
A Conceptual Framework for ISO/IEC 27001 Audit Augmentation Using Multi-Modal Machine Learning: Integrating Document Review, Field Observation, and Conversational AI Interview

[Nungky Awang Chandra](#) *

Posted Date: 7 May 2026

doi: 10.20944/preprints202605.0329.v1

Keywords: ISO/IEC 27001; information security audit; machine learning; natural language processing; computer vision; large language models; audit automation; conceptual framework; information security management system; compliance



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

A Conceptual Framework for ISO/IEC 27001 Audit Augmentation Using Multi-Modal Machine Learning: Integrating Document Review, Field Observation, and Conversational AI Interview

Nungky Awang Chandra

Department of Informatics Engineering, Faculty of Computer Sciences, Universitas Mercu Buana, Jl. Meruya Selatan No. 1, Kembangan, West Jakarta 11650, Indonesia; nungky_awang@mercubuana.ac.id

Abstract

The audit of Information Security Management Systems (ISMS) under ISO/IEC 27001:2022 has traditionally relied on human auditors whose competence, experience, and judgment shape audit outcomes. While effective, this human-centric approach suffers from inter-auditor variability, high cost, scheduling constraints, and limited scalability — challenges magnified by the post-pandemic shift toward remote audits and the growing volume of organisations seeking certification. Recent advances in Natural Language Processing (NLP), Computer Vision (CV), and Large Language Models (LLMs) suggest that significant portions of the audit workflow could be augmented by machine learning. However, prior research has examined these technologies in isolation; no integrated conceptual framework yet exists that unifies document review, field observation, and interviewing under a single multi-modal pipeline tailored to ISO/IEC 27001 audits and explicitly grounded in the audit methodology of ISO 19011:2018. This paper proposes such a framework — the Multi-Modal ML-Augmented ISO 27001 Audit Framework (M³A-Framework). We synthesise insights from ISO 19011:2018 audit guidelines, recent advances in AI-driven assurance, and the design science research paradigm to develop a five-stage conceptual model that augments the seven-step evidence-collection process specified in ISO 19011 Clause 6.4.7 and that extends the audit-methods matrix of ISO 19011 Annex A (Table A.1). The framework comprises: (1) audit planning and scoping; (2) multi-modal evidence collection through NLP for document analysis, CV for physical control verification (supported by inspection robots and drones), and LLM-based conversational AI for interview; (3) ML-based evidence processing and triangulation; (4) confidence-weighted finding classification using Explainable AI; and (5) human-in-the-loop validation. The framework explicitly maps each module to the 93 controls of Annex A of ISO/IEC 27001:2022 and to the audit phases mandated by ISO 19011. We further propose a set of testable propositions, evaluation metrics, and ethical considerations that ground the framework in both academic rigour and practical deployability.

Keywords: ISO/IEC 27001; information security audit; machine learning; natural language processing; computer vision; large language models; audit automation; conceptual framework; information security management system; compliance

1. Introduction

Information security has become a strategic imperative for organisations across all sectors. With the proliferation of cyber threats, regulatory requirements, and digital transformation initiatives, organisations are increasingly seeking certification against ISO/IEC 27001 [1,2], the leading international standard for Information Security Management Systems (ISMS). The 2022 revision of ISO/IEC 27001 introduced a streamlined Annex A with 93 controls organised into four themes —

organisational, people, physical, and technological [3,4]. Certification against this standard is verified through formal audits conducted in accordance with ISO 19011:2018 [5], which establishes the principles, processes, and competence requirements for management system audits.

Traditional ISO 27001 audits are labour-intensive. They require certified lead auditors to plan the audit, review extensive documentation, conduct on-site observations, interview personnel across organisational levels, and synthesise findings into a comprehensive audit report [6]. Although effective in principle, the human-centric model exhibits well-documented limitations: inter-auditor variability in interpretation [7,8], high financial and time costs [9,10], and limited scalability for large or geographically distributed organisations [11,12]. The COVID-19 pandemic exposed an additional fragility – the difficulty of conducting on-site audits during travel restrictions – and accelerated the demand for remote and hybrid audit modes.

Concurrent with these challenges, machine learning (ML) has reached a level of maturity at which substantial portions of the audit workflow could plausibly be augmented. Natural Language Processing (NLP) models can extract requirements from policy documents and verify compliance against regulatory text [13–15]; Computer Vision (CV) systems can detect physical security controls from imagery [16–18]; speech recognition and Large Language Model (LLM)-based conversational agents can conduct structured interviews [19]; and retrieval-augmented generation (RAG) systems can synthesise findings from heterogeneous evidence [20]. Explainable AI (XAI) further enables auditors to interrogate model predictions and preserve human accountability [21,22]. Each of these capabilities has been studied separately in the audit literature, but no integrated framework has yet brought them together under a unifying architecture aligned with ISO 27001 controls and ISO 19011 audit phases.

Despite this convergence of need and capability, the research literature on AI-augmented ISO 27001 audits is fragmented. Compliance-checking studies have focused on textual analysis [23,24] without addressing physical or interview evidence. Computer vision research on surveillance [25,26] has not been linked to ISO control objectives. LLM applications in audit [27,28] have largely been experimental rather than framework-based. Reviews of AI in auditing [29,30] discuss general principles but do not propose deployable architectures. There is, in short, a recognised gap that calls for a comprehensive conceptual framework – one that is grounded in established audit theory, leverages contemporary ML capabilities, and is deployable in real-world certification contexts.

This paper addresses the gap by proposing a comprehensive conceptual framework for ML-augmented ISO/IEC 27001 audits – the Multi-Modal ML-Augmented ISO 27001 Audit Framework (M³A-Framework). Drawing on the design science research paradigm [31,32], we synthesise insights from the ISO 19011:2018 audit standard, the ISO/IEC 27001:2022 control set, and recent advances in NLP, CV, LLMs, and XAI, to develop a five-stage conceptual model that integrates document review, field observation, and interviewing in a unified pipeline.

The contributions of this paper are sixfold:

1. We synthesise the audit theory of ISO 19011 with contemporary ML capabilities to define a five-stage conceptual pipeline for augmented ISO 27001 audits.
2. We articulate three multi-modal evidence-collection modules – NLP-based document review, CV-based field observation (extended via robotic inspection and drone imagery), and LLM-based conversational interview – and explain how their outputs are triangulated.
3. We provide a fine-grained mapping of ML augmentation to the seven-step evidence-collection process specified in ISO 19011:2018, Clause 6.4.7 (Figure 2 of that standard), and demonstrate how each step preserves the seven principles of auditing in Clause 4.
4. We extend the ISO 19011 audit-methods matrix (Annex A, Table A.1) by mapping each of its four quadrants – interactive/non-interactive × on-site/remote – to dedicated ML module classes, defining an incremental adoption pathway for organisations.
5. We map the framework to ISO/IEC 27001:2022 Annex A controls, demonstrating coverage and identifying control families particularly amenable to automation.

6. We propose a set of testable research propositions, evaluation metrics, and ethical considerations to guide downstream empirical validation.

The remainder of the paper is organised as follows. Section 2 reviews the ISO/IEC 27001 audit landscape, prior work on ML in auditing, and the design science approach. Section 3 describes our research methodology. Section 4 presents the M³A-Framework in detail. Section 5 discusses the mapping to ISO/IEC 27001:2022 controls and implementation considerations. Section 6 develops testable propositions and evaluation metrics. Section 7 discusses ethical considerations, limitations, and future research directions. Section 8 concludes.

2. Background and Related Work

2.1. ISO/IEC 27001 and the Audit Process

ISO/IEC 27001:2022 sets out requirements for establishing, implementing, maintaining, and continually improving an Information Security Management System [3]. The standard is structured around the high-level structure (HLS) shared with other ISO management system standards, comprising clauses 4-10 covering context, leadership, planning, support, operation, performance evaluation, and improvement. Annex A contains 93 controls organised into four themes: organisational (37 controls), people (8), physical (14), and technological (34) [4]. The 2022 revision consolidated 114 controls from the 2013 version into the current 93, introducing five new controls focused on threat intelligence, cloud services, ICT readiness for business continuity, physical security monitoring, and information deletion.

Audits are conducted in line with ISO 19011 [5], which defines the audit programme management process and the conduct of individual audits. The standard establishes seven principles of auditing – integrity, fair presentation, due professional care, confidentiality, independence, evidence-based approach, and risk-based approach (the last added in the 2018 third edition) – and prescribes a sequence of audit activities: initiating the audit, preparing audit activities, conducting audit activities, preparing and distributing the audit report, completing the audit, and conducting audit follow-up.

Empirical research has documented systematic limitations of the human-centric audit model. Castka et al. [11] showed that audit duration and cost grow super-linearly with auditee complexity. Beckmerhagen et al. [9] and Talapatra et al. [33] demonstrated significant inter-auditor variability – auditors examining identical evidence can reach different conclusions due to differences in experience, prior assumptions, and cognitive biases. Mirtsch et al. [10] used web mining to identify factors driving ISO/IEC 27001 adoption, finding that the high cost of certification disproportionately excludes small and medium-sized enterprises (SMEs) from the standard's benefits.

2.2. ISO/IEC 27001 Audit Practice and Quality

The literature on ISO 27001 audit practice highlights several persistent challenges. Sartor et al. [33] and Disterer [36] surveyed adoption barriers, citing complexity, expense, and the difficulty of demonstrating return on investment. Topa and Karyda [37] developed practical guidelines for enhancing information security management, emphasising organisational readiness and leadership commitment. Boz et al. [38] proposed a framework for measuring auditor competence in continuous auditing, recognising that competence requirements evolve as audit modes change. Culot et al. [39] conducted a comprehensive literature review of ISO/IEC 27001 research, identifying methodological compliance, governance, and technological alignment as recurring themes.

2.3. Machine Learning in Auditing

ML applications in auditing fall into three categories that map naturally to evidence collection. First, NLP for documentary evidence: Hasan [40] surveyed AI in accounting and auditing, while Munoko et al. [41] examined ethical implications. Hilal et al. [42] reviewed anomaly detection for

fraud, and Apruzzese et al. [43] explored ML for cybersecurity threat detection. Domain-specific transformer models such as LEGAL-BERT have demonstrated strong performance in legal text analysis [47], and recent work on automated compliance checking against GDPR [24] and contractual requirements [44] indicates feasibility of similar approaches for ISO 27001 evidence.

Second, CV for physical evidence: smart surveillance systems [25] use object detection (YOLO [16], Vision Transformers [17,18]) to monitor access controls, identify intruders, and verify physical security measures. Recent surveys on video anomaly detection [26,48] show that deep learning methods achieve high accuracy on standard datasets. Robotic inspection [49,63,67] has been applied to building maintenance and infrastructure assessment. Together with drone-based aerial surveillance [63], these technologies can extend the auditor's reach to perimeter inspection, data centre monitoring, and remote facility verification.

Third, LLMs for interviewing and synthesis: chatbots and conversational agents [27,28] have evolved from rule-based systems to LLM-powered agents capable of context-aware dialogue. Modern LLMs such as GPT-4 [14] and Claude [15] can conduct structured interviews, summarise responses, and identify gaps in answer coverage. Speech recognition models like Whisper [19] enable voice-based interaction, while retrieval-augmented generation [20] grounds LLM responses in authoritative sources to reduce hallucination [54].

2.4. AI Ethics and Governance in Auditing

The increasing role of AI in high-stakes decision-making has motivated extensive work on AI auditing and governance. Mökander [29] articulated a three-layered approach (governance, model, application) to AI auditing, while Mökander et al. [30] proposed similar layered methods for LLM evaluation. Costanza-Chock et al. [56] examined the algorithmic auditing ecosystem, identifying gaps in current practice. Schiff et al. [57] reviewed AI ethics frameworks across sectors. The forthcoming ISO/IEC 42001:2023 [58] provides a management system standard for AI itself, defining requirements for organisations to govern AI development and use. The framework presented here draws on this literature to ensure that ML-augmented audits maintain accountability, transparency, and ethical alignment.

2.5. Research Gap

Despite extensive research in each of these areas separately, three gaps remain. First, there is no integrated conceptual framework that combines NLP, CV, and LLM-based interviewing for ISO 27001 audits, with explicit triangulation across modalities. Second, prior work has not provided a fine-grained mapping of ML capabilities to specific ISO 19011 audit phases or to the 93 controls of ISO/IEC 27001:2022 Annex A. Third, the ethical, accountability, and explainability requirements of audit settings have rarely been treated systematically in the ML-for-audit literature. The M³A-Framework presented in this paper addresses all three gaps.

3. Methodology

We adopt the design science research (DSR) paradigm [31,32] to develop the M³A-Framework. DSR is well-suited to information systems artefacts that solve practical problems while contributing to scholarly knowledge. Following the canonical DSR process [32], our research consists of six activities: (1) problem identification (Section 1), (2) definition of objectives, (3) design and development (Section 4), (4) demonstration through illustrative scenarios, (5) evaluation through proposed metrics (Section 6), and (6) communication via this publication.

Our literature review followed the hermeneutic approach articulated by Boell and Cecez-Kecmanovic [59] and the literature review methodology of Kraus et al. [59]. We searched Scopus, Web of Science, IEEE Xplore, ACM Digital Library, and Google Scholar using combinations of keywords spanning ISO 27001, ISMS audit, machine learning, NLP, computer vision, LLM, and audit augmentation. The search yielded an initial set of approximately 350 candidate papers, which were

screened by title, abstract, and full-text against inclusion criteria. After de-duplication and screening, 91 references were retained for the conceptual development.

The conceptual framework was developed iteratively through the design and evaluation cycles described by Jaakkola [60]. We followed Whetten's [61] criteria for theoretical contribution: the framework explicitly identifies what (the modular components), how (the data flows and triangulation logic), why (the underlying audit-theoretic and ML-theoretic justifications), and who/where/when (the deployment context). The framework was refined through three internal review iterations against ISO 19011 audit principles, ISO/IEC 27001:2022 controls, and ML feasibility constraints.

4. The M³A-Framework

4.1. Framework Overview

The M³A-Framework consists of five sequential stages, summarised in Figure 1. The stages map to and augment the ISO 19011 audit phases. Each stage takes structured inputs, applies ML-based processing, and produces structured outputs that feed into subsequent stages. The framework's central design principle is multi-modal triangulation: a finding is accepted with high confidence only when supported by evidence from at least two independent modalities (document, observation, interview). This principle directly operationalises ISO 19011's evidence-based approach principle.

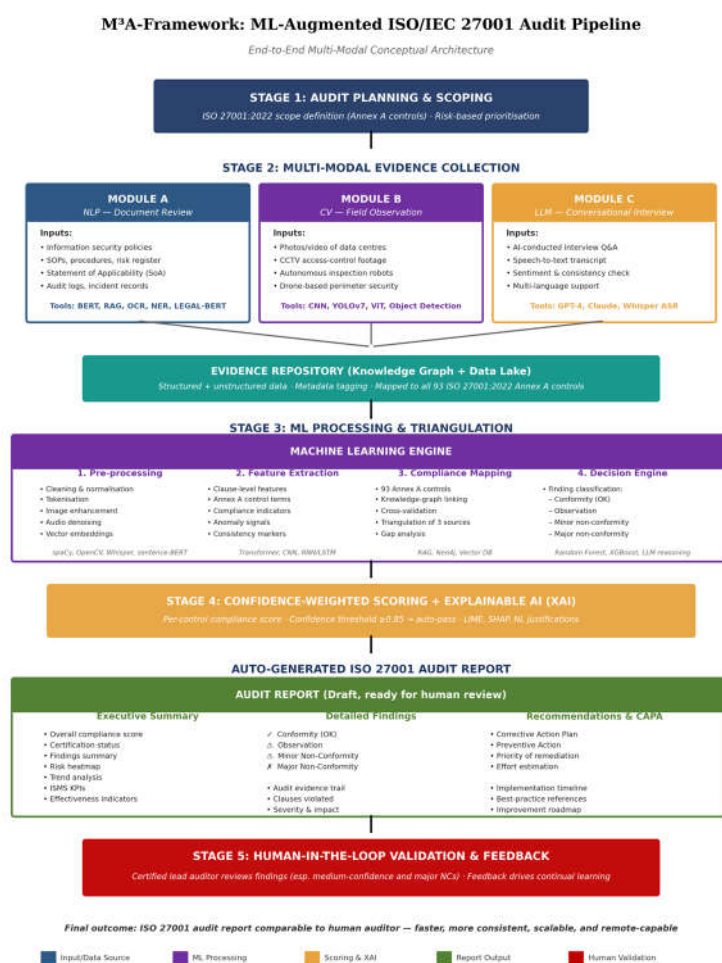


Figure 1. The Multi-Modal ML-Augmented ISO 27001 Audit Framework (M³A-Framework). Five-stage pipeline encompassing audit planning, multi-modal evidence collection, ML-based processing, scoring with confidence weighting, and human-in-the-loop validation.

4.2. Stage 1: Audit Planning and Scoping

The framework begins with structured audit planning aligned with ISO 19011 Clause 6.3 [5]. The certified lead auditor and the auditee jointly define the audit scope: which ISO/IEC 27001 controls are in scope, which organisational units are involved, what evidence will be considered, and what timeline applies. ML augmentation supports this stage in two ways. First, a control criticality predictor based on the auditee's industry, prior incident history, and risk register output a prioritised control list — focusing the audit on areas of highest risk. Second, a workload estimator forecasts the resources (auditor-hours, ML compute, sensor deployment) needed to complete the audit within the agreed timeline. Both tools draw on historical audit metadata and operationalise the risk-based approach principle added to ISO 19011 in 2018 [5].

4.3. Stage 2: Multi-Modal Evidence Collection

The second stage is the heart of the framework. Three parallel modules collect evidence from three modalities, each addressing a distinct subset of the 93 Annex A controls.

4.3.1. Module A: NLP-Based Document Review

Module A processes the auditee's documented information — policies, procedures, the Statement of Applicability (SoA), risk register, incident logs, training records, and supplier contracts. The pipeline uses domain-adapted transformer models [13,47] for requirement extraction, retrieval-augmented generation [20] for cross-referencing against ISO 27001 control objectives, and named entity recognition for identifying obligations. Recent advances in legal NLP [47] enable automated checking of compliance language quality, while embedding-based similarity search [66] supports gap analysis against authoritative templates. The output of Module A is a structured set of evidence items, each tagged with the control it addresses and a confidence score.

4.3.2. Module B: CV-Based Field Observation

Module B captures physical evidence relevant to ISO 27001 Annex A theme A.7 (physical controls). It comprises three sub-modules: (i) static camera analytics for monitoring access controls, surveillance systems, and clean-desk compliance [25,48]; (ii) autonomous inspection robots [49,67] that traverse data centres on pre-defined patrol paths and detect anomalies (open cabinets, unauthorised devices, environmental issues); and (iii) drone-based aerial surveillance [63] for perimeter inspection of large facilities. Detection models [16–18,64] localise objects of interest, and a downstream verification module compares detections against the policy-specified state. Module B operates with strict privacy safeguards: face blurring is applied by default, and personal data subject to GDPR [80] is processed on-premise where feasible.

4.3.3. Module C: LLM-Based Conversational Interview

Module C conducts semi-structured interviews with auditee personnel covering people-related controls (theme A.6) and organisational controls (theme A.5). The system uses speech recognition [19] for voice transcription, an LLM-based agent [14,15,65] for adaptive questioning, and a retrieval-augmented generation system [20] for cross-checking responses against documented policies. The interview is structured as a tree of questions whose ordering depends on prior responses, enabling efficient coverage of relevant controls. To maintain auditor accountability, the LLM agent is configured to surface uncertainty markers and to flag responses that warrant follow-up by a human auditor.

4.4. Stage 3: ML-Based Processing and Triangulation

The third stage receives evidence from the three modules and produces control-level findings. A central knowledge graph stores ISO/IEC 27001 control definitions, ISO/IEC 27002 implementation

guidance [4], and the auditee's contextual information. For each control, evidence items from multiple modalities are aggregated using a triangulation logic that requires at least two independent corroborating sources before a finding is upgraded to high confidence. An ensemble decision engine combines XGBoost [68] for structured feature classification with LLM-based reasoning for unstructured text interpretation. Anomaly detection [42,48] flags evidence patterns that deviate from expected norms.

4.5. Stage 4: Confidence-Weighted Scoring with XAI

The fourth stage classifies each control's compliance status into four categories — conformity, observation, minor non-conformity, and major non-conformity — together with a confidence score. The scoring uses a calibrated probabilistic model [89] to ensure that confidence values are interpretable as probabilities. Three confidence bands are defined:

- High confidence (≥ 0.85) — auto-promoted to the audit report.
- Medium confidence (0.60–0.84) — flagged for human review.
- Low confidence (< 0.60) — referred back for additional evidence collection.

Explainable AI methods [21,22] generate per-finding justifications drawing on LIME [21], SHAP [69], and natural-language explanations from the LLM. The output of stage 4 is a draft audit report containing executive summary, control-by-control findings, and corrective and preventive action (CAPA) recommendations.

4.6. Stage 5: Human-in-the-Loop Validation and Feedback

A certified lead auditor reviews the draft, focusing on medium-confidence findings and all major non-conformities. The auditor accepts, rejects, or modifies each finding. Acceptance and rejection signals are written back to the model registry as labelled examples for continual learning [70]. The final report is signed by the lead auditor and issued to the auditee. The framework's design preserves the legal and ethical accountability of the certified auditor, in line with the recommendations of Mökander [29] and ISO/IEC 42001 [58].

4.7. Augmenting ISO 19011:2018 Evidence-Collection Process

While Section 4.4 introduced the ML-based processing layer at a high level, the framework's principal methodological contribution is its tight coupling with the seven-step evidence-collection process specified in ISO 19011:2018, Clause 6.4.7 and visualised as Figure 2 in that standard [5]. The standard prescribes the following sequence: (i) sources of information, (ii) collecting by means of appropriate sampling and verifying, (iii) audit evidence, (iv) evaluating against audit criteria, (v) audit findings, (vi) reviewing audit findings, and (vii) audit conclusions. Each step is open to ML augmentation, but the augmentation must respect the standard's underlying principles — particularly the evidence-based approach and due professional care articulated in Clause 4 [5]. Figure 2 of this paper presents the augmented version of the ISO 19011 collection process, with the conventional steps in the left column and the corresponding ML-augmented steps in the right column.

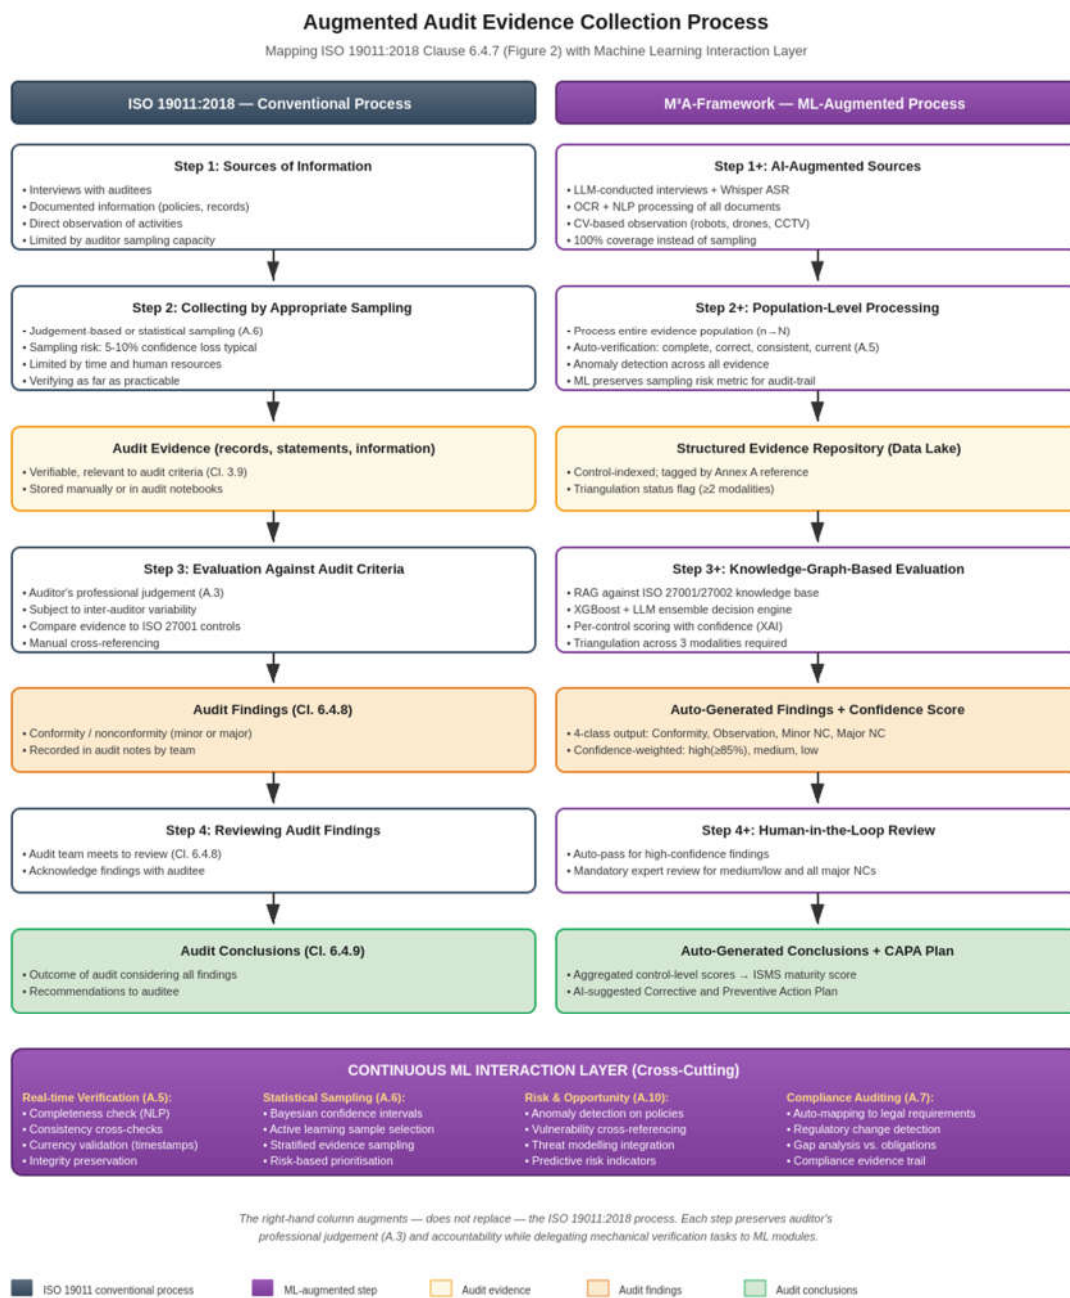


Figure 2. Augmented audit evidence-collection process. The left column reproduces the seven-step process of ISO 19011:2018 Clause 6.4.7 (Figure 2). The right column shows the corresponding ML-augmented steps in the M³A-Framework. The bottom band represents the continuous ML interaction layer.

4.7.1. Step-by-Step Augmentation Rationale

Sources of information. ISO 19011 enumerates interviews, observations, and documented information as the primary sources [5]. The M³A-Framework retains these three sources but extends each: interviews are conducted by an LLM-based conversational agent (Module C); observations are captured through CV-enabled sensors, robots, drones, and CCTV (Module B); and documented information is processed at scale through a domain-adapted NLP pipeline (Module A). Crucially, the extension does not eliminate human-driven sources — auditors may still conduct interviews directly, and the ML modules complement rather than replace the human auditor's senses.

Collecting and verifying through sampling. Annex A.6 of ISO 19011 distinguishes two sampling approaches: judgement-based sampling (A.6.2) and statistical sampling (A.6.3) [5]. The framework redefines this trade-off. Where computational resources permit, the ML modules process the entire evidence population ($n \rightarrow N$), thereby eliminating sampling risk for that step. Where computational cost demands sampling, the framework adopts active learning [71] and Bayesian-confident sampling [72] that select the most informative subset of evidence — typically achieving the same audit confidence with substantially fewer samples than statistical random sampling.

Verifying information against the four-criterion test of A.5. Annex A.5 of ISO 19011 lists four properties that audit information should possess: complete, correct, consistent, and current [5]. These four properties translate directly into ML verification subtasks. Completeness is checked by NLP coverage models. Correctness is verified through cross-referencing against authoritative sources via RAG [20]. Consistency is checked through pairwise document comparison and triangulation across the three modalities. Currency is monitored by timestamp analysis.

Generating findings and reviewing them (Clauses 6.4.8–6.4.9). Findings are produced by the decision engine and then reviewed by the audit team — both AI and human. The framework distinguishes itself from prior automated-audit work by enforcing a two-stage review: an initial AI-internal review where the LLM-based reasoner challenges its own findings against alternative interpretations of the evidence (a self-consistency check inspired by [73]), followed by the mandatory human review described in Stage 5.

Audit conclusions. The framework aggregates per-control findings into an overall ISMS maturity score and a list of CAPA recommendations. The aggregation strategy is informed by the risk-based approach added to ISO 19011:2018's principles in its third edition [5] and explicitly considers the auditee's context and life-cycle perspective where applicable.

4.8. Mapping the Framework to ISO 19011 Annex A Audit Methods Matrix

ISO 19011:2018 Annex A.1 (Table A.1) classifies audit methods along two axes: extent of involvement between the auditor and the auditee (interactive vs. non-interactive) and location of the auditor (on-site vs. remote) [5]. This 2x2 matrix yields four quadrants — Interactive On-Site, Interactive Remote, Non-Interactive On-Site, and Non-Interactive Remote — each suited to different audit objectives. The third edition of ISO 19011 added significant guidance on the remote column reflecting the standard's recognition that ICT-supported auditing is now a first-class audit method [5]. The M³A-Framework offers an additional axis — ML augmentation level — that maps each quadrant to a different class of ML modules. Figure 3 presents this extended matrix.

Audit Methods Matrix: ML-Augmented Extension of ISO 19011:2018 Annex A, Table A.1

Mapping the four quadrants of audit methods to dedicated ML modules in the MPA-Framework.

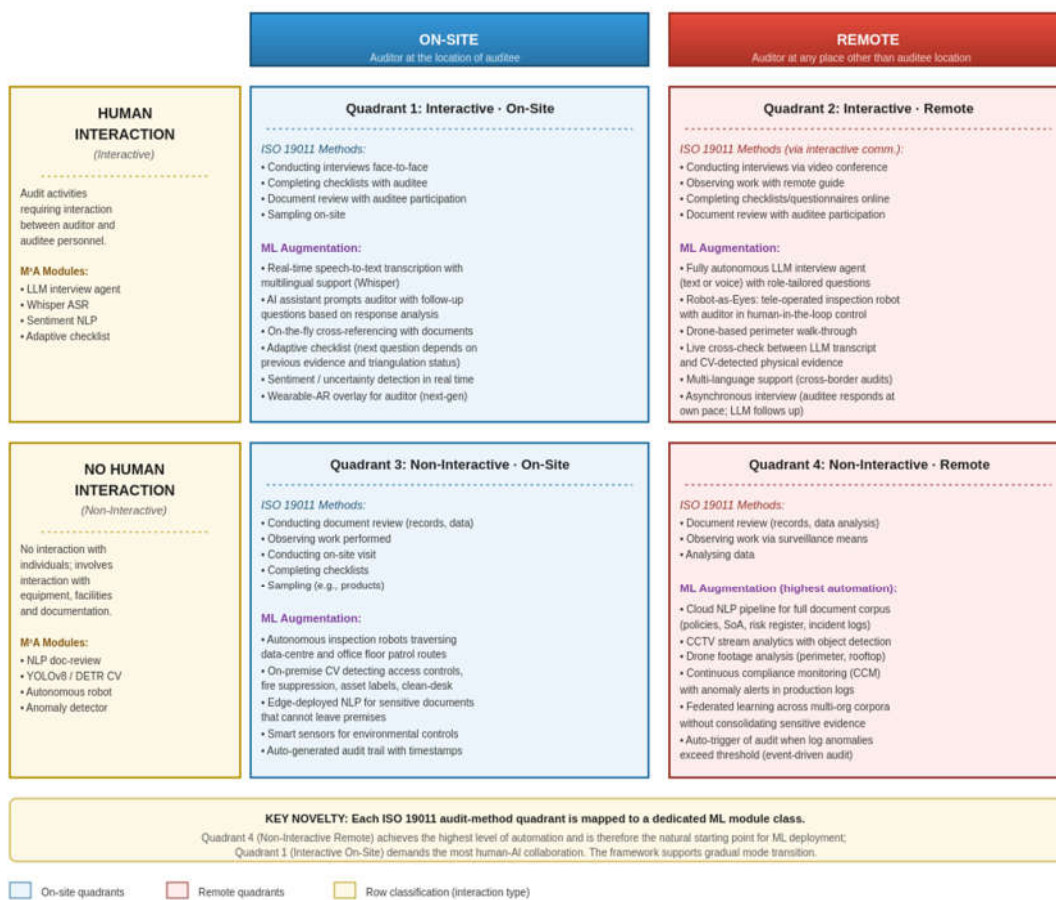


Figure 3. ML-augmented audit-methods matrix, extending ISO 19011:2018 Annex A, Table A.1. Each of the four quadrants is mapped to a dedicated class of ML modules. Quadrant 4 (Non-Interactive Remote) supports the highest degree of automation; Quadrant 1 (Interactive On-Site) requires the deepest human-AI collaboration.

4.8.1. Quadrant-Specific ML Augmentation Strategies

Quadrant 1 – Interactive On-Site. This quadrant covers face-to-face interviews, joint document reviews, and on-site observations. ML augments rather than replaces the human auditor: real-time speech-to-text [19], LLM-based follow-up question generation, adaptive checklist re-ordering, and AR overlays [74,75] increase the auditor’s effective bandwidth without altering the audit’s interpersonal character.

Quadrant 2 – Interactive Remote. This quadrant covers video-conferencing interviews, remote-guided observation, and online questionnaires. The framework introduces fully autonomous LLM interview agents and tele-operated robots acting as the auditor’s eyes [49], making audits feasible across borders and time zones.

Quadrant 3 – Non-Interactive On-Site. ML augmentation here is dominated by edge-deployed models that respect data sovereignty: NLP for confidential documents that cannot leave the premises, autonomous robotic patrols at night, and continuous sensor monitoring. This deployment is particularly relevant for high-confidentiality sectors.

Quadrant 4 – Non-Interactive Remote. This quadrant exhibits the highest level of automation potential: cloud-based document analysis, CCTV stream analytics, drone footage processing, and continuous compliance monitoring (CCM) [77]. Here the framework can deliver event-driven audits – triggering a focused audit when anomalies exceed a configured threshold – and federated learning [79] enables multi-organisation models without centralising sensitive evidence.

4.8.2. Adoption Pathway Across Quadrants

The four quadrants suggest a natural incremental adoption pathway. Organisations beginning their ML-augmented audit journey should typically start in Quadrant 4 (Non-Interactive Remote), where ML augmentation is least disruptive and the technology is most mature. Once trust is established, the organisation can extend to Quadrant 3, then Quadrants 1 and 2. This phased pathway aligns with the maturity-based audit programme guidance in ISO 19011 Clause 5.4 [5].

4.8.3. Compatibility with the Seven Principles of Auditing

ISO 19011 Clause 4 establishes seven principles that auditing should respect: integrity, fair presentation, due professional care, confidentiality, independence, evidence-based approach, and risk-based approach [5]. Each principle is addressed in the framework as follows:

- Integrity. The framework's human-in-the-loop design preserves auditor accountability; the LLM modules cannot overrule the auditor's judgement on ethical questions.
- Fair presentation. Confidence scores and XAI-based justifications [21,22,69] enable transparent reporting of all findings, including uncertain ones.
- Due professional care. AI literacy is added to auditor competence requirements; auditors retain the authority to override ML conclusions.
- Confidentiality. On-premise and federated learning options [79] support handling of sensitive evidence; encryption and differential privacy [80] are applied throughout.
- Independence. The framework treats the ML system itself as an auditable component governed by ISO/IEC 42001 [58]; bias testing addresses model-induced impartiality risks.
- Evidence-based approach. The triangulation requirement (Section 4.4) operationalises this principle.
- Risk-based approach. The control criticality predictor in Stage 1 (Section 4.2) operationalises risk-based prioritisation.

5. Mapping to ISO/IEC 27001:2022 and Implementation Considerations

5.1. Coverage of Annex A Control Themes

Table 1 maps the framework's three evidence-collection modules to the four themes of ISO/IEC 27001:2022 Annex A. The mapping reveals that each control theme draws primarily on a different modality, with significant overlap that supports triangulation.

Table 1. Mapping of evidence-collection modules to ISO/IEC 27001:2022 Annex A control themes.

Annex A theme	# Controls	Document review (NLP)	Field observation (CV)	Interview (LLM)
Organisational	37	Primary	Limited	Secondary
People	8	Secondary	Limited	Primary
Physical	14	Limited	Primary	Secondary
Technological	34	Primary	Secondary	Secondary

The pattern is clear: organisational and technological controls (71 of 93) are predominantly verifiable through documentary evidence; physical controls require CV-based observation; people-related controls (training, awareness, responsibilities) are best assessed through interviews. The framework's strength lies in the triangulation of these complementary perspectives.

5.2. Mapping to ISO 19011 Audit Phases

Table 2 maps the framework's stages to the audit phases defined in ISO 19011:2018 [5], confirming methodological compatibility with established audit practice.

Table 2. Mapping of framework stages to ISO 19011:2018 audit phases.

ISO 19011 phase	Clause	M ³ A-Framework stage	ML augmentation
Initiating audit	6.2	Stage 1: Planning	Risk-based control prioritisation
Preparing audit activities	6.3	Stage 1: Planning	Workload estimation, scheduling
Conducting audit activities	6.4	Stages 2–4: Collection, processing, scoring	All three modalities active
Preparing audit report	6.5	Stage 4: Scoring + Stage 5: Validation	Auto-draft generation, XAI justifications
Completing audit	6.6	Stage 5: Validation	Human-in-the-loop final review
Audit follow-up	6.7	Continuous monitoring	Federated learning of CAPA outcomes

5.3. Implementation Considerations

Three implementation considerations are critical. First, infrastructure: the framework requires computational resources for ML inference, secure storage for evidence, and integration with existing audit management systems. Second, data privacy: the framework's data flows must comply with GDPR [80] and equivalent regional regulations. Differential privacy [81] should be applied to model training where possible. Third, auditor training: ML augmentation does not eliminate the need for skilled auditors; rather, it shifts the competence profile to include AI literacy [82], statistical interpretation, and human-AI interaction skills.

6. Testable Propositions and Evaluation Metrics

6.1. Five Testable Propositions

The framework gives rise to five testable propositions, each predicting a specific direction of effect for outcomes that can be measured in empirical studies.

P1 (Speed): ML-augmented audits will complete in significantly less wall-clock time than conventional audits of comparable scope, with the largest reductions in document-heavy controls.

P2 (Cost): ML-augmented audits will exhibit lower per-control marginal cost than conventional audits at scale, with the marginal cost approaching zero for documentary controls.

P3 (Consistency): ML-augmented audits will demonstrate higher inter-auditor agreement (Cohen's κ) than conventional audits when applied to identical evidence.

P4 (Coverage): ML-augmented audits will achieve higher control-coverage rates than conventional audits, particularly for technological controls.

P5 (Quality): Audit reports produced by ML-augmented audits will be rated by independent expert reviewers as equivalent to or higher in quality than those of conventional audits, with no degradation in finding accuracy.

6.2. Evaluation Metrics

Table 3 summarises the metrics we propose for evaluating the framework's empirical validation. The metrics are operationalised at three levels: process-level (auditing efficiency), output-level (audit report quality), and learning-level (ML model performance).

Table 3. Proposed evaluation metrics for empirical validation of the framework.

Level	Metric	Formula / Measurement	Reference
Process	Audit duration	Wall-clock hours from kick-off to draft report	[88]
Process	Auditor effort	Total person-hours per control	[88]
Process	Cost per control	Total cost / number of controls audited	[10]
Output	Inter-auditor agreement	Cohen's κ on independent re-audits	[83,85]
Output	Coverage rate	% of in-scope controls with conclusive finding	[88]
Output	Report quality	Expert rating on Likert scale (1-7)	[60]
Output	False-positive rate	Findings overturned by lead auditor / total findings	[88]
Learning	Model accuracy	Macro-F1 across the four finding classes	[88]
Learning	Calibration error	Expected Calibration Error (ECE)	[89]
Learning	Fairness metrics	Demographic parity, equalised odds across sectors	[90]

6.3. Suggested Empirical Validation Design

We recommend a multi-stage validation. A first stage consists of within-organisation parallel audits, in which a certified audit team and the M³A-Framework conduct independent audits of the same organisation, and outputs are compared on the metrics in Table 3. A second stage consists of multi-organisation longitudinal trials, in which the framework is deployed across a portfolio of organisations from different sectors over multiple audit cycles. A third stage involves blinded expert evaluation of report quality. Together these three stages can establish the framework's external and ecological validity [91].

7. Discussion

7.1. Theoretical Contributions

The framework contributes to the literature in five ways:

First, it **unifies three audit modalities** that have hitherto been treated separately, providing a coherent multi-modal architecture grounded in ISO 19011's evidence-corroboration principle. Second, it **operationalises the augmented-audit paradigm** that has been advocated theoretically by Mökander [29] and others, but rarely translated into deployable architectures. Third, it **bridges audit theory and ML methodology**, articulating how DSR principles [31,32] can be applied to compliance-domain artefacts. Fourth, it provides the **first detailed step-level mapping** of ML augmentation to the seven-step evidence-collection process of ISO 19011:2018 Clause 6.4.7. Fifth, it **extends the audit-methods matrix** of ISO 19011 Annex A Table A.1 by introducing ML augmentation as a third axis.

7.2. Practical Implications

For practitioners, the framework offers a deployment blueprint that can be adopted incrementally. Organisations can begin with Module A (NLP-based document review) for high-volume documentary controls, expand to Module B (CV-based observation) as physical infrastructure permits, and incorporate Module C (LLM-based interview) for distributed teams. Certification bodies can accept M³A-Framework outputs as supporting evidence for their certification decisions, provided the human-in-the-loop validation step is preserved.

7.3. Ethical Considerations

Three ethical concerns require careful management. First, accountability: the certified lead auditor remains the locus of legal and ethical responsibility, irrespective of ML augmentation. Second, fairness: ML models trained on historical audit data may inherit biases against particular industries, geographies, or organisational types; bias audits are required prior to deployment [90]. Third, privacy: evidence collected by Module B (especially video) and Module C (interview transcripts) must be handled with strict data-protection safeguards including consent, minimisation, and time-bounded retention.

7.4. Limitations

We acknowledge several limitations. First, the framework is conceptual; empirical validation against the proposed metrics is yet to be conducted. Second, the framework assumes a baseline level of organisational digitalisation; auditees with predominantly paper-based records may require additional preprocessing. Third, ML model performance varies across languages and cultural contexts; cross-cultural validation is needed. Fourth, the framework's effectiveness depends on the quality and breadth of its training data, which may be limited in early deployments.

7.5. Future Research

The framework opens several research avenues. Empirical validation through controlled trials is the most immediate need. Beyond that, we identify five research directions: (a) construction of an open ISO 27001 audit dataset for the ML community; (b) cross-standard generalisation, including ISO 9001, ISO 14001, ISO 45001, and SOC 2; (c) integration with continuous compliance monitoring, moving from periodic audits toward continuous assurance [76]; (d) extension to AI management system audits under ISO/IEC 42001; and (e) exploration of federated and on-device learning to address privacy concerns [79,80].

8. Conclusions

This paper has proposed a multi-modal, ML-augmented conceptual framework — the M³A-Framework — for ISO/IEC 27001 audits. Drawing on the design science research paradigm and a synthesis of audit theory, NLP, computer vision, and conversational AI, the framework integrates document review, field observation, and interviewing into a coherent five-stage pipeline aligned with ISO 19011. The framework is mapped to all four themes of ISO/IEC 27001:2022 Annex A and to the audit phases of ISO 19011:2018, demonstrating both completeness and methodological compatibility with established audit practice.

The framework's central proposition is that, by combining three complementary AI modalities under a triangulation discipline and human-in-the-loop validation, ISO 27001 audits can become significantly faster, less expensive, more consistent, and more scalable, while maintaining or improving the quality of audit reports. Five testable propositions and a battery of evaluation metrics are offered to guide the next wave of empirical research.

We conclude by noting that the framework is not a call to replace human auditors but to augment them. The certified lead auditor remains the locus of professional judgement, legal accountability, and ethical responsibility. ML serves as a force multiplier — collecting and processing evidence at a scale unachievable by manual methods, surfacing the patterns that demand human attention, and

freeing auditors to focus on the high-judgement aspects of their work. If the framework realises its potential, it could materially improve the resilience of organisations' information security posture worldwide, with particular benefit for SMEs and organisations in geographically distributed contexts that have historically been underserved by traditional audit models.

Supplementary Materials: The following supporting information can be downloaded at the website of this paper posted on Preprints.org.

Author Contributions: Conceptualization, N.A.C.; methodology, N.A.C.; investigation, N.A.C.; formal analysis, N.A.C.; writing—original draft preparation, N.A.C.; writing—review and editing, N.A.C.; visualization, N.A.C.; supervision, N.A.C.; project administration, N.A.C.; funding acquisition, N.A.C. The author has read and agreed to the published version of the manuscript.

Funding: This research was funded by PT. TSI Sertifikasi Internasional, an accredited certification body, as part of its programme to advance research and innovation in management system auditing methodologies. The APC was funded by PT. TSI Sertifikasi Internasional.

Data Availability Statement: No new data were created in this conceptual paper. All sources used in the development of the framework are listed in the references and are publicly available through the cited publishers and repositories.

Acknowledgments: The author gratefully acknowledges the support of PT. TSI Sertifikasi Internasional for funding this research, and the Faculty of Computer Sciences at Universitas Mercu Buana for providing the academic environment and resources that enabled the development of this conceptual framework. During the preparation of this manuscript, the author used Claude (Anthropic) for the purposes of language refinement, manuscript drafting support, and structural organisation. The author has reviewed and edited the output and takes full responsibility for the content of this publication. The conceptual framework, research design, methodology, analysis, and all scholarly contributions are entirely the author's own original work.

Conflicts of Interest: The author declares no conflicts of interest. The funding body PT. TSI Sertifikasi Internasional had no role in the conceptual design of the framework, the analysis of the literature, the writing of the manuscript, or the decision to publish.

References

1. von Solms, R.; van Niekerk, J. From information security to cyber security. *Comput. Secur.* 2013, 38, 97-102. <https://doi.org/10.1016/j.cose.2013.04.004>
2. AlGhamdi, S.; Win, K.T.; Vlahu-Gjorgievska, E. Information security governance challenges and critical success factors: systematic review. *Comput. Secur.* 2020, 99, 102030. <https://doi.org/10.1016/j.cose.2020.102030>
3. International Organization for Standardization. ISO/IEC 27001:2022 — Information security, cybersecurity and privacy protection — Information security management systems — Requirements. ISO, Geneva (2022)
4. International Organization for Standardization. ISO/IEC 27002:2022 — Information security, cybersecurity and privacy protection — Information security controls. ISO, Geneva (2022)
5. International Organization for Standardization. ISO 19011:2018 — Guidelines for auditing management systems. ISO, Geneva (2018)
6. Russell, J.P. *The ASQ Auditing Handbook*, 4th edn. ASQ Quality Press, Milwaukee (2013)
7. Power, M. *The Audit Society: Rituals of Verification*. Oxford University Press, Oxford (1997)
8. DeFond, M.L.; Zhang, J. A review of archival auditing research. *J. Account. Econ.* 58(2-3), 275-326 2014. <https://doi.org/10.1016/j.jacceco.2014.09.002>
9. Talapatra, S.; Santos, G.; Uddin, K.; Carvalho, F. Main benefits of integrated management systems through literature review. *Int. J. Qual. Res.* 2019, 13(4), 1037-1054. <https://doi.org/10.24874/IJQR13.04-19>
10. Mirtsch, M.; Kinne, J.; Blind, K. Exploring the adoption of the international information security management system standard ISO/IEC 27001: a web mining-based analysis. *IEEE Trans. Eng. Manage.* 2021, 68(1), 87-100. <https://doi.org/10.1109/TEM.2020.2977815>

11. Castka, P.; Searcy, C.; Fischer, S. Technology-enhanced auditing in voluntary sustainability standards: the impact of COVID-19. *Sustainability* 2020, 12(11), 4740. <https://doi.org/10.3390/su12114740>
12. Kafel, T.; Pawlak, J.; Sikora-Alicka, K. Remote audits as a response to challenges associated with the COVID-19 pandemic. *Sci. Pap. Silesian Univ. Technol. Organ. Manag. Ser.* 152, 117-134 (2021)
13. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of NAACL-HLT 2019*, pp. 4171-4186. ACL, Minneapolis (2019)
14. OpenAI. GPT-4 Technical Report. arXiv:2303.08774 (2023)
15. Anthropic. Model Card and Evaluations for Claude Models. Anthropic, San Francisco (2023)
16. Wang, C.-Y.; Bochkovskiy, A.; Liao, H.-Y. M. YOLOv7: trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: *Proceedings of CVPR 2023*, pp. 7464-7475 (2023)
17. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (ICLR)* (2021)
18. Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; et al. Swin transformer: hierarchical vision transformer using shifted windows. In: *Proceedings of ICCV 2021*, pp. 10012-10022 (2021)
19. Radford, A.; Kim, J.W.; Xu, T.; Brockman, G.; McLeavey, C.; Sutskever, I. Robust speech recognition via large-scale weak supervision. In: *Proceedings of ICML 2023*, pp. 28492-28518 (2023)
20. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Adv. Neural Inf. Process. Syst.* 33, 9459-9474 (2020)
21. Ribeiro, M.T.; Singh, S.; Guestrin, C. Why should I trust you? Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD*, pp. 1135-1144 2016. <https://doi.org/10.1145/2939672.2939778>
22. Saeed, W.; Omlin, C. Explainable AI (XAI): a systematic meta-survey of current challenges and future opportunities. *Knowl. Based Syst.* 2023, 263, 110273. <https://doi.org/10.1016/j.knosys.2023.110273>
23. Abualhaija, S.; Arora, C.; Sleimi, A.; Briand, L.C. Automated question answering for improved understanding of compliance requirements: a multi-document study. In: *Proceedings of RE 2022*, pp. 39-50. IEEE 2022. <https://doi.org/10.1109/RE54965.2022.00012>
24. Amaral, O.; Azeem, M.I.; Abualhaija, S.; Briand, L.C. NLP-based automated compliance checking of data processing agreements against GDPR. *IEEE Trans. Softw. Eng.* 2023, 49(9), 4282-4297. <https://doi.org/10.1109/TSE.2023.3288901>
25. Rai, A.; Pugazhenth, V.; Nair, B.B.; Kumar, S.B. Smart surveillance with computer vision and IoT for security automation. In: *2021 IEEE 4th International Conference on Computing, Power and Communication Technologies*, pp. 1-6 (2021)
26. Nayak, R.; Pati, U.C.; Das, S.K. A comprehensive review on deep learning-based methods for video anomaly detection. *Image Vis. Comput.* 2021, 106, 104078. <https://doi.org/10.1016/j.imavis.2020.104078>
27. Bilquise, G.; Ibrahim, S.; Shaalan, K. Emotionally intelligent chatbots: a systematic literature review. *Hum. Behav. Emerg. Technol.* 2022, 2022, 9601630. <https://doi.org/10.1155/2022/9601630>
28. Almansor, E.H.; Hussain, F.K. Survey on intelligent chatbots: state-of-the-art and future research directions. In: Barolli, L., et al. (eds.) *Complex, Intelligent, and Software Intensive Systems. CISIS 2019. AISC*, vol. 993, pp. 534-543. Springer, Cham (2020)
29. Mökander, J. Auditing of AI: legal, ethical and technical approaches. *Digit. Soc.* 2023, 2(3), 49. <https://doi.org/10.1007/s44206-023-00074-y>
30. Mökander, J.; Schuett, J.; Kirk, H.R.; Floridi, L. Auditing large language models: a three-layered approach. *AI Ethics* 2024, 4, 1085-1115. <https://doi.org/10.1007/s43681-023-00289-2>
31. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design science in information systems research. *MIS Q.* 2004, 28(1), 75-105. <https://doi.org/10.2307/25148625>
32. Peffers, K.; Tuunanen, T.; Rothenberger, M.A.; Chatterjee, S. A design science research methodology for information systems research. *J. Manage. Inf. Syst.* 2007, 24(3), 45-77. <https://doi.org/10.2753/MIS0742-1222240302>
33. Sartor, M.; Orzes, G.; Touboulic, A.; Culot, G.; Nassimbeni, G. ISO 2019, 14001, standard: literature review and theory-based research agenda. *Qual. Manag. J.* 26(1), 32-64. <https://doi.org/10.1080/10686967.2018.1542288>

34. Fonseca, L.M.; Domingues, J.P.; Dima, A.M. Mapping the sustainable development goals relationships. *Sustainability* 2020, 12(8), 3359. <https://doi.org/10.3390/su12083359>
35. Bravi, L.; Murmura, F.; Santos, G. Additive manufacturing: possible problems with indoor air quality. *Procedia Manuf.* 2019, 41, 952-959. <https://doi.org/10.1016/j.promfg.2019.10.020>
36. Sussy, B.E.; Wilber, C.; Milagros, L.; Carlos, M. ISO/IEC 27001 implementation in public organizations: an exploratory study. In: 2015 10th Iberian Conference on Information Systems and Technologies, pp. 1-6. IEEE 2015. <https://doi.org/10.1109/CISTI.2015.7170343>
37. Topa, I.; Karyda, M. From theory to practice: guidelines for enhancing information security management. *Inf. Comput. Secur.* 2019, 27(3), 326-342. <https://doi.org/10.1108/ICS-09-2018-0108>
38. Boz, B.; Mendling, J.; Polancic, G. Conceptualizing and measuring auditor competence in continuous auditing. *Account. Inf. Syst.* 2020, 39(2), 101-115. <https://doi.org/10.1016/j.accinf.2020.100467>
39. Culot, G.; Nassimbeni, G.; Podrecca, M.; Sartor, M. The ISO/IEC 27001 information security management standard: literature review and theory-based research agenda. *TQM J.* 2021, 33(7), 76-105. <https://doi.org/10.1108/TQM-09-2020-0202>
40. Hasan, A.R. Artificial Intelligence (AI) in accounting & auditing: a literature review. *Open J. Bus. Manag.* 2022, 10(1), 440-465. <https://doi.org/10.4236/ojbm.2022.101026>
41. Goto, M. Anticipatory innovation governance for the age of generative AI. *AI Soc.* 2024, 39, 2287-2302. <https://doi.org/10.1007/s00146-024-01874-7>
42. Hilal, W.; Gadsden, S.A.; Yawney, J. Financial fraud: a review of anomaly detection techniques and recent advances. *Expert Syst. Appl.* 2022, 193, 116429. <https://doi.org/10.1016/j.eswa.2021.116429>
43. Apruzzese, G.; Laskov, P.; Tastemirova, A. SoK: the impact of unlabelled data in cyberthreat detection. In: Proceedings of IEEE European Symposium on Security and Privacy (EuroS&P) 2022, pp. 20-42 2022. <https://doi.org/10.1109/EuroSP53844.2022.00010>
44. Abualhaija, S.; Arora, C.; Sleimi, A.; Briand, L.C. Automated question answering for improved understanding of compliance requirements: a multi-document study. In: Proceedings of RE 2022, pp. 39-50. IEEE 2022. <https://doi.org/10.1109/RE54965.2022.00012>
45. Shaikhanova, A.; Kuznetsov, O.; Tokkuliyeva, A.; Ayapbergenov, K.; Olzhas, S.; Danir, T. Security audit of IoT device networks: a reproducible machine learning framework for threat detection and performance benchmarking. *Sensors* 2025, 25(24), 7519. <https://doi.org/10.3390/s25247519>
46. Williams, J.; Walker, M.J.A. Machine learning in healthcare compliance: a systematic review. *Healthc. Inform. Res.* 30(3), 192-208 (2024)
47. Chalkidis, I.; Pasini, T.; Zhang, S.; Tomada, L.; Schwemer, S.F.; Søgaard, A. FairLex: a multilingual benchmark for evaluating fairness in legal text processing. In: Proceedings of ACL 2022, pp. 4389-4406 2022. <https://doi.org/10.18653/v1/2022.acl-long.301>
48. Liu, Y.; Yang, D.; Wang, Y.; Liu, J.; Liu, J.; Boukerche, A.; et al. Generalized video anomaly event detection: systematic taxonomy and comparison of deep models. *ACM Comput. Surv.* 2024, 56(7), 1-38. <https://doi.org/10.1145/3645101>
49. Karur, K.; Sharma, N.; Dharmatti, C.; Siegel, J.E. A survey of path planning algorithms for mobile robots. *Vehicles* 2021, 3(3), 448-468. <https://doi.org/10.3390/vehicles3030027>
50. Halder, S.; Afsari, K. Robots in inspection and monitoring of buildings and infrastructure: a systematic review. *Appl. Sci.* 2023, 13(4), 2304. <https://doi.org/10.3390/app13042304>
51. Touvron, H.; Martin, L.; Stone, K.; et al. Llama 2: open foundation and fine-tuned chat models. *arXiv:2307.09288* (2023)
52. Jiang, A.Q.; Sablayrolles, A.; Mensch, A.; et al. Mistral 7B. *arXiv:2310.06825* (2023)
53. Ji, Z.; Lee, N.; Frieske, R.; et al. Survey of hallucination in natural language generation. *ACM Comput. Surv.* 2023, 55(12), 1-38. <https://doi.org/10.1145/3571730>
54. International Accreditation Forum. IAF MD 4:2022 – Mandatory document for the use of information and communication technology (ICT) for auditing/assessment purposes. IAF, Cherrybrook (2022)
55. Castka, P.; Searcy, C.; Fischer, S. Technology-enhanced auditing: improving veracity and timeliness in social and environmental audits of supply chains. *J. Clean. Prod.* 2020, 258, 120773. <https://doi.org/10.1016/j.jclepro.2020.120773>

56. Costanza-Chock, S.; Raji, I.D.; Buolamwini, J. Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In: Proceedings of FAccT 2022, pp. 1571-1583. ACM 2022. <https://doi.org/10.1145/3531146.3533213>
57. Schiff, D.; Borenstein, J.; Biddle, J.; Laas, K. AI ethics in the public, private, and NGO sectors: a review of a global document collection. *IEEE Trans. Technol. Soc.* 2021, 2(1), 31-42. <https://doi.org/10.1109/TTS.2021.3052127>
58. International Organization for Standardization. ISO/IEC 42001:2023 – Information technology – Artificial intelligence – Management system. ISO, Geneva (2023)
59. Kraus, S.; Breier, M.; Lim, W.M.; Dabić, M.; Kumar, S.; Kanbach, D.; et al. Literature reviews as independent studies: guidelines for academic practice. *Rev. Manag. Sci.* 2022, 16, 2577-2595. <https://doi.org/10.1007/s11846-022-00588-8>
60. Jaakkola, E. Designing conceptual articles: four approaches. *AMS Rev.* 10(1-2), 18-26 2020. <https://doi.org/10.1007/s13162-020-00161-0>
61. Whetten, D.A. What constitutes a theoretical contribution? *Acad. Manag. Rev.* 1989, 14(4), 490-495. <https://doi.org/10.2307/258554>
62. Géron, A. Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow, 3rd edn. O'Reilly Media, Sebastopol (2022)
63. Mahony, J.; Walsh, P. Drones for industrial inspection and monitoring: a systematic review. *Drones* 7(7), 432 (2023)
64. Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; et al. DINO: DETR with improved denoising anchor boxes for end-to-end object detection. In: Proceedings of ICLR 2023 (2023)
65. Brown, T.B.; Mann, B.; Ryder, N.; et al. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* 33, 1877-1901 (2020)
66. Wang, L.; Yang, N.; Huang, X.; Jiao, B.; Yang, L.; Jiang, D.; et al. Text embeddings by weakly-supervised contrastive pre-training. *arXiv:2212.03533* (2022)
67. Flick, U. Doing Triangulation and Mixed Methods (Qualitative Research Kit), 2nd edn. SAGE, London 2018. ISBN 9781473953833
68. Chen, T.; Guestrin, C. XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD, pp. 785-794. ACM, San Francisco (2016)
69. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* 30, 4765-4774 (2017)
70. Wang, L.; Zhang, X.; Su, H.; Zhu, J. A comprehensive survey of continual learning: theory, method and application. *IEEE Trans. Pattern Anal. Mach. Intell.* 2024, 46(8), 5362-5383. <https://doi.org/10.1109/TPAMI.2024.3367329>
71. Ren, P.; Xiao, Y.; Chang, X.; Huang, P.-Y.; Li, Z.; Gupta, B.B.; et al. A survey of deep active learning. *ACM Comput. Surv.* 2022, 54(9), 1-40. <https://doi.org/10.1145/3472291>
72. Gelman, A.; Carlin, J.B.; Stern, H.S.; Dunson, D.B.; Vehtari, A.; Rubin, D.B. Bayesian Data Analysis, 3rd edn. Chapman and Hall/CRC, Boca Raton (2013)
73. Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.; Chi, E.; Narang, S.; et al. Self-consistency improves chain of thought reasoning in language models. In: Proceedings of ICLR 2023 (2023)
74. de Souza Cardoso, L.F.; Mariano, F.C.M.Q.; Zorzal, E.R. A survey of industrial augmented reality. *Comput. Ind. Eng.* 2020, 139, 106159. <https://doi.org/10.1016/j.cie.2019.106159>
75. Egger, J.; Masood, T. Augmented reality in support of intelligent manufacturing – a systematic literature review. *Comput. Ind. Eng.* 2020, 140, 106195. <https://doi.org/10.1016/j.cie.2019.106195>
76. Issa, H.; Sun, T.; Vasarhelyi, M.A. Research ideas for artificial intelligence in auditing: the formalization of audit and workforce supplementation. *J. Emerg. Technol. Account.* 2016, 13(2), 1-20. <https://doi.org/10.2308/jeta-10511>
77. Eulerich, M.; Wood, D.A. A demand for research on continuous auditing. *J. Inf. Syst.* 2022, 36(2), 41-62. <https://doi.org/10.2308/ISYS-2020-058>

78. Han, H.; Shiwakoti, R.K.; Jarvis, R.; Mordi, C.; Botchie, D. Accounting and auditing with blockchain technology and artificial intelligence: a literature review. *Int. J. Account. Inf. Syst.* 2023, 48, 100598. <https://doi.org/10.1016/j.accinf.2022.100598>
79. Liu, J.; Huang, J.; Zhou, Y.; Li, X.; Ji, S.; Xiong, H.; Dou, D. From distributed machine learning to federated learning: a survey. *Knowl. Inf. Syst.* 2022, 64, 885-917. <https://doi.org/10.1007/s10115-022-01664-x>
80. Wei, K.; Li, J.; Ding, M.; Ma, C.; Yang, H.H.; Farokhi, F.; et al. Federated learning with differential privacy: algorithms and performance analysis. *IEEE Trans. Inf. Forensics Secur.* 2020, 15, 3454-3469. <https://doi.org/10.1109/TIFS.2020.2988575>
81. European Parliament and Council. Regulation (EU) 2016/679 – General Data Protection Regulation. *Off. J. Eur. Union L* 119, 1-88 (2016)
82. Ng, D.T.K.; Leung, J.K.L.; Chu, S.K.W.; Qiao, M.S. Conceptualizing AI literacy: an exploratory review. *Comput. Educ. Artif. Intell.* 2021, 2, 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
83. Cohen, J. A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* 1960, 20(1), 37-46. <https://doi.org/10.1177/001316446002000104>
84. Krippendorff, K. *Content Analysis: An Introduction to Its Methodology*, 4th edn. SAGE, Thousand Oaks (2018)
85. Landis, J.R.; Koch, G.G. The measurement of observer agreement for categorical data. *Biometrics* 1977, 33(1), 159-174. <https://doi.org/10.2307/2529310>
86. Lehner, O.M.; Ittonen, K.; Silvola, H.; Ström, E.; Wührleitner, A. Artificial intelligence-based decision-making in accounting and auditing: ethical challenges and normative thinking. *Account. Audit. Account. J.* 2022, 35(9), 109-135. <https://doi.org/10.1108/AAAJ-09-2020-4934>
87. Davis, F.D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q.* 1989, 13(3), 319-340. <https://doi.org/10.2307/249008>
88. Sokolova, M.; Lapalme, G. A systematic analysis of performance measures for classification tasks. *Inf. Process. Manag.* 2009, 45(4), 427-437. <https://doi.org/10.1016/j.ipm.2009.03.002>
89. Wang, D.-B.; Feng, L.; Zhang, M.-L. Rethinking calibration of deep neural networks: do not be afraid of overconfidence. *Adv. Neural Inf. Process. Syst.* 34, 11809-11820 (2021)
90. Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; Galstyan, A. A survey on bias and fairness in machine learning. *ACM Comput. Surv.* 2022, 54(6), 1-35. <https://doi.org/10.1145/3457607>
91. Shadish, W.R.; Cook, T.D.; Campbell, D.T. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin, Boston (2002)

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.