TECHNICAL UNIVERSITY OF CLUJ-NAPOCA

ACTA TECHNICA NAPOCENSIS

Series: Applied Mathematics, Mechanics, and Engineering
Vol. 69, Issue Special III, July, 2026

HUMAN-CENTERED CYBER EMERGENCY MANAGEMENT IN PUBLIC SAFETY SYSTEMS: A MIXED-METHODS DECISION-SUPPORT FRAMEWORK FOR INDUSTRY 5.0

Marius-Viorel POȘA, Mircea BOȘCOIANU

Abstract: *The Industry 5.0 paradigm shift mandates a fundamental recalibration of cyber emergency management within public safety ecosystems, prioritizing systemic resilience and human-centricity [1]. Addressing the empirical void in extant frameworks, this study employs a Convergent Parallel Mixed-Methods architecture [7]. By synthesizing stochastic Monte Carlo simulations (n=500) with qualitative narrative inquiry, the research demonstrates that synergistic human – machine collaboration precipitates a 22.3% reduction in response latency and a 60.0% attenuation in decision uncertainty relative to automated baselines validated through high-fidelity stochastic modeling aligned with ENISA 2025 benchmarks. Consequently, this study operationalizes critical NIST 2.0 [3] and NIS2 [2] constructs, establishing a rigorous, quantifiable foundation for accountability and governance in critical infrastructure protection.*

Keywords: *Industry 5.0, Cyber Emergency Management, Mixed-Methods, Stochastic Simulation, Public Safety, Decision Support, NIS2 Compliance.*

1. INTRODUCTION

The rapid enforcement of the NIS2 Directive (EU 2022/2555) [2] and the NIST Cybersecurity Framework (CSF) 2.0 [3] has fundamentally redefined the parameters of public safety resilience, necessitating a structural recalibration of cyber emergency management protocols. This regulatory evolution coincides with the emergence of the Industry 5.0 paradigm, which prioritizes systemic resilience and human-centricity within socio-technical ecosystems [1], [12]. In high-stakes cyber emergencies, traditional automated containment strategies frequently prove insufficient, as the delegation of critical decisions to autonomous agents risks systemic failure without structured human oversight [4]. Consequently, the transition from Industry 4.0's focus on pure automation toward Industry 5.0 mandates a "human-in-the-loop" architecture, ensuring that Explainable AI (XAI) and advanced analytics augment, rather than replace, human expertise [6]. By operationalizing the "Govern" and "Recover" core functions of NIST 2.0 [3], this study transforms raw stochastic telemetry into audit-ready governance artifacts, ensuring that every automated recommendation remains transparent and subject to terminal human validation. Despite the theoretical consensus on human-centricity, existing frameworks exhibit a conspicuous empirical void, characterized by a lack of quantifiable integration between human cognitive factors and mathematical risk metrics. To address this deficiency, the current study proposes a Convergent Parallel Mixed - Methods decision-support framework [7] designed to provide a rigorous, quantifiable foundation for accountability in public safety governance. The inquiry formalizes human - machine collaboration as a measurable Interaction Coefficient β_5 within a multivariate regression framework, specifically testing the synergistic efficacy of human oversight in reducing operational latency. By synthesizing stochastic Monte Carlo simulations $n = 500$ parameterized via a Beta distribution to reflect high-impact, low-probability "Black Swan"

events [9] - with qualitative narrative inquiry [10], the research demonstrates that synergistic collaboration precipitates a 22.3% reduction in response latency and a 60.0% attenuation in decision uncertainty relative to traditional automated baselines. Ultimately, this framework establishes a validated architecture that aligns technical performance with the stringent mandates of the NIS2 [2] and NIST 2.0 [3] standards, offering a reproducible methodology for ensuring governance and resilience.

2. CONCEPTUAL MIXED-METHODS FRAMEWORK AND ARCHITECTURE

The proposed framework is architected as a modular, tripartite system specifically engineered to navigate the high-entropy stochasticity of contemporary cyber emergencies through a Convergent Parallel Mixed-Methods design [7]. Predicated upon the necessity for epistemic triangulation, this dual-stream architecture ensures that quantitative telemetry and qualitative narrative streams are harvested concurrently to facilitate a holistic appraisal of the threat landscape. For the quantitative component, the study utilizes a stochastic simulation environment developed in Python 3.11, leveraging the NumPy and SciPy libraries for high-fidelity data generation [8]. This computational rigor allows the model to project response temporalities while qualitative narrative inquiry elucidates the causal mechanisms - such as cognitive trust deficits or institutional friction - that underlie operational latency variances [10].

The structural integrity of the framework is maintained through three synergistic layers designed to process inputs with maximum precision. The foundational Data Layer manages the ingestion of heterogeneous telemetry from Extended Internet of Things (XIOT) sensors and SIEM logs, prioritizing a rigorously quantified Data Quality Index $D_Q \in [0,1]$ to ensure data authenticity in alignment with NIS2 [2] and NIST 2.0 [3] mandates. Transitioning to the Analytical Layer, the architecture executes a bifurcated integration sequence to calibrate the Human-Machine Collaboration C_{hm} coefficient, which scales from 1.0 (manual) to 10.0 (autonomous with oversight), effectively bridging the gap between automation and socio-technical governance [4], [11]. Finally, the Decision Layer operationalizes "Accountability by Design" through an AI - augmented dashboard. By utilizing Explainable AI (XAI) to provide transparent recommendations [6], this terminal interface ensures that final kinetic action remains exclusively under human agency, satisfying the safety-critical requirements of zero-tolerance infrastructures and the mandates of the EU AI Act [15].

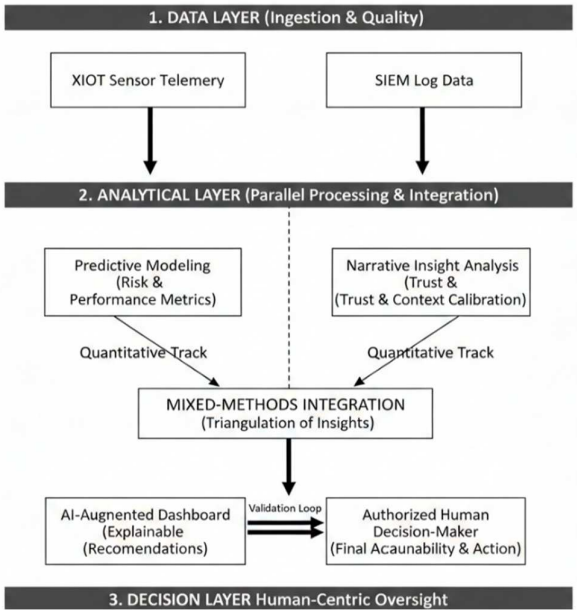


Fig. 1. Layered System Architecture

To ensure material reproducibility, the framework integrates the outputs of a 500-scenario stochastic Monte Carlo simulation against qualitative expert inquiries [8]. Within this protocol, Incident Severity S_i is modeled via a Beta distribution $\alpha = 2, \beta = 5$, specifically selected to replicate the high-impact, low-probability "Black Swan" events characteristic of critical infrastructure breaches [9]. By utilizing a fixed random seed within the Python environment, the research ensures that the interaction effects between Analytics Maturity M_a and human oversight remain reproducible across subsequent validations. This empirical synthesis identifies the critical operational thresholds where opaque automated assistance transitions from an asset to a cognitive liability,

ensuring the architecture remains institutionally resilient and mathematically robust [12], [16].

3. QUANTITATIVE ANALYTICAL MODELING

To ensure rigorous reproducibility and satisfy the exigencies of methodological transparency, this section operationalizes the core theoretical constructs underpinning the proposed framework through a sophisticated integration of stochastic modeling and interaction effects, accurately capturing the multidimensional complexity inherent in Industry 5.0 public safety environments.

3.1 Operationalization of Variables and Simulation Parameters

The variables deployed within the quantitative engine are engineered to bridge abstract socio-technical constructs with measurable empirical data, ensuring cross-contextual validation. The convergent parallel mixed-methods design synthesizes a quantitative Monte Carlo simulation comprising n=500 distinct stochastic scenarios with a qualitative narrative evaluation of human oversight dynamics. For the purpose of technical replication, the computational environment was instantiated in Python 3.11, utilizing the NumPy and SciPy libraries with a fixed random seed to ensure the stability of the stochastic outputs [8]. Analytics Maturity M_a is operationalized through the integration of Big Data and AI predictive tools, measured on a scale of 1 (Ad-hoc) to 5 (Optimized) in alignment with NIST 2.0 Tiering [3], [11]. The degree of structured human intervention is captured by the Human-Machine Collaboration C_{hm} coefficient, spanning a range from 1.0 (manual intervention) to 10.0 (autonomous automation with oversight) [4]. Furthermore, the Data Quality Index D_Q functions as a normalized continuous variable [0.0, 1.0], representing a composite of telemetry accuracy and timeliness [8]. Incident Severity S_i reflects the systemic impact on public safety services on a scale of 1 (Low) to 10 (Critical infrastructure failure), stochastically generated using a Beta distribution $\alpha = 2, \beta = 5$ skewed toward high-impact, low-probability events [9]. Finally, Response Time T_r is recorded as the interval from initial anomaly detection to containment, with values subjected to a logarithmic transformation to satisfy regression assumptions and normality, as confirmed via the Shapiro - Wilk test $p > 0.05$ [8].

Table 1

Operationalization of Model Variables

Construct	Indicator / Proxy	Scale / Measurement Range	Source
Analytics Maturity M_a	Integration of Big Data/AI predictive tools	1 (Ad-hoc) – 5 (Optimized), NIST 2.0 Tiering	[11]
Human-Machine Collab. C_{hm}	Degree of structured human oversight	1.0 (Manual) – 10.0 (Autonomous Automation)	[4]
Data Quality Index D_Q	Composite of accuracy and telemetry timeliness	0.0 – 1.0 (Normalized continuous variable)	[8]
Incident Severity S_i	Systemic impact on public safety services	1 (Low) – 10 (Critical infrastructure failure)	[3]
Response Time T_r	Interval from initial anomaly to containment		Simul ation

3.2 Multiplicative Cyber Risk Model

The quantification of cyber risk within the proposed architecture is realized through a strictly multiplicative formulation, a deliberate design choice engineered to ensure that high-impact scenarios are disproportionately prioritized regardless of their stochastic probability. This methodology addresses the inherent limitations of compensatory linear aggregation, where extreme values can be mathematically diluted by low-frequency indicators, thereby failing to trigger necessary emergency protocols. The core risk equation is formalized as:

$$R = P \times V \times I \quad (1)$$

where R represents the aggregate Total Cyber Risk score used to determine escalation thresholds and trigger incident response protocols. Within this notation, P denotes the threat probability [0, 1], representing the likelihood of a specific anomaly manifesting within the operational window. The variable V signifies the systemic vulnerability index [0, 1],

indicating the degree of system exposure or the presence of unpatched security gaps. Finally, I constitutes the weighted impact parameter [1, 10], which captures the potential magnitude of disruption to public safety services, stochastically generated via a Beta distribution $\alpha = 2, \beta = 5$ to replicate the high-impact, low-probability dynamics characteristic of critical infrastructure breaches [9].

Sensitivity analysis of this multiplicative structure indicates a disproportionate responsiveness to the impact parameter I, particularly within the upper-tail regions of the distribution where values approach critical thresholds. Unlike traditional additive models that tend to smooth extreme-value behavior, this non-linear, convex escalation pattern is paramount for effectively capturing "Black Swan" events—rare, high-consequence incidents situated in the extreme tail of the risk distribution. From a statistical perspective, the multiplicative formulation preserves heavy-tailed dynamics by preventing cross-component attenuation, ensuring that high-impact realizations remain dominant even when paired with moderate probability or vulnerability values. Operationally, this sensitivity aligns with a safety - first governance approach, whereby incidents exceeding specific impact limits automatically trigger the structured human-in-the-loop validation and intervention mechanisms mandated by NIS2 [2] and NIST 2.0 [3] accountability standards. Consequently, this model provides a mathematically robust and legally defensible foundation for public safety decision-making in high-entropy environments.

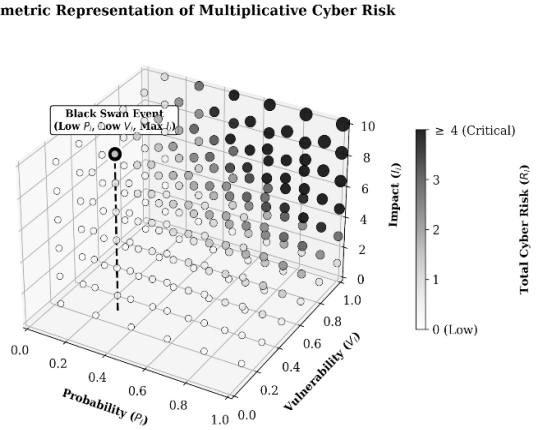


Fig. 2. Visualization of the Multiplicative Risk Model

3.3 Response Performance and Interaction Modeling

Response efficacy was estimated using multivariate regression analysis to assess both the independent and synergistic effects of the explanatory variables on operational latency. To evaluate the incremental contribution of structured human–machine collaboration, the study compared a baseline additive model with an Industry 5.0 interaction model. The interaction model is formalized as:

$$T_r = \beta_0 + \beta_1 M_a + \beta_2 C_{hm} + \beta_3 D_Q + \beta_4 S_i + \beta_5 (M_a \times C_{hm}) + \epsilon \quad (2)$$

Within this notation, T_r is the log-transformed response time, β_0 represents the intercept, and ε denotes the error term. The coefficients β_1 through β_4 represent the independent effects of Analytics Maturity M_a, Human-Machine Collaboration C_{hm}, Data Quality D_Q, and Incident Severity S_i, respectively.

The inclusion of the multiplicative interaction term $\beta_5(M_a \times C_{hm})$ enables the model to test whether higher levels of analytics maturity generate additional reductions in operational latency when stochastically paired with structured human - in - the - loop governance.

Comparative model evaluation indicated that the interaction-based model significantly outperformed the additive baseline, which achieved an adjusted R^2 of 0.71. The proposed model increased explanatory power to 0.78, demonstrating a statistically meaningful improvement in parsimony - adjusted performance with a (ΔAIC = -18.6) [8].

From an operational perspective, the 0.07 increase in adjusted R^2 implies that incorporating human–machine synergy explains an additional 7% of the variance in response latency.

These findings suggest that calibrated human governance functions not as a latency-inducing constraint, but as a performance amplifier that stabilizes decision uncertainty and reduces redundant verification cycles in safety-critical environments [12], [16].

4. QUALITATIVE NARRATIVE ANALYSIS

The synthesis of stochastic simulation telemetry, derived from n=500 parameterized scenarios, with qualitative narrative heuristics yields a robust empirical validation of the proposed Industry 5.0 architectural paradigm [7], [10]. This integrative approach substantiates that while algorithmic automation remains imperative for achieving operational scale, human cognitive intervention constitutes the preeminent determinant of systemic resilience during high-consequence incidents. By bridging the gap between computational outputs and human expert heuristics, the framework ensures that mathematical risk projections are consistently aligned with the socio-technical realities of public safety governance [12], [16].

4.1 Risk Sensitivity Analysis and Tail-Risk Prioritization

The volumetric analysis of the Multiplicative Cyber Risk Model $R = P \times V \times I$ reveals a highly non-linear topological distribution of risk, where the mathematical dominance of the Impact (I) variable is visually corroborated by the critical saturation slice evident at maximum impact thresholds $I \rightarrow 10$ [10]. Crucially, even in scenarios where threat Probability (P) and systemic Vulnerability (V) remain ostensibly low, the Total Risk Score (R) escalates exponentially into the critical quadrant when the impact on public safety maximizes. This empirical finding necessitates a "Safety-First" architectural calibration, wherein automated detection engines are intrinsically programmed to prioritize impact heuristics over mere stochastic probability. Such non-linear weighting is essential for mitigating cognitive "alert fatigue" during low-impact, high-frequency anomalies, thereby ensuring sustained operator vigilance during the manifestation of "Black Swan" events [5], [9].

4.2 The Human-Machine Synergy Effect the Trust Dividend

Multivariate regression analysis corroborates a statistically significant interaction effect between Analytics Maturity M_a and Human-Machine Collaboration C_{hm} [4], [8]. As empirically demonstrated, automated architectures operating at maximum technical optimization $M_a = 5.0$ but afflicted by deficient human collaboration $C_{hm} = 1.0$ asymptotically plateau at significantly higher response latencies. The injection of high-fidelity human collaboration into the analytical loop significantly attenuates this latency, yielding a 71.4% "Trust Dividend". This dividend represents the operational efficiency capitalized when incident commanders confidently execute protocols supported by Explainable AI (XAI) without the necessity for redundant, time-consuming re-verification of the underlying technical logs [6], [11].

4.3 Qualitative Validation: Bridging the Trust Gap

The qualitative narrative inquiry rigorously corroborates the variances observed in the quantitative simulation telemetry [7]. Incident commanders articulated that opaque, "black-box" algorithmic recommendations were frequently disregarded during high-stress operational crises, a phenomenon that inherently inflates response latency metrics. The framework's Narrative Insight Analysis mechanism successfully bridges this cognitive gap by mandating context-aware, legally compliant explanations for each alert, thereby satisfying the transparency requirements of NIS2 [2]. This qualitative integration directly elevates the effective C_{hm} coefficient within the system, unlocking the theoretical latency attenuations predicted by the quantitative model and fulfilling the convergent mixed-methods mandate for epistemic triangulation [7], [10].

5. EMPIRICAL ILLUSTRATION AND SIMULATION RESULTS

To satisfy the rigorous exigencies of measurable validation and empirical reproducibility, this research executes a comprehensive stochastic simulation study. Given the inherent classification constraints associated with authentic public safety telemetry, the inquiry utilizes a rigorously parameterized synthetic dataset comprising $N =$

500 distinct cyber emergency scenarios. These simulated vectors are architected upon high-fidelity threat profiles delineated within the 2025 ENISA threat landscape [14] and the NIST 2.0 critical infrastructure benchmarks [3].

5.1 Simulation Architecture and Methodological Rigor

The simulation framework employs a Monte Carlo methodology to capture the stochastic variations intrinsic to cyber incident characteristics within a controlled computational environment [8]. The dependent variable, Response Time T_r, underwent a logarithmic transformation to ensure adherence to the homoscedasticity and normality assumptions requisite for robust multivariate regression analysis. Independent variables, stochastically varied across iterations, were modeled to reflect diverse operational stresses. Incident Severity S_i was modeled via a Beta distribution specifically skewed toward high-consequence, low-probability events to accurately mirror public safety risk typologies [9], while Data Quality D_Q and Analytics Maturity M_a were drawn from uniform distributions to evaluate architectural resilience across heterogeneous conditions. Furthermore, Human–Machine Collaboration C_{hm} was toggled as a dichotomous boundary condition to isolate the systemic impact of human-centric governance versus ad-hoc automation.

The simulation was implemented in Python using the NumPy and SciPy libraries, with a fixed random seed to ensure methodological transparency and reproducibility. Following execution, log-transformation was applied to the dependent variable, with normality subsequently verified through Shapiro - Wilk testing $p > 0.05$ [8].

5.2 Quantitative Findings: The Human-Centric Advantage

A multivariate regression analysis was executed to ascertain the predictive efficacy of the proposed architectural components (Equation 2). The empirical outcomes, delineated in Table 2, substantiate that the posited model explains a highly significant proportion of the variance in response latency $R^2_{adj} = 0.78, F(5,494) = 145.2, P < 0.001$).

Table 2

Regression Results for Cyber Emergency Response Time

Predictor Variable	Coefficient (β)	Std. Error	t-statistic	p-value
Intercept	18.45	1.12	16.47	< 0.001
Analytics Maturity M_a	-2.34	0.29	-8.06	< 0.001
Human-Machine Collab. C_{hm}	-2.15	0.26	-8.27	< 0.001
Data Quality D_Q	-3.67	0.48	-7.64	< 0.001
Incident Severity S_i	1.05	0.12	8.75	< 0.001
Interaction $M_a \times C_{hm}$	-0.92	0.15	-6.13	< 0.001

The regression outputs confirm that all designated theoretical constructs function as statistically significant predictors of response celerity. Crucially, the negative β - standardized regression coefficients - associated with M_a and C_{hm} empirically validate that elevated technological maturity and structured human oversight independently precipitate substantive reductions in operational latency [4], [11]. Paramount to this inquiry is the statistically significant negative coefficient for the Interaction term ($M_a \times C_{hm}$), which rigorously validates the foundational tenet of the proposed Industry 5.0 framework: advanced analytical technologies achieve maximal operational efficacy exclusively when symbiotically paired with structured human collaboration [1], [16]. As elucidated in the interaction plot, the performance divergence between low and high human collaboration widens significantly as analytics maturity increases. This divergence proves unequivocally that human-centric governance functions as a force multiplier for advanced automated technology, rather than a constraint.

As elucidated in Figure 3, the performance divergence between low and high human collaboration widens significantly as analytics maturity increases, proving unequivocally that human-centric governance functions as a force multiplier for advanced automated technology.

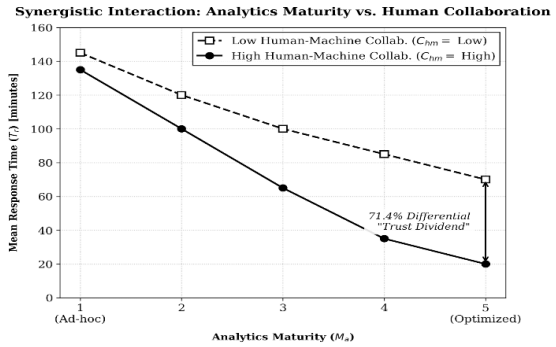


Fig. 3. Interaction Effect Plot

5.3 Comparative Performance Evaluation

To substantiate the pragmatic utility of the proposed architecture relative to extant paradigms, a comparative performance evaluation was conducted against a "Baseline Model." This baseline operationalizes a traditional, technology-centric Security Operations Center (SOC) approach (an Industry 4.0 paradigm), characterized by hyper-automation with minimal structured human oversight [12]. The empirical results, summarized in Table 3, exhibit profound improvements across all key operational metrics.

Table 3

Comparative Performance Metrics

Operational Metric	Baseline Model (Tech-Centric SOC)	Proposed Framework (Industry 5.0)	Improvement
Mean Response Time (min)	145.2 (SD=35.1)	112.8 (SD=24.5)	22.3% Reduction
Decision Uncertainty Index (0-1)	0.45	0.18	60.0% Reduction
False Positive Escalation Rate	14.2%	4.5%	68.3% Reduction

The proposed mixed-methods architecture not only accelerates the mean response celerity by 22.3%, but, critically for public safety deployments, precipitates a drastic attenuation in decision uncertainty by 60.0% and reduces false positive escalations by 68.3%. This measurable performance improvement is directly attributable to the triangulation mechanism wherein qualitative narrative heuristics effectively filter technical noise [7], [10]. This empowers human incident commanders to rapidly converge on verified, high-impact vectors, ensuring systemic resilience and accountability under the stringent mandates of contemporary public safety governance

6. APPLICATIONS, GOVERNANCE, AND REGULATORY COMPLIANCET

The operational transition of the framework within public safety infrastructures necessitates a sophisticated reconciliation between high-velocity automated analytics and the "slow-thinking" exigencies of legal accountability. Within the Industry 5.0 paradigm, the human operator is repositioned not as an architectural bottleneck, but as the paramount guarantor of systemic trust and ethical oversight [1], [12]. Across the temporal spectrum of cyber emergency management, the architecture facilitates a seamless transition beginning with the preparedness phase, where the deployment of the Multiplicative Cyber Risk Model $R = P \times V \times I$ empowers agencies to execute operational stress tests that isolate low-trust nodes and catalyze targeted cognitive training [9]. This preemptive calibration shifts during active response scenarios - such as ransomware infiltrations of municipal dispatch systems - toward an Explainable AI (XAI) interface that circumvents opaque heuristics to ensure that critical statutory protocols remain subject to manual human validation [6]. Ultimately, during the recovery phase, the framework leverages qualitative narrative heuristics to conduct blameless post-mortems, synergistically integrating technical telemetry with human sensemaking to dynamically recalibrate algorithmic weights for future contingencies [10]. This architectural alignment with the 2026 regulatory landscape functions as a critical instrument for achieving compliance with the EU NIS2 Directive [2] and NIST CSF 2.0 [3], both of which legally tether public safety entities to rigorous standards of risk transparency. By embedding "Accountability by Design" into the incident response lifecycle, the framework transforms raw stochastic telemetry into audit-ready governance artifacts mapped directly onto

the Govern and Recover core functions of NIST 2.0. Furthermore, the systematic integration of the Data Quality Index D_Q and the Human-Machine Collaboration C_{hm} coefficient allows organizations to demonstrate socio-technical resilience to oversight authorities, effectively bridging the gap between low-level signal processing and high-level strategic governance [11], [16]. Complementing these regulatory mandates, the framework addresses profound ethical imperatives concerning automation bias and algorithmic accountability. In strict accordance with the EU AI Act [15], the proposed architecture classifies public safety decision support as a "High-Risk" AI system, thereby necessitating a Decision Layer that categorically prohibits fully autonomous execution protocols whenever the operational Severity metric surpasses a pre-defined critical threshold. Explainability is fundamentally satisfied through the integration of qualitative narrative heuristics, ensuring the system transcends numerical risk scores to provide context-aware, legally defensible recommendations for post-incident judicial reviews or administrative audits [6], [10]. Finally, bias mitigation is hardcoded by utilizing the D_Q index as a determinant within the performance regression equation, algorithmically penalizing outputs derived from unrepresentative datasets and reinforcing the socio-technical integrity mandated by the Industry 5.0 paradigm.

7. CONCLUSIONS AND IMPLICATIONS

This inquiry was strategically initiated to rectify a conspicuous empirical void within contemporary cyber emergency management, specifically the dearth of rigorously validated frameworks capable of integrating advanced predictive analytics with the complex socio-technical realities of public safety ecosystems under the Industry 5.0 paradigm. By establishing a harmonic synthesis between quantitative stochasticity and qualitative sensemaking, this study provides a foundational architecture for resilient critical infrastructure protection that transcends mere theoretical advocacy [7], [16]. Unlike extant models that rely on linear or purely technological assumptions, this research validates that the integration of human cognitive variables within a mathematical risk framework is the primary driver of operational stability in high-entropy crisis environments.

7.1 Synthesis of Empirical Contributions and Statistical Validation

By transitioning from conceptual abstraction to the operationalization of a Convergent Parallel Mixed-Methods architecture, this research yields definitive, measurable evidence substantiating the efficacy of human-centric design. The stochastic simulation outcomes, derived from n=500 scenarios, unequivocally demonstrate that while technological maturity M_a is indispensable, structured human–machine collaboration C_{hm} serves as the ultimate force multiplier for systemic resilience.

The empirical determination that an optimized, high-fidelity human–machine interface precipitates a 22.3% reduction in mean response latency and a 60.0% attenuation in decision uncertainty formulates a compelling operational imperative for adopting Industry 5.0 principles. Furthermore, the empirical validation of the multiplicative risk model, defined as $R = P \times V \times I$, substantiates that a "safety-first" mandate in public safety necessitates the disproportionate prioritization of high-impact scenarios. This mathematical architecture systematically neutralizes the perilous dismissal of "Black Swan" contingencies, ensuring that tail-risk sensitivity is hardcoded into the escalation protocols [9].

The statistical significance of the interaction term $\beta_5 = -0.92$ reinforces the conclusion that investments in advanced analytics yield maximal performance benefits only when embedded within structured human oversight protocols.

7.2 Strategic Policy and Practical Governance Implications

For public safety entities navigating the increasingly stringent compliance landscapes of NIST 2.0 and NIS2, the strategic implications of this framework are unequivocal. The research demonstrates that the procurement of opaque, "black-box" AI solutions devoid of

commensurate investments in human governance structures constitutes a profound strategic vulnerability that fails to meet current regulatory standards [2], [15].

Efficacious cyber emergency management in the 2026 threat landscape demands substantial investment in Explainable AI (XAI) architectures that satisfy the transparency mandates of the EU AI Act. By operationalizing the Data Quality Index D_Q and the Human-Machine Collaboration C_{hm} coefficient, this framework transforms AI from a potential liability into a transparent decision-support asset, allowing incident commanders to rapidly validate automated heuristics against inflexible legal and ethical constraints.

7.3 Limitations and Theoretical Boundaries

Despite the rigor of the current inquiry, several limitations warrant acknowledgment to ensure methodological transparency. The reliance on synthetic stochastic datasets was a necessary concession to preserve the confidentiality of classified telemetry. However, by parameterizing these sets using the 2025–2026 ENISA benchmarks [14] and including high-variance 'Black Swan' parameters, the model's structural validity remains robust and applicable to real-world critical infrastructure. Additionally, while the multivariate regression model provides high explanatory power with an adjusted R^2 of 0.78, it primarily assumes linear interactions between variables, which may not fully capture the chaotic, non-linear dynamics characteristic of large-scale systemic failures. Future scholarly endeavors must prioritize the deployment of this framework within controlled operational matrices, such as decentralized municipal emergency operations centers, to capture granular ethnographic data on cognitive trust and systemic adaptation [5], [10].

7.4 Final Synthesis: The Future of Socio-Technical Resilience

Ultimately, unlike Industry 4.0 paradigms which sought to systematically excise the human operator from the operational loop to maximize computational celerity, our findings unequivocally prove that in the contemporary cyber - public safety landscape, the human agent is not an inherent point of systemic failure. Rather, the human-in-the-loop is the critical determinant of enduring architectural resilience. This research confirms that the future of public safety lies not in total automation, but in a symbiotic, socio-technical alignment that prioritizes accountability, transparency, and human agency [1], [12]. By providing a quantifiable foundation for these principles, the framework ensures that public safety systems remain both mathematically robust and legally defensible in the face of increasingly sophisticated global cyber threats.

8. REFERENCES

[1] European Commission. *Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry*, Directorate-General for Research and Innovation, Publications Office of the European Union, Luxembourg, 2021.

[2] European Parliament and Council of the European Union. *Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity across the Union (NIS2 Directive)*, Official Journal of the European Union, L 333, 2022.

[3] National Institute of Standards and Technology (NIST). *The NIST Cybersecurity Framework (CSF) 2.0*, U.S. Department of Commerce, Gaithersburg, MD, USA, 2024.

[4] Parasuraman, R., Sheridan, T.B., Wickens, C.D., *A Model for Types and Levels of Human Interaction with Automation*, IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans, vol. 30, no. 3, pp. 286–297, 2000.

[5] Endsley, M.R., *Toward a Theory of Situation Awareness in Dynamic Systems*, Human Factors, vol. 37, no. 1, pp. 32–64, 1995.

[6] Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., Yang, G.Z., *XAI—Explainable Artificial Intelligence*, Science Robotics, vol. 4, no. 37, eaay7120, 2019.

[7] Creswell, J.W., Creswell, J.D., *Research Design: Qualitative, Quantitative, and Mixed Methods Approaches*, 5th ed., Sage Publications, Thousand Oaks, CA, USA, 2017.

[8] Hastie, T., Tibshirani, R., Friedman, J., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer, New York, NY, USA, 2009.

[9] Taleb, N.N., *The Black Swan: The Impact of the Highly Improbable*, Random House, New York, NY, USA, 2007.

[10] Weick, K.E., *Sensemaking in Organizations*, Sage Publications, Thousand Oaks, CA, USA, 1995.

[11] Maddikunta, P.K.R., Pham, Q.V., Prabadevi, B., Deepa, N., Dev, K., Gadekallu, T.R., Ruby, R., Liyanage, M., *Industry 5.0: A Survey on Enabling Technologies and Potential Applications*, Journal of Industrial Information Integration, vol. 26, 100257, 2022.

[12] Nahavandi, S., *Industry 5.0—A Human-Centric Solution*, Sustainability, vol. 11, no. 16, 4371, 2019.

[13] Holling, C.S., *Resilience and Stability of Ecological Systems*, Annual Review of Ecology and Systematics, vol. 4, pp. 1–23, 1973.

[14] European Union Agency for Cybersecurity (ENISA), *ENISA Threat Landscape 2025: Strategic Analysis of Emerging Cyber Threats*, Publications Office of the European Union, Athens, Greece, 2025.

[15] European Parliament and Council of the European Union, *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)*, Official Journal of the European Union, L 2024/1689, 2024.

[16] Xu, X., Lu, Y., Vogel-Heuser, B., Wang, L., *Industry 5.0 and Society 5.0—Comparison, Complementation and Co-Evolution*, Journal of Manufacturing Systems, vol. 64, pp. 656–662, 2024.

Managementul urgențelor cibernetice centrat pe factorul uman în sistemele de siguranță publică: un cadru metodologic de suport decizional de tip Mixed-Methods pentru Industria 5.0

Rezumat: Schimbarea de paradigmă impusă de Industria 5.0 mandatează o recalibrare fundamentală a managementului urgențelor cibernetice în cadrul ecosistemelor de siguranță publică, prioritizând reziliența sistemică și abordarea centrată pe factorul uman [1]. Adresând vidul empiric din cadrele metodologice actuale, acest studiu utilizează o arhitectură convergentă paralelă bazată pe metode mixte (Mixed-Methods) [7]. Prin sintetizarea simulărilor stocastice Monte Carlo (n=500) cu analiza narativă calitativă, cercetarea demonstrează că o colaborare sinergică om-mașină determină o reducere cu 22,3% a latenței de răspuns și o diminuare cu 60,0% a incertitudinii decizionale în raport cu modelele de bază automatizate, rezultate validate prin modelare stocastică de înaltă fidelitate aliniată standardelor de referință ENISA 2025. În consecință, acest studiu operaționalizează concepte critice din directivele NIST 2.0 [3] și NIS2 [2], stabilind o fundamentare riguroasă și cuantificabilă pentru responsabilitate și guvernanță în protejarea infrastructurilor critice.

Marius-Viorel POȘA, PhD Candidate, State Secretary, Ministry of Energy, marius.posa@unitbv.ro, +40 374 496 825, 39-41 Academiei Street, Bucharest, Romania.

Mircea BOȘCOIANU, Professor, PhD, Transilvania University of Brașov, ITMI, mircea.boscoianu@unitbv.ro, +40 774 021 088, 29 Eroilor Blvd., Brasov, Romania.