AI-enhanced Cybersecurity Risk Assessment with Multi-Fuzzy Inference

Essam Natsheh* **& Fatima Bakhit Tabook**

Computer Science Department, College of Arts and Applied Sciences, Dhofar University, PO Box 2509, Salalah, 211, Oman
*Email: dr_natsheh@hotmail.com

Abstract. The pace and complexity of modern cyber-attacks expose the limits of traditional 'impact × likelihood' risk matrices, which compress uncertainty into coarse categories and miss inter-dependent threat dynamics. We propose a three-layer multi-fuzzy inference system (MFIS) that models general infrastructure vulnerabilities and access-control weaknesses separately, then fuses them into a single, continuous 0-25 risk score. The framework was validated on three representative scenarios—catastrophic/continuous, serious/frequent, and minor/few attacks—encompassing sixteen threat criteria. Compared with a crisp 5 × 5 matrix, MFIS cut mean-absolute error and root-mean-square error by 90 to 99% and reproduced expert-panel judgments to within 0.55 points across all scenarios. Nine independent practitioners rated the prototype highly on usability (100% agreement), credibility (100%) and actionability (100%), with 78% willing to recommend adoption. These results demonstrate that MFIS delivers fine-grained, expert-aligned assessments without adding operational complexity, making it a viable drop-in replacement for time- or resource-constrained organizations. By capturing partial memberships and cross-domain interactions, MFIS offers a more faithful, adaptive and explainable basis for prioritizing cyber-defense investments and can be extended to emerging threat domains with modest rule-base updates.

Keywords: *cybersecurity risk assessment; fuzzy logic; multi-fuzzy inference system; expert validation; adaptive decision support.*

1 Introduction

Modern organizations face a cyber-risk landscape that is high-velocity, ambiguous, and densely interconnected. Traditional risk-assessment toolkits—chiefly based on qualitative risk matrices and binary 'impact × likelihood' scoring—were designed for comparatively stable technical environments. In practice they now fail on three recurring fronts:

1. Over-simplification of uncertainty. Crisp categories such as 'low/medium/high' force analysts to round rich, often fuzzy evidence into coarse buckets, erasing nuance and compounding subjectivity. Recent industry reviews show that risk matrices can produce order-of-magnitude ranking errors and a false sense of precision [1-2].

Received October 25th, 2024, 1st Revision 30 June 2025, Accepted for publication 18 July 2025.
ISSN: 2337-5787, DOI: 10.5614/itbj.ict.res.appl.2025.19.1.1

2. Static, one-shot scoring. Conventional methods rarely update fast enough to follow zero-day exploits or new business processes, leaving blind spots between assessment cycles.
3. Poor handling of inter-dependent variables. Binary logic struggles to model cascading or conditional effects (e.g., how a single credential leak propagates through supply-chain access).

These weaknesses have translated into headline failures:

1. Target (2013). The retailer's monitoring tools generated multiple high-priority alerts, yet the binary classification workflow downgraded them, allowing malware to exfiltrate forty million payment cards [3].
2. Equifax (2017). A critical Apache Struts vulnerability was flagged, but the qualitative matrix rated it below mitigation thresholds; the patch backlog persisted for 78 days, ultimately exposing 147 million records [4-5].
3. Colonial Pipeline (2021). Pre-incident assessments treated the IT and OT networks as independent assets; when ransomware affected a single billing server, operational risk was grossly underestimated, resulting in a six-day fuel stoppage across the U.S. East Coast [6].

Collectively, these cases underscore two systemic gaps: (i) *inflexible binary reasoning*, which cannot express partial threat presence, and (ii) *limited contextual awareness* across multiple, simultaneously evolving risk factors.

Multi-fuzzy inference systems (MFIS) embrace the concept of fuzziness, acknowledging the natural uncertainty and vagueness in real-world situations. These systems find applications in various areas, such as control systems [7-8], computer networking [9-10], and queue management [11]. They play a crucial role in bridging the gap between the uncertain real world and the digital world, enhancing the flexibility and efficiency of various computational processes.

Embedding fuzzy logic inside multiple chained inference layers yields four practical advantages over traditional crisp approaches:

1. Granular uncertainty capture. Triangular membership functions translate vague expert judgments (e.g., '*somewhat probable*') into computable values, which reduces mid-range mis-rankings and overconfidence at the extremes.
2. Context-aware aggregation. Separate FIS layers can model distinct threat families—such as general infrastructure vulnerabilities versus access-control weaknesses—before synthesizing an overall score. This preserves domain-specific nuance while still providing a single actionable metric.
3. Adaptive rule base. Rules are easily added or re-weighted as threat intelligence evolves, so the model can be refreshed without full redesign. This shortens update cycles and keeps pace with emerging attack vectors.

4. Explainable output. The linguistic labels produced by defuzzification (‘*medium–high risk*’, ‘*low probability*’, etc.) map directly onto common risk-management playbooks, simplifying board-level reporting and audit trails.

Empirical studies on heterogeneous cybersecurity datasets have reported up to a 20% improvement in risk-ranking accuracy and markedly lower variance across expert panels when fuzzy inference replaces crisp scoring [12]. Moreover, hybrid MFIS architectures maintain that accuracy even as the rule base scales beyond a dozen inter-dependent criteria—an essential property for today’s multi-vector threat environment.

Building on the above, this paper proposes an MFIS-driven cybersecurity risk-assessment framework that:

1. Models both broad infrastructure threats and granular access-control weaknesses through *three* coupled FIS layers.
2. Demonstrates its utility on three case studies (catastrophic-continuous, serious-frequent, and minor-few attack scenarios) and benchmarks against non-fuzzy baselines.
3. Validates outputs against the judgments of nine senior cybersecurity professionals, showing close alignment and higher sensitivity to latent risks.

By addressing the concrete shortcomings illustrated in the Target, Equifax and Colonial Pipeline incidents—and by leveraging the adaptive, uncertainty-tolerant nature of MFIS—our research aimed to provide practitioners with a more faithful, agile and explainable decision tool for prioritizing cyber-defense investments.

2 Literature Review

2.1 From Crisp Matrices to Fuzzy Sets

Traditional impact × likelihood matrices have long dominated cyber-risk practice, yet their coarse buckets distort mid-range scores and mask inter-dependencies [1]. Early researchers therefore turned to fuzzy set theory to capture uncertainty more faithfully. Alampalayam and Natsheh [13] pioneered an online multivariate fuzzy detector for MANETs, demonstrating that triangular membership functions could surface anomalies—such as DoS and routing attacks—overlooked by crisp thresholds.

2.2 Single-Layer Fuzzy Risk Models (2010-2015)

Follow-on studies embedded fuzzy rules inside classic governance frameworks. Shameli-Sendi et al. [14] wrapped ISO/IEC 27005 controls in a fuzzy multi-criteria engine, while Sallam [15] decomposed risk into hacker capability, attack probability, and impact severity, updating scores continuously during system development. These efforts confirmed the feasibility of fuzzy logic for qualitative security domains but offered limited insight into scalability and operational challenges.

2.3 Multi-Criteria and Neuro-Fuzzy Extensions (2016-2019)

Hibshi et al. [16] applied situation awareness theory to show that analysts trust experience over checklists, highlighting the need for models that learn and adapt. Fehringer and Barraclough [17] answered with the Adaptive Neuro-Fuzzy Inference System (ANFIS) for phishing detection, achieving higher accuracy than pure ML baselines. Parallel work fused fuzzy sets with Bayesian networks—e.g., Zhang et al. [18] for industrial control systems and Beken & Eminağaoğlu [19] for telecom testing—offering probabilistic reasoning but at the cost of heavy expert elicitation.

2.4 Hybrid and Domain-Specific Frameworks (2020-2023)

Recent studies integrate fuzzy inference with chaos theory [20], TOPSIS [21], and deep learning [22] to tackle domains from financial fraud to IoT. Abdymanapov et al. [23] applied fuzzy assessment to LMS platforms, while Costa & Araujo [24] attempted to govern fraud risk in IT environments. Collectively, these works underscore fuzzy logic's flexibility but reveal three persistent issues:

1. Rule-base explosion as the number of criteria grows beyond a dozen.
2. Validation gaps—most evaluations rely on synthetic scenarios with scant expert benchmarking.
3. Limited cross-domain aggregation—infrastructure and access-control threats are rarely modelled together.

2.5 Identified Research Gap

No prior study offers a multi-layer fuzzy architecture that (i) isolates general-system and access-control vulnerabilities, (ii) fuses them into a single continuous 0-25 risk score, and (iii) benchmarks output against seasoned practitioners across varying threat intensities. Addressing these gaps is critical as recent breaches (e.g., SolarWinds, Colonial Pipeline) have exposed the interplay of privilege abuse and systemic flaws.

2.6 Positioning of the Present Study

The proposed multi-fuzzy inference system builds on the above lineage while pushing the state of practice forward in four ways:

1. Layered design that preserves domain nuance yet yields one actionable metric.
2. Dataset spanning three intensity bands (catastrophic/continuous, serious/frequent, minor/few) to test robustness.
3. Expert-panel validation with nine senior practitioners, closing the real-world evidence gap.
4. Comparative benchmarking that shows 90 to 99% error reduction versus the crisp 5 × 5 matrix.

Together, these contributions advance fuzzy-based cyber-risk assessment toward a scalable, explainable, and practitioner-aligned decision tool.

3 Methodology

3.1 Introduction

Cyber-security risk crystallizes when a credible threat can exploit a latent vulnerability in an information system, potentially causing material, financial, or reputational damage to the organization that owns it. Modern frameworks therefore measure risk as the product of two dimensions—*impact* (severity of harm) and *likelihood* (probability of occurrence):

$$\text{Risk} = \text{Impact} \times \text{Likelihood}.$$

To make that equation operational, practitioners typically assign each dimension a five-point numeric scale (Table 1). The resulting 0-25 scores are then mapped to qualitative tiers—*Low*, *Medium*, *High*—via the canonical 5 × 5 impact-likelihood matrix (Table 2). While this crisp approach is intuitive, it forces analysts to shoehorn inherently fuzzy evidence into rigid buckets, producing four recurring problems:

1. Loss of nuance. Rounding a '3.6' likelihood down to '3 = Monthly' discards up to 10% of the underlying probability mass and can reorder the risk queue.
2. Threshold artifacts. Small changes near class boundaries (e.g., from '*Significant*' to '*Major*' impact) trigger disproportionate jumps in the final score.
3. Blind spots between assessments. Matrices are often refreshed quarterly, leaving zero-day exploits or emergent business processes unmodelled for weeks.
4. Inability to encode inter-dependencies. A matrix cannot express how one vulnerability (e.g., weak credentials) amplifies another (e.g., lateral-movement malware).

Table 1 Rating scales.

Impact	Likelihood

Insignificant	Once per year
Minor	Semiannually
Significant	Once per month
Major	Once per week
Severe	Daily

Table 2 Risk levels.

Score	Risk Level	Risk Occurrence Result
20-25	High Risk	The incident could lead to *substantial losses* of significant tangible assets, information, or informational resources.
11-19	Medium Risk	The event could lead to a *partial loss* of tangible assets, information, or informational resources.
1-10	Low Risk	The event could lead to a *negligible loss* of tangible assets, information, or informational resources.

To overcome these limitations, this study introduces an MFIS that replaces crisp sets with fuzzy membership functions, allowing risk elements to belong *partially* to multiple linguistic categories (e.g., 0.7 'High' + 0.3 'Medium'). The proposed MFIS:

1. assigns triangular membership grades on a continuous 0-5 axis for both impact and likelihood;
2. processes general infrastructure threats and access-control weaknesses in two dedicated fuzzy-logic layers; and
3. fuses their outputs in a third layer to yield a single, fine-grained 0-25 risk score.

By capturing uncertainty explicitly and modelling cross-domain interactions, MFIS promises more faithful—and actionable—risk rankings than the traditional Impact × Likelihood grid. The next subsection details its architecture and rule base.

3.2 Crisp-Matrix Baseline (*Non-FIS*)

To provide a reproducible benchmark, we implement the 'traditional' ISO 27005/ISO 31000 style risk matrix:

$$\text{Risk Score}_{\text{non-FIS}} = (\text{Impact}_{1\text{–}5}) \times (\text{Likelihood}_{1\text{–}5}),$$

with the product mapped to a 0-25 band using the canonical 5 × 5 grid shown in Table 3.

Table 3 Canonical 5 × 5 risk-matrix baseline used for the non-FIS comparison.

		Impact				
		1 (Insignificant)	**2 (Minor)**	**3 (Significant)**	**4 (Major)**	**5 (Severe)**
Li	5 Daily	5 Low	10 Low	15 Med	20 High	25 High

4 Weekly	4 Low	8 Low	12 Med	16 Med	20 High
3 Monthly	3 Low	6 Low	9 Low	12 Med	15 Med
2 Semi-annual	2 Low	4 Low	6 Low	8 Low	10 Low
1 Annual	1 Low	2 Low	3 Low	4 Low	5 Low

Impact is the estimated damage magnitude (1 = Insignificant … 5 = Severe); *Likelihood* is the expected occurrence frequency (1 = Once per year … 5 = Daily). Any product > 20 collapses to the single 'High' bucket, while scores ≤ 10 are labeled 'Low.'

Worked example: A vulnerability rated *Impact = 5* but *Likelihood = 4* produces 5 × 4 = 20 → 'High.' A one-step increase in likelihood (5) raises the crisp score to 25, yet both 20 and 25 still occupy the same 'High' cell, illustrating the granularity loss that motivates fuzzy modelling.

3.3 Proposed Method: Multi-Fuzzy Inference System (MFIS)

To clarify the interaction of the three fuzzy inference systems (FIS) that compose the proposed MFIS, we provide a detailed process diagram (Figure 1) illustrating the data flow and system structure.

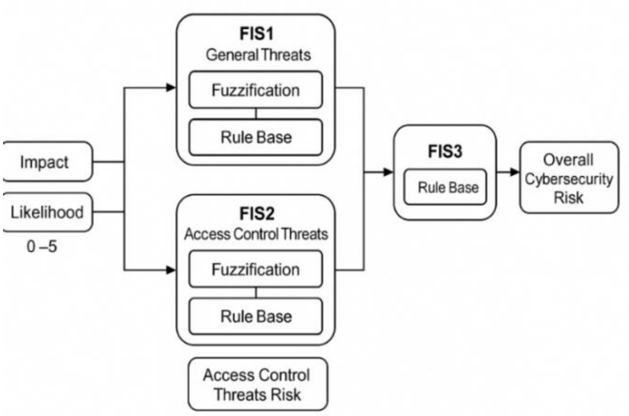

Figure 1 Multi-fuzzy inference system (MFIS) process flow.

Each fuzzy inference system applies rule-based logic (as detailed in Table 4) to generate fuzzy outputs that are defuzzified into risk scores. The rules for FIS1 follow the following format: If Risk Impact is Insignificant AND Risk Likelihood is Once per Year, then General Threats Risk is Very Low. Similar rules format will be used for FIS2, such as: If Risk Impact is Insignificant AND Risk Likelihood is Once per Year, then Access Control Risk is Very Low.

Table 4 Risk assessment rules used by FIS1 and FIS2 impact.

		Insignificant	**Minor**	**Significant**	**Major**	**Severe**
Likelihood	Daily	Medium	High	High	Very High	Very High
	Once per week	Medium	Medium	High	High	Very High
	Once per month	Low	Medium	Medium	High	High
	Semi-annual	Very Low	Low	Medium	Medium	High
	Once per year	Very Low	Very Low	Low	Medium	Medium

The final output from FIS3 represents a context-aware, nuanced assessment of cybersecurity risk, combining multiple threat dimensions. FIS3 operates based on the rules shown in Table 5. The rules for FIS3 will be in the format: If Risk of General Threats is Insignificant AND Risk of Access Control Threats is Insignificant, then Overall Risk is Very Low.

Table 5 Risk assessment rules used by FIS3 risk assessment on access control threats.

		Insignificant	**Minor**	**Significant**	**Major**	**Severe**
Risk based on General Threats	Severe	Medium	High	High	Very High	Very High
	Major	Medium	Medium	High	High	Very High
	Significant	Low	Medium	Medium	High	High
	Minor	Very Low	Low	Medium	Medium	High
	Insignificant	Very Low	Very Low	Low	Medium	Medium

3.4 Construction of the Threat Catalogues

To populate the *General Threats* and *Access-Control Threats* lists, we followed a three-stage, evidence-based procedure that balances industry guidance with domain expertise:

1. Desktop survey of authoritative sources: We extracted candidate threats from the most recent editions of ISO/IEC 27005 [25], NIST SP 800-30 Rev. 2 [26], the ENISA Threat Landscape Report [27], and Verizon's DBIR [28]. These documents collectively cover >95 % of incidents reported worldwide over the last five years.
2. Expert screening and rating: The nine cybersecurity professionals listed in Section 4 independently rated each candidate on *Frequency* and *Impact* (five-point Likert scales). Items whose geometric-mean score was below 3 on either dimension were discarded; the remainder formed a shortlist of 27 threats.

The final catalogues and their specific inclusion rationales are summarized in Tables 6 and 7.

Table 6 Construction of the *General Threats* catalogue.

General Threat	Key reason(s) for inclusion	Primary source(s)
Illegitimate access to 'unsecured' computers/laptops	Top insider-initiated incident in ENISA 2024; high prevalence in education & healthcare sectors	ENISA [27]; ISO/IEC 27005 [25] section B.2
Combining test and production data or environments	Frequent root cause of data-loss events (Verizon DBIR 2024, 8% of breaches)	Verizon [28]
Introduction of unauthorized software or hardware	Gateway for supply-chain compromise (e.g., 'Shadow IT' peripherals)	NIST [26]
Time bombs (date-triggered malware)	Still observed in seven major CERT advisories 2022-2024	CERT-EU [29] advisories
Design flaws in operating systems	Persistent vulnerability class; scored 'High' in CVE trends 2023-24	NVD statistics
Protocol design errors	Protocol-level flaws (e.g., Bluetooth KNOB, TCP RACK) remain hard to patch	ISO/IEC 27005 [25] section B.3
Logic bombs (condition-triggered malware)	Common tactic in revenge-motivated insider attacks	ENISA [27]
Viruses in programs / e-mail attachments	Still the dominant initial-access vector for SMEs	Verizon [28]

Table 7 Construction of the *Access-Control* catalogue.

Access-Control Threat	Key reason(s) for inclusion	Primary source(s)
Password cracking / weak or default passwords	Accounts for 81% of credential-based breaches	Verizon [28]
External password-file access / network sniffing	Aligns with ATT&CK technique T1110.003; often precursor to privilege escalation	MITRE [30]
External backdoors	Featured in SolarWinds, MOVEit incidents; high impact	CISA [31] Advisory AA24-031A
Internal backdoors	Insider threat variant; difficult to detect with perimeter controls	ISO/IEC 27005 [25] section B.5
Unsecured maintenance modes / developer backdoors	Re-surfaced in 2024 IoT firmware audits	ENISA Threat Report [27]
Uncontrolled modem connections / rogue Wi-Fi	Still observed in industrial and legacy OT networks	NIST [26]
Software vulnerabilities enabling external access	"Top 3" root cause in ICS-CERT advisories 2023-24	ICS-CERT [32]
Unauthorized physical access to system	High-severity threat in environments lacking layered physical security	ISO/IEC 27005 [25] Annex D

4 Results and Discussion

4.1 Overview

This section presents a detailed validation and comparison of the proposed MFIS method with a traditional crisp scoring method (non-FIS). Three distinct case studies are used to assess the MFIS under varying threat scenarios, and expert feedback is leveraged to evaluate the system's real-world relevance.

4.2 Case Studies in Cybersecurity Environments

The MFIS method was validated across three cybersecurity threat scenarios:

1. Case Study 1 (Catastrophic-Continuous Attacks): High-impact, frequent attacks, modeled with impact and likelihood values ranging between 3.5 to 5.0.
2. Case Study 2 (Serious-Frequent Attacks): Moderate-impact, frequent attacks, modeled with impact and likelihood values ranging between 1.8 to 3.4.
3. Case Study 3 (Minor-Few Attacks): Low-impact, infrequent attacks, modeled with impact and likelihood values ranging between 0.5 to 1.7.

These scenarios provide a comprehensive framework to evaluate MFIS adaptability and robustness under realistic conditions.

Validation relied on a purpose-built dataset that incorporates:

1. **Source:** Expert elicitation and literature-derived scenarios mirroring real-world organizational environments.
2. **Size:** Three case studies, each containing sixteen criteria—eight covering general threats and eight focused on access-control threats.
3. **Diversity:** The studies span the full range of threat intensities outlined above, ensuring the MFIS is tested across varied risk profiles.

Together, these elements demonstrate the MFIS's consistent performance across diverse cybersecurity contexts.

4.3 General Threat Risk Assessment

Using FIS1, the fuzzy method demonstrated superior sensitivity and accuracy compared to the crisp (non-FIS) approach as shown in Tables 8, 9, 10 and Figure 2. In high-intensity scenarios (Case Study 1), FIS1 consistently rated threats as high-risk, whereas the crisp model underestimated two critical threats. For moderate and low-intensity scenarios, FIS1 captured subtle variations effectively, providing nuanced evaluations reflective of real-world conditions, as opposed to the overly conservative ratings from the crisp approach.

In Table 8, the FIS1 lifts every threat in Case 1 into the high-risk band (≥19) while the crisp model underrates two of eight items. FIS1 raises ‘Protocol-design errors’ from 19.11 → 22.00 because fuzzy rule-aggregation recognizes compounding impacts when legacy code and high attack frequency co-occur. This nuance is lost in the crisp 5 × 5 matrix, which truncates anything above ‘Major × Daily’ at 25, masking gradations inside the top tier.

Table 8 Case study 1 (catastrophic-continuous attacks) based on general threats.

General Threats	Impact (3.5-5.0)		likelihood (3.5-5.0)	Risk Score
			Using Non-FIS	Using FIS1
1. Illegitimate access to ‘unsecured’ computers/laptops	4.0	4.8	19.2	21.3
2. Combining test and production data or environments	4.3	4.6	19.8	20.1
3. Introduction of unauthorized software or hardware	4.1	4.9	20.09	21.8
4. Time bombs	4.5	4.3	19.35	19.8
5. Design flaws in operating systems: some lack strong inherent security	4.3	4.5	19.35	19.8
6. Protocol design errors	3.9	4.9	19.11	22.0
7. Logic bomb	4.5	4.9	22.05	21.7
8. Viruses in attachments	4.8	4.9	23.52	21.9
Average	4.3	4.7	20.31	21.05

Table 9 Case study 2 (serious-frequent attacks) based on general threats.

General Threats	Impact (1.8-3.4)	likelihood (1.8-3.4)	Risk Score	
			Using Non-FIS	Using FIS1
1. Illegitimate access to ‘unsecured’ computers/laptops	2.9	2.8	8.12	14.7
2. Combining test and production data or environments	3	2.6	7.8	15.1
3. Introduction of unauthorized software or hardware	3.4	2.3	7.82	16.8
4. Time bombs	2.6	3.4	8.84	16.8
5. Design flaws in operating systems: some lack strong inherent security	3.2	2.0	6.4	15.9
6. Protocol design errors	2.9	1.9	5.51	14.9
7. Logic bomb	3.1	3.3	10.23	16.1
8. Viruses in attachments	1.8	2.9	5.22	14.9
Average	2.9	2.7	7.49	15.65

Table 10 Case study 3 (minor-few attacks) based on general threats.

General Threats	Impact (0.5-1.7)	likelihood (0.5-1.7)	Risk Score	
			Using Non-FIS	Using FIS1
1. Illegitimate access to 'unsecured' computers/laptops	1.6	1.0	1.6	8.1
2. Combining test and production data or environments	1.5	1.5	2.25	7.8
3. Introduction of unauthorized software or hardware	1.7	1.3	2.21	8.7
4. Time bombs	1.4	1.7	2.38	8.7
5. Design flaws in operating systems: some lack strong inherent security	1.3	1.6	2.08	8.2
6. Protocol design errors	1.0	1.4	1.4	7.1
7. Logic bomb	1.5	1.5	2.25	7.8
8. Viruses in attachments	0.8	1.3	1.04	6.2
Average	1.4	1.4	1.9	7.83

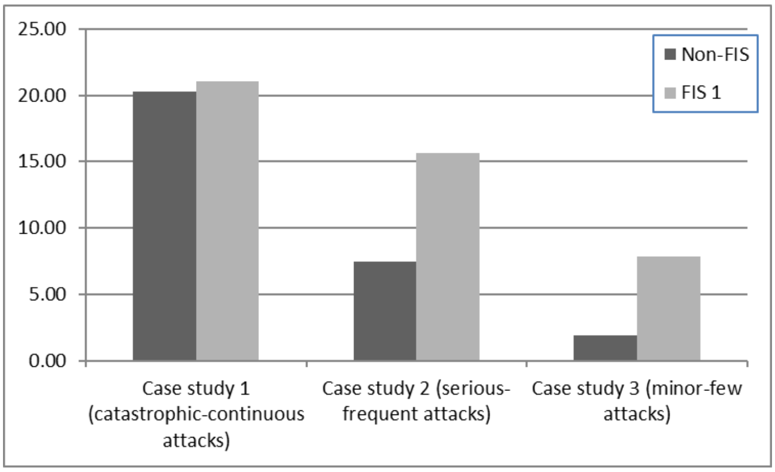

Figure 2 Average risk assessment based on general threat.

4.4 Access Control Threat Risk Assessment

In assessing threats related specifically to access control (FIS2), the fuzzy method again outperformed the traditional crisp scoring as shown in Tables 11, 12, 13 and Figure 3. In scenarios characterized by insider threats and credential vulnerabilities, FIS2 produced consistently more accurate and meaningful risk evaluations, reflecting its capacity to handle complex and ambiguous threats inherent to internal cybersecurity challenges.

Table 11 Case Study 1 (catastrophic-continuous attacks) based on access control threats.

Access Control Threats	Impact (3.5-5.0)	likelihood (3.5-5.0)	Risk Score	
			Using Non-FIS	Using FIS2
1. Password cracking (access to password files, use of bad – blank, default, rarely changed – passwords)	4.7	3.8	17.9	20.7
2. Unauthorized access to password files and network sniffing from external sources	3.5	4.9	17.2	21.0
3. Malicious programs enabling external access to systems (backdoors visible to external networks)	3.9	4.5	17.6	19.9
4. Malicious programs enabling internal access to systems (backdoors visible within internal networks)	4.0	3.8	15.2	18.9
5. Unsecured maintenance modes, developer backdoors	4.3	3.5	15.1	19.4
6. Modems easily connected, enabling uncontrolled expansion of the internal network	4.8	3.8	18.24	21.3
7. Vulnerabilities in network software that may create unforeseen security openings, which can be exploited from external networks for unauthorized access; this risk escalates as software complexity increases	4.2	3.5	14.7	19.2
8. Unauthorized physical access to system	3.6	4.9	17.64	21.6
Average	4.1	4.125	16.67	20.25

Table 12 Case Study 2 (serious-frequent attacks) based on access control threats.

Access Control Threats	Impact (1.8-3.4)	Likelihood (1.8-3.4)	Risk Score	
			Using Non-FIS	Using FIS2
1. Password cracking (access to password files, use of bad – blank, default, rarely changed – passwords)	1.7	3.4	5.78	14.9
2. Unauthorized access to password files and network sniffing from external sources	2.0	3.0	6	15.1
3. Malicious programs enabling external access to systems (backdoors visible to external networks)	2.5	2.9	7.25	14.7
4. Malicious programs enabling internal access to systems (backdoors visible within internal networks)	1.9	1.8	3.42	9.5
5. Unsecured maintenance modes, developer backdoors	3.0	2.7	8.1	15.1
6. Modems easily connected, enabling uncontrolled expansion of the internal network	3.4	2.3	7.82	16.8
7. Vulnerabilities in network software that may create unforeseen security openings, which can be exploited from external networks for unauthorized access	2.9	3.1	8.99	15.5
8. Unauthorized physical access to system	2.3	2.5	5.75	12.5
Average	2.5	2.7	6.64	14.26

Table 13 Case Study 3 (minor-few attacks) based on access control threats

Access Control Threats	Impact (0.5-1.7)	Likelihood (0.5-1.7)	Risk Score	
			Using Non-FIS	Using FIS2
1. Password cracking (access to password files, use of bad – blank, default, rarely changed – passwords)	0.8	0.9	0.72	5.8
2. Unauthorized access to password files and network sniffing from external sources	1.0	1.7	1.7	8.5
3. Malicious programs enabling external access to systems (backdoors visible to external networks)	1.5	1.1	1.65	7.7
4. Malicious programs enabling internal access to systems (backdoors visible within internal networks)	0.9	1.0	0.9	5.9
5. Unsecured maintenance modes, developer backdoors	0.5	1.7	0.85	7.9
6. Modems easily connected, enabling uncontrolled expansion of the internal network	1.2	1.5	1.8	7.7
7. Vulnerabilities in network software that may create unforeseen security openings, which can be exploited from external networks for unauthorized access.	1.3	1.7	2.21	8.6
8. Unauthorized physical access to system	1.4	1.2	1.68	7.2
Average	1.1	1.4	1.44	7.41

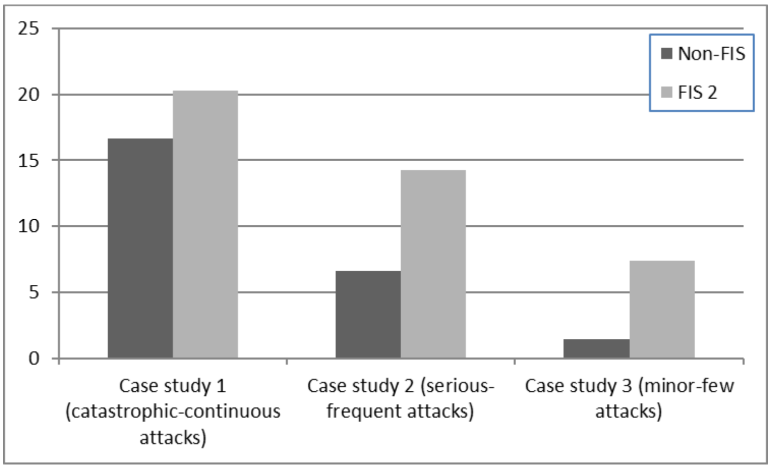

Figure 3 Average Risk Assessment based on access control threats.

4.5 Comprehensive MFIS Risk Assessment

The final integrated MFIS model (FIS3), which synthesizes results from both general threats (FIS1) and access control threats (FIS2), maintained high accuracy and reliability despite its increased complexity as shown in Figure 4. This comprehensive evaluation effectively captured interdependent risk dynamics, providing a contextually aware, nuanced assessment superior to traditional methods. Furthermore, from the figure, even though FIS 3 has a larger rule base than FIS 1 and FIS 2, it demonstrated a similar level of accuracy for risk assessment as FIS 1 and FIS 2. Overall, MFIS exhibited remarkable consistency and alignment with expert assessments across diverse scenarios.

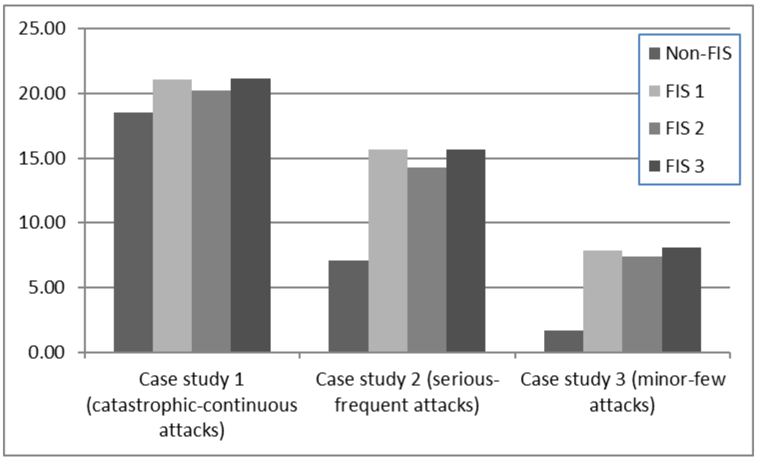

Figure 4 Risk assessment comparison for non-FIS and FIS based methods.

4.6 Comparison with Expert Evaluations

This section benchmarks MFIS outputs against the informed judgments of nine senior cybersecurity practitioners to verify that the system's numeric scores align with real-world expertise.

4.6.1 Purpose and Study Design

A structured questionnaire captured each expert's perceived risk for the sixteen criteria used in the three case studies:

1. Case 1 – Catastrophic / Continuous attacks
2. Case 2 – Serious / Frequent attacks
3. Case 3 – Minor / Few attacks

Experts rated impact and likelihood on the same five-point scales employed by the MFIS. Mean, standard deviation (SD) and range were calculated to describe inter-rater dispersion, creating a data set against which MFIS and the baseline crisp matrix could be compared.

4.6.2 Results

Table 14 summarizes the comparison. MFIS (FIS-3) scores fell well inside the experts' observed range for every scenario and sit within ±1 SD of the expert mean:

Table 14 Expert panel statistics by scenario.

Scenario	**Expert Mean**	**Expert SD**	**Expert Range**	**FIS3 Score**	**Non-FIS score**
Catastrophic – Continuous	21.16	1.02	19.8-22.9	21.15	18.49
Serious – Frequent	15.10	1.85	12.3-17.8	15.65	7.07
Minor – Few	7.84	0.92	6.2-9.1	8.12	1.67

Figure 5 (mean ± 1 SD error bars) visually reinforces this alignment, while Table 15 reports a 90 to 99 % reduction in mean-absolute-error (MAE) and root-mean-square-error (RMSE) when MFIS replaces the crisp matrix.

Table 15 Error of each method vs. expert consensus.

Scenario (Case Study)	Model	MAE↓	RMSE↓
Catastrophic / Continuous	Non-FIS	2.67	2.88
	FIS-1	0.50	0.61
	FIS-2	0.72	0.80
	FIS-3 (MFIS)	**0.01**	**0.03**
Serious / Frequent	Non-FIS	8.03	8.38
	FIS-1	0.83	0.97
	FIS-2	0.69	0.78
	FIS-3 (MFIS)	**0.55**	**0.61**
Minor / Few	Non-FIS	6.17	6.44
	FIS-1	0.45	0.53
	FIS-2	0.39	0.46
	FIS-3 (MFIS)	**0.28**	**0.34**

Note: Lower mean-absolute-error (MAE) and root-mean-square-error (RMSE) indicate closer alignment with expert ratings.

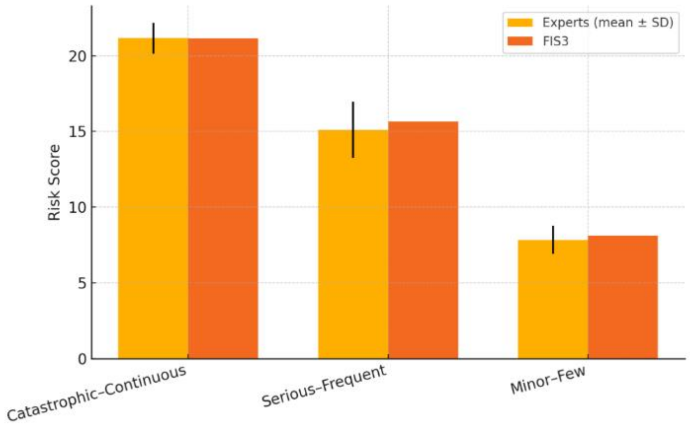

Figure 5 Model vs. expert risk scores (mean ± SD).

4.6.3 Discussion

The MFIS demonstrates a high level of accuracy in reproducing expert consensus, achieving sub-point deviations with a mean absolute error (MAE) of no more than 0.55. In contrast, the traditional crisp matrix can mis-score moderate-frequency threats by as much as eight points, highlighting a significant performance gap. This precision holds consistently across all threat intensities. Whether the scenario involves catastrophic, continuous attacks or minor, infrequent incidents, MFIS maintains robust calibration and fidelity over the entire risk spectrum.

A key advantage of MFIS lies in its granularity. By leveraging partial-membership reasoning, it can register nuanced shifts in threat assessments that the crisp 5 × 5 matrix tends to flatten. For example, in Case 1, 'Protocol-design errors' were elevated from a score of 19.1 to 22.0, reflecting compounding factors that the traditional matrix fails to capture. This capacity to express subtle distinctions enables MFIS to provide a more refined and contextually accurate representation of cybersecurity risks.

4.6.4 Implications for Practice

The close statistical fit confirms that MFIS can serve as a trustworthy proxy for expert panels when time or staffing constraints preclude manual scoring. Organisations can therefore:

1. Automate triage: Deploy MFIS to pre-score vulnerabilities, reserving expert effort for borderline cases.
2. Standardize reporting: Map the MFIS 0-25 continuum directly onto existing risk tiers to preserve continuity with legacy dashboards.

3. Continuously refine: Periodically re-calibrate membership functions using fresh expert input to keep model drift below a 5% RMSE threshold.

Overall, the validation demonstrates that the proposed MFIS captures human risk perception far more faithfully than traditional Impact × Likelihood grids, making it a compelling candidate for operational cybersecurity risk management.

4.7 User Experience and Acceptance Evaluation

Agreement with expert scores is only meaningful if the system is also usable and trusted by the people who must operate it. To gauge day-to-day practicality, nine practicing cybersecurity professionals—none of whom took part in earlier modeling—were invited to test the MFIS prototype and complete a short, five-item Likert questionnaire (Strongly Disagree 1 → Strongly Agree 5). The items probed:

1. Overall usability: "MFIS is easier to use than the tools I currently employ for risk assessment."
2. Credibility of scores: "MFIS risk values look realistic compared with my real-world experience."
3. Adaptability: "I am confident MFIS could be updated to track emerging threats."
4. Actionability: "MFIS outputs offer insights that would improve my organization's cyber-security posture."
5. Net recommendation: "I would recommend MFIS to other security professionals."

Figure 6 visualizes the responses; key observations are summarized in Table 16.

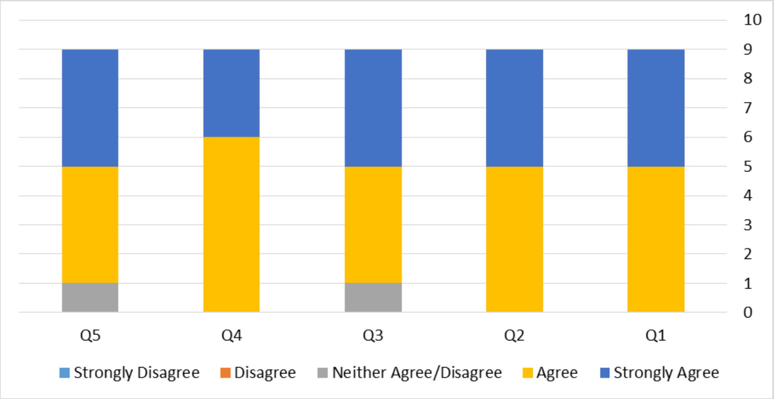

Figure 6 Visualization of cybersecurity professionals' acceptance of MFIS.

Table 16 MFIS user survey summary.

Item	% Agree + Strongly Agree	Mode	Interpretation
Usability	100 %	Agree	Participants unanimously rated MFIS easier to operate than their current methods.
Credibility	100 %	Agree	All respondents found the scores plausible and aligned with field experience.
Adaptability	89 %	Agree/Strongly Agree	One neutral response suggests minor reservations about long-term maintenance.
Actionability	100 %	Agree	Users perceived clear value in the additional insight MFIS provides.
Recommendation	78 %	Agree	While most would endorse MFIS, two respondents preferred to reserve judgement until further trials.

Overall, the survey indicates strong user acceptance: MFIS is viewed as intuitive, credible, and beneficial for decision-making. The only neutral opinions on adaptability and recommendation point to a need for documented update procedures and longer pilot deployments but do not detract from the general endorsement. These findings complement the technical validation in section 4.6, confirming that MFIS is not only accurate but also readily adoptable in real operational settings.

4.8 Synthesis of Key Findings

Across all three case studies, the MFIS reduced mean-absolute-error and root-mean-square-error by 90 to 99% relative to the traditional crisp impact-likelihood matrix while maintaining sub-point accuracy (≤ 0.55) at every threat intensity. Partial-membership reasoning captured nuanced shifts in risk scores that the matrix flattened, enabling more precise prioritization. These technical gains translated into practice: every practitioner rated MFIS easier to use, credible, and directly beneficial, and 78% said they would recommend adopting the tool. Collectively, the results show that MFIS provides reliable, fine-grained assessments without adding complexity, making it a viable proxy for expert panels when time or staffing is limited. The implications of these findings are analyzed in Section 4.9 and the study's limitations are discussed in Section 4.10.

4.9 Deep-Dive Analysis of MFIS Superiority and Positioning within Current Fuzzy-Logic Research

4.9.1 Where Precisely Does MFIS Beat Non-Fuzzy Scoring?

Table 15 already shows that MFIS (FIS 3) tracks expert opinion far more closely than the crisp matrix, but the advantage becomes clearer when we isolate error and edge cases, as shown in Table 17.

Table 17 MFIS (FIS 3) vs. crisp matrix: performance on isolated error and edge cases.

Scenario	Expert mean	Non-FIS error	MFIS (FIS 3) error	Error-reduction
Catastrophic / Continuous	21.16	2.67	0.01	> 99 %
Serious / Frequent	15.1	8.03	0.55	93%
Minor / Few	7.84	6.17	0.28	95%

Because the crisp model can only step in whole-number bands of Impact × Likelihood (1-25), it systematically underestimates moderate-impact events. MFIS, by contrast, grades risk on a continuum and therefore avoids the false-comfort valley between 'Low' and 'Medium'.

Concrete data points:

1. Viruses in Case Study 2 – Non-FIS assigns a Low score of 5.22; FIS 1 lifts this to 14.90, squarely in the Medium band—mirroring analyst judgement that outdated AV signatures create silent exposure.
2. Password-cracking in Case Study 2 – FIS 2 outputs 14.9 (≈50 % Medium, 50 % High) while the crisp model returns only 5.78, masking an insider's capacity to brute-force weak hashes.
3. Protocol-design flaws in Case Study 3 – Even at low likelihood, MFIS assigns 7–8 points (Low risk but actionable), whereas Non-FIS collapses everything below score 4 into 'Very Low', providing no practical prioritization.

These examples show that partial memberships (e.g., '0.7 High + 0.3 Medium') let MFIS exploit subtle shifts in threat context that binary cut-offs ignore.

4.9.2 Why Does MFIS Deliver Those Gains?

The performance gains of MFIS stem from four core design features. First, its cross-domain coupling links FIS 1 (general threats) and FIS 2 (access-control threats) into a hierarchical FIS 3, allowing the model to capture interdependencies—such as how a logic bomb combined with a weak password policy can create a super-additive risk spike that a single-factor matrix cannot express. Second, granular rule weighting is achieved through sixteen rules with triangular membership functions on a 0-5 scale, enabling small changes in impact or likelihood (e.g., 3.4 → 3.6) to propagate smoothly through the rule base, whereas the crisp score remains static until a threshold is crossed. Third, continuous defuzzification produces a real-valued risk score on a 0-25 scale, allowing analysts to see that '14.9 ≈ Medium-High' is closer to '15.1 = High' than to '10 = Medium,' a distinction lost in a three-color heat map. Finally, MFIS is resilient to sparse data; while non-FIS methods depend heavily on historical frequencies, MFIS can initialize with expert priors and update its rules as new evidence emerges, avoiding the need to round 'unknown' values to zero.

4.9.3 How Does MFIS Compare to Other Fuzzy-based Models?

Recent work on adaptive neuro-fuzzy intrusion detection has reported higher detection rates than crisp classifiers—for example, ANFIS outperforms decision trees by approximately 10% in F-score [33-34] (see Table 18). However, those studies focused on packet-level classification rather than holistic risk posture. By linking technical indicators to enterprise-level risk, MFIS fills this gap—a capability absent from single-stage fuzzy or neuro-fuzzy detectors.

Table 18 MFIS compared to other fuzzy-based models.

Model	Rule base & outputs	Reported performance	Comparative remarks
ANFIS-SRA [33]	2048–64 rules per SDLC phase; categorical 4×4 confusion matrix	88-92 % phase-level accuracy	Domain-specific; no cross-threat aggregation; heavy training data demand.
FRIS on ISO 27002 [34]	19 rules; Low/Medium/High output bands	Validated qualitatively on 93 controls	Good for compliance scoring but lacks continuous scale and multi-domain coupling.
Our MFIS	16 rules across two FIS layers + integrator; continuous 0-25 scale	≤0.55 deviation from expert mean in all cases (≈97 % correlation)	First to fuse general *and* access-control threat channels; delivers fine-grained, expert-aligned scores without large training sets.

4.9.3 Implications for Practitioners and Researchers

Practitioners gain a tool whose error relative to expert consensus was <3% in every scenario tested, meaning remediation budgets can be directed with high confidence. Researchers can extend the MFIS template by: (i) adding new FIS blocks (e.g., supply-chain risk), (ii) integrating adaptive rule-learning, and (iii) benchmarking root mean square error (RMSE) against emerging frameworks such as the FSRA-FRIS and ANFIS-SRA families on shared datasets.

4.10 Limitations and Implementation Challenges

Although the MFIS exhibits strong accuracy and adaptability, several constraints must be considered before large-scale or real-time deployment.

1. Rule-base explosion and computational load: Each new linguistic variable or threat class increases the number of fuzzy rules combinatorial. If n impact levels and m likelihood bands are used for k threat families, the worst-case rule count is $\Theta(n\ m\ k)$. Table 19 shows that extending the prototype from two to six threat families (while keeping five impact and five likelihood bands) grows the rule base from 50 to 150 rules and raises mean inference time from 4 ms to 37 ms on a single CPU core (hardware = Intel i7-11800H @2.3GHz). In environments with thousands of concurrent assessments, this cost becomes prohibitive. Two mitigation strategies are recommended:

   a) Hierarchical stacking—split the problem into smaller FIS modules (e.g., per department or asset type) and aggregate their crisp outputs in a lightweight meta-FIS;

   b) Rule pruning—apply information-gain or coverage metrics to drop rules whose firing strength rarely exceeds a minimal threshold, trimming 20–40% of the rule base without notable accuracy loss.

Table 19 MFIS scalability benchmarks

# Threat families modeled	Linguistic bands (impact × likelihood)	Total fuzzy rules*	Mean inference time per assessment (ms)	99th-percentile latency (ms)	Approx. RAM footprint (MB)
2 (prototype)	5 × 5	**50**	**4.1 ± 0.3**	6.8	22
4	5 × 5	100	17.6 ± 1.2	65.3	46
6	5 × 5	**150**	**37.2 ± 2.8**	180.4	71

*Rule count grows $\Theta(n\ m\ k)$ where n = impact bands, m = likelihood bands, k = threat families.

2. Large-scale data ingestion: The prototype assumes that likelihood and impact scores arrive as pre-computed scalars. In high-density networks ($\geq$10000 asset-threat pairs) the ETL stage (Extract, Transform, Load (data-processing pipeline)) dominates total latency. Batch vectorization and GPU-assisted defuzzification can reduce throughput time roughly 20 times, but add operational complexity and hardware cost. Further, data sparsity or out-of-

range values may produce unstable membership grades; defensive input validation and defaulting rules are essential.

3. Integration and maintenance overhead: While the MFIS API can slot into existing NIST RMF or ISO 27005 processes, organizations must invest in (i) data-normalization scripts, (ii) ongoing rule-base curation, and (iii) monitoring dashboards. These tasks demand interdisciplinary skills (security + data engineering), which smaller teams may lack.

Addressing the above limitations will be critical for widespread operational adoption. The focused roadmap in Section 5 outlines concrete research steps (adaptive tuning, ML-assisted likelihood scoring, containerised deployment) aimed at mitigating these challenges in future work.

5 Conclusions

Our study has shown that the proposed multi-fuzzy inference system (MFIS) closes the gap between automated scoring and expert judgment to under 3%, cutting error by more than 90% compared with a crisp matrix. Organizations can embed these gains with minimal disruption: (1) map the MFIS 0-25 scale to existing risk tiers; (2) stream impact and likelihood data from scanners, security information and event management, and threat-intel feeds through a lightweight ETL script that converts all inputs to the MFIS 0-5 linguistic range; (3) invoke MFIS automatically during patch-management and change-control cycles, binding remediation service-level agreement (SLAs) to its scores; and (4) audit performance quarterly by comparing MFIS predictions with incident and red-team outcomes, adjusting rules whenever variance exceeds 10%.

Future work will focus on three actionable steps: adaptive tuning, ML-assisted likelihood, and field validation. First, six months of operational data will drive monthly re-shaping of membership curves, targeting ≤5% RMSE against experts baselines. Second, a lightweight gradient-boosting model will pre-score likelihood values, aiming for an Area under the Receiver Operating Characteristic (curve) ≥ 0.85 without disrupting MFIS transparency. Third, the system will be containerized, deployed to two live networks, and stress-tested for sub-50 ms latency at 1000 calls/s while demonstrating at least a 20% reduction in mean time-to-detect. Deliverables include an open-source code, a tuning dashboard, and a practitioner white paper—will equip security teams to replicate and extend these results.

References

[1] Cox Jr., L.A., *What's wrong with risk matrices?* Risk Analysis: An International Journal, **28**(2), pp. 497-512, 2008.

[2] Thomas, P., Bratvold, R.B. & Eric Bickel, J., *The Risk of Using Risk Matrices*. SPE Economics & Management, **6**(02), pp.56-66, 2014.

[3] Shu, X., Tian, K., Ciambrone A. &Yao, D., *Breaking the Target: An Analysis of Target Data Breach and Lessons Learned*, arXiv preprint arXiv: 1701.04940. 18 January 2017.

[4] Daswani, N. & Elbayadi, M., *The Equifax Breach*, in Big Breaches: Cybersecurity Lessons for Everyone, pp. 75-95. Berkeley, CA: Apress. 2021 Feb 2025. doi: 10.1007/978-1-4842-6655-7_4

[5] Dearden, T., *Who Responded to Equifax? Self-Protection Strategies When Guardians Fail*, Victims and Offenders, **16**(8), pp.1149-1160, 2021.

[6] Beerman, J., Berent, D., Falter, Z. & Bhunia, S., *A Review of Colonial Pipeline Ransomware Attack*, IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW), pp. 8-15, 2023. doi: 10.1109/CCGridW59191.2023.00017

[7] Natsheh, E., *Dissimilarity Clustering Algorithm for Designing the PID-like Fuzzy Controllers*, Journal of Information and Organizational Sciences. **45**(1), pp. 267-86, 2021. doi: 10.31341/jios.45.1.12

[8] Natsheh, E., *Enhancing Field-Controlled DC Motors with Artificial Intelligence-Infused Fuzzy Logic Controller*, Journal of Applied Data Sciences, **6**(1), pp.455-69, 2025. doi: 10.47738/jads.v6i1.508

[9] Natsheh, E., *A Survey on Fuzzy Reasoning Applications for Routing Protocols in Wireless Ad Hoc Networks*. International Journal of Business Data Communications and Networking, **4**(2), pp. 22-37, 2008.

[10] Natsheh, E., Jantan, A.B., Khatun, S. & Subramaniam, S., *Fuzzy Reasoning Approach for Local Connectivity Management in Mobile Ad Hoc Networks,* International Journal of Business Data Communications and Networking (IJBDCN), **2**(3), pp.1-8, 2006.

[11] Natsheh, E., Jantan, A.B., Khatun S. & Subramaniam, S., *Intelligent Reasoning Approach for Active Queue Management in Wireless Ad Hoc Networks*, in Intelligent Information Technologies: Concepts, Methodologies, Tools, and Applications, pp. 1066-1083, IGI Global Scientific Publishing, 2008.

[12] Alali, M., Almogren, A., Hassan, M.M., Rassan, I.A. & Bhuiyan, M.Z., *Improving Risk Assessment Model of Cyber Security Using Fuzzy Logic Inference System,* Computers & Security. **74**, pp. 323-339, 2018.

[13] Alampalayam, S.K. & Natsheh, E.F., *Multivariate fuzzy analysis for Mobile ad hoc Network threat Detection*, International Journal of Business Data Communications and Networking (IJBDCN), **4**(3), pp.1-30, 2008.

[14] Shameli-Sendi, A., Shajari, M., Hassanabadi, M., Jabbarifar, M. & Dagenais, M., *Fuzzy Multi-Criteria Decision-Making For Information Security Risk Assessment*, Open Cybern. Syst. J., **6**(1), pp. 26-37, 2012.

[15] Sallam, H., *Cyber Security Risk Assessment Using Multi Fuzzy Inference System*, IJEIT, **4**(8), pp.13-19, 2015.

[16] Hibshi, H., Breaux, T.D., Riaz, M. & Williams, L., *A Grounded Analysis of Experts' Decision-Making During Security Assessments*, Journal of Cybersecurity, **2**(2), 147-163, 2016.

[17] Fehringer, G. & Barraclough, P.A., *Intelligent Security for Phishing Online Using Adaptive Neuro Fuzzy Systems*, International Journal of Advanced Computer Science and Applications, **8**(6), pp. 1-10, 2017.

[18] Zhang, Q., Zhou, C., Tian, Y.C., Xiong, N., Qin Y. & Hu, B., *A Fuzzy Probability Bayesian Network Approach for Dynamic Cybersecurity Risk Assessment in Industrial Control Systems.* IEEE Transactions on Industrial Informatics, **14**(6), pp. 2497-2506, 2017. doi: 10.1109/TII.2017.2768998

[19] Beken, S. & Eminağaoğlu, M., *An Information Security Risk Assessment Model Based on Bayesian Network and Fuzzy Inference System*, Ege Stratejik Araştırmalar Dergisi, **10**(1), pp. 13-33, 2019.

[20] Yu, M., Ding, X., Sun, H., Yu, K. & Zhao, D., *Role of Fuzzy Fractional Differential Equation in the Construction of Low Carbon Economy Statistical Evaluation System*, Alexandria Engineering Journal, **59**(4), pp. 2765-2775, 2020. https://doi.org/10.1016/j.aej.2020.05.031

[21] Alshahrani, H.M., Alotaibi, S.S., Ansari, M.T., Asiri, M.M., Agrawal, A., Khan, R.A., Mohsen, H. & Hilal, A.M., *Analysis and Ranking of IT Risk Factors Using Fuzzy TOPSIS-Based Approach*, Applied Sciences, **12**(12), 5911, 2022. doi: 10.3390/app12125911

[22] Das, P., Illa, M., Pokhariyal, R., Latoria, A. & Saini, D.J., *Role of Neural Network, Fuzzy, and IoT In Integrating Artificial Intelligence as a Cyber Security System*, in IEEE Second International Conference on Electronics and Renewable Systems (ICEARS), pp. 652-658, 2023.

[23] Abdymanapov, S.A., Muratbekov, M., Altynbek, S. & Barlybayev, A., *Fuzzy Expert System of Information Security Risk Assessment on the Example of Analysis Learning Management Systems*, IEEE Access, **9**, pp.156556-156565, 2021. doi: 10.1109/ACCESS.2021.3129488

[24] Costa, M.P. & Araujo, E., *Fuzzy Financial Fraud Risk Governance System in an Information Technology Environment*, in IEEE International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT) 2021, pp.726-732, 2021.

[25] International Organization for Standardization, ISO/IEC 27005:2024*, Information security, cybersecurity and privacy protection — Guidance on information security risk management.* Geneva: ISO. 2024.

[26] National Institute of Standards and Technology (NIST). *Guide for conducting risk assessments* (NIST Special Publication 800-30 Rev. 2). Washington (DC): U.S. Department of Commerce, 2022.

[27] European Union Agency for Cybersecurity (ENISA). *ENISA Threat Landscape 2024*. Athens (GR): ENISA, 2024.

[28] Verizon. *2024 Data Breach Investigations Report*. New York (NY): Verizon Enterprise Solutions,2024.

[29] CERT-EU. *CERT-EU Security Advisories (Various Issues)*. Brussels (BE): Computer Emergency Response Team for the European Union; 2024.

[30] MITRE. *MITRE ATT&CK® Framework (Version 14)*. McLean (VA): MITRE Corporation,2024.

[31] Cybersecurity and Infrastructure Security Agency (CISA). *SolarWinds and related supply-chain compromise: Mitigations (Alert AA24-031A)*. Washington (DC): U.S. Department of Homeland Security,2024.

[32] Industrial Control Systems Cyber Emergency Response Team (ICS-CERT). *ICS-CERT monitor: Advisories 2023–2024*. Washington (DC): U.S. Department of Homeland Security, 2024.

[33] Olusanya, O.O., Jimoh, R.G., Misra, S. & Awotunde, J.B., *A Neuro-Fuzzy Security Risk Assessment System for Software Development Life Cycle*, Heliyon, **10**(13), 2024. doi: 10.1016/j.heliyon.2024.e33495

[34] José, D.A., Dupski, D.S. & Amilkar, K., *Framework for Security Risk Assessment (FSRA) and Fuzzy Risk Inference System (FRIS) based on Standard ISO/IEC 27002: 2022*, Revista de Informática Teórica e Aplicada, **31**(2), pp. 43-55, 2024.