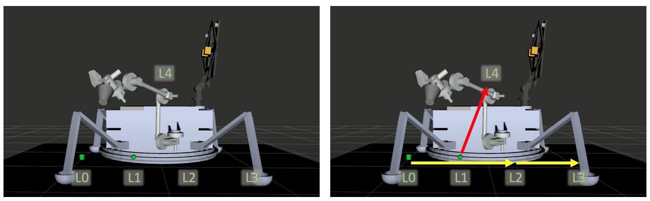

(a) Demonstration Setup (b) Demonstration Task

Figure 8: OceanWATERS setup and task demonstration with 2 objects and 4 location of interest.

ficacy and practicality of the synergistic planning framework, as well as demonstrate the capability of the OceanWATERS simulation test-bed. A visualization of the demonstration setup can be seen in Fig.8a .

In this setup, five discrete locations were defined and labeled $L0$, $L1$, $L2$, $L3$, $L4$. There are two objects that can be moved by the robotic manipulator: a small cylinder and a small sphere. Initially the cylinder is in $L0$ and the sphere is in $L1$. The task specification is given as "*move the cylinder to $L2$ first, then to $L3$, or move the sphere to $L4$*. For the interested readers, the LTLf formula for this task is

$$F\,(c_{L2} \wedge F(c_{L3}))\ \vee\ F(c_{L4}) \tag{1}$$

Here F is an temporal operators that reasons over tasks that need to completed 'Eventually' and $\wedge, \vee$ are binary conjunction and disjunction operators respectively. Symbols c_{L2}, c_{L2}, and c_{L4} are Boolean variables that indicate if object has been placed in that location. Arrows depicting either method of completing the task are shown in Fig.8b. Since the system can complete the task by simply moving the sphere

t

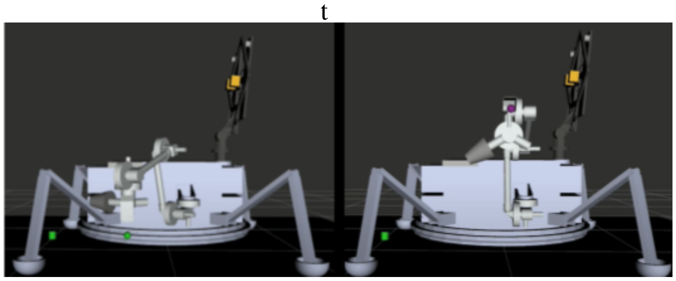

Figure 9: Demonstration Without Obstacle

h

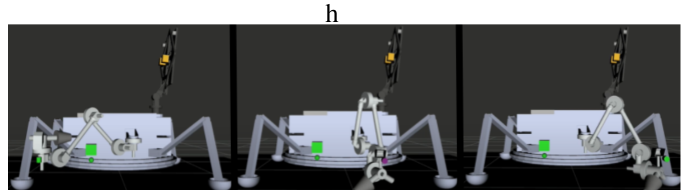

Figure 10: Demonstration With Obstacle

from $L1$ to $L4$, it will initially default to that solution. This execution is shown in Fig.9.

A second scenario is demonstrated where an obstacle is introduced above the sphere object. The discrete model of robotic manipulator does not have any information about the challenge of picking up the sphere. However, as can be seen, the obstacle makes it very challenging, if not impossible to properly pick up the sphere. The Synergistic Planning Framework learns about the challenge of picking up the sphere through iterations, then eventually returns a solution that moves the cylinder to $L2$ and then $L3$. The execution of this task can be seen in Fig.10. This is an equally valid way of completing the task, however it requires more actions, making it less optimal initially.

6. Conclusion

In this paper, we have proposed a new autonomy architecture that is expected to be implemented in deep space exploration mission (e.g. surface investigation of distant bodies such as Europa and Enceladus) in the near future. This architecture consists of three unique components, REASON (on-board robot autonomy executing LTL-based activity segments), RECOURSE (user-friendly interfaces for accelerating science-driven operation), and SSVMs (symbolic expressions with scientists' science preferences for coordination shared by scientists and autonomy), which make it possible to carry out tasks in a way that maximizes science return even in uncertain and dynamic environments where human interaction is very limited due to low-bandwidth, high-latency, and limited mission lifetimes. As parts of the ongoing progress, we present the prototype of the Uplink UI that comprises RECOURSE, and employ Synergistic Planning Framework, which constitutes REASON, on the OceanWATERS testbed, a high-fidelity simulator that mimics the surface environment of Europa, to showcase its robustness with respect to planning with physical constraints of the multiple tools and accommodating for uncertainties in low level motion plans for tools individually or synergistically.

Acknowledgments

The authors thank NASA for support of this project through the COLDTech Program, grant #80NSSC21K1031.

References

[1] Tara A Estlin, Benjamin J Bornstein, Daniel M Gaines, Robert C Anderson, David R Thompson, Michael Burl, Rebecca Castaño, and Michele Judd. Aegis automated science targeting for the MER opportunity rover. ACM Transactions on Intelligent Systems and Technology (TIST), 3(3):50, 2012

[2] R. Francis, T. Estlin, G. Doran, S. Johnstone, D. Gaines, V. Verma, M. Burl, J. Frydenvang,S. Montaño, R. C. Wiens, S. Schaffer, O. Gasnault, L. DeFlores, D. Blaney, and B. Bornstein. Aegis autonomous targeting for chemcam on mars science laboratory: Deployment and results of initial science team use. Science Robotics, 2(7), 2017.

[3] Michael C. Burl, David R. Thompson, Charles deGranville, and Benjamin J. Bornstein., Rockster: Onboard rock segmentation through edge regrouping. Journal of Aerospace Information Systems, 13(8):329–342, 2016.

[4] Li. Jialun, Wu. Zhongchen, Ling. Zongcheng, Cao. Xueqiang, Guo. Kaichen, and Yan Fabao. Autonomous martian rock image classification based on transfer deep learning methods. Earth Science Informatics, 13(3):951–963, Sep 2020.

[5] Hannah R. Kerner, Kiri L. Wagstaff, Brian D. Bue, Danika F. Wellington, Samantha Jacob,Paul Horton, James F. Bell, Chiman Kwan, and Heni Ben Amor. Comparison of novelty detection methods for multi spectral images in rover-based planetary exploration missions. Data Mining and Knowledge Discovery, 34(6):1642–1675, Nov 2020.

[6] S. Chien, G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins, and D. Tran. Aspen - automating space mission operations using automated planning and scheduling. In International Conference on Space Operations(SpaceOps 2000), Toulouse, France, June 2000

[7] S. Knight, G. Rabideau, S. Chien, B. Engelhardt, and R. Sherwood. Casper: space exploration through continuous planning. IEEE Intelligent Systems, 16(5):70–75, 2001.

[8] T. Estlin, D. Gaines, C. Chouinard, R. Castano, B. Bornstein, M. Judd, I. Nesnas, and R. Anderson. Increased mars rover autonomy using ai planning, scheduling and execution. In Proceedings 2007 IEEE International Conference on Robotics and Automation, pages4911–4918, 2007.

[9] Daniel Gaines, Joseph Russino, Ryan Doran, Garyand Mackey, Michael Paton, BrandonRothrock, Steve Schaffer, Ali-Akbar Agha-Mohammadi, Chet Joswig, Heather Justice, Ksenia Kolcio, Jacek Sawoniewicz, Vincent Wong, Kathryn Yu, Gregg Rabideau, Robert Anderson, and Ashwin Vasavada. Self-reliant rover design for increasing mission productivity. In International Conference on Automated Planning and Scheduling (ICAPS), Pasadena, CA, USA, June 2018. Jet Propulsion Laboratory, National Aeronautics and Space Administration.

[10] Akash Arora, P. Michael Furlong, Robert Fitch, Salah Sukkarieh, and Terrence Fong. Multi-modal active perception for information gathering in science missions.Autonomous Robots, 43(7):1827–1853, Oct 2019.

[11] Gregg Rabideau and Ed Benowitz. Prototyping an onboard scheduler for the mars 2020 rover. In 10th International Workshop on Planning and Scheduling for Space (IWPSS 2017), Pasadena,CA, USA, June 2017. Jet Propulsion Laboratory, National Aeronautics and Space Administration.

[12] Jared Strader, Kyohei Otsu, and Ali-akbar Aghamohammadi. Perception-aware autonomous mast motion planning for planetary exploration rovers. Journal of Field Robotics,37(5):812–829, 2020.

[13] Roger C. Wiens, Sylvestre Maurice, Scott H. Robinson et al., The supercam instrument suite on the nasa mars 2020 rover: Body unit and combined system tests.Space Science Reviews, 217(1):4, Dec 2020.

[14] S. Bhaskaran. Autonomous navigation for deep space missions. InSpaceOps 2012, 2012.

[15] Christopher Grasso. VML 3.0 Reactive Sequencing Objects and Matrix Math Operations for Attitude Profiling. InSpaceOps 2012, Reston, Virigina, March 2013. American Institute of Aeronautics and Astronautics.

[16] Eric N Johnson, Anthony J Calise, Michael D Curry, Kenneth D Mease, and J Eric Corban. Adaptive Guidance and Control for Autonomous Hypersonic Vehicles. Journal of Guidance,29(3):725–737, May 2006.

[17] Nikos Mastrodemos, Daniel G Kubitschek, and Stephen P Synnott. Autonomous navigation for the deep impact mission encounter with comet tempel 1. Space Science Reviews,117(1-2):95–121, 2005

[18] J. Riedel, Andrew Vaughan, Robert Werner, Tseng-Chan Wang, Simon Nolet, David Myers,Nickolaos Mastrodemos, Allan Lee, Christopher Grasso, Todd Ely, and David Bayard. Optical Navigation Plan and Strategy for the Lunar Lander Altair; OpNav for Lunar and other Crewedand Robotic Exploration Applications. In AIAA Guidance, Navigation, and Control Conference, Reston, Virigina, 2010. American Institute of Aeronautics and Astronautics

[19] E Riedel, C A Grasso, and W M Owen Jr. A Survey of Technologies Necessary for the NextDecade of Small Body and Planetary Exploration. Technical report, 2009.

[20] Scott D Barthelmy, Louis M Barbier, Jay R Cummings, Ed E Fenimore, Neil Gehrels, DerekHullinger, Hans A Krimm, Craig B Markwardt, David M Palmer, Ann Parsons, Goro Sato, Masaya Suzuki, Tadayuki Takahashi, Makota Tashiro, and Jack Tueller. The Burst Alert Telescope (BAT) on the SWIFT Midex Mission. Space Sci Rev, 120(3-4):143–164, October2005.

[21] A K Jónsson, P H Morris, N Muscettola, and K Rajan. Planning in Interplanetary Space:Theory and Practice. AIPS, 2000

[22] F Teston, R Creasey, J Bermyn, D Bemaerts, and K Mellab. PROBA: ESA's autonomy andtechnology demonstration mission. 1999.

[23] S Chien, R Sherwood, D Tran, R Castano, and B Cichy. Autonomous science on the EO-1mission. 2003

[24] S Chien, R Sherwood, D Tran, and B Cichy. Lessons learned from autonomous sciencecraft experiment.. . . on Autonomous agents . . ., page 11, 2005.

[25] S Chien, D Tran, A Davies, and M Johnston. Lights out autonomous operation of an earth observing sensor web.. . . and Operations (. . ., 2007.

[26] Steve Chien, Gregg Rabideau, Russell Knight, Robert Sherwood, Barbara Engelhardt, DarrenMutz, Tara Estlin, Benjamin Smith, Forest Fisher, T Barrett, et al. Aspen-automating spacemission operations using automated planning and scheduling. In Space Ops 2000, 2000.

[27] Steve Chien, Rob Sherwood, Daniel Tran, Benjamin Cichy, Gregg Rabideau, Rebecca Castano, Ashley Davis, Dan Mandl, Bruce Trout, Seth Shulman, and Darrell Boyer. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. Journal of Aerospace Computing, Information, and Communication, 2(4):196–216, April 2005

[28] HL Enos and DS Lauretta. A rendezvous with asteroid bennu. Nature Astronomy,3(4):363–363, 2019.

[29] Steve Chien, Gregg Rabideau, Daniel Tran, Joshua Doubleday Thompson, Federico Nespoli, Miguel Perez Ayucar, Marc Costa Sitje, Claire Vallat, Bernhard Geiger, Nico Altobelli, Manuel Fernandez, Fran Vallejo, Rafael Andres, and Michael Kueppers. Activity-based scheduling of science campaigns for the rosetta orbiter. Invited Talk, in Proceedings of International Joint Conference on Artificial Intelligence (IJCAI 2015), Buenos Aires, Argentina, July 2015.

[30] M. Costa, M. Perez-Ayucar, M. Almeida, M. Ashman, R. Hoofs, S. Chien, J. Beteta, and M. Kueppers. Rosetta: Rapid science operations for a dynamic comet. In International Conference On Space Operations (Space Ops 2016), Daejeon, Korea, May 2016.

[31] B. A. Bauer and W. M. Reid. Automating the pluto experience: An examination of the new horizons autonomous operations subsystem. In2007 IEEE Aerospace Conference, pages 1–10, 2007

[32] Brian Bauer, Alice F. Bowman, Omar Custodio, Glen Fountain, Sarah A. Hamilton, Helen Hart,Chris Hersman, Adrian Hill, Valerie Mallder, Nickalaus Pinkine, Gabe Rogers, Rebecca Sepan,Karl Whittenburg, and Stephen Williams. Lessons Learned from the New Horizons July 4th Anomaly.

[33] Yanping Guo and Robert W. Farquhar. Baseline design of new horizons mission to pluto and the kuiper belt. Acta Astronautica, 58(10):550–559, 2006

[34] Robert C. Moore. Autonomous safeing and fault protection for the new horizons mission to pluto. Acta Astronautica, 61(1):398–405, 2007. Bringing Space Closer to People, Selected Proceedings of the 57th IAF Congress, Valencia, Spain, 2-6 October, 2006.

[35] Vandi Verma, Dan Gaines, Gregg Rabideau, Steve Schaffer, and Rajeev Joshi. Autonomous Science Restart for the Planned Europa Mission with Lightweight Planning and Execution, The 27th International Conference on Automated Planning and Scheduling (ICAPS), Pittsburgh, Pennsylvania, June 18-23, 2017.

[36] Christel Baier and Joost-Pieter Katoen. Principles of Model Checking, The MIT Press, Cambridge, MA, 2008.

[37] Bastian Seegebarth, Felix Müller, Bernd Schattenberg, and Susanne Biundo. Making Hybrid Plans More Clear to Human Users - A Formal Approach for Generating Sound Explanations, The 22nd International Conference on Automated Planning and Scheduling (ICAPS), Sao Paulo, Brazil, June 25-29, 2012.

[38] Joseph Kim. Bayesian Inference of Linear Temporal Logic Specifications for Contrastive Explanations, Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19), Macao, China, August 10-16, 2019.

[39] N. Koenig and A. Howard. Design and use paradigms for Gazebo, an open-source multi-robot simulator, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566), 2004, pp. 2149-2154 vol.3, doi: 10.1109/IROS.2004.1389727.

[40] David Coleman. "Reducing the Barrier to Entry of Complex Robotic Software: a MoveIt! Case Study," Journal of Software Engineering for Robotics, vol. ArXiv, 2014.

[41] National Aeronautics and Space Administration. Ocean Worlds Autonomy Testbed for Exploration Research & Simulation (OceanWATERS), GitHub Repository, https://github.com/nasa/ow_simulator.

[42] A. Bhatia, M.R. Maly, L. E. Kavraki, and M. Y. Vardi. Motion planning with complex goals. Robotics Automation Magazine, IEEE, 18(3):55–64, Sep. 2011.

[43] K. He, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi, "Towards manipulation planning with temporal logic specifications," in IEEE international conference on robotics and automation. IEEE, 2015,pp. 346–352.

[44] Giuseppe De Giacomo and Moshe Y Vardi. Linear temporal logic and linear dynamic logic on finite traces. In IJCAI'13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence, pages 854–860. Association for Computing Machinery, 2013.

[45] Keling He, Morteza Lahijanian, Lydia E. Kavraki, and Moshe Y. Vardi. Reactive synthesis for finite tasks under resource constraints. In Int. Conf. on Intelligent Robots and Systems (IROS), pages 5326–5332, Vancouver, BC, Canada, Sep. 2017. IEEE.

[46] Keliang He, Morteza Lahijanian, Lydia E. Kavraki, and Moshe Vardi. Automated abstraction of manipulation domains for cost-based reactive synthesis. IEEE Robotics and Automation Letters, 4(2):285–292, Apr. 2019.

[47] I. A.S ucan and L. E. Kavraki, "A sampling-based tree planner for systems with complex dynamics," IEEE Transactions on Robotics, vol. 28, no. 1, pp. 116–131, 2012

[48] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in Int. Conf. on Robotics and Automation, vol. 2. IEEE, 2000, pp. 995–1001

[49] H. Kress-Gazit, G. Fainekos, and G. J. Pappas. Where's waldo? sensor-based temporal logic motion planning. In Int. Conf. on Robotics and Automation, pages 3116–3121, Rome, Italy, 2007. IEEE.

[50] M. Lahijanian, J. Wasniewski, S.B. Andersson, and C. Belta. Motion planning and control from temporal logic specifications with probabilistic satisfaction guarantees,. In Int. Conf. on Robotics and Automation, Anchorage, Alaska, 2010. IEEE.

[51] Morteza Lahijanian, Sean B. Andersson, and Calin Belta. Temporal logic motion planning and control with probabilistic satisfaction guarantees. IEEE Transactions on Robotic, 28(2):396–409, Apr. 2012.

[52] Rebecca Castafio et al. "Rover Traverse Science for Increased Mission Science Return". In:(2003)

[53] Tara Estlin et al. "Enabling Autonomous Science for a Mars Rover". In: American Institute of Aeronautics and Astronautics, May 2008

[54] Srija Chakraborty et al. "Expert Guided Rule Based Prioritization of Scientifically Relevant Images for Downlinking over Limited Bandwidth from Planetary Orbiters". In: Proceedings of the AAAI Conference on Artificial Intelligence33.01 (July 2019), pp. 9440–9445.

[55] D. R. Thompson, T. Smith, and D. Wettergreen. "Information-optimal selective data return for autonomous rover traverse science and survey". In:2008 IEEE International Conference on Robotics and Automation. 2008, pp. 968–973

BIOGRAPHY

Jay McMahon *received his B.S. from the University of Michigan, M.S. from USC and a Ph.D. from the University of Colorado Boulder in 2011. All three degrees in Aerospace Engineering. He is currently an Associate Professor in the Smead Aerospace Engineering Sciences department at the University of Colorado Boulder where his research focuses on spacecraft guidance, navigation and control and asteroid missions and associated science.*

Nisar Ahmed *Nisar Ahmed is an Associate Professor and H.J. Smead Faculty Fellow in the Smead Aerospace Engineering Sciences Department at the University of Colorado Boulder, and holds a courtesy appointment in the Computer Science Department. He is a member of the Research and Engineering Center for Unmanned Vehicles (RECUV) and directs the Cooperative Human-Robot Intelligence (COHRINT) Lab. He conducts research in probabilistic modeling, estimation and control of intelligent autonomous systems for problems involving human-robot/machine interaction, distributed sensor and information fusion, and decision-making under uncertainty.*

Morteza Lahijanian *is an Assistant Professor in the Dept. of Aerospace Engineering Sciences and Dept. of Computer Science (by courtesy) at the University of Colorado Boulder, USA and the director of Assured, Robust, and Interactive Autonomous (ARIA) Systems group. He received his B.S. in Bioengineering from University of California, Berkeley, M.S. in Mechanical Engineering from Boston University, and Ph.D. in Mechanical Engineering from Boston University. He conducted his postdoctoral research in the Dept. of Computer Science at Rice University. Prior to joining CU Boulder, he was a research scientist in the Dept. of Computer Science at University of Oxford, UK.*

Peter Amorese *received his B.S. from the University of Colorado Boulder in 2021 and is currently pursuing a M.S. at the University of Colorado Boulder in Aerospace Engineering. His research interests include formal synthesis, high-level autonomous planning with multiple task preferences, and other formal methods topics.*

Taralicin Deka *is a 3rd year Ph.D. student at the Orbital Research Cluster for Celestial Applications Lab or also called the ORCCA lab at the University of Colorado Boulder. She is originally from Assam, India. She holds a B.Tech degree in Electrical Engineering from the National Institute of Technology Silchar, India, and an MSc. degree in Aerospace Engineering Sciences from the University of Colorado Boulder, USA. Her current research focuses on multi-agent relative-motion planning in perturbed environments.*

Karan Muvvala *is a Ph.D. student in the Aerospace Engineering Sciences department at the University of Colorado Boulder. He received his M.S. in Mechanical Engineering from the University of Colorado Boulder in 2021. His research interest includes Formal Methods, Verification and Planning for Autonomous systems, Game Theory and Control Theory.*

Trevor Slack *received his B.S. from the University of Colorado Boulder in 2021 and is currently pursuing a M.S. at the University of Colorado Boulder. Both degrees in Aerospace Engineering. His research interests includes autonomous risk identification and classification, simulation design, and human-robot interaction.*

Shohei Wakayama *is a Ph.D. student in Aerospace Engineering Sciences department at the University of Colorado Boulder. He received his B.S. in Mechanical Engineering from Kyushu University in 2018. His research interest includes autonomous search and exploration, decision making under uncertainty, and human-robot interaction.*