

Battery State-of-Health Aware Path Planning for a Mars Rover

Mariana Salinas-Camus¹, Chetan S. Kulkarni², and Marcos E. Orchard³

^{1,3} *Department of Electrical Engineering, University of Chile, Av. Tupper 2007, Santiago, Chile*
mariana.i.salinascamus@nasa.gov
morchard@u.uchile.cl

² *KBR Inc., NASA Ames Research Center, Moffet Field 94035 CA, USA*
chetan.s.kulkarni@nasa.gov

ABSTRACT

A rover mission consists of visiting waypoints to gather scientific samples based on set requirements. However, rovers face operational uncertainties during the mission, affecting the performance of its electrical and mechanical components and overall mission success. Hence, it is critical to have a decision-making framework that is aware of the health state of the components when planning the path of the vehicle. In particular, battery degradation, and consequently the battery State of Health (SOH), can affect the optimality of decisions made by the autonomous system in the long term. This paper presents a decision-making system that incorporates information on the energy drawn from the battery (based on the vehicle's velocity), terrain conditions, and model-based prognostic modules to assess the impact on the battery's state of charge (SoC). The decision-making system was formulated as a Markov Decision Process (MDP) to reach the goal destination by sending commands in a determined amount of time while maintaining the battery SoC within the policy stated. The MDP problem was programmed using the open-source framework POMDPs.jl, which has a variety of online and offline solvers. To solve the MDP problem online, we used Monte Carlo Tree Search (MCTS). Results from simulations demonstrate the effect that battery degradation and charging plans have on decision-making.

1. INTRODUCTION

A Mars rover's mission consists of visiting several waypoints. Due to the hostility of the environment and communications delays experienced between Earth and Mars, an autonomous navigation system becomes critical to ensure the success of the mission (Ellery, 2016). One approach that has gained attention in the last years is prognostic decision-

making (PDM), which is used to reduce the risk of failure during a vehicle's mission since it takes into account the conditions of the environment and the performance of the components (Balaban, Alonso, & Goebel, 2012). This approach has been used in studies for autonomous decision-making for planetary rovers (Narasimhan et al., 2012) and other types of vehicles, such as unmanned aerial vehicles (UAVs) for path planning (Quiñones-Grueiro, Biswas, Ahmed, Darrah, & Kulkarni, 2021).

Decision-making system takes into account the health of the components. In particular, for a vehicle that uses batteries as energy storage devices, it becomes crucial to predict the remaining useful life (RUL) and the current State of Charge (SoC) of the batteries to make reliable decisions during the vehicle's operation (Rezvanizanian, Huang, Chuan, & Lee, 2012) (Hogge et al., 2018). Mars rovers have been using Li-ion batteries since the launch of *Spirit* and *Opportunity* in 2003, and the use of this technology has contributed to the long operation life that the explorers have had (Smart et al., 2018). However, Li-ion batteries degrade over time, resulting in capacity loss and power fade. Degradation occurs because of various factors such as extreme temperatures, high/low SoC, and high charge and discharge rate (Han et al., 2019).

The State-of-Health (SOH) of a battery describes the relationship between its current capacity and its nominal capacity. Since we can expect that batteries that operate in a Mars rover degrade throughout a mission, we must consider the SOH when performing path planning. Although there is literature on modeling Li-ion battery degradation with different approaches, such as (Nascimento, Corbetta, Kulkarni, & Viana, 2021) and (Bolun Xu & Kirschen, 2018), there is a gap in the literature on how degradation affects a decision-making system for path planning.

In this paper, our objective is to develop a decision-making system that is aware of the SOH of the battery when planning

Mariana Salinas-Camus et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

the path of the rover, while also taking into account the terrain characterization. We formulated the problem as an MDP since it has been used in similar problems in the literature. The MDP was programmed using the framework POMDPs.jl (Egorov et al., 2017), which is a framework written in Julia that is intended for sequential decision-making under uncertainty. This framework has a variety of functions and solvers for POMDP and MDP problems. In particular, since we want to solve an online MDP problem, we use Monte Carlo Tree Search (MCTS).

The paper is organized as follows. Section 2 provides background material on the battery model used and how the terrain is characterized. Section 3 describes the problem formulation as an MDP. Section 4 shows the results of the system, along with a discussion. Section 5 concludes the paper.

2. BACKGROUND

2.1. Battery Model

In this work, electrochemistry-based battery model proposed in (Daigle & Kulkarni, 2013) for Li-ion cells is being implemented. This model uses only ordinary differential equations and is fast enough for real-time use. We consider the model to predict the End-of-Discharge (EOD) of the battery, which is the moment when the battery voltage is below the voltage cutoff. Equations 1, 2 and 3 define the state vector, input vector, and output vector of the battery model, respectively.

$$x(t) = [q_{s,p}, q_{b,p}, q_{b,n}, q_{s,n}, V'_o, V'_{\eta,p}, V'_{\eta,n}]^T \quad (1)$$

$$u(t) = [i_{app}] \quad (2)$$

$$y(y) = [V] \quad (3)$$

The variables q represent the amount of charge in the electrodes. The subscripts p and n represent the positive and negative sides of the electrode, respectively. The electrodes are split into two volumes; therefore, subscript s is the surface layer, and subscript b is the bulk layer. The states V' represent different voltage drops in the battery. As for the input vector, this corresponds to the applied electric current, where a positive value means that energy is being drawn from the battery, and a negative value means that the battery is being charged. The output vector is the battery voltage, which has an End-of-Discharge (EOD) voltage equal to 3.0 [V].

2.2. Terrain Characterization

Terrain identification is critical when developing an autonomous navigation system so the robot can operate safely and minimize the risk of mission failure. For example, *Spirit* took a path with deformable soil that had not been detected in orbital or terrestrial images. In addition, *Curiosity* experienced mobility difficulties when traversing a terrain of sharp rocks that damaged its wheels. The uncertainty about the ter-

Table 1. Speed used for each terrain in Decision-Maker System.

Terrain	Speed $\frac{m}{hr}$
A	62.33
B	50.01
C	26.62
D	13.57
E	8.35
F	Untraversable

rain's condition makes it difficult to plan the route that the rover should take.

In 2016, NASA JPL developed a new software called Soil Property and Object Classification (SPOC) that is based on Deep Learning and can identify the different types of terrain on the Martian surface using computer vision (Rothrock et al., 2016). This model was trained using the high-resolution images of HiRISE (High-Resolution Imaging Science Experiment). In particular, the images of Columbia Hills — the candidate landing site for mission M2020 — were used to evaluate the performance of SPOC. This site has a size of 40x40 [km²], which corresponds to 25 gigapixels. On Columbia Hills, they identified eight terrain types.

To define a set of feasible and optimal paths, let us consider that the exploration site is divided as a grid, where each cell belongs to a terrain category. Each category considers the properties of the terrain; therefore, a cell can be traversed at a certain estimated velocity. This estimation of the velocity at which the rover can move is calculated by combining maps of Cumulative Fractional Area (CFA), which corresponds to the fraction of the area covered in rocks, the Digital Elevation Model (DEM), and the traversability of the terrain. For purposes of this paper, we use the speed for each category shown in Table 1.

2.3. Monte Carlo Tree Search

There are different methods for online planning, such as forward search, heuristic search, and branch and bound, among others. In this work, we use MCTS since it has been used effectively in various problems, from board games to optimal planning for robots. The advantages of this method are that it avoids exponential complexity and does not require any prior knowledge of the domain, aside from having a policy and termination states (Świechowski, Godlewski, Sawicki, & Mańdziuk, 2022).

The MCTS algorithm finds an optimal solution by using Monte Carlo simulations to calculate the expected value of a potential trajectory in a search tree (Kocsis & Szepesvári, 2006). MCTS is based on four stages: selection, expansion, simulation, and backpropagation.

During the selection stage, the algorithm starts at the root to

the path that leads to the most promising leaf node. The most promising leaf node is selected if it returns the highest value of equation 4, where x_i is the mean value of the node, N is the number of visits for the parent node, and n_i is the number of visits for node i .

$$S_i = x_i + c \frac{\ln(N)}{n_i} \quad (4)$$

In the expansion stage, a node is added to the previously selected node. Later, in the simulation stage, random actions are taken from the last added node until a termination state is reached. Finally, after multiple simulations, during the back-propagation stage, the accumulated rewards calculated in the simulations are updated to the node values.

3. PROBLEM FORMULATION

A rover's mission consists of visiting various waypoints to perform science tasks. A set \mathcal{W} of M waypoints w_i is defined as $\{(x_i, y_i), r_i\}_{i=1}^M$, where (x_i, y_i) is the cartesian coordinate of the waypoint and r_i is the reward associated with the visit to the waypoint i (Balaban et al., 2020).

Consider a map grid \mathcal{M} with length m and width n . The set \mathcal{A} contains all the arcs that connect each point, i.e $\mathcal{A} = (i, j) : \forall i, j \in \mathcal{M}, i \neq j$. The set \mathcal{W} contains M point coordinates (a, b) that correspond to waypoints; each waypoint has an associated reward $r_{a,b}$.

Each arc (i, j) belongs to a certain terrain category, as it was previously defined in Table 1. This category gives information about the velocity with which the vehicle can move and the amount of energy required to do so. Therefore, the arc (i, j) defines the rover's travel time $t_{i,j}$ and the energy $E_{i,j}$ that the battery must provide.

The energy $E_{i,j}$ used for moving the rover through the arc (i, j) depends on the terrain category and the velocity. The rover has a peak energy consumption of 500 [Wh] while driving and a minimum of 150 [Wh]. Therefore, energy consumption is based on the terrain, with rougher terrain resulting in more energy consumed. The range of 150 to 500 [Wh] is evenly divided into 5 terrain categories (category F is not considered a feasible path). Then, a ratio is calculated between the actual velocity at which the rover moves and the nominal velocity for that terrain category. Finally, the energy $E_{i,j}$ is the result of the linear equation $E_{i,j} = x \cdot m + b$, where x is the ratio previously mentioned, m is the slope of the curve, and b is the offset, which is given by each category.

Certain operating factors cause faster degradation of the battery. Thus, a SoC policy is included so the battery operates in a range that preserves its functionality and health and reduces its degradation. The policy is defined by the range $[\alpha^-, \alpha^+]$, where $\alpha^-, \alpha^+ \in [0, 1]$ such that $\alpha^+ > \alpha^-$.

3.1. Markov Decision Process Model

Since the problem of going from a start point to a goal destination is a sequential decision problem, we can formulate it as an MDP (Kochenderfer & Wheeler, 2022) (Puterman, 1990).

In an MDP model, the agent takes an action a_t at time t given the current state s_t . The next state is given by the state transition model $T(s'|s, a)$, which represents the probability of transition from the current state s to state s' when taking action a . The agent receives a reward once an action is taken, which is given by the reward function $R(s, a)$. The reward function returns the expected reward when taking action a in a state s . The rewards accumulated over time are called *utility*. When the horizon of the problem is infinite, the utility is defined by equation 5 where γ is the discount factor. The discount factor affects how much weight future rewards have in the expected utility. For this problem, we chose a discount factor equal to 1.

$$\sum_{t=1}^{\infty} \gamma^{t-1} r_t \quad (5)$$

Lastly, an MDP model has a policy that tells what action to take, given the history of actions and states. Therefore, an optimal policy maximizes the expected utility.

For this problem, we have modeled the MDP as follows:

- **States:** position (x, y) , energy drawn from the battery, battery state (voltage, SoC, temperature, capacity), and time.
- **Actions:** 1) charge, 2) drive-left, 3) drive-right, and 4) drive-straight. Drive actions last for 5 minutes until the decision-making system has to take another action. Charge action lasts for the time taken to fully charge the battery from its current state of charge.
- **Rewards:** The rewards are based on the proximity of the rover to the goal destination. Also, there is a penalty when the battery is near the discharge threshold, and the charging action is not taken. Therefore, we based the penalty either on the SoC or the voltage value.
- **Termination conditions:** 1) maximum operation time was surpassed, 2) SoC is below the threshold, and 3) the Euclidean distance from (x, y) to the goal point is less than a given threshold.

4. RESULTS

To test the decision-making system with different levels of battery degradation, an assumption is made wherein the internal resistance increases by 1%, and the maximum capacity decreases by 0.2% in each cycle. Figure 1 shows discharge curves in four cycles when discharging the battery with a current of 2 [A].

For path planning experiments, we use a map grid of a




Figure 1. Constant-loading discharge curves for battery in different stages of degradation.

Table 2. Probabilities of each terrain for terrain map.

p_a	p_b	p_c	p_d	p_e	p_f
0.3	0.17	0.15	0.15	0.2	0.03

100x100 [m²] where each grid cell size is 1x1 [m²]. The terrain category in each cell was assigned randomly by following the probabilities shown in Table 2. Since the position x, y (which is a state of the MDP) is continuous, while the terrain map is discrete, we take the result of the floor function of the position to find in which cell the rover is. A SoC is introduced to promote the charge action when the battery is near a certain discharge threshold. However, when the battery is severely degraded, beyond cycle 70, SoC is not a good indicator of how discharged the battery is. That is why we also tested the MDP problem using the voltage to indicate when to charge the battery. Finally, as mentioned before, the framework used to program the MDP is called POMDPs.jl where we used MCTS as a solver.

The simulation tests were performed from the start point (3.0, 1.0) to goal points $g_1 = (6.5, 5.0)$ and $g_2 = (10.0, 18.0)$, and we first tested a SoC policy in the range [0.4, 1.0], with different levels of degradation. When testing the system with g_2 more actions are needed and in medium levels of degradation (between cycle 50 and 80). The SoC policy is not always fulfilled to reach the destination point in fewer actions. Although this is reflected negatively in the reward, the rover can reach the goal destination in less operation time.

Figures 2, 3 and 4 show results when the battery is not degraded and the destination is g_2 . The solution consists of 23 actions (or steps), and it takes 11623 [s] for the rover to complete the path. Figure 2 is a visualization of the path that the rover takes, and the green star represents the goal destination. Each dot represents when an action is taken. Remember that driving action lasts for 5 minutes, so the distance that the

rover can traverse during that time depends on the terrain category. Therefore, the overlapping dots in the map represent when the rover is charging the battery or when it is on rough terrain, meaning that the speed is too slow. As the map shows, the rover does not reach the exact destination goal. However, this is because the termination state is reached when the rover's position is within a distance less than 2 [m] from the goal position, which is the case for this example.




Figure 2. Path without battery degradation.

The system allows one to keep track of the variables of the battery, such as voltage, capacity, and temperature, among others. Figures 3 and 4 show the SoC and the battery voltage, respectively. The SoC policy is fulfilled during the entire mission, and therefore, the voltage never reaches EOD.




Figure 3. SoC throughout the mission without battery degradation.

Figures 5, 6 and 7 show the results when testing the system to reach position g_2 with the battery in cycle 50 of degradation. As expected, the solution took more actions when compared to the experiments with less degradation and, there-




Figure 4. Voltage throughout the mission without battery degradation.

fore, the operation took longer to complete. Figure 5 shows the path that the rover took. Notice around the coordinates (3.0, 3.0) that more actions were taken, making the rover go right and left on multiple occasions. This situation also occurs in the previous case (see Figure 2), yet it is more evident here. One factor contributing to this behavior is that around that area, there is a particular rough terrain, which may force the decision-maker system to take several actions to get out of that zone.

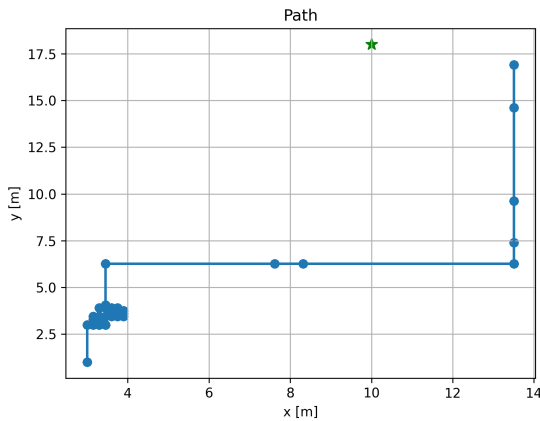


Figure 5. Path with battery in cycle 50.

Figures 6 and 7 show SoC and voltage, respectively. When comparing to the previous case, here we can see that the decision-making system takes the action of charging the battery more often and even charging the battery when the SoC is not near the lower bound of the policy. This behavior may be because the battery is more degraded. Degradation indicates that fewer actions are needed to decrease the SoC value significantly. As a result, the decision-making system takes a more conservative approach to avoid reaching EOD.




Figure 6. SoC throughout the mission with battery in cycle 50.




Figure 7. Voltage throughout the mission with battery in cycle 50.

One issue arises when the battery is severely degraded (after cycle 80). The MDP reaches the voltage termination state (voltage below 3.0 [V]) after three or four driving actions, thus not reaching the goal point. Since the SoC is high for a low voltage, the MDP does not have an incentive to take the charge action because this incentive (or penalization for not taking said action) is based on the SoC; therefore, a reward rule based solely on SoC is insufficient when considering acute levels of degradation. Table 3 illustrates an experiment with the battery in cycle 90 of discharge.

5. CONCLUSION AND FUTURE WORK

In this paper, an approach to include electrochemistry-based battery model degradation in a decision-making system for Mars rovers is being developed and presented. The problem is formulated as an MDP and takes into account terrain characterization. The MDP problem is solved by using MCTS, which is integrated into the framework POMDPs.jl. We in-

Table 3. Example of MDP terminated by End-of-Discharge at Cycle # 90.

Action	Voltage	SoC
drive-right	4.18	1.0
drive-right	3.15	0.87
drive-left	3.06	0.63

roduced battery aging under the assumption of reducing the capacity and increasing the internal resistance of the battery in each cycle of discharge, considering up to 100 levels of degradation.

We tested the framework under different levels of degradation and stated a SoC policy. The results show how degradation influences the ability of the decision-making system to succeed in a given mission and the impact that it has on battery charging plans. Also, from the simulations, we can conclude that using only the SoC as an indicator of when to perform the charging is not sufficient when the battery is severely degraded because the voltage value becomes a more dominant variable.

The system can keep track of the amount of energy that is drawn from the battery. However, in the current implementation, fixed levels of degradation are being used when testing the system. Future work includes a module to calculate the battery's capacity and internal resistance value based on the amount of energy drawn. This will help develop a more robust online decision-making system. Aside from that, an essential part of future work is to validate the results of this theoretical framework by testing a robot in different terrain conditions.

ACKNOWLEDGMENT

This work has been partially supported by FONDECYT Chile Grant Nr. 1210031, and the Advanced Center for Electrical and Electronic Engineering, AC3E, Basal Project FB0008, ANID.

The work by Chetan Kulkarni was supported by the System-Wide Safety (SWS) project under the Airspace Operations and Safety Program within the NASA Aeronautics Research Mission Directorate (ARMD) and under NASA Ames Research Center, Contract No. 80ARC020D0010.

REFERENCES

Balaban, E., Alonso, J., & Goebel, K. (2012). An approach to prognostic decision making in the aerospace domain..
 Balaban, E., Narasimhan, S., Daigle, M. J., Roychoudhury, I., Sweet, A., Bond, C., ... Gorospe, G. (2020). Development of a mobile robot test platform and methods for validation of prognostics-enabled decision making algorithms. *International Journal of Prognostics and*

Health Management, 4(1).
 Bolun Xu, A. U. G. A., Alexandre Oudalov, & Kirschen, D. S. (2018). Modeling of lithium-ion battery degradation for cell life assessment. *IEEE*, 9(2).
 Daigle, M., & Kulkarni, C. (2013). Electrochemistry-based battery modeling for prognostics..
 Egorov, M., Sunberg, Z., Balaban, E., Wheeler, T., Gupta, J., & Kochenderfer, M. (2017). Pomdps.jl: A framework for sequential decision making under uncertainty. In A. Honkela (Ed.), .
 Ellery, A. (2016). *Planetary rovers*. Springer Berlin Heidelberg.
 Han, X., Lu, L., Zheng, Y., Feng, X., Li, Z., Li, J., & Ouyang, M. (2019). A review on the key issues of the lithium ion battery degradation among the whole life cycle. *eTransportation*.
 Hogge, E., Bole, B., Vazquez, S., Kulkarni, C., Strom, T., Hill, B., ... Quach, C. (2018). Verification of prognostic algorithms to predict remaining flying time for electric unmanned vehicles..
 Kochenderfer, M. J., & Wheeler, T. A. (2022). *Algorithms for decision making*. London, England: MIT Press.
 Kocsis, L., & Szepesvári, C. (2006). Bandit based monte-carlo planning. In *Lecture notes in computer science* (p. 282-293). Springer Berlin Heidelberg.
 Narasimhan, S., Balaban, E., Daigle, M., Roychoudhury, I., Sweet, A., Celaya, J., & Goebel, K. (2012, January). Autonomous decision making for planetary rovers using diagnostic and prognostic information. *IFAC Proceedings Volumes*, 45(20), 289–294.
 Nascimento, R. G., Corbetta, M., Kulkarni, C. S., & Viana, F. A. C. (2021). Li-ion battery aging with hybrid physics-informed neural networks and fleet-wide data. *Annual Conference of the Prognostics and Health Management Society 2021*.
 Puterman, M. (1990). Markov decision processes..
 Quiñones-Grueiro, M., Biswas, G., Ahmed, I., Darrach, T., & Kulkarni, C. (2021). Online decision making and path planning framework for safe operation of unmanned aerial vehicles in urban scenarios. *International Journal of Prognostics and Health Management*.
 Rezvanizanian, S., Huang, Y., Chuan, J., & Lee, J. (2012). A mobility performance assessment on plug-in ev battery..
 Rothrock, B., Kennedy, R., Cunningham, C., Papon, J., Heverly, M., & Ono, M. (2016). SPOC: Deep learning-based terrain classification for mars rover missions. In *AIAA SPACE 2016*. American Institute of Aeronautics and Astronautics.
 Smart, M. C., Ratnakumar, B. V., Ewell, R. C., Surampudi, S., Puglia, F. J., & Gitzendanner, R. (2018). The use of lithium-ion batteries for jpl's mars missions. *Electrochimica Acta*, 268, 27-40.
 Świechowski, M., Godlewski, K., Sawicki, B., & Mańdziuk,

J. (2022). Monte carlo tree search: a review of recent modifications and applications. *Artificial Intelligence Review*, 56(3), 2497–2562.

BIOGRAPHIES

Mariana Salinas-Camus is a Ph.D. candidate in Aerospace Engineering at TU Delft. She received her B.S and M.S degrees in Electrical Engineering from Universidad de Chile in 2020 and 2023, respectively. Her current research interests are health management systems and failure prognosis focused on battery management systems, decision-making for unmanned vehicles and development of physics-informed AI models.

Dr. Chetan S. Kulkarni is a staff researcher at the Prognostics Center of Excellence and the Diagnostics and Prognostics Group in the Intelligent Systems Division at NASA Ames Research Center. His current research interests are in Systems Diagnostics, Prognostics and Health Management. Specifically focused in developing physics-based models, prognostics of electronic systems, energy systems and exploration ground systems as well as hybrid systems.

He is KBR Technical Fellow and AIAA Associate Fellow. Associate Editor for IEEE, SAE, IJPHM Journals on topics

related to Prognostics and Systems Health Management. He has been Technical Program Committee co-chair at PHME18, PHM20-22. And co-chairs the Professional Development and Education Outreach subcommittee in the AIAA Intelligent Systems Technical Committee.

Dr. Marcos E. Orchard is Professor with the Department of Electrical Engineering at Universidad de Chile, Principal Investigator at the AC3E Center (Universidad Técnica Federico Santa María) and was part of the Intelligent Control Systems Laboratory at The Georgia Institute of Technology. His current research interest is the design, implementation and testing of real-time frameworks for fault diagnosis and failure prognosis, with applications to battery management systems, mining industry, and finance. His expertise includes statistical process monitoring, parametric/non-parametric modeling, and system identification. His research work at the Georgia Institute of Technology was the foundation of novel real-time fault diagnosis and failure prognosis approaches based on particle filtering algorithms. He received his Ph.D. and M.S. degrees from The Georgia Institute of Technology, Atlanta, GA, in 2005 and 2007, respectively. He received his B.S. degree (1999) and a Civil Industrial Engineering degree with Electrical Major (2001) from Catholic University of Chile. Dr. Orchard has published more than 100 papers in his areas of expertise.