

Automation of Log Collection and Incident Triage for Security Operations: A Review

IFEANYICHUKWU JEFFREY OKWESA¹, DANIEL AJIGA², UZOAMAKA IWUANYANWU³

¹Tek Experts, Lagos, Nigeria

²Alcorn State University, Mississippi, USA

³National Open University of Nigeria, Nigeria

Abstract- Modern security operations centres (SOCs) ingest volumes of telemetry that far exceed the capacity of human analysts to inspect manually, and the resulting flood of alerts has made automation of log collection and incident triage a central concern for defensive practice. This paper reviews the research and engineering literature underpinning that automation, organising a fragmented body of work into a coherent picture spanning log collection and normalization, detection and correlation, automated triage and prioritization, security orchestration and automation, and the datasets used to evaluate these systems. We conduct a structured narrative review of works published up to 2019, complemented by seminal earlier contributions, and we synthesise them into a described taxonomy that maps each stage of the alert pipeline to its dominant techniques, inputs, and failure modes. The review shows that pattern-based log parsing and statistical anomaly detection have matured substantially, that alert correlation and adaptive classification can materially reduce analyst load, and that orchestration platforms increasingly encode triage as executable playbooks. Nevertheless, persistent challenges remain: high false-positive rates rooted in the base-rate fallacy, alert fatigue and analyst burnout, scarcity of labelled operational data, concept drift, and the weak external validity of widely used benchmark datasets. We argue that human factors, evaluation realism, and feedback-driven learning deserve as much attention as detection accuracy. The paper closes with future directions emphasising analyst-in-the-loop design, reproducible benchmarking, and resilience-oriented triage.

Keywords: security operations centre; log analysis; incident triage; SIEM; alert correlation; anomaly detection; alert fatigue; SOAR

I. INTRODUCTION

Security operations depend on the ability to observe what is happening across an organisation's systems and to act on that observation quickly. Observation

begins with logs: the event records emitted by operating systems, network devices, applications, and security controls. Acting on observation begins with triage: the process of deciding which of the many signals warrant investigation, in what order, and by whom. As enterprise environments have grown in scale and heterogeneity, both activities have shifted from manual craft toward automation. A contemporary security operations centre (SOC) may collect billions of events per day, and the security information and event management (SIEM) platforms that aggregate them can generate more alerts in an hour than an analyst can meaningfully examine in a shift [1]. The gap between signal volume and analytic capacity is the practical problem that automation of log collection and incident triage is meant to address.

The difficulty is not merely one of throughput. Intrusion detection operates under an unfavourable statistical regime: because genuine attacks are rare relative to benign activity, even a detector with an excellent false-positive rate can produce a set of alerts dominated by false alarms, a consequence Axelsson [2] analysed as the base-rate fallacy of intrusion detection. The operational manifestation of this arithmetic is alert fatigue, in which analysts, confronted with overwhelming and largely spurious alerts, become desensitised, slower, and more prone to missing true incidents [3]. Automation is therefore attractive not only because it scales but because, done well, it can suppress noise, add context, and route the residual signal to the right responder.

This review examines the body of work that makes such automation possible. Our scope spans the full alert pipeline from raw log ingestion to orchestrated response, and our aim is integrative: to bring together strands of research that are often studied in isolation, log parsing in the software-engineering literature,

intrusion detection in the network-security literature, alert correlation in the systems-security literature, and analyst behaviour in the human-factors literature, into a single account oriented around the practical stages of SOC work. We make three contributions. First, we survey the state of automated log collection and triage as of 2019, situating each technique in its lineage. Second, we present a described taxonomy that organises the pipeline into five stages and characterises the dominant methods, inputs, and failure modes at each. Third, we consolidate the recurring operational challenges (false positives, alert fatigue, labelling scarcity, and concept drift) and derive from them a set of research directions.

The scale and heterogeneity of modern telemetry

To appreciate why automation is unavoidable, it helps to be concrete about the material that a SOC must process. A medium-sized enterprise generates event data from operating-system audit subsystems, authentication services, endpoint detection agents, firewalls, proxies, virtual private network concentrators, cloud control planes, application servers, database engines, and an expanding population of software-as-a-service platforms whose logs arrive only through vendor application programming interfaces. Each of these sources emits records at a different granularity, in a different format, and at a different rate, and the aggregate is measured not in files to be read but in events per second sustained around the clock. The migration of workloads into cloud and container environments has intensified this pressure, because ephemeral infrastructure produces short-lived identifiers, dynamic network addresses, and orchestration events that multiply the volume of telemetry while shortening the window in which any given record remains interpretable. Encryption, ephemeral compute, and the dissolution of the traditional network perimeter have simultaneously pushed defenders toward log-centric monitoring, since the packet-level visibility on which earlier detection relied is increasingly unavailable. The practical result is that the log, rather than the packet, has become the primary substrate of defensive observation, and the ability to collect, structure, and reason over logs at machine speed has become the precondition for any credible detection capability.

Volume is only one axis of the problem; heterogeneity is the other. The same logical event (a user authenticating, a process spawning a child, a connection opening to an external host) may be recorded in a dozen mutually incompatible ways across the systems a SOC monitors, with different field names, different representations of time, and different levels of detail. This semantic fragmentation means that even a perfectly reliable collection pipeline delivers a corpus that cannot be reasoned over uniformly until it has been normalized, a point that motivates the emphasis placed on collection and normalization in Section 3. It also means that the difficulty of automation is not reducible to compute: a faster machine processes noise faster. What is required is a pipeline that transforms raw, heterogeneous, high-volume telemetry into a small number of contextualised decisions, and it is the architecture of that transformation, stage by stage, that this review sets out to map.

Automation as a socio-technical problem

A recurring theme of this paper is that the automation of log collection and triage is not a purely technical undertaking but a socio-technical one, in which the capabilities of software and the cognition of analysts are tightly coupled. The temptation, in a field driven by engineering, is to treat the analyst as an interface to be optimised away and to measure progress by the sophistication of detection algorithms. We resist that framing throughout. The output of the pipeline is consumed by human beings working under time pressure and cognitive load, and the value of any automated stage is realised only insofar as it improves the decisions those humans make. A detector that surfaces more true positives but buries them under a larger volume of false alarms may lower, rather than raise, the rate at which real incidents are caught, because it degrades the human process that must act on its output. This observation reframes the objective of automation from maximising detection to conserving attention and enabling correct decisions, and it explains why we treat human-factors research as central rather than peripheral. The stages of the pipeline are therefore best understood not as a chain of algorithms but as a division of labour between machine and analyst, in which the machine assumes the repetitive, high-volume work of collection,

structuring, filtering, and enrichment, and the analyst retains the interpretive, high-ambiguity work of judgement, a division that shifts as techniques mature but that, on current evidence, cannot be dissolved entirely.

The remainder of the paper is structured as follows. Section 2 describes the review methodology. Section 3 treats log collection and normalization. Section 4 covers detection and correlation. Section 5 addresses automated triage and prioritization. Section 6 examines security orchestration and automation. Section 7 reviews evaluation practice and datasets and presents the taxonomy. Section 8 discusses gaps and operational challenges, Section 9 outlines future directions, and Section 10 concludes.

Contributions on semiconductor, RF/MEMS and microelectronic systems engineering and biomedical imaging inform the framing adopted here [4]. Cognate work addressing health-system strategy, financing, and data-driven transformation provides complementary grounding [5], [6].

II. REVIEW METHODOLOGY

This is a structured narrative review rather than a formal systematic review with quantitative meta-analysis, a choice appropriate to a field whose contributions span heterogeneous venues and problem formulations that resist common effect measures. We nonetheless applied an explicit and repeatable procedure to reduce selection bias.

Sources and search. We searched the digital libraries and indexes most relevant to the subject (IEEE Xplore, the ACM Digital Library, USENIX proceedings, Springer LNCS, and *Computers & Security*) together with general scholarly search engines to locate seminal and cross-disciplinary works. Query terms combined the operational concepts (SOC, SIEM, incident triage, alert correlation, alert fatigue, security orchestration) with technical concepts (log parsing, anomaly detection, intrusion detection, false-positive reduction) and with evaluation concepts (intrusion detection dataset, benchmark).

Inclusion and exclusion. We included peer-reviewed papers, authoritative standards and technical reports, and widely cited system descriptions that address one

or more stages of the log-to-response pipeline. We prioritised works published in or before 2019, the review's cut-off year, while deliberately retaining foundational older works, for example Denning's [7] intrusion-detection model and Roesch's [8] description of Snort, because the field's vocabulary and assumptions descend directly from them. We excluded works that were purely product marketing, that lacked methodological detail, or whose claims could not be verified against a citable venue.

Synthesis. Rather than tabulating results paper by paper, we synthesised the literature thematically along the pipeline stages that structure Sections 3 through 7, extracting for each stage the dominant techniques, their assumed inputs, and their characteristic failure modes. These extractions are consolidated in the taxonomy of Section 7. Throughout, we treated human-factors studies of SOC work [3], [9] as first-class evidence, on the view that the effectiveness of automation cannot be assessed independently of the analysts it is meant to support.

Organising frame. The analytical device that holds the review together is the pipeline itself: the sequence of transformations through which a raw event becomes, at one extreme, a discarded false alarm, and at the other, an executed containment action. We use this pipeline both descriptively, to locate where a given technique operates, and diagnostically, to reason about how a weakness at one stage constrains the performance achievable at the next. Two properties of the pipeline recur and deserve to be named at the outset. The first is dependency: because each stage consumes the output of the stage before it, errors propagate forward and compound, so that the overall reliability of the system is bounded by its weakest link rather than by its strongest. The second is feedback: the outputs of the later stages (the analyst's verdict on an alert, the outcome of a response action) are, in principle, labelled examples that can improve the earlier stages, closing a loop that turns operation into a source of training signal. Much of the analytical interest of the field lies in the tension between these two properties, and we return to both repeatedly.

Interpretive stance. Finally, a word on how we read the evidence. Detection accuracy is the most heavily reported quantity in this literature, and it is also the least reliable guide to operational value, for reasons

the review develops: laboratory accuracy is measured under closed-world assumptions that deployment violates, and it says nothing about the volume of alerts an analyst must process to realise it. We therefore read reported results conservatively, privileging studies that attend to false-positive burden, to the transferability of results across environments, and to the human process that consumes the output. This stance is itself a finding of sorts, in that it reflects the field's own hard-won recognition that the metrics easiest to optimise are not the ones that matter most.

Contributions on foundational information-society, ICT-for-development, digital-learning and forecasting literature inform the framing adopted here [10]–[44]. Cognate work addressing Salesforce platform engineering and CRM governance provides complementary grounding [45]–[47].

III. LOG COLLECTION AND NORMALIZATION

Every downstream capability depends on the quality of collected logs. Log collection is the pipeline stage that gathers event records from distributed sources, transports them reliably to a central store, and renders them into a consistent form. In practice the raw material is deeply heterogeneous: free-text application logs, structured audit records, network flow summaries, and vendor-specific device messages arrive in incompatible formats, with inconsistent timestamps, and with varying levels of detail. Normalization, parsing raw messages into structured fields with a shared schema, is the precondition for correlating events across sources, and it is where much of the engineering difficulty lies.

The central technical problem is log parsing: converting a free-text log line into an event template (the constant part) and a set of parameters (the variable part). Early practice relied on hand-written regular expressions, which are brittle and costly to maintain as software evolves. A substantial line of research has automated template discovery. Xu, Huang, Fox, Patterson, and Jordan [48] showed that console logs, when parsed against source code and mined for numeric and state features, reveal large-scale system problems that are invisible in any single message. Clustering approaches such as LogCluster [49] group similar lines to extract patterns without supervision,

and the Drain parser [50] achieves online parsing by organising candidate templates in a fixed-depth tree, allowing streaming logs to be structured with low latency. The comparative study by Zhu et al. [51] benchmarked thirteen parsers across sixteen datasets and provided open tools, establishing that parser accuracy varies widely by log type and that no single method dominates, a finding with direct operational consequences, since a parser that mislabels the variable part of a message can silently corrupt every feature computed downstream.

Normalization also imposes semantic obligations that parsing alone does not satisfy. Timestamps must be reconciled to a common clock and time zone; identities and addresses must be canonicalised; and events must be enriched with context (asset criticality, user role, threat intelligence) that is not present in the log itself but is essential for later prioritization [1]. The reliability of collection matters as much as its fidelity: dropped or delayed events create blind spots that an adversary can exploit, and the integrity of the log store is itself a security property, since attackers routinely attempt to tamper with or delete the records of their activity. The upshot is that log collection and normalization, though often treated as plumbing, set a ceiling on the accuracy achievable by every subsequent stage.

Transport, buffering, and the reliability of collection

Before a log line can be parsed it must arrive, and the transport layer that carries events from their source to the central store is a frequently underappreciated determinant of detection quality. Collection architectures range from agent-based designs, in which a lightweight process on each host forwards events, to agentless designs that poll devices or subscribe to network protocols, and each carries distinct reliability characteristics. Under bursty load, precisely the condition that accompanies many incidents, when a compromised host generates a flurry of anomalous activity, a naive pipeline may drop events, and the events lost are disproportionately the ones that matter. Well-engineered collection therefore interposes buffering and back-pressure, absorbing spikes without discarding records and degrading gracefully rather than silently when downstream capacity is exhausted. The choice of transport protocol matters too: unreliable datagram-based forwarding,

still common for legacy device logging, offers no delivery guarantee and no ordering, so that events may arrive out of sequence or not at all, corrupting any analysis that depends on the temporal order of events. These are not merely operational concerns; they are security-relevant, because an adversary who understands the collection architecture can time activity to coincide with load, or can generate a flood of benign events precisely to induce the drops that will hide a genuine action.

A related and often neglected property is clock discipline. Correlating events across sources presupposes that their timestamps can be placed on a common timeline, yet hosts drift, time zones differ, and some devices record local time without an offset. A skew of even a few minutes between two sources can defeat a correlation rule that expects a causal event to precede its consequence, producing either missed scenarios or spurious ones. Robust collection therefore treats time synchronisation and timestamp canonicalisation as first-order requirements, not incidental details, and records both the event time asserted by the source and the ingestion time observed by the collector, so that reasoning can distinguish genuine event ordering from artefacts of delayed delivery.

Enrichment and the manufacture of context

Normalization renders events uniform; enrichment makes them meaningful. A normalized authentication event tells us that an account logged in from an address at a time, but it does not tell us whether that account belongs to a database administrator or a temporary contractor, whether the source address sits inside a sensitive network segment or a hostile foreign range, or whether the asset involved is a disposable test machine or a system of record holding regulated data. This context is not present in the log line; it must be joined in from asset inventories, identity directories, vulnerability data, and threat-intelligence feeds. Enrichment is consequential because it is what allows later stages to prioritise: the same raw event may be routine on one asset and alarming on another, and only the joined context distinguishes them. It is also fragile, because the reference data on which it depends are themselves subject to staleness and error, an asset inventory that omits a newly provisioned server, or an identity directory that has not caught up with a role

change, will enrich events incorrectly and mislead everything downstream. The reliability of enrichment thus depends on data-management disciplines that live outside the security pipeline proper, a dependency that couples the effectiveness of the SOC to the general state of an organisation's data governance.

Taken together, transport reliability, clock discipline, and enrichment show that collection and normalization are far from the mechanical prelude they are sometimes assumed to be. They constitute the stage at which the raw, ambiguous, adversarially manipulable material of logs is turned into the structured, contextualised, temporally coherent record that every later stage takes for granted. Investments here are unglamorous and easily deferred, yet because their failures are silent, a mis-parsed field or a dropped event produces no error, only a subtly wrong answer later, they are among the most consequential and least visible determinants of a SOC's true detection capability.

Contributions on enterprise cybersecurity and cyber-defense engineering inform the framing adopted here [52]. Cognate work addressing procurement and supply-chain delivery in energy and infrastructure settings provides complementary grounding [53]–[56].

IV. DETECTION AND CORRELATION

Once events are collected and normalized, detection identifies those that may indicate malicious or anomalous activity, and correlation assembles related detections into a smaller number of meaningful units. The two are conceptually distinct but operationally intertwined, and both have deep roots in the intrusion-detection literature.

Detection. The foundational distinction is between misuse (signature-based) detection, which matches events against known-bad patterns, and anomaly-based detection, which flags deviations from a model of normal behaviour [7]. Signature systems such as Snort [8] remain widely deployed because they are precise, interpretable, and cheap to operate, but they cannot detect novel attacks and require constant rule maintenance. Anomaly-based detection promises coverage of unknown threats, and a large literature applies statistical and machine-learning methods to

network and host data [57], [58]. Flow-based detection, which operates on aggregated connection records rather than full packet payloads, offers scalability advantages for high-speed networks [59]. More recently, log-sequence models such as DeepLog [60] treat system logs as a language and learn expected event sequences with recurrent neural networks, detecting anomalies as departures from the learned grammar; empirical comparisons of such methods on operational logs show both their promise and their sensitivity to data preparation [61].

The recurring caution across this literature, articulated forcefully by Sommer and Paxson [62], is that machine learning is far harder to apply to intrusion detection than to other domains: the cost of errors is highly asymmetric, the notion of “normal” drifts continuously, semantic interpretation of anomalies is difficult, and the closed-world assumption of standard evaluation rarely holds in deployment. Buczak and Guven [63] survey the many data-mining and machine-learning methods proposed for cybersecurity and note the persistent gap between reported laboratory accuracy and operational usefulness.

Correlation. Detection alone produces a stream of isolated alerts; correlation reduces that stream by grouping alerts that share a cause, aggregating duplicates, and reconstructing multi-step attacks. Valeur, Vigna, Kruegel, and Kemmerer [64] proposed a comprehensive correlation framework composed of a sequence of components (normalization, aggregation, verification, and higher-level scenario construction) and showed that correlation quality depends on how these stages are combined. Ning, Cui, and Reeves [65] reconstructed attack scenarios by chaining alerts whose prerequisites and consequences match, turning a mass of low-level alerts into a small number of interpretable attack narratives. Alarm clustering, as developed by Julisch [66], groups alarms by root cause so that a single environmental or configuration issue generating thousands of alarms can be recognised and addressed once rather than repeatedly. Together these techniques attack the volume problem at its source: they compress many detections into few incidents, which is precisely what triage requires.

The precision–recall tension and the asymmetry of error

Underlying every detection design is a trade-off that no algorithm escapes: raising sensitivity to catch more genuine attacks inevitably admits more false alarms, and tightening precision to suppress false alarms inevitably lets some attacks pass. In most application domains this trade-off is balanced by weighing the two kinds of error against each other symmetrically. Intrusion detection is unusual in the severity of its asymmetry. A missed intrusion may cost an organisation a breach with consequences measured in months of remediation and lasting reputational damage, while a false alarm costs minutes of analyst time, yet false alarms arrive in such volume that their aggregate cost, in attention consumed and vigilance eroded, becomes the dominant operational burden. The consequence is that the detection stage cannot be tuned to an optimum in isolation, because the acceptable false-positive rate depends on the capacity of the stages that follow it to absorb and filter the residual. A detector that would be intolerable feeding analysts directly may be perfectly acceptable feeding a correlation engine that compresses its output, and this coupling between detection sensitivity and downstream filtering is a recurring reason why point improvements in detector accuracy so often fail to translate into operational gains.

Hybrid and layered detection in practice

The textbook opposition between signature-based and anomaly-based detection is, in operational reality, a division of labour rather than a competition. Signatures are precise, interpretable, and cheap, and they excel at recognising the known, the specific exploit, the particular malware family, the documented technique. Anomaly methods are the only hope against the genuinely novel, but they pay for that coverage with opacity and false alarms, since a departure from statistical normality is not the same as maliciousness and may reflect nothing more than a new but benign pattern of use. Mature architectures therefore layer the two, using signatures to dispatch the known cheaply and reserving the more expensive and more error-prone anomaly machinery for the residue that signatures cannot explain. Layering also mitigates a structural weakness of each approach: signatures are blind to what has not yet been

catalogued, while anomaly models are blind to attacks that resemble normal behaviour closely enough to evade statistical notice. Neither weakness is fully curable, but their combination narrows the gap through which an adversary must pass. The log-sequence models that treat system logs as a language extend this layering into a third register, learning the expected grammar of an application's behaviour and flagging departures from it, and they are most valuable precisely where signatures and coarse statistics are silent, in the subtle reordering or omission of events that betrays a compromised process behaving almost, but not quite, normally.

Why correlation is the pivotal stage

If detection produces the raw signal, correlation is where that signal first becomes tractable, and it deserves emphasis as the pivot on which the tractability of everything downstream turns. The reason is arithmetic. A detection stage tuned for coverage produces alerts in numbers that no analyst team can examine, and the great majority of those alerts are individually uninformative, a single port scan, a single failed login, a single unusual connection means little on its own. Correlation exploits the fact that meaning in security data is largely relational: it resides not in individual events but in their patterns, sequences, and shared causes. By aggregating duplicates, grouping alerts that stem from a common root cause, and chaining alerts whose prerequisites and consequences link them into a coherent attack narrative, correlation converts a high-volume, low-information stream into a low-volume, high-information one. This compression is what makes human attention viable at all, and it is also what supplies the context on which triage prioritisation depends: an alert embedded in a reconstructed multi-step scenario carries an evidential weight that the same alert in isolation cannot. The corollary is that weaknesses in correlation are especially damaging, because they either fragment genuine attacks into disconnected fragments that each look innocuous, or fabricate spurious scenarios from coincidental co-occurrence, and both failures propagate directly into the triage decisions that follow.

Contributions on cybersecurity, data governance, and IT-risk management inform the framing adopted here [67]–[69]. Cognate work addressing the broader

external academic literature on analytics, machine learning, security, and digital systems provides complementary grounding [70]–[118].

V. AUTOMATED TRIAGE AND PRIORITIZATION

Triage is the decision layer that sits between detection and human investigation. Its task is to rank and route: to estimate which alerts most likely represent genuine, consequential incidents and to present them, with context, in an order that respects analyst time. Because the base-rate fallacy guarantees that most alerts from a sensitive detector are false [2], effective triage is arguably the highest-leverage form of automation in the entire pipeline.

Two broad strategies dominate. The first is false-positive reduction through learning. Pietraszek [119] framed alert classification as a supervised learning problem in which analyst feedback on past alerts trains a classifier to filter future ones, adaptively suppressing the categories of alert that have historically proven benign in a given environment. This analyst-in-the-loop design is significant because it acknowledges that “false positive” is partly environment-specific: an alert that is spurious in one network may be meaningful in another. The second strategy is prioritization by context and impact. Rather than discarding alerts, prioritization scores them using the value of the affected asset, the credibility of the detection, and the stage of the attack it may represent, so that scarce attention is spent where potential damage is greatest [1]. Correlation feeds directly into prioritization: an alert that forms part of a reconstructed multi-step scenario [64], [65] warrants higher priority than the same alert in isolation.

Threat modelling frameworks have begun to structure triage semantically. MITRE ATT&CK [120] provides a shared vocabulary of adversary tactics and techniques, allowing alerts to be mapped to positions in an attacker's likely progression; this mapping supports prioritization by exposing which alerts, taken together, indicate advancement toward an objective. Standards for incident handling, notably the NIST computer security incident handling guide [121], locate triage within a broader lifecycle of preparation, detection and analysis, containment, eradication, and

recovery, and emphasise consistent categorisation and prioritization as prerequisites for effective response.

Critically, triage is where automation most directly meets human cognition. Studies of real analyst work show that triage is not a mechanical filter but an interpretive activity in which analysts assemble fragmentary evidence into hypotheses [9]. Automation that ignores this, that presents opaque scores without supporting context, risks either being distrusted and ignored or, worse, being followed uncritically. The most promising triage systems therefore aim to augment rather than replace analyst judgement, supplying ranked, contextualised, and explainable candidates for human decision.

The environment-specificity of the false positive

A subtlety that distinguishes triage from detection is that the very notion of a false positive is partly local. In detection, correctness is often treated as an objective property of the event: either it was an attack or it was not. In triage, an alert that is technically a true detection may nonetheless be operationally irrelevant in a particular environment, a vulnerability scanner probing hosts that are known and sanctioned, an administrative tool whose behaviour resembles an attack but is authorised, a legacy application whose quirks trip an anomaly model without any adversary involved. What counts as noise is therefore a function of the specific network, its sanctioned activities, and its risk posture, and it changes as the environment changes. This is why the adaptive, feedback-driven approach to alert classification is so consequential: it learns the local meaning of alerts from the decisions analysts make in that environment, rather than assuming a universal ground truth that does not exist. It also explains why triage rules and models transfer poorly between organisations, and why a triage capability, unlike a signature set, cannot simply be purchased and deployed but must be cultivated against the particular telemetry and conventions of the place it serves.

Prioritisation as a decision under uncertainty

Prioritisation is best understood not as classification but as ranking under uncertainty, and framing it this way clarifies what a good triage system must estimate. Every residual alert carries some probability of

representing a genuine, consequential incident, and every alert also carries a potential impact that depends on the asset, the data, and the stage of attack it may represent. The rational allocation of scarce analyst attention weighs these together, prioritising alerts where the product of likelihood and consequence is greatest, so that the analyst's next action is always the one with the highest expected value. This framing exposes why context is not a luxury but a necessity: without the asset criticality, the identity, and the threat intelligence that enrichment supplies, the impact term cannot be estimated, and prioritisation collapses back into ranking by detector confidence alone, which, given the base-rate problem, ranks a torrent of confident false positives above the occasional quiet signal of a real intrusion. It also reveals a subtle failure mode: because impact and likelihood are both estimated, systematic errors in either, an asset inventory that undervalues a critical system, a model that is overconfident on a common benign pattern, silently distort the ranking, sending attention to the wrong places while appearing to function normally. Good prioritisation therefore depends as much on the quality of the contextual inputs as on the sophistication of the scoring, a dependency that again ties the effectiveness of a late stage to the reliability of the early ones.

Contributions on enterprise IT systems, cloud migration, and service management inform the framing adopted here [122], [123]. Cognate work addressing the foundational and directly cited literature underpinning the present study provides complementary grounding [1]–[3], [7]–[9], [48]–[51], [57]–[66], [119]–[121], [124], [125].

VI. SECURITY ORCHESTRATION AND AUTOMATION

Where triage decides what to attend to, orchestration decides what to do about it and executes as much of the doing as possible. Security orchestration, automation, and response (SOAR) refers to platforms that codify response procedures as executable playbooks, integrate the many tools a SOC operates, and automate repetitive actions (enrichment, containment, ticketing, notification) that would otherwise consume analyst time. Although the SOAR acronym is a relatively recent coinage, the underlying idea is a natural extension of the SIEM-centred

workflow described by Bhatt et al. [1], in which the SIEM is the analytic hub and orchestration layers automated action on top of it.

The value proposition of orchestration is threefold. It reduces mean time to respond by removing manual hand-offs; it enforces consistency, so that a given class of incident is handled the same way regardless of which analyst is on shift; and it captures institutional knowledge in playbooks, mitigating the loss of expertise when experienced staff leave, a real risk given the burnout documented in SOC studies [3]. Orchestration also closes the loop back to triage and detection: the outcome of an automated or analyst-led response is itself a labelled example that can retrain classifiers [119] and refine correlation rules [66].

The risks of orchestration mirror its benefits. Automated containment actions (isolating a host, disabling an account, blocking an address) can cause operational harm if triggered by a false positive, so the degree of automation must be calibrated to the confidence of the upstream decision. This calibration reconnects orchestration to the base-rate problem: fully automated response is safe only where triage is highly precise, which is why many deployments keep a human in the loop for consequential actions and reserve full automation for low-risk enrichment. Effective orchestration is thus best understood not as the removal of analysts but as the industrialisation of the routine, freeing human attention for the ambiguous cases that genuinely require it.

The playbook as codified knowledge and its brittleness

The playbook is the central artefact of orchestration, and it is worth examining what a playbook is and why it matters. At its simplest a playbook is an executable encoding of a response procedure: a sequence of conditional steps that gather additional evidence, make decisions, take actions, and record what was done. Encoding procedures this way has a value beyond speed, because it externalises knowledge that would otherwise live only in the heads of experienced analysts. In a field marked by high turnover and burnout, the departure of a senior analyst can take with it an irreplaceable store of tacit judgement about how particular incidents are handled; a playbook captures a portion of that judgement in a form that outlives the individual and can be applied consistently by less

experienced staff. This is the sense in which orchestration is a mechanism of institutional memory as much as of automation. Yet the same property is a source of brittleness. A playbook encodes assumptions about the environment (the tools available, the network topology, the naming of assets, the meaning of alerts) and when those assumptions drift, the playbook continues to execute confidently on premises that no longer hold. Unlike a human, who notices when a situation departs from the expected, an automated playbook follows its branches whether or not they still make sense, and a stale playbook can therefore take precisely the wrong action with mechanical reliability. The maintenance of playbooks against a changing environment is thus a continuing cost, and one that is easy to underestimate at the point of deployment, when the encoded knowledge is fresh and correct.

Calibrating autonomy to confidence

The deepest design question in orchestration is how much autonomy to grant, and the review's framing supplies a principled answer: the degree of automation permitted for an action should be calibrated to the confidence of the decision that triggers it and to the cost of being wrong. Actions divide roughly into two classes by their reversibility and their potential for harm. Enrichment actions (querying a threat-intelligence source, retrieving the history of an asset, gathering related events) are essentially free of downside, because a wrong guess wastes a little compute and nothing more, and these can be automated aggressively regardless of the confidence of the triggering alert. Consequential actions (isolating a host, disabling an account, blocking an address) carry real operational cost when triggered wrongly, because they can interrupt legitimate work and, in the worst case, cause exactly the disruption an adversary hoped to provoke. For these, full automation is justified only where the upstream triage decision is highly precise, and where it is not, the appropriate design keeps a human in the loop to authorise the action while still using automation to prepare and accelerate it. This calibration reconnects orchestration to the base-rate problem that pervades the review: because most alerts from a sensitive detector are false, the confidence required to license autonomous consequential action is high, and it is achieved not by

the detector alone but by the correlation and triage stages that raise precision before the decision to act is reached. Orchestration, in this light, is not a stage that can be designed independently; its safe scope is set by the reliability of everything upstream of it.

Contributions on financial analytics, audit, IT-risk and governance across the wider research portfolio inform the framing adopted here [126]–[152].

VII. EVALUATION, DATASETS, AND A TAXONOMY OF THE PIPELINE

Progress in automated detection and triage is only as credible as the evidence used to demonstrate it, and evaluation has been a persistent weakness of the field. For years the KDD Cup 1999 dataset served as the default benchmark for intrusion detection, but Tavallae, Bagheri, Lu, and Ghorbani [125] documented serious problems, massive record redundancy and an unrealistic class distribution, that inflate reported accuracy and undermine comparability, and they released the NSL-KDD refinement to address the worst of these. Even refined, such datasets capture a dated threat landscape. More

recent efforts, such as the CICIDS2017 dataset [124], generate labelled traffic from realistic profiles and contemporary attacks, and the log-parsing community has produced open benchmarks and tools for reproducible comparison [51]. Nonetheless, the concern raised by Sommer and Paxson [62] endures: laboratory datasets embody a closed world, whereas deployment confronts an open one, and results that do not transfer across environments overstate operational readiness.

To integrate the preceding sections, we describe a taxonomy of the automated log-collection-and-triage pipeline. It organises the field along five sequential stages, characterising each by its purpose, its dominant techniques, its inputs, and its characteristic failure mode. The taxonomy is intended less as a rigid classification than as a map for locating a given contribution and reasoning about how weaknesses at one stage propagate to the next.

Table 1. A taxonomy of the automated log-collection-and-triage pipeline.

Stage	Purpose	Dominant techniques	Primary input	Characteristic failure mode
1. Collection & normalization	Gather, transport, and structure heterogeneous events	Automated log parsing (Drain, LogCluster); schema mapping; enrichment	Raw device, host, application, and network logs	Mis-parsed templates and dropped events silently corrupt all downstream features
2. Detection	Flag events indicating misuse or anomaly	Signature matching (Snort); statistical and ML anomaly detection; log-sequence models (DeepLog)	Normalized events and flows	High false-positive rate; missed novel attacks; concept drift
3. Correlation	Compress and connect related detections	Aggregation and scenario construction (Valeur et al.); prerequisite/consequence chaining (Ning et al.); root-cause clustering (Julisch)	Detection alerts	Fragmented or spurious scenarios when correlation rules are incomplete
4. Triage & prioritization	Rank and route residual alerts for human attention	Adaptive alert classification (Pietraszek); impact- and asset-based scoring; ATT&CK mapping	Correlated alerts plus	Base-rate fallacy; alert fatigue;

Stage	Purpose	Dominant techniques	Primary input	Characteristic failure mode
			asset and threat context	opaque, distrusted scores
5. Orchestration & response	Execute enrichment, containment, and case management	SOAR playbooks; enrichment and feedback capture	automated containment; prioritized incidents	Harmful automated action on false positives; brittle playbooks

Read across, the table makes explicit a property that individual studies tend to obscure: the stages form a dependency chain in which the reliability of each stage is bounded by the reliability of those before it. A superb triage classifier cannot compensate for a parser that mislabels its inputs, and a well-designed playbook is dangerous atop imprecise triage. This systemic view motivates the discussion that follows.

Why detection is over-evaluated and triage under-evaluated

A striking imbalance runs through the field's evidence base: the detection stage is evaluated far more rigorously, and far more often, than any other stage of the pipeline. The reason is partly structural. Detection lends itself to the standard machinery of empirical evaluation because it can be posed as a classification problem with a labelled dataset, a fixed set of inputs, and a scalar accuracy that different methods can be ranked by. Triage, correlation, and orchestration resist this treatment. Their quality depends on context that a static dataset does not contain, the value of assets, the conventions of an environment, the judgement of the analysts they serve, and their outputs are decisions about the allocation of attention rather than binary labels, so there is no agreed scalar against which to score them. The consequence is that the stages with the most direct leverage over operational outcomes are the least served by reproducible benchmarks, and claims about them rest disproportionately on case description and vendor assertion rather than on comparable measurement. This imbalance is not merely an academic inconvenience; it distorts investment and attention toward the stage that is easiest to measure rather than the stage where improvement would matter most, and it leaves practitioners without the evidence

they would need to choose a triage or orchestration approach on anything other than faith.

Reading the taxonomy diagnostically

The taxonomy of Table 1 is intended to be used, not merely contemplated, and its diagnostic use follows from the dependency property it makes visible. Confronted with a SOC that is underperforming (too many missed incidents, too many wasted analyst hours, too slow a response) the natural instinct is to upgrade the most visible component, usually the detector. The taxonomy counsels a different discipline: to trace the pipeline from the front and ask where the binding constraint actually lies. Missed incidents may originate not in a weak detector but in dropped events or mis-parsed fields at collection, which no detector can recover. Wasted hours may originate not in poor detection but in absent correlation, so that analysts drown in individually valid but collectively redundant alerts. Distrust of automation may originate not in the scoring model but in the absence of the explanations analysts need to act on it. In each case the observable symptom appears at one stage while its cause lies at another, and remediation aimed at the symptom disappoints. The value of the taxonomy is that it forces this systemic reading, directing effort to the stage that is genuinely limiting overall performance rather than to the stage that is most conspicuous or most easily replaced.

Contributions on internal audit, risk governance, and financial-sector controls inform the framing adopted here [153]–[156].

VIII. DISCUSSION: GAPS AND OPERATIONAL CHALLENGES

Synthesising the literature exposes several challenges that recur across stages and that, taken together, define the practical limits of current automation.

The base-rate problem and false positives. The statistical asymmetry identified by Axelsson [2] is not an artefact of poor engineering but a structural feature of a domain in which attacks are rare. It implies that raising detector sensitivity, absent commensurate improvements in precision and correlation, simply multiplies false alarms. The literature's most effective responses attack the problem downstream of detection, through correlation that compresses alerts [64], [66] and through adaptive, feedback-driven classification [119], rather than expecting any single detector to solve it.

Alert fatigue and human factors. The human consequence of false-positive-laden alert streams is well documented. Sundaramurthy et al. [3] showed, through embedded fieldwork in operational SOC's, that analyst burnout is a systemic outcome of tooling that generates more work than it removes, and D'Amico and Whitley [9] characterised analytic work as cognitively demanding sense-making rather than rote filtering. These findings reframe automation's goal: the metric that matters is not detections produced but analyst attention conserved and correct decisions enabled. Automation that increases alert volume without increasing decision quality is counterproductive even when its detection accuracy is high.

Labelling scarcity. Supervised triage and detection depend on labelled examples of true and false alerts, but such labels are expensive, environment-specific, and often unavailable at the scale machine learning requires. The feedback loops described in Sections 5 and 6 partially address this by harvesting labels from analyst decisions and response outcomes, but they introduce their own biases: alerts never surfaced to analysts are never labelled, so classifiers can entrench existing blind spots.

Concept drift. Both benign behaviour and adversary technique evolve continuously, so models of "normal" and of "malicious" degrade over time [62]. Static

signatures miss new attacks; static anomaly models raise false alarms as the environment changes. Sustained accuracy therefore demands continual retraining and rule maintenance, an operational cost that laboratory evaluations rarely capture.

Evaluation realism. The dataset problems documented by Tavallae et al. [125] and the closed-world critique of Sommer and Paxson [62] together imply that much reported progress may not transfer to deployment. Newer datasets [124] and open parsing benchmarks [51] improve matters, but the field still lacks widely accepted, continually refreshed benchmarks that reflect current threats and heterogeneous environments, and it especially lacks benchmarks for the triage and orchestration stages, which are evaluated far less rigorously than detection.

Systemic coupling. Finally, the taxonomy of Section 7 highlights that these challenges are not independent. Poor normalization amplifies false positives; false positives drive fatigue; fatigue and drift erode the label quality on which adaptive methods depend. Improvements pursued at a single stage, in isolation, tend to disappoint because the binding constraint lies elsewhere in the chain.

The feedback loop as both remedy and hazard. Several of the challenges above point toward feedback, harvesting labels from analyst decisions and response outcomes, as the most promising general remedy, and it is worth examining why feedback is simultaneously the field's best hope and a subtle trap. Its promise is that it directly attacks labelling scarcity, environment-specificity, and concept drift at once: by learning continuously from the verdicts analysts render on real alerts in the real environment, an adaptive system acquires labels that are local, current, and free, and it can track drift as it happens rather than degrading until a scheduled retraining. The hazard is that the feedback loop observes only the alerts that reach analysts, and the alerts that reach analysts are exactly those the current system chose to surface. Alerts suppressed by the present model generate no labels, so the model never learns that its suppression was wrong, and existing blind spots are quietly entrenched rather than corrected. This selection bias means that a naively closed feedback loop optimises the system's agreement with its own past behaviour, which is not the same as optimising its correctness. Escaping the

trap requires deliberate counter-measures, sampling and surfacing a fraction of alerts the model would otherwise suppress, so that the loop retains a channel through which its own errors can become visible, and the design of feedback that improves rather than merely reinforces is an open and important problem.

The adversary as an active participant. A feature that distinguishes security from most domains in which automation is applied is that the environment is not merely changing but adversarial: some of the drift the system must cope with is deliberately induced by an opponent who studies the defence and adapts to evade it. This changes the character of every challenge discussed above. Concept drift is not only the benign evolution of normal behaviour but the calculated mutation of attacks to slip beneath detection thresholds. False negatives are not only the statistical residue of an imperfect model but the targeted product of an adversary crafting activity to resemble the benign. Even the feedback loop becomes an attack surface, since an opponent who understands that the system learns from analyst decisions may attempt to poison that learning by generating patterns designed to be labelled benign. The practical implication is that robustness in this setting cannot be assessed against a fixed distribution, because the distribution is chosen, in part, by someone trying to defeat the defender. This is a further reason why laboratory accuracy overstates operational readiness: the test set does not fight back, and the deployment does.

Contributions on financial analytics, planning, and enterprise decision systems inform the framing adopted here [157]–[160].

IX. FUTURE DIRECTIONS

Several directions follow directly from the gaps above.

First, analyst-in-the-loop and explainable triage deserve primacy. Given that triage is interpretive sense-making [9] and that adaptive classification already exploits analyst feedback [119], the natural progression is toward systems that present ranked candidates together with the evidence and reasoning behind their scores, so analysts can both trust and correct them. The design objective should be measured in decision quality and attention conserved, not alerts emitted.

Second, reproducible and realistic evaluation must extend beyond detection to the whole pipeline. The community needs continually refreshed datasets in the spirit of CICIDS2017 [124], shared benchmarks for correlation and triage, and open tooling of the kind demonstrated for log parsing [51], so that claims about triage and orchestration become as testable as claims about detection.

Third, drift-aware and feedback-driven learning should be treated as a first-class requirement rather than an afterthought, with architectures that expect models to age, that monitor their own degradation, and that incorporate response outcomes as training signal while guarding against the label biases that closed feedback loops introduce.

Fourth, standardised semantics such as ATT&CK [120] and the incident-handling lifecycle of NIST [121] should be pushed deeper into automation, so that correlation, prioritization, and orchestration reason over a shared model of adversary behaviour rather than over vendor-specific alert codes, improving both interpretability and portability.

Fifth, and cutting across the others, automation should be designed for operational resilience: calibrated to act autonomously only where upstream confidence is high, degrading gracefully when data sources fail or models drift, and preserving institutional knowledge in playbooks so that capability does not evaporate with staff turnover [3].

Sixth, the pipeline should be treated as an object of end-to-end measurement rather than a collection of separately optimised parts. Because the stages form a dependency chain, the quantity that ultimately matters is not the accuracy of any single stage but the quality and cost of the decisions that emerge at the end, per unit of analyst attention consumed. Instrumenting the whole pipeline to expose where signal is lost, where it is amplified into noise, and where analyst time is actually spent would let organisations identify their binding constraint empirically rather than by intuition, and would let researchers report the marginal contribution of a proposed technique to a realistic end-to-end objective rather than to an isolated benchmark. Such whole-pipeline instrumentation is technically modest but conceptually important, because it aligns

measurement with the systemic view that the review argues for throughout.

Seventh, and finally, the community should invest in the transfer and portability of triage and orchestration capability. Because the meaning of alerts is environment-specific, the knowledge encoded in triage models and response playbooks does not move readily between organisations, and this friction slows the diffusion of hard-won improvements. Research into representations that separate the portable core of a capability (the general logic of a response, the structure of a scoring model) from its environment-specific parameters, and into methods for adapting a capability to a new environment with modest local data, would let the field accumulate rather than perpetually rebuild. Standardised semantics such as a shared model of adversary behaviour are one route to this portability; another is the design of capabilities that expect to be specialised, carrying explicit hooks for the local context they must be fitted to. Progress here would compound the returns on every other direction above.

Contributions on public-sector accounting, audit, and financial reporting inform the framing adopted here [161], [162].

X. CONCLUSION

Automation of log collection and incident triage has moved from aspiration to operational necessity, driven by telemetry volumes that no human team can inspect unaided. This review organised the relevant literature into a five-stage pipeline (collection and normalization, detection, correlation, triage and prioritization, and orchestration) and showed that each stage now rests on a substantial body of technique, from automated log parsing and anomaly detection to alert correlation, adaptive classification, and playbook-driven response. The described taxonomy makes plain that these stages form a dependency chain whose overall reliability is set by its weakest link, and that the field's hardest problems (the base-rate fallacy, alert fatigue, labelling scarcity, concept drift, and weak evaluation realism) cut across stages rather than residing in any one. The most durable progress will come not from marginal gains in detection accuracy but from treating the SOC as a socio-technical system: designing automation that conserves analyst attention,

that is evaluated against realistic and reproducible benchmarks along its full length, and that learns continually from the humans it exists to support. For the practitioner and the researcher alike, the guiding question is not how many events a system can flag, but how well it enables correct decisions under an unrelenting flood of signal.

REFERENCES

- [1] Bhatt, S., Manadhata, P. K., & Zomlot, L. (2014). The operational role of security information and event management systems. *IEEE Security & Privacy*, 12(5), 35–41. <https://doi.org/10.1109/MSP.2014.103>
- [2] Axelsson, S. (2000). The base-rate fallacy and the difficulty of intrusion detection. *ACM Transactions on Information and System Security*, 3(3), 186–205. <https://doi.org/10.1145/357830.357849>
- [3] Sundaramurthy, S. C., Bardas, A. G., Case, J., Ou, X., Wesch, M., McHugh, J., & Rajagopalan, S. R. (2015). A human capital model for mitigating security analyst burnout. In *Proceedings of the Eleventh Symposium on Usable Privacy and Security (SOUPS 2015)* (pp. 347–359). USENIX Association. <https://www.usenix.org/conference/soups2015/proceedings/presentation/sundaramurthy>
- [4] Ejofodomi, O. A., Gideon, E. N., Oladipo, G. O., & Oshomah, E. R. (2014). Automated detection of architectural detection in mammograms using template matching. *International Journal of Biomedical Science and Engineering*, 2(1), 1–6.
- [5] Ilodigwe, L., & Adesemoye, A. C. (2019a). A critical review of health insurance financing mechanisms and provider payment reform strategies for balancing coverage expansion and financial protection in emerging markets. *International Journal of Health and Pharmaceutical Research*, 5(2), 31–55. <https://doi.org/10.56201/ijhpr.vol.5.no2.2019.pg31.55>
- [6] Ilodigwe, L., & Adesemoye, A. C. (2019b). From molecules to health systems: A conceptual model explaining why scientific innovation fails to improve patient outcomes

- without effective healthcare delivery infrastructure. *International Journal of Health and Pharmaceutical Research*, 5(2), 56-79. <https://doi.org/10.56201/ijhpr.vol.5.no2.2019.pg56.79>
- [7] Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering*, SE-13(2), 222-232. <https://doi.org/10.1109/TSE.1987.232894>
- [8] Roesch, M. (1999). Snort: Lightweight intrusion detection for networks. In *Proceedings of the 13th USENIX Conference on System Administration (LISA '99)* (pp. 229-238). USENIX Association. https://www.usenix.org/legacy/publications/library/proceedings/lisa99/full_papers/roesch/roesch.pdf
- [9] D'Amico, A., & Whitley, K. (2008). The real work of computer network defense analysts. In J. R. Goodall, G. Conti, & K.-L. Ma (Eds.), *VizSEC 2007: Proceedings of the Workshop on Visualization for Computer Security* (pp. 19-37). Springer. https://doi.org/10.1007/978-3-540-78243-8_2
- [10] Compaine, B. M. (Ed.). (2001). *The digital divide: Facing a crisis or creating a myth?* MIT Press.
- [11] Mossberger, K., Tolbert, C. J., & Stansbury, M. (2003). *Virtual inequality: Beyond the digital divide*. Georgetown University Press.
- [12] Attewell, P. (2001). The first and second digital divides. *Sociology of Education*, 74(3), 252-259. <https://doi.org/10.2307/2673277>
- [13] Helsper, E. J. (2012). A corresponding fields model for the links between social and digital exclusion. *Communication Theory*, 22(4), 403-426. <https://doi.org/10.1111/j.1468-2885.2012.01416.x>
- [14] van Deursen, A. J. A. M., & Helsper, E. J. (2015). The third-level digital divide: Who benefits most from being online? *Communication and Information Technologies Annual*, 10, 29-52. <https://doi.org/10.1108/S2050-206020150000010002>
- [15] Scheerder, A., van Deursen, A., & van Dijk, J. (2017). Determinants of Internet skills, uses and outcomes: A systematic review of the second- and third-level digital divide. *Telematics and Informatics*, 34(8), 1607-1624. <https://doi.org/10.1016/j.tele.2017.07.007>
- [16] Heeks, R. (2010). Do information and communication technologies (ICTs) contribute to development? *Journal of International Development*, 22(5), 625-640. <https://doi.org/10.1002/jid.1716>
- [17] Walsham, G. (2017). ICT4D research: Reflections on history and future agenda. *Information Technology for Development*, 23(1), 18-41. <https://doi.org/10.1080/02681102.2016.1246406>
- [18] Gunkel, D. J. (2003). Second thoughts: Toward a critique of the digital divide. *New Media & Society*, 5(4), 499-522. <https://doi.org/10.1177/146144480354003>
- [19] DiMaggio, P., Hargittai, E., Neuman, W. R., & Robinson, J. P. (2001). Social implications of the Internet. *Annual Review of Sociology*, 27, 307-336. <https://doi.org/10.1146/annurev.soc.27.1.307>
- [20] Wei, K.-K., Teo, H.-H., Chan, H. C., & Tan, B. C. Y. (2011). Conceptualizing and testing a social cognitive model of the digital divide. *Information Systems Research*, 22(1), 170-187. <https://doi.org/10.1287/isre.1090.0273>
- [21] Hilbert, M. (2011). The end justifies the definition: The manifold outlooks on the digital divide and their practical usefulness for policy-making. *Telecommunications Policy*, 35(8), 715-736. <https://doi.org/10.1016/j.telpol.2011.06.012>
- [22] Selwyn, N. (2016). *Is technology good for education?* Polity Press.
- [23] Sharples, M., Taylor, J., & Vavoula, G. (2010). A theory of learning for the mobile age. In *Medienbildung in neuen Kulturräumen* (pp. 87-99). VS Verlag. https://doi.org/10.1007/978-3-531-92133-4_6
- [24] Traxler, J. (2007). Defining, discussing and evaluating mobile learning. *International Review of Research in Open and Distributed Learning*, 8(2). <https://doi.org/10.19173/irrodl.v8i2.346>

- [25] Kukulska-Hulme, A. (2009). Will mobile learning change language learning? *ReCALL*, 21(2), 157–165. <https://doi.org/10.1017/S0958344009000202>
- [26] Donner, J. (2008). Research approaches to mobile use in the developing world: A review of the literature. *The Information Society*, 24(3), 140–159. <https://doi.org/10.1080/01972240802019970>
- [27] Heeks, R. (2008). ICT4D 2.0: The next phase of applying ICT for international development. *Computer*, 41(6), 26–33. <https://doi.org/10.1109/MC.2008.192>
- [28] Brown, J. S., & Adler, R. P. (2008). Minds on fire: Open education, the long tail, and Learning 2.0. *EDUCAUSE Review*, 43(1), 16–32.
- [29] Kozma, R. B. (2003). Technology and classroom practices: An international study. *Journal of Research on Technology in Education*, 36(1), 1–14. <https://doi.org/10.1080/15391523.2003.10782399>
- [30] Pittinsky, M. S. (2003). *The wired tower: Perspectives on the impact of the Internet on higher education*. Financial Times Prentice Hall.
- [31] Box, G. E. P., & Jenkins, G. M. (1976). *Time series analysis: Forecasting and control* (Rev. ed.). Holden-Day.
- [32] Armstrong, J. S. (Ed.). (2001). *Principles of forecasting: A handbook for researchers and practitioners*. Kluwer Academic. <https://doi.org/10.1007/978-0-306-47630-3>
- [33] De Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, 22(3), 443–473. <https://doi.org/10.1016/j.ijforecast.2006.01.001>
- [34] Gardner, E. S. (2006). Exponential smoothing: The state of the art: Part II. *International Journal of Forecasting*, 22(4), 637–666. <https://doi.org/10.1016/j.ijforecast.2006.03.005>
- [35] Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5–10. <https://doi.org/10.1016/j.ijforecast.2003.09.015>
- [36] Fildes, R., Nikolopoulos, K., Crone, S. F., & Syntetos, A. A. (2008). Forecasting and operational research: A review. *Journal of the Operational Research Society*, 59(9), 1150–1172. <https://doi.org/10.1057/palgrave.jors.2602597>
- [37] Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5), 495–503. <https://doi.org/10.1057/palgrave.jors.2601841>
- [38] Ahmed, N. K., Atiya, A. F., Gayar, N. E., & El-Shishiny, H. (2010). An empirical comparison of machine learning models for time series forecasting. *Econometric Reviews*, 29(5–6), 594–621. <https://doi.org/10.1080/07474938.2010.481556>
- [39] Crone, S. F., Hibon, M., & Nikolopoulos, K. (2011). Advances in forecasting with neural networks? Empirical evidence from the NN3 competition. *International Journal of Forecasting*, 27(3), 635–660. <https://doi.org/10.1016/j.ijforecast.2011.04.001>
- [40] Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and applications* (3rd ed.). Wiley.
- [41] Fildes, R., & Goodwin, P. (2007). Against your better judgment? How organizations can improve their use of management judgment in forecasting. *Interfaces*, 37(6), 570–576. <https://doi.org/10.1287/inte.1070.0309>
- [42] Chatfield, C. (2000). *Time-series forecasting*. Chapman & Hall/CRC.
- [43] Hyndman, R. J., Koehler, A. B., Snyder, R. D., & Grose, S. (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, 18(3), 439–454. [https://doi.org/10.1016/S0169-2070\(01\)00110-8](https://doi.org/10.1016/S0169-2070(01)00110-8)
- [44] Norris, P. (2001). Digital divide: Civic engagement, information poverty, and the

- Internet worldwide. Cambridge University Press.
<https://doi.org/10.1017/CBO9781139164887>
- [45] Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. *Iconic Research and Engineering Journals*, 2(6).
- [46] Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2019a). A conceptual model for ETL design and data integration in Salesforce-centric enterprise architectures. *Iconic Research and Engineering Journals*, 3(5).
- [47] Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2019b). A governance framework for Salesforce platform management in regulated healthcare environments. *Iconic Research and Engineering Journals*, 3(6).
- [48] Xu, W., Huang, L., Fox, A., Patterson, D., & Jordan, M. I. (2009). Detecting large-scale system problems by mining console logs. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (SOSP '09)* (pp. 117–132). ACM.
<https://doi.org/10.1145/1629575.1629587>
- [49] Vaarandi, R., & Pihelgas, M. (2015). LogCluster: A data clustering and pattern mining algorithm for event logs. In *Proceedings of the 2015 11th International Conference on Network and Service Management (CNSM)* (pp. 1–7). IEEE.
<https://doi.org/10.1109/CNSM.2015.7367331>
- [50] He, P., Zhu, J., Zheng, Z., & Lyu, M. R. (2017). Drain: An online log parsing approach with fixed depth tree. In *Proceedings of the 2017 IEEE International Conference on Web Services (ICWS)* (pp. 33–40). IEEE.
<https://doi.org/10.1109/ICWS.2017.13>
- [51] Zhu, J., He, S., Liu, J., He, P., Xie, Q., Zheng, Z., & Lyu, M. R. (2019). Tools and benchmarks for automated log parsing. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)* (pp. 121–130). IEEE.
<https://doi.org/10.1109/ICSE-SEIP.2019.00021>
- [52] Dosunmu, A. A., & Ogundele, P. O. (2019). Security audit and enterprise risk assessment frameworks for resilient information systems. *Iconic Research and Engineering Journals*, 3(5), 434-447.
<https://doi.org/10.64388/IREV3I5-1713225>
- [53] Agbabiaka, J., Okonkwo, C. S., Ogunwole, O., Mayo, W., & Okeke, O. T. (2019). Supply chain risk management model for EPC and gas processing projects. *Iconic Research and Engineering Journals*, 3(2), 968-980.
<https://doi.org/10.64388/IREV3I2-1713124>
- [54] Okonkwo, C. S., Ogunwole, O., & Okeke, O. T. (2018b). Framework for strategic procurement optimization in oil and gas operations. *Iconic Research and Engineering Journals*, 1(7), 153-168.
<https://doi.org/10.64388/IREV1I7-1713119>
- [55] Okonkwo, C. S., Ogunwole, O., & Okeke, O. T. (2018c). Model for inventory availability and plant uptime improvement in energy facilities. *Iconic Research and Engineering Journals*, 2(4), 160-172.
<https://doi.org/10.64388/IREV2I4-1713120>
- [56] Okonkwo, C. S., Ogunwole, O., Okeke, O. T., & Mayo, W. (2019). Conceptual framework for cost reduction through contract negotiation and vendor governance. *Iconic Research and Engineering Journals*, 2(9), 468-482.
<https://doi.org/10.64388/IREV2I9-1713121>
- [57] García-Teodoro, P., Díaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers & Security*, 28(1–2), 18–28.
<https://doi.org/10.1016/j.cose.2008.08.003>
- [58] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15.
<https://doi.org/10.1145/1541880.1541882>
- [59] Sperotto, A., Schaffrath, G., Sadre, R., Morariu, C., & Pras, A. (2010). An overview of IP flow-based intrusion detection. *IEEE Communications Surveys & Tutorials*, 12(3), 343–356.
<https://doi.org/10.1109/SURV.2010.032210.00054>

- [60] Du, M., Li, F., Zheng, G., & Srikumar, V. (2017). DeepLog: Anomaly detection and diagnosis from system logs through deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17)* (pp. 1285–1298). ACM. <https://doi.org/10.1145/3133956.3134015>
- [61] He, S., Zhu, J., He, P., & Lyu, M. R. (2016). Experience report: System log analysis for anomaly detection. In *Proceedings of the 2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE)* (pp. 207–218). IEEE. <https://doi.org/10.1109/ISSRE.2016.21>
- [62] Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy* (pp. 305–316). IEEE. <https://doi.org/10.1109/SP.2010.25>
- [63] Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176. <https://doi.org/10.1109/COMST.2015.2494502>
- [64] Valeur, F., Vigna, G., Kruegel, C., & Kemmerer, R. A. (2004). A comprehensive approach to intrusion detection alert correlation. *IEEE Transactions on Dependable and Secure Computing*, 1(3), 146–169. <https://doi.org/10.1109/TDSC.2004.21>
- [65] Ning, P., Cui, Y., & Reeves, D. S. (2002). Constructing attack scenarios through correlation of intrusion alerts. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS '02)* (pp. 245–254). ACM. <https://doi.org/10.1145/586110.586144>
- [66] Julisch, K. (2003). Clustering intrusion detection alarms to support root cause analysis. *ACM Transactions on Information and System Security*, 6(4), 443–471. <https://doi.org/10.1145/950191.950192>
- [67] Mbonu, I. S., Aliliele, C., Iwuanyanwu, U., & Oluoha, O. M. (2018). A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. *Iconic Research and Engineering Journals*, 2(2), 207–226. <https://doi.org/10.64388/IREV2I2-1714911>
- [68] Mbonu, I. S., Aliliele, C., Uzoka, E., & Oluoha, O. M. (2019a). A review of comparative data protection regulations and secure cloud implementation strategies across jurisdictions. *Iconic Research and Engineering Journals*, 2(9), 482–501. <https://doi.org/10.64388/IREV2I9-1714912>
- [69] Mbonu, I. S., Iwuanyanwu, U., Uzoka, E., & Oluoha, O. M. (2019b). Advances in enterprise log analytics and automated incident response architectures using Python and SIEM platforms. *Iconic Research and Engineering Journals*, 3(2), 1000–1019. <https://doi.org/10.64388/IREV3I2-1714915>
- [70] Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510. <https://doi.org/10.1016/j.ymssp.2005.09.012>
- [71] Susto, G. A., Schirru, A., Pampuri, S., McLoone, S., & Beghi, A. (2015). Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3), 812–820. <https://doi.org/10.1109/TII.2014.2349359>
- [72] Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. *Journal of Manufacturing Systems*, 48, 144–156. <https://doi.org/10.1016/j.jmsy.2018.01.003>
- [73] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
- [74] Stouffer, K., Lightman, S., Pillitteri, V., Abrams, M., & Hahn, A. (2015). Guide to Industrial Control Systems (ICS) Security. NIST SP 800-82.

- [75] Dragos, & Inc. (2017). TRISIS Malware: Analysis of Safety System Targeted Attack. Dragos Intelligence.
- [76] Papernot, N., McDaniel, P., Sinha, A., & Wellman, M. P. (2018). SoK: Security and privacy in machine learning. in Proc. IEEE EuroS&P, 399-414.
- [77] Sproul, W. D., Christie, D. J., & Carter, D. C. (2005). Control of reactive sputtering processes. Thin Solid Films, 491(1), 1–17. <https://doi.org/10.1016/j.tsf.2005.05.022>
- [78] Strijkmans, K., Schelfhout, R., & Depla, D. (2018). Tutorial: Hysteresis during the reactive magnetron sputtering process. Journal of Applied Physics, 124(24), 241101. <https://doi.org/10.1063/1.5042084>
- [79] Musil, J., Baroch, P., Vlček, J., Nam, K. H., & Han, J. G. (2005). Reactive magnetron sputtering of thin films: present status and trends. Thin Solid Films, 475(1), 208–218. <https://doi.org/10.1016/j.tsf.2004.07.041>
- [80] Sarakinos, K., Alami, J., & Konstantinidis, S. (2010). High power pulsed magnetron sputtering: A review on scientific and engineering state of the art. Surface and Coatings Technology, 204(11), 1661–1684. <https://doi.org/10.1016/j.surfcoat.2009.11.013>
- [81] Anders, A. (2017). Tutorial: Reactive high power impulse magnetron sputtering (R-HiPIMS). Journal of Applied Physics, 121(17), 171101. <https://doi.org/10.1063/1.4978350>
- [82] Anders, A. (2014). A review comparing cathodic arcs and high power impulse magnetron sputtering (HiPIMS). Surface and Coatings Technology, 257, 308–325. <https://doi.org/10.1016/j.surfcoat.2014.08.043>
- [83] Negri, E., Fumagalli, L., & Macchi, M. (2017). A Review of the Roles of Digital Twin in CPS-based Production Systems. Procedia Manufacturing, 11, 939–948. <https://doi.org/10.1016/j.promfg.2017.07.198>
- [84] Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital Twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnLine, 51(11), 1016–. <https://doi.org/10.1016/j.ifacol.2018.08.474>
- [85] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics, 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [86] Brunton, S. L., Proctor, J. L., & Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proceedings of the National Academy of Sciences, 113(15), 3932–3937. <https://doi.org/10.1073/pnas.1517384113>
- [87] Benner, P., Gugercin, S., & Willcox, K. (2015). A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. SIAM Review, 57(4), 483–531. <https://doi.org/10.1137/130932715>
- [88] Evensen, G. (2003). The Ensemble Kalman Filter: theoretical formulation and practical implementation. Ocean Dynamics, 53(4), 343–367. <https://doi.org/10.1007/s10236-0030036-9>
- [89] Mayne, D. Q. (2014). Model predictive control: Recent developments and future promise. Automatica, 50(12), 2967–2986. <https://doi.org/10.1016/j.automatica.2014.10.128>
- [90] Odejebi, O. D., Hammed, N. I., & Ahmed, K. S. (2019). Approximation complexity model for cloud-based database optimization problems. IRE Journals, 2(9), 1–10.
- [91] Ahmed, K. S., Odejebi, O. D., & Oshoba, T. O. (2019). Algorithmic model for constraint satisfaction in cloud network resource allocation. IRE Journals, 2(12), 516–532.
- [92] Dagodzo, D. (2018a). A conceptual framework for UAV integration into national power grid inspection programs. IRE Journals, 2(5), 391–412. <https://doi.org/10.64388/IREV2I51716082>
- [93] Pflug, A., Siemers, M., Melzig, T., Schäfer, L., & Bräuer, G. (2014). Simulation of linear magnetron discharges in 2D and 3D. Surface and Coatings Technology, 260, 411–416. <https://doi.org/10.1016/j.surfcoat.2014.09.042>

- [94] Nyberg, T., Högberg, H., Greczynski, G., & Berg, S. (2019). A simple model for non-saturated reactive sputtering processes. *Thin Solid Films*, 688, 137413. <https://doi.org/10.1016/j.tsf.2019.137413>
- [95] Berg, S., Särhammar, E., & Nyberg, T. (2014). Upgrading the 'Berg-model' for reactive sputtering processes. *Thin Solid Films*, 565, 186–192. <https://doi.org/10.1016/j.tsf.2014.02.063>
- [96] Elebe, O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. and digital identity verification framework for.
- [97] Srinidhi, B., Yan, J., & Bhargava, H. K. (2015). Effect of information security investments on firm performance. *Decision Support Systems*, 74, 1–15.
- [98] McKone, K. E., Schroeder, R. G., & Cua, K. O. (2001). The impact of total productive maintenance practices on manufacturing performance. *Journal of Operations Management*, 19(1), 39–58. [https://doi.org/10.1016/S0272-6963\(00\)00030-9](https://doi.org/10.1016/S0272-6963(00)00030-9)
- [99] Dal, B., Tugwell, P., & Greatbanks, R. (2000). Overall equipment effectiveness as a measure of operational improvement: A practical analysis. *International Journal of Operations & Production Management*, 20(12), 1488–1502. <https://doi.org/10.1108/01443570010355750>
- [100] Bamber, D., Sharp, C. J., & Hides, M. T. (1999). Factors affecting successful implementation of total productive maintenance: A UK manufacturing case study perspective. *Journal of Quality in Maintenance Engineering*, 5(3), 162–181. <https://doi.org/10.1108/13552519910282601>
- [101] McKone, K. E., Schroeder, R. G., & Cua, K. O. (1999). Total productive maintenance: A contextual view. *Journal of Operations Management*, 17(2), 123–144. [https://doi.org/10.1016/S0272-6963\(98\)00039-4](https://doi.org/10.1016/S0272-6963(98)00039-4)
- [102] Tezel, A., Koskela, L., & Tzortzopoulos, P. (2016). Visual management in production management: A literature synthesis. *Journal of Manufacturing Technology Management*, 27(6), 766–799. <https://doi.org/10.1108/JMTM-08-2015-0071>
- [103] Bagavathiappan, S., Lahiri, B. B., Saravanan, T., Philip, J., & Jayakumar, T. (2013). Infrared thermography for condition monitoring: A review. *Infrared Physics & Technology*, 60, 35–55. <https://doi.org/10.1016/j.infrared.2013.03.006>
- [104] Lee, J., Wu, F., Zhao, W., Ghaffari, M., Liao, L., & Siegel, D. (2014b). Prognostics and health management design for rotary machinery systems: Reviews, methodology and applications. *Mechanical Systems and Signal Processing*, 42(1), 314–334. <https://doi.org/10.1016/j.ymssp.2013.06.004>
- [105] Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing*, 104, 799–834. <https://doi.org/10.1016/j.ymssp.2017.11.016>
- [106] Saxena, A., Goebel, K., Simon, D., & Eklund, N. (2008). Damage propagation modeling for aircraft engine run-to-failure simulation. in *Proc. Int. Conf. Prognostics and Health Management (PHM)*, 1–9. <https://doi.org/10.1109/PHM.2008.4711414>
- [107] Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., & Gao, R. X. (2019). Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115, 213–237. <https://doi.org/10.1016/j.ymssp.2018.05.050>
- [108] Lee, J., Kao, H.-A., & Yang, S. (2014a). Service innovation and smart analytics for Industry 4.0 and big data environment. *Procedia CIRP*, 16, 3–8. <https://doi.org/10.1016/j.procir.2014.02.001>
- [109] Tao, F., Zhang, M., Liu, Y., & Nee, A. Y. C. (2018b). Digital twin driven prognostics and health management for complex equipment. *CIRP Annals*, 67(1), 169–172. <https://doi.org/10.1016/j.cirp.2018.04.055>
- [110] Selcuk, S. (2017). Predictive maintenance, its implementation and latest trends. *Proceedings of the Institution of Mechanical Engineers*, 231(9), 1670–1679. <https://doi.org/10.1177/0954405415601640>

- [111] Okonkwo, C. S., Ogunwale, O., & Okeke, O. T. (2018a). Model for inventory availability and plant uptime improvement in energy facilities. *IRE Journals*, 2(4), 160–172. <https://doi.org/10.64388/IREV2I4-1713120>
- [112] Peng, Y., Dong, M., & Zuo, M. J. (2010). Current status of machine prognostics in condition-based maintenance: A review. *The International Journal of Advanced Manufacturing Technology*, 50(1), 297–313. <https://doi.org/10.1007/s00170-009-2482-0>
- [113] Liu, H.-C., Liu, L., & Liu, N. (2013). Risk evaluation approaches in failure mode and effects analysis: A literature review. *Expert Systems with Applications*, 40(2), 828–838. <https://doi.org/10.1016/j.eswa.2012.08.010>
- [114] Antoni, J. (2007). Fast computation of the kurtogram for the detection of transient faults. *Mechanical Systems and Signal Processing*, 21(1), 108–124. <https://doi.org/10.1016/j.ymssp.2005.12.002>
- [115] Lei, Y., Lin, J., Zuo, M. J., & He, Z. (2014). Condition monitoring and fault diagnosis of planetary gearboxes: A review. *Measurement*, 48, 292–305. <https://doi.org/10.1016/j.measurement.2013.11.012>
- [116] Lei, Y., Jia, F., Lin, J., Xing, S., & Ding, S. X. (2016). An intelligent fault diagnosis method using unsupervised feature learning towards mechanical big data. *IEEE Transactions on Industrial Electronics*, 63(5), 3137–3147. <https://doi.org/10.1109/TIE.2016.2519325>
- [117] Wen, L., Li, X., Gao, L., & Zhang, Y. (2018). A new convolutional neural network-based data-driven fault diagnosis method. *IEEE Transactions on Industrial Electronics*, 65(7), 5990–. <https://doi.org/10.1109/TIE.2017.2774777>
- [118] Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018a). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, 94(9), 3563–3576. <https://doi.org/10.1007/s00170-0170233-1>
- [119] Pietraszek, T. (2004). Using adaptive alert classification to reduce false positives in intrusion detection. In E. Jonsson, A. Valdes, & M. Almgren (Eds.), *Recent Advances in Intrusion Detection (RAID 2004)* (Lecture Notes in Computer Science, Vol. 3224, pp. 102–124). Springer. https://doi.org/10.1007/978-3-540-30143-1_6
- [120] Strom, B. E., Applebaum, A., Miller, D. P., Nickels, K. C., Pennington, A. G., & Thomas, C. B. (2018). *MITRE ATT&CK: Design and philosophy* (Technical Report MP180360). The MITRE Corporation. https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf
- [121] Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). *Computer security incident handling guide* (NIST Special Publication 800-61 Revision 2). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-61r2>
- [122] Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2018). Lessons learned from offline assessment of security-critical systems: The case of Microsoft Active Directory. *Iconic Research and Engineering Journals*, 2(6), 277–299. <https://doi.org/10.64388/IREV2I6-1717205>
- [123] Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2019). Implementation of Active Directory for efficient management of enterprise networks. *Iconic Research and Engineering Journals*, 3(4), 608–627. <https://doi.org/10.64388/IREV3I4-1717206>
- [124] Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP)* (pp. 108–116). SciTePress. <https://doi.org/10.5220/0006639801080116>
- [125] Tavallaee, M., Bagheri, E., Lu, W., & Ghorbani, A. A. (2009). A detailed analysis of the KDD CUP 99 data set. In *Proceedings of the 2009 IEEE Symposium on Computational Intelligence for Security and Defence Applications (CISDA)* (pp. 1–6). IEEE. <https://doi.org/10.1109/CISDA.2009.5356528>

- [126] Bola-Sadipe, D., Aliliele, C., & Eboh, E. E. (2019). Conceptual model for strategic accounting systems strengthening financial transparency and regulatory compliance. *IRE Journals*, 3(6), 535-564. <https://doi.org/10.64388/IREV3I6-1715327>
- [127] Akeju, B., Edivri, J., Ogbale, J. I., Okoruwa, P. O., Fadayomi, O., & Abolaji, T. O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. *Iconic Research and Engineering Journals*, 1(9), 476-492. <https://doi.org/10.64388/IREV1I9-1713778>
- [128] Fadayomi, O., Abolaji, T. O., Edivri, J., Ogbale, J. I., Okoruwa, P. O., & Akeju, B. (2019). Risk-based cybersecurity assurance and data availability limitations, advances and future research opportunities. *Iconic Research and Engineering Journals*, 2(12), 602-617. <https://doi.org/10.64388/IREV2I12-1713779>
- [129] Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2018). Developing sustainable diagnostic laboratory infrastructure models for emerging and resource constrained health systems. *Iconic Research and Engineering Journals*, 1(8), 118-132. <https://doi.org/10.64388/IREV1I8-1713586>
- [130] Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2019). Capital project delivery models for high risk healthcare infrastructure in developing national health systems. *Iconic Research and Engineering Journals*, 2(10), 626-649. <https://doi.org/10.64388/IREV2I10-1713588>
- [131] Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2018). Optimizing laboratory spatial planning strategies to improve diagnostic accuracy, safety, and clinical throughput. *Iconic Research and Engineering Journals*, 2(1), 87-113. <https://doi.org/10.64388/IREV1I7-1713587>
- [132] Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2019). Regulatory compliant design systems for molecular and pathology laboratories in highly controlled environments. *Iconic Research and Engineering Journals*, 3(4), 607-631. <https://doi.org/10.64388/IREV3I4-1713589>
- [133] Arumosoye, O. M., & Obriki, O. D. (2018). Development of an integrated heat stress risk conceptual model for industrial operations in extreme environments. *Iconic Research and Engineering Journals*, 1(12), 141-160. <https://doi.org/10.64388/IREV1I12-1714415>
- [134] Arumosoye, O. M., & Obriki, O. D. (2019). Systematic review of near-miss and hazard observation data utilization in industrial safety management. *Iconic Research and Engineering Journals*, 3(2), 981-999. <https://doi.org/10.64388/IREV3I2-1714417>
- [135] Obogo, S. F., Ozobu, C. O., & Uduokhai, D. O. (2019a). Advances in leadership driven safety culture transformation in large construction workforces. *Iconic Research and Engineering Journals*, 3(5), 507-523. <https://doi.org/10.64388/IREV3I5-1715497>
- [136] Obogo, S. F., Ozobu, C. O., & Uduokhai, D. O. (2019b). Development of a construction safety risk governance model for multi contractor high rise projects. *Iconic Research and Engineering Journals*, 2(9), 502-522. <https://doi.org/10.64388/IREV2I9-1715495>
- [137] Obogo, S. F., Uduokhai, D. O., & Ozobu, C. O. (2019c). Systematic review of hazard identification and risk control practices in urban construction projects. *Iconic Research and Engineering Journals*, 2(12), 666-681. <https://doi.org/10.64388/IREV2I12-1715496>
- [138] Obriki, O. D., & Arumosoye, O. M. (2018). Conceptual modeling of data-driven occupational safety risk control in large-scale energy infrastructure projects. *Iconic Research and Engineering Journals*, 1(7), 169-189. <https://doi.org/10.64388/IREV1I7-1714414>
- [139] Obriki, O. D., & Arumosoye, O. M. (2019). A conceptual framework linking management safety walkthrough frequency and coverage to safety culture outcomes in mega projects. *Iconic Research and Engineering Journals*, 2(8), 355-374. <https://doi.org/10.64388/IREV2I8-1714416>
- [140] Dagodzo, D. (2018b). A conceptual framework for UAV integration into national power grid inspection programs. *Iconic Research and Engineering Journals*, 2(5), 391-412. <https://doi.org/10.64388/IREV2I5-1716082>

- [141] Dagodzo, D. (2018c). A review of UAV applications in electrical transmission line inspection: Methods, technologies, and challenges. *Iconic Research and Engineering Journals*, 2(6), 234-254. <https://doi.org/10.64388/IREV2I6-1716083>
- [142] Sunday, E. A., & Omoegun, G. O. (2018). Integrating solar power solutions in small-scale manufacturing industries in Nigeria. *International Journal of Scientific Research in Science, Engineering and Technology*, 4(8), 832-853.
- [143] Sunday, E. A., & Omoegun, G. O. (2019). Optimizing electrical load distribution for hybrid solar installations in developing economies. *International Journal of Scientific Research in Mechanical and Materials Engineering*, 3(6), 27-47.
- [144] Sunday, E. A., Omoegun, G. O., Essien, M. A., & Oluokun, O. A. (2019). Thermodynamic efficiency and control strategies in residential air conditioning systems. *International Journal of Scientific Research in Civil Engineering*, 3(3), 52-76.
- [145] Bobga, M. A., Boakye, K., Ogbona, C. S., & Yeboah, T. J. (2018). Pedagogical strategies for teaching students with learning difficulties in resource-constrained schools. *Iconic Research and Engineering Journals*, 2(4), 173-194. <https://doi.org/10.64388/IREV2I4-1714176>
- [146] Yeboah, T. J., Bobga, M. A., Boakye, K., & Ogbona, C. S. (2019). Professional development and teacher competence in managing exceptional learners: A Ghanaian perspective. *Iconic Research and Engineering Journals*, 2(12), 618-636. <https://doi.org/10.64388/IREV2I12-1714177>
- [147] Amayo, E. B. (2017). Multi algorithm of loss objective function of a single phase induction motor. *International Journal of Development Research*, 7(5), 12805-12810.
- [148] Amayo, E. B., & Popoola, T. T. (2017a). Multi algorithm of weight objective function of a single phase induction motor. *International Journal of Trend in Research and Development*, 4(2), 184-188.
- [149] Amayo, E. B., & Popoola, T. T. (2017b). Scientific study of telecommunication deployment in achieving quality network. *International Journal of Trend in Research and Development*, 4(5), 311-314.
- [150] Amayo, E. B., Ubeku, E., & Oshevire, P. (2015). Multi algorithm of a single objective function of a single phase induction motor. *Journal of Multidisciplinary Engineering Science and Technology*, 2(12), 3400-3403.
- [151] Oshevire, P., Eyenubo, O. J., & Amayo, B. (2017). Voltage control in the presence of distributed generation. *ATBU Journal of Science, Technology and Education*, 5(2), 165-173.
- [152] Adeyelu, O. O. (2019). An adaptive safety management system implementation model for aerodromes in developing economy contexts. *Iconic Research and Engineering Journals*, 3(4), 628-654. <https://doi.org/10.64388/IREV3I4-1718479>
- [153] Akomolafe, O., & Agu, M. U. (2018). A conceptual model for enhancing internal audit quality through technology-enabled risk assessment frameworks. *Iconic Research and Engineering Journals*, 1(9), 458-475.
- [154] Akomolafe, O., & Agu, M. U. (2019a). A conceptual framework for developing risk-based internal control models in the insurance and banking sectors. *Iconic Research and Engineering Journals*, 2(8), 335-354.
- [155] Akomolafe, O., & Agu, M. U. (2019b). A review of data-driven risk evaluation models for emerging market financial institutions. *Iconic Research and Engineering Journals*, 3(6), 433-448.
- [156] Akomolafe, O., & Agu, M. U. (2019c). Advances in financial resilience through integrated governance and compliance strategies. *Iconic Research and Engineering Journals*, 2(10), 607-620.
- [157] Lawal, O. A., & Oduleye, T. E. (2018a). A conceptual model for financial analytics driven enterprise value creation in technology firms. *Iconic Research and Engineering Journals*, 2(2).
- [158] Lawal, O. A., & Oduleye, T. E. (2018b). A review and conceptual framework for tax governance and cross-border compliance

- analytics. *Iconic Research and Engineering Journals*, 2(5).
- [159] Lawal, O. A., & Oduleye, T. E. (2019a). A conceptual risk assessment model for transfer pricing in multinational corporations. *Iconic Research and Engineering Journals*, 2(12).
- [160] Lawal, O. A., & Oduleye, T. E. (2019b). Conceptualizing data driven executive decision systems for strategic financial planning. *Iconic Research and Engineering Journals*, 3(3).
- [161] Dogbatsey, E. A., & Ebhojie, O. (2018). Budget compliance and statutory reporting in Sub-Saharan Africa: A systematic review of frameworks, gaps, and reform pathways. Zenodo.
<https://doi.org/10.5281/zenodo.20446859>
- [162] Dogbatsey, E. A., Ebhojie, O., & Oyeleye, A. O. (2019). Reconciliation control and financial workflow redesign in emerging market organizations: A conceptual framework. Zenodo.
<https://doi.org/10.5281/zenodo.20447103>