Integrating Machine Learning for Planetary Science: Perspectives for the Next Decade

A White Paper to the NRC Planetary Science and Astrobiology Decadal Survey 2023-2032

Primary author: Abigail R. Azari, Space Sciences Lab, University of California, Berkeley
Contact information: azari <at> berkeley <dot> edu, 669.244.2996

Co-authors[†]: John B. Biersteker¹, Ryan M. Dewey², Gary Doran³, Emily J. Forsberg⁴, Camilla D. K. Harris², Hannah R. Kerner⁵, Katherine A. Skinner⁶, Andy W. Smith⁷

Secondary Contributors: Rashied Amini³, Saverio Cambioni⁸, Victoria Da Poian⁹, Tadhg M. Garton¹⁰, Michael D. Himes¹¹, Sarah Millholland¹², Suranga Ruhunusiri¹³

Co-signers: Andrew M. Annex¹⁴, K.-Michael Aye¹⁵, Matthew R. Argall¹⁶, Marissa E. Cameron³, Xin Cao¹³, Frank Crary¹⁵, Daniel J. Crichton³, Matthew Cross¹⁷, Mario D'Amore¹⁸, Julien de Wit¹, Raymond Francis³, Jörn Helbert¹⁸, Don R. Hood¹⁹, Catriona M. Jackman²⁰, Anthony Lagain²¹, Raven E. Larson¹⁵, Chuhong Mai²², Lukas Mandrake³, Andrew Needham²³, Victor Pankratius¹, James Parr²⁴, Louise M. Prockter²⁵, Fernando Pérez²⁶, Jani Radebaugh²⁷, Bethany P. Theiling⁹, Matthew Tiscareno²⁸, Andrew Vanderburg²⁹, Jon Vandegriff³⁰, Josh Vander Hook³, Indhu Varatharajan¹⁸, Kiri Wagstaff³, Ingo Waldmann³¹

Affiliations:

¹ Massachusetts Institute of Technology

² University of Michigan

³ Jet Propulsion Laboratory, California Institute of Technology

⁴ University of Idaho

⁵ University of Maryland at College Park

⁶ Georgia Institute of Technology

⁷ Mullard Space Science Laboratory, University College London (UCL)

⁸ Lunar and Planetary Laboratory, University of Arizona

⁹ NASA Goddard Space Flight Center

¹⁰ University of Southampton

¹¹ University of Central Florida

¹² Princeton University

¹³ University of Iowa

¹⁴ Johns Hopkins University (JHU)

¹⁵ Laboratory for Atmospheric and Space Physics, University of Colorado

¹⁶ Space Science Center,

University of New Hampshire

¹⁷ Institute for Earth and Space Exploration, Western University

¹⁸ German Aerospace Center (DLR)

¹⁹ Texas A&M University

²⁰ Dublin Institute for Advanced Studies

²¹ Space Science and Technology Centre, Curtin University

²² Arizona State University

²³ Jacobs, NASA Johnson Space Center

²⁴ Frontier Development Laboratory

²⁵ Universities Space Research Association, Lunar and Planetary Institute

²⁶ University of California, Berkeley, Lawrence Berkeley National Laboratory

²⁷ Brigham Young University

²⁸ SETI Institute

²⁹ University of Wisconsin, Madison

³⁰ JHU Applied Physics Laboratory

³¹ Centre for Space Exochemistry Data, UCL

[†]We would like to recognize the extraordinary effort which this decadal has taken and the members of our community who were unable to participate in this work. We would also like to acknowledge conversations with white paper teams on data management, automation, and other machine learning relevant contributions and we encourage you to review these additional data science relevant papers. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

1 Introduction

In the last decade, machine learning has emerged as one of the most promising complements to traditional data analysis and modeling methods in scientific fields. In one of the most profound examples, the first image of a black hole was captured by applying a machine learning algorithm to petabytes of data collected from eight telescopes [1]. Since planetary science’s last decadal survey, the use of machine learning has increased in each division of NASA’s Science Mission Directorate (SMD). However, even though the number of planetary science publications involving machine learning has grown exponentially over the last ten years, it lags in both percent share and growth rate compared to heliophysics, astrophysics, and Earth science (Figure 1). In this white paper, we assert that planetary science, similar to related disciplines, has the opportunity to leverage machine learning methods for scientific advancement in our field.

Planetary science is rapidly transforming into a data rich field with missions in the next ten years set to collect larger and more detailed datasets than ever before. For example, in the fourteen years orbiting Mars, the Mars Reconnaissance Orbiter has already returned more data than all other interplanetary missions combined, including missions currently active (see [MRO By The Numbers](#)). These missions increase the total data available to our community and increase the rate at which we receive new observations. This can pose challenges to data analysis efforts that machine learning can help address. For example, the Transiting Exoplanet Survey Satellite (TESS) generates ~ 1 million new light curves every month, pushing the exoplanet community toward automated candidate identification [2]. Integrating machine learning into planetary science has the potential to enhance current endeavors and lead to new discoveries in our field.

The broad applicability of data science methods is considered a new branch of scientific investigation [3]. Data from planetary science missions and the field’s interest in interpretable models and reproducible science pose challenges to immediate application of off-the-shelf machine learning methods [e.g., 4]. However, these challenges are shared by other fields, including Earth science [5]. We can learn from demonstrated funding programs and researchers in sibling disciplines to adopt field-specific machine learning methods. These efforts should embrace interdisciplinary partnerships across multiple application fields and industry to address unique challenges in planetary science as they manifest in the context of machine learning. Our vision for the next decade is one rich with applications of machine learning at every stage of planetary missions. We envision machine learning enhancing current activities by by operating alongside established methods.

In recognition of the novelty of these methods to the field, a concerted effort is needed to ensure integration and applicability of machine learning methods in planetary science.

Figure 1: Yearly trends of machine learning across NASA SMD topics as a percentage of published literature. For comparison, in 2019 the term ‘modeling’ had near identical occurrence rates within publications in both planetary science and Earth science ($>50\%$). Data from [Elsevier’s Scopus](#) publication database.

The purpose of this white paper is two-fold. First, to introduce and demonstrate to the reader the applicability of these methods at every stage of a mission (Section 2). Second, to offer recommendations to ensure the next decade of planetary science capitalizes on the benefits of machine learning for scientific advancement (Section 3). **To support these applications we recommend a sustained commitment to the ecosystem of machine learning in planetary science, independent of individual mission funding.**

2 Machine Learning at Every Stage of the Mission Lifecycle

2.1 Mission Development

Mission development requires surveys of existing datasets and the development of models to translate science goals into discrete engineering requirements. Data science techniques are already informing the design of missions currently under development. In this section we highlight several examples of how ML can be incorporated to efficiently develop future missions.

Planning trajectories for missions is critical for meeting requirements. Trajectory design is constrained by vehicle configuration and fuel capacity, resulting in a non-trivial problem that can be computationally expensive [6]. ML can be used to improve accuracy and reduce computational cost of optimal trajectory design. Recent work [7] has shown the potential for ML to address current challenges by demonstrating the feasibility of training a deep neural network on the ground to enable real-time onboard spacecraft guidance.

Earth and laboratory analog studies can be combined with machine learning to develop mission-like training datasets for missions exploring new environments. This style of transfer learning, in which a trained machine learning model is used on a different dataset, could offer significant advantages from the bevy of Earth data we have toward mission planning.

Initial landing site selection is labor intensive and can be subject to human errors, with gigabytes of terrain maps being labeled by hand and locations evaluated with respect to predicted landing ellipses. ML and data-driven techniques can efficiently provide high-level classification of surface terrain across large areas to automate preliminary selection of potential sites. Recent work to aid in landing site selection of the Mars 2020 mission [8] leveraged ML to enable automated classification of orbital images by terrain type, allowing scientists to focus on choosing landing sites with the highest potential for scientific return.

Evaluation of mission architecture and instrument design against scientific requirements cycles through design, testing, and improvement. Data science techniques can improve and accelerate this cycle through robust uncertainty propagation and surrogate models that approximate the re-

What is Machine Learning?

Machine learning (ML) is a type of artificial intelligence that refers to computer algorithms that improve (or learn) through inclusion of data. ML methods are a highly flexible class of algorithms designed to adjust their parameters to the structure of the data they have access to and produce models that make predictions. This flexibility makes them suitable for various uses and the scientific community is still exploring and discovering novel ways in which they can become part of our toolbox. In the context of this white paper, ML ranges in complexity from linear regression to neural networks, with applications ranging from image classification to pattern recognition. Within this white paper we use data science as an umbrella term for methods that take advantage of large data sets, such as machine learning.

sults of high-fidelity simulations. For example, during hardware testing, an iterative approach using Bayesian inference can improve model validation and reduce uncertainty in instrument performance [9]. Probabilistic models of individual components can be combined into estimates for the performance envelope of an entire system, resulting in more reliable benchmarking against design requirements early in the design cycle [10]. Probabilistic models combined with models of the physical system can generate simulated datasets that data pipelines and analysis techniques can be tested against for satisfaction of requirements.

These examples show how data science methods impact the development of missions. In the next decade we expect machine learning techniques to improve mission designs by enabling the use of larger datasets and enhancing mission development.

2.2 Mission Operations

Machine learning shows promise for increasing scientific return from planetary missions by addressing limitations in science operations and data return. While planetary missions have increasingly generated and returned more data to Earth over the last several decades, physical limitations on communication bandwidth and latency combined with an increasing number of missions that must share Deep Space Network resources mean that much less data can be returned compared to Earth-orbiting missions. Data science techniques can automate or assist with science operations. The use of ML in tandem with onboard planning enables new mission architectures and exploration of high science value, high risk destinations. In this section, we describe opportunities for integrating ML into operations onboard spacecraft and on the ground.

Automated content-based alerts can accelerate science planning by notifying scientists of interesting observations. Scientists might have as little as 5 hours to decide which commands to send to the spacecraft at the next communication opportunity [11]. Using automated classification, science planners could filter new observations to review only those containing features of interest, or could “subscribe” to be notified when an onboard algorithm detects such features. Current examples include implementing a convolutional neural network to classify geologic features and spacecraft hardware within Mars rover and orbital imagery [12] and automatically generating image captions for content-based search or notification capabilities [13].

Prioritization of novel observations can also help focus scientists’ attention on unusual phenomena and respond with any follow-up observations while there is still an opportunity to do so. Applications have prioritized the order of spectra shown to an expert, improving the accuracy of spectrum classification [14] and have identified novel geology in multi-spectral images to flag observations for mission team members [15]. See Figure 2 (left) for a successful example on detecting Martian geologic features.

Onboard detection and summarization (i.e., high-level observation summaries) of features can greatly increase the number of observations containing these phenomena returned from a mission. It is nearly impossible to sequence observations of transient, unpredictable events (e.g., dust devils, storms, plumes) in advance. Searching for phenomena onboard a spacecraft enables downlinking only those candidate observations likely to contain the desired target or returning more data for targets [16]. This approach has been used on several missions, including the delivery of observation content summaries from EO-1 [17] and the automatic detection of dust devils on Mars Exploration Rover Opportunity [18]. Even imperfect algorithms will increase the fraction of valuable observations in the downlink.

Autonomous targeting can recover mission time spent waiting for ground-in-the-loop targeting decisions. Such methods are becoming essential for certain operations. For example, the AEGIS system (Figure 2), used to identify geologic targets for follow-up observations onboard Mars Exploration Rovers and Mars Science Laboratory, has increased observations by $\sim 28\%$ per day [19].

Onboard autonomy will be increasingly important for missions at far distances (icy moons and icy giants), with multiple spacecraft, and with short duration and limited communication opportunities (Europa and Venus landers). Applications of autonomy are extremely valuable, but face challenges in hardware and software design and in producing explainable and interpretable results. Long-term success will require addressing these challenges (Section 3), but promises to improve the scientific return of planetary missions.

2.3 Analyses

Statistical techniques have long been used to analyze planetary data to advance scientific understanding. Recent increases in data volume returned by planetary missions have expanded the applicability of data science techniques. We present a high level overview of leading edge applications that these methods and data volumes enable.

Detection, characterization, and classification with ML can systematically and autonomously identify features of interest and distinguish real objects and events from false positives. Applications range from standardizing and identifying rare plasma intensifications in a time series [21] to vetting large datasets for exoplanet candidate identification [22]. Compared to compiling such events or features by eye, ML can be employed as a tool to find and reduce bias while improving reproducibility for scientific analyses.

Grouping features (clustering) into categories with statistics-based assumptions has the potential to be a major method leading to groundbreaking findings in planetary science. Such techniques have been implemented in Bayesian clustering for mineral identification on Mars [23]. ML also offers the opportunity to develop techniques that can be quickly and feasibly applied to much larger data volumes. For example, the development of PlanetNet, a tool to “quickly and accurately map spatial and spectral features across large, heterogeneous areas” illustrated in Figure 3, uses a dataset that is too large for manual identification [24].

Resolving (and integrating across) multi-dimensional non-linear trends is of great interest in studying planetary environments. Traditional statistical analysis struggles to account for these multivariate trends, particularly when the complete parameter space has not been thoroughly sampled. Data science allows for the identification and description of simultaneous, multivariate, and nonlinear trends; for example using temporal variability to estimate spatial scales [25].

Figure 2: Examples of automated feature detection on Mars with Mars Science Laboratory. Left: Mastcam images of a meteorite (top) and light-toned material (bottom) corresponding to red pixels detected by a ML algorithm [20]. Right: Navcam image of the Martian surface highlighting targets for follow-up analysis from the Autonomous Exploration for Gathering Increased Science (AEGIS) [19].

Describing large physical systems with spatiotemporal varying datasets can be challenging. Planetary missions can only capture small snapshots of large-scale systems. Statistical techniques can use these snapshots to construct global descriptions of the system. This can be done, for example, by combining observations into useful proxies [26].

ML integration in physical models can build on the physics-based modeling already embraced by the field. Large-scale physical models have been a mainstay for interpreting planetary science datasets for decades and ML techniques are quickly gaining use in enhancing physics-based model performance and returns [27].

Advanced statistical techniques are adding a new branch of methods with which to engage in data-driven scientific pursuit. Similar to other mission stage applications, interpretability of these methods' results is an active challenge. To fully enable ML applications, it is therefore of critical importance to develop these techniques within planetary science.

3 Recommendations

To reach our future vision of the field, we have crafted ten recommendations for integration of machine learning into planetary science for the next decade.

1. Invest in post-mission and non-mission related funding and activities. Due to the structure of the NASA Planetary Science Division, funding opportunities are often tied to missions or demonstrable scientific goals. Because ML techniques are best suited to larger datasets, we expect many ML proposals will focus on missions considered retired or ended. Increased funding for post-mission studies would assist in these applications. Beyond missions, research is needed for how ML techniques will address specific challenges in planetary science data. ML applications to physical problems are still a major research area. Recent findings highlight the need for incorporating domain knowledge to increase interpretability of ML [4] and propose a machine assisted pipeline in scientific discovery [28]. Basic investigations into domain-specific concerns of ML techniques are needed and are not funded by missions.

2. Expand programs for data analysis. ML applications currently fit into data analysis programs, such as the existing NASA Planetary Data Archiving, Restoration, and Tools and Data Analysis Programs. We recommend a review of these programs for inclusion of ML.

3. Increase support for planetary data infrastructures. The NASA Planetary Data System (PDS) forms the backbone of NASA's planetary data. We recommend increased discussions for PDS and other data and tool hubs (e.g. [MAPSIT](#)) to support ML applications.

4. Invest in onboard software and hardware. A maturation pipeline is needed to mitigate risk of including instruments and spacecraft capable of onboard data processing in mission proposals. We recommend including ML in [PICASSO](#) and [MatISSE](#) or developing new programs for increasing technology readiness levels of onboard software and automation.

5. Invest in the continued development of data return facilities. For ML users interested in data from missions traveling far from Earth (like the upcoming Dragonfly and Europa Clipper missions), data return and the cost of transmission limit data sizes. In addition to onboard solutions

Figure 3: Stormy regions (red) in Saturn's clouds detected by Planet-Net. Reprinted from [24] by permission from [Nature Astronomy](#).

(see previous item), we support investments for updates to the Deep Space Network and other solutions for higher data return from the outer planets.

6. Review changing demands for computational resources on the ground. Greater usage of NASA’s high performance computing facilities is expected as ML applications increase the demand on computational resources. We recommend that these facilities receive continued support.

7. Invest in education. Educational resources for new users of ML methods are essential. Research programs like the **Frontier Development Laboratory** work to integrate artificial intelligence with NASA research and provide training for students and early career professionals. We recommend expansion of these programs and investigation into models that would educate larger groups (see **AI4ESS** for a potential model). On longer timelines we see planetary science degrees incorporating data rich courses [see 29, for necessity of such courses in education] and hiring the staff required to teach such courses (see **NSF FDSS** for an program that encouraged investment through faculty sponsorship).

8. Support community groups and early career researchers. While the community of ML users in planetary science is growing, in our observation it is less cohesive than other subsets of the field and is primarily composed of early career participants. On this white paper for example, all of the authors are early career professionals or students. Supporting early career researchers and groups like **OpenPlanetary** that focus on planetary data access will help connect researchers. The NSF EarthCube **Intelligent Systems for Geosciences** provides an ideal example.

9. Actively support underrepresented members of our community. Geoscience and computer science are source fields for ML in planetary science. These fields, respectively, are the least racially diverse of all STEM fields [30] and have some of the lowest shares of women degree recipients with declining participation in recent years [31]. Opportunities must actively include the participation, experiences, and leadership of underrepresented members of our communities or they will reinforce existing inequalities and propagate injustices into the next decade.

10. Support interdisciplinary work. For the continued success and growth of ML in domain applications, scientific fields must engage in interdisciplinary work [32], such as with ML industry researchers and physical scientists of other fields. This furthers Strategy 1.3 in the NASA SMD 2020-2024 **strategy plan** to “Advance discovery in emerging fields by identifying and exploiting interdisciplinary opportunities”.

A Single Solution: ‘SpaceCube’

Our recommendations represent different aspects of a single target: the sustained commitment to people, applications, and infrastructure outside of mission funding. We have looked across fields for examples for our future, referencing the joint program between the NSF Directorate for Geosciences and the Division of Advanced Cyberinfrastructure called **NSF EarthCube**. For Earth research, EarthCube envisions almost every aspect of our recommendations. EarthCube funded projects have supported community groups, worked to create data infrastructures, and defined new research areas. Given NASA’s excellence in planetary research, and NSF’s EarthCube program, an NSF-NASA partnership would further our vision for the field to engage in purposeful work in machine learning for planetary science.

We propose a ‘SpaceCube’ partnership between NASA and NSF to target areas of shared advancement to address these recommendations.

4 Conclusions

The future of planetary science is exciting, with each new discovery changing our understanding of the solar system and beyond. Machine learning and data science are poised to advance these discoveries and add value at each step of the mission lifecycle. In this white paper we have envisioned a data rich future for our field achieved by integrating these methods. For the full incorporation of machine learning into planetary science, investigations into domain-specific applications for our field must be supported. The successes of adjacent fields have informed our recommendations, which will enable planetary science to join them in benefiting from machine learning techniques for scientific insight. The NSF EarthCube program exemplifies many of our recommendations; but for Earth science. As such, we urge the survey to consider cross-agency collaborations to leverage current activities outside of NASA. The [last decadal survey](#) stated that planetary missions have benefited from “laboratory studies, theoretical studies, and modeling”. We look forward to seeing data science and machine learning methods included in the next decade.

References

[1] [Bouman et al., 2016](#), “Computational Imaging for VLBI Image Reconstruction” | [2] [Yu et al., 2016](#), “Identifying Exoplanets with Deep Learning III: Automated Triage...” | [3] [Hey et al. 2009](#), “The Fourth Paradigm: Data-Intensive Scientific Discovery” | [4] [Azari et al. 2020](#), “Incorporating Physical Knowledge into Machine Learning for Planetary Space Physics” | [5] [Karpatne et al 2019](#), “Machine Learning for the Geosciences: Challenges and Opportunities” | [6] [Mereta et al. 2017](#), “Machine learning of optimal low-thrust transfers...” | [7] [Izzo et al. 2019](#), “Machine Learning...in Interplanetary Trajectory Design” | [8] [Ono et al. 2016](#), “Data-driven surface traversability analysis for Mars 2020...” | [9] [Stout 2015](#), “Bayesian-based simulation model validation for spacecraft thermal systems” | [10] [Uebelhart et al. 2012](#), “Uncertainty Characterization in Integrated Modeling” | [11] [Wilson et al. 2017](#), “NASA Mars 2020 landed mission development” | [12] [Wagstaff et al. 2018](#), “Deep Mars: CNN Classification of Mars Imagery...” | [13] [Ono et al. 2019](#), “Make Planetary Images Searchable: Content-based search for PDS...” | [14] [Wagstaff et al. 2013](#), “Guiding scientific discovery with explanations using DEMUD” | [15] [Kerner et al. 2020](#), “Comparison of novelty detection methods for multispectral images...” | [16] [Wagstaff et al. 2019](#), “Enabling Onboard Detection of Events...for the Europa Clipper...” | [17] [Chien et al. 2008](#), “Results from Automated Cloud and Dust Devil Detection...” | [18] [Thompson et al. 2013](#), “Autonomous Spectral Discovery...Onboard the EO-1 Spacecraft” | [19] [Francis et al. 2017](#), “AEGIS autonomous targeting for ChemCam...science team use” | [20] [Kerner et al. 2019](#), “Novelty Detection...Application to Planetary Exploration” | [21] [Azari et al. 2018](#), “...Statistical Survey of Energetic H⁺ Sudden Flux Intensifications” | [22] [Millholland & Laughlin 2017](#), “Supervised Learning Detection...Hot Jupiter Candidates” | [23] [Dundar & Ehlmann](#), “Rare jarosite detection in CRISM...non-parametric Bayesian clustering” | [24] [Waldmann & Griffith](#), “Mapping Saturn using deep learning” | [25] [Smith et al. 2018](#), “Evaluating Single-Spacecraft...Simple Monte Carlo Simulations...” | [26] [Ruhunusiri et al. 2018](#), “An Artificial Neural Network for ...Solar Wind Proxies at Mars” | [27] [Cambioni et al. 2019](#), “...Machine Learning Applied to Simulations of Giant Impacts” | [28] [Pankratius et al. 2016](#), “Computer-Aided Discovery: Toward Scientific Insight Generation with Machine Support” | [29] [NASEM 2018](#), “Data Science for Undergraduates: Opportunities and Options” | [30] [Dutt 2020](#), “Race and racism in the geosciences” | [31] [NSF, NCES 2019](#), “Women, Minorities, and Persons with Disabilities in Science...” | [32] [Ebert-Uphoff et al. 2019](#), “Thoughtfully using artificial intelligence in Earth science”