Toward Risk-Informed Operation of Autonomous Vehicles to Increase Resilience in Unknown and Dangerous Environments

Conference Paper · August 2016

DOI: 10.1115/DETC2016-60002

Douglas Lee Van Bossuyt
Naval Postgraduate School

Proceedings of the ASME 2016 International Design Engineering Technical Conference
IDETC 2016
August 21-24, 2016, Charlotte, North Carolina, USA

IDETC2016-60002

TOWARD RISK-INFORMED OPERATION OF AUTONOMOUS VEHICLES TO INCREASE RESILIENCE IN UNKNOWN AND DANGEROUS ENVIRONMENTS

Zachary Mimlitz
Colorado School of Mines
Golden, Colorado 80401
Email: zmimlitz@mines.edu

Adam Short
Colorado School of Mines
Golden, Colorado 80401
Email: ashort@mines.edu

Douglas L. Van Bossuyt*
Colorado School of Mines
Golden, Colorado 80401
Email: dvanboss@mines.edu

ABSTRACT

Operation of autonomous and semi-autonomous systems in hostile and expensive-to-access environments requires great care and a risk-informed operating mentality to protect critical system assets. Space exploration missions, such as the Mars Exploration Rover systems Opportunity and Curiosity, are very costly and difficult to replace. These systems are operated in a very risk-averse manner to preserve the functionality of the systems. By constraining system operations to risk-averse activities, scientific mission goals cannot be achieved if they are deemed too risky. We present a quantifiable method that increases the lifetime efficiency of obtaining scientific goals via the implementation of the Goal-Oriented, Risk Attitude-Driven Reward Optimization (GORADRO) method and a case study conducted with simulated testing of the method. GORADRO relies upon local area information obtained by the system during operations and internal Prognostics and Health Management (PHM) information to determine system health and potential localized risks such as areas where a system may become trapped (e.g.: sand pits, overhangs, overly steep slopes, etc.) while attempting to access scientific mission objectives through using an adaptable operating risk attitude. The results of our simulations and hardware validation using GORADRO show a large increase in the lifetime performance of autonomous rovers in a variety of environments, terrains, and situations given a sufficiently tuned set of risk attitude parameters. Through designing a GORADRO behavioral risk attitude set of parameters, it is possible to increase system resilience in unknown and dangerous environments encountered in space exploration and other similarly hazardous environments.

*Address all correspondence to this author.

NOMENCLATURE

α Reward Potential: Output of the GORADRO scoring algorithm (Eqn. 1)

ζ_n Risk Attitude Parameter: Optimization parameter from the GORADRO scoring algorithm (Eqn. 1) used to control the behavior of the model

V Immediate Value: Used in the GORADRO scoring algorithm

ρ Value Density: Used in the GORADRO scoring algorithm

Ψ Lifetime Returns: The performance indicator for trials representing the predicted scientific returns of the rover over its lifespan (Eqn. 2)

$\mathscr{E}$ Scientific Efficiency: The rate at which the rover obtains scientific targets (Eqn. 3)

L Lifespan: how the long the rover is operational (Eqn. 6)

ΔS Accumulated Science: The number of scientific targets the rover reached

Δt Elapsed Time: The amount of time the rover has operated

λ Failure Rate: the statistical rate at which the rover breaks down due to terrain hazards (Eqn. 4)

h The discrete hazard value for a given point on the map

Π Average Failure Rate: the time average failure rate (λ) of the rover

F Probability of Failure: The probability that the rover will fail after a given period of time (Eqn. 5)

INTRODUCTION

In August of 2014, the Mars Science Laboratory and its operators at the National Aeronautics and Space Administration (NASA)'s Jet Propulsion Laboratory (JPL) faced a difficult decision. Preliminary drilling at a site called Bonanza King had shown the potential for significant discoveries [1], but the rover was still a long way from its mission objective at Mount Sharp and in a state of uncertain, though clearly deteriorating, health. The decision was made by mission controllers and engineers to leave Bonanza King and the discoveries that may have waited there, and continue on to Mount Sharp [2]. But what if the results gleamed from further drilling at Bonanza King would have surpassed the rover's findings at Mount Sharp? What if the rover was at a healthy enough state to accomplish both missions? Situations like this and the questions that arise from them are common in deep space missions and, as we explore farther into space, they will only become more prevalent.

One possible solution to this problem comes from psychology where the concept of the risk attitude has been thoroughly studied in human behavior [3]. A risk attitude is defined as how willing people are to accept risks. By applying these same ideas to autonomous vehicles, we can regulate autonomous decision-making process and, through that, control autonomous system behavior. Designing risk attitude-based control algorithms and the behavior models to fit them requires large amounts of data and optimization. This paper details the process for optimization of these behaviors through a case study on a simulated rover, based on physical hardware.

The model presented in this paper was developed to empower rovers and similar exploration vehicles on extraterrestrial worlds to autonomously make critical mission decisions towards the purpose of maximizing the scientific returns of the mission without the need for constant human intervention. This is necessary as we explore deeper into space where research platforms will encounter more situations that need to be addressed in short order without waiting for Earth-based controllers to intervene in order to keep the mission going. The average communication delay between operators on Earth and scientific platforms on Mars is 20 minutes [4], and with new missions going even deeper into space, to targets such as Titan [5] or Europa [6], the communications delay will make the current practices used on Mars rovers of operator guidance and intervention unfeasible [4].

To solve this problem, the Goal-Oriented, Risk Attitude-Driven Reward Optimization (GORADRO) model for autonomous behavior presented in this paper uses a set of parameters representing a risk attitude to drive decision-making. Using this process, the GORADRO algorithm can weigh the relative impact of terrain-based hazards and potential scientific targets both in the immediate future and in the long term, in order to maximize long term mission returns without requiring supervision from Earth-based operators. This level of autonomy will be critical to future deep space missions.

BACKGROUND

The GORADRO method builds upon several topics from the fields of risk analysis, Prognostics and Health Management (PHM), decision theory, and autonomous mobile platform controls. Previous risk attitude informed path finding methods have only considered a single risk attitude parameter. The GORADRO method builds upon this foundation by considering a more complex risk attitude in relation to hazards, scientific returns, and topographical features.

Prior Work

GORADRO was developed to fill the need for a lightweight, yet versatile algorithm capable of making autonomous decisions which are not only risk-informed, but also consider the reward of all possible actions. The simple nature of this method allows it to be run on a simple processor with very low energy consumption or heat dissipation. Prior studies have been conducted to test the validity of GORADRO against other operational models [7] and to qualify the relationship between terms in the model with respect to overall performance [8].

Risk Analysis

Risk analysis serves as the basis for risk-informed design and decision making. Many techniques exist and are being developed in the fields of PHM and Decision-Based Design (DBD) in order to make risk-informed decisions. The process of decision making within PHM is known as Prognostic-Enabled Decision Making (PDM). Several methods including Failure Flow Decision Function (FFDF) [9], Active Mission Success Estimation (AMSE) [10], Uncoupled Failure Flow State Reasoning (UFFSR) [11], and Failure Flow Identification and Propagation (FFIP) [12] use PHM or DBD concepts to identify risks in the system's – in this case the rover's – health and make informed decisions to extend the system's lifespan while maximizing its efficiency.

FFIP models how failure can propagate through a system leading to system failure. FFIP builds upon the structure of a Function Basis for Engineering Design (FBED) [13], which is a method of representing a complex system as a series of functional blocks representing functionality in the system and flow lines representing the passage of energy, matter, and information through the system. Failure in FFIP initiates at one function and then propagates along flow lines to other functions potentially leading to system failure. An example of a FFIP failure path in component terms (functional terms) for a planetary rover is a solar panel (collect energy) failing, leading to a battery draining (store energy), leading to loss of power (electrical energy flow), leading to loss of the entire rover.

UFFSR [11] builds upon the foundation of FFIP, but accounts for failures that do not follow flow lines and instead propagate through physical space. An example of a UFFSR failure path using a Unmanned Aerial Vehicle (UAV) is a battery (store energy function) overheating and combusting and then fire (energy flow) spreading throughout the craft and damaging the CPU (process function) leading to loss of the entire UAV.

FFDF [9] is a method based on PDM that analyzes the way failure propagates through a system along FFIP and UFFSR paths, and attempts to make control decisions that force the failure down the paths that have the lowest probability of causing critical system failure. An example of FFDF would be the decision made by a person when falling on whether or not to reach out and catch themselves, potential injuring their wrist, but stopping their head from hitting the ground. AMSE [10] builds on the FFDF method by putting it into a mission structure and analyzing a decision's effect on total mission success as the mission progresses.

Risk Attitudes

The concept of a risk attitude is taken from psychology where it is used to quantify a person's willingness to take risks [3, 14]. Risk in this sense is broken down into three components: chance, uncertainty, and reward. Chance is defined as the effect of uncertainty on goals [15] where uncertainty is the potential for multiple outcomes given a single event with poorly defined probabilities [16]. Finally, reward is taken as the worth or value of the outcome [16].

Between these three concepts, we can construct a robust model for risk-informed decision making and use it guide the rovers through a risk-inherent environment. Risk attitudes can be placed on a spectrum between total risk aversion and total risk tolerance. A risk averse attitude is an attitude that is less likely to accept increase reward in exchange for increased risk. Using a person crossing the street as an example, a risk averse attitude may manifest as the person standing at the cross walk and waiting for the appropriate light and walking signal before crossing the street. A risk tolerant attitude is an attitude that is more likely to accept increased risk in exchange for increased reward. An example of a risk tolerant behavior using the street crossing example is a person who is more likely to jaywalk in the interest of saving time.

Prior experiments with the Risk-Attitude Informed Route planning (RAIR) method have shown that rovers in such an environment cannot be entirely risk-adverse while continuing to complete mission objectives [17]. GORADRO uses a spread of risk attitudes in order to balance accepting risks and preserving rover health while continuing to accomplish mission objectives. The results of this paper show how the model's drive to reach scientific targets is balanced against its aversion to hazardous terrains in order to maximize mission returns.

Path Finding

Path Finding utilizing heuristic evaluation of terrain has been performed computationally since the 1960s [18] and is used currently to solve problems within logistics, infrastructure, aeronautics, navigation, and robotics [19]. Existing robotics applications of path finding include warehouse item management, security patrolling, tour guiding, and exploration. The RAIR method is an existing method that attempts to quantify an appropriate risk attitude for an environment while path finding in a hazardous environment. The GORADRO method attempts to better address the problem of risk attitudes for path finding in a hazardous environments by breaking down the risk attitude into multiple parameters representing specific classes of hazards and rewards. This multivariate approach allows for the study and generation of more complex risk attitudes and reduction of risk of failure faced by the system.

METHODOLOGY

The GORADRO model for autonomous behavior uses a set of parameters which represent a risk attitude to weigh the relative cost and benefit of terrain hazards and scientific rewards in order to determine the best course of action throughout the life of the mission. Using the GORADRO algorithm the rover or other autonomous system can easily pathfind through risk-inherent terrain, such as that found on Mars, the Moon, and many other areas of interest for autonomous scientific missions.

Development

GORADRO was developed to fill the need for an autonomous behavior that considered not only the risk of failure from decision, but also the potential reward from that decision. This trade off of risk and reward is crucial to making fully informed decisions with regard to mission success. Utility theory provides a framework in which rover health and scientific return can be compared through the GORADRO model in order to find a balance which allows for better performing rover missions on extraterrestrial worlds.

Behavioral Model

The GORADRO model uses a weighted parameter model to compare the relative benefits and risks of scientific targets and hazards in the vicinity of the rover [7]. The algorithm, shown in Eqn. 1, is used to calculate the reward potential (α) of a specific point of interest. The point of interest with the highest reward potential is chosen as the next destination as discussed below.

$$\alpha = \zeta_1 \cdot V_{\text{science}} - \zeta_2 \cdot V_{\text{hazard}} + \zeta_3 \cdot \rho_{\text{science}} - \zeta_4 \cdot \rho_{\text{hazard}} \qquad (1)$$

The other variables which appear in this equation are V, the immediate value of either the scientific target or the hazard level of the terrain at the point of interest; and ρ which is the average density of scientific targets or terrain hazards in the general area about the point of interest. Finally, the set of $\zeta_1 - \zeta_4$ in this equation constitute the risk attitude.

Risk Attitude Breakdown. Each of the four terms and their respective risk attitudes has a specific purpose and address both an environmental and a time concern of the analysis. ζ_1 and ζ_3 look at the reward from an action while ζ_2 and ζ_4 address the risk of an action. Similarly, ζ_1 and ζ_2 are applied to values from the immediate spacial area and so near future while ζ_3 and ζ_4 act on measurements of areas in order to predict long term future conditions. Other environmental or internal state factors could be added into the model for operations in unique environments but these four where found to be the most robust basis for decision making [8].

Autonomous Operations

If every possible point in the environment was evaluated for the reward potential (α) at that point and compiled, that set would constitute a Surface of Reward Potential (SoRP) showing the reward potential for any arbitrary point. On this SoRP, we would expect scientific targets to generally appear as local maximums and hazardous regions to generally appear as local minimums. In order for GORADRO to follow the most efficient path between points of maximum reward potential, it would merely need to follow the gradient of the SoRP.

However, generating this surface and performing advanced mathematical operations on it is unreasonably complex for a deployed rover with limited computational resources. As such, in order to find a path through the environment in which the autonomous system is exploring, the GORADRO model regularly calculates the reward potential of several discrete, evenly spaced points equidistant from the rover as shown in Figure 1, and chooses the point with the highest reward potential as the next destination. This process is analogous following the gradient of the SoRP and while there is an error associated with the resolution at which the model is sampling surrounding area, it can be effectively reduced by increasing the number of sampling points [7].

Input to Algorithm

The values for V_{science}, V_{hazard}, ρ_{science}, and ρ_{hazard} in Eqn. 1 come from information about the environment the model is operating in. Since these parameters need to be numerical representations of physical concepts, the information acquired by sensors on the physical rover needs to be pre-processed by risk analysis techniques in order to obtain these numerical representations. These types of algorithms are very well established and not discussed in this paper which will continue to focus on how the GORADRO algorithm employees the results of these analyses in order to make decisions.

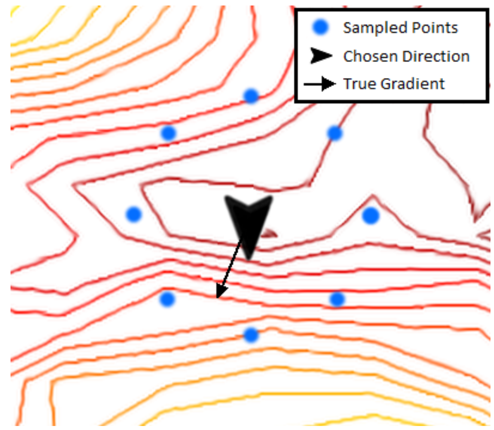

FIGURE 1. A REPRESENTATION OF GORADRO'S SAMPLING PROCEDURE TO FIND THE POINT OF GREATEST REWARD POTENTIAL. THE TOPOGRAPHICAL LINES REPRESENT THE SURFACE OF REWARD POTENTIAL (SoRP). THE SECOND ARROW SHOWS HOW THE MODEL'S SELECTION COMPARES TO THE TRUE GRADIENT OF THE SoRP.

Parameter Tuning

The risk attitude of the model represents the preference of the model for scientific targets over terrain hazards or vise-versa and quantitatively controls the choices made by the behavior. It is this risk attitude that we seek to optimize in this analysis in order to show the ability of the GORADRO model to increase the lifetime scientific returns of a rover in various hazardous environments.

This optimization is performed through a case study which utilizes a high resolution parameter sweep. The trials with best overall performance are marked as the best trials and further analysis or investigation can be conducted from there. The process for determining the performance of a rover in a trial is given below.

Implementation

In order to successfully implement the GORADRO algorithm for control of an autonomous vehicle, several factors must be addressed before the algorithm can fully function. The first factor is noise, which is addressed through the use of time averaging. Small changes in the rover's location can have large and inconsistent effects on the reward potential, so the potentials for each are averaged over the last several measurements before a decision is made. The second is repetition. In order to prevent the rover from returning to the same favorable point repeatedly, a list of all visited targets is kept, and they are excluded from future analysis. Another factor is tied potentials; however, such ties are unlikely given the volatile nature of the inputs and averaging processes, and statically insignificant under the recalculation cycle time and so are merely handled by the processing order of options. The final concern that must be address is potential wells. In a potential well, the SoRP has a geometry such that when at point A it is most favorable to travel to point B, and at point B it is most favorable to travel to point A. In order to prevent the rover from being trapped in such a state, a separate process from the GORADRO algorithm checks at periodic intervals that the rover has traveled a significant distance from the the point is was at during the last check. If the rover has not, it is assumed to be stuck and forced to drive outside the region without regard to risk or rewards. As this clearly posses a threat to the rover, the time interval and drive distance are carefully tuned to disrupt rover operations as mush a possible.

Determining Performance

The overall performance of the GORADRO model is determined using Eqn. 2.

$$\Psi = \mathscr{E} \cdot L \tag{2}$$

Where Ψ is defined as the total gain or score of scientific return, $\mathscr{E}$ as efficiency of reaching targets, and L as the estimated lifespan of the rover. The processes for determining $\mathscr{E}$ and L from the data are described below.

Determining the Scientific Efficiency. The scientific efficiency of the rover is determined using the simple rate equation shown in Eqn. 3. Where ΔS is the science gained during the mission/trial and Δt is the elapsed time of the mission/trail.

$$\mathscr{E} = \frac{\Delta S}{\Delta t} \tag{3}$$

Determining the Lifespan. A statistical failure distribution is used to find the average lifespan (L) of a rover behavior based on the average terrain stress from hazards it encounters during the simulation [17]. This distribution is driven by the failure rate due to terrain hazards given in Eqn. 4, where λ is the failure rate for a discrete terrain hazard value of h. This equation is normalized for a 90 day mission under average terrain stress [4, 8].

$$\lambda(h) = 1.583 \cdot \mathrm{e}^{-\left(\frac{h-11.03}{4.511}\right)^2 - 0.00399} \tag{4}$$

The chance of the rover failing after an elapsed time, t, is built as F in Eqn. 5 and based on the failure rate where Π is the time average hazard rate (λ).

$$F(t) = 1 - \mathrm{e}^{-\Pi t} \tag{5}$$

Finally, the average lifespan of a rover for the purposes of this analysis is taken to be when the chance of failure is equal to 95%.

$$F(L) = 0.95 \tag{6}$$

While deployed rovers may fail long before this point in time, this analysis assumes that hazards contribute to wear on the rover without risking immediate failure on exposure such that the buildup of wear will follow statistical averages.

In Summary

GORADRO represents a new approach to making autonomous decisions based on both the potential rewards and inherent risk for any action. This methodology allows for greater autonomy in scientific rover platforms deployed in risk-inherent terrain by providing a framework for comparing the relative benefits and risks of scientific targets and terrain-based hazards using the concept of a risk attitude.

CASE STUDY

In order to quantify the effectiveness of risk attitudes in controlling autonomous behavior, we developed a case study to test a wide range of risk attitudes. Using the simulated environment described below, we examined the performance of a rover operating using the GORADRO model with 4,096 unique risk attitudes each on two maps in search of trends in the response of the model across both maps. Through this study we found that there was a cluster of risk attitudes which consistently performed above average.

Simulation Model Environment

The simulator used in this case study is the Simulated Physics and Environment for Autonomous Risk Studies (SPEARS) developed by the Van Bossuyt Research Group at the Colorado School of Mines. The SPEARS simulator allows rovers with custom-written autonomous behavioral models to be placed in artificially-generated landscapes and allowed to explore under a variety of conditions. While it is possible to simulate communication with the rovers, the trials presented in this paper have the rovers explore an unknown environment completely autonomously.

In order to provide mission conditions for the simulated rovers, SPEARS generates both scientific targets and terrain hazards for the rovers to seek out and avoid. The trials developed as

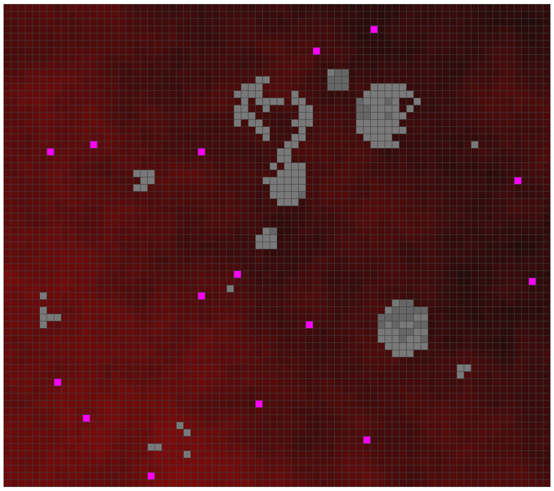

FIGURE 2. A SCREEN CAPTURE OF THE SPEARS SIMULATOR SHOWING A MAP. IN THIS IMAGE, HEIGHT IS SHOWN AS THE COLOR GRADIENT FROM BLACK TO RED, SCIENTIFIC TARGETS ARE SHOWN IN MAGENTA, AND THE DARK GRAY SECTIONS ARE SEVERE TERRAIN HAZARDS. THE GRIDLINES ARE 1 METER SPACING.

part of this research used discrete, binary targets and a continuous range of hazards of varying magnitudes. It is also possible to use targets of varying magnitudes and discrete regions of binary hazards for the purposes of experimentation. However, for clarity in this paper we have limited the case study to binary targets and a continuous range of hazards. A screen capture of SPEARS displaying a region of one of the maps used in the trials is shown in Figure 2.

Experimental Setup

Targets and Hazards. SPEARS provides scientific targets and terrain hazards to the autonomous behavior as binary or integer values which are well-suited for the algorithm to process and respond to. Again, this simplification to integers masks pre-processing that would come from processing external inputs and sensor data through risk analysis techniques in a deployed system. For the purposes of testing the autonomous behaviors, these processes and their inaccuracies are currently ignored so that the focus of this paper remains on the performance of the autonomous GORADRO algorithm.

Parameter Choices. In order to identify trends in the correlation between GORADRO's risk attitude and its lifetime scientific gain, and in order to find the ideal behavioral parameters, a parameter sweep was conducted over ζ_2, ζ_3, and ζ_4 with ζ_1 held fixed. ζ_1 is held fix to avoid analyzing non-unique attitudes. Based on the linear nature of the algorithm, the attitude $\{2,6,8,4\}$ will represent the same behavior as the attitude $\{1,3,4,2\}$ due to even scaling not effecting the comparative weighting. The ranges used in the parameter sweep of the other parameters were taken from prior work in the optimization of this model [8] and are shown in Table 1. This parameter sweep provides 4,096 data points over which GORADRO's performance can be analyzed.

TABLE 1. VALUES FOR THE PARAMETER SWEEP ACROSS THE RISK ATTITUDE

Parameter	Low Value	High Value	Step
ζ_1	1000	1000	0
ζ_2	15	30	1
ζ_3	50	90	2.666
ζ_4	0	1	0.066

Table 1 shows that the magnitude of the parameters being examined differed greatly. It is important to note that not all of this difference comes from the influence of that parameter on the model. For instance, it is inaccurate to say that ζ_1 is 10 to 20 times more important than ζ_3 because the scale of their coupled terms is also different. The V_{science} term coupled with ζ_1 in Eqn. 1 will never be more than 1, while the V_{hazard} term with ζ_3 may be any number from 0 to 10, and so the difference in magnitudes of the input parameters is inversely shown in the magnitude of the risk attitude parameters.

Other Variables. In order to account for the diversity in the environments GORADRO may face on deployment, the parameter sweep is performed on two different maps. Terrain elevation, magnitudes of surface hazards, and spacing between scientific targets can all dramatically affect GORADRO's performance for a given risk attitude. Since we are interested in investigating GORADRO's performance in different environments with similar conditions, each of the two maps has the same average density of scientific targets and hazards, and follow similar topography profiles.

RESULTS AND DISCUSSION

Experimental Results

The results from the trials from both maps show clustering of high performing risk attitudes, shown in Figure 3 with blue circles. These two bubble charts show the estimated lifetime scientific return, Ψ from Eqn. 2, with the better performing trials appearing red and larger while the lower performing trials appear small and yellow-green. The results are scaled non-linearly to accent the grouping.

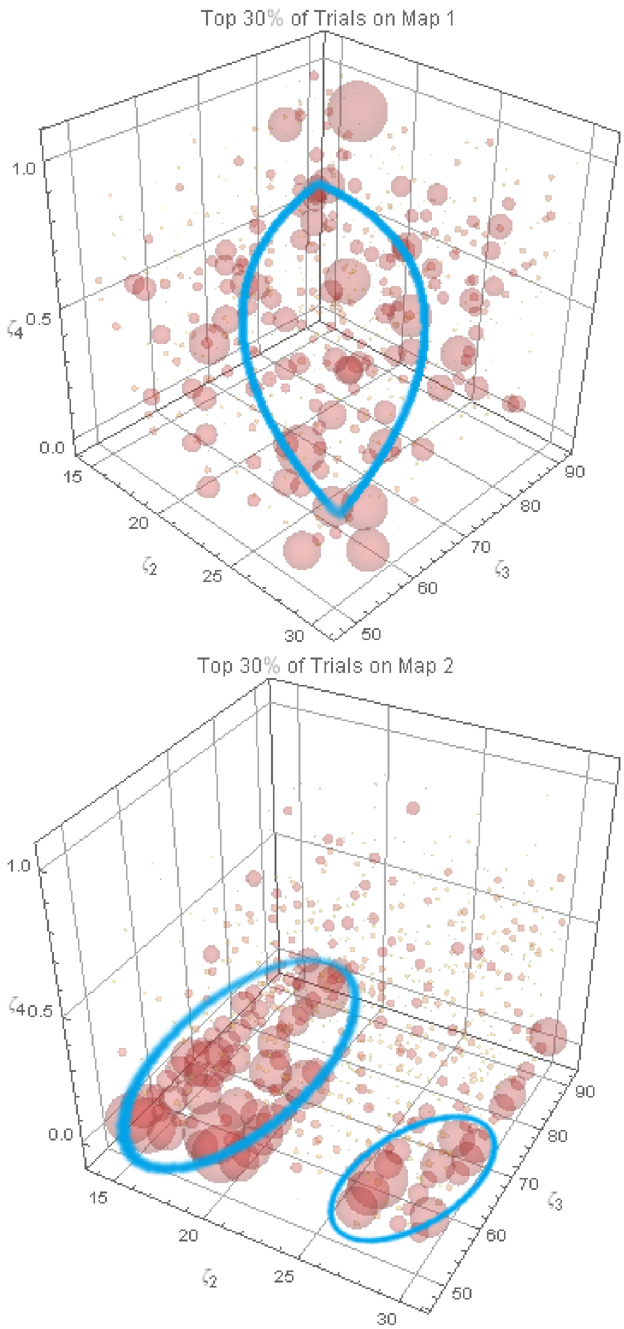

FIGURE 3. THREE-DIMENSIONAL PLOTS OF THE TOP 30% OF TRIAL RESULTS FROM BOTH MAPS SHOWING THE PERFORMANCE OF TRIALS, Ψ. THE BETTER TRIALS ARE SHOWN AS BEING LARGER AND REDDER.

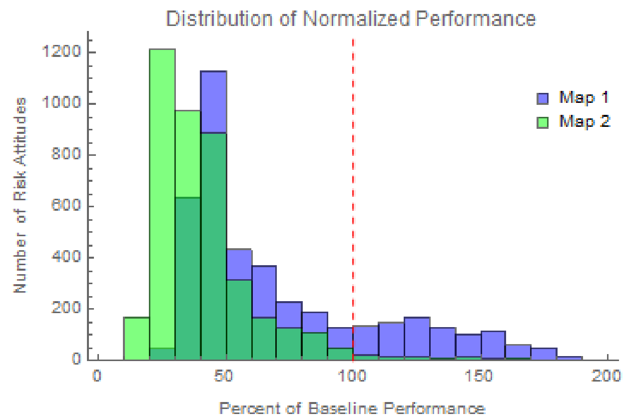

FIGURE 4. THE SPREAD OF THE PERFORMANCE OF ALL RISK ATTITUDES ACROSS BOTH MAPS, NORMALIZED TO THE BASELINE PERFORMANCE.

Analysis of Data

In order to demonstrate the increase in lifetime efficiency from the GORADRO method, a baseline was calculated for each map by simulating a rover which drove from target to target, always choosing the closest target as the next destination without any of the environmental considerations used by GORADRO. Using this baseline, the results of each map were scaled to find the performance of each risk attitude as a percentage of the baseline. A histogram showing the distribution of the performance for all risk attitudes across both maps is shown in Figure 4.

While many of the risk attitudes performed below the baseline (with the greatest concentration around 50%), a considerable number of attitudes on both maps performed well above the baseline (Map 2 does have over 90 risk attitudes above this line), increasing the effectiveness of the rover by nearly 100% in some cases.

However, as can be seen in the difference of the distributions of the bubble charts in Figure 3, the high performing attitudes across the two maps are very distinct with little to no overlay. This difference in the underlying nature of high performing risk attitudes speaks to a high environmental dependence in the optimal values of the the risk attitude that should be explored further.

Another notable trend can be seen in the two-dimensional cross section of the data. Figure 5 shows a range of cross sections of the values from the trials on the first map (which are shown in Figure 3). As can clearly be seen, the data forms ridges of higher performance consistent across all the cross sections of the data.

Interpreting Results

In order to more clearly visualize trends in the data, the 100 highest scoring trials from each map were mapped onto parallel

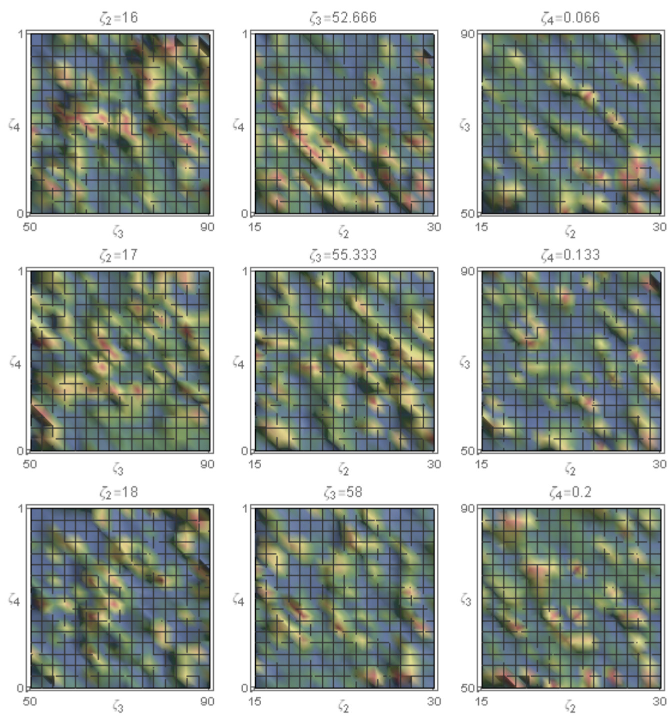

FIGURE 5. A SMALL SAMPLING OF CROSS SECTIONS OF THE TRIAL RESULTS FROM MAP 1 WHICH SHOW A CLEAR PATTERN OF RIDGES OF HIGHER PERFORMANCE.

axis plots which are shown in Figure 6[1]. In the plot for map 1 we see a clear trend of high ζ_2 and low ζ_3 values, thought the values for ζ_4 are far more varied. In the plot for map 2, there are two clusters of values for ζ_2 and a very strong trend to low values of ζ_4.

Effect of Environmental Differences

In order to generally explain the trends seen in the parallel axis plots (Figure 6), we look to differences in the environment of the individual maps. Each map was generated using the same parameters but evolved differently throughout the generation. Table 2 shows the differences in basic properties of each map's environment.

The higher average hazard level on the first map can be attributed to the low value trend in ζ_4 which accounts for the area hazard, not due to a lower importance of this factor, but a numerically larger input for ρ_{hazard}. This increase in the average map hazard also caused all of the rovers to see a significantly lower operational lifetime resulting in a lower baseline and making it harder for rovers to achieve higher performance levels. This increased difficulty resulted in the distinction between the large

[1] It is important to note that the color scheme for this plot is opposite of the other figures presented here for readability purposes.

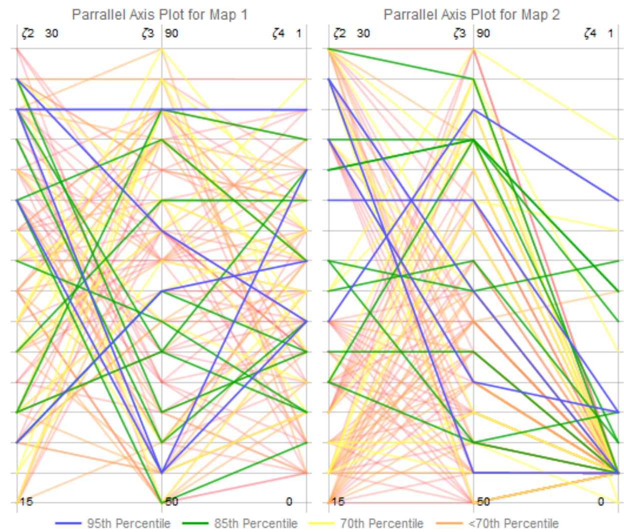

FIGURE 6. A PARALLEL AXIS PLOT SHOWING THE TRENDS BETWEEN THE HIGH PERFORMING RISK ATTITUDES ON BOTH MAPS.

TABLE 2. LISTING OF MAJOR ENVIRONMENTAL PROPERTIES FOR EACH OF THE MAPS USED IN THE CASE STUDY.

Property Name	Map 1 Value	Map 2 Value
Average Hazard:	0.903	0.825
Number of Targets:	747	750
Target Cluster Density:	0.125 trgts/m^2	0.108 trgts/m^2
Elevation Change:	2.047 m	1.214 m
Roughness:	0.014 m	0.012 m

number of high performing attitudes and dispersed grouping in the map 1 data, and the lower number of high performing attitudes and tight clustering of the map 2 data.

Similar correlations can be noted in the map 1 data. The lower hazard rate and better clustering of targets resulted in a map that was much easier to perform well, which allowed a greater range and number of risk attitudes to achieve performances above 100%. Likewise, this cluster could be attributed as the cause for the trend in low values of ζ_3 for the same reasons the map 2 trials saw low values of ζ_4.

The exact and quantitative relationship between environmental factors and the risk attitude parameters requires further and more detailed study. However, the results of this case study provide a good starting point for the advancement of that research. Furthermore, despite the difference in what risk attitudes performed well in each map, the results of this case study clearly show that the implementation of properly tuned risk attitudes to the behavior of autonomous rovers can dramatically improve their lifetime performance.

Influences of Design

The physical design of the rover plays a large part in determining appropriate risk attitudes. It is important to acknowledge that the robustness of the design and scientific equipment of the rover contribute heavily into the value calculations for the environmental hazards and scientific rewards. For instance, fragile vehicles will consider most terrain hazards to be of greater risk to the rover than a more robust vehicle would. Similarly, a point of scientific interest would only be considered a valid target for a rover that has the appropriate equipment to take the required measurements. The design decisions that control the final rover configuration play a critical, if largely silent, role in the GORADRO method, coming into play in the risk analyses which provide the input parameters for Eqn. 1.

CONCLUSIONS AND APPLICATIONS

From the case study and results presented above, it is clear that properly designed risk attitudes can dramatically affect the overall performance of an extraterrestrial rover. Properly tuned risk attitudes on both maps showed an increased performance of up to 200% of the baseline. If applied to deployed systems, GORADRO may convey that same increase in the efficiency of deep space mission without contributing to payload weight and while saving money on Earth-based operators and time on deep space communication channels. These kinds of savings and efficiencies are critical to future missions as we push farther into the reaches of space and explore even more distant worlds.

Future Work

Future research to be conducted on the GORADRO model includes detailed testing and quantifying how the impacts of environmental features, such as target and hazard density, impact the magnitudes of the risk attitude parameters as well as the related impacts on the uncertainty of values such as the terrain risk and location of scientific targets. A separate line of questioning to be explored is how risk attitudes evolve with time over the course of the mission lifetime, as failures early in the rover's life as a result of high risk decisions have a much greater impact on overall performance and mission goals than the same failure near the end of the mission. Finally, as a separate performance validation, the GORADRO method should be compared against existing risk mitigation methods such as RAIR which do not consider reward in decisions.

ACKNOWLEDGMENT

This research was partially supported by United States Nuclear Regulatory Commission Grant Number NRC-HQ-84-14-G-0047 and the Colorado School of Mines Undergraduate Research Fellowship Program. Any opinions or findings of this work are the responsibility of the authors, and do not necessarily reflect the views of the sponsors or collaborators. The authors wish to acknowledge the work of the undergraduate research assistants in the Van Bossuyt lab and specifically wish to thank the following students for their contributions: Alexis Humann, Robin Coleman, and David Hodge. The authors wish to thank LeVar Burton, Fred Rogers, Gene Roddenberry, Carl Sagan, and their individual middle school and high school science and technical arts teachers who inspired them to pursue careers in the sciences and engineering, and instilling in them a sense of purpose and compassion.

ACRONYMS

AMSE	Active Mission Success Estimation
DBD	Decision-Based Design
FBED	Function Basis for Engineering Design
FFDF	Failure Flow Decision Function
FFIP	Failure Flow Identification and Propagation
GORADRO	Goal-Oriented, Risk Attitude-Driven Reward Optimization
JPL	Jet Propulsion Laboratory
NASA	National Aeronautics and Space Administration
PDM	Prognostic-Enabled Decision Making
PHM	Prognostics and Health Management
RAIR	Risk-Attitude Informed Route planning
SoRP	Surface of Reward Potential
SPEARS	Simulated Physics and Environment for Autonomous Risk Studies
UAV	Unmanned Aerial Vehicle
UFFSR	Uncoupled Failure Flow State Reasoning

REFERENCES

[1] Grotzinger, J., Gupta, S., Malin, M., Rubin, D., Schieber, J., Siebach, K., Sumner, D., Stack, K., Vasavada, A., Arvidson, R., et al., 2015. "Deposition, exhumation, and paleoclimate of an ancient lake deposit, gale crater, mars". *Science,* ***350***(6257), p. aac7575.

[2] Nassivera, J., 2014. "Curiosity rover moves on to mount sharp after cancelled drilling operation". In HNGN.

[3] Weber, E. U., Blais, A.-R., and Betz, N. E., 2002. "A domain-specific risk-attitude scale: Measuring risk percep-

tions and risk behaviors". *Journal of behavioral decision making,* ***15***(4), pp. 263–290.
[4] Leger, P. C., Trebi-Ollennu, A., Wrigh, J. R., Maxwell, S. A., Bonitz, R. G., Biesiadecki, J. J., Hartman, F. R., Cooper, B. K., Baumgartner, E. T., and Mainmone, M. W., 2005. "Mars exploration rover surface operations: Driving spirit at gusev crater". *Jet Propulsion Laboratory*.
[5] Oleson, S., 2014. "Titan submarine: Exploring the depths of kraken". *NASA*.
[6] Bergin, C., 2015. "Nasa budget proposal to initiate europa mission plan". *NASA Space Flight*, feb.
[7] Mimlitz, Z., and Van Bossuyt, D., 2015. "Goal oriented, risk mitigating autonomous behavior for extraterrestrial rovers". In Colorado Space Grant Consortium Space Research Symposium.
[8] Mimlitz, Z., 2015. "Optimizing goal oriented autonomous risk attitudes for extraterrestrial rovers in an non-homogeneous environment". In American Institute of Aeronautics and Astronautics Region V Student Conference.
[9] Short, A. R., and Van Bossuyt, D. L., 2015. "Rerouteing failure flows using logic bocks in functional models for improved system robustness: failure flow decision functions". *ASME IDETC*.
[10] Short, A. R., and Van Bossuyt, D. L., 2015. "Active mission success estimation through phm-informed probabilistic modeling". *Annual Conference of the Prognostics and Health Management Society*.
[11] Van Bosuyt, D. L., 2015. "Modeling of function failure propagation across uncoupled systems". *Annual Reliability and Maintainability Symposium*.
[12] Jensen, D., Tumer, I. Y., and Kurtogula, T., 2009. "Flow state logic (fsl) for analysis of failure propagation in early design". *ASME IDETC and CIEC*, pp. 1033–1043.
[13] Hirtz, J., Stone, R. B., McAdams, D. A., Szykman, S., and Wood, K. L., 2002. "A functional basis for engineering design: reconciling and evolving previous efforts". *Research in engineering Design,* ***13***(2), pp. 65–82.
[14] Van Bossuyt, D., Dong, A., Tumer, I., and Carvalho, L., 2013. "On measuring engineering risk attitudes". *Journal of Mechanical Design*, dec.
[15] Standards Australia New Zealand, 2009. *Risk Management: Principles and Guidelines*. AS/NZS ISO 13000:2009.
[16] Hubbard, D., 2007. *How to Measure Anything: finding the Value of Intangibles in Business*. Jon Wiley & Sons Inc., Hoboken, NJ.
[17] Short, A., and Van Bossuyt, D., 2015. "Risk attitude informed route planning in a simulated planetary rover". *ASME International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*, aug.
[18] Lim, L. Y., 1968. *A pathfinding algorithm for a myopic robot*. Jet Propulsion Laboratory, California Institute of Technology.
[19] Liu, F., Narayanan, A., and Bai, Q., 2000. "Real-time systems".

