Optimization to Minimize Risks Using Continuous Asymmetric Risk Analysis

Zachry Engel
Lone Star Analysis
Dallas, TX
zengel@lone-star.com

Nickalus Harrill
Lone Star Analysis
Dallas, TX
nharrill@lone-star.com

Nicolas Velez
Lone Star Analysis
Dallas, TX
ncamacho@lone-star.com

Jacob Ediger
Lone Star Analysis
Frisco, TX
jediger@lone-star.com

Randal Allen
Lone Star Analysis
Orlando, FL
rallen@lone-star.com

ABSTRACT

The current Risk Reporting Matrix used by the Department of Defense is a useful tool for determining the risk associated with an action, but it's limited due to the discrete nature of the metric. Currently, the Risk Reporting Matrix is built using two Likert Scales, one for the likelihood of the event, and one for its consequence. This gives the user 25 possibilities for a likelihood-consequence combination. While this gives a concise and easy to understand answer, it does not fully inform the user of the true risk level for the decision.

At I/ITSEC 2023, Engel presented the Continuous Asymmetric Risk Assessment (CARA) which aids in alleviating these shortcomings by transforming the discrete risk matrix to a continuous gradient field. It is designed to provide the user with infinite combinations of likelihood and consequence which more accurately describe the risk associated with the decision in question. Furthermore, by leveraging the use of asymmetric Gaussian distributions, CARA creates confidence intervals around nominal risk, displaying likely outcomes and its variation.

In this paper, we will present an extension of CARA by pairing the existing risk analysis tool with optimization. By leveraging optimization algorithms, we will demonstrate how better-quality decisions can be made using CARA in conjunction with optimization. Through this process, we will demonstrate how a user of CARA can optimize the best decision to reduce risk based on factors such as cost, time, variability, and resources available. This reduces the human error in a decision making process by analyzing risk decisions analytically through optimizing the risk reduction strategy.

ABOUT THE AUTHORS

Zachry Engel is a Senior Analyst for Lone Star Analysis. He has held this position since January 2023, before which he worked as a Research and Development Engineer after completing an internship with Lone Star Analysis in June 2020. Responsible for the development of new technologies for Lone Star, Zachry's primary interests within the company involve studying Artificial Intelligence and Optimization algorithms. He also works for Lone Star's Operational Optimization Solutions business unit to create risk analysis models. Zachry completed his doctorate in Mathematics in August of 2020 at the University of Texas at Arlington under the supervision of Dr. Suvra Pal. Zachry's research interests involved Biostatistics and Optimization algorithms. Zachry received a bachelor's degree in mathematics in May 2016 from the University of Texas at Arlington.

Nickalus Harrill is the Program Manager for Planning and Analysis. He has been with Lone Star Analysis for five years concentrating on Analysis of Alternatives (AoA) and Business Case Analysis (BCA). He has worked with several different military entities and their counterparts solving problems centered on mitigating risk for programs involved in Aircraft Acquisitions, Work Force Selections, Program Sustainment and Cost Analysis. Nickalus received his undergraduate degree in Public Policy from the University of Texas at San Antonio and his M.B.A in Data Analytics from Texas A&M Corpus Christi.

Nicolas Velez Camacho is an Associate Analysis at Lone Star Analysis currently focusing on developing and integrating data analysis systems using Python He joined the company in June 2022 after graduating with a bachelor's degree in Aerospace Engineering May 2022.

Jacob Ediger is an Associate Analyst at Lone Star Analysis. He graduated with a bachelor's degree in mathematics and physics from Central Methodist University in 2021, and he is currently pursuing his Master of Science in Mathematics from Central Methodist University.

Randal Allen is the Chief Scientist of Lone Star Analysis. He is responsible for applied research and technology development across a wide range of M&S disciplines and manages intellectual property. He maintains a CMSP with NTSA. He has published and presented technical papers and is co-author of the textbook *Simulation of Dynamic Systems with MATLAB and Simulink*. He holds a Ph.D. in Mechanical Engineering (University of Central Florida), an Engineer's Degree in Aeronautical and Astronautical Engineering (Stanford University), and an M.S. in Applied Mathematics and a B.S. in Engineering Physics (University of Illinois, Urbana-Champaign). He serves as an Adjunct Professor/Faculty Advisor in the MAE department at UCF where he has taught over twenty aerospace-related courses.

Optimization to Minimize Risks Using Continuous Asymmetric Risk Analysis

Zachry Engel	Nickalus Harrill	Nicolas Velez	Jacob Ediger	Randal Allen
Lone Star Analysis	Lone Star Analysis	Lone Star Analysis	Lone Star Analysis	Lone Star Analysis
Dallas, TX	Dallas, TX	Dallas, TX	Frisco, TX	Orlando, FL
zengel@lone-star.com	nharrill@lone-star.com	ncamacho@lone-star.com	jediger@lone-star.com	rallen@lone-star.com

INTRODUCTION

In 1988, the United States Department of Defense (DoD) developed the Risk Management Framework through the Electronics System Center for the United States Airforce (Garvey and Lindsdowne, 1988). This framework led to the development of the risk matrix tool to visualize the likelihood and impact of a given risk event. Since its development, the risk matrix has been a commonly used tool to visualize risk due to the relative simplicity of the creation and explanation of the risk matrix. The traditional risk matrix, shown in Figure 1 below, is a 5x5 grid with two axes: one axis to represent the likelihood of the event and one axis for the consequence of the event should it occur.

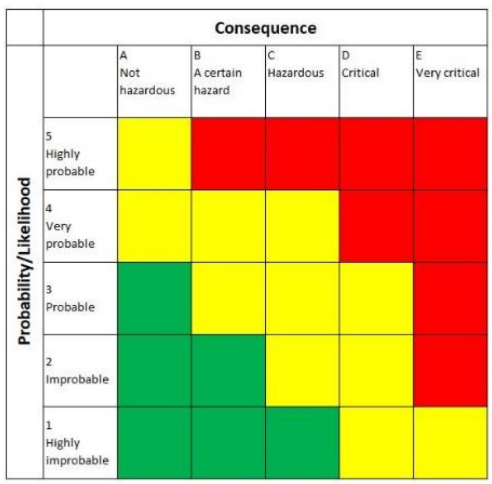

Figure 1 – A general risk matrix.

While the simplicity of this tool is a key reason for its popularity, the simplicity and discrete nature of the risk matrix causes issues for the interpretation of the risk being evaluated. For example, if two separate risk events are in the same region in the risk matrix, it is impossible to assess if the events are truly equally risky. In reality, one event could have a higher degree of risk than the other.

To alleviate the issues with the risk matrix, Engel presented the Continuous Asymmetric Risk Analysis (CARA) at I/ITSEC 2023. CARA is a tool designed as a substitute to the risk matrix used by the DoD. This tool changes the discrete grid of the risk matrix to a continuous gradient field. This allows the user to have an infinite number of options to represent the likelihood of the risk event occurring and the consequence should it occur. Furthermore, by using asymmetric confidence intervals centered around the nominal value of the risk, CARA allows its user to see the nominal value of the risk being evaluated and its variability. Through these methods, CARA provides a more detailed and precise calculation and representation of a risk event.

This paper will demonstrate an extension of CARA, by leveraging optimization to provide better decision-making quality. By using optimization, CARA will show its user how to best allocate available resources to minimize the risk based on factors such as cost, time, variability, and effectiveness. This allows for a decision-making process that alleviates bias in the interpretation of the optimal resource allocation, which in turn leads to strategies optimally designed to minimize risk.

CONTINUOUS ASSYMETRIC RISK ANALYSIS

Overview

A**s** demonstrated by Engel at I/ITSEC 2023, CARA is a tool used to alleviate the issues related to the risk matrix. Like the risk matrix, CARA is based on two axes to determine the risk level of an event: one axis to represent the likelihood of an event and one axis to represent the consequence of the event should it occur. CARA is different from the risk matrix in two key ways that allow for more accurate representation of a risk event: CARA is a continuous gradient of options as opposed to a discrete grid, and it shows an asymmetric two-dimensional confidence interval centered around the nominal risk value. By using a continuous gradient to represent the possible values for the likelihood and consequence instead of a discrete grid, as is used in the risk matrix, CARA allows the user to select any value in the range of options to represent the likelihood and consequence. By showing the user the two-dimensional confidence intervals centered around the nominal risk value, which is the value used for the risk matrix, CARA shows the user how variable the possible outcomes are for the risk event being evaluated. By using these two features of CARA, the user of the tool t is impossible to assess if the events are truly equally risky. In reality one event could have a higher degree of risk than the other.. Figure 2 below shows an example of CARA with its continuous gradient and the asymmetric confidence interval centered around the median value.

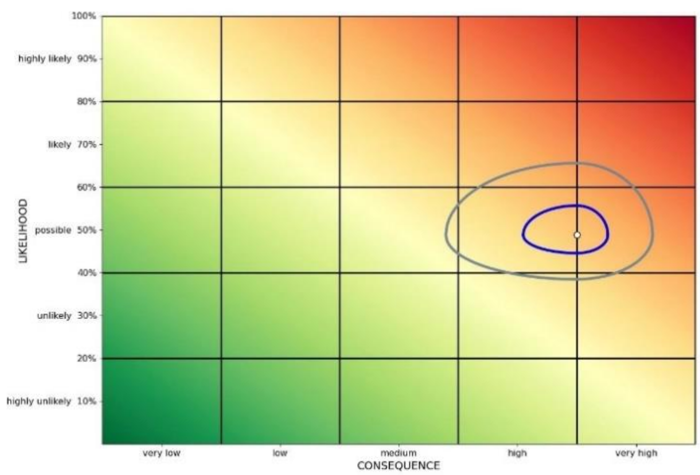

Figure 2 - Example of the output CARA.

Data Collection: Likelihood, Consequence, Mitigation, and Prevention

CARA utilizes a distinct approach to risk evaluation, featuring two independent axes: one for likelihood and another for the impact of events. The collection methodology used in CARA is underpinned by research demonstrating a negligible Pearson correlation coefficient between these factors, thereby enabling separate data collection without mutual influence. CARA accommodates data input through Likert Scale values or continuous metrics, akin to traditional matrices, with Likert scales providing a structured five-level format suitable for both likelihood and impact assessments. For instance, Likert scale values from Department of Defense guidelines illustrate standard criteria used for quantifying likelihood and consequence. What sets CARA apart is its capability to also integrate assessments of prevention and mitigation measures alongside event likelihood and impact. Prevention measures aim to lower both the probability and severity of potential events, whereas mitigation measures focus solely on reducing event severity. This holistic approach allows CARA to offer a comprehensive decision-making tool surpassing the scope of traditional risk matrices, thereby enhancing clarity and effectiveness in risk management practices.

Asymmetric Distribution

CARA introduces a distinctive feature with its two-dimensional confidence interval, crucial for visually assessing risk variability. Initially considering a bivariate normal distribution, the authors found it too symmetrical for the diverse data analyzed in CARA plots, where Likert scale responses and discrete polling methods often create asymmetric distributions. To address this, they adopted a bivariate asymmetric Gaussian distribution, allowing flexibility around the median to better fit the data's asymmetry. CARA calculates median, 10th, and 90th percentile values for both likelihood and consequence, then employs a quantile-parameterized distribution (QPD) to create these distributions. These distributions are independent for likelihood and impact, eliminating the need for a covariance matrix as verified by the Spearman correlation coefficient. This approach enhances CARA's ability to provide accurate and nuanced risk analysis across various scenarios.

Monte Carlo Simulation for Data Generation

After data collection, regardless of its format, CARA employs Monte Carlo simulations to generate data for asymmetric confidence intervals, effectively capturing inherent uncertainties in decision-making. This approach accommodates the variability in Likert Scale values used by the Department of Defense, reflecting the uncertainty associated with each value. Monte Carlo simulations also enable CARA to analyze how prevention and mitigation measures influence both the likelihood and consequence of risks, evaluating numerous scenarios to provide deeper insights into risk assessment. By aggregating outcomes from thousands of simulations, CARA produces a comprehensive risk plot that surpasses the simplicity of traditional risk matrices, where a single decision-maker places a point. This method condenses extensive trial data into a single, informative graph, enhancing the precision and utility of risk analysis in diverse contexts.

Generation of Median Values Asymmetric Confidence Intervals

The CARA tool aggregates data from Monte Carlo simulations to generate a nominal risk value and asymmetric likelihood regions. CARA's output employs a grid that resembles traditional risk matrices to aid familiarity. Background colors represent risk levels, transitioning from green (low risk) to red (high risk) in a gradated manner. The white dot inside the asymmetrical boundaries denotes the median risk value, determined by the median likelihood and consequence. CARA calculates asymmetric bounds using percentiles (10th, 30th, 70th, and 90th) from Monte Carlo results, creating confidence intervals (40% and 80%) around the median to show likely values and their variability.

Applications

CARA aims to evaluate risk in a manner consistent with the DoD protocol while surpassing the limitations of traditional risk matrices by providing a more informative, objective, and flexible decision-making tool. Additionally, CARA can visually represent bowtie analyses, which assess likelihood and impact while integrating prevention and mitigation measures. In a bowtie model, the analysis progresses from assessing the risk event to evaluating prevention measures to reduce its likelihood. Subsequently, mitigation strategies are considered to minimize the event's impact if it occurs, culminating in an overall assessment of likelihood and impact. CARA adopts a similar methodology to assess risk levels but presents visual representations of likelihood and consequence instead of raw data, enhancing accessibility and comprehension. Unlike traditional bowtie models, CARA allows for flexible analysis of various prevention and mitigation strategies in any sequence desired by the user. Leveraging Microsoft Power BI, CARA offers a user-friendly interface for real-time scenario simulations and includes features to filter through a comprehensive array of preventive measures and mitigation strategies. Outputs are visualized as integrated uncertainty clouds on a DoD risk reporting matrix, providing a comprehensive view of risk scenarios. Figure 3 below shows an example of a CARA plot with a single prevention and mitigation measure added to the initial risk statement.

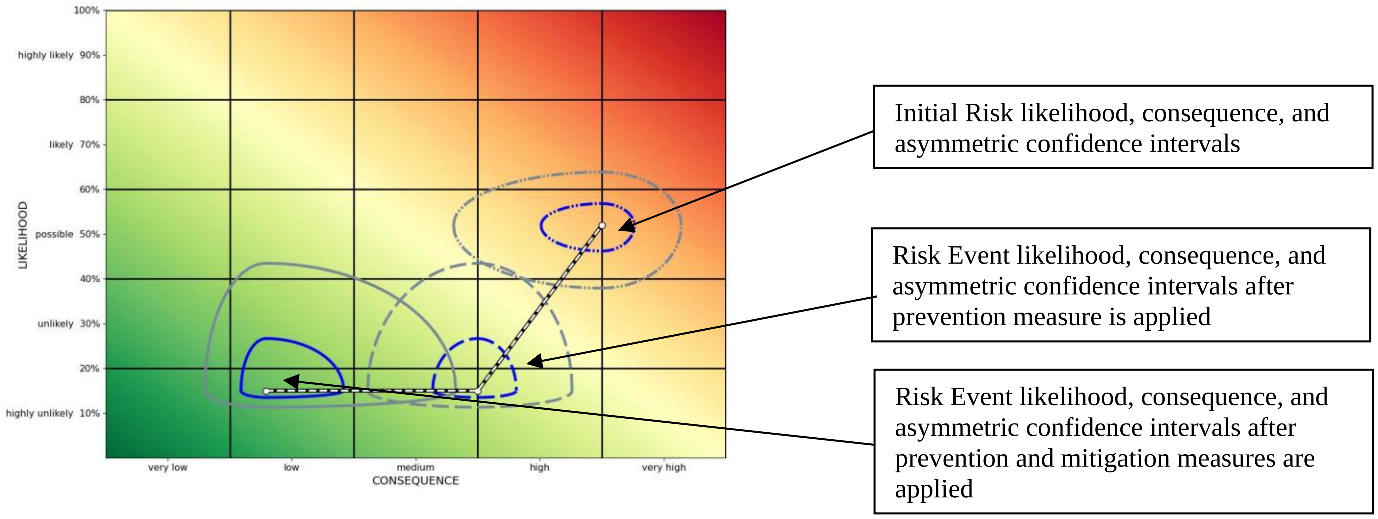

Figure 3 – An initial risk statement showing (1) the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied, (2) a secondary plot showing the application of a single mitigation measure, and (3) a tertiary plot showing the combination of prevention and mitigation measures

OPTIMIZATION

Once the analysis for CARA is completed, the user must be able to make decisions based on the information provided by the analysis and plots. While the plots generated by CARA are easily interpreted, an issue arises from the number of possible outcomes of CARA plots, which is correlated to the number of potential prevention and mitigation measures that are available for that particular risk event. Every CARA plot begins with a single item which represents the original risk statement without any prevention or mitigation measures added to it to reduce the likelihood or consequence of the events. However, CARA will evaluate and generate all possible combinations of prevention and mitigation measures, including the order in which they are implemented. For example, if the risk statement being evaluated has 10 potential prevention and/or mitigation measures, but can only implement 6 of them, the number of possible combinations of those prevention and mitigation measures is 10 choose 6, or 210. With so many possible options, the correct decision becomes difficult to determine and is only exacerbated as the number of possible prevention and mitigation measures increases.

Another issue when deciding which of the strategies to implement is resource availability. When evaluating the results of the CARA plots, a user may instinctually select the combination of prevention and mitigation steps that moves the risk statement furthest to the bottom left of the plot, which would indicate that the risk being evaluated is both unlikely to occur and would have minimal impact should it occur. However, the implementation of those prevention and mitigation measures may exceed the resources available to the decision maker. For example, a single preventative measure may reduce the likelihood of a risk event occurring by 40%, but this measure would also take billions of dollars and decades to create. If the budget of a mission, project, or event being evaluated does not have enough time and money to implement the measure, then the user of CARA would want to consider a separate decision that can be implemented quicker and for less money. The resources that are referenced in this paper can be anything, such as money, time, reduction in variability of the CARA plots, change to productivity, manpower, or any other resource that is important to the user of CARA.

To alleviate these issues, we will use optimization to determine the optimal combination of prevention and mitigation measures. This method will also show the user in what order these measures should be implemented to minimize the likelihood and consequence of the risk event. To accomplish this, we will use an adaptive non-convex stochastic optimization algorithm which will allow CARA to show the user of the tool the optimal data-driven strategy based on available resources and their constraints which allows for a concise decision-making tool.

The goal of utilizing optimization in conjunction with CARA is to minimize the risk of an event. By this we mean the goal is to minimize the likelihood and consequence values of the event being analyzed. However, as mentioned previously, CARA will minimize the risk value of the event in question subject to available resources which will be considered as constraints. If we do not consider any constraints on the optimization problem, then the solution becomes trivial; namely, the optimal solution to minimize the risk of an event is to apply all possible prevention and mitigation measures, regardless of the resources needed to implement these measures.

To begin the process of minimizing the risk in CARA, the user must first identify which available resources are important to the event in question. In most cases, the most pertinent resources are time to implement, and the cost associated with a particular prevention or mitigation measure. However, the user of CARA with optimization can reduce the risk subject to any constraint metric they desire. For example, the user may wish to minimize the risk subject to time to implement and reduction of variation.

By leveraging optimization techniques, the user of CARA can make a decision on the optimal allocation of resources available to minimize the likelihood and consequence of a risk event. This process is agnostic of the subjective inputs a user may have in regards to the correct implementation of the possible prevention and mitigation measures available. The data driven method only evaluates the constraints and resources available and produces the combination of prevention and mitigation measures that lead the maximum reduction in risk.

ANALYSIS AND VERIFICATION

To demonstrate the optimization process within CARA, we conducted a risk analysis for an outside company to evaluate the risk of transitioning their workforce from a direct workforce to a contracted workforce. To accomplish this, we evaluated the risk of transitioning the workforce subject to its impact on cost, performance, and schedule to the company in question. Through this analysis, we also evaluated several possible prevention and mitigation measures that may be implemented to reduce the likelihood and consequence of a given event.

For the authors to conduct this risk analysis, the authors began by evaluating multiple risks corresponding to cost, performance, or schedule. A single risk event may impact cost, performance, and schedule or a subset of these three areas. While any risk event may have an impact on all three areas, some of the impacts were inconsequential and therefore did not warrant investigation. To collect the data, the authors polled many Subject Matter Experts (SMEs) using a five-point Likert scale to gather data similar to the five-point Likert scale used by the DoD to evaluate risks to cost, performance, and schedule. To gather the data on the likelihood and consequence of a prevention or mitigation measure, the authors also used a five-point Likert scale. The SMEs that were polled were asked to give their opinion on:

- The probability that a particular event will occur.
- The impact on cost, performance, or schedule that risk event would have should it occur.
- The likelihood that a prevention or mitigation measure can be applied.
- The impact the prevention or mitigation measure would have on reducing the likelihood or consequence of the risk event if it were applied.

Once the data was collected, the authors used the responses and Monte Carlo methods to determine the values of the median value of the likelihood and consequence. These values are used as the nominal risk value, and the boundary values of the asymmetric confidence intervals. For the optimization study, the authors decided to minimize the nominal risk value subject to two constraints: the time and cost associated with implementing a particular prevention or mitigation measure. To demonstrate the capability of the optimization process, we chose to optimize a risk event that had two possible prevention and three possible mitigation measures. However, for the purpose of the demonstration, we chose to only allow for three of these events to be chosen. Each of the prevention and mitigation measures had an associated cost and time to implement. The goal of the optimization problem was to minimize the nominal risk value, while taking the least amount of time and money. Figures 4, 5, and 6 below show some of the results of the CARA plots and the optimal strategy.

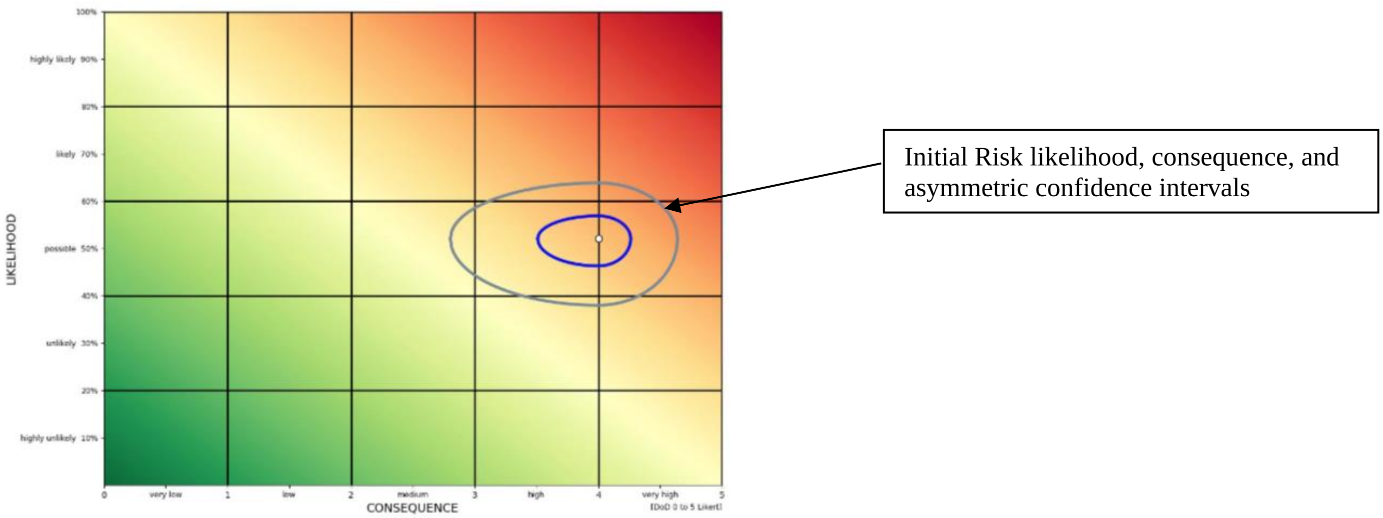

Figure 4 – An initial risk statement showing the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied.

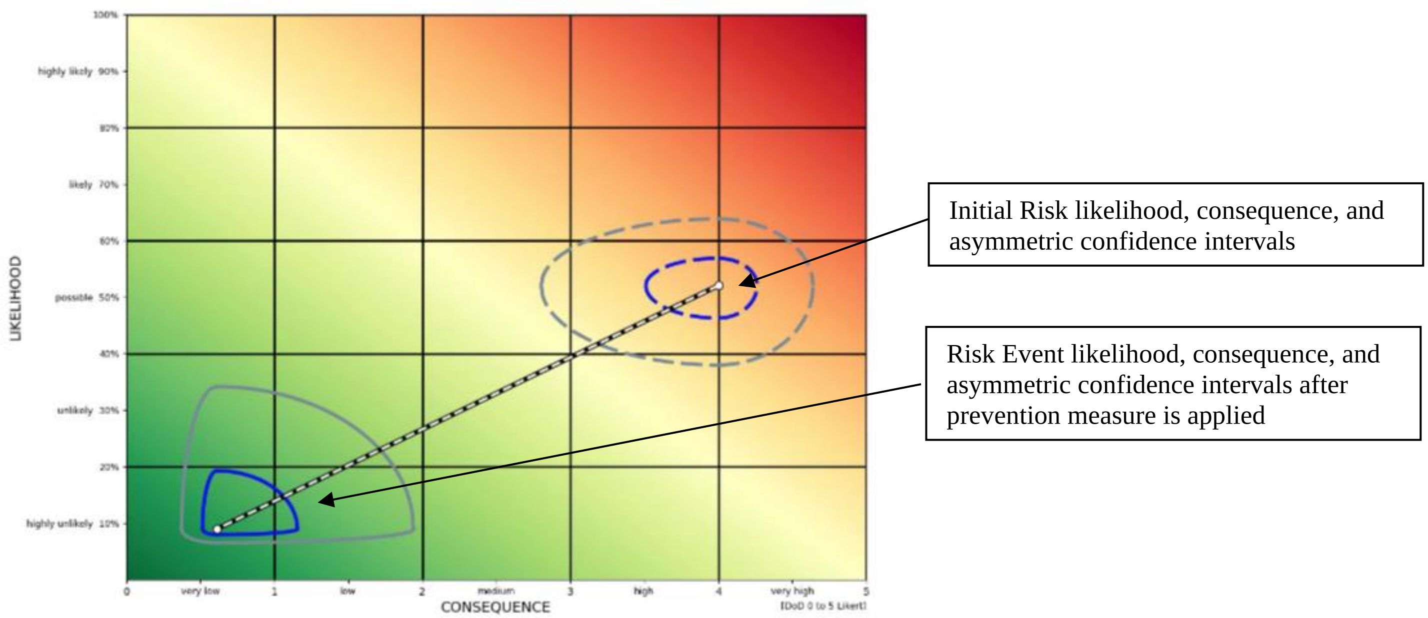

Figure 5 – An initial risk statement showing the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied and a secondary plot showing the application of a single possible prevention measure. This result is not the optimal result of the minimization process.

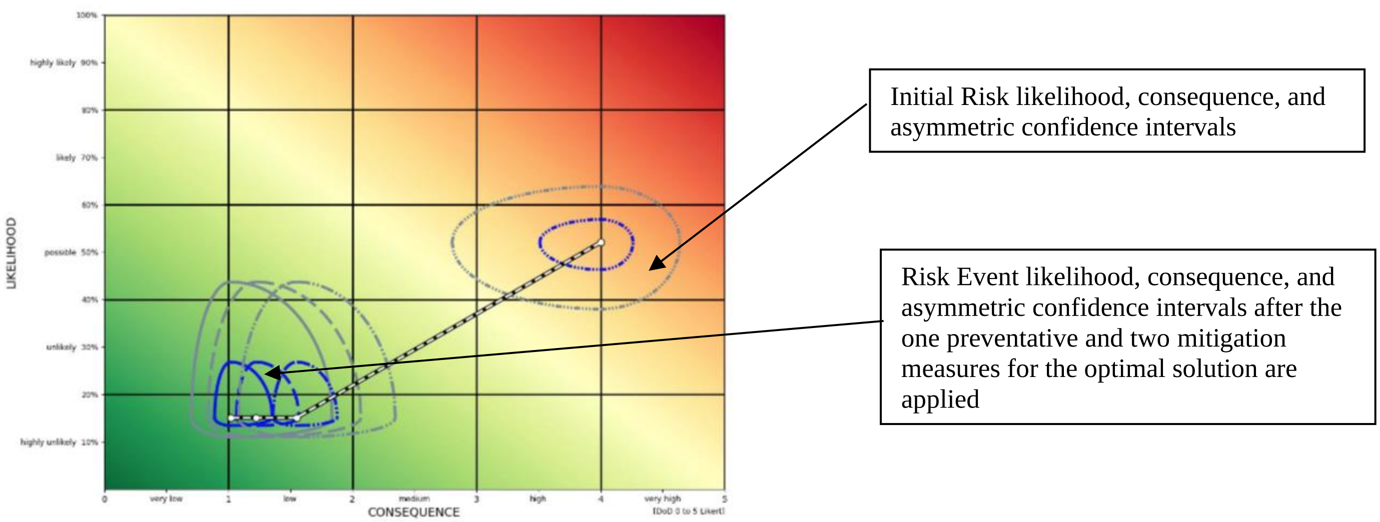

Figure 6 – An initial risk statement showing the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied and three additional plots showing the application of one preventative and two mitigation measures that lead to the optimal solution.

The CARA plots shown in Figures 4-6 show the likelihood, consequence, and confidence intervals for the risk statement and the effects of applying prevention or mitigation measures. As a prevention or mitigation measure is applied, we can see that the plot shifts and the previous plot fades into the background. In this way, the user of CARA can see the result of the application of the prevention and mitigation measures as well as how much the plot moved after their respective implementations.

To a user of CARA who does not consider the time and cost to implement the preventative and mitigation measures, it appears that the single preventative measure applied to create Figure 5 should be implemented since it reduces the probability and likelihood more than the one preventative and two mitigation measures used for Figure 6. However, the preventative measure used in Figure 5 would take 3 years to implement at the cost of 2.5 million dollars. The one preventative and two mitigation measures used to create the plot in Figure 6 would take less than 1 year at the cost of 1.32 million dollars. From this analysis, we can conclude that the application of the one preventative and two mitigation measures in Figure 6 is the optimal decision to reduce the risk of this risk event.

CONCLUSIONS

CARA is a risk analysis tool that leads to in-depth and understandable results from a risk study. The continuous nature of the gradient used in CARA allows for precise values to represent the likelihood and consequence of a risk event. By accepting any data as the input values for the risk study, whether that is a discrete value like a Likert scale or continuous data, CARA provides clear insight into the risk event. Furthermore, by allowing the user to evaluate the impacts of prevention and mitigation measures, CARA shows the user how to allocate available resources.

Since CARA considers all possible combinations of prevention and mitigation measures, the number of possible actions the user can implement grows exponentially as we consider more prevention and mitigation measures. By leveraging optimization in tandem with CARA, the user is now able to find the optimal allocation of available resources to minimize the probability and likelihood of a risk event occurring. This allows for concise decision making by showing the user which combination and order of prevention and mitigation measures to implement to reduce the overall risk of the event without requiring the user to evaluate all possible scenarios.

FUTURE WORK

Real-Time Implementation of CARA

While the use of bowtie models allows a user to predict future events given past information, it does not provide actionable information in real-time. Current implementations of bowtie models are static since they rely solely on past information. By accepting real-time data, CARA would be able to dynamically update the risk values of the events in question. In this way, CARA would be able to provide actionable information given the ever-changing environment of a real-world scenario. Furthermore, by using the optimization strategy outlined in this paper, CARA can provide the optimal real-time strategy to a decision maker.

Use of Other Distributions to Generate the Asymmetric Confidence Intervals

CARA utilizes an asymmetric Gaussian distribution centered around the nominal, or median, value of the risk statement to generate the asymmetric confidence intervals. While this distribution is very useful and easy to implement, further investigation into other asymmetric distributions is warranted. One such distribution is the metalog distribution (Keeling, 2016). This distribution is a continuous distribution that has been shown to have the ability to fit bounded, unbounded, and semi-bounded distributions with almost unlimited flexibility in its shape. While the metalog distribution is more difficult to implement (for example the number of parameters used in the distribution can range from 1 to infinity), its flexibility to fit any data potentially allows for a more robust approximation of the data being input into CARA.

REFERENCES

1. Allen, R. (2019). Adaptive Nonconvex Optimization for AI and Machine Learning, I/ITSEC paper 19109.
2. Allen, R., Engel, Z., Haney, E., (2021), Evolved AI for Model-Based Reinforcement Learning, I/ITSEC paper 21199.
3. Allen, R., Engel, Z., Volpi, M., (2021), Evolved AI for the Neural Net Enthusiast, I/ITSEC workshop ID9.
4. Bao, C., Wu, D., Wan, J., Li, J., & Chen, J. (2017). Comparison of different methods to design risk matrices from the perspective of applicability. Procedia Computer Science, 122, 455–462.
5. Department of Defense Risk Management Guide for Defense Acquisition Programs, 7th ed. 2014. (Interim Release). Washington, D.C.: Office of the Deputy Assistant Secretary of Defense for Systems Engineering.
6. Engel, Z., Harrill, N., Velez, N., Ediger, J., Allen, R. (2023) Continuous Asymmetric Risk Analysis, I/ITSEC 23270.
7. Elmontsri, M. (2014). Review of the strengths and weaknesses of risk matrices. Journal of Risk Analysis and Crisis Response, 4(1), 49–57. https://doi.org/10.2991/jrarc.2014.4.1.6
8. Garvey, P.R. and Lansdowne, Z.F. (1998) Risk Matrix: An Approach for Identifying, Assessing, and Ranking Program Risks. Air Force Journal of Logistics, 22, 18-21.
9. Keelin, Thomas W. (2016) The Metalog Distribution. Decision Analysis, 13(4): 243-277.
10. Landell, Hanna (2016), The Risk Matrix as a Tool for Risk Analysis (Master's thesis, University of Gävle, Gävle, Sweden). Retrieved from https://www.diva-portal.org/smash/get/diva2:944825/FULLTEXT01.pdf
11. Lemmens, S. M., Lopes van Balen, V. A., Röselaers, Y. C., Scheepers, H. C., & Spaanderman, M. E. (2022). The risk matrix approach: A helpful tool weighing probability and impact when deciding on preventive and diagnostic interventions. BMC Health Services Research, 22(1). https://doi.org/10.1186/s12913-022-07484-7
12. Sutherland, H., Recchia, G., Dryhurst, S. and Freeman, A.L.J. (2022), How People Understand Risk Matrices, and How Matrix Design Can Improve their Use: Findings from Randomized Controlled Studies. Risk Analysis, 42: 1023-1041. https://doi.org/10.1111/risa.13822.
13. Tong, Y. L. (1990). The multivariate normal distribution. Springer Series in Statistics. New York: Springer-Verlag.