A Concept for Planetary Drilling Autonomy

Sarah Boelter[1], Ebasa Temesgen[1], Greta Brown[1], Mario Jerez[1], Elsa Forberger[1], Brian J. Glass[2], and Maria Gini[1]

***Abstract*— In planetary environments, robotic agents must reason through faults before they escalate to mission-critical failures. No broadly applicable solution exists to give a specialized agent like The Regolith and Ice Drill for Exploring New Terrain (TRIDENT) awareness for when a situation may escalate to a drilling fault. We propose building on our previous work with online time-series subspace analysis methods and percussive beat frequency detection techniques to define a Markov Decision Process to autonomously avoid drilling faults.**

I. Introduction

In April 2024, the NASA Space Technology Mission Directorate (STMD) released an internal and external feedback opportunity to solicit feedback on 187 civil space shortfalls. STMD defines shortfalls as technology areas requiring development to meet future mission needs [1]. Currently there is a need for robotic autonomy for sensing and mitigating risks from robotic operation in challenging environmental conditions, often with hardware and software limitations.

This paper focuses on a robotic drill, known as "The Regolith and Ice Drill for Exploring New Terrain" (TRIDENT) [2], shown in Figure 1. It is a 1-meter rotary percussive drill manufactured by Honeybee Robotics. We collected 30 instances of fault data with the TRIDENT drill with NASA Ames and the Goddard Institute Field Team (GIFT) in the Bishop Tuff in Fall 2023 and at a Haughton Crater Analog site in 2024. TRIDENT gives indicators, visually or otherwise, of a failure well before it happens. Data indicates a failure has a high likelihood of prediction, but the necessary components to intelligently monitor actions do not exist yet. We propose building on previous work with online time-series subspace analysis methods and percussive beat frequency detection to develop a Markov Decision Process to autonomously predict and avoid drilling faults. This work is conceptual, and focuses on analysis and developing theoretical solutions for implementation during summer lab testing.

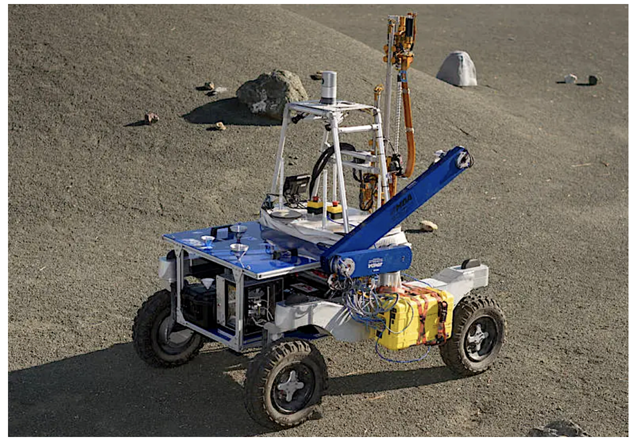

Fig. 1. Trident drill on Rover during ARADS project [3]

II. Related Work

A. Autonomy Under High-Risk Scenarios

Autonomy under high-risk and little oversight conditions is not a new concept and many ideas have been proposed on how to best solve this over the decades. Muscettola et al. [4] discuss issues surrounding the Sojourner rover and its mission to Mars in 1997. Since the rover lacked onboard autonomy software, operating it for several months was incredibly taxing on the ground crew, substantiating the need for onboard autonomy with minimal human oversight in future missions. They described three requirements for space operations: autonomy for long duration, methods with guaranteed success, and high reliability. Paulsen et al. [5] go into detail about U.S. and Soviet drilling missions as far back as the early 1970's. Specifically the about issues encountered with human-operated rotatory percussive drilling with Apollo 15 and the difficulty with the drill extraction. The Soviet lander Luna 16, 20 and 24 missions in the 1970's could acquire lunar soil autonomously through drilling to a maximum depth of 2 meters. The main difference is that the landers in question weighed several thousand kilograms, and could provide much higher downward force on the drill, unlike modern landers and rovers which average around 200 kg. Paulsen et al. also go into some of the challenges of drilling in space and the two broad aspects that need to be considered. First, the need for automation and what must be considered to have drilling autonomy, like vertical force, and limited computing capacity. Second, the environmental conditions that differ from Earth.

B. Autonomy in Petroleum and Mining

Recently the Norwegian Research Center conducted research on autonomous drilling for petroleum. Specifically, de Wardt et al. [7] published a taxonomy for autonomous drilling. Due to the complexity, uncertainty, and poor observability in drilling, it is suggested automation technologies

[1]Sarah Boelter, Ebasa Temesgen, Greta Brown, Mario Jerez, Elsa Forberger, and Maria Gini are with the Department of Computer Science and Engineering at the University of Minnesota in Minneapolis, MN. {boelt072, temes021, brow6802, jerez005, forbe140, gini}@umn.edu

[2]Brian J. Glass is with NASA Ames Research Center in Moffett Field, California brian.glass@nasa.gov

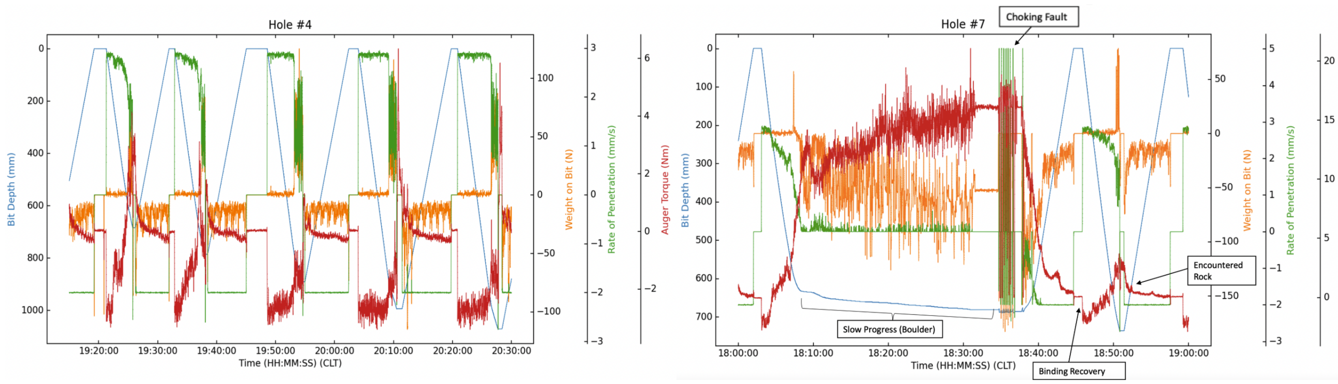

Fig. 2. The left image shows normal drilling operations during testing. The right image illustrates choking fault during testing. Both testing done in Devon Island, Canada [6]

applicable to other industries like aviation may not apply to drilling [7]. When compiling perspectives for the taxonomy, de Wardt et al. discuss the Autonomous Systems Taxonomy developed at NASA by Fong et al. [8], specifically its focus on collaboration and interaction as one of its main axes of analysis for classifying autonomous systems. This is especially important because in drilling systems, autonomy is not simply an all-machine environment. Due to the high level of uncertainty, the sparsity of data, and the complex problems and environment, it is often necessary for autonomous agents to interact and collaborate with human agents. Plan generation requires evaluating consequences of actions in the short, medium, and long term [7].

Another complication of autonomy inherent to petroleum and relevant to planetary drilling is weak environmental observability. The uniqueness of sparse data observed from a distance makes autonomy very challenging. To estimate the subsurface environment and drilling operation state, questions difficult to answer accurately about the drilling operation must be asked [7].

To achieve autonomous operations, methods must be utilized that estimate the internal state of the drilling system. Cayeux et al. [9] developed a method based on a Markov Decision Process to optimize the drilling operation, so the drill can complete operations while minimizing time lost to unexpected drilling events. Specifically, actions are decided based on factors like hole depth, thresholds are adjusted and updated based on operational conditions, and states are defined based on friction coefficients [9]. Some issues stated by Cayeux et al. are not relevant to our work, specifically the flow rate, as we are not drilling for oil, and the mechanisms in a scaled-down planetary drill are much simpler than those in a larger drilling rig.

III. Previous Work

The Atacama Rover Astrobiology Drilling Studies (ARADS) project [2] aimed to explore and understand the mobility and distribution of salts, compounds, and biosignatures down to a 1 meter depth in planetary environments using a rover with an attached TRIDENT drill. This system was tested in Chile in 2019. Automation software was prototyped. While The ARADS project helped automate tasks, the drill still requires constant oversight and environmental assessment. Glass et al. [2] detail how it is impossible to know without prior geological surveying what exists below the surface, especially in planetary environments.

Time series data points like in Figure 2 can give us indications faults like choking faults, hard-materials faults, corkscrewing faults, bit inclusions, and binding faults as represented in Table I. The data is non-static over time, and depending on conditions faults could span as little as a few seconds to several minutes.

Previous work predicting situational risk factors for critical errors for TRIDENT used change point analysis techniques, building off [10]. Subspace-based methods are best at detecting instantaneous changes while not requiring an extensive collection of historical sequences. They also do not extract statistical properties but mainly characteristic shapes and ignore noise, as the decomposition acts as a filter for random signal contributions. We think that these methods are promising for the unique properties of our application scenario. We now use a new time-series change analysis technique, the Enhanced Singular Spectrum Transformation (ESST) [11]. Instead of computing the decomposition and comparing the extracted characteristics separately, ESST combines both trajectory matrices in one matrix and computes the score directly after decomposition, using more characteristics than the SST [12]. We conducted testing with offline data, and saw that the severity of the change score can help determine if the drill is entering a faulty state.

Techniques from the field of structural health management (SHM) have been used on previous models of extraterrestrial drills [13]. This previous work excelled at detecting oncoming faults, but the diagnostic rate was too slow to catch those that onset in under 20 seconds. In addition, previous SHM systems were developed on drills that did not percuss. Our current exploratory work builds on previous SHM research by providing an unsupervised ensemble-learning method of identifying percussive beats in drill vibration, and shows that the frequencies present during percussive strikes can be used to perform high-frequency health diagnostics. These frequencies can be coupled with techniques like ESST to help inform machine learning models of oncoming drill faults.

IV. Fault Types and Drill Variables

Before analyzing the drill data, it is important to understand the data types and how data must be processed. We analyzed data collected at field events over the last decade to determine what data would be the best candidates for processing initially. We examined both electronic data and notes from field testing. Most of the time-series data has a similar length, however, there were some examples both exceptionally long or short, potentially due to different types of material being drilled into.

It is important to understand the types of faults encountered in the drilling process, illustrated in Table I. We combed through data collected after 2016 and notes in search of data that primarily resembled choking and binding faults.

TABLE I
Fault Types for planetary drilling [13]

Fault Type	Description
Binding Fault	Increased torque to to friction on drill string
Choking Fault	Cuttings caught in borehole increasing torque
Hard-Materials Fault	Stalled ROP with increased torque
Corkscrewing Fault	Flutes caught on protruding rock
Bit Inclusion	Gravel caught in drill flutes increasing torque

Time-series examples from two different boreholes shown in Figure 2. The first shows normal drilling operation and the second illustrates a fault. We inferred that faults tend to raise variables shown in Table II like Torque and Weight On Bit while the Rate of Penetration is roughly zero. This indicates we need methods able to detect shifts under 30 seconds. From observing plots and notes, we inferred choking and binding faults tend to raise Torque and Weight On Bit while the Rate of Penetration is roughly zero. This indicates that multiple variables can be used to predict faults.

TABLE II
Drill Variables

Variable	Description
Torque (Nm)	Force by motor for drill string rotation
Rate of Pen. (mm/s)	Velocity of the drill string relative to Y axis
Weight On Bit (N)	Percussion force applied by the drill string
Depth (m)	Depth of drill string from 0-1m

V. Problem Statement

Drilling is a notoriously unobservable problem. We built off previous work of Cayeux et al. [9] to optimize the time needed to drill to the total desired depth. We can decompose this problem into two parts: the time to execute the series of actions, and the time to mitigate any drilling faults if they occur, which can be expressed as:

$$\Delta t_{TD}(A) = \sum_{a \in A} t_a(a) + \sum_{f \in F} P(f) t_m(f) \tag{1}$$

where Δt_{TD} is the estimated time to reach the total depth (TD), A is a set of actions a, t_a is a function that estimates the duration of an action, f is a possible drilling event from the set of possible drilling faults F, $P(f)$ is the occurrence probability of a drilling fault, and t_m is the estimated time to mitigate a drilling fault. We want to minimize the number of actions to reach the total depth. The objective function is:

$$\underset{A \in \mathcal{A}}{\operatorname{argmin}} \ \Delta t_{TD}(A) \tag{2}$$

where $\mathcal{A}$ is a set of all possible series of actions to reach the total depth TD.

VI. Decision Making While Drilling

A. Defining the MDP

We can represent our MDP for TRIDENT as a five-tuple $\langle S, A, T, R, \gamma \rangle$ where

1) S is a state as defined by equation 3,
2) A is a set of all possible actions in S, which, for example, is defined as $A = \{continue, retract\}$, where *continue* equates with continuing to drill, and *retract* equates to retracting before drilling to the specified depth.
3) T is the transition model, which represents for each state s and action a the probability of reaching another state s'. We write T as $T(s'|s, a)$.
4) R is defined as a reward function $R(s, a)$, where R represents the reward for taking action a in state s.
5) The discount factor γ is used to favor a more immediate reward in the next state versus a more distant reward several states in the future.

We make similar considerations of fault detection and depth stated by Cayeux et al. [9], but we must consider fault detection techniques available to us for this hardware. We can define our state as a combination:

$$s = (C_T, C_W, R, P, B_{skip}, D) \tag{3}$$

- C_T is the normalized torque change score computed by the ESST change detection algorithm,
- C_W is the normalized weight on bit change score computed by the ESST change detection algorithm,
- R is rate of penetration from avionics equiptment,
- P is a binary variable stating active percussion,
- B_{skip} states if skipped percussive beat was detected,
- D is depth in meters.

Using fault detection techniques like ESST [11], we can monitor the Torque and Weight on Bit change score, which often spike to indicate faulty conditions for C_T and C_W. We can also consider R, the Rate of Penetration (ROP), to monitor stalled or slowed progress approaching a velocity of zero. P indicates hard or differing material composition with active percussion. If our percussion is active, we then also begin monitoring B_{skip}. B_{skip} represents training from different beat detection methods. These begin with calculating a novelty curve corresponding to an increase in spectral energy of the vibration signal, not different from change point detection methods [10]. Then, an algorithm is used to pick which points in the novelty curve correspond to beats. Beat bounds are found from the local minima of the novelty curve. After applying the Fourier transform to beats, features like frequency peaks, and overall energy are used to determine if the beat denotes a fault or nominal operation. Analysis of test data shows that non-faulting beats typically have a frequency peak around 6 kHz, and faulting beats tend to have higher frequency peaks and more overall energy. We can use this to help detect potentially faulty beats.

In order to define our transition functions, we need to understand how different key variables change under different drilling depths up to a meter. Since we have extensive time-series fieldwork data for lunar and martian analog drilling conditions in 10 cm increments, we can accumulate statistics on hours of past drilling data for a Gaussian distribution of non-fault prone drilling data and compute the mean and standard deviation of those data points to come up with the necessary transition functions for incremental depth ranges.

B. Solving the MDP

To solve this MDP, we will use a reinforcement learning framework. Our goal is to find an optimal policy

$$\pi^* = \arg\min_{\pi} \mathbb{E}\left[\Delta t_{TD}(A) \mid \pi\right] \tag{4}$$

By minimizing Δt_{TD}, we minimize the total amount of actions as well as possible faults. We parameterize a stochastic policy $\pi(a \mid s; \theta)$ using a neural network that outputs the probability of selecting each action given the current state s. The reinforcement learning agent interacts with its environment where at each timestep it selects an action $a \sim \pi(a \mid s; \theta)$, receives a reward $R(s, a)$, which combines a term for depth progress with a time and fault-penalty term and is defined as

$$R(s,a) = w_1\,\Delta D - w_2 \left[t_a(a) + \sum_{f \in F} P(f \mid s,a)\, t_m(f) \right] \tag{5}$$

where ΔD is the incremental depth of the drilled, and w_1 and w_2 are the weights for the reward. In space drilling, where faults are extremely costly, w_2 should be set relatively high to ensure that any action that increases fault risk is heavily penalized and the time penalty incurred from a potential drilling event is t_m.

We aim to maximize the total drilled depth by accumulating incremental depths ΔD, while minimizing the total time spent, represented by action durations $t_a(a)$. Additionally, $P(f \mid s, a)$ is the probability of encountering a fault f, e.g., choking, binding, over-torque, etc. We will build on Cayeux et al. and use the Gauss-Seidel Value Iteration method [9], to converge to the optimal value function:

$$V(s) \leftarrow \max_a \left[R(s,a) + \gamma \sum_{s'} P(s' \mid s,a) V(s') \right], \tag{6}$$

where $V(s)$ is the value function for state s, $R(s,a)$ is the immediate reward received by taking action a in state s, $\gamma \in [0,1]$ is the discount factor for future rewards, $P(s' \mid s,a)$ is the probability of transitioning to state s' given the current state s and action a. We only have to keep one copy of state values in memory instead of two because the values are updated in place. When variables within the current state veer into uncertainty, we use the optiml policy π^* to determine the best course of actions for the current state.

VII. Conclusions and Future Work

Future work will include the implementation and testing of our MDP on the TRIDENT drill during in-lab testing in April and May of 2025. After tuning our solution, we will proceed with testing our lab-perfected methods during analog field testing in Alaska in September 2025.

Acknowledgments

This work is supported by the Minnesota Space Grant Consortium, NASA PSTAR Mars Exploring through Analog-site Drilling (MEAD), and National Science Foundation CSGrad4US Fellowship under Grant No. 2240197.

References

[1] S.T.M.D., “Stakeholder webinar: Civil space shortfall ranking,” 2024.
[2] B. Glass, D. Bergman, V. Parro, L. Kobayashi, C. Stoker, R. Quinn, A. Davila, P. Willis, W. Brinckerhoff, K. Warren-Rhodes, M. Wilhelm, L. Caceres, J. DiRuggiero, K. Zacny, M. Moreno-Paz, A. Dave, S. Seitz, A. Grubisic, M. Castillo, R. Bonaccorsi, and the ARADS Team, “The Atacama Rover Astrobiology Drilling Studies (ARADS) Project,” *Astrobiology*, vol. 23, pp. 1245–1258, Dec. 2023.
[3] Astrobiology, “TRIDENT on rover,” astrobiology.com/wp-content/uploads/2024/04/The-Atacama-Rover-Astrobiology-Drilling.png.
[4] N. Muscettola, P. P. Nayak, B. Pell, and B. C. Williams, “Remote agent: to boldly go where no AI system has gone before,” *Artificial Intelligence*, vol. 103, no. 1–2, pp. 5–47, Aug. 1998. [Online]. Available: https://doi.org/10.1016/S0004-3702(98)00068-X
[5] G. Paulsen, K. Zacny, P. Chu, E. Mumm, K. Davis, S. Frader-Thompson, K. Petrich, D. Glaser, P. Bartlett, H. Cannon, and B. Glass, “Robotic drill systems for planetary exploration,” in *Space 2006*.
[6] B. Glass, C.Fortuin, H. Battah, and I. King, “TRIDENT Drill Validation Testing in Haughton Crater, Devon Island, Canada,” ASCE Earth and Space 2024, 2024.
[7] J. Wardt, E. Cayeux, R. Mihai, J. Macpherson, P. Annaiyappa, and D. Pirovolou, “Taxonomy describing levels of autonomous drilling systems: Incorporating complexity, uncertainty, sparse data, with human interaction,” in *IADC/SPE International Drilling Conference and Exhibition*, 2024.
[8] T. Fong, J. Frank, J. Badger, I. Nesnas, and M. Feary, “Autonomous systems taxonomy,” *Autonomous Systems CLT Meeting*, 2018.
[9] E. Cayeux, B. Daireaux, A. Ambrus, R. Mihai, and L. Carlsen, “Autonomous decision-making while drilling,” *Energies*, vol. 14, p. 969, 2021.
[10] S. Boelter, E. Temesgen, B. Glass, and M. Gini, “Understanding drill data for autonomous application,” in *International Workshop on Autonomous Agents and Multi-Agent Systems for Space Applications at AAMAS*, 2024.
[11] L. Weber, “”changepoynt”,” 2024. [Online]. Available: https://github.com/Lucew/changepoynt
[12] T. Idé and K. Inoue, “Knowledge discovery from heterogeneous dynamic systems using change-point correlations,” in *Proceedings of SIAM International Conference on Data Mining*, 2005, pp. 571–575.
[13] S. Statham, “Autonomous structural health monitoring technique for interplanetary drilling applications using laser doppler velocimeters,” Ph.D. dissertation, Georgia Institute of Technology, 2011.