A Flight-Ready Software Platform for Deploying Science Autonomy Onboard Spacecraft

MARK WRONKIEWICZ
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

GARY DORAN
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

JACK LIGHTHOLDER
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

ZAKI HASNAIN
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

REX STOCKHAM
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

ROBERT J. BROWN
TMC Technologies, NASA Independent Verification and Validation (IV&V), Jon McBride Software Test and Research (JSTAR) Team, Fairmont, WV 26554, USA

SCOTT A. ZEMERICK, IEEE
TMC Technologies, NASA Independent Verification and Validation (IV&V), Jon McBride Software Test and Research (JSTAR) Team, Fairmont, WV 26554, USA

LUKAS MANDRAKE
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011, USA

***Abstract*— Spacecraft exploring Earth and other bodies in the solar system are producing increasingly large and complex volumes of scientific data. Recently, the field of science autonomy has begun to develop data analysis algorithms to select, prioritize, and characterize science instrument data directly onboard the spacecraft. These data-derived insights can then inform downstream autonomous decisions like data prioritization to mitigate data bandwidth constraints. However, science autonomy development has largely focused on bespoke, instrument-specific algorithms; there are no generalized platforms for sharing reusable utilities or centralizing the operation of multiple science autonomy algorithms. To accelerate future development of flight-ready science autonomy, we developed a software package called Science Yield Improvement via Onboard Prioritization and Summary of Information System (SYNOPSIS). SYNOPSIS manages the execution of one or more science autonomy algorithms, stores their outputs in a centralized database, and prioritizes data even across multiple instruments. We simulated the use of SYNOPSIS with science autonomy algorithms to analyze and prioritize data collected by the Curiosity rover on Mars and integrated it into two popular flight software packages — core Flight System (cFS) and F-Prime — to ease infusion of future science autonomy algorithms. As opportunities for onboard data analysis grow with the volume and complexity of data, we expect that platform technologies like SYNOPSIS will facilitate future exploration of our solar system.**

***Index Terms*— Space technology, Data compression, Knowledge representation, Software tools, Space vehicle communication**

Manuscript received XXXXX 00, 0000; revised XXXXX 00, 0000; accepted XXXXX 00, 0000.
The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology through the Advanced Multi-Mission Operations System (AMMOS) program under a contract with the National Aeronautics and Space Administration (80NM0018D0004).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

I. INTRODUCTION

A. Science Autonomy for Upcoming Missions

Missions to explore other bodies in our solar system face challenging communication scenarios because of the physics-based constraints that come with operating over interplanetary distances. These constraints limit data bandwidth – the rate at which information can be transmitted back to Earth – as well as spacecraft response time – the time delay involved when downlinking data to ground teams, generating a command dictating a response, and uplinking that command to the spacecraft. Therefore, autonomous decision making is a valuable capability that can allow the spacecraft to take immediate actions based on all data available and in real time. In this vein, Onboard Science Instrument Autonomy (OSIA) encompasses the field of algorithms designed to analyze science instrument data *onboard* to inform autonomous decisions. Instead of commanding spacecraft via a series of imperative commands, OSIA permits ground teams to pre-specify how spacecraft should use science data to drive goal-based behavior. Spacecraft can then generate actions based on their in-situ representation of the science environment – a sometimes privileged perspective compared to operators and scientists on Earth due to the aforementioned bandwidth restrictions and time delays. Broadly, OSIA aims to codify science intent into onboard software as a means to mitigate communication constraints and improve science return.

Over the past two decades, science autonomy algorithms were developed for a handful of science instruments and even deployed on a couple missions. In use on both the Curiosity and Perseverance Mars rovers is Autonomous Exploration for Gathering Increased Science (AEGIS) – a platform for autonomous rock selection capable of generating targets without the ground team's input [1]. Other OSIA algorithms have been developed for the Ocean Worlds Life Surveyor (OWLS) instrument suite to identify likely biosignatures in microscopy and mass spectrometry data [2], [3], Mars Organic Molecule Analyzer (MOMA) on the ExoMars rover Rosalind Franklin to find biosignatures also in mass spectrometry data [4], [5], and Phoenix lander to summarize meteorological data [6]. Finally, algorithms designed to identify ephemeral phenomena like the cloud and dust devil detectors used on the MER rovers [7], or ongoing efforts to detect dust vortices in pressure sensor data [8], offer a path to adapting spacecraft behavior based on science data. While these technologies have already improved science return or are in development to do so, their relatively novel approach means that most science autonomy is decidedly not "heritage." Thus, significant research, development, and testing are required for infusion onto new missions.

B. Challenges to Infusion

The infusion of new technology into missions depends on many factors, but broadly it is judged by mission teams working to satisfy core mission requirements while minimizing costs and risks due to technology. To assess a technology's maturity and risk, many rely on NASA's Technology Readiness Level (TRL) scale [9]. For NASA's major mission classes (e.g., Flagship, Discovery, New Frontiers), TRL 5 is required for any new technologies included in a mission proposal and that maturity must be increased to TRL 6 by pre-phase A. This creates a barrier to infusion as the cost required to reach these maturity levels is beyond the scope of most research funds; TRL 5 requires end-to-end implementation with evaluation in a simulation environment and TRL 6 necessitates integration with surrounding software/hardware systems as well as testing in full-scale, realistic environment. Research grants typically focus on proving feasibility through TRL 3. Even if a new technology reaches TRL 5, mission teams may still hesitate to include them in their proposal if formulation leads do not have ample time to evaluate how the technology could affect their mission concept's ability to stay on budget and schedule. These barriers combine to form a valley of death where **many TRL 3-4 technologies (including science autonomy) tend to stall.**

C. Facilitating Mission Infusion

To bridge this TRL gap for science autonomy, we created Science Yield Improvement via Onboard Prioritization and Summary of Information System (SYNOPSIS). SYNOPSIS is an open-source[1] software package that manages the onboard execution of one or more OSIA algorithms and the data products they produce. Science autonomy development is naturally a bespoke process as the software must be built to couple with data from a specific instrument for a precise science objective. This poses a unique challenge when building OSIA for competed NASA science missions, which are often novel by definition in order to target new scientific discoveries. Because the science questions and instruments change for each new opportunity, OSIA developers typically need to restart at TRL 1 for each target mission and often cannot begin in earnest until the target science instrument has been matured to the point that sample data is available. However, while the core algorithms of science autonomy often must be tailored to match well-defined use case, some of the higher level software components need not be replicated. For example, OSIA has generic needs including integration with flight software, execution of data processing algorithms onboard, and management of any generated data products or metadata. It is the handling of these shared needs where SYNOPSIS provides value. It coordinates the execution of OSIA algorithms and their data products in a flight-integrated software environment thus permitting science autonomy experts to focus on the bespoke, instrument-specific aspects of their algorithms. SYNOPSIS encapsulates these capabilities into a flight-ready software library to **facilitate the advancement of science autonomy algorithms from TRL 3 to TRL 5** for inclusion in competed mission proposals.

SYNOPSIS comprises three specific capabilities: First, it directs the processing of science data with OSIA, which can be configured to trigger according to ConOps needs (e.g., immediately after data collection or later once adequate compute is available). SYNOPSIS does so using a common software interface meaning that if a mission concept includes multiple OSIA algorithms, they can be developed with a unified output format. Second, SYNOPSIS stores the resulting data products and metadata in an onboard database. This database retains the spacecraft's representation of the scientific environment to inform arbitrary downstream autonomous actions (e.g., data prioritization), which is becoming especially important as commercial space missions and upcoming Artemis missions will lead to further rationing of DSN communications. Third, SYNOPSIS can leverage the onboard database to generate a rank-ordered data prioritization list such that the most scientifically valuable data products are transmitted to Earth first. While not the focus of this work, the onboard database could also support other decisions like science-aware navigation and sample site selection, adaptive data compression based on the bandwidth available for an upcoming downlink pass, or optimizing planning and scheduling for a given set of scientific activities. This could involve integration with systems like Multi-mission EXECutive (MEXEC), which includes dynamic planning and scheduling capabilities [10], to fold science-based information into existing systems that make autonomous decisions.

To further facilitate infusion, the open-source SYNOPSIS library is written in C and integrated into core Flight System (cFS) and F-Prime – two popular flight software packages. For new algorithms to be SYNOPSIS compatible, we require that they generate quantitative metadata about the relative value and characteristics of the instrument data, as will be described in Sections III A and III B. As the OSIA field grows, we expect that

[1] https://github.com/nasa-ammos/synopsis

some science autonomy developers will release open-source versions of their SYNOPSIS-compatible software. Future work will focus on this community development as a way to promote reuse of science autonomy algorithms, especially for common instruments (e.g., cameras or mass spectrometers).

The remainder of this paper will describe the technical implementation of SYNOPSIS. We will motivate and demonstrate its capabilities using science data from MSL. We show that SYNOPSIS:

- Facilitates the deployment and execution of science autonomy algorithms
- Quantifiably improves science return through the use of metrics capturing science utility and sample diversity
- Enables mission architects to specify rule-based relationships across instruments to leverage multi-instrument data prioritization schemes

II. SYSTEM ARCHITECTURE

This section provides an intuitive overview of SYNOPSIS and its intended role in future missions. It describes the flow of science data through the onboard system, basic ConOps considerations, and the motivation for including data prioritization as a core capability.

A. Overview and CONOPs

SYNOPSIS manages the onboard processing of instrument data for informing downstream autonomous decisions (Figure 1). Mission science begins on the spacecraft with instruments collecting observational data. SYNOPSIS then applies the appropriate OSIA algorithm to the raw instrument observations. This may include processing that, for example, identifies high-value snippets in time series data, extracts scientifically interesting portions of an image, quantifies background signal in mass spectrometry data, etc. The OSIA associated with each instrument could be simple or sophisticated depending on the science needs of the mission. This generates metadata (described in Section III A) characterizing the data and its scientific value. Metadata is stored in an onboard database to provide a centralized onboard representation of the scientific environment. To enable data prioritization, the metadata should include a quantitative assessment of each observation's scientific utility as well as features indicating diversity relative to other observations (as described in Section III A). Collectively, this information is useful to inform SYNOPSIS or other algorithms when making downstream autonomous actions. While this work focuses on the act of prioritizing data directly before a communications pass to transmit products in order of estimated scientific importance, SYNOPSIS's data record could inform other tasks (as discussed in Section I C) related to collecting new samples or optimizing the planning and scheduling of autonomous tasks.

The SYNOPSIS framework is designed to support a wide variety of ConOps. SYNOPSIS processing commands can either be integrated directly into instrument data acquisition commands (calling relevant OSIA processing directly following data acquisition), triggered asynchronously after specified conditions are met, or commanded independently by ground teams. The latter approaches support a decoupling of data acquisition and OSIA data processing, which is useful during critical mission phases, such as landings or target flybys, when onboard compute resources must be fully dedicated to data acquisition or vehicle health. In such cases, bulk OSIA processing can be commanded to run after critical events have concluded. Thus, the SYNOPSIS library provides flexibility for how mission teams choose to integrate SYNOPSIS into their ConOps and instrument commanding.

During each ground in the loop cycle, operators can also exercise the option to reconfigure or override SYNOPSIS. We expect that reconfiguration will be an important feature of future science autonomy packages since science teams will need to change the behavior of the onboard system as their understanding of the science environment grows throughout a mission. SYNOPSIS also includes commands to manually override the autonomous prioritization (e.g., to move data products to the front or end of the downlink queue). This is especially crucial as a safety valve for nascent science instruments with limited or no flight heritage.

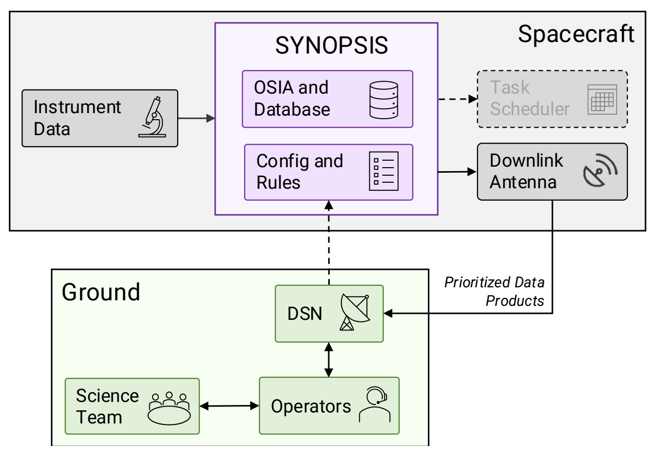

Fig. 1. **Overview of SYNOPSIS and its relationship to the spacecraft and ground teams.** As data is collected by the spacecraft, SYNOPSIS directs its processing by science autonomy algorithms, and stores the results in preparation for downstream autonomous actions. While we focus on using this information for downlink prioritization, it could enable other autonomous actions as well. Ground teams can update parameters (via configuration files) to tune the spacecraft's desired behavior.

In the case of data prioritization, SYNOPSIS rank orders raw and/or processed data products such that the highest priority data products are transmitted first. Such behavior is valuable for spacecraft capable of collecting much more data than can be transmitted – an increasingly common situation for modern spacecraft. It also reduces risks for spacecraft with timeliness constraints where the availability of high-priority science data (or lack thereof) can inform immediate decisions on what actions to take next. Finally, data prioritization improves robustness to Deep Space Network (DSN) rationing as the number of other high-priority missions increase (e.g., Commercial Lunar Payload Services (CLPS), Commercial Mars Payload Services (CMPS), and Artemis). Using OSIA-produced metadata about the original data inputs, SYNOPSIS applies a content- and rules-based algorithms to generate a prioritized queue of data to downlink. As discussed in Section III A, prioritization is configurable to permit science utility, sample diversity, or some mixture of the two to drive the ranking.

III. SYSTEM IMPLEMENTATION

This section provides the full technical description of SYNOPSIS including the ingredients needed to quantify science value, the application of rules and constraints in multi-instrument prioritization settings, and the prioritization algorithm. We also detail its integration with existing Flight Software (FSW) frameworks to facilitate infusion.

A. Prioritization Precursor: Quantifying Scientific Value

SYNOPSIS applies OSIA software to process observed data onboard and assess its science content. The two core components of this assessment include: **(1) a Science Utility Estimate (SUE) float scalar**, which provides a quantitative estimate of an observation's scientific value to the science team, and **(2) a multidimensional Diversity Descriptor (DD) vector**, which describes an observation's content along multiple dimensions each quantifying an axis of the data useful for comparing observations against each other. Together, these two quantities allow data to be assessed onboard in terms of their relevance to the science mission and novelty relative to previously observations. The SUE and DD, along with pointers to the processed data and any additional metadata, are stored in SYNOPSIS's onboard database.

Both the SUE and DD *must be developed in collaboration with scientific experts* to ensure the onboard assessment extracts meaningful information from each observation. The development of these quantities is OSIA-specific and therefore out of scope of this work, but see [3] for an example of an algorithm using this approach to prioritize cellular biosignatures for astrobiology; there, particles were tracked in water sample flowing through a microscopic chamber. OSIA algorithms assessed the particle movement and composition looking for signs of cellular motility and biological structure to form the basis of SUE calculations. Other extracted properties like particle size and average speed were two of the nine elements in the DD vector. Again, the SUE and DD should be defined with scientists' input and SYNOPSIS is flexible to any definition that follows the basic format defined above. These quantities formed the basis for data prioritization using a procedure described in Section II C.

B. Prioritization Precursor: Cross-Instrument Rules and Constraints

SYNOPSIS also permits prioritization assessments across *multiple* instruments as well as constraints to set hard bounds on certain prioritization behaviors. Cross-instrument rules exist to autonomously identify situations where the whole is bigger than the sum of the parts — when multiple instruments capture different perspectives of the same high-value scientific phenomena, those data products should be collectively prioritized. These rules permit science teams to encode arbitrary logic that can adjust SUE values. Constraints allow science teams to specify hard limits on the prioritization behavior, such as guardrails to cap the transmitted data volume or number of data products from a particular instrument that can be transmitted. Together, the rules and constraints enable the specification of inter-instrument relationships during prioritization, which is valuable for some multi-instrument spacecraft. For example, a pair of cameras that collect both high-resolution detail and low-resolution contextual information for the same scene, such as on Lunar Reconnaissance Orbiter (LRO) and Mars Reconnaissance Orbiter (MRO), could leverage such a capability to upweight an image pair if OSIA identifies a valuable feature in the high-resolution image. As another example, an astrobiology mission carrying instruments to measure both molecular- and biological-scale biosignatures (e.g., chemical metabolites through mass spectra and cellular structures through fluorescent labeling) could drastically increase the prioritization of data for any samples where both instruments identified positive evidence for life. In Section IV C, we explore this capability in ChemCam data where two instruments provide complementary information for the same geologic sample.

Rules and constraints are implemented using a custom, generic "SYNOPSIS Rule Declaration (SRD)" language. Each rule defined in the language has up to four clauses (see Figure 3):

1) a **rule variable declaration clause**, which declares either one or two variables to which data products are bound during evaluation,
2) a **rule application clause**, which consists of a logical expression to determine whether the rule should be applied based on metadata properties, similar to a `WHERE` clause in a SQL statement,
3) a **utility adjustment clause**, which specifies how to adjust the overall scientific utility of the queued data products if the rule is applicable, and
4) an optional **maximum applications clause**, which specifies the maximum number of times the rule should be applied to a given set of queued products.

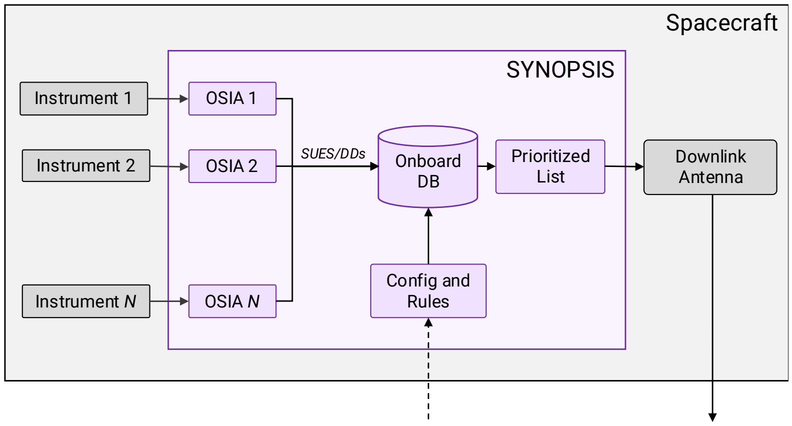

Fig. 2. **Detailed view of SYNOPSIS illustrating how raw science data is processed and prioritized for downlink.** SYNOPSIS manages how and when instrument data is processed by executing OSIA algorithms and then stores results (SUEs/DDs) in an onboard database. When a data prioritization command is received, SYNOPSIS generates a list of data products prioritized by the desired configuration in order to maximize science return.

```
RULE([variable declaration]):
APPLIES [application expression]
ADJUST UTILITY [adjustment expression]
MAXIMUM APPLICATIONS [integer];
```

Fig. 3. Rule definition components. Each element in square brackets is replaced with a code block. The `MAXIMUM APPLICATIONS` component is optional.

```
CONSTRAINT([variable declaration]):
APPLIES [application expression]
SUM [aggregation expression]
LESS THAN [comparison expression];
```

Fig. 4. Constraint definition components. Each element in square brackets is replaced with a code block as described below. The `SUM` component can also be replaced with `COUNT` in which case there is no aggregation expression.

For a given candidate prioritization, rules are applied in order of specification. For each rule with a single variable declaration, SYNOPSIS attempts to apply it with each data product in the queue bound to the variable, in prioritized order. Then, for each rule with two variables declared, it is applied over all pairs of data products in the queue (again in prioritized order, cycling through all bindings of the second variable before moving to the next product being bound to the first variable). Rules with two variables enables application logic considering pairs of products, such as two observations of the same target by different instruments. Higher-order rules involving relationships between more than two data products is not currently allowed to limit computational complexity.

Similarly, constraint definitions have the following clauses (see Figure 4):

1) a **variable declaration clause**, which declares a single variable to which data products are bound during evaluation,
2) a **constraint application clause**, which as with rule definitions determines whether the constraint applies to the bound data product,
3) a **sum or count clause**, which accumulates some value used by the constraint, and
4) a **less than clause**, which specifies the upper bound on the accumulated quantity.

C. Prioritization Algorithm

Once SYNOPSIS's database is populated with observations and derived metadata, downstream algorithms can query it to guide autonomous decisions. Prioritization would typically be triggered as part of a command sequence preceding a downlink pass. The prioritization used here is driven by the two factors previously described: a configurable trade between science utility and diversity as well as any cross-instrument rules and/or constraints dictating how per-instrument ranking should be combined to achieve the final prioritization order. This choice supports scenarios where science data collection far outpaces downlink bandwidth. While we explore this scenario to quantitatively evaluate a specific instantiation of SYNOPSIS, other prioritization frameworks may be more appropriate depending on the specific science objectives, instrument configuration, mission phase, data bandwidth, etc.

To leverage both the SUEs and DDs during prioritization, we adapt the **Maximum Marginal Relevance (MMR) algorithm** [11]. Intuitively, MMR picks data with high science value (i.e., large SUEs) while discounting or "downweighting" observations that are highly similar (i.e., with nearby vectors in the DD hyperspace). It does so via a greedy approach that iteratively selects observations which maximize the total science utility after applying this discount factor based on the most similar observation previously downlinked or already selected for downlink. Given a similarity metric sim (e.g., Gaussian similarity between DD vectors), the discount factor for the i^{th} observation (df_i) is computed as follows:

$$\text{df}_i = (1 - \alpha) + \alpha \left(1 - \max_j \text{sim}\left(\text{DD}_i, \text{DD}_j\right)\right) \quad (1)$$

This discount factor is controlled by setting MMR's α parameter. When α is close to 0, the left hand side of Eq. 1 dominates resulting in a discount factor for the i^{th} observation close to 1 and leaving the SUE mostly unchanged. When α is close to 1, the right hand side dominates and any existing observation j with a similar DD will lead to a discount factor close to 0 (downweighting the i^{th} observation's SUE). Favoring the SUE indicates that SYNOPSIS should prioritize data that resembles known science targets of interest while favoring the DD results in data products that are novel compared to previous observations. While MMR uses a greedy approach to iteratively select the next best data product rather than computing a globally optimal solution, this is a computationally light approach amenable to onboard compute constraints. Overall, this provides science teams a mechanism to tune the prioritization to behavior over the course of a mission.

After generating a preliminary ranked list of candidate data products, SYNOPSIS applies any rules and constraints to exercise within- or cross-instrument relationships (as previously described in Section III B). These act to adjust the final ranking order according to a ground-team-specified configuration. Once completed, SYNOPSIS saves a text file containing the ranked data products, their file path on disk, and relevant metadata including their SUE and DD. Finally, the downlink ordering is provided to the onboard downlink system for transmission to Earth. At any point, operators on the ground can interact with the prioritization system by updating OSIA parameters, changing SYNOPSIS's configuration, or manually specifying data priorities as in traditional prioritization schemes (e.g., to override the autonomous prioritization order).

D. Flight-compatible Software Architecture

SYNOPSIS was designed to integrate with mission flight software software environments thereby reducing the effort needed to deploy new science autonomy algorithms. It is implemented in C to ensure maximum efficiency and retain portability to most flight software environments. In addition, it has a built-in message management designed to report messages at multiple severity levels, including INFO, WARN, and ERROR, emulating those of traditional flight software event reporting. Thus, messages generated by SYNOPSIS can be easily captured by higher level flight software and passed into the event reporting framework in use by the broader spacecraft's flight software. Further, SYNOPSIS uses best practices of onboard memory management; the library itself performs no memory allocation, instead expecting prefixed allocated memory blocks to be provided at library initialization. Abstracting memory allocation and cleanup away from SYNOPSIS allows for cleaner integration with flight software, centralizing SYNOPSIS specific memory management to functions initializing and de-initializing the library. SYNOPSIS was also designed to follow the Power of 10 (P10) rules for developing safety critical code, which prescribes the avoidance of certain C coding practices that make software hard to analyze and test [12]. The library also includes unit tests and code coverage reporting to ensure robustness and maintainability. Additionally, dependencies such as the SQLite3 database were selected for their lightweight footprint, reliability, and ease of integration with existing spacecraft systems and constraints. These design decisions were made to balance the novel aspects of onboard science autonomy capabilities while still respecting the constraints associated with integrating new software into flight software environments.

E. Flight Software Integration

Flight Software (FSW) is the software that executes on spacecraft, satellites, and other vehicles to manage and control their operation. It performs core functions such as: (1) command and control of the vehicle, (2) navigation, guidance, and attitude control, (3) system health monitoring and fault management, (4) communication with ground systems, and (5) data acquisition and storage from onboard sensors and instruments. Since SYNOPSIS is meant to serve as a flight-compatible platform for OSIA, we developed it with an eye toward FSW integration. Besides its implementation in C, we included unit tests and coverage reports with Linux Coverage Visualizer (LCOV), a built-in logger to track system messages at various severity levels (e.g., INFO, WARN, ERROR) for ground teams, and functions to generate EVRs from logs. SYNOPSIS's prioritization scheme also supports simple prioritization bins that are drained in order and the ability to manually override the autonomous prioritization system. This makes SYNOPSIS backwards compatible with the standard data prioritization scheme in use by current Mars rovers.

To reduce the effort required for OSIA infusion, **we integrated SYNOPSIS with two popular FSW packages.** The core Flight System (cFS) is an open-source platform- and project-independent, reusable flight software framework [13] that has flight heritage on numerous large and small NASA missions including the Roman Space Telescope, the Global Precipitation Measurement (GPM), and the Lunar Atmosphere Dust and Environment Explorer (LADEE) [14]. Additionally, it was also selected for use on the lunar Gateway [15]. The cFS application software layer includes a set of reusable modules to facilitate development. These applications are meant to be tailored to the mission requirements using tables, but new applications can also be developed for any mission specific requirement that is not directly provided by cFS.

F-Prime (or F') is another open-source flight software framework, developed and maintained by NASA's Jet Propulsion Laboratory. It is a relatively new platform and was designed for small-scale space missions such as SmallSats and small landed vehicles [16], [17]. Its key features include: (1) a component-based architecture, (2) emphasis on small spacecraft including SmallSats and rovers and (3) built-in testing tools that help in ensure the reliability of the software. One significant mission that utilized F-Prime is the Mars Helicopter Ingenuity, which made history as the first rotorcraft to fly in the Martian atmosphere [18]. Both cFS and F-Prime are well-known packages in the FSW ecosystem, serving as platforms upon which mission-specific functionalities can be developed.

We integrated SYNOPSIS into both FSW packages and simulated its use in a manner consistent with eventual mission deployment. Our primary goal was to demonstrate SYNOPSIS in a realistic spacecraft ConOps environment with human-in-the-loop commanding and telemetry. We sent SYNOPSIS commands from the "ground station" to the "spacecraft" FSW, thus demonstrating the normal uplink/downlink processes for typical spacecraft. For example, we tested commands including *Perform Prioritization, Acquire Data,* and *Re-prioritize with Different Parameters*. These commands would be issued via ground teams as part of the standard operation of SYNOPSIS. After cFS integration, we also demonstrated the SYNOPSIS within the open-source NASA Operational Simulator for Small Satellites (NOS3) platform. NOS3 is meant to SmallSat software framework that simulates many aspects of a mission making it useful for development, testing, and personnel training [19]. Within NOS3, the flight software executes as if operating in space providing the flight software with representative real-world simulated data inputs expected during actual operations. A link to the video demonstrating the NOS3 integration is available online[2] (and also via the SYNOPSIS GitHub repository).

While certain components of the SYNOPSIS codebase and FSW integration will require tailoring to any given mission opportunity, we expect that the library will substantially reduce the effort needed to infuse new science autonomy. This will permit science autonomy developers to focus more time on building and evaluating bespoke OSIA algorithms for their target use case knowing that external software exists to expedite the FSW integration. Other components such as the downlink planner and configuration of data product relationships and constraints are designed to be adaptable to specific mission requirements, but a small number of implementation options can likely serve the needs of many missions. Overall, we expect this open-source tool will lower the barrier for science autonomy's inclusion on more mission concept proposals and future mission deployments.

[2]https://www.youtube.com/watch?v=41cvM5iSn2Y

IV. DEMONSTRATION AND RESULTS

To demonstrate the built-in prioritization capabilities of SYNOPSIS, we simulated a mission scenario involving real data from two instruments on Mars Science Laboratory (MSL).

A. Motivating Mission Scenario

We demonstrate SYNOPSIS using data from the Chemistry and Camera (ChemCam) instrument suite on the MSL rover (Curiosity) [20], [21], which includes two instruments used to sample the composition and morphology of geologic targets: the Laser-Induced Breakdown Spectrometer (LIBS) and Remote Micro-Imager (RMI) as shown in Figure 5. LIBS allows chemical analyses of rock and soil samples by focusing intense laser pulses on geologic targets, creating plasma. The emitted light from this plasma is collected with a fiber optic cable and routed to three spectrometers. By analyzing the resulting spectra, the elemental composition of the observed sample can be assessed by comparing measurements to known spectral signatures. For simulation purposes, we assume these compositions are available onboard via comparison with a spectral library, even though they are typically computed on the ground. **Of particular interest is evidence of manganese (Mn) enrichment,** which can indicate past interactions with liquid water and conditions favorable for habitability.

Designed to complement the LIBS, the RMI is a camera that captures contextual images about the specific locations sampled by LIBS. It uses uses the same telescope as the LIBS laser to image geologic targets with its 1024 x 1024 pixel CCD detector. The RMI permits analysis of submillimeter LIBS imprints on geologic materials, thereby allowing for precise localization of the LIBS sample spots.

Upon receiving commands from ground teams, MSL initiates its drive sequence, navigating over the Martian terrain until it reaches a predetermined or algorithmically-selected location. There, the ChemCam suite can be targeted autonomously via Autonomous Exploration for Gathering Increased Science (AEGIS) or manually after a ground-in-the-loop cycle. AEGIS is typically allotted resources to collect LIBS measurements for 1-2 geologic targets post-drive and these targets are selected autonomously based on ground-specified parameters (e.g., rock brightness, size, etc.; [1]). For simplicity, we analyze the performance of SYNOPSIS independently of whether targets are selected manually or autonomously, although in principle, metadata about the targeting method could be used by rules to guide data prioritization. Typically, an RMI "before" image is acquired, then LIBS samples are taken, and a follow on RMI "after" image ends the data collection sequence. Following this measurement sequence, the onboard computer packages collected data and operational telemetry into a downlink bundle. Leveraging Mars orbiters as relay stations, the rover will

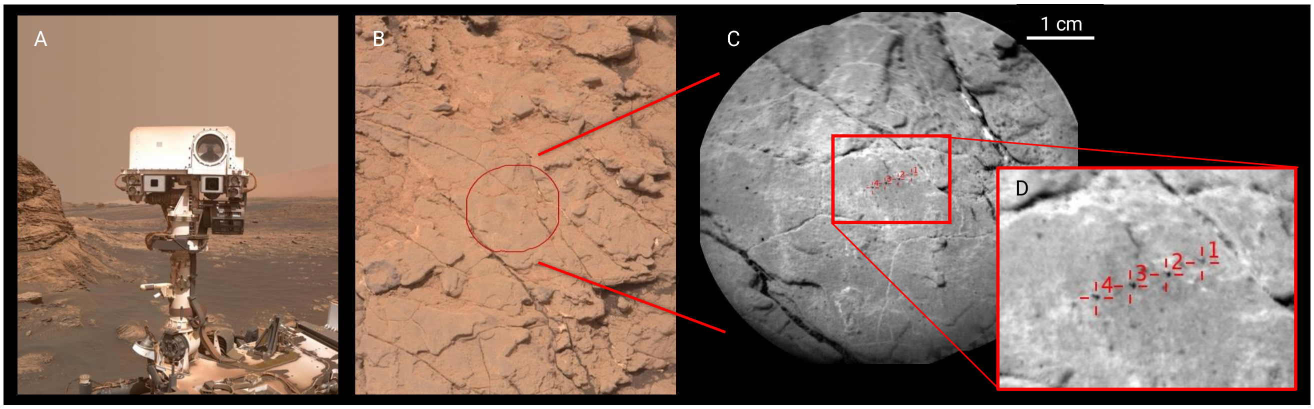

Fig. 5. **The LIBS and RMI instruments within ChemCam help conduct geologic science at the surface of Mars. (A)**: Selfie image taken near the Mont Mercou outcrop showing the ChemCam instrument's large circular aperture at the top of MSL's mast. **(B)**: MastCam M100 image showing the rock sampled with the red contour outlining the RMI mosaic field of view. **(C)** RMI mosaic with LIBS sample points indicated with red crosshairs. **(D)** Same as (C) but expanded to visualize the laser marks where geochemical spectra are sampled. Image credits: NASA/JPL-Caltech/MSSS.

transmit the collected science data back to the Earth-based mission teams for spectral and image analysis as well as further mission planning.

Under normal conditions, ChemCam ConOps proceeds at a regular cadence; planning cycles generally span one to three sols, and the data bandwidth typically permits the transmission of all ChemCam data to Earth during a downlink cycle on that same sol or within a few sols. For demonstration, we hypothesize a synthetic ConOps cycle where downlink bandwidth is substantially degraded. This could arise if there's a rationing of DSN (e.g., due to an ongoing crewed Artemis mission) or one of the aging relay orbiters malfunctions. In this hypothetical situation, there may be little bandwidth allotted for science data as ground teams will be prioritize engineering telemetry to support navigation and safety-critical rover commanding. Through this synthetic context, we explore SYNOPSIS's capability to use (1) data prioritization and (2) explicit rules and constraints to ensure high scientific yield, even when confronted with degraded communication scenarios. We use real data from MSL collected from 11 sample sites from sols 2828 thru 2837 available on the Planetary Data System (PDS) (https://pds-geosciences.wustl.edu/msl/msl-m-chemcam-libs-4_5-rdr-v1/mslccm_1xxx/data/, accessed June 5, 2024). These measurements comprise a good demonstration set as a few of the sample locations have been previously identified to contain Mn, a rare trace element of interest in the study of past habitability [22].

B. Prioritization: MSL LIBS Data

The LIBS instrument collects spectra that can be used to quantify the abundance of different elements in the host rock (often reported as oxide abundance). To determine a sample's scientific importance, we normalize each oxide's concentration such that measurements from the 1^{st} to 99^{th} percentile fill the interval $[0, 1]$. Because Mn enrichment is a particularly interesting paleoenvironmental indicator, we directly map its normalized concentration to the SUE value. The remaining oxides are similarly normalized to a unit interval and mapped to an 8-dimensional DD vector. Such a configuration will prioritize samples with (1) a high MnO concentration and (2) a diverse set of other oxides relative to previous observations. Figure 6 demonstrates the prioritization outcome where the top-5 LIBS measurements are shown when setting the MMR's $\alpha = 1$ (assigning equal balance to science utility and diversity). SYNOPSIS's top ranked measurements all contain high SUE values, but note the variation in other oxides for each row. With this configuration, ground scientists receive observations that both contain elements of interest and represent the chemical diversity of measured samples.

In Figure 7, we further investigate how the balance of utility and diversity compares to a naive First In, First Out (FIFO) prioritization queue. By changing the α parameter, ground teams can focus on accumulating samples that match the science targets of interest or more quickly exploring the full diversity space. Compared to a simple FIFO queue, we found that SYNOPSIS is always more efficient at accumulating utility and often more efficient at accumulating diverse samples. We explicitly compute the advantage of each approach over the naive FIFO approach in Figure 8. With this MSL data subset, there is a clear advantage in quickly accumulating samples with high utility (i.e., samples with MnO). As expected, this advantage wanes as the data bandwidth increases toward an unconstrained amount (Figure 8C, right). We assume unit data volume for all data products to facilitate our quantitative assessment, but SYNOPSIS can rank samples based on utility per bit data volume. When favoring diversity, our data suggests that ground teams can still realize an improvement, but it is smaller and more variable (Figure 8D). These figures demonstrate the range of prioritization capabilities available to ground teams depending on the type and stage of a mission.

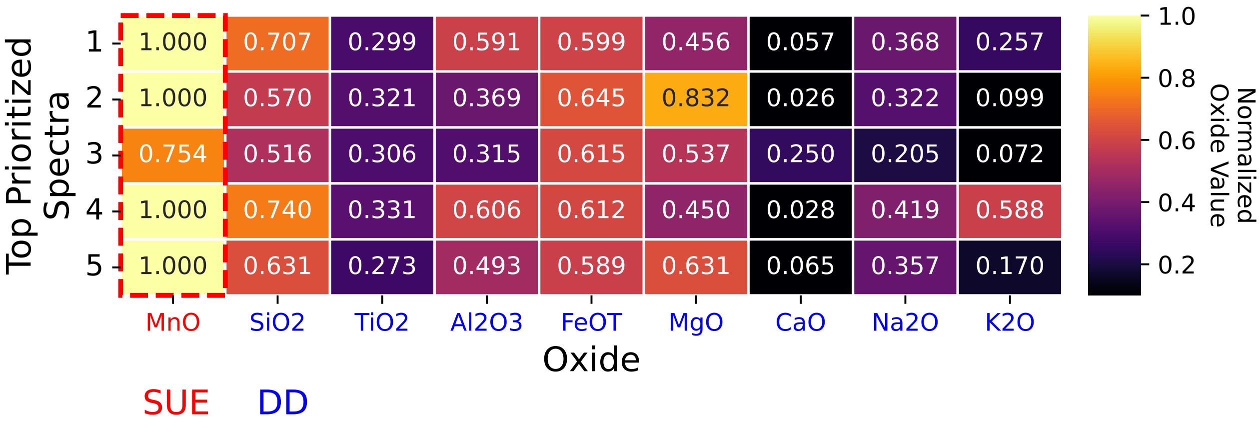

Fig. 6. **Data prioritization example illustrating the trade off between utility and diversity for ChemCam LIBS data.** SYNOPSIS balances utility and diversity of samples based on the SUE and DD. Here, the top five spectra are shown, which all have high MnO concentrations (indicated in red; a potential indicator of past water). The third spectra, however, was prioritized above the fourth and fifth spectra despite having lower concentration of MnO. This is because the other oxide levels (labeled along the horizontal axis in blue) were substantially different (i.e., 5 of the 8 oxides included in the DD were more extreme than in the other samples chosen). Looking across the full set of (non-MnO) oxide levels, SYNOPSIS tends to select samples with a diverse composition.

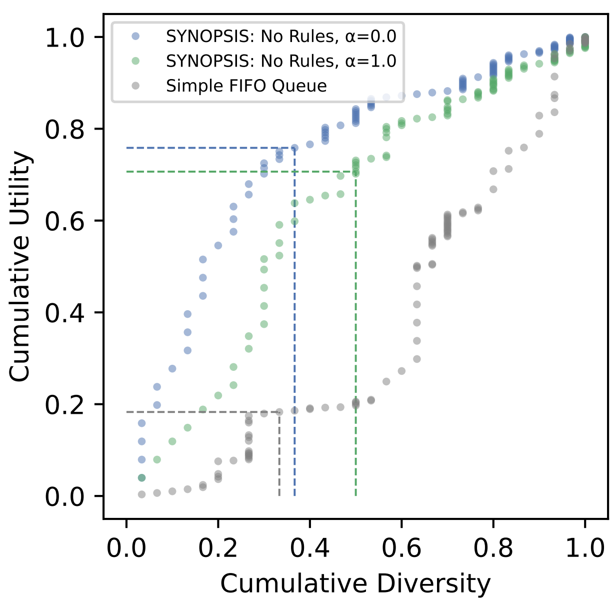

Fig. 7. **SYNOPSIS out-performs the baseline FIFO approach, and operators can tune the alpha parameter to select the desired balance between utility and diversity.** Each point represents one observation and illustrates the growth of cumulative utility and diversity (two proxies for science yield) as more data is transmitted to ground. The same data is prioritized using three different strategies: SYNOPSIS with $\alpha = 0$ (blue; favoring utility), SYNOPSIS with $\alpha = 1$ (green; favoring a balance between utility and diversity), and a simple FIFO queue (gray) where observations are transmitted in the order they are collected. The thin dashed lines indicate the cumulative diversity and utility after 25% of the LIBS data is transmitted.

C. Relationships and Constraints: MSL LIBS+RMI Data

To demonstrate SYNOPSIS's ability to execute multi-instrument data prioritization, we also tested cross-instrument rules to customize how LIBS and RMI observations were ranked in concert with each other. If a geologic sample with a high Mn concentration is identified, the associated before/after RMI images can provide valuable visual context for the LIBS sampling locations (Figure 5). Otherwise, their importance is diminished. To evaluate this logic, we simulated three situations (Figure 9): First, a simple FIFO queue is used to simulate the downlink in the order data was observed. Second, a greedy transmission scheme was tested, which, which transmits data in order of highest adjusted SUE (using diversity-based discounting described in Section III C). For this case and the following, we simply assume RMI images have a static SUE value of 0.1 while LIBS measurements are evaluated based on their MnO content as described previously. (Note that alternative approaches, such as using TextureCam to detect veins within RMI images, could be used to derived SUEs based on directly image content [23].) Finally, we codified the notion that the value of RMI observations should increase when an associated LIBS observation is also of high value. We specified that any RMI SUE value should be tripled (from 0.1 to 0.3) if at least one associated LIBS measurement had a SUE value of ≥ 0.75. We also include a constraint that the first RMI image from a sample pair (i.e., the one before LIBS sampling) should be marked as the lowest priority if no high-value LIBS measurements were identified. Without valuable LIBS observations, it is unlikely both RMIs will be relevant, so we give the tiebreaker to the post-LIBS image as it shows the sample spots.

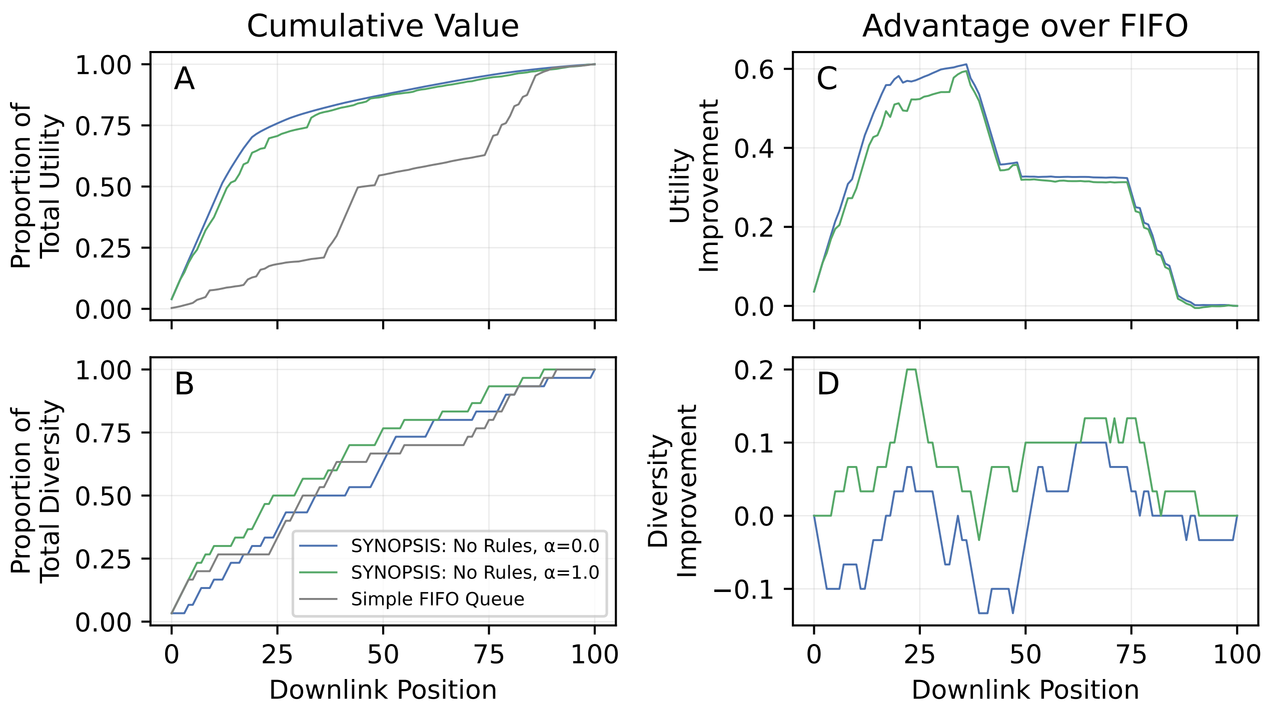

Fig. 8. **Prioritizing LIBS data based on content provides ground teams options to surpass a simple FIFO-based downlink approach.** Cumulative total utility (a) is higher for both values of alpha, and the relative advantage (c) indicates this is especially true if downlink is limited to the first 20-40 products (blue and green traces). Cumulative total diversity assessed as volume represented in the DD hyperspace (b) grows marginally better when balancing utility and exploration (green trace), but the performance across all three methods are more similar. The relative advantage (d) is still highest early in the downlink cycle for an approach that maximizes diversity, but slightly detrimental when favoring utility (blue trace).

In this data, a subset of LIBS samples from "aegis post 2830b", "Mary Anning ccam", and "Ayton" contain high Mn concentrations [22], [24] corresponding to at least one SUE value ≥ 0.75. The FIFO approach leads to suboptimal prioritization as time-of-measurement is clearly not indicative of geologic value here (Figure 9, top-left). Both the second and third cases lead to a more quickly increasing cumulative SUE (Figure 9, top-center and top-right), but note that the third case also transmits contextual RMI images for those with high Mn concentrations earlier in the queue (Figure 9, bottom-right). By combining SYNOPSIS's rules and constraints, science teams would gain faster access to high value chemical measurements along with their geologic context. Such a prioritization scheme is useful under a degraded communications situation because valuable science data reaches ground faster to inform tactical decision making. As an example, SYNOPSIS could rapidly help ground teams decide whether to continue sampling geologic targets at a set of local sites or command the rover to leave in search of better sample locations.

V. CONCLUSIONS

A. Summary

Past work on science autonomy has demonstrated its ability to improve the impact of exploration missions by analyzing science instrument data onboard to make autonomous decisions. It offers a path to increase the yield of future missions especially as the number and complexity of future space exploration missions grows. To facilitate the deployment of science autonomy into future missions, we developed the open-source SYNOPSIS platform and integrated it into the cFS and F-Prime FSW packages. SYNOPSIS manages OSIA execution, stores data products to maintain science knowledge about an environment, and has built in data prioritization functionality. With these capabilities, SYNOPSIS enables downstream autonomous decisions; data prioritization was quantitatively assessed here, but integration with other systems like planning and scheduling software to choose sample sites, adaptive compression of raw data immediately before a downlink pass, or incorporation of science understanding into navigation decisions are all possible. We demonstrated that SYNOPSIS could improve the science return of MSL's ChemCam data if deployed in a situation with degraded communications.

B. Future Directions

Going forward, we plan to improve the SYNOPSIS code base and identify new science autonomy developers who could benefit and/or contribute to related flight-ready software tools. In particular, building a library of off-the-

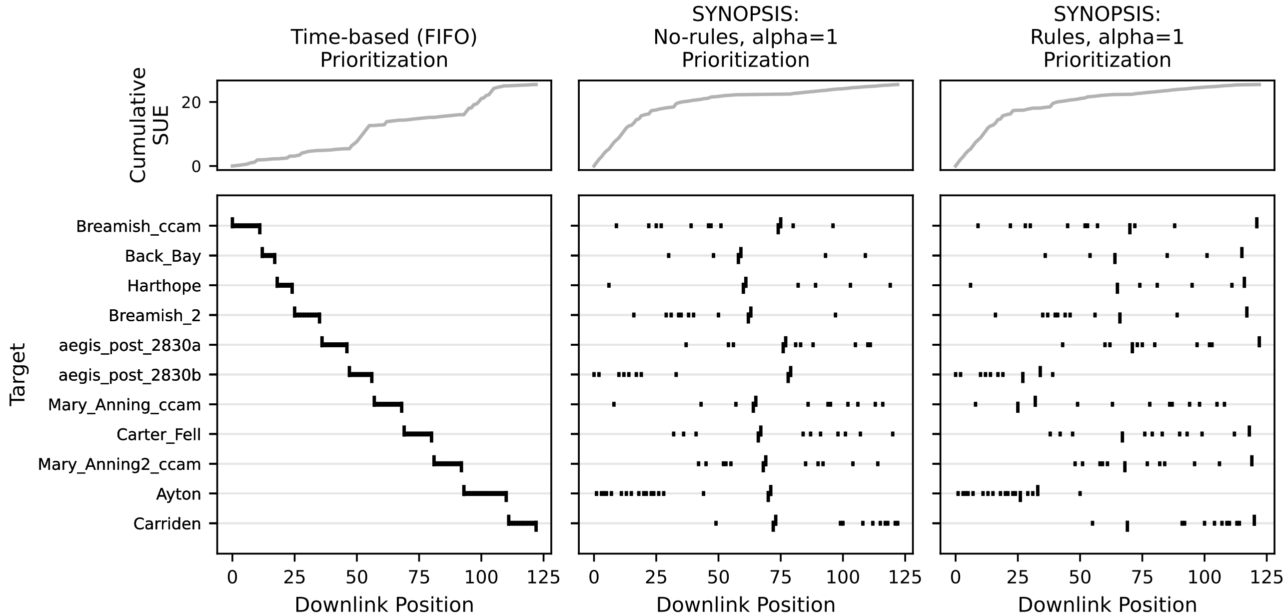

Fig. 9. **Cross-instrument prioritization rules enable sophisticated data prioritization for multiple instruments observing the same sample.** Lower plots show the transmission order for both before/after RMI (long upward and downward ticks, respectively) and LIBS observations (short ticks) from 11 sample sites. Above each plot is the cumulative SUE for downlinked data products. Left: FIFO based prioritization transmits products in the same order they were acquired. Middle: Including content-based assessment order prioritizes data products simply by their SUE values. Here, we assume a static utility for each RMI image, so both before and after RMIs are prioritized at roughly the same position. Right: Including content- and rules-based assessment allows for most RMI images to be deprioritized except when they correspond to high-utility LIBS measurements. The cumulative SUE increases faster for both of the SYNOPSIS-based prioritizations.

shelf OSIA algorithms for fast development and testing is necessary to improve OSIA adoption. This is analogous to a "model zoo" in the machine learning community where well-tested neural networks are built into standard software packages. For the core SYNOPSIS code base, we plan to continue maturing the library to reach TRL 5 for inclusion in mission proposals and technology demonstrations on late stage missions. This will necessitate further simulation and evaluation with hardware test beds as well as potentially terrestrial rover or airborne campaigns. We also plan to make technical improvements to reduce the compute time needed to calculate the MMR discount factors during prioritization, which is slow for large datasets (greater than approximately 100 data products). Of particular interest are the recent advances in vector databases to support Large Language Models (LLMs), which are purpose built to enable fast vector calculations. We also plan to test on flight-like compute platforms including the RAD750, Snapdragon, LEON, and (once available) High Performance Spaceflight Computer (HPSC). Finally, we plan to explore other data prioritization schemes beyond MMR to quantify prioritization performance vs. resources used for the scenario outlined in Section IV. This includes algorithms from information search including diverse near neighbor problem, diversity-aware search, and variations on the knapsack problem. Performance aside, alternative prioritization algorithms deserve investigation to better support missions facing constraints differing from those explored here.

VI. Acknowledgments

We wish to thank Abigail Fraeman, Paul Ramirez, Raymond Francis, John Lucas, Zachary Lynch, Jeff Levison, and Luke Dahl for helpful comments, discussion, and support. In addition, we thank the two anonymous reviewers for their constructive comments that improved this manuscript.

VII. Disclosures

The authors declare no conflicts of interest.

REFERENCES

[1] 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*, vol. 2, no. 7, p. eaan4582, 2017. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.aan4582

[2] S. Mauceri, J. Lee, M. Wronkiewicz, L. Mandrake, G. Doran, J. Lightholder, Z. Cieslarova, M. Kok, M. F. Mora, and A. Noell, "Autonomous ce mass-spectra examination for the ocean worlds life surveyor," *Earth and Space Science*, vol. 9, no. 10, p. e2022EA002247, 2022. [Online]. Available: https://doi.org/10.1029/2022EA002247

[3] M. Wronkiewicz, J. Lee, L. Mandrake, J. Lightholder, G. Doran, S. Mauceri, T. Kim, N. Oborny, T. Schibler, J. Nadeau, J. K. Wallace, E. Moorjani, and C. Lindensmith, "Onboard science instrument autonomy for the detection of microscopy biosignatures on the ocean worlds life surveyor," *The Planetary*

Science Journal, vol. 5, no. 1, p. 19, Jan. 2024. [Online]. Available: http://dx.doi.org/10.3847/PSJ/ad0227

[4] V. Da Poian, E. Lyness, R. Danell, B. Theiling, and W. Brinckerhoff, "Science autonomy and planetary missions: Ml and data science applied to the exomars mission," in *2023 IEEE Aerospace Conference*, 2023, pp. 1–7.

[5] V. Da Poian, E. I. Lyness, J. Y. Qi, I. Shah, G. Lipstein, P. D. Archer Jr., L. Chou, C. Freissinet, C. A. Malespin, A. C. McAdam, C. A. Knudson, B. P. Theiling, and S. M. Hörst, "Leveraging open science machine learning challenges for data constrained planetary mission instruments," *RAS Techniques and Instruments*, vol. 3, no. 1, p. 156–165, Jan. 2024. [Online]. Available: http://dx.doi.org/10.1093/rasti/rzae009

[6] P. A. Taylor, D. C. Catling, M. Daly, C. S. Dickinson, H. P. Gunnlaugsson, A. Harri, and C. F. Lange, "Temperature, pressure, and wind instrumentation in the phoenix meteorological package," *Journal of Geophysical Research: Planets*, vol. 113, no. E3, Mar. 2008. [Online]. Available: http://dx.doi.org/10.1029/2007JE003015

[7] A. Castano, A. Fukunaga, J. Biesiadecki, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castano, and S. Chien, "Automatic detection of dust devils and clouds on mars," *Mach. Vis. Appl.*, vol. 19, no. 5-6, pp. 467–482, Oct. 2008.

[8] S. Diniega, M. Wronkiewicz, and B. Jackson, "Quantifying science gain through strategic data acquisition strategies: A pilot study of martian dust devils," presented at the Tenth International Conference on Mars 2024, 2024. [Online]. Available: https://www.hou.usra.edu/meetings/tenthmars2024/pdf/3096.pdf

[9] S. R. Hirshorn, L. D. Voss, and L. K. Bromley, "Mars concurrent exploration science analysis group (mce-sag) final report," NASA, Tech. Rep. 20170001761, 2017. [Online]. Available: https://ntrs.nasa.gov/citations/20170001761

[10] M. Troesch, F. Mirza, K. Hughes, A. Rothstein-Dowden, A. Donner, R. Bocchino, M. Feather, B. Smith, L. Fesq, B. Barker *et al.*, "Mexec: An onboard integrated planning and execution approach for spacecraft commanding," 2020.

[11] J. Carbonell and J. Goldstein, "The use of MMR, diversity-based reranking for reordering documents and producing summaries," in *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '98. New York, NY, USA: Association for Computing Machinery, 1998, p. 335–336. [Online]. Available: https://doi.org/10.1145/290941.291025

[12] G. J. Holzmann, "The power of 10: Rules for developing safety-critical code," *Computer*, vol. 39, no. 6, pp. 95–99, 2006.

[13] A. Prajapati, "Evolution of Core Flight System (cFS)," nTRS Author Affiliations: Goddard Space Flight Center NTRS Meeting Information: IAIAA SciTech Forum; 2025-01-06 to 2025-01-10; undefined NTRS Document ID: 20240004389 NTRS Research Center: Goddard Space Flight Center (GSFC). [Online]. Available: https://ntrs.nasa.gov/citations/20240004389

[14] R. C. Elphic, G. T. Delory, B. P. Hine, P. R. Mahaffy, M. Horanyi, A. Colaprete, M. Benna, and S. K. Noble, "The Lunar Atmosphere and Dust Environment Explorer Mission," in *The Lunar Atmosphere and Dust Environment Explorer Mission (LADEE)*, R. C. Elphic and C. T. Russell, Eds. Cham: Springer International Publishing, 2015, pp. 3–25. [Online]. Available: https://doi.org/10.1007/978-3-319-18717-4_2

[15] D. Connell and K. Halloran, "Goddard's Core Flight Software Chosen for NASA's Lunar Gateway - NASA," Feb. 2021, section: Earth's Moon. [Online]. Available: https://www.nasa.gov/solar-system/goddards-core-flight-software-chosen-for-nasas-lunar-gateway/

[16] R. L. Bocchino, T. K. Canham, G. Watney, L. Reder, and J. Levison, "F prime: An open-source framework for small-scale flight software systems," 2018. [Online]. Available: https://api.semanticscholar.org/CorpusID:55610828

[17] The F´ Framework Team, "F´: A Flight-Proven, Multi-Platform, Open-Source Flight Software Framework." [Online]. Available: https://github.com/nasa/fprime

[18] T. Tzanetos, M. Aung, J. Balaram, H. F. Grip, J. T. Karras, T. K. Canham, G. Kubiak, J. Anderson, G. Merewether, M. Starch, M. Pauken, S. Cappucci, M. Chase, M. Golombek, O. Toupet, M. C. Smart, S. Dawson, E. B. Ramirez, J. Lam, R. Stern, N. Chahat, J. Ravich, R. Hogg, B. Pipenberg, M. Keennon, and K. H. Williford, "Ingenuity mars helicopter: From technology demonstration to extraterrestrial scout," in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 01–19.

[19] D. M. Geletko, M. D. Grubb, J. P. Lucas, J. R. Morris, M. Spolaor, M. D. Suder, S. C. Yokum, and S. A. Zemerick, "Nasa operational simulator for small satellites (nos3): the stf-1 cubesat case study," 2019. [Online]. Available: https://arxiv.org/abs/1901.07583

[20] R. C. Wiens, S. Maurice, B. Barraclough, M. Saccoccio, W. C. Barkley, J. F. Bell, S. Bender, J. Bernardin, D. Blaney, J. Blank, M. Bouyé, N. Bridges, N. Bultman, P. Caïs, R. C. Clanton, B. Clark, S. Clegg, A. Cousin, D. Cremers, A. Cros, L. DeFlores, D. Delapp, R. Dingler, C. D'Uston, M. Darby Dyar, T. Elliott, D. Enemark, C. Fabre, M. Flores, O. Forni, O. Gasnault, T. Hale, C. Hays, K. Herkenhoff, E. Kan, L. Kirkland, D. Kouach, D. Landis, Y. Langevin, N. Lanza, F. LaRocca, J. Lasue, J. Latino, D. Limonadi, C. Lindensmith, C. Little, N. Mangold, G. Manhes, P. Mauchien, C. McKay, E. Miller, J. Mooney, R. V. Morris, L. Morrison, T. Nelson, H. Newsom, A. Ollila, M. Ott, L. Pares, R. Perez, F. Poitrasson, C. Provost, J. W. Reiter, T. Roberts, F. Romero, V. Sautter, S. Salazar, J. J. Simmonds, R. Stiglich, S. Storms, N. Striebig, J.-J. Thocaven, T. Trujillo, M. Ulibarri, D. Vaniman, N. Warner, R. Waterbury, R. Whitaker, J. Witt, and B. Wong-Swanson, "The chemcam instrument suite on the mars science laboratory (msl) rover: Body unit and combined system tests," *Space Science Reviews*, vol. 170, no. 1–4, p. 167–227, Jun. 2012. [Online]. Available: http://dx.doi.org/10.1007/s11214-012-9902-4

[21] A. R. Vasavada, "Mission overview and scientific contributions from the mars science laboratory curiosity rover after eight years of surface operations," *Space Science Reviews*, vol. 218, no. 3, Apr. 2022. [Online]. Available: http://dx.doi.org/10.1007/s11214-022-00882-7

[22] A. H. Treiman, N. L. Lanza, S. VanBommel, J. Berger, R. Wiens, T. Bristow, J. Johnson, M. Rice, R. Hart, A. McAdam, P. Gasda, P.-Y. Meslin, A. Yen, A. J. Williams, A. Vasavada, D. Vaniman, V. Tu, M. Thorpe, E. D. Swanner, C. Seeger, S. P. Schwenzer, S. Schröder, E. Rampe, W. Rapin, S. J. Ralston, T. Peretyazhko, H. Newsom, R. V. Morris, D. Ming, M. Loche, S. Le Mouélic, C. House, R. Hazen, J. P. Grotzinger, R. Gellert, O. Gasnault, W. W. Fischer, A. Essunfeld, R. T. Downs, G. W. Downs, E. Dehouck, L. J. Crossey, A. Cousin, J. M. Comellas, J. V. Clark, B. Clark, S. Chipera, G. Caravaca, J. Bridges, D. F. Blake, and R. Anderson, "Manganese-iron phosphate nodules at the groken site, gale crater, mars," *Minerals*, vol. 13, no. 9, p. 1122, Aug. 2023. [Online]. Available: http://dx.doi.org/10.3390/min13091122

[23] G. Doran, U. Rebbapragada, E. Song, K. Wagstaff, D. Gaines, R. Anderson, and A. Vasavada, "Scientist-guided autonomy for self-reliant rovers," in *International Joint Conference on Artificial Intelligence*, 2017.

[24] J. Frydenvang, K. Nikolajsen, E. Dehouck, P. Gasda, S. Clegg, A. Bryk, L. Edgar, B. Tutolo, W. Rapin, O. Forni, E. Cloutis, A. Fraeman, A. Vasavada, S. Maurice, R. Wiens, O. Gasnault, and N. Lanza, "Geochemical variations observed by the chemcam instrument as the curiosity rover transitioned from clay- mineral to sulfate-mineral rich regions on aeolis mons in gale crater, mars." presented at the Tenth International Conference on Mars 2024, 2024. [Online]. Available: https://www.hou.usra.edu/meetings/tenthmars2024/pdf/3267.pdf

Mark Wronkiewicz received his Ph.D. in Neuroscience from the University of Washington in 2017. He is a member of the Machine Learning and Instrument Autonomy Group at NASA's Jet Propulsion Laboratory. Mark's work focuses on science autonomy development, leading planetary science research related to the Martian atmosphere, and supporting agricultural monitoring for food security decisions.

Gary Doran received his Ph.D. in Computer and Information Sciences from Case Western Reserve University in 2015. He is currently a member of the Science Data Modeling and Computing Group at the Jet Propulsion Laboratory, California Institute of Technology, where his research focuses on processing data onboard spacecraft to identify science targets for autonomous instrument targeting or downlink prioritization.

Jack Lightholder received a B.S. in Computer Science from Arizona State University in 2016. He is a member of the Machine Learning and Instrument Autonomy Group at NASA's Jet Propulsion Laboratory. Jack supports a variety of missions in concept, development and operational phases, including instrument flight software development for the Near Earth Asteroid CubeSat mission and engineering camera operations for the Mars Science Laboratory, Curiosity.

Zaki Hasnain is a data scientist in NASA JPL's Systems Engineering Division where he participates in and leads research and development tasks for space exploration. His research interests include physics informed machine learning and system health management for autonomous systems. He has experience developing data-driven, game-theoretic, probabilistic, physics-based, and machine learning models and algorithms for space, cancer, and autonomous systems applications. He has developed planetary protection models for the Europa Clipper mission.

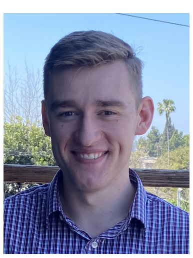

Rex Stockham is a cybersecurity engineer at NASA's Jet Propulsion Laboratory in Pasadena, California. As a member of the Cyber Defense Engineering and Research Group, Rex supports a variety of missions in development and operational phases. Currently, Rex acts as the development lead for a suite of data collection, threat analytics, and visualization tools for mission engineers. Rex is also supporting vulnerability management and cybersecurity compliance efforts on the Mars 2020 mission.

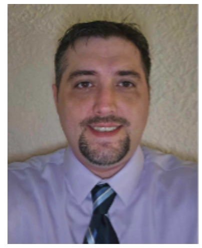

Robert J. Brown is a Systems Engineer of TMC Technologies. Robert is a member of the NASA Independent Verification and Validation (IV&V) Jon McBride Software Test and Research (JSTAR) Team with specialization in the development and modeling of digital twins of NASA spacecraft and SmallSats. Robert is also the JSTAR team's Information Systems Security Officer (ISSO), responsible for facilitating access to JSTAR digital twins and other demonstration capabilities.

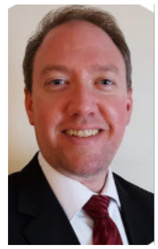

Scott A. Zemerick (Member, IEEE), is TMC Technologies' Chief Scientist and Engineer. Scott is a lead member of the NASA Independent Verification and Validation (IV&V) Jon McBride Software Test and Research (JSTAR) Team where he specializes in the development and modeling of digital twins of NASA spacecraft and SmallSats. Scott is also the Program Manager for the West Virginia SmallSat Center (WVSSC), which focuses on SmallSat mission assurance and orbital debris remediation technologies.

Dr. Lukas Mandrake has been at NASA JPL for 22 years and is the program manager for the Office of Earth Data Science and Technology and the ML Capability Lead for the Office of Technology Infusion. His research accelerates high-impact work of physical scientists through innovative data-driven approaches including novel machine learning deployed both on the ground and onboard spacecraft. He also co-leads the Science Understanding through Data Science (SUDS) and Onboard Science Instrument Autonomy (OSIA) communities.