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 its effectiveness is limited due to the discrete nature of its metric. Currently, the Risk Reporting Matrix is built from two Likert Scales: one for the likelihood of the event, and one for the severity of its consequences. This gives the user 25 possible likelihood-consequence combinations. While these results provide concise and easy-to-understand answers, they do not fully inform the user of the "true risk level" for a given decision.

The Continuous Asymmetric Risk Assessment (CARA) solution minimizes this shortcoming by transforming the discrete risk matrix into a continuous gradient field. It is designed to provide the user with infinite combinations of likelihood and consequence that more accurately describe the risk associated with the decision in question. Furthermore, by leveraging the use of asymmetric distributions, CARA creates confidence intervals around the nominal risk value, which allows the user to see the likely outcomes and their variability.

To describe how CARA identifies risk and creates asymmetric confidence intervals, this paper will evaluate a risk scenario to demonstrate its efficiency in the creation of novel and informative risk assessment. It will also demonstrate CARA's capability to evaluate risk using objective data or Subject Matter Expert (SME) opinions via Likert scale values and/or continuous data for more flexibility. Finally, it will also demonstrate how the risk level and variability is reduced as mitigation and prevention steps are added and how this leads to more accurate and descriptive risk analysis over a continuum.

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.

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

INTRODUCTION

The Department of Defense (DoD) Risk Management Framework was first developed by the Electronics System Center, US Air Force (Garvey and Lindsdowne, 1988) as a method to easily visualize the likelihood and impact of a potential risk. Since the initial development of the risk matrix tool, it has been a mainstay in risk analysis, especially across DoD. Risk matrices have become ubiquitous due to the relative simplicity of creating and explaining them. However, their simplicity also leads to an incomplete picture that warrants improvement.

The traditional risk matrix is a 5x5 grid with two axes: one axis for the likelihood of an event, and one axis for the impact or consequence of that event occurring. It is colored green on the lower left corner, yellow across the diagonal, and red on the upper right corner. Figure 1 below shows an example of a risk matrix.

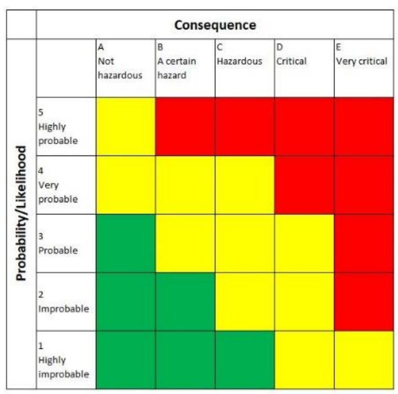

Figure 1 – A general risk matrix.

Its colors and discrete nature are easy to understand. However, because the values are discrete, the 25 possible outcomes restrict the information that can be gained from the assessment. For example, if two unrelated events map to the same discrete square, this does not mean that these two events are equally likely or that they have the same consequences. Furthermore, the risk matrix does not show the variability of potential future events.

To solve these issues, we introduce the Continuous Asymmetric Risk Analysis (CARA) solution. CARA is a continuous gradient risk assessment tool that solves the issues of a discrete tool and the lack of variation in a traditional risk matrix. CARA allows for any possible combination of likelihood and impact instead of only the 25 possible outcomes in the risk matrix. Furthermore, by using Monte Carlo simulations, CARA creates asymmetric confidence intervals around the median (i.e. nominal) value. In this way, CARA can show the possible outcomes of the event with more accuracy and in greater detail. By combining these two aspects, CARA users can account for the median value and the risk's variability, which leads to a more complete decision-analysis tool. Finally, CARA enables users to consider prevention and mitigation measures to reduce the likelihood and impact of a potential risk.

As an example, consider the likelihood and impact of a car accident. There is a baseline assessment of the risk from the common driving conditions, driver behavior, and vehicle characteristics. One preventative measure would be ensuring proper care. This will reduce the likelihood of a crash since a well-maintained car is less likely to break down on the road and will be more structurally sound than a poorly maintained vehicle. Another possible mitigation measure would be putting on a seatbelt. While this does not reduce the likelihood of the accident occurring, it will reduce the severity of any injuries should a crash occur.

CONTINUOUS ASSYMETRIC RISK ANALYSIS

Overview

To fix some of the issues with the traditional risk matrix, we propose the tool Continuous Asymmetric Risk Analysis (CARA) as an alternative method to evaluate risk. Like risk matrices, CARA is a two-dimensional assessment that evaluates a potential risk's likelihood of occurrence and impact should it occur. However, CARA differs from the risk matrix in two key ways: it is continuous rather than discrete, and it contains asymmetric two-dimensional confidence intervals. The continuous nature of CARA alleviates the restriction and ambiguity inherent in the discrete 25 choices in traditional risk matrices by allowing an infinite number of possible combinations of likelihood and impact. Furthermore, by utilizing two-dimensional confidence intervals, CARA users can evaluate the nominal value of the risk and its potential outcomes. By combining these two advantages, CARA shows not only the level likelihood and impact but also the risk's level of variability. By using true numerical values to show the variation and the median value, also called the nominal value, CARA eliminates the subjective interpretation of the traditional risk matrix. Since the traditional risk matrix only allows for 25 discrete values, two points in the same region will be interpreted as having the same risk value unless the user provides bias or their own subjective interpretation. By producing a range of values rather than a fixed region, CARA eliminates the subjective nature of the discrete regions.

Furthermore, CARA's method of data collection and simulation avoids the subjective nature of the traditional risk matrix. Typically, in the current risk matrix, the decision of where to place the point representing the risk is made by forced consensus or by a single person (usually the person in charge). Additionally, when the point has been placed, the interpretation of the risk is highly subjective since each block of the risk matrix encompasses 4% of the total risk matrix area. CARA avoids the bias of a single person by harnessing the collective input of multiple, equally-weighted subject matter experts. Furthermore, by allowing a continuous gradient of values the nominal risk can be objectively interpreted. This objectivity leads to a more succinct and informative risk decision tool.

Data Collection: Likelihood, Consequence, Mitigation, and Prevention

As with all other forms of the Risk Matrix, CARA has two axes: one to represent the likelihood of the event occurring and one to represent the impact of that event should it occur. The two axes are independent of each other. While this may seem counterintuitive, this claim has been supported by checking the Pearson's correlation coefficient. Since the Pearson correlation coefficient is always approximately zero in CARA, there is no correlation between the likelihood and consequence of the event, so we can collect the data for each variable separately without needing to consider the other.

To collect the data, CARA can use Likert Scale values (like a traditional risk matrix) or continuous inputs. The Likert scale is a five-level psychometric scale that is commonly used for quantitative human analysis but also for the collection of risk matrix data, since the five-level format easily plugs into the standard 5x5 risk matrix. Since CARA allows for any values, it can accept the same Likert scale values as the risk matrix. For example, Table 1 below shows the typical likelihood criteria used in the Department of Defense Risk Management guide to quantify the likelihood of an event.

Table 1: Typical Likelihood Criteria

Level	Likelihood	Probability of Occurrence
5	Near Certainty	> 80% to ≤ 99%
4	Highly Likely	> 60% to ≤ 80%
3	Likely	> 40% to ≤ 60%
2	Low Likelihood	> 20% to ≤ 40%
1	Not Likely	> 1% to ≤ 20%

Likert Scale values can also be used to quantify the impact of an event if it occurs. Table 2 shows a simplified version of the Consequence Criteria Likert scale values used in the Department of Defense (DoD) Risk Management Guide. The metrics in Table 2 are used by the DoD to quantify consequence.

Table 2: Sample Consequence Criteria

Level	Cost	Schedule	Performance
5 - Critical Impact	• 10% or greater increase over Acquisition Program Baseline (APB) objective values	• Schedule slip will require a major schedule rebase lining	• Degradation precludes system from meeting a Key Performance Parameters (KPP) or key technical/supportability threshold; will jeopardize program success.
4 – Significant Impact	• 5% - <10% increase over APB objective values	• Schedule deviations will slip program to within 2 months of approved APB threshold schedule date	• Degradation impairs ability to meet a (Kery System Attributes) KSA. Technical design or supportability margin exhausted in key areas
3 – Moderate Impact	• 1% - <5% increase over APB objective values	• Can meet APB objective schedule dates, but other non-APB key may slip.	• Unable to meet lower tier attributes, design or supportability margins reduced.
2 – Minor Impact	• Costs that drive unit production cost increase of <1% over budget	• Some schedule slip, but can meet APB objective dates and non-APB key event dates	• Reduced technical performance or supportability; can be tolerated with little impact on program objectives
1 – Minimal Impact	• Minimal impact. Costs expected to meet approved funding levels	• Minimal schedule impact	• Minimal consequences to meeting technical performance or supportability requirements.

While the traditional risk matrix only considers the likelihood and impact of the event, CARA allows the user to also consider the impact of prevention and mitigation measures. Prevention measures are applied to a potential risk to reduce the probability and impact of an event that could occur. In contrast, mitigation measures reduce the impact of the event if it occurs, but they do not reduce the probability of the event occurring. Similar to the collection of data for likelihood and impact, CARA can also collect prevention and mitigation data using a Likert Scale or a continuous value. By investigating both the likelihood and impact of the prevention and mitigation measures, CARA can provide a more succinct and informative decision-making tool than the risk matrix.

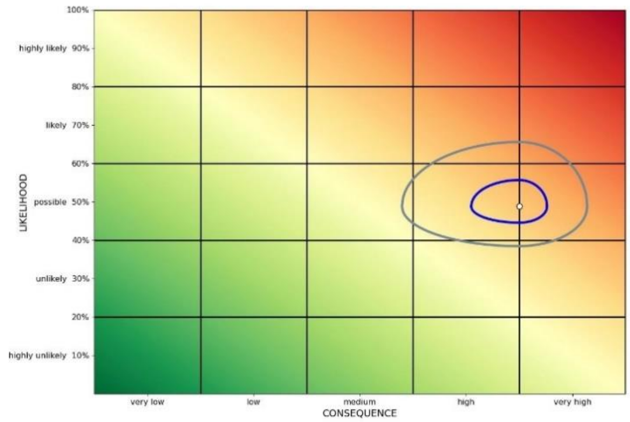

Figure 2 - Example of the output CARA.

Asymmetric Distribution

One of CARA's defining features is the two-dimensional confidence interval, which allows users to visually analyze the variability of a possible risk. The authors considered using a bivariate normal distribution, but the normal distribution was too symmetrical for the kinds of data being analyzed . For data close to the edges of the CARA plot, a symmetric confidence interval could exceed the bounds of the CARA plot. Furthermore, responses using a Likert scale format or any other discrete polling method are rarely symmetric.

To solve the issues with the bivariate normal distribution, the authors instead employed a bivariate asymmetric Gaussian distribution. By allowing asymmetry with respect to the median, CARA can more accurately fit a distribution to data that has asymmetry. To create these distributions, CARA finds the median, 10^{th} and 90^{th} percentile values for both consequence and likelihood and fits a quantile-parameterized distribution (QPD) to the three values. By combining the two QPDs, CARA can create the two-dimensional confidence intervals around the median value. The two asymmetric Gaussian distributions used to create the confidence intervals (one for the likelihood and one for the impact) are independent of each other, which precludes the need for a covariance matrix. This independence has been tested using the Spearman correlation coefficient, which proved the independence of the two dimensions.

Monte Carlo Simulation for Data Generation

Once the data is collected, regardless of format, CARA uses Monte Carlo simulations to generate data for the asymmetric confidence intervals. By using Monte Carlo simulations, CARA can reflect the uncertainty inherent in the decision-making process. The values used in the DoD Likert Scale values all have uncertainty, which can be captured by using Monte Carlo simulations. Furthermore, the use of Monte Carlo simulations allows us to investigate the how the prevention and mitigation measures impact the likelihood and consequence of a risk statement. Using the Monte Carlo simulations, CARA can evaluate thousands of possible scenarios for the potential risk, thus providing more insight into the risk statement. The outcomes of all these possible scenarios are aggregated to create a single CARA plot. This generated plot is a much more informative risk assessment than a single person deciding where a point belongs on a risk matrix. Although thousands of possible scenarios are evaluated using Monte Carlo simulations, CARA aggregates the thousands of trials and condenses them into a single CARA graph.

Generation of Median Values Asymmetric Confidence Intervals

Once the Monte Carlo simulations are completed, CARA aggregates the data in order to generate the nominal risk value and the asymmetric likelihood regions. Figure 2 below displays an example of the output of CARA after the Monte Carlo simulations have been completed.

The grid shown in Figure 2 does not mean the CARA plots are discreet. Rather, this grid is intended to suggest the CARA plots' similarity to the original risk matrix. By overlaying a grid over the plot, those who have used the traditional risk matrix are more comfortable with the new tool. Furthermore, the color values in the background correspond to the colors of the risk matrix, except they are gradated to represent continuous values:

- The bottom left corner is green, representing an event with a low level of risk.
- The upper right corner of the plot is red, representing events with a higher degree of risk.
- The white dot located inside the two asymmetric boundaries represents the nominal value of the risk statement, which is calculated by finding the median value of likelihood and consequence.

To find the two asymmetric bounds, CARA finds the 10^{th}, 30^{th}, 70^{th}, and 90^{th} percentile values for the likelihood and consequence based on the results of the Monte Carlo iterations. By combining the 30^{th} and 70^{th} percentile values for both dimensions, CARA creates the interior region, which represents the 40% confidence interval about the median. Similarly, CARA can create the 80% confidence interval about the median by combining the 10^{th} and 90^{th} percentile for both dimensions. By showing the confidence intervals and the median value, the CARA users can see the most likely value and the variation around that value.

APPLICATIONS

The main purpose of CARA is to evaluate risk in a similar (and compatible) manner as the current protocol used by the DoD. However, the ultimate goal of CARA is to create a decision-making tool for risk analysis that is more informative, objective, and flexible than the current risk matrix approach. For example, CARA can also be applied as a visual representation of a bowtie analysis. Figure 3 below demonstrates a graphical representation of a bowtie model.

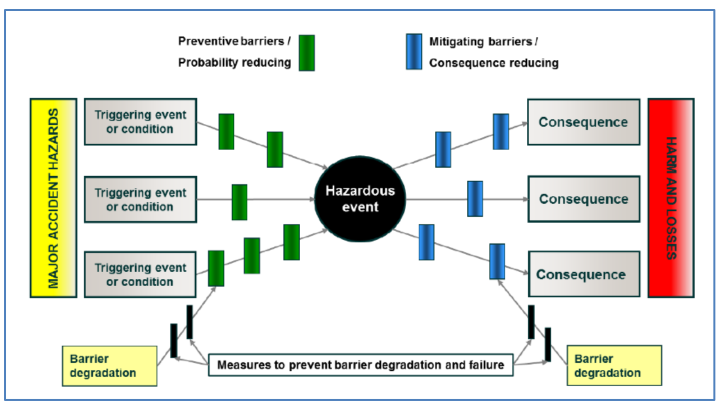

Figure 3 – Graphical representation of a bowtie model.

A bowtie model is used to assess the likelihood and impact of a risk while taking prevention and mitigation measures into consideration. Working from left to right, the bowtie model starts by analyzing a possible risk event. Then, prevention measures are introduced to evaluate the reduction in the likelihood of an event occurring. The likelihood and impact of the event is then analyzed. After the initial evaluation of the impact of the risk, a bowtie model will then apply prevention and mitigation measures to reduce the impact of the event should the event occur. Finally, the overall likelihood and impact of the event is calculated.

CARA accomplishes this similar methodology to evaluate the risk level of a possible event. Unlike a traditional bowtie model, however, CARA provides a visual representation of the risk's likelihood and consequence instead of raw numbers, though those numbers are readily available. This provides a more easily digestible risk analysis tool than the bowtie model. Furthermore, CARA can analyze all possible combinations of mitigation and prevention measures in any order the user desires. CARA leverages Microsoft Power BI, enabling a user-friendly customer interface that facilitates real-time "what-if" scenario generation. Along with this, CARA provides the capability to filter through a comprehensive list of preventative measures and risk mitigation strategies. Outputs are delivered as comprehensive clouds of uncertainty (both likelihood and consequence) plotted on an integrative DoD risk reporting matrix.

ANALYSIS AND VERIFICATION

To demonstrate the capabilities of CARA, the authors conducted a risk analysis study for a company to evaluate the risk associated with a transition from a direct to organic workforce. Like the DoD risk analysis outline in the Risk Management Guide, the authors evaluated the risks associated with cost, performance, and schedule while considering possible prevention and mitigation measures for each risk along with the likelihood of the event occurring. This evaluation was used to determine the risk associated with transitioning to new employees, which in turn was used to decide if the benefits of a new workforce outweighed the inherent risk of a full workforce transition.

To accomplish this, the authors began by evaluating over forty possible risks corresponding to cost, performance, or schedule along with possible prevention and mitigation strategies. To collect the data, multiple Subject Matter Experts (SMEs) with over 200 years of combined experience in the field were polled using a five-point Likert scale. All SMEs were asked to give their professional opinion on:

- The likelihood of an event occurring.
- The likelihood of a prevention or mitigation step being applied to the risk assessment.
- How that prevention or mitigation step would reduce the likelihood or impact of the event.

While a degree of bias is inherent in any polling-based data collection, CARA reduces this bias by considering thousands of combinations of different SMEs' input, each of which is valued equally. Furthermore, CARA does not evaluate a scenario based on one SME's assessment; rather, it takes the opinion of different SMEs for each risk, mitigation, and prevention and then aggregates the combined responses. Once the data was collected, the authors used Monte Carlo simulations to evaluate the risk of each risk statement. Over 5,000 possible strategies were evaluated for the risk assessment. Once all the risks were evaluated, the plots organized such that any combination of risks, preventions, and mitigations could be evaluated. Figures 4, 5, 6, and 7 display plots generated from this workforce transition analysis.

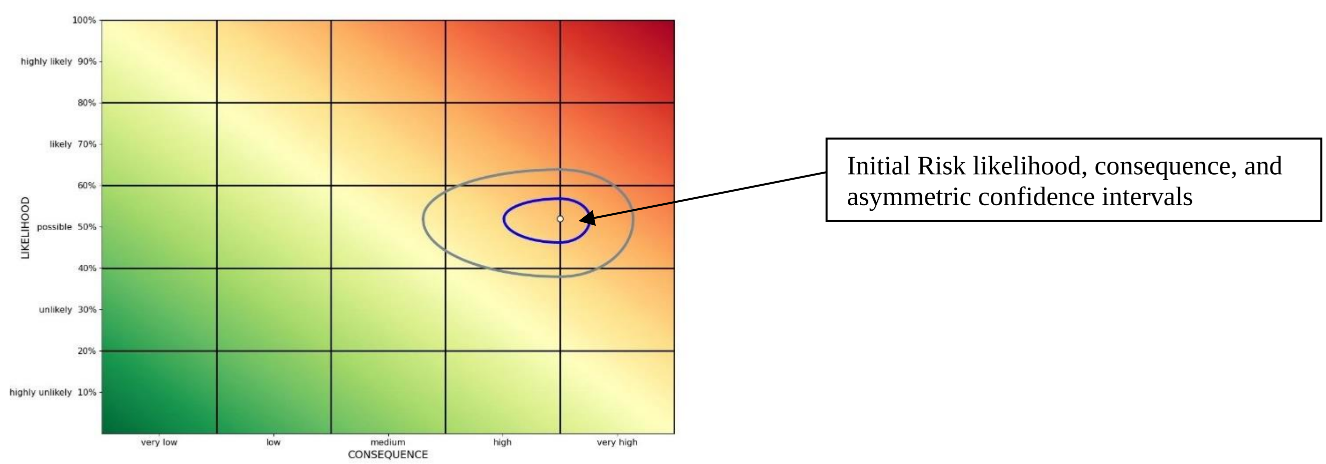

Figure 4 – An initial risk statement showing the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied.

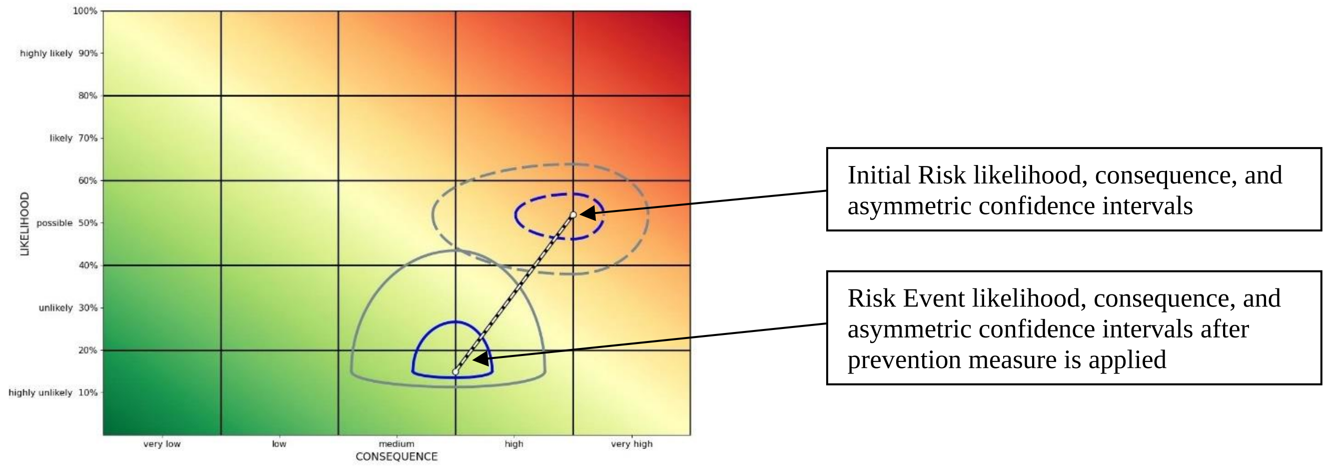

Figure 5 – An initial risk statement showing (1) the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied and (2) a secondary plot showing the application of a single prevention measure.

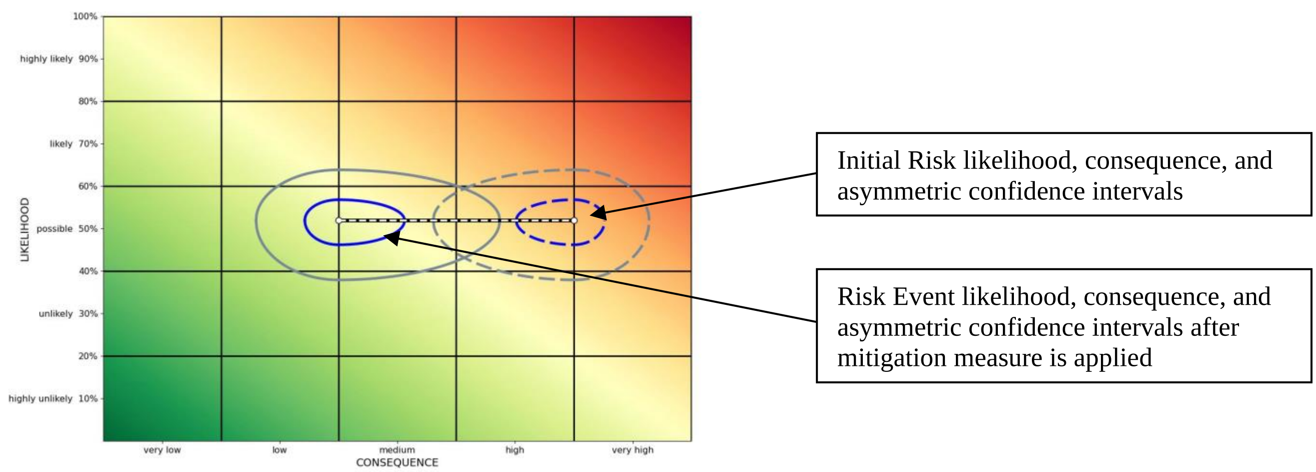

Figure 6 – An initial risk statement showing (1) the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied and (2) a secondary plot showing the application of a single mitigation measure.

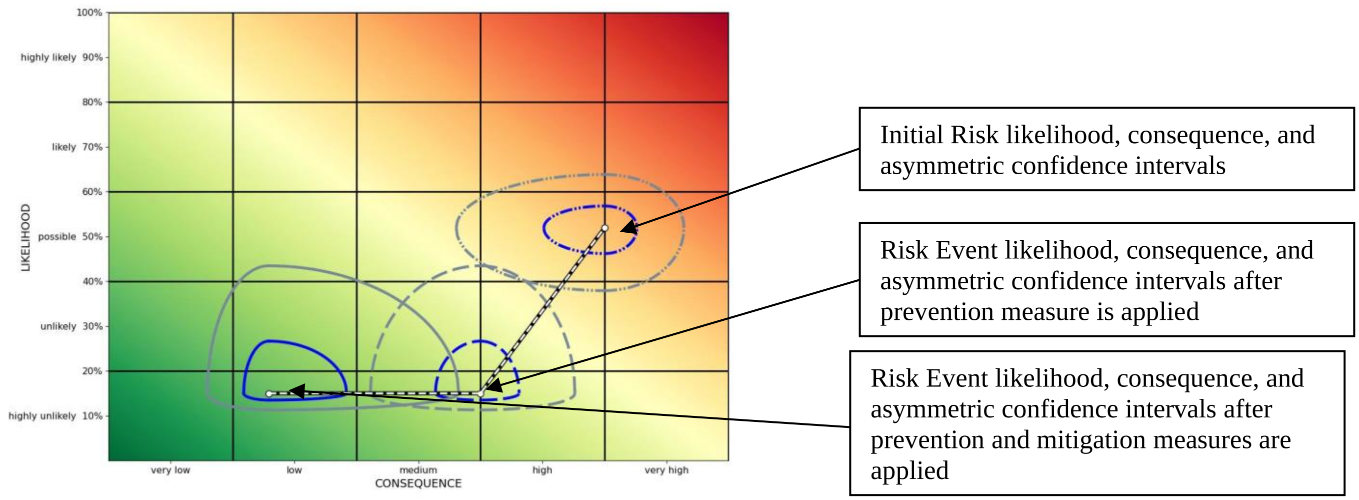

Figure 7 – An initial risk statement showing (1) the probability, likelihood, and variation of the risk with no prevention or mitigation measures applied, (2) a secondary plot showing the application of a single mitigation measure, and (3) a tertiary plot showing the combination of the same prevention and mitigation measures from Figures 5 and 6.

As we can see in Figures 4-7, the values of the likelihood and consequence change when a prevention and/or mitigation measure is applied to the initial risk statement. Using these plots, the user of CARA can quickly visually analyze risks and how preventions and mitigations can change the values. Furthermore, CARA retains the previous risks and shows how they change when new measures are applied to the risk statement. As more prevention measures are applied, the previous steps will slowly fade away into the background, so the user of CARA can clearly see the most recent step but also reference previous steps as needed.

Figure 5 clearly shows a reduction in both likelihood and consequence when a prevention measure is applied. In Figure 6, only consequence is reduced when a mitigation measure is applied. However, in both cases, we can see that while the application of a mitigation or prevention measure increases the variability of likelihood and consequence of the risk statement in question, even as it reduces the nominal values of the risk. By examining the possible outcomes, the user can decide if the reduction in the nominal value is acceptable given the increase of the variation.

CONCLUSIONS

Continuous Asymmetric Risk Analysis (CARA) is an easily understandable alternative to the traditional risk analysis matrix, and it offers clear advantages due to its continuous nature. Furthermore, CARA accepts discrete data (such as Likert scale values) or continuous data, and it allows for a clear visual analysis of a bowtie model. By leveraging the use of Monte Carlo simulations, CARA can show the risk value and asymmetric confidence intervals around the nominal value, both of which are novel approaches that demonstrate clear advantages over the current risk matrix.

By considering thousands of possible scenarios (through using of Monte Carlo simulations and evaluating all possible combinations of risks with preventions and mitigations), CARA provides an in-depth decision analysis. CARA users can evaluate all risks and select a scenario that reduces the nominal value of the risk and its variation. By being able to see all possible combinations of the preventions and mitigation measures, users can consider the time, effort, and money required and select the optimal strategy for their situation.

Through all these calculations, CARA makes risk management easier to perform, more objective, and more easily defendable. By allowing for any value in the gradient, as opposed to the discrete matrix traditionally used, CARA removes the subjectivity of interpretations inherent in the risk matrix. Additionally, CARA removes most polling bias through Monte Carlo analysis and an aggregation of multiple SME inputs. By showing the user of CARA not only a nominal value of the risk, but also confidence intervals to show the likely outcomes of the risk, the user of CARA can also consider the variability of the event. Through all of these advantages, CARA leads to more confident, accurate, and objective risk analysis.

FUTURE WORK

Real-Time Implementation of CARA

A key disadvantage of bowtie analysis models is their static nature. Bowtie analyses are based on the current projection of the likelihood and impact as well as the effect of a prevention and/or mitigation step. Most bowtie analysis are used to predict the impact of future events, which by their very nature are constantly changing. A real-time implementation of CARA can help address this issue. By accepting real-time data and by dynamically updating the effects of the impact, preventions, mitigations, and the likelihood of the event, CARA circumvents this issue entirely. For example, if we evaluate the risk of a war occurring, the likelihood and impact of that war occurring is constantly changing based on current global events. Furthermore, the possible prevention and mitigation steps that can be implemented to reduce the likelihood and impact of the war will change along with current technologies, strategies, leaderships, geopolitical atmospheres, and more. Complex future events like these would be much easier to evaluate and respond to with a graphic representation of a bowtie analysis that updates in real-time.

Integrate CARA with Artificial Intelligence to Optimize Decision Strategies

In conjunction with a real-time implementation, CARA could be even more effective if it were integrated with artificial intelligence (AI). By using AI, CARA could learn and apply the risk preferences of its users. For example, does the user value a lower nominal value of the risk’s nominal with a higher degree of uncertainty, or would they value a higher nominal value with less variability? Furthermore, based on the information that the artificial intelligence learns about the user, it could optimize a strategy based on factors such as cost, time, nominal risk values, and variability of the risk.

Use of Other Distributions to Generate the Asymmetric Confidence Intervals

As discussed previously, CARA uses an asymmetric Gaussian distribution to generate the confidence intervals around the median value of the risk. However, the use of other distributions to create these confidence intervals warrants further investigation. By using other distributions, CARA may be able to create asymmetric confidence intervals that model the data more accurately than the asymmetric gaussian distribution. One such distribution is the metalog distribution (Keelin, 2016). The metalog distribution is a continuous probability distribution that allows for bounded, unbounded, and semi-bounded distributions with virtually unlimited shape flexibility. Like a Taylor series, the metalog distribution can be fit to almost any dataset by using as many terms as necessary to fit the data.

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. 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
7. 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.
8. Keelin, Thomas W. (2016) The Metalog Distribution. Decision Analysis, 13(4): 243-277.

9. 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
10. 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
11. 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.
12. Tong, Y. L. (1990). The multivariate normal distribution. Springer Series in Statistics. New York: Springer-Verlag.