International Journal of Science and Research Archive

eISSN: 2582-8185

Cross Ref DOI: 10.30574/ijsra

Journal homepage: https://ijsra.net/

(RESEARCH ARTICLE)

Check for updates

AI-Driven Threat Intelligence Adoption in Security Operations Centers: Human Skill Erosion or Analyst Augmentation? *A Comparative Case Study of Resource-Constrained SOCs in Emerging Markets*

Kwaku Gyamfi Boamah *

College of Computing, Grand Valley State University, Allendale, MI 49401, USA.

International Journal of Science and Research Archive, 2026, 19(03), 025-033

Publication history: Received on 21 April 2026; revised on 27 May 2026; accepted on 30 May 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.19.3.1226

Abstract

The integration of artificial intelligence (AI) into Security Operations Centers (SOCs) is widely promoted as a remedy for the global cybersecurity workforce deficit, yet the dominant literature draws almost exclusively on large, well-resourced enterprises in Western contexts. This paper addresses a significant empirical gap by examining whether AI-driven threat intelligence tools augment analyst capabilities or accelerate professional deskilling within resource-constrained SOCs operating in Sub-Saharan Africa's telecommunications and banking sectors. Using a comparative case study design across three SOC environments varying in AI adoption maturity, the study employs semi-structured interviews with 31 participants, 180 hours of behavioral observation, and pre/post analyst skills assessments. Findings demonstrate that augmentation and deskilling are not technologically determined outcomes but are mediated by governance structures, training investment, analyst seniority, and whether AI systems are embedded within human decision-making frameworks rather than positioned as substitutes for human judgment. Junior analysts in high-AI-maturity environments exhibit a mean 23% decline in manual log analysis proficiency, while senior analysts at the same sites show marginal skill improvement when effective human-in-the-loop governance is maintained. SOCs with low AI maturity preserve foundational analyst skills but carry structural detection coverage gaps that AI deployment could address. The paper proposes a governance and training model for resource-constrained emerging market contexts and identifies directions for future empirical validation.

Keywords: Security Operations Centers; AI-Driven Threat Intelligence; Human-AI Collaboration; Analyst Deskilling; Skill Augmentation; Emerging Markets; Cybersecurity Governance

1. Introduction

Security Operations Centers (SOCs) serve as the institutional fulcrum of enterprise cybersecurity - the human and technical architecture through which organizations monitor threats, investigate alerts, and coordinate incident response. The past decade has witnessed the rapid integration of artificial intelligence into SOC workflows, driven by the twin pressures of an expanding global threat surface and a cybersecurity workforce deficit estimated at 4.8 million unfilled positions in 2024 [1]. AI-driven systems - encompassing machine learning-based anomaly detection, natural language processing for threat intelligence aggregation, and automated response orchestration - have demonstrated measurable throughput benefits in controlled and well-resourced deployment environments [2]. Vendors and industry analysts have consequently positioned AI adoption as a near-mandatory strategic direction for organizations seeking to maintain operational security posture in the face of accelerating adversarial sophistication [3].

A second-order consequence of this adoption trajectory, however, has received comparatively limited scholarly attention: the risk that sustained reliance on automated detection and triage tools may progressively erode the

* Corresponding author: Kwaku Gyamfi Boamah

foundational analytical skills upon which SOC resilience ultimately depends. Gartner has articulated this concern explicitly, warning that cybersecurity leaders who subordinate workforce skill development to technological investment risk producing SOC teams with diminished capacity to critically interrogate the outputs of their own tools [4]. This warning resonates with a broader body of human factors research demonstrating that over-reliance on automation produces complacency, atrophied manual proficiency, and degraded situational awareness in high-stakes operational domains [5]. The problem is not that AI-driven tools are inherently skill-degrading; it is that their effects on human analytical capability depend substantially on the governance and training conditions under which they are deployed - conditions that vary widely and are rarely examined in the published literature.

The existing research on AI adoption in SOC environments exhibits a pronounced geographic and organizational bias. Peer-reviewed empirical studies overwhelmingly draw on enterprise SOCs in North America, Western Europe, and Australia - environments characterized by deep staffing pools, dedicated AI training programs, mature security governance frameworks, and vendor support ecosystems [2, 6]. This bias is consequential because resource-constrained SOCs - disproportionately prevalent in the telecommunications and banking sectors of Sub-Saharan Africa and other emerging market regions - face a structurally distinct version of the human-AI tradeoff. Prior work by the present author on AI integration in cyber incident response teams demonstrated that AI-driven analytics can enable earlier threat detection and prioritization of high-severity alerts [7]. Separately, research on cloud security posture management in multi-cloud environments documented that automated risk identification tools achieve significant reductions in misconfiguration incidents and threat detection times, but that their effectiveness depends critically on operator expertise and governance frameworks [8]. A third strand of related work on usability standards for security configurations confirmed that interface complexity and inadequate operator guidance are primary contributors to configuration errors, underscoring the irreducibly human dimension of automated security systems [9]. Taken together, these findings from the author's prior research converge on a proposition that this paper is designed to test empirically: that the human-AI collaboration outcome in resource-constrained operational environments is mediated by governance and training conditions, and cannot be read off from the technical capability of AI tools alone.

This paper addresses the identified gap through a comparative case study of three operational SOC environments in Sub-Saharan Africa. The study investigates whether AI-driven threat intelligence tools augment analyst capabilities or accelerate deskilling, and what governance and training interventions mediate that outcome. The research contributes both theoretical and practical value: theoretically, by extending human-AI collaboration theory to resource-constrained emerging market contexts that the existing literature has not examined; practically, by generating governance and training recommendations that are calibrated to the specific structural constraints of these environments. The remainder of the paper is organized as follows: Section 2 reviews the relevant literature; Section 3 develops the theoretical framework; Section 4 presents the methodology; Section 5 reports findings; Section 6 discusses implications; Section 7 concludes.

2. Literature Review

2.1. The Human Analyst in Security Operations: Expertise and Situated Knowledge

The foundational scholarship on SOC analyst work practice establishes that effective threat detection is not reducible to automated pattern matching. Goodall, Lutters, and Komlodi's field study demonstrated that expert-level intrusion detection depends on deeply situated knowledge - an accumulated, context-specific understanding of network behavior, organizational topology, and adversarial tradecraft built through repeated manual engagement with security telemetry [10]. This situated expertise enables experienced analysts to recognize anomalous patterns that deviate from the behavioral baseline of their specific environment, a capability that generalized machine learning models trained on external datasets may not replicate with equivalent fidelity. The social dimensions of this expertise are equally significant: analysts develop proficiency through discussion of ambiguous alerts with colleagues, participation in post-incident reviews, and the sharing of pattern-matching heuristics across teams [10]. These collaborative knowledge-building practices are disrupted when AI systems resolve ambiguous alerts automatically before analysts have the opportunity to engage with them as learning objects.

The concept of alert fatigue has been extensively documented in this context, understood not merely as a product of alert volume but as a symptom of misalignment between meaningful threat analysis and the attentional demands of monitoring AI-generated alert queues [11]. A systematic review by Piki and Stavrou found that 51-54% of SOC analysts report being overwhelmed by alert volumes, and that IBM data indicate teams resolve fewer than half of their assigned alerts within a standard shift [11]. While AI-driven alert prioritization reduces raw alert volumes, it introduces a different cognitive challenge: analysts must evaluate AI-generated verdicts without access to the raw evidence needed for independent verification - a condition that predicts automation complacency [5].

2.2. Automation, Skill Erosion, and the Augmentation Debate

The relationship between automation and professional skill is a well-established theoretical problem with an empirically equivocal literature. Parasuraman and Riley's taxonomy of human responses to automation - comprising use, misuse, disuse, and abuse - identifies the cognitive and organizational conditions under which appropriate automation reliance or complacency prevails, providing a foundational framework for analyzing human-AI collaboration dynamics [5]. Research on cognitive offloading documents that individuals who consistently delegate structured tasks to automated systems show measurable reductions in independent task performance when automation is withdrawn, consistent with the neural adaptation mechanisms underlying expert skill formation [12].

In the cybersecurity domain, the deskilling risk is concentrated at the triage layer, where AI automation reduces the frequency of manual engagements through which junior analysts would otherwise build the pattern-recognition skills foundational to advanced investigation. Perkins, Canning, and Flowerday's scoping review confirmed that increased automation amplifies the risk of automation bias and complacency whereby analysts become over-reliant on automated outputs, failing to seek confirmatory or contradictory information independently [13]. The countervailing augmentation literature establishes that skill outcomes are not technologically determined. Singh et al.'s longitudinal study of LLM usage by 45 SOC analysts found that AI tools function as flexible, on-demand cognitive aids supporting sensemaking and context-building when analysts retain primary decision authority, producing augmentation rather than displacement [6]. Lee and See's foundational work on trust in automation identifies decision authority, feedback loops, and trust calibration as the principal governance variables determining whether human-AI collaboration produces augmentation or complacency [14].

2.3. AI and Cybersecurity in Resource-Constrained Emerging Market Contexts

The cybersecurity literature specific to Africa documents a structural tension between accelerating digital exposure and lagging security operations capacity. Analysis by Grant Thornton found that 28% of African organizations experienced breaches despite having security controls in place, and identified the cybersecurity skills gap as a critical structural constraint, with 55% of African firms planning cybersecurity recruitment in 2023 yet still unable to maintain fully operational SOCs [15]. Research by the Carnegie Endowment for International Peace on cybersecurity and financial inclusion in Africa highlighted the specific threat surface of the mobile money and banking ecosystems that represent the primary operational context for two of the case study sites in this research [16]. The Africa cybersecurity market, valued at USD 4.84 billion in 2024 and growing at a CAGR exceeding 20%, reflects investment growth concentrated in commercial tooling acquisition rather than workforce development - a pattern that intensifies the human-AI tradeoff by deploying AI systems into environments where the human analytical infrastructure needed to govern them is underdeveloped [17].

Notwithstanding this context, peer-reviewed research examining AI adoption dynamics within African or broadly emerging market SOC environments is substantially absent from the literature. This gap is not incidental: SOC environments in Sub-Saharan Africa characteristically combine small analyst teams, high staff turnover, limited vendor support, heterogeneous legacy and contemporary tooling, and exposure to regionally specific threat actors - factors that collectively alter the human-AI collaboration equation in ways that published research on well-resourced enterprise SOCs does not capture [18].

3. Theoretical Framework

This study is grounded in a socio-technical systems perspective, which treats SOC performance as an emergent property of the interaction between human actors, technical artefacts, organizational structures, and the broader institutional environment [10]. AI-driven threat intelligence tools are not studied as isolated technical objects but as components of a socio-technical configuration whose properties - including the skill outcomes of analysts who use them - are shaped by organizational governance, training investment, role structure, and operational context. The framework integrates three theoretical traditions.

First, Parasuraman and Riley's automation interaction model provides a taxonomy of human responses to automation and identifies the cognitive and organizational conditions under which appropriate reliance or complacency prevails [5]. This framework is employed to categorize the human-AI collaboration modes observed across the three case study sites and to analyse the governance conditions that differentiate augmentation from complacency. Second, the concept of situated expertise developed by Goodall et al. [10], anchored in Lave and Wenger's communities of practice framework [19], is used to examine how AI adoption affects the conditions under which expert analytical knowledge is built, transmitted, and maintained within SOC teams. Third, the contingency perspective on automation skill outcomes established in the human-computer interaction literature holds that augmentation or deskilling depends on whether

automated systems handle tasks in their entirety or handle sub-tasks while leaving higher-order judgment to the human [12]. This perspective frames the comparative analysis across case study sites differentiated by AI maturity level.

These three frameworks converge on a set of propositions that guide the empirical inquiry: that skill outcomes in AI-augmented SOC environments are contingent rather than technologically determined; that governance and training investment are the primary mediating variables; that junior and senior analysts face qualitatively different risk profiles with respect to deskilling; and that resource-constrained environments create distinctive vulnerabilities not captured by the dominant literature's focus on well-resourced enterprise contexts.

4. Materials and Methods

4.1. Research Design

This study employs a comparative case study design following Yin's established protocols for multiple-case research [20]. The case study approach is appropriate when research questions require contextual understanding of a contemporary phenomenon within its real-world setting and when the boundary between phenomenon and context is not clearly specified - conditions that characterize the investigation of human-AI collaboration dynamics in specific organizational environments. The comparative design across three cases with varying AI maturity levels enables cross-case pattern analysis while preserving the contextual richness of each site. Three sites were selected through purposive theoretical sampling to represent the range of AI adoption maturity characteristic of the telecom and banking sectors in Sub-Saharan Africa. Table 1 provides a summary profile of each site.

Table 1 Case Study Site Profiles

Case / SOC Profile	AI Maturity Level	Analyst Team Size	Primary Sector	AI Tools in Use
SOC-A: Large Telecom Operator (West Africa, 21-country operations)	High - AI-SIEM with ML triage, automated playbooks, behavioral analytics	18 - 24 analysts (tiered L1-L3)	Telecommunications	AI-augmented SIEM, SOAR platform, ML threat scoring, automated playbook orchestration
SOC-B: Regional Commercial Bank (East/West Africa, 6-country footprint)	Moderate - Signature-based detection with ML anomaly modules; limited automation	8 - 12 analysts (L1-L2 dominant)	Banking / Financial Services	Hybrid SIEM (rule-based + ML add-on), basic SOAR workflow, threat intelligence feed integration
SOC-C: National Telecommunications Regulator / CERT (Sub-Saharan Africa)	Low - Traditional rule-based SIEM; AI piloted but not operationalized	4-6 analysts	Government / Critical Infrastructure	Legacy SIEM (rule-based), open-source threat intelligence aggregation, manual incident tracking

Note: SOC designations are pseudonymous to protect operational confidentiality. AI maturity classifications are adapted from the SANS 2022 SOC Survey framework for emerging market conditions [18].

4.2. Data Collection

Data collection proceeded across three complementary methods. Semi-structured interviews were conducted with 31 participants across all three sites - comprising SOC analysts (L1 through L3), SOC managers and team leads, and information security governance officers. Interview guides were developed following Braun and Clarke's established protocols for qualitative inquiry [21], with questions designed to explore analysts' perceptions of AI tool impacts on investigative practice, their assessments of their own skill development trajectories, and their understanding of governance and training arrangements. Interviews ranged from 45 to 75 minutes, were conducted with informed consent, and were audio-recorded with participant permission.

Behavioral observation was conducted over a combined 180 hours across the three sites during active monitoring shifts. Observation focused on analyst-tool interaction patterns, the degree of active engagement with AI system outputs versus passive acceptance, escalation and consultation behaviors, and knowledge-sharing practices. Observation notes

were recorded using a structured field protocol adapted from D'Amico and Whitley's cognitive task analysis framework for SOC analyst work practice [22]. Skills assessments using standardized threat analysis exercises - requiring analysts to investigate synthetic security scenarios from raw telemetry logs without AI assistance - were administered both prior to and six months following the primary data collection phase, enabling pre/post comparison of independent analytical proficiency.

4.3. Data Analysis

Interview transcripts were analyzed using Braun and Clarke's six-phase thematic analysis procedure: familiarization, initial code generation, theme searching, theme review, theme definition, and reporting [21]. Coding proceeded inductively in the first instance and deductively in a second pass testing emergent themes against a priori propositions from the theoretical framework. Intercoder reliability was established through independent coding of a 20% transcript sample, with disagreements resolved by discussion. Observation data were coded thematically with attention to behavioral sequences corroborating or contradicting interview themes. Skills assessment data were analyzed using descriptive statistics, with within-case and cross-case comparison of score distributions reported as percentage differences from baseline to follow-up assessment.

4.4. Ethical Considerations

The present research work was conducted in full compliance with applicable ethical standards. Written informed consent was obtained from all individual participants. Institutional affiliations and personal identifiers have been suppressed through pseudonymous site designations and role-level attribution to protect participant and organizational confidentiality. The study did not involve animal subjects.

Statement of Informed Consent: Informed consent was obtained from all individual participants included in the study.

Statement of Ethical Approval: Ethics review was conducted in accordance with institutional requirements. The research did not constitute a regulated intervention with human subjects requiring formal committee review beyond standard research ethics compliance.

5. Results and Discussion

5.1. Cross-Case Overview

Table 2 presents a structured cross-case comparison of the principal human-AI collaboration dimensions observed across the three sites. Findings reported below draw on this matrix and are developed in greater analytical depth in the subsections that follow.

Table 2 Cross-Case Comparison of Human-AI Collaboration Dimensions Across SOC Sites

Dimension	SOC-A (High AI Maturity)	SOC-B (Moderate AI Maturity)	SOC-C (Low AI Maturity)	Cross-Case Pattern
Alert Triage Speed	Significantly faster (AI handles ~70% of L1 triage automatically)	Moderate improvement; analysts validate ML flags but retain primary triage responsibility	No AI acceleration; manual triage creates throughput bottlenecks	AI maturity strongly predicts triage throughput; gains most pronounced at L1
Analyst Confidence in Independent Detections	Mixed: junior analysts accept AI verdicts with limited critical scrutiny; senior analysts actively interrogate outputs	Moderate; analysts retain investigative ownership but exhibit early signs of alert dependency	High analytical ownership; deep contextual knowledge of network baselines	Inverse relationship between AI dependency and analyst confidence in independent detection judgment
Skill Retention / Atrophy Signals	L1 analysts show reduced manual log analysis proficiency (mean -23% on assessment); L2–L3	Skills broadly stable; threat hunting skills underdeveloped due	Strong foundational skills; limited exposure to advanced AI-	Skill divergence by role: AI augments senior analysts while creating deskilling

	show marginal improvement (+4%)	to alert-driven workflow	assisted threat correlation	risk for junior analysts
Training and Governance Practices	Structured AI-literacy program (quarterly); periodic manual-analysis exercises; documented escalation governance	Ad hoc AI training; governance relies on informal norms; no structured AI-skill maintenance regime	No AI training program; general cybersecurity CPD; regulatory training aligned to national frameworks	Governance investment is the strongest mediating variable; structured training counters deskilling across maturity levels
Observed Human-AI Collaboration Mode	Augmentation with emerging automation bias risk at L1 [12]	Assisted investigation with healthy analytical engagement	Independent analysis with AI disuse; skills preservation but detection coverage gaps	Collaboration mode is contingent on maturity, role, and governance - not AI capability alone [13]

Note: Cells reflect predominant patterns from combined interview, observation, and skills assessment data. Internal variation by analyst seniority and team is detailed in the narrative. Superscript citation numbers indicate the theoretical constructs anchoring each cross-case pattern.

5.2. SOC-A: High AI Maturity - Augmentation with Emerging Deskilling Risk

SOC-A operated the most advanced AI integration profile, with machine learning-driven triage handling approximately 70% of L1 alert classification automatically. Senior analysts expressed broadly positive assessments of AI impact on investigative effectiveness, describing AI tools as enabling earlier threat correlation and freeing cognitive capacity for higher-order hypothesis generation and adversarial intent assessment - a dynamic consistent with Singh et al.'s characterization of AI tools as on-demand cognitive aids for sensemaking and context-building [6]. However, the junior analyst population presented a starkly divergent profile. Six of eight L1 analysts interviewed expressed difficulty articulating the manual log analysis reasoning underlying the AI system's triage decisions, and observation data confirmed that the majority of L1 alert interactions consisted of reviewing AI-generated summaries and accepting or escalating recommended dispositions with limited engagement with underlying telemetry evidence.

The pre/post skills assessment data provided quantitative support for this qualitative finding: L1 analysts at SOC-A showed a mean decline of 23% in performance on manual log analysis tasks between baseline and follow-up assessments, compared to a 4% mean improvement for L2 and L3 analysts at the same site. This divergence by seniority is analytically significant and has direct implications for organizational resilience: the AI system was simultaneously augmenting analysts with sufficient expertise to engage with it critically and eroding the foundational skill-building opportunities of analysts who needed those opportunities most. This pattern resonates with the cognitive offloading literature, which documents that delegating structured tasks to automated systems reduces independent analytical proficiency for less experienced practitioners [12], and with Perkins et al.'s finding that automation increases the risk of complacency most acutely among less experienced operators [13].

The AI literacy program at SOC-A - the most developed governance investment among the three sites - was effective in building conceptual understanding of the ML detection models among senior analysts but did not include structured manual analysis exercises. Several senior analysts independently identified this gap as consequential. Their assessments align with prior work by the present author on AI integration in incident response, which found that AI-driven analytics require supporting human oversight frameworks to deliver their effectiveness benefits without undermining analytical governance capacity [7]. The absence of structured manual practice at SOC-A represents the most actionable governance finding of the study.

5.3. SOC-B: Moderate AI Maturity - Assisted Investigation with Stable Skill Profile

SOC-B operated a configuration where ML anomaly modules supplemented but did not replace rule-based detection, with analysts retaining primary triage responsibility while using AI flags as one investigative input among several. This configuration produced the most stable skill profile of the three sites. Pre/post assessment data showed no statistically significant change across the analyst population, and interview data described a workflow in which analysts experienced AI modules as providing useful investigative signals rather than definitive verdicts - preserving the active analytical engagement that the theoretical framework associates with skill maintenance [14]. Analysts at SOC-B articulated a clear understanding of their AI tools' limitations, including the high false positive rate of anomaly detection modules and the tendency to flag routine mobile money reconciliation patterns as suspicious. This contextual

knowledge - the ability to distinguish genuine anomalies from AI artefacts - represents precisely the kind of situated expertise that Goodall et al. identified as the core of expert intrusion detection practice [10].

The most significant governance deficit at SOC-B was the absence of structured threat hunting investment. Because the analyst workflow was primarily reactive, driven by alert queue management, there was limited organizational investment in proactive exercises that would develop the advanced investigative competencies required for L3-level work. This finding is consistent with the author's prior research on cloud security posture management, which identified continuous skills investment as a prerequisite for maintaining the human analytical capacity needed to govern automated security systems effectively [8]. SOC-B's AI tools were functioning well within their detection scope, but the absence of a broader skills development strategy meant that organizational analytical capacity was not advancing in pace with the threat landscape.

5.4. SOC-C: Low AI Maturity - Skills Preservation with Detection Coverage Gaps

SOC-C presented a contrasting profile in which the absence of operational AI integration preserved broad analytical skills at the cost of detection throughput. Analysts at this site demonstrated the strongest performance on manual log analysis and threat investigation assessments of the three case groups. Skills assessment data showed a 6% improvement in mean performance between baseline and follow-up across the SOC-C analyst group, attributed by the site manager to a structured mentoring program pairing experienced analysts with junior staff on complex incident cases. This low-cost governance mechanism - more easily maintained in the small SOC teams characteristic of African telecom and government environments than in large enterprise contexts - represents a contextual advantage that resource-constrained organizations can leverage [19].

However, the absence of ML-assisted anomaly detection at SOC-C created a detection coverage problem with direct security consequences. Manual triage bottlenecks generated monitoring gaps during peak alert periods, and certain categories of behavioral attack patterns - low-and-slow lateral movement, insider threat indicators, coordinated distributed probing - were generating no alerts for manual review, because only statistical behavioral analysis can surface them reliably. This finding illustrates the dual risk of the resource-constrained AI adoption decision: under-investment preserves analyst skill but creates structural detection gaps, while over-investment without governance produces deskilling. The usability research finding that security configuration complexity directly contributes to human error and system misconfiguration applies analogously here: when AI systems are adopted without the interface and governance scaffolding that enables humans to engage with them effectively, the promised security benefits are not reliably realized [9].

5.5. Cross-Case Synthesis

Three cross-case patterns emerge from the comparative analysis. First, the augmentation-versus-deskilling outcome is contingent rather than technologically determined. The same category of AI-driven threat intelligence tool produced augmentation for senior analysts and deskilling for junior analysts at SOC-A, while producing neither effect at SOC-B, and was absent from SOC-C's operational environment altogether. Technical capability does not determine human skill outcomes; governance, training investment, role structure, and the degree of human analytical engagement with AI outputs are the decisive variables [5, 14]. Second, governance quality is the strongest cross-case mediating variable. The sites with more developed governance frameworks - SOC-A's AI literacy program and SOC-C's mentoring structure - produced better skill outcomes than the governance deficit at SOC-B, regardless of AI maturity level. Third, the deskilling risk is concentrated at the entry point of the analyst pipeline. Junior analysts are most exposed to skill atrophy because automated triage removes the manual engagements through which foundational detection knowledge is built, and their existing skill base is insufficient to compensate through critical engagement with AI outputs. This distribution of risk creates a compounding organizational vulnerability: as senior analysts with the expertise to govern AI systems eventually depart, the organization's capacity to maintain effective human oversight of its own tooling erodes.

6. Conclusion

This paper has examined whether AI-driven threat intelligence tools in resource-constrained SOC environments augment analyst capabilities or accelerate professional deskilling. Through comparative case study evidence from three operational SOCs in Sub-Saharan Africa, the research demonstrates that neither outcome is technologically predetermined. Augmentation and deskilling are contingent products of governance structures, training investment, analyst seniority, and the conditions of human analytical engagement with AI system outputs. The findings challenge vendor narratives of unqualified AI-enabled performance enhancement while also challenging the categorical scepticism that equates automation with skill degradation. The more accurate characterization is that AI adoption creates divergent skill trajectories: augmentation for analysts who already possess the expertise to engage critically

with AI outputs, and deskilling risk for junior analysts whose skill formation depends on the operational engagements that automation displaces.

The geographic contribution of this research is as significant as its theoretical contribution. Resource-constrained SOCs in Sub-Saharan Africa face the human-AI tradeoff in its most demanding form: small teams, limited training budgets, high turnover, and exposure to a rapidly evolving threat surface. The governance and training recommendations that emerge from this research - tiered analyst AI interaction protocols, structured manual analysis exercises, phased AI adoption strategies, and integrated skills development frameworks - are designed for this context's specific constraints. Prior research by the author establishing the effectiveness of AI integration in incident response [7], the security benefits of automated risk identification in cloud environments [8], and the importance of human-centered design in security configuration [9] collectively underscore that the question is not whether to adopt AI-driven tools but how to do so in a manner that preserves and develops the human analytical substrate that gives those tools their operational value.

Several limitations of the current study should be acknowledged. The comparative case study design generates analytical depth at the expense of statistical generalizability, and the findings should be treated as empirically grounded propositions warranting larger-scale validation rather than as established conclusions. The skills assessment instrument, while adapted from established protocols, has not been independently validated across all SOC role profiles represented. Future research should prioritize longitudinal quantitative studies tracking analyst skill trajectories over multi-year AI adoption periods, and should expand the geographic scope of emerging market SOC research to include South and Southeast Asia, the Middle East, and Latin America. The human-AI collaboration dynamics documented here are not unique to Africa; they represent a general challenge whose specific expression is shaped by institutional and operational context.

Compliance with ethical standards

Acknowledgments

The author acknowledges the operational security teams and professional communities across Sub-Saharan Africa whose collaborative expertise informed the framing of this research. The author also acknowledges the College of Computing at Grand Valley State University for institutional support during the preparation of this manuscript. No external funding was received for this study.

Disclosure of conflict of interest

The author declares no financial or non-financial conflict of interest in the preparation or submission of this manuscript.

References

[1] ISC2. 2024 ISC2 Cybersecurity Workforce Study. ISC2; 2024. Available from: https://www.isc2.org/Insights/2024/10/ISC2-2024-Cybersecurity-Workforce-Study

[2] Singh R, Tariq S, Jalalvand F, Baruwal Chhetri M, Nepal S, Paris C, et al. LLMs in the SOC: An Empirical Study of Human-AI Collaboration in Security Operations Centres. arXiv preprint arXiv:2508.18947. 2025. Available from: https://arxiv.org/abs/2508.18947

[3] Gartner. Top Cybersecurity Trends CISOs Must Act on in 2026. Gartner Research; 2026. Available from: https://www.gartner.com/en/articles/top-cybersecurity-trends-2026

[4] Gartner. Cybersecurity for CIOs: Build Resilience and Mitigate Risk in the AI Age. Gartner; 2026. Available from: https://www.gartner.com/en/articles/cybersecurity-for-cios

[5] Parasuraman R, Riley V. Humans and automation: Use, misuse, disuse, abuse. Human Factors. 1997;39(2):230–53. https://doi.org/10.1518/001872097778543886

[6] Mohsin A, Janicke H, Ibrahim A, Sarker IH, Camtepe S. A Unified Framework for Human-AI Collaboration in Security Operations Centers with Trusted Autonomy. arXiv preprint arXiv:2505.23397. 2025. Available from: https://arxiv.org/abs/2505.23397

[7] Boamah KG, Asante A, Tieman A, Okai KF. Artificial intelligence integration in cyber incident response teams to enable faster containment, forensic accuracy, and resilient business continuity. International Journal of Science and Research Archive. 2025;17(01):1263–80. https://doi.org/10.30574/ijsra.2025.17.1.2933

[8] Boamah KG. Cloud Security Posture Management: A Comprehensive Analysis of Automated Risk Identification and Mitigation in Multi-Cloud Environments. International Journal of Research and Innovation in Social Science. 2025;9(11). https://doi.org/10.47772/IJRISS.2025.91100349

[9] Boamah KG. Usability Standards for Privacy-Preserving Security Configuration in IoT Devices. International Journal of Research Publication and Reviews. 2024;5(9):3758–70. Available from: https://www.researchgate.net/publication/397982760

[10] Goodall JR, Lutters WG, Komlodi A. Developing expertise for network intrusion detection. Information Technology and People. 2009;22(2):92–108. https://doi.org/10.1108/09593840910962186

[11] Piki A, Stavrou E. Alert Fatigue in Security Operations Centres: Research Challenges and Opportunities. ACM Computing Surveys. 2025. https://doi.org/10.1145/3723158

[12] Baigelenov A, Parsons P. De-skilling, Cognitive Offloading, and Misplaced Responsibilities: Potential Ironies of AI-Assisted Design. In: Extended Abstracts of the ACM CHI Conference on Human Factors in Computing Systems (CHI EA '25). ACM; 2025. https://doi.org/10.1145/3706599.3719921

[13] Perkins E, Canning A, Flowerday S. Automation Bias and Complacency in Security Operation Centers. Computers. 2024;13(7):165. https://doi.org/10.3390/computers13070165

[14] Lee JD, See KA. Trust in automation: Designing for appropriate reliance. Human Factors. 2004;46(1):50–80. https://doi.org/10.1518/hfes.46.1.50_30392

[15] Grant Thornton South Africa. Emerging Cybersecurity Challenges and Opportunities in Africa. Grant Thornton; 2024. Available from: https://www.grantthornton.co.za/insights2/emerging-cybersecurity-challenges-and-opportunities-in-africa/

[16] Carnegie Endowment for International Peace. Security and Trust in Africa's Digital Financial Inclusion Landscape. Carnegie Endowment; 2024. Available from: https://carnegieendowment.org/research/2024/03/security-and-trust-in-africas-digital-financial-inclusion-landscape

[17] Market Data Forecast. Africa Cyber Security Market Size, Share and Growth. Market Data Forecast; 2024. Available from: https://www.marketdataforecast.com/market-reports/africa-cyber-security-market

[18] SANS Institute. 2022 SANS SOC Survey. SANS Institute; 2022. Available from: https://www.sans.org/white-papers/2022-soc-survey/

[19] Lave J, Wenger E. Situated Learning: Legitimate Peripheral Participation. Cambridge University Press; 1991. https://doi.org/10.1017/CBO9780511815355

[20] Yin RK. Case Study Research and Applications: Design and Methods. 6th ed. SAGE Publications; 2018.

[21] Braun V, Clarke V. Using thematic analysis in psychology. Qualitative Research in Psychology. 2006;3(2):77–101. https://doi.org/10.1191/1478088706qp063oa

[22] D'Amico A, Whitley K. The real work of computer network defense analysts. In: Proceedings of the Workshop on Visualization for Computer Security (VizSec 2008). Springer; 2008. p. 19–37. https://doi.org/10.1007/978-3-540-85933-8_3

[23] Goodall JR, Lutters WG, Komlodi A. I know my network: Collaboration and expertise in intrusion detection. In: Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2004). ACM; 2004. p. 342–5. https://doi.org/10.1145/1031607.1031671

[24] Creswell JW. Qualitative Inquiry and Research Design: Choosing Among Five Approaches. 3rd ed. SAGE Publications; 2013.