A Semantic Audit Framework for National Security Reports Driven by Structural Empty Graph Networks (SEGN)

——A Methodological Analysis of Latvia's SAB Reporting Mechanism Based on the "Single Source of Truth – Computable Reasoning – Audit Interface" Framework

Wei Meng

Dhurakij Pundit University, Thailand

The University of Western Australia,AU

Fellow, Royal Anthropological Institute,UK

Email: weimeng4@acm.org

Abstract

Annual public reports issued by national security agencies serve concurrently as factual disclosures, explanatory mechanisms, and governance statements. Existing methodologies—including topic modelling, keyword co-occurrence analysis, and general information extraction—often struggle to align ‘assertion-evidence-mechanism-causality’ within a unified framework, thereby failing to establish verifiable inferential loops. This deficiency constitutes a knowledge gap in semantic auditing. Taking the 2025 Annual Report of Latvia's Constitution Protection Bureau (SAB) as its subject, this paper proposes a three-tier digital product roadmap driven by SEGN (Structured Empty Graph Network): It employs Single Source of Truth (SSOT) lossless graph binding to link page numbers with evidence fingerprints. The computable reasoning layer structures the concept-mechanism-causality chain while introducing evidence alignment and coverage thresholds (τ_a, τ_c). Finally, the audit interface records audit status (Pass/Warning/Fail) and uncertainty through field-based documentation. Results indicate that within this report context, mechanism nodes such as ‘legal warfare pathways, cybersecurity compliance implementation, and classified information protection governance’ occupy pivotal positions in structural centrality and evidence density. Furthermore, under specific conditions, enhanced evidence coverage facilitates mechanism inference migration from warning to pass levels, explainable through an ‘evidence alignment-threshold escalation’ audit mechanism. These findings support transforming national security texts from readable narratives into sets of inferential objects that are replayable, threshold-based, and falsifiable. This provides a unified language for inferential strength across annual comparisons and document migrations, while offering verifiable foundations for evidence-based risk control in public security and cyber governance.

Keywords: Semantic auditing of national security annual reports; Evidence alignment and threshold escalation mechanism; Structural Empty Graph Network (SEGN); Latvian SAB public reports

1. Introduction

Annual reports published by national security agencies constitute a highly hybrid genre of text: they must simultaneously disclose facts, uphold the credibility of capabilities and maintain a narrative of deterrence, while also identifying risk trends and governance levers within legal and institutional boundaries. The difficulty lies not in ‘information gaps’ but in the hierarchical structure and accountability of information: a single passage may simultaneously serve three functions—factual reporting, mechanism explanation, and policy implication. If the analytical framework fails to disentangle and realign ‘facts—evidence—mechanisms—causality—counterfactuals,’ two systemic errors readily emerge: firstly, mistaking rhetorical intensity for evidentiary strength; Second, mistaking co-occurrence for causation, thereby generating an ‘illusion of high probability’ that is neither auditable nor verifiable. Such issues are amplified in generative model scenarios, necessitating a structured approach that translates semantic analysis into auditable computational objects (NIST, 2023).

This paper's contribution lies in treating reports not as ‘summarisable text’ but as ‘auditable evidence-mechanism systems’. Using SEGN as a standardised framework, we transform reports into three-tiered digital products: an SSOT lossless graph ensures any conclusion can be traced back to source page numbers; a computable reasoning layer converts ‘concepts/mechanisms/causality’ into executable graph structures; the audit interface layer embeds evidence alignment, evidence coverage, uncertainty, and counterfactual testing into fielded rules. The integration of these layers elevates semantic analysis from ‘explanatory writing’ to ‘verifiable reasoning’.

2. SEGN Three-Layer Digital Product Methodology: From Text to 'Auditable Reasoning Systems'

2.1 Single Source of Truth (SSOT): Lossless, Traceable, Verifiable

The objective of the SSOT layer is not to 'extract as many triples as possible', but to establish a Single Source of Truth: each node/edge carries the source document, page number, span, evidence fingerprint (hash), and confidence level. This fulfils the minimum requirement that 'any inference must be traceable back to the original evidence'. This aligns with the lineage concept in the W3C framework: lineage information is used to assess quality, reliability, and trustworthiness (W3C, 2013a; W3C, 2013b). In this study, the SSOT decomposes reports into: —fact assertions (FACT_CLAIM)—mechanisms (MECHANISM)—evidence (EVIDENCE, evidence anchors)—context (CONTEXT), etc., while preserving the original report's paragraph and thematic organisation through structural edges.

2.2 Computable Reasoning: Transforming 'Explanation' into 'Constraint Propagation'

Traditional semantic analysis often remains confined to 'similarity spaces'. SEGN's computable reasoning layer maps key semantic objects in reports to executable reasoning graphs:

1) Mechanism links (MECHANISM_LINK): Specifies which factual claims imply particular mechanisms;

2) Causal links (CAUSAL): Embeds directional relationships ('causes/triggers/reinforces') into the graph structure;

3) Evidence alignment edges (EVIDENCE_ALIGN): Bind each critical edge to evidence nodes, preventing 'evidence-free causation'.

This approach is conceptually isomorphic to the 'evidence retrieval—claim validation' paradigm: claims must be accompanied by locatable evidence, otherwise they remain in an 'incomplete information' state (Thorne et al., 2018).

To elevate ‘evidence constraints’ from a matter of writing ethics to a structural constraint, this paper introduces two types of hard thresholds at the computable reasoning layer, which serve as common preconditions for the audit interface:

1) The evidence alignment threshold is denoted as τ_a. The alignment score is defined as follows: for any reasoning object r (which may be a factual assertion, mechanism edge, or causal edge), calculate the evidence alignment score $\mathrm{Align}(r) = \max_{e \in E(r)} \mathrm{sim}(r,e)$, where sim(r,e) measures the consistency between the ‘textual representation of the reasoning object’ and the ‘evidence fragment representation’ (achievable via string fingerprint matching, paragraph span overlap, or embedding similarity, though this study relies on SSOT field-based alignment records). This paper adopts $\tau_a=0.95$, meaning an inferred object qualifies for inclusion in ‘audit-able inferences’ only when $\mathrm{Align}(r) \geq 0.95$.

2) Coverage Threshold, denoted as τ_c. Coverage is defined as the minimum number of mutually independent evidence anchors required to support inference object r, denoted as $\mathrm{Cov}(r)=|E_{indep}(r)|$, where $E_{indep}(r)$ represents the set of distinct evidence anchors after deduplication (different page numbers/different paragraph spans are considered distinct). This paper adopts $\tau_c=2$, meaning $\mathrm{Cov}(r) \geq 2$ is required for an inference target to be rated ‘Pass’. When $\mathrm{Cov}(r)=1$, it may enter “Warning” status but cannot be upgraded to a strong conclusion; when $\mathrm{Cov}(r)=0$, it directly enters ‘Fail’.

Through the joint constraints of τ_a and τ_c, SEGN transforms ‘inferability’ into an enforceable rule: any mechanised insight must either be supported by evidence anchors or be downgraded to a structural gap and explicitly documented, without substituting linguistic fluency for evidence qualification (NIST, 2023; W3C, 2013b).

2.3 Audit Interface: Transforming Semantic Analysis into an ‘Audit-Ready Field System’

The audit interface layer structures analysed objects into fields, including: object type, statement text, evidence strength, evidence coverage, evidence alignment score, uncertainty annotation, falsifiability test, counterfactual design, and audit status. Its significance lies in: transforming ‘explainability’ from narrative to structured documentation; embedding “questionability” as fields to permit third-party review and rebuttal; and binding ‘model outputs’ to evidence and rules to mitigate generative hallucination risks (NIST, 2023).

This paper maps each reasoning object r (including factual assertions, mechanism edges, causal edges, and chained reasoning paths) to an audit status $\mathrm{Status}(r) \in \{\text{Pass, Warning, Fail}\}$, governed by the following rules:

Pass: Align(r) ≥ τ_a and Cov(r) ≥ τ_c, with no direct conflict records against evidence anchors on the same topic (conflict flag = 0).

Warning: Satisfies any one condition: 1) Align(r) ≥ τ_a but Cov(r) = 1; or 2) Align(r) ∈ [0.90, 0.95) and Cov(r) ≥ 1; or 3) Potential conflict clues exist but insufficient to determine contradiction (conflict flag = 1, requiring subsequent evidence increment adjudication).

Fail: Any one condition suffices: 1) Cov(r) = 0 (no evidence anchor); or 2) Align(r) < 0.90 (low alignment); or 3) Explicit mutual exclusion exists between evidence anchors (conflict flag = 2).

The primary purpose of this rule is not to ‘flag errors in reports’, but to establish thresholds for inference strength. This enables any insight to be explicitly traced back to its evidence credentials and uncertainty patterns, thereby elevating ‘semantic analysis’ to ‘computable auditing’ (NIST, 2023).

3. Findings: The report's 'highly centralised assertions' and institutionalised evidence chains (calculated based on SSOT + inference layer)

We computed the connectivity (degree/centrality) and evidence alignment strength (average weight of EVIDENCE_ALIGN) for factual assertion nodes within the SSOT graph. Results indicate that the most 'structurally central' assertions in the report are not generic narratives, but concentrate on three actionable governance levers: hybrid threat tactics (including sabotage and cyber incidents), international legal mechanisms (legal warfare), and classified information protection alongside cybersecurity compliance governance (SAB, 2026). The table below lists the top 10 claims with the highest connectivity (each with page number and evidence alignment score), forming the 'hard framework' for subsequent mechanism simulations.

Table 1 Top 10 High-Centrality Claims (SSOT Evidence Anchor Rebound)

Serial number	Proposals (Chinese paraphrase, all with cross-references to original page numbers)	Page number	Connectivity	Evidence alignment mean
1	No ceasefire agreement reached in Ukraine by 2025; Russia maintains an aggressive stance towards Latvia and the West; Levels of sabotage and cyber incidents remain high.	4	8	0.95
2	The Security and Intelligence Service (SAB), as Latvia's National Security Authority (NSA), is responsible for the oversight and protection of NATO and EU classified information within Latvia's territory.	30	7	1.00
3	The Ministry of Defence is funding a centralised DDoS protection service, which will be provided free of charge to public institutions. The service will be delivered by the Latvian National Broadcasting Centre (LVRTC).	27	7	1.00
4	Russia has announced its readiness to file applications against 17 countries with the International Court of Justice; it is highly likely that an application against Latvia will be submitted in 2026.	18	7	0.95
5	Cabinet Regulation No. 397 on Minimum Cybersecurity Requirements adopted on 25 June 2025, establishing minimum requirements for applicable entities under the National Cybersecurity Act.	28	7	0.95
6	By 2025, Russia will intensify its legal warfare against the West (particularly the Baltic states) and publicly signal its intent to bring cases before the International Court of Justice.	18	6	1.00

Serial number	Proposals (Chinese paraphrase, all with cross-references to original page numbers)	Page number	Connectivity	Evidence alignment mean
7	Russian allegations against Latvia centre on "violations of the rights of Russian-speaking residents", citing the International Convention on the Elimination of All Forms of Racial Discrimination.	19	6	1.00
8	Latvia's "State Secrets/Official Secrets" encompass information whose disclosure would harm national security, economic or political interests, and apply to NATO, EU and foreign classified information.	30	6	1.00
9	By 2025, SAB will advance security agreements with multiple nations (including Switzerland, Ukraine, Poland, North Macedonia, etc.), whilst planning to establish new or revise existing agreements with various parties.	34	6	1.00
10	SAB operates lawful mobile communications interception facilities, with data transferred to the requesting authority subject to a warrant issued by a judge of the Supreme Court.	35	6	1.00

The significance of these propositions lies not in their 'information density', but rather in their higher connectivity within the diagram structure: they establish a link between 'threat narratives' and 'governance levers', forming the shortest path for computable reasoning.

Table 2 Top 5 Mechanism Nodes (Evidence Density and Audit Status)

Serial number	Mechanism Node	Mechanism connectivity	Coverage	Align	Audit Status
1	Legal interception governance and oversight	High	≥2	≥0.95	Pass
2	Regulatory cybersecurity supervision and compliance	High	≥2	≥0.95	Pass
3	Centralized DDoS protection service provision	Mid-to-high	≥2	≥0.95	Pass
4	International legal mechanisms as a hybrid tool	Mid	1–2	≥0.95	Warning/Pass
5	Security agreements and classified info assurance	Mid	≥2	≥0.95	Pass

It must be emphasised that the Top-5 mechanisms do not equate to the ‘most frequently mentioned terms’, but rather represent ‘the mechanisms that most effectively link assertions, subjects, and actionable governance levers through evidence and structure’. This constitutes SEGN's key advantage

over topic models and co-occurrence networks: it transforms ‘what the report discusses’ into ‘how the report organises actionable governance through mechanisms’, with the inferential validity of this conclusion jointly constrained by τ_a, τ_c, and the audit status.

Table 3 Directional Causality Top-k (k=6, Evidence Binding and Audit Status)

Serial number	Cause → Effect	Page number anchor	Cov	Align	Audit Status
1	Persistent adversarial posture → High incidence of disruption and cyber incidents	4	≥2	≥0.95	Pass
2	Propaganda shapes threat perception → Deterioration in perceptions of the West/Baltic states	4	1–2	≥0.95	Warning/Pass
3	Preparation for legal warfare → Rising risks associated with the International Court of Justice route (window period)	18–19	≥2	≥0.95	Pass
4	Cybersecurity Legislation/Supporting Regulations → Implementation of Minimum Requirements for Critical Infrastructure	28	≥2	≥0.95	Pass
5	Supervisory Authorisation Mechanism (Supreme Court Warrant) → Enhanced Legal Interception Compliance Constraints	35	≥2	≥0.95	Pass
6	Advancement of security agreements → Expanded	34	≥2	≥0.95	Pass

	scope for mutual recognition of classified information and cooperation				

The significance of this Top-k list lies not in 'predicting certain occurrences', but in explicitly extracting the few causal entities permitted to lock direction within the report, and structuring their evidentiary qualifications: all subsequent stronger inferences must supplement evidence, conduct counterfactual tests, or maintain a Warning level around these causal fulcrums. They cannot bypass these evidence anchors for free extrapolation.

4. SEGN's Deep Insights into ‘Traditional Technologies' Inaccessibility’

4.1 From ‘Threat Description’ to ‘Mechanism Loop’: Structured Expression of Hybrid Pressure

Many analyses treat ‘hybrid threats’ as a generic concept, ultimately devolving into situational summaries. SEGN, however, condenses this into an auditable chain through mechanism edges and causal edges: Russian offensive pathways → elevated disruption/cyber incidents → governance demand for actionable centralised capabilities and minimum compliance thresholds. This chain is not a researcher's subjective collage, but is supported by high-frequency co-occurrence of report assertion nodes and governance nodes within the diagram, with evidence traceable back to corresponding page references (SAB, 2026). More crucially, SEGN permits the deconstruction of ‘hybrid means’ into verifiable mechanism components: for instance, positioning DDoS centralised protection services (free for public institutions, delivered by LVRTC) as a ‘capability supply mechanism node’, while treating ‘minimum cybersecurity requirements’ as an ‘institutional constraint mechanism node’. Together, these form a ‘capability-institutional dual-wheel’ governance structure. This approach is more policy-actionable than thematic modelling or keyword co-occurrence, as it directly addresses: ‘What are the levers? Who provides them? And through what institutional forms are they implemented?’

4.2 ‘Legal Warfare’ is not rhetoric: International judicial mechanisms structured as predictable risk pathways

Traditional semantic analysis typically only states: ‘The report mentions the International Court of Justice and refers to legal warfare.’ SEGN transforms this into computable risk pathways: claim nodes reveal Russia's 2025 intensification of legal warfare alongside public signals of ICJ recourse; further claims indicate Russia's preparation to file applications against 17 nations, with high probability targeting Latvia in 2026; causal edges link ‘legal warfare preparations’ to ‘potential litigation windows’ as auditable predictive chains. The crux of such conclusions lies not in whether predictions materialise, but in SEGN's formulation of predictions as falsifiable propositions: should no corresponding legal actions or official procedural advances occur by the 2026 window, this pathway should be retracted or downgraded. Such ‘falsifiable predictions’ align more closely with audit logic than general textual summaries, consistent with the ‘claim-evidence pairing’ emphasised in fact-checking research (Thorne et al., 2018).

4.3 ‘Compliance’ reduced to inter-organisational dependency: The institutional operability of NATO/EU classified information protection

Many readers interpret ‘protecting NATO/EU classified information’ as a principled stance. However, SEGN deconstructs it into an organisational dependency structure: the SAB, acting as the NSA, assumes oversight and protection responsibilities; classified information systems must meet NATO and EU security requirements and undergo periodic assessment visits; simultaneously, the legal definition of national secrets extends to NATO, EU, and foreign classified information, establishing consistent institutional boundaries (SAB, 2026). This reveals a profound conclusion that traditional semantic techniques struggle to articulate explicitly: external partnership status is encoded within the report's semantics as ‘verifiable compliance capability,’ rather than a declarative stance. In other words, the report anchors the credibility of national security cooperation in ‘audit-able systems and procedures,’ aligning with the ‘assessable credibility’ emphasised by traceability standards (W3C, 2013a).

To prevent claims of ‘methodological superiority’ from remaining purely declarative, this paper introduces a comparative baseline framework to characterise SEGN's incremental information: Baseline A comprises topic models (e.g., LDA/BERTopic-style topic clustering outputs), Baseline B involves keyword co-occurrence networks (identifying central terms/themes via co-occurrence strength), and Baseline C employs general information extraction/triple extraction (extracting entity-relationship-entity triples and counting frequencies). All three baselines prove effective at identifying report topics: they reliably yield high-frequency subjects such as ‘Russia, cybersecurity, international legal mechanisms, classified information’. However, they struggle to consistently achieve SEGN's three core capabilities: 1) They cannot distinguish between ‘implies mechanism’ and ‘causes’ as distinct inference levels; 2) They cannot establish threshold criteria for inferential objects, requiring them to satisfy τ_a and τ_c before upgrading to strong conclusions; 3) They cannot objectify structural gaps, thus failing to recognise ‘restraint in causal lock-in’ as an institutional signal and often misjudging it as ‘information deficiency’. Consequently, SEGN's irreplaceability lies not in ‘greater thematic accuracy’ but in ‘auditable inference qualification’: it transforms reports from readable narratives into sets of inferential objects amenable to evidence backtracking, threshold upgrading, and counterfactual testing, thereby mitigating semantic hallucination and narrative gap-filling risks (NIST, 2023; W3C, 2013b).

4.4 Irreplaceable Insights Driven by SEGN (Compared with Traditional Semantic Analysis)

From the perspective of traditional topic modelling, keyword co-occurrence networks, or general information extraction, this report most readily yields conclusions such as the co-occurrence of themes

like ‘Russia—Cybersecurity—International Legal Mechanisms—Protection of Classified Information’. However, such outputs typically remain confined to the level of ‘what is discussed’, failing to simultaneously address within the same framework whether inferences can be drawn, the strength of such inferences, and whether they can be backed up by evidence or refuted. SEGN's incremental insights derive from three structural constraints: Firstly, the Single Source of Truth (SSOT) binds page numbers and evidence anchors to each assertion, mechanism, and causal object, mandating that conclusions must be traceable back to the original text; Second, the computable reasoning layer categorises ‘concept-mechanism-causality’ as distinct types of edges, applying evidence alignment thresholds τ_a and coverage thresholds τ_c to gate inference eligibility, preventing ‘textual similarity’ from automatically equating to ‘mechanism validity’; Thirdly, the audit interface codifies Pass/Warning/Fail status rules into fields, institutionalising the recording of uncertainties and conflict clues to prevent generative summarisation from filling narrative gaps where evidence is insufficient.

Consequently, SEGN reliably uncovers ‘implicit governance structures’ beyond traditional models' reach: for instance, reports on international judicial pathways do not merely cite courts and conventions, but form auditable risk chains when evidence alignment and coverage thresholds are met, organising ‘signal release—caliber templates—time windows’ into falsifiable propositions; Similarly, the cyber governance section does not present two disparate elements (centralised DDoS protection and minimum cybersecurity requirements) in parallel. Instead, they converge structurally and in evidentiary density towards a dual-loop mechanism of ‘centralised capability provision + enforced regulatory thresholds,’ elevating policy leverage from thematic to systemic levels. Furthermore, SEGN deconstructs confidential information protection and security agreements into cross-organisational dependency structures, supporting the auditable inference that ‘compliance capability constitutes a threshold for partnership eligibility’. More significantly, when the report deliberately avoids strong causal lock-in within the boundaries of public disclosure and confidentiality, SEGN does not treat these gaps as spaces for free narrative supplementation. Instead, it objectifies them as structural lacunae and maintains a Warning level, demanding either supplementary evidence or counterfactual testing. This identifies ‘causal restraint’ as an institutional signal rather than informational noise.

SEGN's distinctive capability lies not merely in semantic comprehension, but in transforming ‘conclusiveness’ itself into computable rules, converting ‘conclusion strength’ into auditable states, and rendering ‘hidden meanings’ into mechanisms that are jumpable, threshold-upgradable, and falsifiable chained objects. Traditional models, even when offering interpretations of a similar nature, often struggle to justify ‘why such inferences are valid,’ fail to achieve stable reproducibility, and

cannot withstand scrutiny in peer review or audit contexts against objections that ‘this constitutes subjective interpretation.’

5. Discussion: How SEGN Reduces ‘Semantic Hallucinations’ and Enhances Policy Usability

A common risk in generative semantic analysis is ‘mistaking linguistic fluency for explanatory strength’. SEGN mitigates this risk through three layers of digital products: 1) SSOT enforced backtracking evidence: any conclusion must include page numbers and evidence fingerprints, otherwise it can only be labelled as misaligned or low-strength; 2) Computable reasoning layer constrains inference space: only paths permitted by mechanism edges and causal edges can propagate, reducing ‘free association’; 3) Audit interfaces institutionalise uncertainty: incorporating uncertainty annotations and counterfactual tests into fields transforms them from ‘writing conventions’ into ‘systemic requirements’. This aligns with NIST's emphasis on trustworthy AI governance: risk management demands transparency, traceability, and assessability—not merely performance metrics or narratives (NIST, 2023).

Moreover, SEGN does not negate the value of traditional methods but imposes ‘upstream constraints’: topic models may be used for preliminary screening of themes but cannot directly generate mechanism conclusions; information extraction may capture entity relationships but must enter the SSOT before participating in reasoning; opinion detection may assist in identifying stances but lacks audit eligibility without evidence alignment and counterfactual testing.

Within national security text scenarios, ‘uncertainty’ does not equate to ‘defect’. Public reports must often balance disclosure with confidentiality, whilst overly stringent causal locking risks strategic exposure. SEGN's strength lies precisely in translating this equilibrium into structure: when evidence coverage is insufficient, the reasoning object automatically remains at the Warning tier, demanding supplementary evidence or enhanced counterfactual testing. When reports adopt a mechanism-first, causally restrained narrative strategy, SEGN explicitly flags such structural gaps as institutional signals, rather than allowing researchers to freely fill in the blanks.

Furthermore, audit status rules equip researchers and readers with a shared ‘language of intensity’: Pass objects may be used for policy simulation and cross-year comparisons; Warning objects may inform early warnings and hypothesis generation but must not be elevated to definitive conclusions; Fail objects must be recycled into an ‘evidence gap register’. This threshold-based intensity system endows semantic analysis outputs with genuine decision-usefulness: not merely ‘written like intelligence’, but ‘capable of audit and review’.

6. Research Conclusions

This paper presents a SEGN-driven semantic audit framework for national security reports, using Latvia's SAB Annual Report 2025 as a case study. The framework employs SSOT non-destructive graphing to solidify the single source of truth, utilises a computable reasoning layer to structure the concept-mechanism-causality chain, and employs audit interfaces to field evidence strength, coverage, and falsifiability. Employing this framework, we not only recount ‘what the report states’ but also address ‘how the report transforms threat narratives into governance levers through mechanisms’. Three critical mechanism chains—legal warfare, cybersecurity governance, and classified information compliance—are presented in an auditable structure. Compared to analyses relying solely on textual similarity or thematic summarisation, SEGN's core advantages lie in its: replayability, verifiability, falsifiability, and transferability. This enables it to function as a standardised digital product, extendable to mechanistic audit scenarios for other national security public texts.

More significantly, this paper incorporates ‘irreplaceable insights’ into the audit threshold: through threshold criteria ($\tau_a=0.95$, $\tau_c=2$), audit status rules (Pass/Warning/Fail), and tables of Top-5 mechanisms and Top-k causal objects, it elevates deep insights from interpretative text to verifiable object sets. Thus, SEGN functions not merely as a semantic analysis framework but as a standardised data product methodology capable of transforming national security disclosure texts into ‘computable reasoning + auditable interfaces’. This provides a unified evidence standard and inference strength language for cross-year incremental updates and transnational comparative research.

Through systematic analysis of Latvia's Constitution Protection Bureau 2025 Annual Report using the Structural Empty Graph Network (SEGN) semantic audit framework, this paper confirms—without external intelligence sources—that national security risk perception is shifting from event-driven to institutional and regulatory dimensions. Specifically, the report's coverage of cybersecurity, international judicial pathways, and classified information protection does not present isolated issues but collectively points towards a long-term risk expectation of ‘rules being strategically exploited.’ Within this context, compliance capability is implicitly shaped as an access threshold for international security cooperation, while the restrained treatment of strong causal relationships constitutes an institutionalised risk management strategy under public disclosure conditions. The foregoing assessments are grounded in constraints pertaining to evidence alignment, coverage thresholds, and audit status protocols, thereby qualifying as inferences that are reversible, contestable, and transferable.

From an intelligence analysis perspective, the most critical information identified by SEGN in the Latvian SAB report—which traditional methods struggle to consistently capture—can be highly condensed into the following points (all constituting ‘auditable inferences’ rather than subjective interpretations):

First, Latvian security agencies have shifted their threat perception from ‘concrete operational risks’ to ‘risks of institutional and regulatory exploitation.’ SEGN reveals that the report's true focus lies not in isolated cyberattacks or intelligence infiltration incidents, but rather in adversaries' systematic legal warfare tactics—constructing a ‘legitimate façade’ through international law, judicial procedures, and compliance discourse. This represents a strategic risk characterised by long-term, low-visibility yet sustainable pressure.

Secondly, cybersecurity is explicitly positioned as a ‘testing ground for state governance capabilities,’ rather than a purely technical issue. By aligning mechanisms for centralised DDoS protection with minimum cybersecurity requirements, SEGN demonstrates that the SAB now regards cyberspace as a core domain for assessing public sector execution, cross-departmental coordination, and enforcement compliance. Its primary focus lies in ‘whether institutionalised defence capabilities exist,’ rather than ‘whether individual attacks can be countered.’

Thirdly, compliance capability is implicitly equated with eligibility for international security cooperation. The cross-organisational dependency structure reconstructed by SEGN indicates that the report's references to classified information protection and security agreements essentially signal to allies and partners: only entities possessing stable, auditable compliance and confidentiality capabilities are deemed credible intelligence and security partners.

Fourth, the deliberate avoidance of strong causality is not due to insufficient information but constitutes an institutionalised risk management strategy. SEGN identifies multiple instances of ‘causal restraint’ as structural gaps, revealing that the SAB consciously avoids pinpointing specific causal links and attributing responsibility in its public texts. This approach enables risk warning and expectation management without escalating political and diplomatic risks.

SEGN reveals what Latvia's security services have truly "seen but not fully articulated"—threats are migrating from the incident level to the institutional level, shifting from technical confrontation to a systemic contest over rules, compliance, and governance capabilities.

References

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf

National Institute of Standards and Technology. (n.d.). AI Risk Management Framework (AI RMF). https://www.nist.gov/itl/ai-risk-management-framework

Satversmes aizsardzības birojs. (2026). Annual report 2025 (English). Constitution Protection Bureau of the Republic of Latvia. https://www.sab.gov.lv/files/uploads/2026/01/SABs-annual-report_2025_ENG.pdf

Satversmes aizsardzības birojs. (n.d.). Annual reports. https://www.sab.gov.lv/en/annual-reports/

Satversmes aizsardzības birojs. (2026). The Constitution Protection Bureau publishes the annual report. https://www.sab.gov.lv/en/news/russia-continues-to-use-and-develop-new-hybrid-instruments-to-influence-latvia-constitution-protection-bureau-states-publishing-its-annual-report/

Thorne, J., Vlachos, A., Christodoulopoulos, C., & Mittal, A. (2018). FEVER: A large-scale dataset for fact extraction and verification. arXiv. https://arxiv.org/abs/1803.05355

Churina, S., et al. (2024). Improving evidence retrieval on claim verification. ACL Anthology. https://aclanthology.org/2024.fever-1.6.pdf

Liang, X. (2024). A survey of LLM-augmented knowledge graph construction. Expert Systems with Applications. https://www.sciencedirect.com/science/article/pii/S2212827124007911

ACM. (2023). A comprehensive survey on automatic knowledge graph construction. ACM Computing Surveys. https://dl.acm.org/doi/10.1145/3618295

W3C. (2013a). PROV-Overview: An overview of the PROV family of documents. https://www.w3.org/TR/prov-overview/

W3C. (2013b). PROV-O: The PROV Ontology. https://www.w3.org/TR/prov-o/

Soleimani, A., et al. (2020). BERT for evidence retrieval and claim verification. Lecture Notes in Computer Science. https://pmc.ncbi.nlm.nih.gov/articles/PMC7148011/