A Decision Support System for Extravehicular Operations Under Significant Communication Latency

Timothy McGrath[1]
JES Tech, Houston, TX 77058

Jason Norcross[2]
KBR, Houston, TX 77058

Jon Morris[3], Federico Piatti[4],
D2K Technologies, Oceanside, CA 92056

Fernando Figueroa[5]
NASA Stennis Space Center, John C. Stennis Space Center, MS 39529

and

Brianna Sparks[6], Jeffrey T. Somers[7]
NASA Johnson Space Center, Houston, TX 77058

Humanity hopes to perform extravehicular activities (EVAs) on the surface of Mars; however, several technical and operational challenges must first be overcome. Foremost among these challenges is managing a significant communication latency between Earth and Mars. Current and historical paradigms of EVA operations have required near-real-time communication between the crewmember(s) and Earth-based mission control. Next-generation operational paradigms for supporting deep space exploration will necessitate a distributed decision authority system, including delayed Earth-based mission control, the on-planet extravehicular crewmember(s), and intermediate mission support from intravehicular (IV) crewmember(s) within real-time communication range. This latter group is of particular interest: they must provide operations support without the plentiful resources available to mission control on Earth. Thus, NASA is developing the Personalized EVA Informatics and Decision Support (PersEIDS) software platform. PersEIDS is designed to bolster operator situational awareness and offload operator workload by automatically tracking and projecting consumables usage over an EVA timeline, providing real-time probabilistic safety assessments and recommending alternative EVA timeline(s) when the active timeline is not expected to be completed under consumables limits. The PersEIDS concept of operations, use cases, and models will be presented. A limited version of PersEIDS was demonstrated during a three-day-long study where each day a roughly four-hour-long simulated Martian EVA was performed in virtual reality at the NASA Johnson Space Center. The first day was a control trial without PersEIDS support; the second and third days represented different levels of decision support provided by PersEIDS to the IV crewmember acting as mission control. With PersEIDS support, the IV crewmember was able to manage the mission to completion faster

---

[1] Human Performance Engineer; Human Physiology, Performance, Protection & Operations (H-3PO) Laboratory, JSC/SK3
[2] Senior Scientist; H-3PO Laboratory, JSC/SK
[3] VP of Engineering; Autonomous Systems Laboratory, SSC
[4] Staff PHM Engineer; Autonomous Systems Laboratory, SSC
[5] Electrical Engineer, Autonomous Systems Laboratory, SSC
[6] Flight Operations Engineer, Flight Operations, JSC/CX3
[7] Co-Lead, H-3PO Laboratory, JSC/SK3

and with more remaining consumables; however, additional testing is required to understand confounding factors, e.g., training bias.

I. Introduction

FUTURE human spaceflight to Mars will present an incredible set of challenges. In addition to the difficulties of assembling and launching infrastructure to support human life on Mars, the Mars environment is notably inhospitable[1]. Humans will venture from the relative safety of established habitats to build further base infrastructure and accomplish science objectives. Historically, these so-called extravehicular activities have been tightly managed by an expert team of mission controllers on Earth. NASA's mission control center (MCC) employs dedicated subteams to manage different aspects of spaceflight: EVA operations, biomedical monitoring, science objectives, medical operations, habitat power, and communications, amongst others. Throughout the entire history of human spaceflight, these teams have managed missions in near real-time. During Lunar EVA, communication latency (one-way light time or OWLT) was approximately 1.25 seconds, enabling meaningful communication between the EVA crew and MCC[2]. However, the distance to Mars will extend this communication latency significantly. Depending on the relative positions of Earth and Mars in their respective orbits, the expected communication latency will be between 4 and 22 minutes OWLT[3]. While ground analogs (e.g. the BASALT research program[4–6]) have studied the technological and operational limitations that arise from Mars-like communication latencies, no program has yet studied biomedical operations for EVA under similar limitations. Ultimately, this Mars communication latency will preclude Earth-based mission control from making the types of time-sensitive, safety-critical decisions that have characterized all prior EVAs. On Mars, a small group of crewmembers within the established Martian habitat (referred to as intravehicular or *IV* crewmembers) must make timely decisions to ensure the safety of the extravehicular (EV) crewmembers exploring the surface. These decisions will be informed, in part, by the biomedical data of the EV crew. In this paper, the phrase *biomedical data* will be used to refer to the relevant physiological states of the human crewmember as well as the state of any suit systems which support the health, safety, and performance of the human operator.

The resource limitations and communication latency to Mars will force humans to adopt a new paradigm of operations to support productive science and habitat building. NASA has begun the process of studying these ideas and enumerating gaps. In 2020, NASA published a comprehensive concept of operations (ConOps) document in conjunction with future Earth-independent human space exploration. Due to the large communication latency, EVA operations will be directed by an IV crewmember in-situ on the surface of Mars[7]. The Martian crew will also use their more-detailed knowledge of the Martian surface to recommend EVA plans or operations suggestions to the Earth-based mission support team. However, mission control on Earth will retain decision authority over developing detailed EVA timelines with objectives. These detailed EVA timelines will then be communicated back to the IV crewmember on Mars before an EVA begins. Future EVA operations are organized into three discrete categories: (1) pioneering tasks to assemble base infrastructure, (2) maintenance tasks on infrastructure, and (3) science tasks such as geological sampling. The first two categories are within the historical expertise of operations-focused flight controllers. The last category will necessitate the development of a new subteam within MCC to advise on science matters. This subteam is currently in development and will be utilized during the Artemis program for Lunar EVAs. This delayed-advisory mission support paradigm from Earth will likely lead to new approaches to EVA operations: for example, perhaps science-oriented EVAs will originally be planned with an initial set of science objectives, with built-in intermediate tasks which act as communication breaks so that the Earth-based MCC may understand, confirm, and respond to the current status of any performed science objectives and advise new ones. A science EVA plan may have initial goals and then "stretch" goals, executed at the discretion of the EV/IV pair with recommendations from MCC.

Decision support systems (DSS) are any system which facilitate a decision-making process[8]. A DSS may support decisions in several ways. For example, they could automate the calculation of intermediate states, rank solutions, visualize states, etc. Regardless of its product definition, a DSS must be designed to fit the associated cognitive system that its human operators employ when performing work[9, 10]. During Martian EVA, the IV crewmember will be tasked with a monumental amount of cognitive burden previously performed by dedicated teams of numerous people on Earth. The broad number of critical decisions to be made, the time sensitivity of each, and the knowledge scope required to make them will inherently necessitate the use of DSS to facilitate Martian EVA. One set of decisions, among many the IV crewmember must perform, is the set which alter a planned EVA to ensure the safety of the EV crewmember. Adding to the complexity, biomedical data alone is insufficient to fully inform these decisions: there are other operational reasons that one might decide to alter an EVA. Therefore, a biomedical decision support platform alone will be insufficient to fully support decisions related to altering EVA timelines in-situ. However, as an early research and engineering problem, such a biomedical support platform can be a critical system to study, optimize, and train future EVA operations.

NASA is developing the *Personalized EVA Informatics and Decision Support* (PersEIDS) system to support the time-sensitive, safety-critical decisions that must be made by crew on the surface of Mars. PersEIDS operates as a companion software component of EVA mission scheduling and timeline management software, e.g., Maestro[11] or Playbook[12–14]. The software is designed to leverage observed biomedical data during previous missions and analog studies to project the biomedical state of the crewmember/suit system over any given EVA timeline. This DSS then performs inference about the projected biomedical state over a timeline against any user-input flight rules or constraints to generate a safety assessment of a given timeline. This same reasoning routine is performed simultaneously against potentially hundreds of EVA timelines, including the planned timeline and any feasible alternatives. The IV crewmember can then use these safety assessments of EVA timelines, in context of other operational constraints, to make decisions to alter the planned timeline to ensure the safety of the EV crewmember.

This paper details the initial considerations to develop the DSS, including the simplified model of the IV crewmember's relevant decision process (Section II). Section III will discuss the software architecture currently utilized by PersEIDS in ground analog demonstrations. Subsequent sections will detail initial demonstrations: their setup, results, and lessons learned. Finally, a discussion follows to put these results in context as well as describe the limitations of the current PersEIDS architecture and assumptions.

II. Simplified EVA Biomedical Monitoring Decision Process

The IV crewmember will be tasked with making many decisions simultaneously to ensure EVA safety. Traditionally, the monitoring of an EVA based on real-time biomedical data has relied heavily on the decades of learned experience of the mission support team. Teams of flight surgeons and biomedical engineers within the MCC are collectively able to reason against many flight constraints using dozens of data points. Some of these flight rules are more critical than others, and some are merely warnings. Some flight rules are grounded in peer-reviewed etiology, and some have arisen as observational heuristics. Adding to the complexity, the prioritized rule set is situation-dependent: certain operations or system states may cause certain biomedical constraints to increase or decrease in criticality. In short, the human cognitive processes that underlie EVA monitoring are complex and leverage operational contexts that may not be apparent to the naive software aid. To the authors' knowledge, no comprehensive model of the biomedical monitoring decision process has been published.

As an initial simplification of this complex cognitive process, PersEIDS has identified two general decisions which must be made continuously during the monitoring of an EVA to ensure the safety of the EV crewmember, and these decisions apply in all EVA operational contexts.

1. Is the current planned EVA safe if performed nominally?
2. Is the current planned EVA nearing a “point of no return?”

These decision trees and their subsequent actions are detailed in Fig. 1.

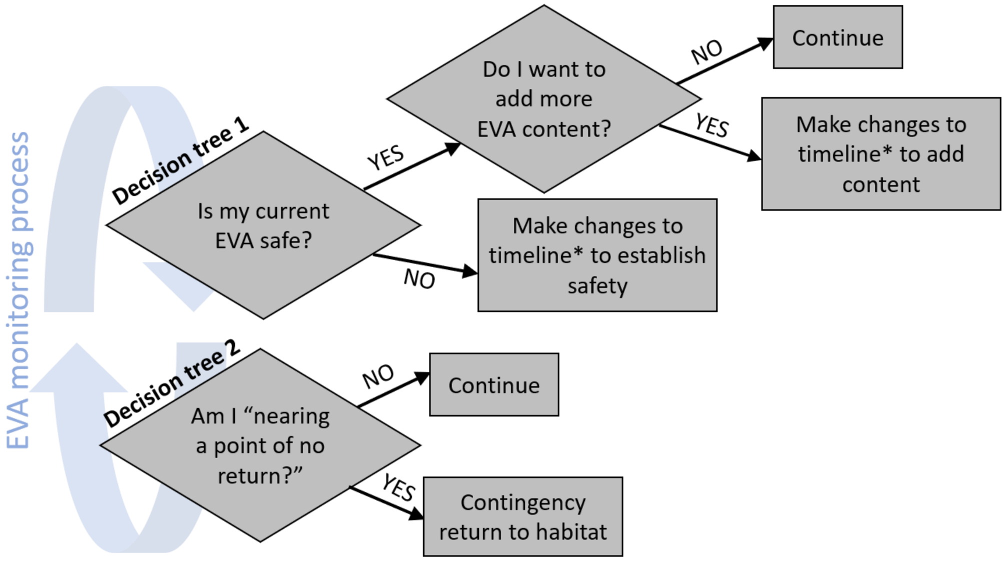

Figure 1. A simplified model of the EVA monitoring process as a decision tree. Making changes to the timeline means changing future tasks, including their ordering, existence, and/or expected duration.

The phrases *current*, *planned*, or *active* are used interchangeably to mean the timeline that an EV crewmember is currently executing during an EVA. The current EVA may have historical content—tasks which have already been performed and biomedical data aggregated during them. A timeline will typically also have future content—tasks which have not yet begun. As the EVA progresses, the IV crewmember must consider the current biomedical state of the EV crewmember and all future timeline content.

A. Decision Tree 1: Assessing the Safety of the Planned EVA

The planned EVA may be assessed as *safe* or *not safe*. A safe timeline, in this context, is nominally completed by the EV crewmember without violating any flight rules. It is desired that the current EVA is expected to be safe; however, many circumstances can cause this expectation to change. For example, an EVA may begin with a traverse task which is expected to be short, but in reality is longer and requires higher metabolic load to complete. After this task is completed, the crewmember will have dipped further into their consumables than expected. If they have dipped far enough into these consumables, the EVA will no longer be considered safe to complete as planned, and intervention will be necessary. From the perspective of this decision process, an intervention manifests as a change to the future content of the timeline—changing the order, existence, or duration of planned tasks in the EVA.

In order to monitor an EVA for safety, the IV crewmember must maintain an internal model of what constitutes a safe EVA. There are obvious final conditions which would signal an unsafe EVA; for example, running out of oxygen before the EVA is complete, but the IV crewmember cannot wait until the end of an EVA to decide if the EVA was unsafe; they are tasked with commanding an intervention well before the critical point of an unsafe EVA occurs. The IV crewmember considers the current biomedical state of the EV crew as well as all planned EVA content. The biomedical state must then be projected forward to the relevant temporal point in the EVA where the flight rule applies—for many, e.g., oxygen consumption, this temporal point is the end of the EVA. This projected biomedical state must then be compared against the relevant flight rule, and a Boolean decision (safe/not safe) is reasoned. In this paper, the process of determining the safety of an EVA is referred to as *safety assessment*.

As the EVA progresses, the uncertainty around this safety assessment decreases. At the beginning of an EVA, uncertainty is at its maximum: all EVA content is yet to be performed and many things can go wrong. At the end of an EVA, this uncertainty (conceptually) has asymptotically converged to zero: there is no future content, so the projection of the biomedical state is simply the current biomedical state at the end of the EVA.

The IV crewmember will constantly perform this safety assessment during an EVA. If the planned EVA is unsafe, they must make a change to the planned timeline to ensure safety. They will consider available options for making an EVA safe: removing or reordering future tasks or reducing the duration of these tasks. For any considered alternative plan, a safety assessment is also then performed; if an alternative plan is unsafe, it will not be considered. Through this iterative cognitive process, the IV crewmember will determine an alternative plan which reestablishes the expectation of safety while minimally satisfying productivity needs. Note that safety is more important than productivity: EVA content will be removed until an EVA is deemed safe. If this causes the entire future EVA to be removed, resulting in an immediate return to the habitat, it is an acceptable outcome to ensure the safety of the EV crewmember. When the IV crewmember determines the satisfactory alternative, they will switch the planned EVA to this alternative, including the set of tasks, their order, and expected durations of each task in the timeline.

On the contrary, if—when performing the safety assessment—the IV crewmember determines the current timeline is safe, they then perform a slightly different decision process. They may consider adding content to the planned EVA while maintaining the expectation of safety. This choice is strictly optional and is dependent on the type of EVA being performed. Some EVAs, such as those performing maintenance, may have goals strictly based on the completion of those goals: once the goal is complete, the EV crewmember should return to the habitat. Other types of EVAs, such as science EVAs, may have built-in operational directives to perform additional content if deemed safe (e.g., sample additional rocks if possible). If the IV crewmember chooses to not add any EVA content, the decision process restarts. If, however, additional EVA content is required, the IV crewmember must then consider their alternatives much like described above: perform a safety assessment of the alternative, determine its suitability to current safety and productivity needs, and switch if/when the appropriate alternative is found.

B. Decision Tree 2: Anticipating the “Point of No Return”

During Martian EVAs, a requirement will exist to, colloquially, never venture to the point where the EV crewmember can no longer return to the safety of the habitat within available constraints, e.g., consumables. Call this the *point of no return*. It is not a literal geographic point, but it can be. In the context of an EVA timeline, the point of no return can be defined by a combination of factors, including long elapsed durations of EVA activity, large distances travelled from the habitat, or infeasible return traverse paths due to terrain constraints. This point of no return will vary based on mission architecture[7]. In missions where a rover is available, this point will simply be the point at which a rover cannot rescue the crewmember. In EVAs where no rover is available, the EV crewmember will have to return to the habitat under their own locomotive abilities. There are additional nuances to the definition of the point of no return, e.g., in conjunction with mission architectures for rescuing an incapacitated crewmember. A full discussion of these nuances is beyond the scope of this paper. As a first approximation, PersEIDS currently defines the point of no return as the point in an EVA where the EV crewmember can no longer return to the habitat within available consumable constraints.

Just as the IV crewmember must monitor the expected safety of an EVA (decision tree 1), so must they monitor the risk of passing the point of no return. This is a different class of risk management with different consequences. Decision tree 1 monitors the expectation of performance in an EVA, whereas decision tree 2 manages the tail risk. If an EVA is considered unsafe in decision tree 1, there are (usually) many available interventions. The most extreme of which is the contingency return EVA, the specific operational paradigm where the EV crewmember is instructed to stop whatever they are doing and return to the habitat immediately. As the contingency return EVA is the final available intervention under standard operations, the risk management of the point of no return (decision tree 2) corresponds directly to the problem of monitoring if the contingency return EVA is feasible. If this contingency return is feasible, the point of no return has not been passed. The corresponding decision tree is quite simple: if the IV crewmember determines they are nearing the point of no return, they intervene to order a contingency return EVA. *Nearing*, in this sense, is an assessment that prescribes the IV crewmember's internal model of risk tolerance onto the EVA. If the IV crewmember is particularly risk-averse, they might feel they are nearing the point of no return earlier during an EVA.

While decision tree 2 is simple, there exists a high degree of cognitive difficulty in performing the point-of-no-return assessment. The feasibility of the contingency return must be projected over the entire future state of the EVA. This contingency return may be feasible during any given point in an EVA, but this does not mean that it will not become infeasible at a future point in the EVA. The uncertainty in making this assessment may be manifested as risk intolerance: if the IV crew feels they do not understand the uncertainty around a point-of-no-return assessment, they may choose to decrease their risk tolerance to manage the tail risk of a catastrophic event occurring.

III. PersEIDS Overview

PersEIDS, as a DSS, must support this decision process by estimating the safety of an EVA timeline based on integrated biomedical models which are used to reason against any relevant flight rules.

A. Embedded Biomedical Models, Associated Flight Rules, and Automated Safety Assessment

1. *Metabolic Consumables*

PersEIDS currently supports two types of flight rules to reason against when making safety assessments for an EVA timeline, each with a dedicated mathematical model. The first of the two is a constraint on maximum metabolic consumption. The oxygen supply for the suit is designed to support a specified amount of consumption before needing to be refilled. This amount of consumption is specified as a metabolic load in units of BTUs. As the crewmember works inside their suit, their metabolic rate (typically specified in units of BTU/hr) will vary based on the physical workload required to perform an EVA task. This metabolic rate is directly proportional to the consumption of oxygen from the atmosphere within the suit[15]. The harder the crewmember works, the higher their metabolic rate, and the faster they consume this oxygen supply.

Over decades of spaceflight and ground analog operations, NASA has collected data on crewmember metabolic rate as a function of the task they were performing. These data are contextualized according to many factors: the individual crewmember, environment (e.g., flight mission or ground analog), suit model, etc. A separate but tangential effort to the PersEIDS project is focused on the capture, organization, and modeling of these data. PersEIDS uses this database to build prior beliefs of metabolic rate over a given task for a given crewmember. These priors may assume any arbitrary distribution, but initial approximations in this work are approximated as Gaussian.

As an EVA progresses, metabolic rate measurements from the suit are taken and discretely integrated to track the actual metabolic consumption $\widetilde{M}_t$ at current time t over an EVA. PersEIDS performs a projection of metabolic consumption $\widehat{M}_T$ at EVA completion time T by integrating this current metabolic state with the prior belief of all yet-to-be-completed tasks:

$$\widehat{M}_T \sim \widetilde{M}_t + R_c\left(d_c - \tilde{d}_c\right) + \sum_{i=c+1}^{N-1} (R_i d_i) \quad \forall\, t < T, \left(d_c - \tilde{d}_c\right) \geq 0 \tag{1}$$

where N is the number of tasks, indexed as $i = 0,1,\ldots,N-1$ with current task index c. d is the expected duration to complete a task, and $\tilde{d}$ is the actual duration elapsed in any task. By definition, $\left(d_c - \tilde{d}_c\right)$ is always strictly non-negative: the expected time to complete the current task should always be greater than the current time elapsed in that task. R_i is the prior belief of metabolic rate that will be observed in task i. Eq. (1) relates that the projection of metabolic consumption at the end of a timeline $\widehat{M}_T$ is distributed as the sum of current metabolic consumption $\widetilde{M}_t$, expected metabolic consumption over the remainder of the current task $R_c\left(d_c - \tilde{d}_c\right)$, and the expected metabolic consumption of all future tasks. The metabolic rate prior R of each task is assumed to be independent.

The cumulative distribution function (CDF) of M_{proj} is then interpreted as the probability of completing the EVA with a total metabolic consumption less than a given amount. PersEIDS then assesses this CDF at the point of the maximum allowable metabolic consumption.

2. *Inspired partial pressure of carbon dioxide (PiCO$_2$)*

During an EVA, just as a crewmember inspires oxygen, they expire carbon dioxide. The buildup of carbon dioxide in the suit is a risk which must be monitored and mitigated. The symptoms of carbon dioxide poisoning at non-lethal levels are much like those of carbon monoxide poisoning: nausea, disorientation, etc.[16, 17]. Even in the case that the re-inspiration of expired carbon dioxide does not pose directly fatal consequences, the degradation of a crewmember's performance due to acute effects may lead to unsafe operational scenarios, and repeated exposure may have chronic adverse effects. Space suits are designed with systems to "wash out" the expired CO$_2$ from the suit; however, these systems have an effective maximum rate of CO$_2$ scrubbing from the suit. Additionally, complex fluid dynamics in partial gravity may lead to imperfect CO$_2$ washout, where some portion of expired CO$_2$ stays near the mouth and nose and is inevitably re-inspired.

NASA studies the washout capability of space suits. The carbon dioxide generated through exhalation is captured via experimental methods, then the *partial pressure of inspired CO$_2$* (PiCO$_2$ for short) is modeled as a function of metabolic rate to characterize the risk to astronauts[18]. In flight, where no face-mounted oronasal mask may be used, this model is applied to approximate the PiCO$_2$ level within the suit. The model varies per suit, dependent on factors such as the suit helmet geometry and the design of the CO$_2$ scrubbing system. The risk to the crewmember is subsequently managed by flight surgeons and biomedical engineers through the prescription of a maximum allowable time a crewmember may spend at any given PiCO$_2$ level. There are multiple of these flight rules tracked during EVA: crewmembers may spend longer times at lower PiCO2 levels. If the PiCO$_2$ level within the suit is high, a mission controller may suggest a crewmember to lower their metabolic rate (i.e., reduce workload). Within the construct of the $PiCO_2$ mathematical model, this reduces the expected level of $PiCO_2$ in the suit. Within the actual suit engineering system, this would allow the CO_2 scrubbing subsystem to "catch up," removing excess carbon dioxide from the suit atmosphere.

As an EVA progresses, the same metabolic rate measurements mentioned previously are transformed into instantaneous $PiCO_2$ approximations via the suit-specific $PiCO_2$ model. These measures are assumed to be independent as metabolic rate may vary wildly between consecutive samples. Each $PiCO_2$ estimate then contributes the entire inter-sample duration to any associated flight rule(s). For example, there may be two individual rules: one for a maximum time between ranges of 0.0 and 5.0 mmHg and one between 0.0 and 8.0 mmHg. If an instantaneous $PiCO_2$ sample of 6.5 mmHg is recorded, the duration since the last sample will be added to both timers corresponding to these flight rules.

Synthesizing the probabilistic estimate of each of these timers at the end of an EVA is a complex problem: the accumulation of samples within a specific $PiCO_2$ range is a process which, in general, does not follow a Gaussian distribution. The probability of any individual estimate contributing to a $PiCO_2$ timer may be modeled as a Bernoulli process: for each sample, there is some probability that a sample contributes to a $PiCO_2$ timer. Within a task, this probability is constant, based on the static metabolic rate prior. However, across tasks, this probability varies because the expected metabolic rate varies. Under a constant probability model, the process would follow a binomial distribution. With arbitrary probabilities, this model is better defined by the Poisson binomial distribution[19]. In general, evaluating this distribution is a computationally hard problem[20]; traditional approaches leverage discrete Fourier transforms to perform convolution of multiple binomial distributions. The discussion of this computation is beyond the scope of this paper. PersEIDS leverages publicly available software libraries for this purpose. PersEIDS computes the convolution of the current $PiCO_2$ timer state with the future probability of the timer state (defined by the Poisson binomial distribution) to generate the probabilistic estimate of the $PiCO_2$ timers at the end of an EVA. Similar to the metabolic consumption model/rule, the CDF of this distribution can be directly interpreted as the probability of completing the EVA within the maximum allowable time at a given $PiCO_2$ range.

3. *Automated Safety Assessments*

In order to support the safety assessments that must be performed by an IV crewmember, PersEIDS takes these probabilistic projections of crew state against flight rules and provides a synthesized safety metric. PersEIDS defines the safety of a timeline as the probability of completing the EVA without violating *any* flight rules. This probability is presented as a scalar between 0 and 100 percent. This automated safety assessment is performed for each timeline that PersEIDS is monitoring simultaneously, and all are displayed to the human operator of the PersEIDS system.

IV. Initial Demonstration

In September 2022, the PersEIDS software application was initially demonstrated during a three-day study in the APACHE virtual reality (VR) environment at NASA's Johnson Space Center. Each day, a roughly four-hour simulated Martian EVA was performed. A single subject acted as the EV crewmember on all three days. A single subject acted as the IV crewmember on each day; however, this subject was not the same across all three test days. Together, in virtual reality, the EV-IV crewmember pair performed Martian-like EVA operations and the EV crewmember (in virtual reality) performed an EVA with typical tasks: traverse, base maintenance, science payload deployments, geological sampling, etc. The IV crewmember (not in VR) was responsible for managing the EVA, including operational goals and EV crewmember safety. Six total consumables constraints were required to be managed by the IV: one constraint on total allowable metabolic consumption and five separate constraints on maximum allowable time in separate $PiCO_2$ bands. If any of these six constraints were violated, the EVA was considered failed. An EVA began when the EV crewmember exited the habitat and concluded when the EV crewmember returned to the habitat. The IV crewmember was not in virtual reality. Metabolic rate measurements were estimated at 1 Hz by performing a simple transformation of heart rate measurements measured by a wireless Polar H10 heart monitor (Polar Electro, Finland). These metabolic rate estimates were sent to the PersEIDS application in real-time to monitor consumables progression and provide projections when applicable (day three).

The APACHE environment, the location of this study, has a treadmill integrated into the VR environment so the EV crewmember can perform long traverse tasks without breaking VR. APACHE also has a sandbox where all non-traverse tasks are performed. This sandbox is meant to simulate the feel and difficulty of the Martian regolith. A composite photo of the APACHE physical environment, VR environment, and PersEIDS user interface overlay can be seen in Fig 2.

Figure 2. Initial PersEIDS demonstration in the APACHE environment.

The test conditions were different across all three days. On day one, a nominal EVA was performed without the help of PersEIDS. The EVA was easy to complete, and the IV crewmember had no decision support available. On day two, a "challenge" EVA was performed: during an initial long traverse task out to the science site, the resistance of the treadmill was increased, and the path to the site was longer than the subjects had anticipated. The IV crewmember still had no decision support available and was tasked with managing this EVA manually. On day three, a similar challenge EVA was performed (though the path and operational details were different), but this time the IV crewmember used PersEIDS as a decision support tool.

As a software demonstration, no measures of workload or attention were taken. Subjective feedback was taken from the EV crewmember and IV crewmember to understand the usability of the demonstrated software capabilities.

V. Preliminary Results

On day one, where the EVA was typical and had no off-nominal events, the EV/IV pair were able to complete all EVA goals and safely return to the habitat without violating any constraints.

On day two, the EV crewmember began the off-nominal traverse, where the traverse required a longer duration and a higher metabolic workload than anticipated. This higher metabolic workload resulted in higher $PiCO_2$ estimates accumulating in their associated allowable time bands. The IV crewmember, it is inferred, noticed the buildup of $PiCO_2$ in the virtual suit too late. As a result, a mitigation strategy was implemented for the rest of the EVA, limiting the maximum allowable metabolic rate of the EV crewmember so as to not spend additional time at high $PiCO_2$ values. Ultimately, this mitigation strategy significantly slowed down the pace of the EVA, and certain scientific objectives were subsequently omitted in order to return to the habitat without violating any constraints. In a strict sense, the EVA was completely safe; however, the failure to complete all scientific objectives is notable.

On day three, The EV/IV pair performed a different EVA, although the existence of a harder-than-expected traverse task remained the same. In this condition, the IV crewmember did have access to the PersEIDS end-of-timeline consumables projections. The EV/IV pair were able to complete all science and operational objectives effectively, without implementing a risk mitigation strategy, and returned to the habitat easily.

VI. Discussion

The overall success of the EVA without modification when using the PersEIDS system suggests that *some* portion of the PersEIDS DSS is facilitating better decision making. Users commended the usability and interface of the software. With a very minimal amount of introduction to the program, IV crewmembers seemed to fully understand the information conveyed to them. Through simple questions/answers, IV crewmembers communicated that they understood how to use the information conveyed to manage the consumables of an EVA.

The decision making of the EV/IV pair highlighted an emergent behavior of the EVA management task: there are a plethora of mitigation strategies that human beings use to bolster safety during an EVA. In this case, the mitigation strategy involved limiting the maximum metabolic rate of the EV crewmember. PersEIDS does not explicitly model this intervention; the only intervention modeled within PersEIDS is to change the active timeline (which could mean removing or modifying future tasks). The overall preliminary success of the system, combined with this emergent use case, supports the idea that PersEIDS (in its current form) best supports the decision-making process merely by providing safety and consumables projections. When under nominal, safe EVA operations, no intervention is needed. However, when an unsafe active timeline is realized, the human operator will consider many intervention strategies before fully committing to a different active timeline. The positive feedback on safety projections is a positive development, but further investigation will be needed to assess the cognitive fit of the PersEIDS intervention recommendations to those an actual mission controller may consider.

This initial demonstration, while successfully illustrating the initial software capabilities of the platform, was severely limited in scope, so much so that it is hard to generalize conclusions about the overall fit of the PersEIDS DSS to the cognitive processes used by the IV crewmember to support a mission. For example, the IV crewmember maintains certain internal mental models, and it is unknown how well the transparently-implemented PersEIDS models fit the IV's internal mental models of risk tolerance (e.g., is a safety score from 0 to 100 percent a proper representation? Can the user comprehend that?) and risk uncertainty. Nor is it well known if the designed decision support aid supports optimally (i.e., to maximize the situational awareness of the operator). Decision tree 1, the safety assessment of a given timeline, expressly fits within level 3 (projection) of the Endsley model of situational awareness[21] because it synthesizes safety over a projection of the entirety of the future timeline. Decision tree 2, on the other hand, is a level 2 situational awareness support aid because it only provides safety of the contingency return EVA *in the current moment*, and the user is expected to monitor this in order to anticipate when the EVA is nearing the "point of no return." Future work is needed to assess these levels of decision support to each relevant decision and optimize the engineering design thereto.

Biomedical context is not the only context necessary to assess the safety of a planned timeline. For example, atmospheric or Martian weather events may endanger the health and safety of the crew. Similarly, decisions to change an active timeline may be informed by non-biomedical data; operational reasons like broken equipment may force a change to a timeline. In both of these cases, PersEIDS is not a necessary or sufficient decision support tool. Further work is needed to understand the entire scope of the EVA work domain, particularly the work of the IV crewmember managing the mission. PersEIDS will likely exist as one decision support tool within a large ecosystem of tools supporting EVA.

Further work is also needed in future demonstrations to address a number of confounding factors. The IV crewmember was not a trained mission controller, although those subjects did work in and around the EVA community. It is unknown how similarly actual mission controllers would utilize PersEIDS or if the prescribed mitigations (in this case, limiting metabolic rate) are actual mitigation strategies to be used in flight. Similarly, it is unknown if there are other mitigation strategies to be employed before micromanaging a crewmember's metabolic rate. The contrived demonstration test condition of a harder-than-expected traverse task does represent a real risk in flight, but there are other risks which may be more appropriate to demonstrate PersEIDS capabilities. Additionally, the use of a similar EVA test condition on day two and day three may induce a training bias to the EV crewmember: if they expect that they may have a harder-than-usual traverse, they may consciously or subconsciously manage their workload so as to protect the margin of consumables needed to perform such a traverse.

VII. Conclusion

PersEIDS represents a DSS for biomedical information during EVA. In the initial demonstration of the system, an EV/IV crewmember pair in VR was able to successfully navigate an off-nominal EVA scenario which was unable to be addressed on a prior day without the decision support aid. This may suggest positive decision support; however, limitations in the demonstration design preclude the ability to generalize conclusions about system performance. Future work is needed to understand the true fit and performance of the system in EVA-like scenarios. Martian EVA is at least a decade away; as future understandings of EVA needs and constraints are elucidated, the PersEIDS system will undergo architecture changes to accommodate these needs. PersEIDS is currently a research toolbox but will be used to inform future flight software to power these decisions on Mars. PersEIDS, in conjunction with EVA timeline management software and other companion DSS, will power the future of human decision-making during space exploration operations.

Acknowledgments

This work was funded by NASA Exploration Capabilities.

References

[1] B. Belobrajdic, K. Melone, and A. Diaz-Artiles, “Planetary extravehicular activity (EVA) risk mitigation strategies for long-duration space missions,” *npj Microgravity*, vol. 7, no. 1. 2021. doi: 10.1038/s41526-021-00144-w.

[2] K. H. Beaton *et al.*, “Extravehicular activity operations concepts under communication latency and bandwidth constraints,” in *IEEE Aerospace Conference Proceedings*, 2017. doi: 10.1109/AERO.2017.7943570.

[3] A. H. Stevens *et al.*, “Tactical Scientific Decision-Making during Crewed Astrobiology Mars Missions,” *Astrobiology*, vol. 19, no. 3, 2019, doi: 10.1089/ast.2018.1837.

[4] D. S. S. Lim *et al.*, “The BASALT Research Program: Designing and Developing Mission Elements in Support of Human Scientific Exploration of Mars,” *Astrobiology*, vol. 19, no. 3, 2019, doi: 10.1089/ast.2018.1869.

[5] K. H. Beaton *et al.*, “Using Science-Driven Analog Research to Investigate Extravehicular Activity Science Operations Concepts and Capabilities for Human Planetary Exploration,” *Astrobiology*, 2019, doi: 10.1089/ast.2018.1861.

[6] S. E. Kobs Nawotniak *et al.*, “Opportunities and Challenges of Promoting Scientific Dialog throughout Execution of Future Science-Driven Extravehicular Activity,” *Astrobiology*, vol. 19, no. 3, 2019, doi: 10.1089/ast.2018.1901.

[7] D. Coan, “Exploration EVA system concept of operations,” 2020.

[8] P. G. W. Keen, “DECISION SUPPORT SYSTEMS: A RESEARCH PERSPECTIVE,” in *Decision Support Systems: Issues and Challenges*, 1980. doi: 10.1016/b978-0-08-027321-1.50007-9.

[9] I. Vessey and D. Galletta, “Cognitive fit: An empirical study of information acquisition,” *Information Systems Research*, vol. 2, no. 1, 1991, doi: 10.1287/isre.2.1.63.

[10] M. J. Miller and K. M. Feigh, “Assessment of Decision Support Systems for Envisioned Human Extravehicular Activity Operations: From Requirements to Validation and Verification,” *J Cogn Eng Decis Mak*, vol. 14, no. 1, 2020, doi: 10.1177/1555343419871825.

[11] M. J. Miller *et al.*, “Extravehicular Activity Mission System Software (EMSS)-Enabling Human Planetary Exploration Data Within The Broader Planetary Data Ecosystem,” in *5th Planetary Data Workshop (PDW) and 2nd Planetary Science Informatics & Data Analytics (PSIDA) meeting*, 2021.

[12] S. Hillenius, “Designing interfaces for astronaut autonomy in space,” in *CanUX 2015 Conference*, 2015.

[13] J. J. Marquez *et al.*, “Supporting real-time operations and execution through timeline and scheduling aids,” in *43rd International Conference on Environmental Systems*, 2013. doi: 10.2514/6.2013-3519.

[14] J. J. Marquez, S. Hillenius, I. Deliz, J. Zheng, B. Kanefsky, and J. Gale, “Enabling Communication between Astronauts and Ground Teams for Space Exploration Missions,” in *IEEE Aerospace Conference Proceedings*, 2019. doi: 10.1109/AERO.2019.8741593.

[15] D. J. Macfarlane, “Open-circuit respirometry: a historical review of portable gas analysis systems,” *European Journal of Applied Physiology*, vol. 117, no. 12. 2017. doi: 10.1007/s00421-017-3716-8.

[16] N. J. Langford, “Carbon dioxide poisoning,” *Toxicological Reviews*, vol. 24, no. 4. 2005. doi: 10.2165/00139709-200524040-00003.

[17] K. Permentier, S. Vercammen, S. Soetaert, and C. Schellemans, “Carbon dioxide poisoning: a literature review of an often forgotten cause of intoxication in the emergency department,” *International Journal of Emergency Medicine*, vol. 10, no. 1. 2017. doi: 10.1186/s12245-017-0142-y.

[18] K. J. Kim *et al.*, “The Partial Pressure of Inspired Carbon Dioxide Exposure Levels in the Extravehicular Mobility Unit,” *Aerosp Med Hum Perform*, vol. 91, no. 12, 2020, doi: 10.3357/AMHP.5608.2020.

[19] M. Fernández and S. Williams, “Closed-form expression for the poisson-binomial probability density function,” *IEEE Trans Aerosp Electron Syst*, vol. 46, no. 2, 2010, doi: 10.1109/TAES.2010.5461658.

[20] Y. Hong, “On computing the distribution function for the Poisson binomial distribution,” *Comput Stat Data Anal*, vol. 59, no. 1, 2013, doi: 10.1016/j.csda.2012.10.006.

[21] M. R. Endsley, “Toward a theory of situation awareness in dynamic systems,” *Human Factors*, vol. 37, no. 1. 1995. doi: 10.1518/001872095779049543.