Research Report

GEORGE NACOUZI, OSONDE A. OSOBA, JONATHAN TRAN, SASCHA ISHIKAWA

Artificial Intelligence and Machine Learning Applications for Defensive Counterspace

A Decision Support Tool Capability Demonstration

This publication has completed RAND's quality-assurance process but was not edited.

For more information on this publication, visit **www.rand.org/t/RRA582-1**.

About RAND

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.
© 2024 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

About This Report

This report builds on a prior-year effort to harness artificial intelligence and machine learning (AI/ML) for decision support in defensive counterspace (DCS) missions. This updated effort aims to improve the AI/ML model for DCS missions, to improve the physics-based modeling and simulation tool, and to incorporate a user interface. The intended audience for this report includes U.S. Space Force leaders, requirements staff, space operators, and AI/ML experts who are interested in the broader context for which the models were developed.

The research reported here was commissioned by the Department of the Air Force's United States Space Force and conducted within the Force Modernization and Employment Program of RAND Project AIR FORCE.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the Department of the Air Force's (DAF's) federally funded research and development center for studies and analyses, supporting both the United States Air Force and the United States Space Force. PAF provides the DAF with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future air, space, and cyber forces. Research is conducted in four programs: Strategy and Doctrine; Force Modernization and Employment; Resource Management; and Workforce, Development, and Health. The research reported here was prepared under contract FA7014-22-D-0001.

Additional information about PAF is available on our website: www.rand.org/paf/

This report documents work originally shared with the Department of the Air Force (DAF) on August 3, 2021. The draft report, dated February 2023, was reviewed by formal peer reviewers and DAF subject-matter experts.

Acknowledgments

This research was made possible through the support of our original sponsors, Lt Gen William Liquori and Col Jay Fulmer (USSF S5/9), who were instrumental in helping us define the focus and path of the research. We want to give special thanks to Col Fulmer, who also took on the role of action officer, for the consistent support, assistance, and leadership he provided throughout the term of the project; his contributions were key to making this project successful. We also want to thank the many government subject-matter experts (SMEs)—representing the National Space Defense Center (NSDC), Air Force Research Laboratory (AFRL), National Reconnaissance Office (NRO), Space Security and Defense Program (SSDP), and Defense

Advanced Research Projects Agency (DARPA)—who provided valuable and insightful feedback and suggestions.

We thank Jim Chow, the director of RAND's Project AIR FORCE, for his guidance and support for this research, and our RAND colleagues, Bonnie Triezenberg and Li Ang Zhang, for their careful and thoughtful reviews of the report. This report has been much improved by their efforts. Finally, we are grateful to Sydne Newberry and John Godges for their crucial assistance in the preparation of this document.

Abstract

This report summarizes the results from a FY20 add-on project doing a detailed examination of the applicability of artificial intelligence and machine learning (AI/ML) as an enabler of decision support in defensive counterspace missions. This examination builds on work developed on the previous project by improving the AI/ML model for decision support as well as the modeling and simulation tools used for training. This work also goes further by incorporating a user interface for the decision support tool demonstration. The aim is to use our concrete development and demonstration to identify the key or relevant complexities that the United States Space Force (USSF) will need to contend with as it aims to apply AI/ML tools to operations.

Our work highlights two key dimensions of complexity in the use of AI/ML for USSF missions. The first dimension is the *technical* complexity in the kinds of AI/ML models that can be useful for operational missions in contested adversarial settings. Any AI/ML decision aid will need reliable access to situational awareness data and signals. Thus, we argue that infrastructure investments in robust data pipelines will pay dividends, especially when AI/ML tools are brought to bear. The second dimension is the complexity of *integrating* the technical prowess of AI/ML tools into the operational decision-making process. The development of our demonstration tool highlighted the complexity of designing effective interfaces between operators and computer tools. We recommend that USSF both clarify its doctrine of operational AI/ML and invest in improving the science of the human-computer interface.

Summary

The United States Space Force (USSF) is implementing a Space Warfighting Construct to prepare the United States for a conflict that extends into space. One sub-function of this construct involves preparing warfighters for combat by giving them tools for defensive counterspace (DCS) operations—tools that help to rapidly develop effective space-based courses of action (CoAs) in response to adversarial activities against friendly space assets. This project builds on our FY19 project, which examined the value of artificial intelligence and machine learning (AI/ML) by developing and demonstrating a sample AI/ML-based decision support tool for DCS.[1] This report describes the improved AI/ML-based decision-support workflow we developed and the technical details of the AI/ML models driving the tool discussed here.

Background Context

DCS operations require the rapid orchestration of a set of activities, including maintaining continuous awareness of an opponent's activities while coordinating multiple actions by U.S. forces. Efficient DCS operations will likely require advanced computing and decision support tools that can assist space operators in planning, assessing, selecting, and executing complex CoAs in highly time-constrained conditions.

The USSF asked RAND's Project AIR FORCE to assess the possibility of using AI/ML tools as a decision aid for a space warfighter facing an attack on a friendly space system. The earlier, FY19, project established the potential viability of AI/ML for DCS with an early proof-of-concept. However, questions remained about how to *improve* the tool's performance and how to *integrate* such tools into DCS operational workflows.

Approach

RAND conducted a study with these objectives and related methods for fulfilling them:

- Improve the AI/ML model for decision support and the modeling and simulation tools.
- Develop and demonstrate our implementation of a DCS decision support tool, which includes both an upgraded AI/ML decision architecture and a web-based user interface intended to be integrated into the operator's DCS decision workflow (see Figure S.1).

---

[1] Nacouzi, George et al., "Assessment of Artificial Intelligence for Defensive Counterspace Missions," RAND Corporation, RR-4400/1-AF, July 2020, Not available to the general public; Osoba, Osonde et al., "Assessment of Artificial Intelligence for Defensive Counterspace Missions," RAND Corporation, RR-4400/2-AF, July 2020, Not available to the general public.

- Use the demonstration process to identify important issues that USSF may encounter as it integrates AI/ML into its operations—and recommendations for dealing with the issues.

Figure S.1. Sample View of User Interface for DCS Decision Support Tool

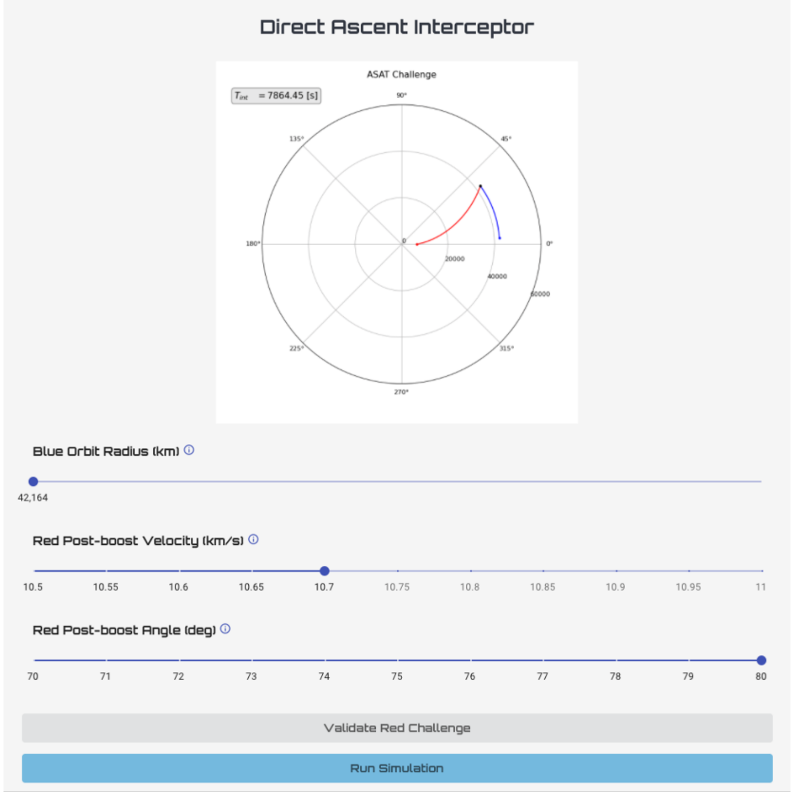

Blue Orbit Radius (km)
42,164

Red Post-boost Velocity (km/s)

Red Post-boost Angle (deg)

Conclusions

- AI/ML models can support the selection of high-quality CoAs in DCS missions.
- Any AI/ML decision aid will need reliable access to space situational awareness (SSA) data and signals. Robust data pipelines are infrastructure investments that will pay efficiency dividends, especially when AI/ML tools are brought to bear.

Recommendations Based on This Research

Based on our experience designing and implementing this proof-of-concept demonstration, we recommend that the USSF take these next steps:

- Develop a roadmap for an operational AI/ML decision-support tool for DCS missions.
- Prioritize establishing SSA data pipelines.
- Deploy AI/ML decision aids when training USSF operators.

Recommendations for Future Research

- Investigate the doctrinal implications of using AI/ML in space warfighting.
- Invest in the science of human-computer interface (HCI) and the use of game-theoretic models for AI/ML algorithms.

Contents

Figures and Tables

Figures

Tables

Chapter 1. Introduction and Objectives

Introduction

Satellite systems provide essential services to military and civilian societies worldwide. The United States (U.S.) military relies on space systems to provide essential support to its missions during peacetime and conflict. Space services including position, navigation, and timing; satellite communication; intelligence, surveillance, and reconnaissance (ISR); and missile warning are all critical parts of military operations and as such are critical components of the joint war-fight. U.S. space capabilities are currently unmatched by any other nation and provide a significant advantage to all-domain military operations. Consequently, potential adversaries have been developing anti-satellite ground and space-based capabilities to disrupt, deny, or destroy U.S. space systems to gain an advantage in future military operations.[2] [3] These actions have made it clear that we cannot expect a future of unimpeded action in space; space is now a war-fighting domain.[4]

In response, the U.S. is implementing a Space Warfighting Construct to prepare the DAF for a conflict that extends into space.[5] The construct, which is intended to prepare the space enterprise to fight through a conflict, involves preparing the warfighter for combat, improving the resilience of the architecture, leveraging materiel and non-material solutions, and integrating with allies and the commercial sector. This report focuses on the first capability, i.e., preparing the warfighter for combat. This involves enabling the warfighter to rapidly develop effective space-based courses of action (CoAs) in response to adversary activities against friendly space assets.

Space operations are complex and require highly precise timing due to the high speed and intricacies involved in orbital mechanics. These operations are further complicated when they involve defensive counterspace (DCS) operations, which require the rapid orchestration and integration of various activities, including maintaining continuous awareness of an opponent's activities while coordinating multiple actions by U.S. and friendly forces. Examples of space defense operations are discussed by Harrison and include CoAs such as maneuvering friendly satellites, deploying decoys, and using bodyguard satellites to protect the attacked space system.[6]

---

[2] Harrison, Todd et al., "Space Threat Assessment 2021," Center for Strategic & International Studies, March 2021.

[3] National Air and Space Intelligence Center, "Competing in Space," NASIC Public Affairs Office, December 2018.

[4] Thompson, Lt Gen David "D. T.," Col Gregory J. Gagnon, and Maj Christopher W. McLeod, "Space as a War-fighting Domain," Air & Space Power Journal, 2018.

[5] Air Force Space Command, "Space Warfighting Construct," April 2017. As of June 14, 2021.

[6] Harrison, Todd et al., "Defense Against The Dark Arts in Space," Center for Strategic & International Studies, February 2021.

Selecting and assessing the best defensive CoAs that meet a commander's intent is a challenging and time-consuming endeavor that is further complicated by the likelihood of multiple simultaneous attacks, limited insight into the attacker's intents, and the limited available response time. This can result in some fleeting options, i.e., a delay in the execution of the defensive CoA can result in lower effectiveness, possibly eliminating that CoA option.

In the last few years, technological advances have lowered the barriers to space, and over 80 nations now have some space-based capabilities.[7] Commercial companies such as OneWeb and SpaceX plan to deploy proliferated constellations of hundreds to thousands of satellites in low earth orbit (LEO) by the early 2020s. These new ventures could increase the total number of active satellites in orbit by an order of magnitude by 2030.[8] The shrinking barrier to space access means any nation can build a presence in space. This will likely lead to the space environment being more congested, thereby making it a more challenging and probably contested domain in which to operate.

Potential adversaries have also benefited from the technological advances in space systems. They have noted the utility and advantages derived by the U.S. military from space services and have been developing both space and counterspace capabilities with the goal of gaining space superiority. Both China and Russia are developing a range of kinetic and non-kinetic counterspace capabilities, both ground and space based.[9] Iran and North Korea have also been developing capabilities to affect space systems, mostly as ground-based jammers. However, as spacefaring nations, they could develop more sophisticated counterspace weapons in the future.[10]

Earlier RAND Work

With this context in mind, in 2019 tasked RAND's Project AIR FORCE to assess the possibility of using artificial intelligence/machine learning (AI/ML) tools as a decision aid for a space warfighter faced with a potential attack on a friendly space system. The intent was to assess whether AI/ML can help the warfighter rapidly select among existing Blue space based CoAs to defend and protect the targeted space system.

In 2019, we completed a project in which we developed a custom AI/ML-based decision tool as a proof of concept to specify several high-performing notional U.S. (Blue) DCS CoAs against two example adversary (Red) CoAs.[11] The first Red CoA involved a direct ascent interceptor, and the second Red CoA assumed an orbital based interceptor attack against a friendly satellite. We constrained our model to these two CoAs for the purpose of demonstrating the approach;

---

[7] European Space Policy Institute, *ESPI Report 79- Emerging Spacefaring Nations*, June 2021.

[8] Cates, G.R. et al., "Launch Uncertainty: Implications for Large Constellations," The Aerospace Corporation, November 2018.

[9] National Air and Space Intelligence Center, 2018.

[10] National Air and Space Intelligence Center, 2018.

[11] Nacouzi, George et al., 2020; O. Osoba et al., 2020.

however, other Red CoAs are possible as delineated by the National Air and Space Intelligence Center (NASIC) (2019). We assumed that the AI/ML model has a direct data feed from the Blue space situational awareness (SSA) system. This simplification does not impact the performance of the AI/ML algorithm; however, a more realistic representation would include some delay and uncertainties associated with the data, such as limited sensor and processing capabilities. The Blue CoAs were similar to the options discussed by Harrison (Harrison et al., 2021b). The 2019 report included a detailed workflow summarizing the activities involved in a DCS operation. We also recommended an overall architecture centered around the use of an AI/ML decision support tool which is maintained and continuously improved and reviewed by the relevant stakeholders including DCS operators and SMEs.

One of the products of the FY19 effort was a working AI/ML decision support prototype tool based on a customized implementation of the canonical *generative adversarial network* (GAN) architecture. We also explored different model deployment options (e.g., disaggregated models for responding to specific Red threats or to provide best Blue responses of a specific type). The AI/ML-based decision support tool showed promising results and demonstrated the feasibility and value of the approach. We refer the reader to the 2019 project report for a detailed discussion of the findings and recommendations.[12]

Several improvements were suggested in the 2019 project, both for the AI/ML decision support algorithm and the modeling and simulation (M&S) approach that we used to generate the training data for the algorithm.[13] The primary limitation of the FY19 AI/ML prototype was the use of repeated model sampling to approximate optimal Blue responses (as opposed to making the effectiveness optimization intrinsic to the model). We addressed this limitation in this work with a newer hybrid AI/ML model architecture. The sponsor also asked for a more detailed demonstration with an intuitive user interface to give leaders more insight into the feasibility and utility of AI/ML for the DCS mission.

Project Scope

The DCS mission involves numerous steps, including collecting and analyzing SSA data, gathering ISR of opponents' space and counterspace systems, developing "patterns of life," and identifying possible indicators and warnings (I&W). These activities, although critical to the mission, are outside the scope of both the 2019 and the current project. The RAND team focused only on the identification and selection of CoAs, as shown in Figure 1.1, in response to possible hostile activities against a friendly space system. The decision support tool we developed

---

[12] Nacouzi, George et al., 2020.

[13] Nacouzi, George et al., 2020.

provides the candidate CoAs to the user along with their estimated effectiveness, allowing users to make an informed decision and select the CoA that best fits their needs.[14]

Figure 1.1. Study Scope Constrained to Courses of Action

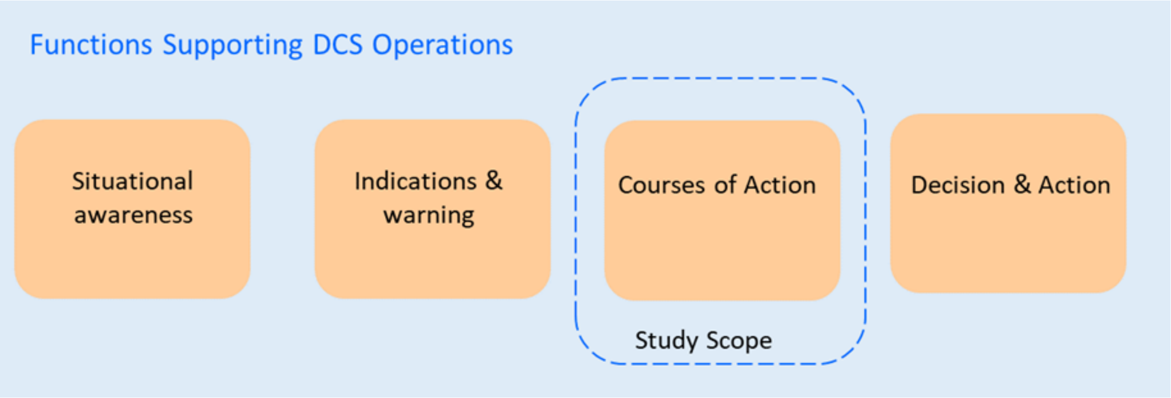

The follow-on project implemented many improvements (suggested and new) for the AI/ML decision tool. The project also developed a user interface (UI). The purpose of the UI is to help with the demonstration of the AI/ML tool, but not to represent an operational UI.

How We Improved and Used the Tool

This report describes technical improvements to the AI/ML model that underpins the decision support tool (in Chapter Two), shows a sample run of the tool (in Chapter Three), and offers recommendations for further improvement (in Chapter Four).

As a preview of Chapter Two, the most salient technical improvements include the following:

- **Improvement of M&S**: The M&S tool used to generate the training data for the AI/ML algorithm underwent significant improvements including:
  - code refactoring (or the restructuring of code to create smaller or more-manageable components), which improved extensibility and readability and enabled interfaces to couple the M&S tools directly with the AI/ML algorithm and UI.
  - enhanced fidelity of the generated training data set though improved algorithms (which provide greater realism and increased resolution) and an updated experimental design (which allows for a more comprehensive data set).
- **Assessment of Impact of Decision Delay:** Our presentation of decision options in the decision tool allows users to get an estimate of the impact of delaying the execution of

---

[14] We recognize that the CoAs considered in the research are simpler than what one would expect in real operations, which likely affects how quickly the tool identifies possible candidate solutions. However, we find that the tool identifies and assesses possible viable CoAs much quicker than a human operator.

the suggested Blue CoA. This feature is intended to provide the user with insight into the fleeting nature of some of the Blue CoA options.

- **Improvement of AI/ML Model Architecture**: The 2019 decision tool used a Generative Adversarial Network (GAN) model,[15] which was successful at identifying an acceptable set of responsive Blue CoAs. We updated the learning model by augmenting the generative loss function with explicit reward terms (adapting ideas from actor-critic reinforcement learning[16] and imitation learning[17] models). The new model, the Reward-Augmented GAN Planner (RAGAN-P), is more capable at identifying and selecting the highest performing Blue CoA. This improvement yields CoAs of higher estimated effectiveness than were generated with the initial GAN approach.

Report Organization

Chapter Two discusses the technical improvements in greater detail. Chapter Three walks through the UI of the DCS demonstration tool, showing how to execute it. Chapter Four summarizes the findings and provides some suggestions and recommendations.

---

[15] A generative adversarial network (GAN) is a ML model architecture designed to learn to imitate complex features of a data set and generate similar synthetic samples. We trained a GAN architecture in FY19 to generate synthetic Blue CoA responses by "learning to imitate" good Blue CoAs found in the physics simulation database.

[16] Konda, Vijay R., and John N. Tsitsiklis, "Actor-Critic Algorithms," *Advances in Neural Information Processing Systems*, pp. 1008-1014, 2000.

[17] Ho, Jonathan, and Stefano Ermon. "Generative Adversarial Imitation Learning," *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4572-4580, 2016.

Chapter 2. Technical Improvements to the Decision Support Tool

The most salient technical improvements include the following:

- improvement of M&S.
- assessment of impact of decision delay.
- improvement of AI/ML model architecture.

This chapter addresses each in turn.

Improvement of M&S

The simulation model for this work elaborates and improves on our prior models built for the original report (Nacouzi et al., 2019; Osoba et al., 2019). This section discusses some of these improvements.

Code Refactoring

In anticipation of future developments, we first restructured the existing model's code to improve its maintainability and extensibility. We ported the prior model from Octave to Python to accommodate the interfacing between the M&S framework, the AI/ML algorithm, and the UI. Without modifying the external behavior of the M&S framework, we modularized its functions, both simplifying the underlying logic and eliminating code redundancy. During this process, the M&S framework underwent periodic output verification to ensure capability matching with the previous project. After fully refactoring the M&S framework from the prior study, we implemented numerous improvements described in the following sections.

Physics Capabilities

The implementation of a modular M&S framework allowed us to expand the physics modeling capabilities from the prior study. The higher-fidelity physics, in turn, allowed for a more accurate assessment of mission effectiveness across all Blue responses and a reduction of the problematic edge cases that had been generated as part of the synthetic data set; most notably, the inclusion of hyperbolic Kepler equations allowed our model to assess a wider range of challenge and response trajectories. We also made minor adjustments to how the mission effectiveness is scored for each Blue response. As a result, the synthetic training data set (described below) has undergone valuable improvements from the prior study.

Expansion of Action Spaces

The description of action spaces remains consistent with the prior study. We maintain the same CoAs for Red and Blue but implemented improvements with the new M&S framework.

Firstly, we can model all combinations of notional Red and Blue CoAs. Secondly, we expanded the number of parameters for generating the training data set. Most notably, we included a Blue parameter that accounts for a delay in the Blue response to a Red challenge. Thirdly, we reevaluated the parameter ranges for generating the training data set. In many cases, we updated the parameter ranges to provide a more realistic training data set.

Synthetic Training Data Set

With these improvements, the synthetic training data set is now generated using a full factorial experiment design in which each relevant Red and Blue parameter is discretized (that is, assigned a discrete quantity) and all possible combinations of discrete values across all parameters are explored. It should be noted that the training data set does not parameterize the Blue response delay time (that happens later), and all data points assume immediate action in response to the Red challenge. The resulting sample count for the synthetic data set is provided in Table 2.1 for three Blue CoAs and two Red CoAs. Compared with the prior study, there are more sample points and a redistribution of sample points such that each Red and Blue action is more equitably represented in the training data set.

Table 2.1. Sample Count Statistics in Synthetic Data Set (n = 30250)

Blue CoA	Red CoA1	Red CoA2	Total
Blue CoA1	7986	7986	15972
Blue CoA2	1331	1331	2662
Blue CoA3	5808	5808	11616

Inline Simulation Interface

The updated M&S framework allows for directly linking the AI/ML model with the web UI applications. This provides two benefits. The most relevant one is that it allows users on the web UI to evaluate and visualize the effects of time delays of candidate Blue response CoAs in real time. The other benefit is that directly interfacing the AI/ML model with the M&S framework allows us as well as users to receive training feedback on candidate scenario solutions. (We did not take full advantage of this benefit because of the costs of software integration and computation time for querying the online simulation model during AI/ML training.)

Assessment of Impact of Decision Delay

Time is an important factor in decision-making for DCS. There are strategic implications if Blue incorrectly assesses Red intent or acts before Red's intent is clear, but there is also significant risk by waiting for additional information before responding to a potential attack.

While we may have more response options if we act early, doing so could increase the risk of choosing a less appropriate, or even escalatory, CoA, along with potentially incurring an associated political cost, by affording Red the opportunity to claim its action was benign.

Since a Blue CoA also has a time-dependent effectiveness or availability that decreases over time, decision-making in this context relies on identifying the maximum amount of uncertainty that can be tolerated while still executing an effective defense. A noted strength of decision analysis is the ability to calculate the value of additional information to the decisionmaker.[18] This is key to scenarios in which Blue would like to know the "cost" of additional information, as paid by the potential loss of CoA availability and effectiveness. More specifically, Blue would like to know the maximum acceptable time "price" for procuring further information about Red's intent. We can examine what happens as the waiting time increases and identify the point at which Blue is *indifferent* between waiting and acting with the best option available at that time. The point at which Blue is indifferent indicates the maximum time/cost Blue would be willing to pay for additional information.

Notional plots of the time-dependency of effectiveness appear below. Representing one Blue CoA against one Red attack, these plots illustrate how chosen, designated thresholds—based on risk tolerance for effectiveness (yellow lines) and probability of attack against Blue asset (orange lines)—will determine whether this CoA could be a viable option given time constraints. **Error! Reference source not found.** illustrates a scenario in which a time window exists in which Blue is willing to risk executing this CoA given the confidence level of attack characterization at that time. **Error! Reference source not found.** then shows what happens if the rate of effectiveness degradation is increased, illustrating that this CoA may not be ideal for this scenario given the inherent risk associated with its use under uncertainty.[19]

---

[18] See, for example, La Valle, 1968; Merkhofer, 1977. In operational contexts, this evaluation may not be feasible to do, especially under intense time pressure. There might be value in developing rapid heuristics or approximations for estimating the value of additional information in real time. Or at least identifying clear operational policies on decision delays, regardless of what the theoretically optimal "time price" might be.

[19] Note that for each possible Red state, there will be a threshold probability for Blue action as well as a threshold timing for this action. From this we see an increasing probability as time progresses. This of course is only the case if Red is truly in that state and we are acquiring information that confirms this truth. A full evaluation of the decision space requires an examination of each possible Red state as time progresses, and the probability of that state as perceived by Blue increases. From this evaluation, we can identify an optimal action time window for each branch and identify where, or if, the branches overlap.

Figure 2.1. Effectiveness of CoA Is Still Above Threshold When Attack Is Sufficiently Attributed

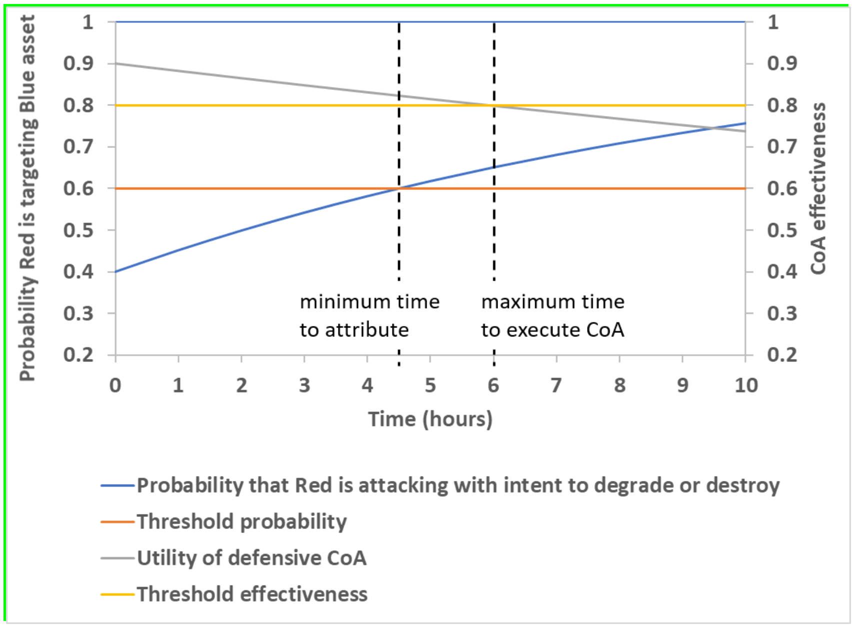

Figure 2.2. Degradation of CoA Effectiveness Occurs Too Rapidly to Execute Once Blue Has Sufficiently Attributed the Attack

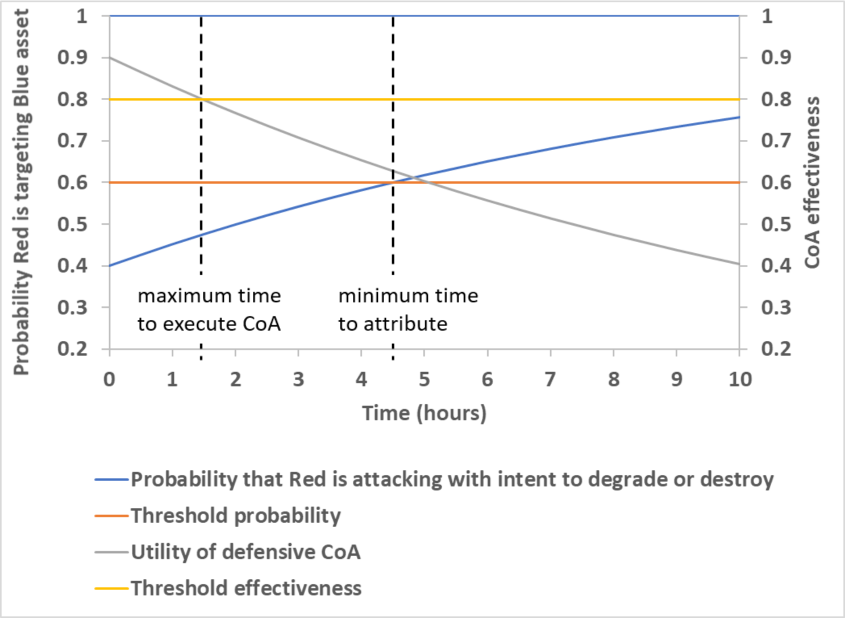

Improvement of AI/ML Model Architecture

The FY19 AI/ML decision support model (see Figure 2.3) was based on a GAN architecture (Goodfellow et al., 2014). The architecture consists of a generator model, G(.), interacting with a discriminator model, D(.), such that the tuned generator achieves useful imitative behavior. In other words, the key idea behind the FY19 implementation was that *if* the model learned to imitate "good" responses found in the data set, then the final trained model would be an effective policy model for Blue. This strategy turned out to be only moderately successful. The model learned to present *valid* responses, but the responses were not always highly *effective*.

Figure 2.3. Schematic for Previous FY19 Generative Adversarial Network (GAN) Model

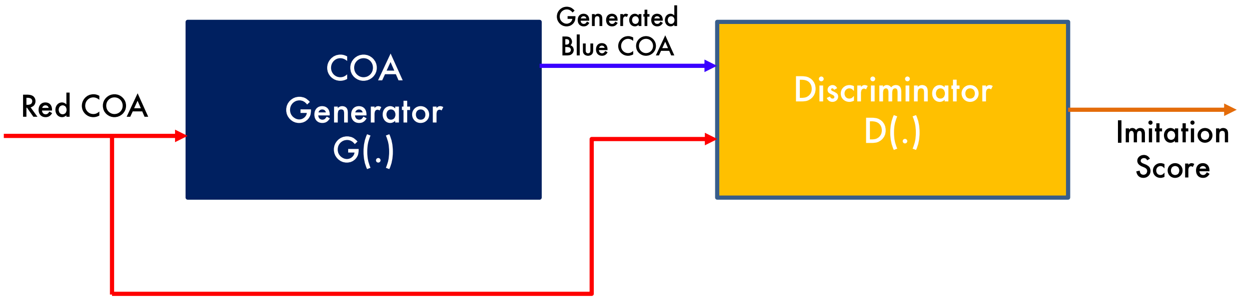

We addressed this shortcoming by adding a repeated sampling step and filtering the sampled responses subject to a quality criterion. In this way, we incorporated an intrinsic motivation for optimal responses directly into the ML model. A successfully trained policy model will thus produce good or near-optimal Blue responses by default.

To achieve this outcome, we switched from a primarily GAN-based architecture to a single-time-step, *Actor-Critic* reinforcement learning (RL) architecture.[20] The new model architecture (see Figure 2.4) retains an adversarial sub-component path (red line) to train the model to learn valid internal representations of scenarios and Blue responses. But we augmented the base model structure with a second signal path (blue line) that estimates Blue's reward if a generated/proposed Blue response is carried out in the real world. Using this "score estimator" model allows us to update Blue's response policy model via an iterative mathematical algorithm known as gradient descent. Thus, we modify the learning objective by adding an extra term that motivates the policy model to maximize the expected critic score. For this reason, we call the updated model the RAGAN-P.

---

[20] The structure of this hybrid GAN/actor-critic model is also based on insights from Ho and Ermon's Generative Adversarial Imitation Learning (GAIL) (Ho and Ermon, 2016) and InfoGAIL (Li, Song, and Ermon, 2017). Our use of a score critic is similar to the use of reward augmentation in InfoGAIL. But our setup bears a stronger similarity to actor-critic models, with the primary differences being (1) our use of pre-training for the critic instead of training both the actor and critic concurrently, and (2) our use of a GAN-style generator instead of a standard stochastic neural network to represent the actor model. Our use of a pre-trained critic is more efficient than having the critic model learn values via temporal-difference learning from sample data in real time. It represents a form of *partially* model-based learning.

Figure 2.4. Schematic for Current RAGAN-P Model with Both Effectiveness and Imitation Scores

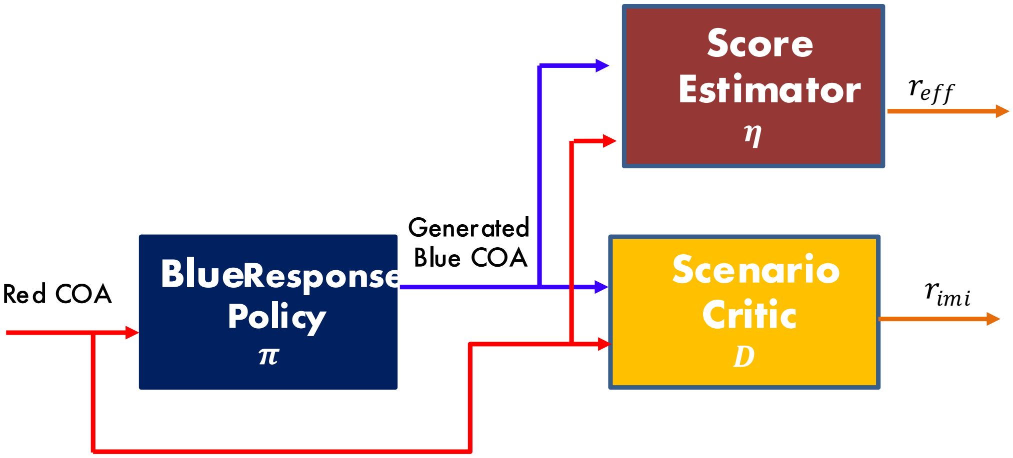

The goal of our training efforts is to produce a Blue response policy that can generate successful responses to a Red challenge. To achieve that end, we use the model structure outlined in Figure 2.4. The model provides two kinds of evaluative feedback for each designated Red challenge and Blue response. The first is an imitation score, r_{imi}, that indicates how statistically similar the generated scenario is to samples in the simulation database. The second score, r_{eff}, estimates how well the suggested Blue response would fare in the real world.

Chapter 3. Demonstration of Decision Support Tool User Interface

This chapter addresses the complexity of integrating the technical prowess of AI/ML tools into the operational decision-making process. The development of our DCS demonstration tool highlighted the complexity of designing effective interfaces between operators and computer tools. Our tool represents just some preliminary attempts to manage this complexity.

The DCS tool consists of two parts. The first part is an accessible web-based UI for walking through a counterspace decision-making context. The second part is a backend AI/ML decision support model. The UI communicates with the backend model that provides optimized Blue responses to Red DCS challenges specified in the UI. The backend model also includes simulation elements that can provide the user with information about how decision delays can impede response effectiveness.

In the following sections, we walk through the expected steps for using the DCS tool, from the specification of Red challenges to the analysis of Blue response options. We end with an exploration of potential improvements and future directions. We emphasize that the UI described here was designed for and intended to run with the backend AI/ML demonstration tool, thus it is not intended, by any means, to represent an operational UI.

Figure 3.1 shows the landing page of the UI.

Figure 3.1. DCS Tool User Interface "Landing Page"

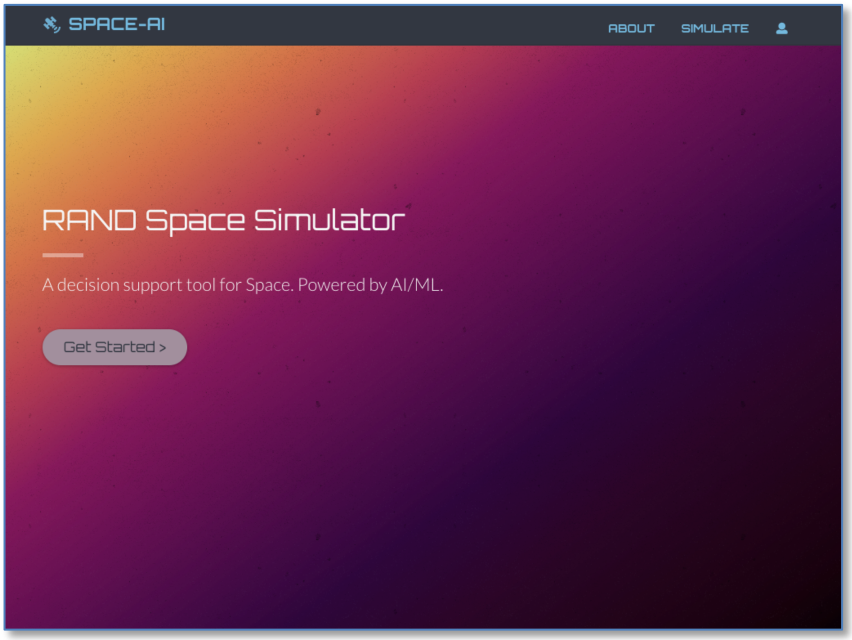

RAND Space Simulator

A decision support tool for Space. Powered by AI/ML.

DCS Tool Workflow

We conceive the use of the DCS tool occurring in three steps, as reflected in the UI of the tool:

1. specifying the Red challenge
2. sampling Blue responses from the AI/ML tool
3. time-delay characterization of individual Blue responses.

Step One: Specifying the Red Challenge

We ask the user to choose the Red challenge currently in play. The Red CoA options consist of either direct ascent or sub-orbital attack, as Figure 3.2 shows. An operational deployment of this kind of tool would need to include any other relevant Red threats as selectable options. Such deployments would also rely on SSA data pipelines to automate this Red CoA selection. Without a direct feed of timely and appropriate SSA data—e.g., Red interceptor trajectory—the tool's effectiveness would likely be greatly diminished. (Timely and relevant SSA data is needed to develop Blue CoAs regardless of the type of tool or approach used.)

Figure 3.2. On the "Simulate" Page, a User Is Prompted to Select a Red CoA

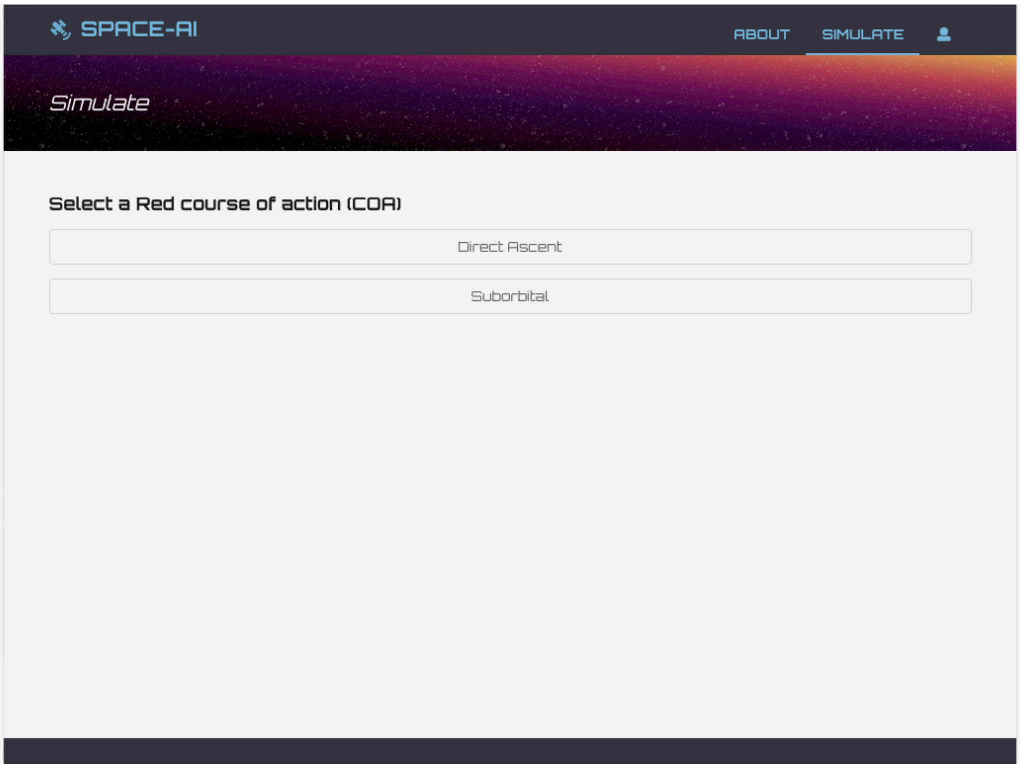

Once a Red CoA is chosen, it must be defined by selecting the detailed parameters for that CoA (e.g., burnout velocity and burnout angle, which appear as "post-boost" velocity and angle in Figure 3.3). Before running the simulation, the parameters are validated to ensure the Red challenge is viable. Clicking the “Validate Red Challenge” button (near the bottom of Figure 3.3) checks if the selected Red CoA parameters represent a valid Red challenge. If so, an intercept trajectory is plotted. If not, the tool indicates the failure, and a Blue CoA is not offered. A valid challenge leads to the next step (by clicking the “Run Simulation” button at the bottom of the figure).

Figure 3.3. A User Is Prompted to Specify Red Parameters for a Direct Ascent Intercept

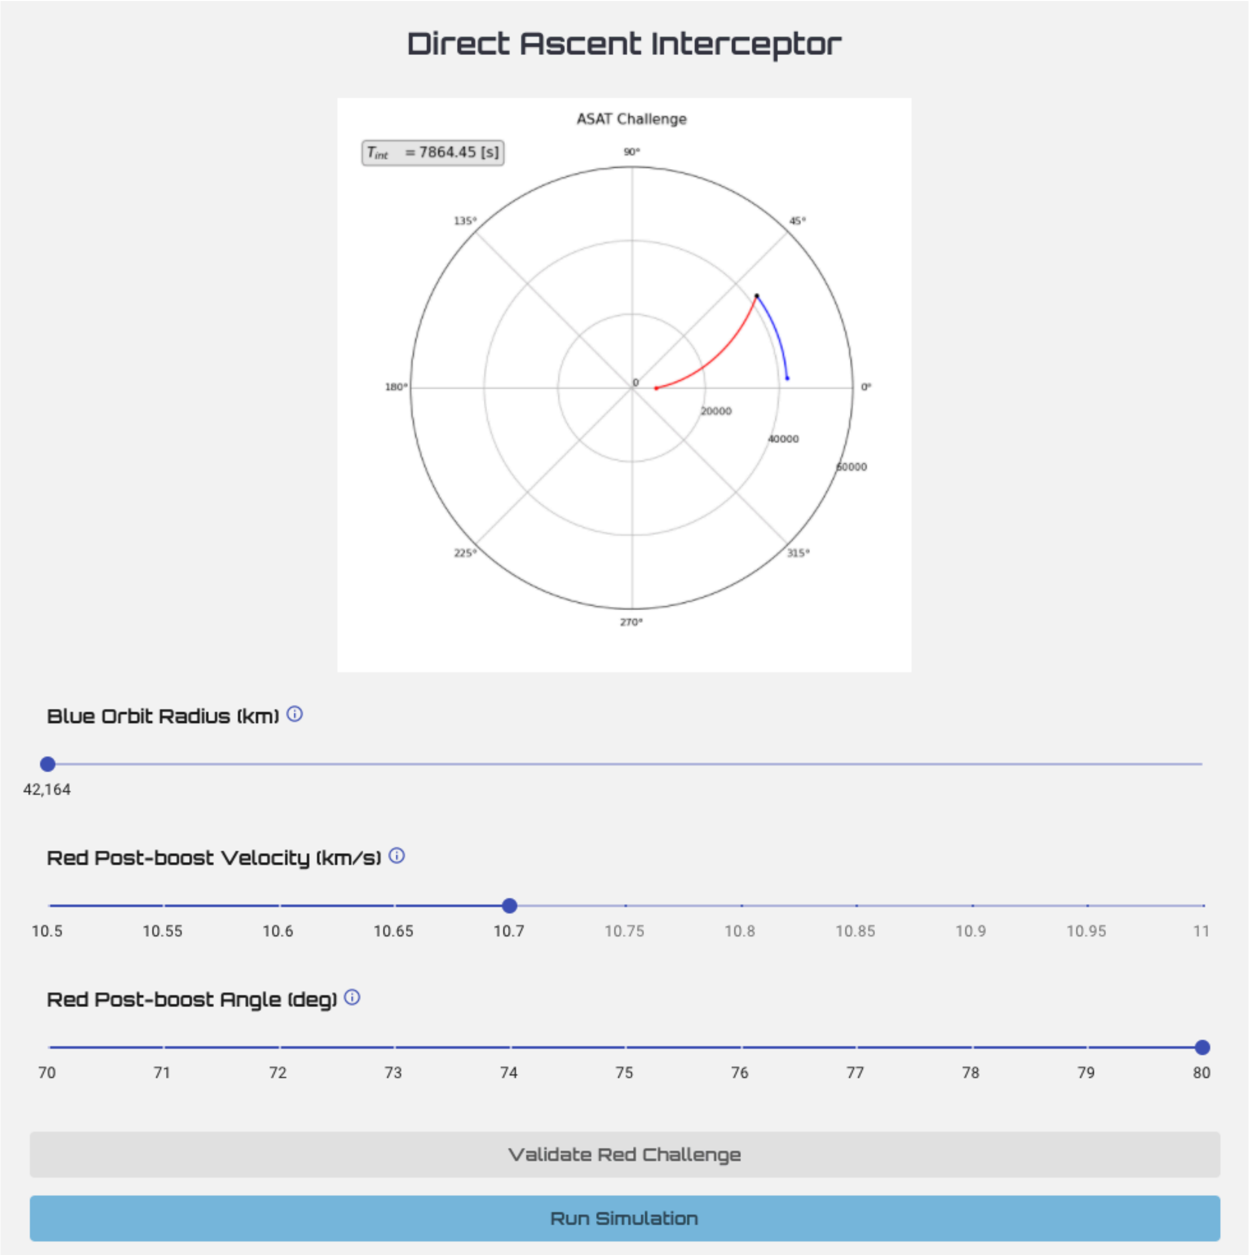

Blue Orbit Radius (km) ⓘ 42,164

Red Post-boost Velocity (km/s) ⓘ 10.7

Red Post-boost Angle (deg) ⓘ 80

Step Two: Sampling Blue Responses from the AI/ML Tool

Clicking the "Run Simulation" button submits the Red CoA parameters to the backend AI/ML decision model. The model generates a portfolio of Blue responses. Each Blue response is a valid, potential CoA in response to the same Red challenge. Examples of Blue actions include maneuvering an attacked satellite, deploying decoys, or launching a guardian satellite[21] to intercept the attacking missile, similar to the Blue CoAs discussed by Harrison and colleague. Each Blue CoA receives a score for mission effectiveness, a *delta-v* (or change in velocity) required for the response, and additional CoA-specific information such as the number of decoys deployed.

To give the user a sense of the overall effectiveness of the model's recommendations, the model generates a histogram of effectiveness for the entire portfolio of Blue responses. Figure 3.4 shows a sample distribution of effectiveness scores for 100 AI/ML-generated Blue responses to a single Red CoA.

Figure 3.4. Distribution of 100 Blue COA Effectiveness Scores for Responding to a Red CoA

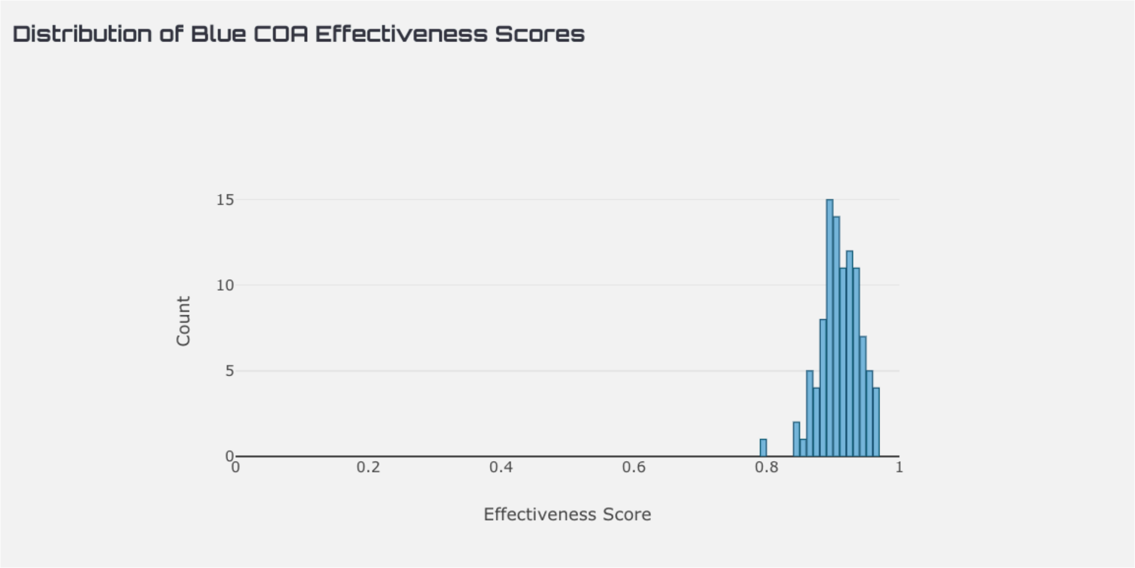

To give the user detailed insights into each recommended alternative, Table 3.1 lists a sampling of the Blue CoAs and the parameter values corresponding to each level of mission effectiveness. This list provides the user-operator with a "dashboard" of CoA options, which appear in a descending order of their effectiveness scores.

---

[21] A guardian satellite is an orbital-based interceptor used to 'guard' another satellite by intercepting an incoming physical threat.

Table 3.1. Sample List of Blue Responses Generated by the AI/ML Model

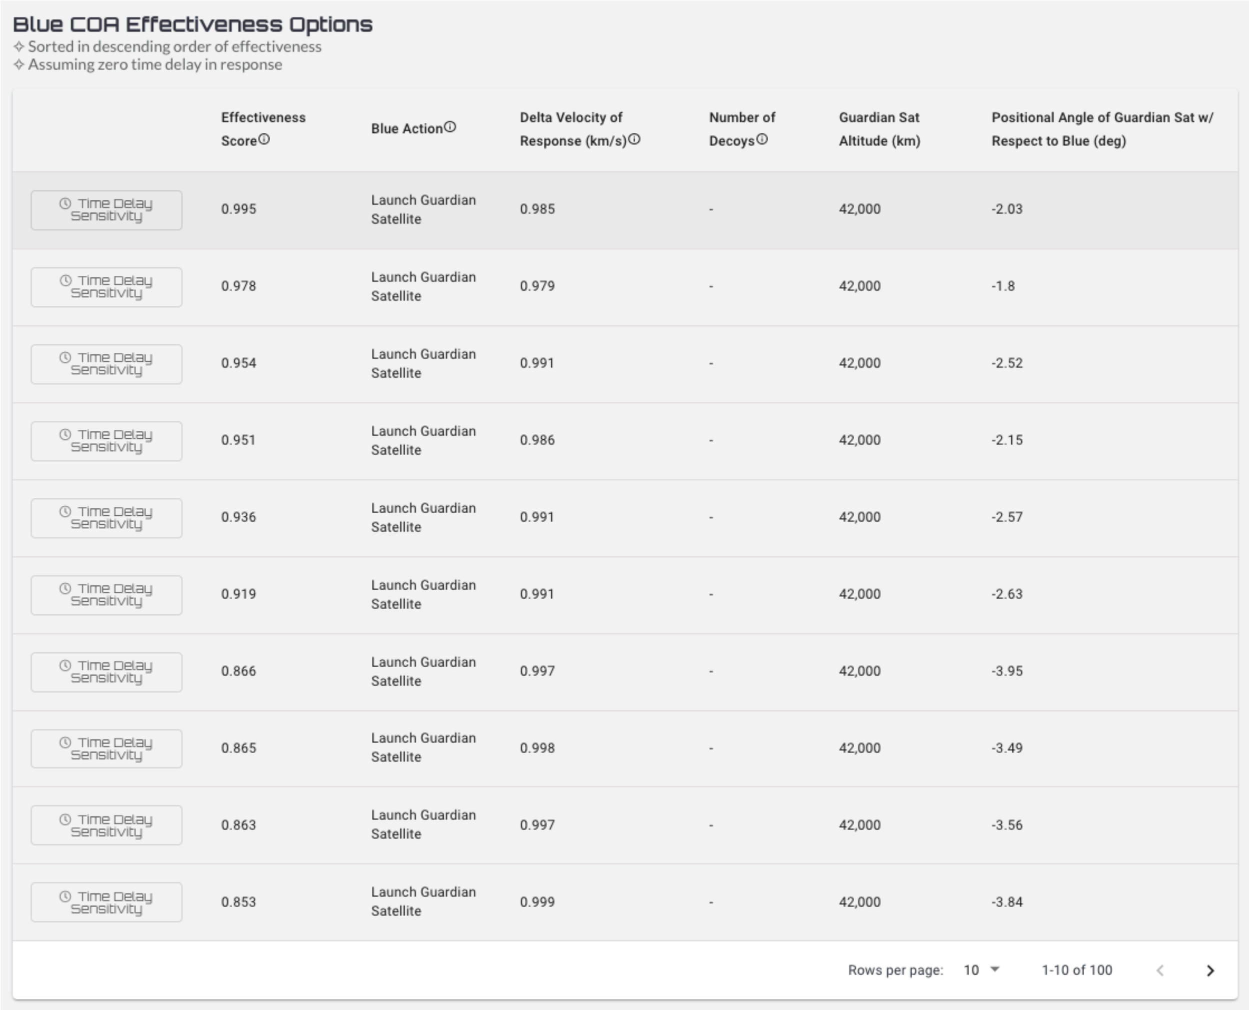

Blue COA Effectiveness Options

✧ Sorted in descending order of effectiveness
✧ Assuming zero time delay in response

	Effectiveness Score ⓘ	Blue Action ⓘ	Delta Velocity of Response (km/s) ⓘ	Number of Decoys ⓘ	Guardian Sat Altitude (km)	Positional Angle of Guardian Sat w/ Respect to Blue (deg)
Time Delay Sensitivity	0.995	Launch Guardian Satellite	0.985	-	42,000	-2.03
Time Delay Sensitivity	0.978	Launch Guardian Satellite	0.979	-	42,000	-1.8
Time Delay Sensitivity	0.954	Launch Guardian Satellite	0.991	-	42,000	-2.52
Time Delay Sensitivity	0.951	Launch Guardian Satellite	0.986	-	42,000	-2.15
Time Delay Sensitivity	0.936	Launch Guardian Satellite	0.991	-	42,000	-2.57
Time Delay Sensitivity	0.919	Launch Guardian Satellite	0.991	-	42,000	-2.63
Time Delay Sensitivity	0.866	Launch Guardian Satellite	0.997	-	42,000	-3.95
Time Delay Sensitivity	0.865	Launch Guardian Satellite	0.998	-	42,000	-3.49
Time Delay Sensitivity	0.863	Launch Guardian Satellite	0.997	-	42,000	-3.56
Time Delay Sensitivity	0.853	Launch Guardian Satellite	0.999	-	42,000	-3.84

These sample Blue CoAs are produced by the back-end AI/ML model specifically designed for this task. The updated RAGAN-P model architecture (shown again in Figure 3.5) combines its traditional actor-critic reinforcement learning architecture for checking validity (leading to the yellow box) with a reward-augmented generative adversarial network architecture for scoring effectiveness (leading to the maroon box). This hybrid architecture can better capture the internal structure of CoAs and scenarios while also generating high-scoring Blue response CoAs.

Figure 3.5. RAGAN-P Model Contains Two Evaluators: Effectiveness Score and Imitation Score

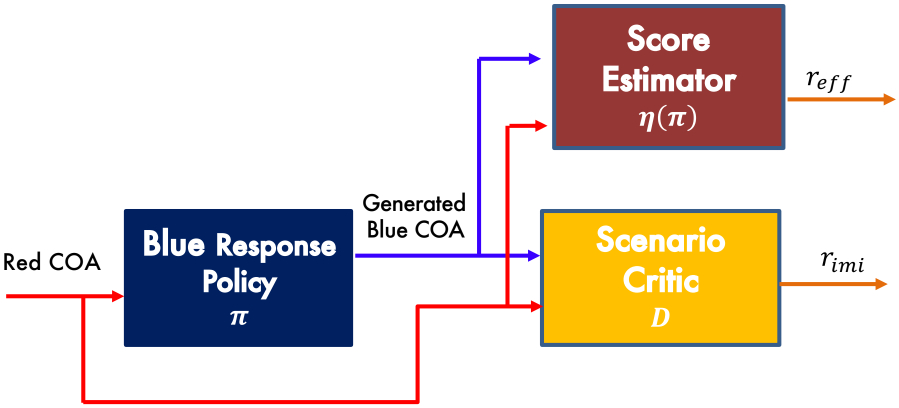

Step Three: Time-Delay Characterization of Individual Blue Responses

Clicking a Blue CoA on the dashboard (as in Table A.1) produces a time-delay sensitivity chart (as in Figure 3.6). The aim is to provide detailed insight on the consequences for delaying the execution of a response CoA, i.e., waiting too long may render the suggested CoA ineffective.

Figure 3.6. Time-Delay Sensitivity Plot of a Blue CoA

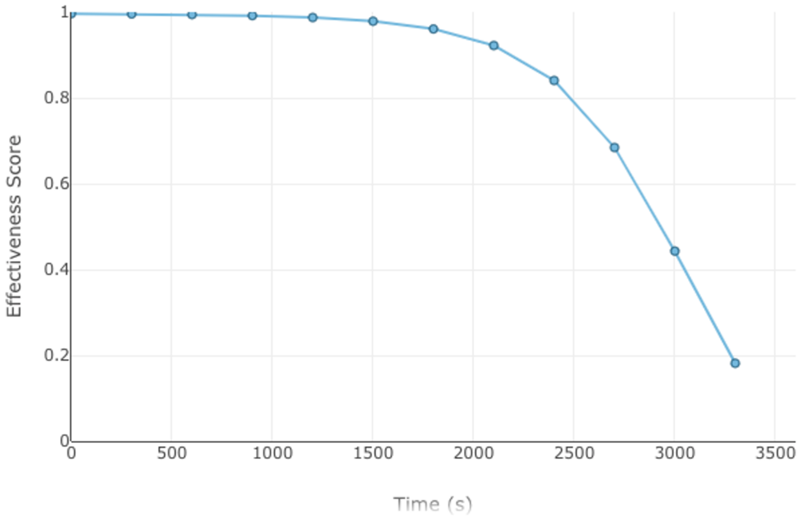

Potential Improvements

The goal of the tool is to better inform responses to DCS challenges. The tool is limited by the fact that the user must specify the current Red CoA manually. This additional step takes time and is prone to manual errors. An operational version of the tool would use a more direct and effective approach by providing the model with a feed of appropriately vetted sensor data. Upon receiving the vetted details of a Red CoA, the DCS tool could then respond in a more timely and automated manner to provide the user(s) with options that are both valid and effective.

Chapter 4. Findings, Recommendations, and Conclusion

Findings

The outcome of this project confirms our FY19 findings that using AI/ML for DCS decision support can provide significant value to an operator needing to make timely and impactful defensive decisions.[22] The decision support tool quickly provides the operator with multiple effective CoA options that can be executed to defeat the Red attack.[23] Our detailed demonstration of how AI/ML-based decision-support tools may be tested in an operational setting has yielded two key takeaways:

1. **Although AI/ML-assisted CoA selection in DCS operations is a tractable computational problem, confident and timely situational awareness signals are crucial:** Our work on this as well as the FY19 project provides reasons for optimism but reaping the benefit of AI/ML tools in real-world settings will require a proper integration of these decision support tools into the DCS Observe-Orient-Decide-Act (OODA) loop. The most basic consideration for this integration is having timely access to situational awareness information, e.g., tracking a hostile system as exhibited by the sensitivity of its response to time delays.
2. **Further fundamental research is needed:** The model powering the AI/ML DCS demonstration tool has several limitations. More fundamental research is needed to improve the model capability and its practicality in real-world settings. For example, how can we solve the DCS game when we have a proliferated constellation of Blue assets? How do resource constraints (e.g., on *delta-v,* or change in velocity) and/or mission prioritizations change Blue strategy? How do we best account for the uncertainties associated with the Red action, e.g., uncertain intents as well as targets.

Recommendations Based on This Research

1. **Develop a roadmap for an operational AI/ML decision-support tool for DCS missions.** This roadmap could include evaluation by users of a more comprehensive and integrated version of the tool developed in this project—with a data pipeline to an appropriate sensing grid and data source—and deployment of the tool as a prototype in a DCS operations center.

---

[22] The FY19 project provided a preliminary assessment of the value of using an AI/ML based approach to support a DCS decision tool. The assessment was based on an AI/ML based tool the RAND team developed. The tool was designed to suggest effective Blue DCS CoAs in response to two possible Red counterspace attacks. In FY20 the team modified the AI/ML algorithm and improved the tool's ability to identify better Blue CoA responses, i.e., the effectiveness of the Blue CoAs as measured by the probability of defeating the Red CoA attack was improved.

[23] Developing and evaluating even a limited number of Blue CoAs without the aid of a fully automated tool would be time consuming, require subject matter experts and may not identify the most effective CoAs.

2. **Prioritize establishing SSA data pipelines:** Decision support tools function within an OODA loop. They require relevant, timely, and trustworthy SSA to perform their roles. The requisite signals and data pipelines might already exist within DAF and the IC. However, they might be siloed and placed under varying authorities and classification levels. The USSF should develop new or leverage existing space data and signal sources to serve as inputs into diverse USSF decision-support tools and pipelines to reap the benefits of effective AI/ML-assisted decision-making. This is a basic but possibly existential barrier for the use of AI/ML in operational DCS settings: Without the availability of timely and relevant SSA data to run the AI/ML tool, its viability would be undermined. (The ability to develop *any* type of responsive CoA, with or without the use of AI/ML, requires the availability of timely SSA data.)
3. **Deploy AI/ML decision aids when training USSF operators:** Develop a more comprehensive version of the tool as a training instrument for DCS operators. It would give trainees an opportunity to select different CoA responses to various hostile activities against space systems and develop insight into the CoAs' strengths and limitations.

Recommendations for Future Research

1. **Investigate the doctrinal implications of using AI/ML in space warfighting:** There is a need to develop a more detailed doctrine for AI/ML use in space warfighting. The 2018 National Defense Strategy[24] identifies a broad need to invest more in the "military application of autonomy and AI/ML." While this statement helps clarify national priorities, it does not translate into a specified doctrine for AI/ML use in the domain of space. There are several possible postures on AI/ML's use and autonomy, e.g., includes human-in-the-loop, human-on-the-loop, context-dependent autonomy and fully autonomous. Each posture has strategic implications that need to be considered within the context of space warfighting. This is especially relevant for certifying the approved uses of AI and autonomy against the DoD's adopted AI ethical principles.[25] A future study is needed to consider the doctrinal implications for using AI/ML in space operations, specifically regarding DCS.
2. **Invest in the science of HCI and the use of game-theoretic models for AI/ML algorithms.**

Conclusion

Modern AI/ML can enable a powerful construct for improving or optimizing decision-making. There are performance and efficiency benefits to be gained from proper integration of AI/ML in DCS operational decision-making as the USSF implements its Space Warfighting Construct. This integration requires care. Our goals in this work were to further assess the

---

[24] Mattis, Jim, *Summary of the 2018 National Defense Strategy of the United States of America*, Department of Defense Washington United States, 2018.

[25] U.S. Department of Defense, "DOD Adopts Ethical Principles for Artificial Intelligence," webpage, February 24, 2020.

feasibility of using AI/ML specifically for DCS CoA selection and to identify key lessons that would help USSF operationalize these kinds of tools. Our work indicates that an appropriately designed AI/ML-enabled decision support tool is feasible and can offer significant value by quickly providing an operator with valid and effective CoA options that otherwise would require significant time and effort to formulate.

Abbreviations

AI/ML	artificial intelligence/machine learning
CoA	course of action
DAF	Department of the Air Force
DCS	defensive counterspace
HCI	human-computer interface
HITL	human-in-the-loop
I&W	indicators and warnings
ISR	intelligence, surveillance, and reconnaissance
M&S	modeling and simulation
SSA	space situational awareness
UI	user interface

References

Air Force Space Command, “Space Warfighting Construct,” April 2017. As of June 14, 2021.

Defense Innovation Board. “AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense,” Supporting document, *Defense Innovation Board*, 2019.

European Space Policy Institute, *ESPI Report 79- Emerging Spacefaring Nations*, June 2021.

Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. “Generative Adversarial Nets,” *Advances in Neural Information Processing Systems*, pp. 2672-2680, 2014.

Grabo, C.M., *Anticipating Surprise: Analysis for Strategic Warning,* December 1, 2002. As of September 26, 2023: https://www.ni-u.edu/ni_press/pdf/Anticipating_Surprise_Analysis.pdf

Harrison, Todd, Kaitlyn Johnson, and Makena Young, “Defense Against The Dark Arts in Space,” Center for Strategic & International Studies, February 2021.

Harrison, Todd, Kaitlyn Johnson, Makena Young, and Joe Moye, “Space Threat Assessment 2021,” Center for Strategic & International Studies, March 2021.

Ho, Jonathan and Stefano Ermon. "Generative Adversarial Imitation Learning," *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4572-4580, 2016.

Konda, Vijay R. and John N. Tsitsiklis. “Actor-critic algorithms,” *Advances in Neural Information Processing Systems*, pp. 1008-1014, 2000.

Mattis, Jim. *Summary of the 2018 National Defense Strategy of the United States of America*. Department of Defense Washington United States, 2018. As of September 26, 2023: https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf

Nacouzi, George, Osonde Osoba, Krista Langeland, Nicholas Martin, Bruce McClintock, Natalie Crawford, and Jonathan Tran, “Assessment of Artificial Intelligence for Defensive Counterspace Missions,” RAND Corporation, RR-4400/1-AF, July 2020. Not available to the general public.

National Air and Space Intelligence Center, “Competing in Space,” NASIC Public Affairs Office, December 2018. As of February 25, 2019: https://media.defense.gov/2019/Jan/16/2002080386/-1/-1/1/190115-F-NV711-0002.PDF

Osoba, Osonde, George Nacouzi, Krista Langeland, Nicholas Martin, Bruce McClintock, and Natalie Crawford, "Assessment of Artificial Intelligence for Defensive Counterspace Missions," RAND Corporation, RR-4400/2-AF, July 2020. Not available to the general public.

Thompson, Lt Gen David "D. T.," Col Gregory J. Gagnon, and Maj Christopher W. McLeod, "Space as a War-fighting Domain," Air & Space Power Journal, 2018.

U.S. Department of Defense, "DOD Adopts Ethical Principles for Artificial Intelligence," webpage, February 24, 2020.

Wong, Yuna H., John M. Yurchak, Robert W. Button, Aaron Frank, Burgess Laird, Osonde A. Osoba, Randall Steeb, Benjamin N. Harris, and Sebastian J. Bae, *Deterrence in the Age of Thinking Machines*, RAND Corporation, RR-2797-RC, 2020.