The Eastasouth Journal of Information System and Computer Science
Vol. 2, No. 02, December 2024, pp. 305-322
ISSN: 3025-566X, DOI: 10.58812/esiscs.v2i02

Intelligent Cybersecurity Frameworks for Data Protection in Cloud-Integrated Management Information Systems

Kesavan Sundara Mudaliyar[1], S. Satish Kumar[2]
[1] Department of Computer Science and Engineering, AMET University, Chennai, India
[2] Associate Professor, AMET University, Chennai, India

Article Info

Article history:

Received Dec 01, 2024
Revised Dec 18, 2024
Accepted Dec 30, 2024

Keywords:

Artificial Intelligence;
Cloud Computing;
Cybersecurity Framework;
Data Governance;
Federated Learning;
Machine Learning;
Management Information Systems;
Systematic Literature Review;
Threat Intelligence

ABSTRACT

Enterprise workloads keep moving to public, private, and hybrid cloud environments, and that shift is widening the attack surface available to would-be intruders just as organizations lean harder on Management Information Systems (MIS) to run day-to-day decisions. This paper works through fifteen recent studies sitting at the crossing point of artificial intelligence (AI), cybersecurity, cloud computing, and MIS governance, and tries to say something coherent about what they add up to. Rather than proposing and testing one new tool, the review pulls out the themes that keep resurfacing, AI and machine-learning-based threat detection, cyber threat intelligence, big-data analytics, data governance, federated and privacy-preserving learning, sustainable data-center design, and the human side of security that technical papers tend to skip and organizes them into a five-layer conceptual framework meant to help later empirical work. The review follows an explicit search, screening, and synthesis process, described in Section II; each proposed framework layer is traced back, in the discussion itself, to the specific literature that motivates it, and a thematic distribution chart shows how attention is split across sub-topics in the corpus. What comes out of this is that detection capability and MIS governance are comparatively well covered, while a handful of cross-cutting issues, explainability at the implementation level, how employees behave once AI is in the loop, energy-aware security operations, and the practical limits of federated learning, are thinner than their real-world importance would suggest. The paper closes by naming its own limitations and laying out where empirical work still needs to happen.

Corresponding Author:

Name: Kesavan Sundara Mudaliyar
Institution: Department of Computer Science and Engineering, AMET University, Chennai, India
Email: kesevan2007@gmail.com

1. INTRODUCTION

Cloud computing is no longer a peripheral IT option for most organizations; it is the default. Firms lean on Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS) offerings not just to host applications but to run the Management Information Systems (MIS) that day-to-day decisions depend on. Early work on cloud computing framed it as a long-promised utility model finally becoming practical [20], [21], and that framing still holds, but it comes with a cost: distributed, multi-tenant, API-driven cloud environments look nothing like the on-premises data centers

the security playbook was originally written for. The result is a data-protection problem that is technical and organizational at once. Technical, because the attack surface has genuinely changed shape. Organizational, because the very data that needs protection from outside and inside threats is also the data decision-makers are relying on to do their jobs.

A fair amount of recent research picks at pieces of this problem from different angles. Some of it stays close to the technical question of how AI and machine learning (ML) can catch intrusions, anomalies, and emerging threats in cloud and big-data settings [1], [6], [7], [10]. Other work treats cybersecurity as one input into a wider MIS governance and decision-support picture, asking how explainable AI, predictive analytics, and data governance practices shape outcomes like supply-chain resilience, project delivery, and sustainable growth [2], [3], [4], [5], [9], [11], [14]. A smaller slice looks at the human and infrastructural sides that purely algorithmic papers tend to leave out, employee training and behavior [13], and the energy trade-offs of running secure data centers at scale [8]. And a newer strand applies federated learning so organizations can build shared threat intelligence without pooling sensitive data in one place, which speaks directly to the tension between collective defense and data privacy [15].

Each of these studies is useful on its own but narrow. Put together, though, they start to sketch a more joined-up way of thinking about cybersecurity in cloud-integrated MIS, one where detection, governance, human behavior, infrastructure efficiency, and privacy-preserving collaboration are treated as layers of the same problem rather than five separate research programs. Broader cloud-security surveys and broader MIS-governance surveys tend to keep these apart, each writing for its own audience; what this review tries to add is a joint reading, using a fixed set of fifteen recent (2023–2024) sources as a manageable sample to test whether that joint reading is actually worth doing.

Concretely, the paper does four things: (1) works through fifteen recent peer-reviewed and conference sources from 2023–2024 on AI-driven cybersecurity, cloud infrastructure, and MIS governance, following the search and screening approach laid out in Section II; (2) sorts the literature into five thematic clusters and measures how research attention is actually distributed across them; (3) builds a five-layer conceptual framework out of the architectural ideas that keep recurring across the sources, tracing each layer back to the literature behind it; and (4) points to specific gaps, implementation-level explainability, human factors, energy-aware operations, and the real limits of federated learning, that the existing literature has not really settled.

The rest of the paper runs as follows. Section II lays out the review methodology. Section III covers foundational cloud-security concepts that the reviewed literature tends to assume rather than explain. Section IV is the review itself, split into five thematic subsections. Section V pulls together a comparative discussion and a thematic distribution figure. Section VI presents the proposed conceptual framework. Section VII discusses open challenges. Section VIII lays out future directions. Section IX is honest about this review's own limitations, and Section X closes things out.

One more thing worth saying up front: none of this is happening in a vacuum. The economic case for cloud computing, pay-as-you-go pricing, elastic scaling, no large up-front capital outlay for hardware, was made well before AI-driven security tooling existed, and it is precisely that economic case that pulled workloads, and the sensitive data attached to them, into shared, multi-tenant environments in the first place [20], [21]. The fifteen sources reviewed here are, in a sense, downstream of that earlier economic shift: they are trying to catch up, on the security and governance side, to a migration that happened mostly for reasons that had nothing to do with security.

2. METHODOLOGY

This review follows a scaled-down, transparent version of standard systematic-review practice, sized to fit the fixed corpus of fifteen sources supplied for review. The steps below explain how that corpus was scoped, screened, and pulled together, so the thematic clusters and framework that follow read as the output of a defined process rather than an impressionistic summary.

2.1 Scope and Corpus

The corpus behind this paper is fifteen sources published in 2023 and 2024, spread across peer-reviewed journals (Journal of Information Systems Engineering and Management, Journal of Computer Science and Technology Studies, Educational Administration: Theory and Practice, Computer Fraud & Security, Fuel Cells Bulletin, among others), one IEEE conference proceedings volume, and an open-access regional journal (The Eastasouth Journal of Information System and Computer Science). All fifteen touch on some mix of artificial intelligence, cybersecurity, cloud computing, big data, or management information systems, which is what sets the scope. For background and terminology, this corpus is supplemented by seven independently verified foundational references on cloud computing, cloud security, and federated learning [16] – [22], cited throughout Section III; these seven sit outside the fifteen-source review corpus and are not folded into the thematic counts in Fig. 1. Because the primary corpus was handed over as a fixed reference list rather than built through an open-ended database search, this review makes no claim to database-level exhaustiveness, it does not claim to have covered IEEE Xplore, ACM Digital Library, or Scopus in full for the relevant window. That constraint is picked back up as a limitation in Section IX.

2.2 Screening and Inclusion

Each source was read for its title, its abstract-level focus, and its venue, to work out (i) whether its main concern is technical detection, MIS/governance, human factors, or infrastructure, and (ii) which of the other categories, if any, shows up as a secondary concern. All fifteen supplied sources stayed in the review, none were dropped, because all of them sit within the defined scope. Screening here functioned as categorization rather than filtering, which fits the review's purpose: synthesizing and building a framework across a bounded, pre-specified corpus rather than trying to cover an open literature exhaustively.

2.3 Thematic Synthesis Procedure

After categorization, sources sharing a primary or secondary focus were grouped into thematic clusters inductively, the clusters were allowed to emerge from what was in the corpus rather than being decided in advance. That process produced the five clusters used in Section IV: AI/ML-driven detection and intelligence; cloud infrastructure and sustainability; MIS governance and decision support; human and organizational factors; privacy-preserving and federated approaches. Where a source clearly spans two clusters, a paper on federated learning that is also about cloud-based threat intelligence, say, it is discussed under its primary cluster and cross-referenced under the secondary one, which is why the counts behind Fig. 1 allow a single source to be counted more than once.

2.4 Framework Derivation

The five-layer conceptual framework in Section VI came out of looking for architectural pieces that show up, in some form, across more than one cluster, several sources independently describe a data-collection stage, an analytics stage, and a decision/response stage, even while using different terminology and working in different application domains. Fig. 2 is essentially the union of those recurring pieces, arranged into one layered structure and traced back, layer by layer, to the specific

sources that motivate it, as described in Section VI. No new experimental data went into validating this framework; it is offered as a conceptual synthesis, and testing it is left to the future work discussed in Section VIII.

3. BACKGROUND AND FOUNDATIONAL CONCEPTS

Before getting into the review proper, it helps to walk through four ideas that keep showing up, explicitly or implicitly, across the reviewed sources: the shared-responsibility model and baseline controls; the move from signature-based to AI-driven detection; the reframing of security data as an MIS asset; and the tension between centralization and privacy. Because the fifteen-source corpus described in Section II is aimed squarely at AI, MIS, and cloud governance, this section leans on several widely cited foundational references on cloud computing and cloud security more broadly, early framing papers on cloud computing as a computing utility [20], [21], the NIST definition of cloud computing [16], well-known surveys of cloud security issues [17], [18], [22], and a foundational reference on federated learning [19], to ground terms that the domain-specific corpus mostly assumes rather than defines.

3.1 The Shared-Responsibility Model and Baseline Controls

Cloud computing is most often defined, following the National Institute of Standards and Technology, as a model for convenient, on-demand network access to a shared pool of configurable computing resources that can be provisioned quickly with little management overhead, organized around five essential characteristics, three service models (IaaS, PaaS, SaaS), and four deployment models (public, private, community, hybrid) [16]. Early framing work described this as computing finally becoming available the way electricity or water is, a utility you draw on rather than infrastructure you own and pointed out that this shift changes who is responsible for what [20], [21]. Cloud data protection is usually discussed through the shared-responsibility model: the cloud provider secures the underlying infrastructure (physical data centers, host virtualization, core networking), while the tenant organization is on the hook for identities, data, and application-layer configuration, a division examined in detail by early foundational work on cloud security architecture [18]. Within that tenant side of the line, a handful of controls come up again and again in the literature: identity and access management (IAM) and least-privilege enforcement, so a compromised credential cannot be used to move laterally beyond its intended scope; encryption at rest and in transit, so intercepted or exfiltrated data stays useless without the right keys; network segmentation, which limits how far an intrusion can spread once a single host or container is compromised; and continuous monitoring through Security Information and Event Management (SIEM) platforms that pull log data from across the environment into one place for analysis. Misconfigured versions of these controls, an overly permissive IAM role, an unencrypted storage bucket, remain, across the broader cloud-security literature, one of the most common root causes of cloud data exposure, which is part of why several of the reviewed sources treat governance and configuration discipline as a precondition for anything more advanced [3], [4].

3.2 From Signature-Based to AI-Driven Detection

Sitting on top of those baseline controls, the reviewed literature keeps returning to a shift from rule-based or signature-based detection toward AI- and ML-based anomaly and intrusion detection, driven by the fact that novel and polymorphic attacks routinely slip past static signature databases [1], [6], [10]. Signature-based systems compare observed traffic or file characteristics against a database of known-bad patterns. They are cheap to run and rarely produce false positives on known threats,

but they simply cannot catch something that does not match an existing signature. AI/ML-based approaches take a different route: they learn a model of "normal" behavior, through supervised classification trained on labeled attack data, or unsupervised anomaly detection trained only on normal traffic and flag departures from that model as suspicious. That trade swaps signature matching's low false-positive rate for broader coverage of novel threats, at the cost of needing careful tuning so analysts are not buried in false alarms, a trade-off that several reviewed sources tackle head-on through cognitive- or intelligence-layer designs meant to prioritize and contextualize alerts rather than just flagging them [7], [12].

Worth noting: the reviewed literature does not treat AI/ML-based detection as a wholesale replacement for signature-based methods so much as a complement to them. Kaur et al.'s survey of the shifting threat landscape implicitly backs a layered detection posture, where cheap signature matching filters out the large volume of well-understood threats and frees the more expensive AI/ML analysis for whatever survives that first pass [10]. That layered posture also lines up with older, service-model-specific work cataloguing distinct risk profiles across IaaS, PaaS, and SaaS deployments, published well before AI-driven detection became a common proposed fix work the newer, AI-oriented sources in this review's corpus arguably extend rather than replace [17]. This layered posture is built into Layer 3 of the framework proposed in Section VI, which is described as covering multiple detection techniques rather than one model type.

3.3 Cybersecurity Data as an MIS Asset

A second recurring move in the literature is treating cybersecurity telemetry not just as an operational log for security analysts to review, but as an input feeding the same decision-support dashboards, predictive models, and governance processes used for other business analytics [2], [3], [5], [9]. That matters because it implies cybersecurity maturity and MIS/analytics maturity are not independent of each other: an organization with weak data governance is unlikely to get much out of even a technically solid AI threat-detection system, because its output will not reliably reach the people who need to act on it, will not get reconciled against other operational data, and will not be kept in a form that survives an audit. Several reviewed sources make, in effect, the same point, that an organization's MIS governance maturity is a decent leading indicator of how much value it will get from investing in AI-driven security tooling [3], [4], [11].

This reframing changes how success gets measured, too. A narrow technical view of cybersecurity might track detection accuracy or mean time to respond. An MIS-oriented view, in line with how Goffer et al. and Hossin et al. treat analytics as a driver of broader organizational outcomes [5], [9], would also ask whether security telemetry is reaching and being acted on by, the decision-making processes it is supposed to inform. That is different, and this review would argue under-measured, question.

3.4 Centralization Versus Privacy

A third idea worth flagging is the tension between centralization and privacy. Pooling security telemetry across business units, or across cooperating organizations, generally makes ML-based detection models statistically stronger, since rarer attack patterns become easier to spot in a larger dataset. But pooling also concentrates sensitive data in a way that can itself become an attack target or a regulatory liability, particularly for organizations bound by data-residency or sector-specific privacy rules concerns already visible in early public-cloud security analyses [22]. Federated learning training is a shared model across distributed data sources without moving the raw data, as laid out in foundational

work distinguishing horizontal, vertical, and transfer federated-learning architectures [19] is offered in the literature as one way to resolve that tension [15].

This is not a purely technical trade-off; it has a governance side too, which loops back to Section III-C. Choosing a federated architecture over a centralized one is, in effect, a data-sharing policy decision made by governance stakeholders rather than a modeling choice made solely by a data-science team, and several sources in the governance cluster covered in Section IV-C seem to assume this kind of cross-functional decision-making capacity already exists inside the organizations they describe [2], [3]. With that groundwork laid, Section IV turns to the fifteen sources themselves.

4. SYSTEMATIC REVIEW OF LITERATURE

The fifteen reviewed sources were sorted into five non-exclusive thematic clusters based on their stated focus, methodological orientation, and application domain, following the process described in Section II-C and quantified in Fig. 1. Several sources cut across more than one cluster; where that happens, they are discussed under their primary emphasis and cross-referenced elsewhere.

4.1 AI- and ML-Driven Threat Detection and Cyber Threat Intelligence

The biggest cluster of reviewed work applies AI and ML techniques directly to threat detection and threat intelligence in cloud and big-data settings. Chakraborty et al. put forward an intelligent cybersecurity framework aimed squarely at cloud data protection, treating AI-based analysis as the core detection mechanism inside a broader protective architecture [1]. That framing detection as one piece of a layered architecture rather than a standalone tool echoes across most of the other sources in this cluster and shapes the layered framework proposed in Section VI.

Hasan et al. look at how big-data analytics can be built into MIS platforms to sharpen cybersecurity threat detection and response, arguing that the sheer volume and speed of modern telemetry make manual triage a bottleneck without automated, data-driven support [6]. This connects directly back to Section III-B's point about the shift toward AI-driven detection: as long and traffic volume grows, purely manual or rule-based triage stops being a workable option, and analytics-driven prioritization goes from nice-to-have to structural necessity for keeping response times reasonable.

Hossain et al. push this further with what they call "cognitive cyber defense," combining AI and MIS through big-data and cloud frameworks to support next-generation digital resilience, and their point is that detection alone does not get you very far without a cognitive layer to contextualize and prioritize what it finds [7]. Naming that layer explicitly is one of the clearer articulations in the reviewed corpus of the distinction this review draws between Layer 3 (Intelligent Analytics) and Layer 4 (Decision and Orchestration).

Kaur et al. takes a wider-angle view of the threat landscape itself, cataloguing advanced cyber threats and surveying strategic and emerging approaches to cybersecurity innovation, which gives useful context for why AI-based detection has become necessary rather than just a nice upgrade [10]. Their contribution is methodologically different from the rest of this cluster, which tends toward proposing specific frameworks; here it functions as background for the motivation laid out in Section I rather than as a source for a specific architectural piece.

Orthi et al. push the AI-driven detection literature toward the intelligence side of things, treating AI-driven cyber threat intelligence as itself a form of management information system,

one that ties cybersecurity governance to IT project management to support organizational resilience, rather than treating threat intelligence as a purely technical, analyst-facing product [12]. This is one of the clearest bridges in the corpus between the detection cluster and the governance cluster in Section IV-C, and it shows up in how Layer 4 of the proposed framework is designed to cut across both.

Finally, Tuhin et al. link this cluster to the privacy-preserving one discussed below by proposing federated machine learning specifically for privacy-preserving cyber threat intelligence across global cloud-integrated MIS platforms [15]. Taken together, these six sources point to a sub-field that is maturing but still fragmented. Detection accuracy and model architecture get plenty of attention, but the sources vary a lot in how explicitly they deal with model explainability, false-positive management, and the handoff between an AI system's output and whatever human or automated response comes next, a gap picked up again in Section VII.

It is worth setting this cluster against the older, non-AI-specific cloud-security literature summarized in Section III. Foundational surveys like Subashini and Kavitha's and Zissis and Lekkas's catalogued cloud-specific risks data breaches, insecure interfaces, account hijacking, multi-tenancy exposure, mostly in terms of architectural controls and policy rather than adaptive, learning-based detection [17], [18]. The AI/ML-driven detection cluster reviewed here reads as this older literature's technical answer to those earlier-flagged risks: where the foundational surveys establish what can go wrong in a cloud environment, the sources in this cluster try to operationalize continuous, adaptive detection of when it actually is going wrong, which supports the broader argument running through this review that the newer, domain-specific literature builds on the older architectural and policy-oriented cloud-security work rather than displacing it.

4.2 Cloud Infrastructure, Data-Center Security and Sustainability

A second, smaller cluster looks at the infrastructural substrate underneath all of this. Hossain et al., in a companion piece to their cognitive-defense work, take on green and secure data centers directly, arguing that energy efficiency and advanced cybersecurity measures get treated as competing goals when they should really be designed together workload consolidation that cuts energy use can also shrink the attack surface by reducing the number of active, monitorable endpoints, while poorly designed energy-saving tricks like aggressive hardware power-cycling can make continuous monitoring and log continuity harder [8]. This angle is comparatively rare in broader literature, which tends to keep sustainability and security on separate tracks with separate budget owners.

The federated learning framework Tuhin et al. propose is relevant here too, since it is explicitly built for "global cloud-integrated MIS platforms" implying a distributed infrastructure where threat intelligence must move across geographically and organizationally distinct cloud deployments without one central data store [15]. Read next to Hossain et al.'s cognitive-defense framework, which is likewise built around big-data and cloud infrastructure [7], this sub-cluster points toward an infrastructure design philosophy where security, data locality, and resource efficiency get co-optimized rather than handled one after another, a philosophy this review adopts directly in Section VI by treating infrastructural efficiency as a cross-cutting constraint on the proposed framework rather than a box off to the side.

4.3 Management Information Systems, Governance and Decision Support

The single largest cluster by source count is the MIS and governance

layer sitting above and, the literature argues, inseparable from the technical security layer. Chakraborty et al., in a separate 2024 paper from their cybersecurity framework work, look at autonomous decision intelligence, tying explainable AI to scalable decision-support-system (DSS) architectures in modern MIS, and argue that as more decisions get handed to automated systems, explainability stops being a nice-to-have and becomes a governance requirement [2]. That argument matters a lot for the framework proposed later in this review, because it means Layer 4 (Decision and Orchestration) cannot just pass model output upward untouched; it must translate that output into something auditable and explainable to the humans and governance processes downstream.

Chy et al. study the relationship between data governance and business analytics success through a case study of global corporations, and their empirical grounding backs up a claim echoed all through the reviewed literature: analytics and security initiatives underperform when data governance is weak, no matter how sophisticated the underlying algorithms are [3]. Because this is one of the few reviewed sources with an actual case-study methodology rather than a purely conceptual framework, it carries extra weight for Section III-C's argument that governance maturity is a leading indicator of realized value from AI-driven security investment.

Das et al. connect MIS to project management practice, comparing organizational success factors when MIS supports agile IT project management, which matters for cybersecurity because security initiatives are usually delivered as IT projects bound by the same governance and delivery constraints [4]. Goffer et al. look at predictive analytics within MIS specifically for supply-chain resilience and mitigating economic disruption, showing how the same MIS/analytics infrastructure discussed for security purposes elsewhere in the literature is also expected to serve broader enterprise-resilience goals and hinting that investment in that shared infrastructure may pay off well beyond the security domain narrowly defined [5].

Hossin et al. take the MIS/analytics discussion toward sustainable economic growth, smart manufacturing, and Industry 4.0, positioning cybersecurity and data protection as one input among several into an organization's broader digital-transformation and sustainability strategy [9]. Mahmud et al. zero in on big data and cloud computing specifically within IT project management, proposing a framework for using those technologies to improve performance and decision-making in project delivery again reinforcing the theme that security and analytics capabilities are only as good as the project and governance structures used to deliver them [11].

Siddiqa et al., in a conference contribution, propose AI-driven project management systems that improve IT project efficiency through MIS integration, offering one of the more concrete architectural proposals in the governance cluster [14]. Across this cluster, a consistent picture emerges technical security capability (Section IV-A) and infrastructural design (Section IV-B) are necessary but not sufficient on their own; without governance structures, explainability, and project-delivery discipline behind them, they are unlikely to add up to organizational resilience.

This governance emphasis also connects naturally to the deployment-model vocabulary the NIST definition of cloud computing lays out [16]: choosing among public, private, community, and hybrid deployment models is itself a governance decision with direct security consequences, since it decides how much of the shared-responsibility boundary from Section III-A the organization actually controls. None of the seven governance-cluster sources reviewed here get into deployment-model choice

explicitly, which suggests connecting MIS governance frameworks to the specific deployment-model decisions organizations face in practice is an open opportunity, related to but distinct from the research directions discussed in Section VIII.

4.4 Human and Organizational Factors

One reviewed source takes on the human side directly. Shan-A-Alahi et al. look at how cybersecurity training shapes employee behavior in business settings, offering a counterpoint to the technology-heavy focus of most of the other reviewed sources [13]. This study stands out in the reviewed set precisely because it is the only one to put human behavior front and center as the main variable of interest rather than as a background assumption.

Given how often human error and social engineering get cited in the broader security literature as the proximate cause of breaches, how thin the human-factors research is in this set of sources is itself worth flagging. It also matters directly for the framework proposed in Section VI: Layer 5 (Response and Governance) explicitly builds in human-in-the-loop review, not because the technical literature in Section IV-A demands it, but because Shan-A-Alahi et al.'s finding implies that automated detection and orchestration layers cannot be assumed to work well on their own without accounting for how trained or untrained human operators actually interact with them [13]. This gap gets picked up again in Section VII.

It is a little striking, honestly, that so much of the reviewed literature builds elaborate technical and governance architectures around the assumption of a competent, attentive human somewhere at the end of the pipeline, while only one source studies whether that assumption holds. Shan-A-Alahi et al.'s finding that training shapes behavior is intuitive, but it also means the other fourteen sources in this corpus are, to varying degrees, resting on an unexamined premise which is not a criticism of any one of them individually so much as a pattern worth naming across the set as a whole.

4.5 Privacy-Preserving and Federated Approaches

Finally, Tuhin et al.'s proposal for federated machine learning in privacy-preserving cyber threat intelligence is a distinct methodological approach within the reviewed literature: instead of centralizing data to boost model performance, it deliberately trades away some of that performance in exchange for data locality and privacy guarantees, on the premise that organizations in regulated or competitive settings may be unable or unwilling to share raw security telemetry even when there is a collective defensive upside [15]. This source is the clearest bridge in the corpus between the technical detection cluster (Section IV-A) and the governance cluster (Section IV-C), since federated architectures are as much a governance and data-sharing-policy decision as a technical one, and its influence on Layer 3 of the proposed framework is spelled out in Section VI.

5. COMPARATIVE ANALYSIS

Rather than laying the fifteen reviewed sources out in a single comparison table, this section pulls together what the review found across them in narrative form, then quantifies it in Fig. 1. Read across Section IV, the fifteen sources split cleanly into two dominant themes and three smaller ones. AI/ML-driven detection and cyber threat intelligence, the work by Chakraborty et al. [1], Hasan et al. [6], Hossain et al. [7], Kaur et al. [10], Orthi et al. [12], and Tuhin et al. [15] and MIS governance and decision support, the second Chakraborty et al. study [2], Chy et al. [3], Das et al. [4], Goffer et al. [5], Hossin et al. [9], Mahmud et al. [11], and Siddiqa et al. [14], together account for the large majority of the corpus. Cloud infrastructure and sustainability shows up mainly through Hossain et al.'s green-data-center work [8] and, secondarily, Tuhin et al.'s federated architecture [15]; human and organizational factors rests almost entirely on Shan-A-Alahi et al. [13]; and privacy-preserving, federated

approaches is carried by Tuhin et al. alone [15], though its implications ripple into both the detection and governance clusters.

Fig. 1 turns that qualitative picture into numbers, and because several sources address more than one theme, the counts sum to more than fifteen. Two things stand out. First, MIS/governance and AI-driven detection are the two most heavily represented themes, which fits with these being the two poles, technical capability and organizational capability, the broader research area currently organizes itself around. Second, the comparatively thin representation of human factors and infrastructure sustainability suggests these are promising areas for future work rather than areas the existing literature has already worked through in depth; that observation directly motivates the framework proposed in Section VI and the research directions discussed in Section VIII.

The framework proposed in Section VI is built directly out of this clustering, and each of its five layers is traced back, in the discussion there, to the specific reviewed sources that motivate its inclusion and design Layer 1 to [1] and [8], Layer 2 to [3], [4], and [11], Layer 3 to [1], [6], [7], [10], and [15], Layer 4 to [2], [12], and [14], and Layer 5 to [13], [2], and [9]. The idea is to let a reader trace any piece of the proposed framework back to the literature it comes from, rather than treating the framework as a free-standing contribution disconnected from the review that precedes it. One more thing worth flagging from the corpus timeline: nine of the fifteen sources were published in 2024, against six in 2023, a split that points to fast-growing research attention rather than a settled, mature field. Combined with the venue spread noted in Section II-A, established information-systems journals alongside newer regional venues and one IEEE conference proceedings volume, this suggests a field still finding its shared terminology and methodology, which feeds into the cross-study comparability concern raised later in Section VII-E and is part of why this review attempts, in the framework proposed next, to offer one possible shared point of reference.

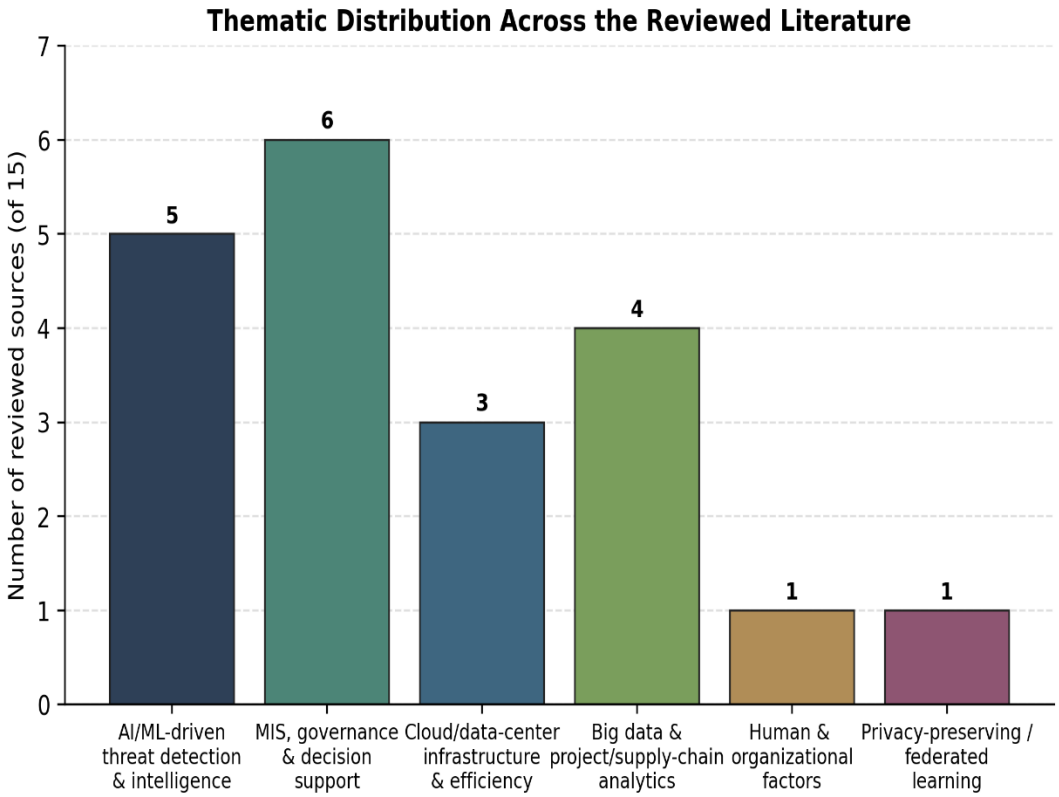

Figure. 1. Thematic distribution across the reviewed literature (counts exceed 15 because several sources span more than one theme).

6. PROPOSED CONCEPTUAL FRAMEWORK

Building on the five thematic clusters from Section IV and the process described in Section II-D, this section puts together a five-layer conceptual framework, shown in Fig. 2. It is offered as a conceptual synthesis of recurring architectural ideas across the reviewed literature, not as a system anyone has tested, no performance claims are attached to it. The point is to give researchers and practitioners a shared vocabulary for locating their own work within the broader problem space, and to make explicit the dependencies between layers that individual studies in the reviewed literature tend to treat as separate.

Layer 1 (Data and Resource Layer) is the raw substrate discussed in Section III-A and Section IV-B: virtual machines, containers, storage services, network traffic, identity and access logs, API call records. This is where baseline controls IAM, encryption, segmentation get enforced, and it is the layer most directly touched by the energy-efficiency concerns Hossain et al. raise, since the number and configuration of active resources here sets both the attack surface and the energy footprint of the environment [8].

Layer 2 (Collection and Preprocessing) reflects the governance concerns running through Section IV-C log aggregation, normalization, and feature extraction are only as reliable as the data-governance practices Chy et al. flag as a precondition for downstream analytics success [3]. A system that collects thorough telemetry but normalizes or labels it inconsistently across business units ends up with Layer 3 models that are only as good as the noisiest source feeding them, which is why this review treats Layer 2 as architecturally distinct from, and a prerequisite to, Layer 3, rather than folding preprocessing into the analytics layer the way some of the reviewed sources implicitly do.

Layer 3 (Intelligent Analytics) maps to the technical detection cluster in Section IV-A, covering supervised and unsupervised ML classifiers, anomaly detection, and following Tuhin et al. federated learning architectures for cross-organizational model training without centralized raw data [15]. This layer is deliberately kept narrow, scoped to model training and inference; interpreting and prioritizing model output is left to Layer 4, mirroring the distinction Hossain et al. draw between raw detection and a "cognitive" contextualization layer [7].

Layer 4 (Decision and Orchestration) reflects the explainable-AI and DSS integration Chakraborty et al. argue for, plus the AI-driven project and threat-intelligence management systems Siddiqa et al. and Orthi et al. propose [2], [12], [14]. This is where raw model output, an anomaly score, say gets translated into something organizations can actually act on, like a prioritized, explained alert with a recommended action. It is also the layer most directly caught up in the explainability gap discussed in Section VII-A, since it is where a governance requirement (explainability) and a technical capability (model output) must be reconciled.

Layer 5 (Response and Governance) closes the loop with automated mitigation, human-in-the-loop review motivated directly by the human-factors finding in Section IV-D [13] and compliance and audit reporting, feeding back into Layer 1 through the continuous monitoring loop on the left side of Fig. 2. Hossain et al.'s infrastructural-efficiency concerns [8] are treated as a cross-cutting constraint on all five layers rather than a discrete layer of their own, since energy and cost considerations shape design decisions at every level from data collection through response automation, in line with the joint-optimization argument made in Section IV-B.

6.1 Illustrative Walkthrough

To make the five layers concrete, it helps to trace a single hypothetical event through the framework, illustratively, not as a validated case study. Say an attacker gets hold of a valid but low-privilege cloud API credential and starts issuing an unusual sequence of storage-read calls outside the account's normal usage pattern. At Layer 1, that activity just generates raw API and access logs like any other call; nothing marks it

as suspicious yet. At Layer 2, those logs get aggregated and normalized alongside logs from other accounts and services, and basic features, call frequency, time of day, resource scope, get extracted in a governed, consistently labeled form, in line with the governance precondition Chy et al. emphasize [3].

At Layer 3, an anomaly-detection model, possibly one trained partly through a federated arrangement across business units, following Tuhin et al. [15], flags the sequence as statistically unusual against the historical baseline and produces a numerical anomaly score. At Layer 4, that raw score does not get dumped on an analyst by itself; in line with Chakraborty et al.'s argument for explainable AI and DSS integration [2], it gets translated into a prioritized, human-readable alert ("unusual storage-read volume from credential X, inconsistent with 90-day baseline, recommended action: verify credential owner") and ranked against whatever else is happening concurrently.

At Layer 5, a human analyst, whose effectiveness here depends on the training and workload factors Shan-A-Alahi et al. highlights [13] reviews the prioritized alert, confirms or dismisses it, and, if confirmed, triggers a mitigation step such as credential revocation. Either way, the outcome gets logged for audit and compliance purposes and fed back into Layer 1 as an extra labeled data point for future model training, closing the loop shown in Fig. 2. This walkthrough exists purely to illustrate how the layers are meant to interact; it is not a report of an actual observed incident or of measured system performance.

6.2 Design Trade-Offs: Centralized Versus Federated Instantiation of Layer 3

Section III-D and Section IV-E both note that Layer 3 (Intelligent Analytics) can be built either as a centralized analytics pipeline or as a federated one, and that the choice carries governance weight as well as technical weight. Five dimensions come up repeatedly across the reviewed literature and supplementary foundational references when weighing that choice. On model performance, the centralized option is generally stronger, since it trains on the full pooled dataset [1], [6], [7], [10], while the federated option is weaker or more variable, depending on how heterogeneous the data is across participants [15], [19]. On data privacy and regulatory fit, that ranking flips: centralizing raw data brings higher exposure risk, while federated training keeps raw data local to each participant [15]. On implementation complexity, centralized pipelines can lean on conventional ML/DL tooling, while federated setups need coordination and secure-aggregation infrastructure on top of that [19]. On vulnerability to poisoning, a single trusted centralized pipeline is harder to corrupt than a federated one, where malicious participants can submit manipulated updates [19]. And on governance ownership, a centralized pipeline is typically a single-organization decision, while a federated one is inherently a cross-organizational data-sharing policy decision [15]. None of this is a scored or weighted decision matrix; organizations weighing the choice should read these five dimensions against their own regulatory and risk context rather than treating either instantiation as universally better.

The centralized instantiation, which fits most of the AI/ML-driven detection cluster reviewed in Section IV-A, generally offers stronger raw model performance because it trains on the full pooled dataset, at the cost of the privacy and regulatory exposure discussed in Section III-D. The federated instantiation, following Tuhin et al. [15] and the foundational federated-learning framework of Yang et al. [19], flips that trade-off, and adds the poisoning and communication-overhead concerns discussed in Section VII-D. Neither instantiation wins on every dimension, which is exactly why this review treats the

choice between them as a governance decision requiring the kind of cross-functional judgment discussed in Section III-D2 and Section IX-C, rather than as a purely technical default.

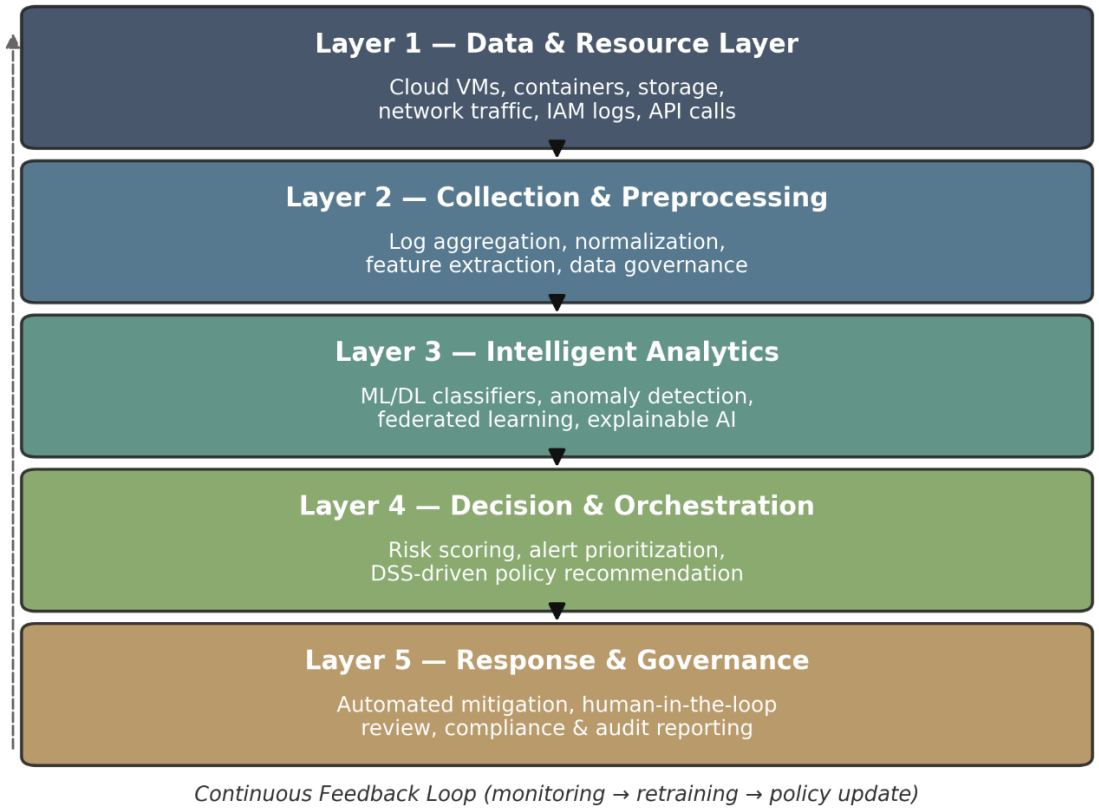

Figure 2. Proposed five-layer conceptual framework synthesized from the reviewed literature.

7. OPEN CHALLENGES AND RESEARCH GAPS

A handful of gaps come out of this review that no single reviewed source fully closes. This section works through five of the more significant ones.

7.1 The Explainability Implementation Gap

Chakraborty et al. treat explainability as a governance requirement [2], but few of the detection-focused sources in Section IV-A say much about how that requirement gets met at the model level, through inherently interpretable model classes, post-hoc explanation techniques, or structured natural-language justification of individual alerts. That leaves a gap between what the governance literature asks for and what the technical literature actually delivers: it is not clear from the reviewed corpus how an organization following Chakraborty et al.'s DSS-integration guidance [2] would even go about checking whether a specific detection system proposed elsewhere in the corpus [1], [6], [7] among them actually meets that bar.

7.2 The Human-Factors Gap

The human-factors literature in this review comes down to a single source [13], even though the technical and governance literature keeps assuming alerts, dashboards, and DSS outputs will get acted on appropriately by human operators. How training, workload, alert fatigue, and behavior play out in AI-augmented environments specifically as opposed to general security-awareness settings is not well covered by the reviewed sources. That is a meaningful gap given that both Layer 4 and Layer 5 of the proposed framework depend, implicitly, on competent and attentive human review somewhere in the pipeline.

7.3 The Sustainability/Security Integration Gap

The sustainability/security trade-off Hossain et al. raise [8] gets treated as a standalone concern in the reviewed literature rather than being folded into the AI-driven detection frameworks discussed in Section IV-A; none of the detection-focused sources reviewed here report the computational or energy cost of the AI/ML techniques they propose, which is a notable gap given that inference-time cost matters directly to whether a proposed technique can run continuously across large-scale cloud telemetry without itself becoming a meaningful contributor to an organization's energy footprint.

7.4 The Federated-Learning Limitations Gap

Federated learning is offered as a fix for the centralization/privacy tension [15], but the reviewed literature does not get very far into the known limitations of federated approaches vulnerability to poisoning by malicious or compromised participants, weaker model performance relative to centralized training on non-independent, non-identically distributed data, and the communication overhead of coordinating updates across distributed participants. That leaves open how organizations should weigh federated learning's privacy benefits against those limitations in a security-specific context, where a successful poisoning attack arguably costs more than it would in most other federated-learning applications.

7.5 The Cross-Study Comparability Gap

A fifth, more methodological gap cuts across the other four: the reviewed sources use noticeably different language for closely related ideas, "cognitive cyber defense" [7], "AI-driven cyber threat intelligence as MIS" [12], and "intelligent cybersecurity framework" [1] all describe architectures that, per the layer-by-layer mapping in Section VI, sit in overlapping positions across Layers 3 and 4 and none of the reviewed sources report results against a shared benchmark dataset, a shared metric set, or a shared terminology standard. That makes it genuinely hard, from the reviewed literature alone, to tell whether two differently named frameworks reflect real architectural differences or largely the same idea described independently by different research groups, a problem this review has tried to partly address through the layer-mapping exercise in Section VI, but one the underlying literature does not resolve on its own.

8. FUTURE RESEARCH DIRECTIONS

Building on the gaps identified in Section VII, six directions for future work are followed directly from this review.

8.1 Joint Evaluation of Detection Performance and Explainability

Empirical studies that measure detection performance and explainability together comparing analyst response time and accuracy when using explainable versus black-box AI outputs on matched alert sets, for instance would help close the gap between the governance and technical literatures identified in Section VII-A. That kind of study would give organizations following guidance like Chakraborty et al.'s [2] something concrete to base a choice on among the detection approaches proposed elsewhere in the corpus [1], [6], [7].

8.2 Human-Factors Research in AI-Augmented Operations Centers

Human-factors research aimed specifically at AI-augmented cybersecurity operations centers, rather than general security-awareness training, would extend Shan-A-Alahi et al.'s contribution [13] into the AI-specific context this review has focused on. Useful questions here include how alert-prioritization interfaces affect analyst trust and override behavior, and how training should differ for analysts working mainly from AI-generated recommendations versus analysts working from raw telemetry.

8.3 Energy-Aware Detection Benchmarking

Future technical proposals in the AI-driven detection space should report energy and computational cost alongside detection accuracy, so the kind of joint optimization Hossain et al. argue for [8] can be evaluated quantitatively rather than just conceptually. Standardized benchmarking that reports detection accuracy per unit of computational cost would make this trade-off visible across studies in a way the current literature does not support.

8.4 Adversarial Evaluation of Federated Threat Intelligence

Further empirical testing of federated threat-intelligence architectures under adversarial conditions participant poisoning, non-independent and non-identically distributed data across organizations would strengthen the practical case for the approach Tuhin et al. propose [15], and would help organizations weigh the privacy benefits identified in Section III-D against the limitations identified in Section VII-D.

8.5 Layered Framework Validation

More broadly, this review suggests that future studies would benefit from explicitly locating their contribution within a layered framework like the one proposed in Section VI, making the dependencies between technical, governance, human, and infrastructural layers an explicit part of the study design rather than an unspoken assumption. Longitudinal or comparative case studies that instrument all five layers of a single deployed environment would provide the empirical validation this conceptual framework does not itself supply.

8.6 Standardized Terminology and Reporting

Directly addressing the comparability gap in Section VII-E, the research community here would benefit from a shared minimum reporting standard for AI-driven cybersecurity frameworks something like a requirement that any new framework proposal state which of the five layers in Fig. 2 (or an equivalent breakdown) it addresses, which it assumes are handled elsewhere, and what terminology it uses for each. This would not require every research group to adopt this review's particular five-layer framework, only to report their contribution in a way that makes cross-study comparison tractable, the way structured abstracts and reporting checklists have improved comparability in other empirical fields.

9. PRACTICAL IMPLICATIONS

Beyond the research-oriented gaps discussed in Sections VII and VIII, the synthesis in this review carries some practical implications for three groups of stakeholders: security teams, MIS and IT governance leaders, and policy/compliance functions. These are offered as reasonable inferences from the reviewed literature, not as prescriptions this review has itself tested.

9.1 Implications for Security Teams

For security operations teams, the clustering in Section IV suggests AI/ML detection capability (Section IV-A) should not be judged in isolation from the decision-and-orchestration layer that consumes its output. A detection system with highly accurate but unexplained alerts may, following the argument in Section IV-C, end up underperforming a slightly less accurate system whose output is easier to fold into existing DSS and audit workflows [2]. Teams evaluating vendor or in-house detection tooling might reasonably score candidate systems on explainability and integration effort alongside raw detection accuracy, rather than accuracy alone.

9.2 Implications for MIS and IT Governance Leaders

For MIS and IT governance leaders, this review's central argument that governance maturity is a leading indicator of realized security value implies data-governance investment (Section III-C) may be a prerequisite for AI-driven security investment rather than something to pursue in parallel with it [3], [4]. Organizations considering an AI-driven detection or threat-intelligence

initiative could reasonably use their existing data-governance maturity, characterized in sources like Chy et al. [3], to decide how to sequence that initiative relative to other governance investments, instead of chasing at once without regard for the dependency between them.

9.3 Implications for Policy and Compliance Functions

For policy and compliance functions, the human-factors gap in Section VII-B and the federated-learning limitations in Section VII-D both point to specific due-diligence questions worth adding to internal AI-governance checklists: has a proposed AI-driven security system been evaluated for its effect on analyst behavior and workload, not just its standalone detection accuracy, and has a proposed federated or collaborative threat-intelligence arrangement been evaluated for adversarial robustness before sensitive telemetry-sharing agreements get signed. Neither question is fully answered by the reviewed literature, which is exactly why they are framed here as due-diligence questions rather than settled best practice.

9.4 Implications for Academic Researchers

For researchers, the comparability gap in Section VII-E points to a low-cost way to raise the field's cumulative value: adopt, or at least report against, a shared breakdown like the one proposed in Section VI when describing a new framework or empirical result. Doing so would let future reviews of this kind lean less on the kind of title-and-abstract-level reconstruction this review has had to do (see Section II-B) and more on a direct, structurally comparable reading of each contribution's scope and claims.

10. LIMITATIONS OF THIS REVIEW

This review has a few limitations worth being upfront about. First, the corpus of fifteen sources was handed over as a fixed reference list rather than built through an original, open-ended database search across platforms like IEEE Xplore, ACM Digital Library, Scopus, or Web of Science; as Section II-A notes, this review does not claim database-level exhaustiveness, and the thematic distribution in Fig. 1 should be read as a description of this specific corpus rather than a claim about how the broader published literature on AI-driven cybersecurity actually breaks down.

Second, some of the reviewed sources appear in venues regional or newer journals, or non-security-specialist venues publishing security-adjacent work who's editorial and peer-review rigor this review has not independently audited beyond the bibliographic record it was given. Readers building specific technical claims from any single reviewed source are encouraged to go check that source directly rather than leaning solely on this review's characterization of it.

Third, because this review works from each source's stated title, abstract-level focus, and publication metadata rather than a full independent re-analysis of its underlying data or code where applicable, the categorizations discussed in Section IV and the thematic groupings behind them reflect this review's interpretation of each source's stated contribution rather than an independently verified read on methodological quality. And the conceptual framework proposed in Section VI is, by design, untested; it should be read as a hypothesis-generating synthesis rather than a validated architecture, pending the empirical work outlined in Section VIII.

One more limitation concern generalizability across organizational contexts. Several of the reviewed sources draw their motivating examples from large enterprises or global corporations [3], [5], [9]; how the five-layer framework proposed in Section VI would need to be adapted, or simplified, for smaller organizations with thinner security and data-governance staffing is not something the reviewed corpus addresses, and so it is not something this review's synthesis of it addresses either. Readers in resource-constrained organizational settings should probably treat

the framework's five layers as a conceptual checklist for prioritization rather than a specification demanding dedicated tooling or staffing at every layer.

11. CONCLUSION

This paper has worked through fifteen recent studies on AI-driven cybersecurity, cloud infrastructure, and management information systems, organizing them into five thematic clusters: AI/ML-driven threat detection and intelligence, cloud infrastructure and sustainability, MIS governance and decision support, human and organizational factors, and privacy-preserving federated approaches. A thematic distribution chart, together with the layer-by-layer mapping back to the literature woven through Section VI, quantified how these themes are actually represented across the reviewed literature, showing that technical detection capability and MIS governance are comparatively well covered, while human factors and the security/sustainability trade-off are noticeably thinner.

Building on that synthesis, the paper proposed a five-layer conceptual framework meant to spell out the dependencies between these themes for future researchers and practitioners, tied back explicitly to the literature that motivates each layer. The framework is offered as a conceptual synthesis rather than a validated system, and this paper's discussion of its own limitations, open challenges, and future research directions is meant to point toward the empirical work still needed to validate, extend, or revise it.

REFERENCES

[1] P. Chakraborty, H. U. Himel, M. Bashir, M. S. Sikder, H. Kaur, M. T. Bin Ansar, J. Kaur, and M. K. Tuhin, "An intelligent cybersecurity framework for data protection in cloud computing environments," Spectrum of Engineering Sciences, vol. 2, no. 3, pp. 657–677, 2024. DOI: https://doi.org/10.5281/zenodo.21538159. URL: https://thesesjournal.com/index.php/1/article/view/3573.

[2] P. Chakraborty, R. A. M. Rashed, M. Bashir, H. Imam, M. A. Siam, M. A. Miah, K. B. Siddiqa, and A. Islam, "Toward autonomous decision intelligence: Integrating explainable AI and scalable DSS architectures in modern management information systems," Journal of Information Systems Engineering and Management, vol. 9, no. 4s, 2024, DOI: https://doi.org/10.52783/jisem.v9i4s.14591. URL: https://jisem-journal.com/index.php/journal/article/view/14591.

[3] M. A. R. Chy, E. Rozario, M. H. Rijvi, S. M. M. Uddin, M. M. Bakhsh, E. Hossain, M. J. Hossain, U. S. Saha, and M. I. Faruk, "Understanding the relationship between data governance and business analytics success: A case study of global corporations," Journal of Information Systems Engineering and Management, vol. 9, no. 4s, 2024, DOI: https://doi.org/10.52783/jisem.v9i4s.14807. URL: https://www.jisem-journal.com/index.php/journal/article/view/14807.

[4] N. Das, J. Hassan, H. Rahman, K. B. Siddiqa, S. M. Orthi, C. R. Barikdar, and M. A. Miah, "Leveraging management information systems for agile project management in information technology: A comparative analysis of organizational success factors," Journal of Business and Management Studies, vol. 5, no. 3, pp. 161–168, 2023, DOI: https://doi.org/10.32996/jbms.2023.5.3.17. URL: https://al-kindipublisher.com/index.php/jbms/article/view/8230.

[5] M. A. Goffer, P. Chakraborty, H. Rahman, C. R. Barikdar, N. Das, S. Hossain, and M. E. Hossin, "Leveraging predictive analytics in management information systems to enhance supply chain resilience and mitigate economic disruptions," Educational Administration: Theory and Practice, vol. 30, no. 4, pp. 11134–11144, 2024. DOI: https://doi.org/10.53555/kuey.v30i4.9641. URL: https://kuey.net/index.php/kuey/article/view/9641.

[6] S. N. Hasan, J. Hassan, C. R. Barikdar, P. Chakraborty, U. Haldar, M. A. R. Chy, E. Rozario, N. Das, and J. Kaur, "Enhancing cybersecurity threat detection and response through big data analytics in management information systems," Fuel Cells Bulletin, vol. 2023, no. 12, 2023, DOI: https://doi.org/10.52710/fcb.137. URL: https://fuelcellsbulletin.org/index.php/journal/article/view/137.

[7] M. D. Hossain, M. S. Sikder, M. S. Uddin, R. M. Ahsan, B. Uddin, and T. Hossen, "Cognitive cyber defense: AI–MIS integration through big data and cloud frameworks for next-generation digital resilience," The Eastasouth Journal of Information System and Computer Science, vol. 1, no. 2, pp. 140–152, 2023, DOI: https://doi.org/10.58812/esiscs.v1i02.764. URL: https://esj.eastasouth-institute.com/index.php/esiscs/article/view/764.

[8] M. D. Hossain, M. S. Uddin, M. S. Sikder, T. Hossen, B. Uddin, and R. M. Ahsan, "Green and secure data centers: Balancing energy efficiency with advanced cybersecurity measures," Journal of Computer Science and Technology Studies, vol. 6, no. 5, pp. 300–315, 2024, DOI: https://doi.org/10.32996/jcsts.2024.6.5.24. URL: https://al-kindipublisher.com/index.php/jcsts/article/view/11146.

[9] M. E. Hossin, J. Hassan, M. A. R. Chy, S. Hossain, E. Rozario, F. B. Khair, and M. A. Goffer, "Harnessing business analytics in management information systems to foster sustainable economic growth through smart manufacturing

and Industry 4.0," Educational Administration: Theory and Practice, vol. 30, no. 10, pp. 730–739, 2024, DOI: https://doi.org/10.53555/kuey.v30i10.9643.URL: https://kuey.net/index.php/kuey/article/view/9643.

[10] J. Kaur, S. N. Hasan, S. M. Orthi, M. A. Miah, M. A. Goffer, C. R. Barikdar, and J. Hassan, "Advanced cyber threats and cybersecurity innovation: Strategic approaches and emerging solutions," Journal of Computer Science and Technology Studies, vol. 5, no. 3, pp. 112–121, 2023, DOI: https://doi.org/10.32996/jcsts.2023.5.3.9. URL: https://al-kindipublisher.com/index.php/jcsts/article/view/8238

[11] F. Mahmud, S. M. Orthi, A. S. M. Saimon, M. Moniruzzaman, M. A. Miah, M. K. Ahmed, F. B. Khair, M. S. Islam, and M. M. T. G. Manik, "Big data and cloud computing in IT project management: A framework for enhancing performance and decision-making," Fuel Cells Bulletin, vol. 2023, no. 9, 2023, DOI: https://doi.org/10.52710/fcb.166. URL: https://fuelcellsbulletin.org/index.php/journal/article/view/166

[12] S. M. Orthi, P. Chakraborty, M. A. Siam, A. Shan-A-Alahi, A. Al Zaiem, S. N. Hasan, J. Kaur, F. Mahmud, and M. A. Goffer, "AI-driven cyber threat intelligence as a management information system: Integrating cybersecurity governance and IT project management for organizational resilience," The Eastasouth Journal of Information System and Computer Science, vol. 1, no. 2, pp. 194–214, 2023, DOI: https://doi.org/10.58812/esiscs.v1i02.873. URL: https://esj.eastasouth-institute.com/index.php/esiscs/article/view/873.

[13] A. Shan-A-Alahi, M. Mustafizur, K. M. R. Hossan, A. Al Zaiem, and M. M. Rahman, "Cybersecurity training and its influence on employee behavior in business environments," Computer Fraud & Security, vol. 2024, no. 12, 2024, DOI: https://doi.org/10.52710/cfs.689. URL: https://computerfraudsecurity.com/index.php/journal/article/view/689.

[14] K. B. Siddiqa, H. Rahman, C. R. Barikdar, S. M. Orthi, M. A. Miah, and R. Rahman, "AI-driven project management systems: Enhancing IT project efficiency through MIS integration," in Proc. 2024 Int. Conf. on Progressive Innovations in Intelligent Systems and Data Science (ICPIDS), 2024, pp. 114–119, DOI: https://doi.org/10.1109/ICPIDS65698.2024.00027. URL: https://ieeexplore.ieee.org/document/10974128

[15] M. K. Tuhin, M. T. Bin Ansar, M. S. Sikder, H. U. Himel, H. Kaur, H. Imam, and A. Shan-A-Alahi, "Federated machine learning for privacy-preserving cyber threat intelligence in global cloud-integrated MIS platforms," The Eastasouth Journal of Information System and Computer Science, vol. 2, no. 2, pp. 284–293, 2024, DOI: https://doi.org/10.58812/esiscs.v2i02.1123. URL: https://esj.eastasouth-institute.com/index.php/esiscs/article/view/1123.

[16] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology, Special Publication 800-145, Gaithersburg, MD, USA, 2011, DOI: https://doi.org/10.6028/NIST.SP.800-145.

[17] S. Subashini and V. Kavitha, "A survey on security issues in service delivery models of cloud computing," Journal of Network and Computer Applications, vol. 34, no. 1, pp. 1–11, 2011, DOI: https://doi.org/10.1016/j.jnca.2010.07.006.

[18] D. Zissis and D. Lekkas, "Addressing cloud computing security issues," Future Generation Computer Systems, vol. 28, no. 3, pp. 583–592, 2012, DOI: https://doi.org/10.1016/j.future.2010.12.006.

[19] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 12:1–12:19, 2019, DOI: https://doi.org/10.1145/3298981.

[20] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," Communications of the ACM, vol. 53, no. 4, pp. 50–58, 2010, DOI: https://doi.org/10.1145/1721654.1721672.

[21] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, pp. 599–616, 2009, DOI: https://doi.org/10.1016/j.future.2008.12.001.

[22] K. Ren, C. Wang, and Q. Wang, "Security challenges for the public cloud," IEEE Internet Computing, vol. 16, no. 1, pp. 69–73, 2012, DOI: https://doi.org/10.1109/MIC.2012.14.