Check for updates

RESEARCH ARTICLE

Automated regulatory compliance for AI systems in the security domain: The case of dual-use deployment

[version 1; peer review: 1 approved, 3 approved with reservations]

Giedre Sabaliauskaite[1], R. Andrew Paskauskas[2], Evaldas Bružė[2], Raminta Matulytė[2], Tomas Lavišius[2]

[1]Faculty of Public Governance and Business, Mykolas Romeris University, Vilnius, 08303, Lithuania
[2]Security Research Laboratory, Mykolas Romeris University, Vilnius, 08303, Lithuania

v1 **First published:** 28 Mar 2026, **6**:83
https://doi.org/10.12688/openreseurope.23137.1
Latest published: 28 Mar 2026, **6**:83
https://doi.org/10.12688/openreseurope.23137.1

Abstract

The rapid expansion of the European Union's digital regulatory framework and recent adoption of the Artificial Intelligence Act (AI Act) has introduced complex, overlapping compliance obligations for AI systems, and consequently, their developers and users. These challenges are amplified for dual-use AI systems, which may be developed for civilian markets yet deployed in military contexts, requiring alignment with both EU regulatory instruments and defence-specific frameworks such as NATO policies. Existing approaches to regulatory compliance remain largely manual, fragmented, and difficult to scale, particularly when legal requirements must be translated into actionable, system-level specifications for dual-use contexts.

This paper proposes a novel ontology-based methodology for automated regulatory compliance requirements specification for dual-use AI systems. The methodology systematically integrates legal and technical perspectives by structuring compliance obligations across deployment domains (civilian and defence), system lifecycle phases, and requirement categories, including privacy- and security-related obligations. Implemented as a machine-readable knowledge graph using RDF/Turtle, the approach enables executable compliance modelling, where regulatory obligations are formalised as machine-interpretable entities that can be queried, validated, and deployment-specifically configured through semantic reasoning and SPARQL-based analysis.

The methodology is validated through a detailed case study of an AI-powered espionage detection system, demonstrating how context-

Open Peer Review

Approval Status ? ? ✓ ?

	1	2	3	4
version 1 28 Mar 2026	? view	? view	✓ view	? view

1. **Hewa Majeed Zangana**, Duhok Technical College, Duhok Polytechnic University, Duhok, Iraq
2. **Stefka Schmid**, Aalto University, Helsinki, Finland
3. **Florentina-Loredana Dragomir-Constantin**, National Defense University Carol I, Bucharest, Romania
4. **Aryendra Dalal**, Middle Georgia State University, Macon, USA

Any reports and responses or comments on the article can be found at the end of the article.

aware semantic reasoning and SHACL-based validation can ensure that regulatory requirements are consistently specified and mapped to concrete system components. The proposed framework advances the state of the art by providing a rigorous, extensible foundation for compliance-by-design and automated analysis, thereby reducing compliance risk and supporting responsible AI engineering in complex omni-use regulatory environments.

Keywords
Regulatory compliance automation; dual-use technologies; ontology-based requirements specification; security-by-design; privacy-by-design; AI Act; GDPR; NATO cybersecurity framework

This article is included in the Horizon 2020 gateway.

This article is included in the Artificial Intelligence and the Social Sciences and Humanities collection.

Corresponding author: Giedre Sabaliauskaite (giedre.s@mruni.eu)

Author roles: Sabaliauskaite G: Conceptualization, Formal Analysis, Investigation, Methodology, Supervision, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Paskauskas RA**: Conceptualization, Data Curation, Formal Analysis, Methodology, Resources, Software, Validation, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Bružė E**: Formal Analysis, Funding Acquisition, Investigation, Project Administration, Resources, Writing – Review & Editing; **Matulytė R**: Formal Analysis, Investigation, Writing – Review & Editing; **Lavišius T**: Formal Analysis, Investigation, Writing – Review & Editing

Competing interests: No competing interests were disclosed.

Grant information: This research was prepared as part of the project "Implementation of Mission-Based Science and Innovation Programs" (Project No. 02-002-P-0001). The project is financed by the European Union through the Recovery and Resilience Facility (2021–2027 programming period, New Generation Lithuania plan), and co-funded by the state budget of the Republic of Lithuania administered via the Research Council of Lithuania (Lietuvos Mokslo Taryba). These combined mechanisms ensure alignment with both EU-level resilience and recovery objectives and Lithuania's national science and innovation priorities. The funding received for this project is restricted to supporting research and writing activities; it does not cover publication or printing fees.
The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Sabaliauskaite G *et al*. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Sabaliauskaite G, Paskauskas RA, Bružė E *et al.* **Automated regulatory compliance for AI systems in the security domain: The case of dual-use deployment [version 1; peer review: 1 approved, 3 approved with reservations]** Open Research Europe 2026, **6**:83 https://doi.org/10.12688/openreseurope.23137.1

First published: 28 Mar 2026, **6**:83 https://doi.org/10.12688/openreseurope.23137.1

1. Introduction

Artificial intelligence (AI) systems are increasingly embedded in socio-technical infrastructures across sectors such as transport, healthcare, public administration, finance, and defence. As they are becoming more autonomous, adaptive, and scalable, there is an urgent need to establish the harmonised rules for their development, placing on the market, and use.

In June 2024, the European Union (EU) has adopted the Artificial Intelligence Act (AI Act) – the world's first comprehensive regulatory framework for AI (European Commission, 2024a). The AI Act forms part of a broader and rapidly expanding digital regulatory ecosystem, that includes General Data Protection Regulation (GDPR), the Data Act, the Digital Services Act (DSA), the Digital Markets Act (DMA), the Cyber Resilience Act (CRA), and the NIS2 Directive among other regulations, aimed at ensuring trust, safety, and competitiveness of Europe's digital economy (European Parliament, 2025). Together, these regulations impose complex, overlapping, and evolving compliance obligations on developers, providers, and deployers of AI systems.

The AI Act defines technology-neutral legal requirements related to risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness, and cybersecurity. In addition, to operationalize their implementation, technical standards are being developed by the CEN-CENELEC JTC 21 (AI), a joint committee by the European Committee for Standardization (CEN) and the European Committee for Electrotechnical Standardization (CENELEC). The Committee is planning to develop circa 35 technical standards to support the AI Act (Kilian et al., 2025).

The AI standardization in EU faces numerous challenges, such as the overwhelming volume and fragmentation of relevant standards, high compliance costs, and limited participation of start-ups and Small and Medium-sized Enterprises (SMEs) in standardisation processes (Kilian et al., 2025). AI system developers must navigate the cumulative and sometimes ambiguous interaction between multiple digital laws, where concepts such as risk, accountability, transparency, and security are defined differently across legal instruments. This creates substantial interpretative burdens and increases the risk of inconsistent or incomplete compliance, particularly where requirements must be translated into technical and organisational system specifications.

The growing complexity, dynamism, and interdependence of EU digital regulation call for more systematic, scalable, and adaptive approaches to compliance. Automated regulatory compliance requirements specification, which uses computational methods to identify, structure, formalise, and update regulatory obligations applicable to AI systems, offers a promising pathway to address these challenges. By enabling machine-readable representations of legal and standardisation requirements, automation can support early compliance-by-design, reduce ambiguity, improve traceability between legal norms and technical controls, and lower barriers to compliance for resource-constrained actors, such as SMEs.

The importance of automated compliance specification is further amplified in the context of dual-use AI systems, where systems developed for civilian purposes may also be deployed in defence or security contexts. In this article, dual-use systems are the systems, which can include software and hardware, and which can be used for both civil and military purposes.

Recent scholarship (Kremidas-Courtney, 2025) suggests that the traditional binary classification of dual-use is increasingly inadequate, advocating instead for an "omni-use by design" philosophy where AI systems operate simultaneously across civilian and security domains. AI systems used for surveillance, target recognition, or decision support often fall at the intersection of civilian AI regulation, sector-specific safety standards, and defence-specific legal frameworks. From an engineering standpoint, this creates "additive" complexity: while military applications may benefit from partial exemptions under the AI Act, dual-use systems often remain subject to a cumulative burden where defence standards (e.g., NIAP, NSP) must layer on top of a civilian GDPR and NIS2 baseline (Cichocki, 2025; Ndiaye et al., 2026).

This paper investigates how automated approaches to regulatory compliance requirements specification can support compliance with the EU AI Act and the wider EU digital regulatory framework for AI systems. Building on insights from European AI standardisation and the interplay between the AI Act and other digital legislation, the paper argues that automation is not merely a technical optimisation, but a necessary governance mechanism to operationalise abstract legal norms within complex AI systems.

Moreover, this paper proposes a novel ontology-based six-step methodology which provides a systematic approach to specifying regulatory compliance requirements for dual-use AI systems. The methodology systematically integrates legal and technical perspectives by structuring compliance obligations across deployment domains, system lifecycle phases, and requirement categories. Unlike prior ontology-based compliance models limited to single regulatory frameworks,

this study introduces a cross-domain executable compliance architecture capable of representing layered EU and defence governance structures within a unified semantic model.

Proposed methodology leverages semantic reasoners and formal constraint languages to address the "interpretative burden" of 2025 regulatory guidelines, providing a structured foundation for real-time compliance automation in high-risk environments like 5G infrastructure and maritime cybersecurity (Khan et al., 2025; Paskauskas, 2025). Furthermore, a case study on an espionage detection system is included to validate the proposed methodology through the integration of ontology-driven privacy obfuscation and re-identifiability metrics (Topalli & Badii, 2025).

The remainder of this paper is structured as follows. Section 2 provides the background information on AI systems, including their development lifecycle, regulatory landscape, current practices and related work in the area of automating regulatory compliance. Section 3 presents a novel methodology for dual-use AI system regulatory compliance requirement specification. Further details on how the regulatory compliance requirements are specified are provided in Section 4. A case study is included in Section 5, aimed at validating the proposed methodology. Finally, section 6 summarizes the paper, discusses its major contributions and limitations, identifies directions for future work, and concludes the paper.

2. Background and related work

2.1. AI system and its development lifecycle

European Commission Guidelines define AI system as a "machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments" (European Commission, 2025a).

This definition comprises seven main elements (European Commission, 2025a):

1) A *machine-based system*, that includes both the hardware and software components that enable the AI system to function.

2) *Autonomy*: AI system is designed to operate with 'some degree of independence of actions from human involvement and of capabilities to operate without human intervention'.

3) *Adaptiveness after deployment*: self-learning capabilities, allowing the behaviour of the system to change while in use. The new behaviour of the adapted system may produce different results from the previous system for the same inputs.

4) *Objectives*: AI system is designed to operate according to one or more objectives, that could be explicit or implicit. Explicit objectives refer to clearly stated goals that are directly encoded by the developer into the system. Implicit objectives refer to goals that are not explicitly stated but may be deduced from the behaviour or underlying assumptions of the system. These objectives may arise from the training data or from the interaction of the AI system with its environment.

5) *Output generation*: Ability of a system to derive outputs from given inputs.

6) *Outputs that can influence physical or virtual environments*: The ability of a system, to generate outputs, such as predictions, content, and recommendations, based on inputs it receives and using machine learning and logic and knowledge-based approaches, is fundamental to what AI systems do and what distinguishes those systems from other forms of software.

7) *Interaction with the environment*: AI systems are not passive – they actively impact the environments in which they are deployed. The influence of an AI system may be both to tangible, physical objects (e.g., autonomous vehicle) and to virtual environments, including digital spaces, data flows, and software ecosystems.

Similar to other machine-based systems, the development of an AI system follows a structured lifecycle that spans from initial conception to eventual decommissioning and encompasses multiple distinct stages. AI system development lifecycle, defined by ISO/IEC ISO/IEC 22989 standard (ISO/IEC, 2022), comprises of several phases, including (1) Inception, (2) Design and development, (3) Verification and validation, (4) Deployment, (5) Operation and monitoring, (6) Re-evaluation, and (7) Retirement. Some of these stages can be repeated throughout lifecycle. E.g.,

Design and development can be repeated to implement updates and required changes to AI system, followed by Verification and validation and Deployment stages.

The lifecycle stages are often grouped into two main phase – pre-deployment or 'building' phase of the system (stages 1–3), and post-deployment or 'use' phase (stages 4–7) (European Commission, 2025a).

AI systems are frequently integrated into broader software ecosystems that follow comparable development lifecycles. These typically encompass the phases of requirements analysis, system design, development and implementation, testing and acceptance, deployment and integration, as well as ongoing maintenance and eventual decommissioning (ENISA, 2019).

2.2. Regulatory landscape

Developing AI systems within the EU requires the integration of legal considerations from the earliest stages of the technology lifecycle. Achieving both operational and regulatory compliance necessitates a comprehensive understanding of EU-level legislation, national legal frameworks, sector-specific requirements, and applicable international standards. This section includes a non-exhaustive overview of the illustrative regulatory frameworks that can affect the development of dual-use AI systems.

At the EU level, for many years, global discussions on AI regulation lacked clear direction. The EU finally addressed this gap in 2024 by adopting the first comprehensive AI regulation globally—the AI Act (European Commission, 2024a). The Act imposes obligations on AI providers and deployers, categorized by risk level: (1) unacceptable risk; (2) high risk; (3) limited risk; and (4) minimal risk (European Commission, n.d.). However, its scope excludes AI systems used exclusively for military, defence, or national security purposes (Art. 2(3)). The AI Act foresees a number of requirements that systems and organisations developing or deploying them need to comply with – for example, for high-risk systems, establishing a human in the loop factor, conducting a fundamental rights impact assessment and others.

Another EU regulation, GDPR (European Commission, 2016a), applies whenever an AI system processes personal data, wholly or partly by automated means (Art. 2(1)), which effectively includes most AI systems. GDPR establishes a comprehensive framework for personal data processing, covering principles such as data minimisation and storage limitation, legal bases, and other operational requirements. For example, under the GDPR, a Data Protection Impact Assessment (DPIA) must be carried out for any new project that is likely to pose a high risk to individuals' personal data. Similar to the AI Act, GDPR excludes certain processing activities from its scope—for example, processing necessary for common foreign and security policy (Art. 2(2)(b)) or by competent authorities for law enforcement purposes (Art. 2(2)(d)). In the latter case, for law enforcement authorities processing data for law enforcement purposes, the Law Enforcement Directive (LED) applies (European Commission, 2016b). As related to the personal data processed for national security purposes, member states usually establish their own rules. For example, Lithuania has adopted the Law of the Republic of Lithuania on the Legal Protection of Personal Data Processed for the Purposes of the Prevention, Investigation, Detection or Prosecution of Criminal Offences, the Execution of Sentences or National Security or Defence (Lietuvos Respublikos Seimas, 2011) which implements the LED as well as sets additional rules in relation to data processing for the purposes of national security and defence.

To strengthen EU cyber resilience, a comprehensive cybersecurity framework has been introduced – the NIS2 Directive. It sets specific management and technical requirements for critical infrastructure organizations (entities) operating in designated sectors (European Commission, 2022a). Although AI system providers are not always classified as cyber-security subjects, supply chain obligations may require compliance with NIS2. Additionally, the CRA may apply to certain systems (European Commission, 2024b). Article 2(1) states that the CRA covers "products with digital elements made available on the market, the intended purpose or reasonably foreseeable use of which includes a direct or indirect logical or physical data connection to a device or network." However, Article 2(7) exempts products developed or modified exclusively for national security or defence purposes, or those designed to process classified information.

A key EU-level regulation governing technologies with both civilian and defence applications is the Dual Use Regulation (European Commission, 2021a). It defines "dual-use items" as goods, software, and technology that can serve both civil and military purposes (Art. 2(1)). The Regulation establishes a Union-wide framework for controlling exports, brokering, technical assistance, transit, and transfer of such items.

In addition to EU-level regulations, national legislation can significantly influence AI system development and deployment. For instance, NIS2, as a directive, requires transposition into national law, such as Lithuania's Cybersecurity Law (Lietuvos Respublikos Seimas, 2024). Other general laws may also affect compliance. While GDPR provides a

comprehensive EU-wide data protection framework, Member States often adopt supplementary national data protection laws introducing additional requirements beyond GDPR.

Depending on the way that AI system is trained and further used, intellectual property considerations may be relevant. OECD, for example, has issued the report on intellectual property issues where AI systems are trained on scraped data (OECD, 2025). While the report covers many jurisdictions, it emphasises that in relation to the EU, such matters are essentially governed by the Directive on Copyright and Related rights in the Digital Single Market (European Commission, 2019) and then further to a certain extent repeated in the AI Act.

In addition, sector-specific regulations may apply. For example, the financial sector is mostly exempt from NIS2 as it is governed by the Digital Operational Resilience Act (DORA), the primary cybersecurity framework for financial entities (European Commission, 2022c). DORA imposes strict cybersecurity requirements on ICT service providers. Consequently, if an AI provider qualifies as an ICT service provider under DORA, its systems must comply with these requirements. Other sectoral regulations—such as those related to healthcare, medical devices, and employment—may similarly impact AI systems.

As the EU lacks exclusive competence in defence regulation, frameworks governing AI systems in this domain are established at the national level. General principles may derive from international public law. Military alliances such as NATO issue guidance and recommendations rather than binding requirements, which are regarded as best practice standards. However, while certain guidelines are important but rather more of a recommended standard in the defence sector, the applicable international humanitarian law (IHL) standards still bind the development and applicability of AI technologies. The IHL legal framework today consists mainly of the four Geneva Conventions of 1949, their additional protocols (ICRC, 2026a) and customary law (ICRC, 2026b) and lays down which actions can be taken during the warfare, including when AI is used during armed conflict (Cuypers, 2023).

Defence deployments do not automatically displace EU legal obligations. Where a system is marketed or made available within the EU, EU regulatory requirements apply to that placement. Where the same system is deployed within NATO infrastructures, additional accreditation and security control requirements arise. The resulting compliance structure is cumulative rather than substitutive: EU market obligations operate alongside defence accreditation requirements, subject to applicable exemptions under EU law.

The International Committee of Red Cross (ICRC) has highlighted the following areas in which AI is being developed for use in warfare, which, among others, can raise significant legal questions: (a) integration in weapon systems, particularly autonomous weapon systems, (b) use in cyber and information operations, (c) underpinning military 'decision support systems' (ICRC, 2023). According to the ICRC, international humanitarian law requires that humans, not machines remain in control of AI used in warfare. While AI may support military operations, human combatants should retain responsibility for the actions made in warfare (ICRC, 2021). In other words, where AI is used in warfare, its design must, from the beginning, ensure the compliance with IHL obligations.

Finally, the development and deployment of AI, including for military purposes, should always consider the protection of fundamental rights which are primarily established in the European Convention on Human Rights (European Court of Human Rights, 1950), Charter of Fundamental Rights of the EU (European Commission, 2012) and national constitutions.

Today, developing software and AI systems requires navigating an increasingly intricate regulatory landscape. Prior to initiating development, it is essential to clearly define: (a) the intended end user, (b) the categories of data the system will process, (c) the system's primary purpose, and (d) the geographical location in which the system will be developed and deployed.

Based on these parameters, a comprehensive compliance matrix for AI systems can be developed, enabling the identification and resolution of potential regulatory conflicts. For instance, if an AI system is intended for defence use only, many civilian regulatory requirements may not apply. Conversely, for civilian AI systems, the jurisdiction where the AI system is manufactured and (or) placed determines the applicability of specific international, EU-level and national legislation. When personal data is processed, it is essential to assess both the developer and the end-user to determine whether the GDPR, the Law Enforcement Directive or other pieces of data protection legislation apply.

The compliance matrix becomes increasingly complex as additional contextual and operational factors are introduced. It must be noted, therefore, that in the legal scholarship, it is emerging that real applicability of the exemptions included

under the certain EU regulations such as the GDPR or AI Act is uncommon in practice, one of the reasons being that governments and militaries often tend to rely on private companies to develop AI. As a result, it is argued that dual-use AI systems may be covered by GDPR/LED (as well as the AI Act for non-military use case scenarios), and only a small part of in-house defence/national security-related processing is purely exempt from the said regulations (Powell, 2024; Vogiatzoglou, 2024; Juric, 2024).

2.3. AI Act risk classification in dual-use contexts

The regulatory classification of AI systems under the AI Act depends primarily on their intended purpose and deployment context (European Commission, 2024a). For dual-use AI systems, classification requires careful distinction between civilian market placement and defence-only application.

Under Article 6 of the AI Act, an AI system is classified as high-risk where it is either:

(a) intended to be used as a safety component of a product subject to third-party conformity assessment under Union harmonisation legislation, or

(b) listed in Annex III.

Annex III includes, inter alia, AI systems used in:

- biometric identification and categorisation of natural persons,
- management and operation of critical infrastructure,
- law enforcement,
- migration and border control,
- access to essential private and public services.

An espionage detection system deployed in civilian corporate or public-sector environments would require analysis under Annex III, particularly if it performs biometric identification or is used within critical infrastructure contexts. However, a system that detects unauthorised photography attempts without performing identity recognition or categorisation may not automatically fall within the biometric identification category. The classification therefore depends on functional scope rather than technological capability alone.

In any case, it is necessary to assess whether the designed system does not automatically fall under the list of prohibited AI practices under the AI Act. The AI Act prohibits certain practices under Article 5, including specific forms of remote biometric identification in publicly accessible spaces (subject to limited exceptions). Dual-use systems must therefore ensure that civilian variants do not implement prohibited functionality.

Article 2(3) excludes AI systems developed or used exclusively for military, defence, or national security purposes. This exemption is narrow. Where a system is placed on the EU market for civilian use, even if technically capable of defence deployment, the Regulation applies to that civilian deployment. Providers of dual-use AI systems must therefore design compliance mechanisms that accommodate both regulated civilian contexts and defence accreditation environments.

In addition to high-risk classification, providers of high-risk AI systems must comply with obligations including:

- risk management systems (Art. 9),
- data governance measures (Art. 10),
- technical documentation (Art. 11),
- logging capabilities (Art. 12),
- transparency and provision of information to users (Art. 13),

- human oversight (Art. 14),
- accuracy, robustness, and cybersecurity requirements (Art. 15).

These obligations map naturally onto lifecycle phases addressed in the proposed methodology. Risk management and data governance correspond primarily to the design and development phase, while logging, oversight, and monitoring obligations extend into deployment and operational phases.

For dual-use AI systems, the AI Act therefore operates as the civilian compliance baseline. Defence deployments may fall outside the Regulation where exclusively military; however, when personal data is processed, GDPR obligations may continue to apply subject to applicable exemptions. The additive nature of compliance in dual-use systems reinforces the need for structured, lifecycle-based modelling capable of distinguishing domain-specific obligations while preserving traceability.

Companies developing AI systems with potential dual-use applications (e.g., drones, advanced materials software, certain sensors) should consider the AI Act's compliance requirements from the design phase, particularly if civilian market access is desired (RAND, 2024). In addition, it may be necessary to clearly separate the military and civilian use cases and ensure that the civilian variants meet all relevant AI Act and GDPR requirements, which may demand separate documentation and conformity assessments.

2.4. Current practices in regulatory compliance

Although the above regulations are mandatory, the evolving market and regulatory environment increasingly demands clearer implementation guidance. This section describes recent practices and supporting measures for achieving regulatory compliance.

Technologies subject to the AI Act face several regulatory challenges stemming from ambiguities in the legislation. To support effective implementation, the EU and NATO have introduced policy measures that clarify requirements, refine definitions, and provide guidance essential for developing and deploying relevant technological solutions.

Furthermore, EU regulation, such as CRA, the NIS2 Directive, and the EU AI Act, along with sector-specific safety and cyber security standards, such as IEC 61508 (IEC, 2010) and ISA/IEC 62443 (IEC, 2018), use risk-based approaches. However, they differ considerably in scope, language, and application level, resulting in regulatory and technical gaps between horizontal legislation and sector-specific standards (Ndiaye et al., 2026).

In this study, an additional challenge arises when the technologies under development fall within a dual-use domain, requiring compliance not only for civilian (EU and national level) but also for defence and security (NATO and national level) contexts, where compliance obligations arise through accreditation and governance mechanisms rather than supranational legislation. Within the NATO framework, AI-enabled systems deployed in Alliance environments are subject to security accreditation requirements defined in instruments such as the NATO Information Assurance Policy (NIAP), the NATO Security Policy (NSP), and the NATO Cyber Defence Policy.

These instruments function as binding governance frameworks within the NATO, implemented through procurement conditions, security clearance regimes, certification processes, and operational directives. While not legislative acts comparable to EU regulations, they establish mandatory requirements for systems intended for deployment within NATO environments, including controls related to encryption, access management, incident reporting, secure data exchange, and personnel vetting.

Compliance in defence deployments is, therefore, mainly achieved through accreditation processes and security authorisation rather than conformity assessment under EU market legislation. In a simplified way, for dual-use systems, this creates a layered compliance environment: EU regulatory obligations apply where the system is placed on the EU market or processes personal data within EU jurisdiction, while NATO accreditation requirements apply where the system is deployed within Alliance-controlled infrastructures.

This distinction is important for methodological modelling. The ontology must represent both legally binding EU instruments and defence accreditation frameworks while preserving their different normative statuses.

AI is inherently dual-use, and NATO, therefore, places significant emphasis on its governance through key policy and implementation documents that regulate its application in the defence domain while encouraging collaboration with the civilian sector to support innovation (NATO's Strategic Warfare Development Command, n.d.). Several high-level documents outline the core principles and strategic direction guiding NATO's use of AI.

The NATO AI Strategy, which was adopted in 2021 and revised in 2024, emphasizes the need for reliable implementation of AI in NATO structures, while adhering to key principles: Lawfulness, Responsibility and Accountability, Explainability and Traceability, Reliability, Governability, and Bias Mitigation. This means that any AI technologies implemented in this domain must satisfy accreditation criteria derived from NATO AI governance principles. The strategy identifies that any data adapted to AI must be reliable and free from potential biases. In this sense, the dual-use nature of AI focuses number of challenges, such as "integrating civilian-market solutions into defence and the need to mitigate risks from adversarial use, such as AI-enabled disinformation and information operations that threaten democratic trust" (NATO, 10 July 2024a).

Therefore, responsible AI certification standards, assessment templates, constant monitoring of technology development trends should be integrated into policy practices. Moreover, the Strategy highlights increased cooperation with the EU as a key objective to facilitate information exchange and the sharing of best practices in AI safety and accreditation. Within the scope of the Strategy, any technology applied in the NATO or dual-use domain, must ensure the application of security-by-design principles, to confirm that AI applications are safe and compliant with international law before deployment.

In addition, Data Strategy for the Alliance (NATO, 5 May 2025) and NATO's Digital Transformation Implementation Strategy (NATO, 17 October 2024b) demonstrate NATO's strong commitment to becoming a data-driven organization:

- The Data Strategy for the Alliance emphasizes that NATO transitions into a data-centric organization. The strategy clearly highlights security-by-design and privacy-by-design principles through the implementation of a "Data Centric Reference Architecture" and the use of modern standards like STANAG 5636 to ensure semantic interoperability (NATO, 2022). These are important prerequisites to implement any AI-driven technology across the Alliance. Moreover, the strategy refers to necessity of "adhering to international agreements and approved NATO policies regarding data protection and legal restrictions" (NATO, 5 May 2025). It should be implemented by secure shared data spaces, which could help to mitigate risks associated with dual-use technologies.

- The NATO's Digital Transformation Implementation Strategy emphasizes that data-driven decision-making within the Alliance has to possess interoperability across different domains. This requires the commitment to "security and interoperability engineered by design" (NATO, 17 October 2024b). Thus, relevant standards and coherent governance are needed to ensure that capabilities remain resilient and future-proof. Moreover, the strategy emphasizes the necessity to integrate cybersecurity and cyber defence practices into operational domain, connecting them to NATO and Allied initiatives to reinforce the protection of trusted data and infrastructure. This underscores the shared responsibility of Allies to address data management and protection requirements. Any technology, which has the potential to be disseminated among the members of Alliance, must be fitted to these requirements.

The NATO policy documents referenced above highlight a significant strategic shift in the technological domain—particularly in relation to AI—while also acknowledging the challenges that dual-use technologies must meet in order to comply with both defence-sector requirements and the broader AI and data protection obligations applicable in the civilian domain.

The following are the EU policy practices relevant to dual-use application domain. The Action Plan on Synergies between Civil, Defence and Space Industries (European Commission, 2021b) emphasizes the use of dual-use technologies like AI for both commercial and military applications to boost EU innovation. The Commission Communication "Roadmap on critical technologies for security and defence" (European Commission, 2022b) defines strategic framework for the EU to achieve technological sovereignty, reduce dependencies on third countries for technologies, including AI, semiconductors, and cybersecurity. The European Defence Industrial Strategy (European Parliament, 2024) promotes the integration of civilian innovation (often AI-driven) into the defence sector, further blurring the lines the AI Act seeks to regulate.

In addition, the European Policy Centre report, "From Dual-Use to Omni-Use: Securing Europe's Future" (Kremidas-Courtney, 2025), argues that traditional dual-use classifications are no longer relevant in the technological advancement

era, because most of the advanced technologies (including AI) operate simultaneously across different domains, like defence, health, industry, creating an "omni-use" approach. The article argues that strict binary logic of dual-use slows innovation and imposes unnecessary compliance burdens. To address this, it recommends a shift toward "omni-use by design" in research and innovation frameworks. This is supported in Cichocki (2025), where the application of AI in maritime cybersecurity both as an offensive weapon and a defensive tool is explored.

Commission Guidelines on Prohibited AI Practices (European Commission, 2025b) explicitly emphasize that "unacceptable risk" (e.g., remote biometric identification) must be removed from AI-driven technologies, which aim to be disseminated in the EU whether in civil or dual-use domains. The guidelines specifically address practices such as harmful manipulation, social scoring, and real-time remote biometric identification, among others.

Ensuring the consistent, effective, and uniform application of the AI Act across the EU is essential; therefore, the scope for divergent interpretations of its prohibitions should be minimized. The accompanying guidelines provide legal clarification and practical examples to assist stakeholders in understanding and implementing the Act's requirements. Nevertheless, authoritative interpretations under EU primary law remain exclusively within the competence of the Court of Justice of the European Union (CJEU). Namely, the CJEU cases set the tone for further implementation AI Act and GDPR related practices. Several cases already emerged, showing case law in what ways and to what extent AI technologies can correlate with GDPR requirements, including aspects of national security, data protection, automation, and data anonymization.

In the case "La Quadrature du Net", CJEU ruled that Member States cannot cite "national security" to issue blanket mandates for mass data retention (Court of Justice of the European Union, 2020). This is crucial for the AI Act's dual-use scope: it prevents governments from easily labelling a civilian AI surveillance tool as "national security" just to bypass EU regulations. Moreover, the Court established a critical exception for national security: when a Member State faces a "genuine and present or foreseeable" serious threat to its national security, it may temporarily order the general retention of such data. This preventive measure must be strictly limited in time and subject to effective review by a court or an independent administrative body to prevent abuse.

The following are a few examples of CJEU rulings related to automated decisions:

- In the case "Ligue des droits humains", CJEU emphasized that automated processing systems must have reliable, non-discriminatory criteria and human review (Court of Justice of the European Union, 2022). This sets the judicial standard for the "human oversight" or "human in the loop" requirements in the AI Act.

- In "SCHUFA Holding AG", the Court clarified that any organization providing automated risk-based scores-including those for identity verification or fraud detection-may be subject to the strict requirements of Article 22 (Court of Justice of the European Union, 2023). Under this provision, such automated decisions are generally prohibited unless they are necessary for a contract, authorized by law.

- Finally, the CJEU ruling in "Dun & Bradstreet Austria GmbH" established that individuals have a right to know the "logic involved" in an automated decision (Court of Justice of the European Union, 2025). This directly impacts providers of high-risk dual-use AI who might try to hide their algorithms.

In summary, regardless of the domain (whether it is NATO policy and accreditation frameworks, or EU legal regulations), to a certain extent, AI-enabled technologies must balance and not exceed the limits of the AI Act and GDPR requirements, which is essentially confirmed by both policy documents and currently emerging case law.

2.5. Related works on automating regulatory compliance

A recent study on requirements engineering for regulatory compliance in software systems (Kosenkov et al., 2025) identifies numerous persistent challenges, including the abstract nature of regulations, conflicts with existing practices, the need for new expertise, lack of established principles and methods, regulatory complexity, resource demands, enforcement difficulties, evolving system dynamics, overlapping regulatory frameworks, organizational barriers, and regulatory gaps. Out of the 280 publications reviewed in this study, 255 (91%) reported at least one of these challenges. Thus, there is an urgent need for establishing foundational principles and practices that can support development of systematic methodologies and software tools for automating regulatory compliance. Currently, the interpretation and application of regulatory frameworks remain heavily dependent on the specialized knowledge of legal experts (Elahi-doost, 2024).

Generative AI (GenAI) has the potential to enhance regulatory compliance (Imperial et al., 2025). However, due to above-mentioned challenges – especially, the overlapping regulatory frameworks, regulatory gaps, and conflicts with existing practices – we cannot rely on the AI models yet. Furthermore, questions of liability and accountability persist, particularly regarding responsibility for misinterpretations or incomplete outputs generated by GenAI. Additionally, frequent regulatory updates require continuous model retraining to avoid outdated interpretations. Current standards and regulations are also not fully aligned with GenAI capabilities, and their effective application still demands substantial domain expertise (Imperial et al., 2025).

The first step toward automating regulatory compliance is the abstraction and formal structuring of legal provisions into a machine-interpretable representation (Abualhaija et al., 2025). It can be achieved through extracting the concepts – entities, their relationships and attributes – from the legal texts to build a model or ontology. These concepts then serve as the basis for specifying compliance criteria. Once the legal knowledge is appropriately represented, automated methods can support analysts in various compliance related tasks. Abualhaija et al. (2025) identified two possible scenarios for using automated tools for determining regulatory compliance: (i) retrieving compliance relevant information using question answering techniques, and (ii) applying analysis (e.g., Machine Learning) methods to assist in compliance assessment by classifying provided text according to predefined concepts.

An ontology is a formal specification of a conceptualization, defining the relevant concepts and relationships. Ontologies provide explicit meanings for each concept and specify the constraints governing their consistent use. Their reliance on strict logical formalisms makes them more complex to construct as compared to other conceptual models but enables automated validation of knowledge consistency and the use of semantic reasoners to derive new insights. Ontologies are typically expressed in languages that abstract away from data structures and implementation details, making them well suited for integrating heterogeneous databases and supporting interoperability across diverse systems (Abualhaija et al., 2025).

Thus, Ontology-based Requirements Engineering (ObRE) is a promising way forward for advancing regulatory compliance automation (Amaral et al., 2021; Guizzardi et al., 2023). The following are several examples of using ontological frameworks to automate regulatory compliance:

- Hosseini et al. (2025) proposed an ontological framework to support security-by-design of industrial control systems by conforming with the IEC 62443 standard.
- Hernandez et al. (2024) introduced a Trustworthy AI Requirements (TAIR) ontology, which provides the basis for mapping the concepts and requirements from the AI Act.
- Guizzardi et al. (2023) proposed an ontological approach for specifying so called “ethicality requirements” – requirements for developing ethical systems, including AI systems.
- Paskauskas (2025) developed an ontology for 5G network security ensuring compliance with EU cybersecurity requirements.
- Orlando & Santoro (2025) proposed a framework that integrates natural language processing (NLP) with semantic and topic modelling techniques for automated analysis of GDPR enforcement decisions.

The major limitations of currently available ontological frameworks for supporting regulatory compliance include:

a) Limited scope: not taking into consideration multiple standards and regulations.

b) Insufficient for dual-use: not addressing the needs of developing dual-use or omni-use technologies.

3. Methodology for dual-use AI system regulatory compliance requirement specification

3.1. Methodology steps

This section describes the steps of the proposed methodology for ontology-based regulatory compliance requirements specification for dual-use AI systems. Its scope is presented in Figure 1. It used regulations (both, for civilian and defence use) as an input to perform regulatory compliance analysis and specifies requirements for pre-deployment and post-deployment phases. A list of regulation, shown in Figure 1, is only an example and does not include a complete list of regulations relevant for dual-use AI systems.

Methodology comprises six steps, progressing from problem definition through validated knowledge graph deployment.

Step 1: Problem Framing and Regulatory Scoping

The first step establishes the dual-use context and identifies the applicable regulatory landscape across both civilian and defence domains.

Activities:

- Define the technology under consideration and its potential deployment contexts (civilian, defence, or both);
- Identify applicable civilian regulations (e.g., GDPR, NIS2 Directive, EU AI Act, Cyber Resilience Act, national cybersecurity laws);
- Identify applicable defence policy and accreditation frameworks (e.g., NATO Information Assurance Policy, NATO Security Policy, NATO Cyber Defence Policy, TEMPEST standards, STANAG agreements);
- Document the dual-use compliance challenge: where requirements overlap, diverge, or potentially conflict.

Outputs:

- Regulatory scope document listing all applicable frameworks by domain;
- Initial mapping of regulatory hierarchies (international → EU/NATO → national);
- Identification of cross-domain requirements (e.g., GDPR may still apply where personal data processing falls outside national security/defence derogations).

Rationale:

Dual-use technologies must satisfy regulatory requirements from fundamentally different legal regimes. Early scoping prevents gaps in compliance coverage and identifies potential conflicts that must be resolved in system design.

Step 2: Conceptual Framework Development

The second step translates the regulatory scoping into an explicit conceptual model that structures requirements across domains, lifecycle phases, and categories.

Activities:

- Create a visual representation of the regulatory compliance framework showing:
  - Two top-level compliance domains (civilian use, defence use);
  - Regulatory source hierarchies feeding into each domain;
  - Lifecycle phase decomposition (design/development vs. usage/maintenance);
  - Requirement categorization (privacy-related vs. security-related).
- Map design principles to lifecycle phases:
  - Privacy-by-design (GDPR Article 25) → primarily governs design/development phase;
  - Security-by-design (GDPR Article 32) → primarily governs usage/maintenance phase;
- Define the structural parallelism: both domains follow identical phase and category organization to enable systematic comparison.

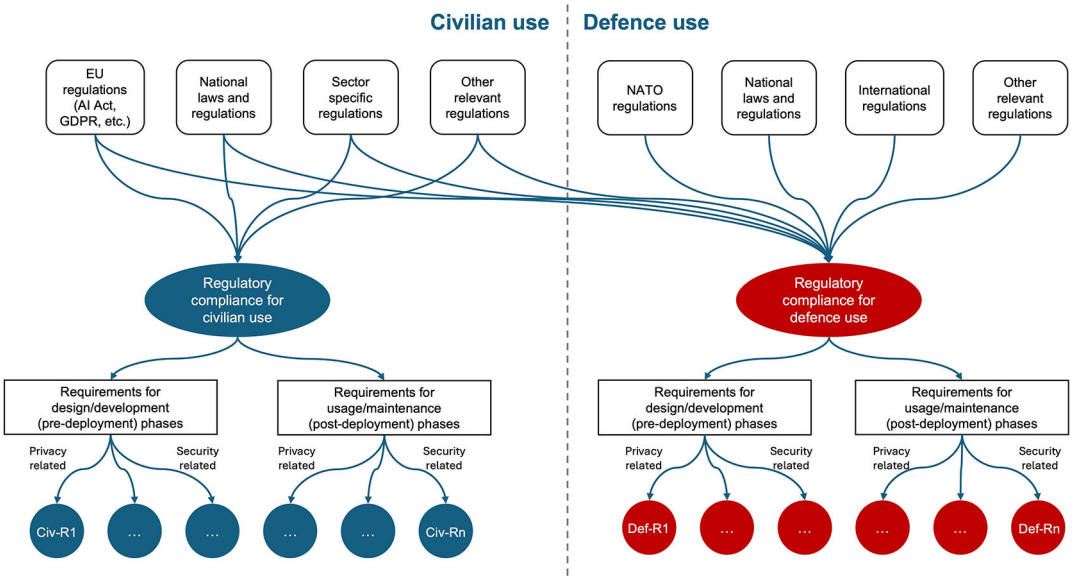

Figure 1. Methodology scope. (Methodology uses civilian and defence regulations as an input to perform regulatory compliance analysis for civilian and defence use respectively and specifies requirements for pre-deployment and post-deployment phases. Requirements can be classified by their type, e.g., privacy-related, security-related).

Outputs:

- Conceptual diagram showing dual-domain regulatory framework (as shown in Figure 1);
- Design principle mapping to lifecycle phases;
- Structural template for requirement specification (domain × phase × category).

Rationale:

A visual conceptual model serves as a shared reference between legal experts, ontology engineers, and system developers. The structural parallelism ensures that civilian and defence requirements can be systematically compared and that gaps in either domain are immediately visible.

Step 3: Ontology Design and Implementation

The third step formalizes the conceptual model as a machine-readable ontology expressed in RDF/Turtle.

Activities:

- Define namespace and prefixes (e.g., rc: <http://l3ce.lt/ontology/regulatory-compliance>)
- Create class hierarchy:
  - Root class: RegulatoryCompliance;
  - Domain classes: RegComplyforCivilUse, RegComplyforDefenceUse;
  - Regulatory framework classes: EURegulations, NATORegulations, NatLawsAndRegs;
  - Specific regulations: GDPR, NIS2Directive, NATOSecurityPolicy, etc.;
  - Structural classes: LifecyclePhase, DesignPrinciple, RequirementCategory, Requirement.

- Define object properties:
  - hasPhase — links compliance domain to lifecycle phases;
  - hasRequirement — links phases to specific requirements;
  - derivedFrom — links requirements to source regulations;
  - hasRequirementCategory — classifies requirements as privacy-related or security-related;
  - appliesToDomain — tags requirements for civilian or defence context;
  - hasDefenceEquivalent/hasCivilianEquivalent — maps cross-domain relationships.
- Define datatype properties:
  - requirementIdentifier — unique identifier string (e.g., "CIV-DESIGN-R1");
  - articleReference — citation to specific regulatory articles;
  - legalBasis — substantive description of the compliance obligation.

Outputs:

- RDF/Turtle ontology file defining classes, properties, and structural relationships;
- Documented namespace and prefix conventions;
- Class hierarchy diagram.

Rationale:

RDF/Turtle provides a standardized, interoperable format that can be loaded into knowledge graph platforms (e.g., Stardog, GraphDB) and queried using SPARQL. The formal semantics enable automated consistency checking and reasoning. The visual representation is depicted in the STARDOG Knowledge Graph in Figure 2.

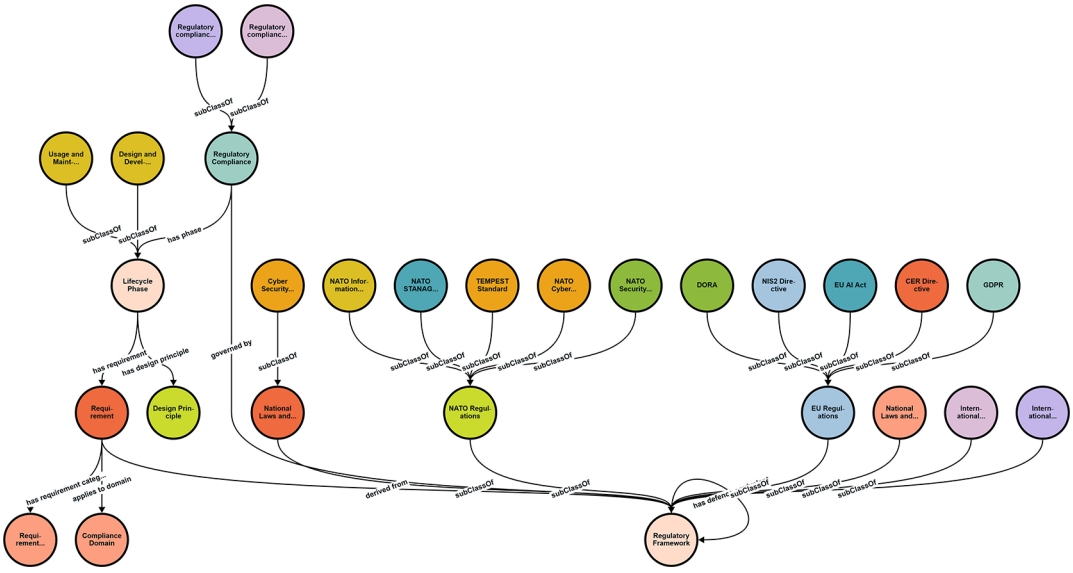

Figure 2. Regulatory compliance ontological framework for dual-use technologies.

Step 4: Domain-Specific Requirements Population

The fourth step instantiates concrete regulatory requirements within the ontological framework, populating both civilian and defence modules.

Activities:

- For each domain (civilian, defence), we create requirement instances following the structural template:
  - Design/Development phase: one privacy-related requirement (R1), one security-related requirement (R2);
  - Usage/Maintenance phase: one privacy-related requirement (R3), one security-related requirement (R4);
- Populate each requirement instance with:
  - requirementIdentifier: structured ID reflecting domain, phase, and sequence (e.g., "DEF-USAGE-R4");
  - rdfs:label: human-readable requirement name;
  - derivedFrom: link to source regulation(s);
  - articleReference: specific article/section citations;
  - legalBasis: substantive text describing the compliance obligation;
  - hasRequirementCategory: classification as PrivacyRelated or SecurityRelated;
  - appliesToDomain: explicit domain tag (CivilianDomain or DefenceDomain);
- Where regulations support decomposition, we create sub-requirements linked to parent (e.g., DEF-USAGE-R4 → DEF-USAGE-R4-ART3, DEF-USAGE-R4-ART5, DEF-USAGE-R4-ART7).

Outputs:

- Populated civilian requirements module (Civ-R1 through Civ-R4);
- Populated defence requirements module (Def-R1 through Def-R4);
- Sub-requirements where applicable;
- Complete RDF/Turtle file ready for knowledge graph deployment.

Rationale:

The parallel structure across domains (R1–R4 in both civilian and defence) enables systematic comparison and gap analysis. Explicit domain tagging via appliesToDomain provides a reliable partitioning mechanism that supplements class-based inference.

Step 5: Knowledge Graph Deployment

The fifth step loads the ontology into a triplestore platform and verifies structural integrity.

Activities:

- Create database instance in target platform (e.g., Stardog Cloud);
- Load RDF/Turtle ontology file;

- Verify successful loading by checking triple count;
- Configure visualization settings for knowledge graph exploration;
- Test basic navigation: browse class hierarchy, inspect individual requirement instances.

Outputs:

- Deployed knowledge graph accessible for querying;
- Visual representation of regulatory compliance framework;
- Confirmation of structural integrity (all expected classes, properties, and instances present).

Rationale:

Knowledge graph deployment enables interactive exploration of the regulatory framework by stakeholders with varying technical backgrounds. Visual navigation supports validation with domain experts (legal, compliance, engineering) who may not be familiar with SPARQL.

Step 6: SPARQL-Based Validation and Cross-Domain Analysis

The sixth step validates the framework through structured queries and demonstrates its analytical capabilities for dual-use compliance planning.

Activities:

- Design validation queries to verify:
  - All requirement instances are correctly typed and populated;
  - Domain tagging is consistent (appliesToDomain values);
  - Lifecycle phase associations are complete;
  - Regulatory source linkages (derivedFrom) are present.
- Execute cross-domain comparison queries:
  - Retrieve all requirements grouped by domain;
  - Pair civilian and defence requirements by structural position (e.g., CIV-DESIGN-R1 ↔ DEF-DESIGN-R1);
  - Identify shared regulatory sources;
  - Identify domain-specific sources.
- Iteratively refine ontology based on query results:
  - Correct missing or malformed property values;
  - Add missing cross-references;
  - Enhance requirement decomposition where needed.

Outputs:

- Validated query results demonstrating framework completeness;
- Cross-domain comparison tables (see Tables 1–4);
- Documentation of any refinements made during validation;
- SPARQL query library for reuse in future compliance assessments (viz. GitHub repo).

Rationale:

- SPARQL-based validation ensures that the ontological framework is complete, consistent, and queryable. The cross-domain comparison queries directly support dual-use compliance planning by revealing where civilian and defence requirements align (shared obligations) and where they diverge (domain-specific security frameworks), as identified in Step 1.

Table 1. Summary of six steps of the methodology.

Step	Focus	Key Output
1	Problem Framing & Regulatory Scoping	Regulatory scope document
2	Conceptual Framework Development	Visual model and structural template
3	Ontology Design & Implementation	RDF/Turtle ontology file
4	Domain-Specific Requirements Population	Populated civilian and defence modules
5	Knowledge Graph Deployment	Deployed, navigable knowledge graph
6	SPARQL-Based Validation & Analysis	Validated results and query library

Table 2. Civilian domain requirements and legal basis.

Req. ID	Requirement	Category	Article Reference
CIV-DESIGN-R1	DPIA and Privacy by Design	Privacy-related	GDPR Article 25 (Data protection by design), Article 35 (DPIA)
Legal basis	Data Protection Impact Assessment required before deployment. Data minimization principles must govern image capture design: lens detection must avoid unnecessary recording, capture only what is essential for security verification. Retention periods must be defined during design phase, not post-deployment. System architecture must embed privacy safeguards from inception.		
CIV-DESIGN-R2	Security Architecture by Design	Security-related	GDPR Article 32 (Security of processing), NIS2 Article 21 (Cybersecurity risk-management measures)
Legal basis	Encryption architecture must be designed into system from inception. Secure communication protocols required for data transmission to remote monitoring station. Hardened firmware and secure boot mechanisms must be specified during development. Access control frameworks must be architected before deployment.		
CIV-USAGE-R3	Breach Notification and Data Subject Rights	Privacy-related	GDPR Article 33 (Notification to supervisory authority), Article 34 (Communication to data subject)
Legal basis	Personal data breaches must be reported to supervisory authority within 72 hours of becoming aware. High-risk breaches require communication to affected data subjects without undue delay. Organizations must implement data subject rights: right of access (Art. 15), right to erasure (Art. 17), right to restriction (Art. 18). Transparency notices required in deployment locations informing individuals of monitoring.		
CIV-USAGE-R4	Operational Cybersecurity Measures	Security-related	NIS2 Article 21 (Cybersecurity risk-management measures), GDPR Article 32
Legal basis	Encrypted data transmission required for all communications to remote monitoring server. Access controls must restrict image viewing to authorized personnel only. Comprehensive audit logging of all system access and data processing activities. Incident response procedures must be documented and tested. Vulnerability management and patch deployment processes required. Regular security assessments and penetration testing.		

Table 3. **Defence domain requirements and legal basis.**

Req. ID	Requirement	Category	Phase	Regulatory Source	Article Reference
DEF-DESIGN-R1	DPIA and Privacy by Design	Privacy-related	Design & Development	Law of the Republic of Lithuania on the Legal Protection of Personal Data Processed for the Purposes of the Prevention, Investigation, Detection or Prosecution of Criminal Offences, the Execution of Sentences or National Security or Defence*	Article 25 Article 18
Legal basis	Data Protection Impact Assessments identify risks, evaluate safeguards, and ensure compliance before systems go live. Embedded in privacy by design and by default principle.				
DEF-DESIGN-R2	NATO Security Controls	Security-related	Design & Development	NIAP, NSP	—
Legal basis	NIAP: Multi-layered security controls including encryption, access controls, and incident response. NSP: Access controls, physical security, & personnel training for classified info.				
DEF-USAGE-R4	NATO Cyber Defence Compliance	Security-related	Usage & Maintenance	NATO Cyber Defence Policy	—
Legal basis	Protection of digital assets from cyber threats, including unauthorised data capture through non-networked means such as cameras.				

*Applicable only in Lithuania and might vary in other Member States.

Table 4. **Example of refining requirements into sub-requirements.**

Req. ID	Sub-Requirement	Article Reference	Legal Basis
DEF-USAGE-R4-ART3	Detection Mechanisms	Article 3, Paragraph 2	Entities must deploy defence mechanisms that include the detection and reporting of unauthorised data capture.
DEF-USAGE-R4-ART5	Secure Incident Reporting	Article 5, Paragraph 8	Detected incidents must be reported through secure, encrypted channels.
DEF-USAGE-R4-ART7	Threat Intelligence Sharing	Article 7, Paragraph 6	Threat intelligence must be shared across NATO entities.

The following are two example queries: A – retrieves all requirements across civilian and defence domains; B – pairs civilian and defence requirements side-by-side. More details on query outputs are provided in Section 4.

Query A: Retrieve all requirements across both domains

```
PREFIX rc: <http://l3ce.lt/ontology/regulatory-compliance>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema>
SELECT ?identifier ?label ?domain ?category ?article ?legalBasis
WHERE {
       ?req a rc:Requirement;
          rdfs:label? label;
          rc:requirementIdentifier? identifier;
          rc:hasRequirementCategory? category;
          rc:appliesToDomain? domain.
       OPTIONAL {?req rc:articleReference? article}
```

```
        OPTIONAL {?req rc:legalBasis? legalBasis}
        FILTER(?domain IN (rc: CivilianDomain, rc: DefenceDomain))
}
ORDER BY? domain? identifier
```

Query B: Side-by-side cross-domain pairing

```
PREFIX rc: <http://l3ce.lt/ontology/regulatory-compliance>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema>
SELECT? civId? civLabel? civArticle? defId? defLabel? defArticle
WHERE {
        ?civReq a rc: Requirement;
                rc:requirementIdentifier? civId;
                rdfs:label? civLabel;
                rc:appliesToDomain rc: CivilianDomain.
        OPTIONAL {?civReq rc:articleReference? civArticle}
        ?defReq a rc: Requirement;
                rc:requirementIdentifier? defId;
                rdfs:label? defLabel;
                rc:appliesToDomain rc: DefenceDomain.
        OPTIONAL {?defReq rc:articleReference? defArticle}
        FILTER (REPLACE(?civId, "CIV", "") = REPLACE(?defId, "DEF", "")).
}
ORDER BY ?civId
```

3.2. Methodology summary

The six-step methodology provides a systematic approach to specifying regulatory compliance requirements for dual-use AI systems. The focus and key outputs of each step are summarized in Table 1.

The methodology addresses the limitations identified in Section 2.5 by: (a) explicitly considering multiple regulations across both civilian and defence domains, and (b) providing a structured approach for dual-use technology compliance that maintains traceability between legal requirements and technical specifications.

3.3. Methodology innovation: from static mapping to executable compliance logic

Existing regulatory compliance approaches in AI governance predominantly rely on static documentation, checklist-based conformity assessments, or manual legal interpretation. Even ontology-based models frequently function as conceptual mapping tools without executable validation capacity.

The methodological contribution of this study lies in transforming regulatory requirements into executable compliance logic. The framework operationalises legal obligations as structured RDF entities linked to lifecycle phases, deployment domains, and regulatory sources, enabling:

1) Query-based compliance verification - SPARQL queries retrieve and compare applicable obligations dynamically depending on deployment context.

2) Cross-domain alignment logic - Structural pairing of civilian and defence requirements supports automated gap analysis.

3) Hierarchical decomposition with traceability - Article-level sub-requirements enable fine-grained mapping between system components and specific legal provisions.

4) Conditional applicability modelling - Domain tagging and potential restriction flags allow the representation of context-dependent regulatory scope (e.g., GDPR scope exemptions and Article 23 restrictions, AI Act military exemption).

Table 5. Aligning civilian and defence domain requirements.

Civilian ID	Civilian Requirement	Civilian Article Reference	Defence ID	Defence Requirement	Defence Article Reference
CIV-DESIGN-R1	DPIA and Privacy by Design	GDPR Article 25, Article 35	DEF-DESIGN-R1	DPIA and Privacy by Design	Law of the Republic of Lithuania on the Legal Protection of Personal Data Processed for the Purposes of the Prevention, Investigation, Detection or Prosecution of Criminal Offences, the Execution of Sentences or National Security or Defence, Article 25, Article 18
CIV-DESIGN-R2	Security Architecture by Design	GDPR Article 32, NIS2 Article 21	DEF-DESIGN-R2	NATO Security Controls	—
CIV-USAGE-R3	Breach Notification and Data Subject Rights	GDPR Article 33, Article 34			
CIV-USAGE-R4	Operational Cybersecurity Measures	NIS2 Article 21, GDPR Article 32	DEF-USAGE-R4	NATO Cyber Defence Compliance	—

This shifts compliance engineering from static documentation toward an executable knowledge representation capable of supporting iterative system design, accreditation preparation, and deployment-specific compliance configuration.

4. Specification of regulatory compliance requirements using proposed methodology

This section demonstrates how the methodology works and what results it generates. Currently, ontology captures a subset of applicable regulations (GDPR, NIS2, NATO policies) and their selected features to show how the methodology works. In the future, it should be expanded to include all relevant provisions of the selected regulations as well as the EU AI Act risk classification, Cyber Resilience Act, sector-specific and other relevant regulations and standards.

Tables 2 through 5 present the SPARQL query results showing regulatory compliance requirements across both civilian and defence deployment contexts, organized by domain, lifecycle phase, and requirement category. The results shown in Tables 2–5 are illustrative examples and do not constitute a complete set of requirements.

4.1. Regulatory compliance requirements for civilian use

Table 2 shows non-exhaustive examples of civilian domain requirements derived from EU regulations (GDPR, NIS2 Directive), organised by lifecycle phase (DESIGN = design/development; USAGE = usage/maintenance) and category (privacy-related, security-related). Legal basis for each requirement is included in Table 2 as well. It captures the substantive compliance obligations derived from regulatory frameworks; however, the obligations in the following table are included for exemplary purposes only and should be expanded to capture more detailed obligations in real-life scenarios. This enables automated compliance checking and traceability between system design decisions and regulatory mandates.

4.2. Regulatory compliance requirements for defence use

Table 3 presents the ontological representation of regulatory compliance requirements for defence-use deployments, organized by lifecycle phase and requirement category. It provides the substantive compliance content for each requirement, showing how the ontological model captures actionable regulatory obligations. Regulatory requirements are assigned to specific lifecycle phases (pre-deployment vs. post-deployment) and categories (privacy-related vs. security-related). In the provided example, requirements are derived from national regulations and NATO frameworks (NIAP, NSP, NATO Cyber Defence Policy), reflecting the dual-use nature of the system under consideration.

The SPARQL code looks like this:

```
PREFIX rc: <http://l3ce.lt/ontology/regulatory-compliance>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema>
PREFIX skos: <http://www.w3.org/2004/02/skos/core>
SELECT? identifier? label? domain? category? article? legalBasis
WHERE {
       ?req a rc: Requirement;
          rdfs:label? label;
          rc:requirementIdentifier? identifier;
          rc:hasRequirementCategory? category;
          rc:appliesToDomain? domain.
       OPTIONAL {?req rc:articleReference? article}
       OPTIONAL {?req rc:legalBasis? legalBasis}
       FILTER(?domain IN (rc: CivilianDomain, rc: DefenceDomain))
}
ORDER BY? domain? identifier
```

It must be noted, however, that identification of certain requirements may differ depending on the particular system to which the ontological framework is applied. As briefly indicated previously, GDPR includes certain exemptions, including for defence data processing operations. Therefore, for example, the DPIA requirement may not be attributed to defence domain if the system is intended to be used solely by military for military operations which fall outside the EU competences scope. However, as included in. the table, certain data protection obligations as related to defence domain can be established in national laws – for example, the Law of the Republic of Lithuania on the Legal Protection of Personal Data Processed for the Purposes of the Prevention, Investigation, Detection or Prosecution of Criminal Offences, the Execution of Sentences or National Security or Defence, includes the obligation to perform a DPIA (Article 25). When the competent authorities process personal data for the purposes of national security or defence, the Law applies unless otherwise provided for in other laws. In other words, while the GDPR may exclude military operations from its scope, similar obligations still may be required under the national military-related laws. And vice versa – not all data processing operations performed by defence organisations necessarily fall outside the EU regulation, e.g., if defence organisations process data not for exclusively military purposes. Therefore, even regulatory domain distinctions must be done with extreme cautiousness to capture all possible scenarios so that no AI system component slips through the regulatory gaps.

The ontological framework supports hierarchical requirement decomposition, as demonstrated by the granular specification of NATO Cyber Defence Policy obligations, and can be evoked by executing the code below:

```
PREFIX rc: <http://l3ce.lt/ontology/regulatory-compliance>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema>
SELECT? identifier? label? category? article? legalBasis
WHERE {
       ?req a rc: Requirement;
          rdfs:label? label;
          rc:requirementIdentifier? identifier;
          rc:hasRequirementCategory? category.
       OPTIONAL {?req rc:articleReference? article}
       OPTIONAL {?req rc:legalBasis? legalBasis}
       FILTER (CONTAINS(?identifier, "DEF-"))
}
ORDER BY? identifier
```

Table 4 illustrates the framework's capacity for requirement decomposition. The parent requirement DEF-USAGE-R4 (NATO Cyber Defence Compliance) defined in Table 3 is refined into article-level sub-requirements, enabling precise traceability between system components and specific regulatory provisions.

4.3. Aligning civilian and defence requirements

Table 5 presents direct SPARQL query output demonstrating the framework's ability to programmatically align civilian and defence requirements by structural position. It includes results from a SPARQL query that automatically pairs civilian and defence requirements by matching their structural position (e.g., CIV-DESIGN-R1 pairs with DEF-DESIGN-R1). This demonstrates the framework's support for automated cross-domain compliance analysis.

The SPARQL query looks like this:

```
PREFIX rc: <http://l3ce.lt/ontology/regulatory-compliance>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema>
SELECT? civId? civLabel? civArticle? defId? defLabel? defArticle
WHERE {
       ?civReq a rc: Requirement;
              rc:requirementIdentifier? civId;
              rdfs:label? civLabel;
              rc:appliesToDomain rc: CivilianDomain.
       OPTIONAL {?civReq rc:articleReference? civArticle}
       ?defReq a rc: Requirement;
              rc:requirementIdentifier? defId;
              rdfs:label? defLabel;
              rc:appliesToDomain rc: DefenceDomain.
       OPTIONAL {?defReq rc:articleReference? defArticle}
       FILTER (REPLACE(?civId, "CIV", "") = REPLACE(?defId, "DEF", ""))
}
ORDER BY? civId
```

4.4. Regulatory compliance requirement specification summary

In summary, The SPARQL query results, presented in Tables 2-5, demonstrate three key capabilities of the ontological framework:

1) *Lifecycle-phase organization:* Requirements are systematically allocated to either the design/development phase (where privacy-by-design and security-by-design principles are embedded) or the usage/maintenance phase (where operational compliance obligations apply).

2) *Dual regulatory coverage*: The framework simultaneously captures EU civilian regulations (GDPR) and NATO defence-specific frameworks (NIAP, NSP, Cyber Defence Policy), reflecting the dual-use nature of the system/technology.

3) *Hierarchical decomposition*: High-level requirements can be decomposed into granular, article-specific sub-requirements, supporting both strategic compliance planning and detailed technical validation.

5. Case study

This section validates the proposed ontology-based methodology by applying it to a sample dual-use system – espionage detection system, which uses AI technology designed to detect and report unauthorised photographing of classified or confidential information displayed on screens. The system does not perform biometric identification, categorisation, or profiling of individuals; it detects behavioural indicators (e.g., camera lenses directed toward protected displays) without associating detected persons with identifiable biometric templates.

5.1. System description

The espionage detection system comprises several main components (see Figure 3). It includes: a monitor, which displays sensitive/classified information; an AI system, connected to monitor via Ethernet, which includes software and hardware components to monitor analyse the behaviour of people in the room; a camera connected to AI system via USB for visual monitoring of the room; a remote monitoring station, which communicates with the AI system via 5G network. Beyond basic hardware, the system's "intelligence" is grounded in an ontology-driven context-aware framework. While the camera and AI module identify unauthorised photography, recent 2025 research suggests that semantic reasoning is required to distinguish between legitimate mobile use and an actual "espionage" event. By integrating a Re-Identifiability Index (RII), the system can mitigate identity leakage during visual monitoring, ensuring that only the "threat" is processed while preserving bystander privacy through adaptive obfuscation (Topalli & Badii, 2025). Furthermore, AI system could

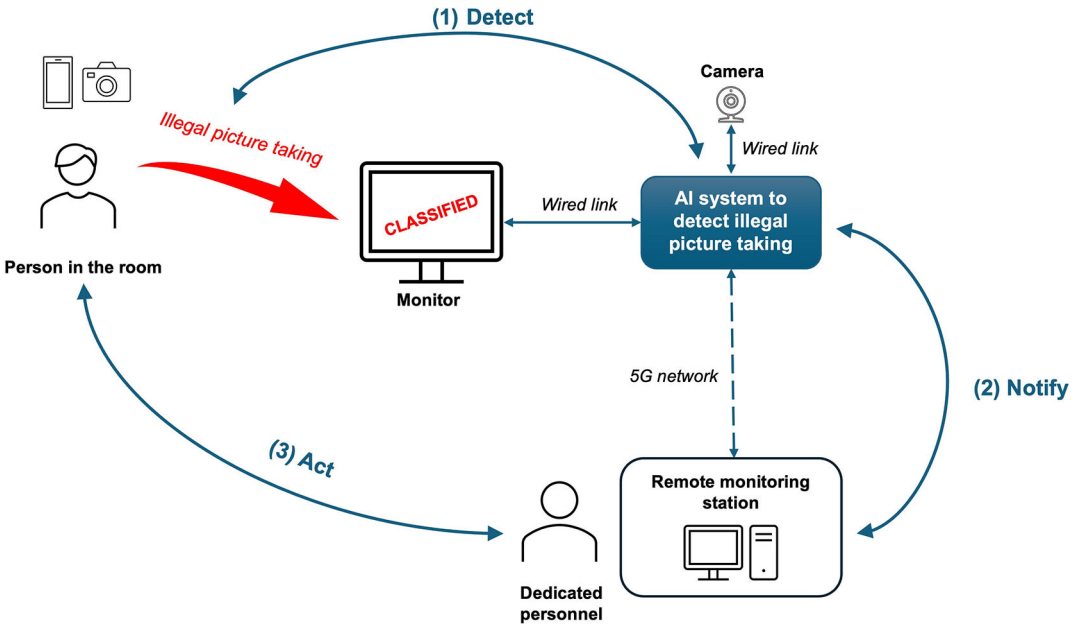

Figure 3. Espionage detection system architecture and operational flow.

implement additional features such as scene interpretation (Gómez-Romero et al., 2011) and deepfake and biometric security (Pn et al., 2024).

Operational workflow of espionage detection consists of three phases (see Figure 3):

(1) *Detect.* Camera monitors the room, where classified information is displayed on a screen. The AI system analyses the video feed to detect attempts at unauthorised picture-taking, such as the presence of camera lenses or smartphone screens pointed at the display.

(2) *Notify.* Upon detection of a potential security breach, the system transmits an alert notification to the remote monitoring station. Notification includes captured image (potentially containing faces of individuals present), AI system's confidence level, timestamp, information about what was displayed on screen at the time of detection, and monitor's location data (remote monitoring station could be receiving information from multiple espionage detection AI systems at the same time).

(3) *Act.* Dedicated security personnel at the remote monitoring station receive the alert and take appropriate response actions.

The system is intended for deployment in both civilian contexts (corporate boardrooms, government offices, financial institutions, healthcare facilities) and defence contexts (NATO command centres, military briefing rooms, classified research facilities). This dual-use nature creates overlapping and sometimes divergent compliance obligations that the ontological framework must address.

5.2. Component-to-requirement mapping

To validate the ontological framework, we systematically map each espionage detection system component and data flow to the applicable regulatory requirements across both civilian and defence domains. The results are shown in Table 6, which demonstrates how the ontological framework enables systematic mapping of system components to applicable regulatory requirements. Each component is traced to specific requirements in both civilian and defence domains, supporting dual-use compliance planning. The results presented in Section 5 tables (Tables 6–17) are illustrative and do not represent an exhaustive set of requirements.

5.3. Analysis of cross-domain compliance patterns

The component mapping reveals two distinct patterns in how regulatory requirements apply across civilian and defence deployment contexts.

Pattern 1: Divergent Requirements (Domain-Specific Security Frameworks)

Security-related requirements (R2 and R4) diverge significantly between domains. Civilian deployments rely on EU cybersecurity frameworks, while defence deployments must satisfy NATO-specific standards (see Table 7).

Key finding: A dual-use system cannot simply apply civilian security standards to defence deployments. The ontological framework captures this divergence by linking R2 and R4 to different regulatory sources depending on the domain.

Pattern 2: Additive Requirements (Defence Builds Upon Civilian)

In systems that are developed for use in both civilian and defence domains, defence deployments do not replace civilian requirements – they add to them. An Espionage detection system deployed in a NATO facility must satisfy:

- National law (R1)
- All applicable NIS2 requirements (civilian baseline) plus
- NATO-specific requirements (NIAP, NSP, Cyber Defence Policy)
- Potential TEMPEST certification for electromagnetic emission security

Table 6. Espionage detection system Component Mapping to Regulatory Requirements.

System Component	Data/Activity	Applicable Civilian Requirements	Applicable Defence Requirements
USB Camera	Captures images of room occupants (potential biometric data)	CIV-DESIGN-R1: DPIA required; data minimisation in image capture design	DEF-DESIGN-R1: DPIA required; privacy safeguards embedded in design
AI Detection Module	Processes visual data to identify unauthorised photography attempts	CIV-DESIGN-R2: Secure architecture; hardened firmware; secure boot	DEF-DESIGN-R2: NATO security controls; multi-layered access controls
Screen Interface	Ethernet connection; displays classified content	CIV-DESIGN-R2: Encryption architecture for data in transit	DEF-DESIGN-R2: NSP physical security requirements; TEMPEST considerations
5G Transmission	Transmits personal data (images, location) to remote server	CIV-USAGE-R4: Encrypted transmission; audit logging	DEF-USAGE-R4: Secure encrypted channels (Art. 5 [8])
Remote Monitoring Station	Receives alerts; stores images; enables personnel access	CIV-USAGE-R3: Breach notification within 72 hours; data subject rights	
		CIV-USAGE-R4: Access controls; audit logging; vulnerability management	DEF-USAGE-R4-ART3: Detection and reporting mechanisms
Dedicated Personnel	View captured images; make response decisions	CIV-USAGE-R3: Transparency notices; purpose limitation	DEF-USAGE-R4-ART7: Threat intelligence sharing across NATO entities

Table 7. Divergent requirements regarding design and usage security.

Requirement	Civilian Framework	Defence Framework	Key Differences
R2 (Design, Security)	GDPR Article 32, NIS2 Article 21	NATO NIAP, NATO NSP	Civilian: commercial encryption standards (TLS, AES). Defence: NATO-approved cryptography; personnel security clearances; physical security requirements
R4 (Usage, Security)	NIS2 Article 21, GDPR Article 32	NATO Cyber Defence Policy	Civilian: report to national supervisory authorities. Defence: secure encrypted channels; threat intelligence sharing across NATO entities

Table 8. Requirements mapped to compliance layer.

Compliance Layer	Requirements
EU Data Protection	CIV-R1, CIV-R3 (GDPR)
EU Cybersecurity	CIV-R2, CIV-R4 (NIS2, GDPR Art. 32)
NATO Information Assurance	DEF-R2 (NIAP, NSP)
NATO Operational Security	DEF-R4, DEF-R4-ART3, DEF-R4-ART5, DEF-R4-ART7
Certification Standards	TEMPEST, STANAG (where applicable)

Table 9. Regulatory requirements across civilian and defence domains.

Metric	Result
Civilian requirements populated	4 (CIV-R1 through CIV-R4)
Defence requirements populated	4 (DEF-R1 through DEF-R4)
Defence sub-requirements	3 (DEF-R4-ART3, DEF-R4-ART5, DEF-R4-ART7)
Total requirements captured	11
Regulatory sources represented	7 (GDPR, NIS2, NIAP, NSP, NATO Cyber Defence Policy, national laws, TEMPEST)

Table 10. Systems development life cycle phases supported by design principles.

Lifecycle Phase	Privacy-Related	Security-Related	Design Principle
Design and Development (pre-deployment)	R1 (both domains)	R2 (both domains)	Privacy-by-design (GDPR Art. 25, national law)
Usage and Maintenance (post-deployment)	R3 (civilian domain)	R4 (both domains)	Security-by-design (GDPR Art. 32)

Table 11. Civilian and defence requirements compared.

Position	Cross-Domain Query Result	Finding
R1 (Design, Privacy)	Civilian: GDPR Articles 25, 35; Defence: DPIA and Privacy by Design	DPIA is required, but legal basis differs
R2 (Design, Security)	Civilian: GDPR/NIS2; Defence: NIAP/NSP	Divergent frameworks – different security standards apply
R3 (Usage, Privacy)	Civilian: GDPR Article 33	Applicable only in civilian domain
R4 (Usage, Security)	Civilian: NIS2; Defence: NATO CDP	Divergent frameworks – different reporting and operational security

Table 8 effectively delineates these distinctions.

Key finding: The regulatory burden for defence deployment is cumulative, not substitutive. The ontological framework supports this through explicit domain tagging (*appliesToDomain*), enabling queries that retrieve all applicable requirements for a given deployment context.

5.4. Validation of framework capabilities

The case study validates four key capabilities of the ontological framework, supported by recent advancements in semantic modelling and automated compliance.

Table 12. **Decomposition of high-level requirements: traceability.**

Parent Requirement	Sub-Requirements	Regulatory Source
DEF-USAGE-R4 (NATO Cyber Defence Compliance)	DEF-R4-ART3: Detection mechanisms	Article 3, Paragraph 2
	DEF-R4-ART5: Secure incident reporting	Article 5, Paragraph 8
	DEF-R4-ART7: Threat intelligence sharing	Article 7, Paragraph 6

Table 13. **Validation results.**

Capability	Validated	Evidence
Dual-use coverage	✓	11 requirements across 7 regulatory sources
Lifecycle phase organization	✓	Requirements correctly mapped to design vs. usage phases
Cross-domain comparison	✓	SPARQL queries successfully pair and compare requirements
Hierarchical decomposition	✓	NATO CDP decomposed to article-level sub-requirements

Table 14. **Espionage detection system as a single technical platform.**

Development Decision	Civilian Deployment	Defence Deployment
Encryption standard	Commercial TLS/AES (GDPR Art. 32, NIS2 Art. 21)	NATO-approved cryptographic modules
Access control	Role-based access with audit logging	Security clearance verification; need-to-know enforcement
Incident reporting	National supervisory authority (within 72 hours)	Secure encrypted channels to NATO entities
Data retention	Defined by DPIA; subject to erasure requests	May be extended for security evidence; erasure rights potentially limited
Certification	CE marking; NIS2 compliance attestation	TEMPEST certification; STANAG compliance

Table 15. **Compliance documentation as regards Civilian-Defence deployment.**

Documentation	Civilian Deployment	Defence Deployment
Regulatory basis	GDPR, NIS2 Directive	NIAP + NSP + NATO CDP
Applicable requirements	CIV-R1, CIV-R2, CIV-R3, CIV-R4	DEF-R1, DEF-R2, DEF-R4 (+ sub-requirements)
Certification evidence	NIS2 compliance; GDPR accountability documentation	TEMPEST certificate; NATO accreditation
Audit trail	Access logs; breach register; DPIA	Above + threat intelligence sharing records

Capability 1: Dual-Use Coverage

The framework demonstrates feasibility of Dual-Use Coverage by successfully capturing 11 requirements across seven regulatory sources (see Table 9). This dual-domain approach is increasingly critical as AI-driven maritime and critical

Table 16. Market positioning.

Market Segment	Compliance Requirements	Supervisory Pathway
Corporate sector	GDPR, NIS2 (if critical infrastructure)	National authority registry control
Government (civilian)	GDPR, NIS2, national cybersecurity law	National authority accreditation
Financial institutions	GDPR, DORA	Regulatory compliance supervision
NATO/Defence	GDPR + full NATO framework	TEMPEST certification; NATO accreditation

Table 17. Benefits of a validated ontological framework.

Application	Benefit
Component-to-requirement mapping	Traceability from system architecture to legal obligations
Deployment-specific queries	Generate civilian or defence compliance checklists from single knowledge base
Conflict identification	Reveals tensions between GDPR transparency and defence classification
Certification pathway planning	Identifies required certifications for each market segment
DPIA scoping	Universal entry point for both civilian and defence deployments even if derived from different legal sources

Table 18. A summary of contributions.

Contribution	Description
Dual-use compliance framework	A structured ontological approach that explicitly models regulatory requirements for technologies deployable in both civilian and defence contexts, validated across maritime, energy, and telecommunications sectors (Cichocki, 2025; Ndiaye et al., 2026)
Lifecycle-phase organization	Integration of privacy-by-design (GDPR Article 25) and security-by-design (GDPR Article 32) principles, technically supported by ontology-driven context-aware frameworks for real-time video obfuscation (Topalli & Badii, 2025)
Cross-domain analysis capability	SPARQL queries that systematically compare civilian and defence requirements, enhanced by hybrid ANN-ISM modelling to identify security interdependencies in 5G networks (Khan et al., 2025)
Validated methodology	A replicable six-step process, including SHACL constraints and logical reasoners (e.g., HermiT) to automate the detection of compliance gaps in high-risk infrastructure (Paskauskas, 2025; Ndiaye et al., 2026)
Practical compliance guidance	Actionable findings for developers, including Re-Identifiability Indices (RII) for vision systems and conflict resolution strategies for “omni-use” deployment tensions (Topalli & Badii, 2025; Cichocki, 2025)
Emerging Tech Roadmap	Extensions for 6G security (ES3A), decentralized Blockchain compliance (BAML), and federated learning, addressing the shift toward autonomous, agentic AI governance (Duan et al., 2025; Khan et al., 2025)

infrastructure systems now operate in “omni-use” environments where offensive and defensive capabilities overlap (Cichocki, 2025).

Validation evidence: SPARQL Query A (Section 3.1, Step 6) successfully retrieved all 8 primary requirements across both domains in a single query, confirming structural completeness.

Capability 2: Lifecycle Phase Organization

The framework allocates privacy-by-design to the design phase (see Table 10). This is enabled by ontology-driven context-aware frameworks that dynamically adjust privacy decisions and obfuscation levels based on entity sensitivity, such as masking bystander faces in real time while maintaining the security intelligence of the monitor's surroundings (Topalli & Badii, 2025).

Validation evidence: The espionage detection system component mapping (Table 10) demonstrates that design-phase requirements (R1, R2) apply to system architecture decisions (camera design, AI module, encryption architecture), while usage-phase requirements (R3, R4) apply to operational activities (transmission, monitoring, personnel access).

The use of four primary requirement nodes per domain (R1–R4) in this study reflects a minimal structural template designed to demonstrate methodological feasibility rather than a limitation of the ontological framework. The symmetrical organisation across domains (design/privacy, design/security, usage/privacy, usage/security) provides a controlled baseline that enables systematic cross-domain comparison and validation through SPARQL queries. In practice, the framework is not constrained to four requirements per lifecycle phase or domain. Additional requirement instances can be introduced at any level of granularity, including article-level, paragraph-level, or control-level decomposition, without altering the structural integrity of the ontology. The current configuration serves as a proof-of-concept instantiation that illustrates how regulatory obligations can be categorised, decomposed, and aligned across civilian and defence contexts. The scalability of the model derives from its class-property architecture rather than from the number of instantiated requirement nodes.

Capability 3: Cross-Domain Comparison

The framework enables systematic alignment of civilian and defence requirements (see Table 11). Research in 2025 demonstrates that while security frameworks diverge, hybrid ANN-ISM modelling can improve the accuracy of 5G network security systems by identifying interdependencies between disparate regulatory process areas (Khan et al., 2025).

Validation evidence: SPARQL Query B (Section 3.1, Step 6) successfully paired civilian and defence requirements by structural position, enabling the comparative analysis shown in Table 11.

Capability 4: Hierarchical Decomposition

High-level requirements, such as NATO Cyber Defence Compliance, are decomposed into article-specific sub-requirements (see Table 12). This decomposition facilitates the use of SHACL constraints and logical reasoners (e.g., HermiT) to automate the validation of system specifications against complex regulatory mandates, such as the High-Risk Supplier management required by the Connecting Europe Facility (Paskauskas, 2025).

Validation evidence: The sub-requirements map directly to espionage detection system functions:

- DEF-R4-ART3 → AI detection module and camera system
- DEF-R4-ART5 → 5G encrypted transmission to remote monitoring station
- DEF-R4-ART7 → Integration with NATO threat intelligence networks (for defence deployments)

Summary of validation results is presented in Table 13.

The case study confirms that the ontological framework can represent the regulatory compliance requirements for a dual-use AI system, support systematic comparison across deployment contexts, and enable traceability from legal obligations to system components.

5.5. Practical implications for system development

The validated framework yields concrete guidance for developers navigating the "additive" burden of dual-use compliance.

Implication 1: Single Design, Dual Compliance Paths

A system can be developed as a single technical platform but must support subset configurations. In 5G contexts, this requires semantic interoperability layers that allow the system to interpret information exchange consistently across heterogeneous cloud and edge environments (Brabra et al., 2016). Thus, the espionage detection system can be developed as a single technical platform, but it must satisfy different compliance paths depending on the deployment context (see Table 14).

Recommendation: Design the system architecture to accommodate the more stringent defence requirements, enabling civilian deployment as a subset configuration.

Implication 2: Unified Knowledgebase

Developers can use a single RDF-based knowledge base to generate deployment-specific documentation (see Table 15). This reduces the "interpretative burden" on smaller entities, who often face disproportionate GDPR enforcement risks despite the lack of direct correlation between violation severity and fine amounts (Orlando & Santoro, 2025).

Recommendation: Use SPARQL queries to generate deployment-specific compliance checklists. The query FILTER(?domain = rc: CivilianDomain) or FILTER(?domain = rc: DefenceDomain) retrieves only the applicable requirements for each context.

Implication 3: Procurement and Commercialisation Strategy

For the espionage detection system developer, the framework informs market positioning, as shown in Table 16.

Recommendation: Develop compliance packages for each market segment, using the ontological framework to generate segment-specific requirements documentation. The civilian package represents the baseline; defence package adds NATO-specific requirements.

5.6. Key case study findings

The espionage detection system case study validates the proposed ontology-based methodology for regulatory compliance requirements specification in dual-use AI systems. The key findings are summarised below:

1) *The methodology produces an integrated and queryable compliance framework.* The six-step methodology yielded an RDF/Turtle ontology containing 11 regulatory requirements across civilian and defence domains, derived from 7 regulatory sources. The framework was successfully deployed in Stardog Cloud and validated through SPARQL queries that retrieved, compared, and analyzed requirements across domains.

2) *Structural parallelism enables systematic cross-domain analysis.* The consistent organization of requirements by lifecycle phase (design/development vs. usage/maintenance) and category (privacy-related vs. security-related) in both domains enables direct comparison. SPARQL queries can automatically pair civilian and defence requirements by structural position (e.g., CIV-DESIGN-R1 ↔ DEF-DESIGN-R1), revealing where requirements align and where they diverge.

3) *Security frameworks diverge by domain.* Security-related requirements (R2, R4) are sourced from different regulatory frameworks depending on deployment context. Civilian deployments rely on GDPR Article 32 and NIS2 Directive; defence deployments require compliance with NATO Information Assurance Policy, NATO Security Policy, and NATO Cyber Defence Policy. The framework captures this divergence through domain-specific derivedFrom relationships.

4) *Defence compliance is additive, not substitutive.* For the systems that are deployed in both civilian and defence domains, defence deployments must satisfy all applicable civilian requirements plus NATO-specific obligations. The ontological framework supports this cumulative model through explicit domain tagging (appliesToDomain), enabling queries that retrieve the set of requirements for a given deployment context.

5) *The framework supports practical compliance planning.* The validated ontology yields actionable guidance for system developers, enabling them to trace system architecture to legal obligations, identifying required certification, etc. More details are provided in Table 17.

5.7. Case study summary

The espionage detection system case study demonstrates that the proposed ontology-based methodology can accomplish the following:

- Represent regulatory requirements from multiple, overlapping legal frameworks.
- Organize requirements by lifecycle phase and category to support privacy-by-design and security-by-design.
- Enable systematic comparison across civilian and defence deployment contexts.
- Support hierarchical decomposition for precise traceability to regulatory articles.
- Generate practical compliance guidance for dual-use technology developers.

These capabilities address the limitations identified in Section 2.5: the framework explicitly considers multiple regulations and provides a structured approach for dual-use technology compliance that maintains traceability between legal requirements and technical specifications.

6. Summary and future work

This section provides the summary of the paper, highlighting major contributions and limitations, and directions for future work.

6.1. Summary

This paper addressed the challenge of specifying regulatory compliance requirements for AI systems operating in dual-use contexts, where civilian and defence deployment scenarios impose overlapping, divergent, and sometimes conflicting obligations. As the EU digital regulatory ecosystem expands – encompassing the AI Act, GDPR, NIS2 Directive, and sector-specific frameworks – and as dual-use technologies increasingly operate at the intersection of civilian and defence domains, systematic approaches to compliance specification become essential.

We proposed a six-step ontology-based methodology that enables: Structured scoping of applicable regulations across civilian and defence domains (step 1); Conceptual modelling using lifecycle phases and requirement categories (step 2); Formal ontology design and implementation in RDF/Turtle (step 3); Systematic population of domain-specific requirements (step 4); Deployment in knowledge graph platforms for visualisation and exploration (step 5); SPARQL-based validation and cross-domain analysis (step 6).

The methodology was validated through a case study applying the framework to the espionage detection system – an AI-powered dual-use technology designed to detect unauthorised photographing of classified information.

6.2. Contributions

This paper makes several contributions to the emerging literature on AI governance, regulatory technology (RegTech), and ontology-based requirements engineering, described in Table 18. It proposes a regulatory compliance framework for dual-use technologies, which integrates privacy-by-design and security-by-design principles. Proposed framework has been validated using a case study. Furthermore, practical guidance for developers has been provided.

6.3. Scope of automation

The proposed framework automates the structuring, alignment, and retrieval of regulatory compliance requirements rather than the legal interpretation process itself. Regulatory texts are initially translated into ontological representations through expert-driven abstraction. Once formalised, however, compliance applicability, cross-domain comparison, lifecycle allocation, and requirement decomposition become computationally executable. The automation therefore operates at the level of compliance configuration and validation, not at the level of autonomous legal reasoning.

6.4. Limitations

The current work has several limitations that should be acknowledged:

- *Scope of regulatory coverage.* The ontology captures a subset of applicable regulations (GDPR, NIS2, NATO policies) and only a very limited number of selected requirements for illustrative purposes. Additional frameworks – including the EU AI Act risk classification, Cyber Resilience Act, sector-specific regulations (e.g., DORA for financial services), and national transposition laws – were identified but not fully modelled. Additionally, the broader framework of international law, including international humanitarian and human rights law, was referenced but not fully analysed to be operationally included in the proposed framework.

- *Granularity of requirements.* Requirements were specified at a summary level (R1–R4 per domain) with selective decomposition (DEF-R4 sub-requirements). A production compliance system would require more granular specification at the article and paragraph level across all applicable regulations.

- *Single case study.* Validation was conducted using one dual-use technology (espionage detection system) and only on a very high level without going into the details of factual deployment in defence scenarios. Generalizability to other dual-use AI systems – such as autonomous vehicles, medical AI, or critical infrastructure monitoring – requires additional case studies.

- *Static representation.* The current ontology captures requirements at a point in time. Regulatory frameworks evolve; the AI Act implementing standards are still under development by CEN-CENELEC JTC 21. Mechanisms for ontology versioning and update propagation were not addressed.

- *Expert validation.* While the framework was validated through SPARQL queries demonstrating structural completeness and cross-domain comparison capability, formal validation with legal experts, compliance officers, and certification bodies was not conducted within the scope of this study.

6.5. Future work

Building on the contributions and addressing current limitations, the following directions for future research are identified:

- *Extended Regulatory Coverage and Granularity*: Expand the ontology to include article-level mappings for the EU CRA and CEN-CENELEC JTC 21 standards. This includes integrating co-assurance frameworks that harmonize safety standards (e.g., IEC 61508) with EU-wide legislative requirements (Ndiaye et al., 2026).

- *Automated Compliance with SHACL and Reasoners*: Further develop SHACL (Shapes Constraint Language) rules to validate system specifications. Future iterations will leverage logical reasoners to automate the detection of compliance gaps, particularly for high-risk infrastructure components like SDN and cloud-based 5G subsystems (Paskauskas, 2025).

- *Integration of Decentralized and Privacy-Preserving Technologies*: Investigate the use of Blockchain Hyperledger and Federated Learning (FL) to enable secure, privacy-preserving financial and security compliance (Khan et al., 2025). This addresses the challenge of data silos while ensuring a consensus-driven feedback loop for security models.

- *Evolution Toward 6G and Edge Intelligence*: Adapt the framework for 6G networks by incorporating Smart Service-based Security Architectures (ES3A) (Duan et al., 2025). Research will focus on edge–cloud IoT frameworks that perform real-time, privacy-preserving analytics (Xia, 2025).

- *Semantic Interoperability in Multi-Cloud Environments*: Enhance the ontology to support federated-fog computing. This includes developing semantic layers that mitigate latency and optimize computational resources for real-time IoT applications across diverse organizations (Huaranga-Junco et al., 2024).

- *Generative AI and Large Language Model Support*: Explore the use of LLMs and Retrieval-Augmented Generation (RAG) to assist in extracting requirements from emerging 2025/2026 regulations. Future work will evaluate if these systems can overcome the lack of true semantic understanding in automated governance contexts (Wang, 2025).

6.6. Concluding remarks

The growing complexity of the EU digital regulatory ecosystem, combined with the increasing prevalence of dual-use AI technologies, demands systematic approaches to compliance specification. Manual interpretation of overlapping regulations is error-prone, resource-intensive, and difficult to maintain as frameworks evolve.

Ontology-based methods offer a pathway to a structured, machine-readable compliance specification that supports traceability, cross-domain analysis, and automated validation. By representing regulatory requirements in a formal knowledge graph, organizations can navigate the intersection of civilian and defence obligations more systematically – reducing compliance risk while enabling innovation in dual-use technology development.

The methodology and framework presented in this paper represent an initial step toward automated regulatory compliance for AI systems. As the EU AI Act implementation progresses and technical standards mature, ontological approaches will become increasingly valuable for operationalising abstract legal norms within complex AI systems.

Software availability

Source code available from: https://github.com/SecOntologyLab/regulatory-compliance-ontology

License: MIT.

Archived software available from: https://doi.org/10.5281/zenodo.18763491 (Paskauskas, 2026)

The knowledge graph was hosted and SPARQL queries were executed using Stardog Cloud (https://www.stardog.com/), latest stable release accessed February 2026.

Interactive visualisations were developed using D3.js (https://d3js.org/).

Ethics and consent

Ethical approval and consent were not required.

Data availability

Underlying data

The data that supports the findings of this study are openly available in repository: https://doi.org/10.5281/zenodo.18763491 (Paskauskas, 2026)

Data is available under the terms of the Creative Commons Attribution 4.0 International.

Repository contains ontology files (TTL), SPARQL query library, methodology documentation, and interactive visualisations required to reproduce the reasoning patterns described in the manuscript.

Files are provided in open, non-proprietary formats (Turtle/TTL, SPARQL, HTML) conforming to W3C Semantic Web standards. All artefacts required to reproduce the ontological reasoning patterns described in the manuscript are included in the repository.

Extended data

No additional extended data beyond the repository contents.

Acknowledgements

In preparing this manuscript, we acknowledge the limited use of a generative AI tool (*Microsoft Copilot – no version number provided by the tool*) solely to improve readability, grammar, and language clarity, and to troubleshoot citation manager formatting and bibliography consistency issues (e.g., Zotero). All AI-assisted use occurred under direct human oversight.

No generative AI tool was used to generate scientific content, perform conceptual analysis, interpret evidence, determine inclusion or exclusion decisions, or contribute to the theoretical, methodological, or empirical arguments of the manuscript.

All AI-assisted text was reviewed and edited using a human-in-the-loop approach, and the authors take full responsibility for the manuscript's content, reasoning, and conclusions.

References

Abualhaija S, Ceci M, Briand L: **Legal Requirements Analysis: A Regulatory Compliance Perspective.** Ferrari A, Ginde G, editors. *Handbook on Natural Language Processing for Requirements Engineering.* Springer Nature Switzerland; 2025; 209–242.
Publisher Full Text

Amaral G, Guizzardi R, Guizzardi G, *et al.*: **Trustworthiness Requirements: The Pix Case Study.** Ghose A, Horkoff J, Silva Souza VE, *et al.*: *Conceptual Modeling.* Springer International Publishing; 2021; **13011**: 257–267.
Publisher Full Text

Brabra H, Mtibaa A, Sliman L, *et al.*: **Semantic Web Technologies in Cloud Computing: A Systematic Literature Review.** *2016 IEEE International Conference on Services Computing (SCC).* 2016; 744–751.
Publisher Full Text

Cichocki R: **Artificial Intelligence in Maritime Cybersecurity: Dual-Use Applications for Defense and Offense in the Age of Digital Seas.** *TransNav Int J Mar Navig Saf Sea Transp.* 2025; **19**(2): 617–623.
Publisher Full Text

Court of Justice of the European Union: **CJEU Case C-511/18. La Quadrature du Net and Others.** Court of Justice of the European Union; 2020.
Reference Source

Court of Justice of the European Union: *CJEU Case C-817/19. Ligue des droits humains.* Court of Justice of the European Union; 2022.
Reference Source

Court of Justice of the European Union: *CJEU Case C-634/21. SCHUFA Holding AG.* Court of Justice of the European Union; 2023.
Reference Source

Court of Justice of the European Union: *CJEU Case C-203/22. Dun & Bradstreet Austria GmbH.* Court of Justice of the European Union; 2025.
Reference Source

Cuypers M: *Artificial Intelligence and International Humanitarian Law: Brothers in Arms or Rather the Opposite?.* KU Leuven: Blog; 2023.
Reference Source

Duan Z, Nan G, Li R, *et al.*: **Agile Orchestration at Will: An Entire Smart Service-Based Security Architecture Towards 6G.** *IEEE Wirel. Commun.* 2025; **32**: 87–94.
Publisher Full Text

Elahidoost P: **Towards a Tool Supported Approach for Regulatory Requirements Engineering.** *2024 IEEE 32nd International Requirements Engineering Conference (RE).* 2024; 520–524.
Publisher Full Text

ENISA: *Good practices for security of IoT: secure software development lifecycle.* European Network and Information Security Agency; 2019.
Publisher Full Text

European Commission: *Charter of Fundamental Rights of the European Union. 2012/C 326/02.* Brussels: European Commission; 2012.
Reference Source

European Commission: *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).* Brussels: European Commission; 2016a.
Reference Source

European Commission: *Directive (EU) 2016/680 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data by competent authorities for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and on the free movement of such data, and repealing Council Framework Decision 2008/977/JHA.* Brussels: European Commission; 2016b.
Reference Source

European Commission: *Directive (EU) 2019/790 of the European Parliament and of the Council of 17 April 2019 on copyright and related rights in the Digital Single Market and amending Directives 96/9/EC and 2001/29/EC.* Brussels: European Commission; 2019.
Reference Source

European Commission: *Consolidated text: Regulation (EU) 2021/821 of the European Parliament and of the Council of 20 May 2021 setting up a Union regime for the control of exports, brokering, technical assistance, transit and transfer of dual-use items (recast).* Brussels: European Commission; 2021a.
Reference Source

European Commission: *Action Plan on synergies between civil, defence and space industries, COM(2021) 70 final.* Brussels: European Commission; 2021b.
Reference Source

European Commission: *Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive).* Brussels: European Commission; 2022a.
Reference Source

European Commission: *Roadmap on critical technologies for security and defence, COM(2022) 61 final.* Brussels: European Commission; 2022b.
Reference Source

European Commission: *Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011.* Brussels: European Commission; 2022c.
Reference Source

European Commission: *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).* Brussels: European Commission; 2024a.
Reference Source

European Commission: *Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013 and (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act).* Brussels: European Commission; 2024b.
Reference Source

European Commission: *Commission Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act), C(2025) 5053 final.* Brussels: European Commission; 2025a.
Reference Source

European Commission: *Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act).* Brussels: European Commission; 2025b.
Reference Source .

European Commission: **AI Act.** *Shaping Europe's digital future Directorate-General for Communications Networks, Content and Technology.* Brussels: European Commission; n.d.
Reference Source

European Court of Human Rights: **European Convention on Human Rights.** 1950.
Reference Source

European Parliament: **Joint Communication on a new European Defence Industrial Strategy: Achieving EU readiness through a responsive and resilient European Defence Industry.** 2024.
Reference Source

European Parliament: *Interplay between the AI Act and the EU digital legislative framework.* Policy Department for Transformation, Innovation and Health; 2025.
Reference Source

Gómez-Romero J, Patricio MA, García J, *et al.*: **Ontology-based context representation and reasoning for object tracking and scene interpretation in video.** *Expert Syst. Appl.* 2011; **38**(6): 7494–7510.
Publisher Full Text

Guizzardi R, Amaral G, Guizzardi G, *et al.*: **An ontology-based approach to engineering ethicality requirements.** *Softw. Syst. Model.* 2023; **22**(6): 1897–1923.
Publisher Full Text

Hernandez J, Golpayegani D, Lewis D: **Ontology-Based Approach for Mapping Concepts and Requirements from Regulations and Standards: The Case of the EU AI Act and International Standards.** Savelka J, Harasta J, Novotna T, *et al.*, editors. *Frontiers in Artificial Intelligence and Applications.* IOS Press; 2024.
Publisher Full Text

Hosseini AM, Kastner W, Sauter T: **Ontology Framework Supporting Security-By-Design of Industrial Control Systems.** *IEEE Trans. Industr. Inform.* 2025; **21**(9): 7188–7197.
Publisher Full Text

Huaranga-Junco E, González-Gerpe S, Castillo-Cara M, *et al.*: **From cloud and fog computing to federated-fog computing: A comparative analysis of computational resources in real-time IoT applications based on semantic interoperability.** *Futur. Gener. Comput. Syst.* 2024; **159**: 134–150.
Publisher Full Text

ICRC: *Artificial intelligence and machine learning in armed conflict: A human-centred approach. ICRC Position Paper.* International Committee of the Red Cross; 2021.
Reference Source

ICRC: *What you need to know about artificial intelligence in armed conflict.* International Committee of the Red Cross; 2023.
Reference Source

ICRC: *The Geneva Conventions and their Commentaries.* International Committee of the Red Cross; 2026a.
Reference Source

ICRC: *Customary IHL.* International Committee of the Red Cross; 2026b.
Reference Source

Imperial JM, Jones MD, Tayyar Madabushi H: **Standardizing Intelligence: Aligning Generative AI for Regulatory and Operational Compliance.** *SSRN.* 2025.
Publisher Full Text

IEC: *Functional safety of electrical/electronic/programmable electronic safety-related systems (IEC 61508:2010, Edition 2.0).* International Electrotechnical Commission; 2010.

IEC: *Industrial communication networks–Network and system security—Part 4–1: Secure product development lifecycle requirements (IEC 62443-4-1).* International Electrotechnical Commission; 2018.

ISO/IEC: *ISO/IEC 22989:2022: Information technology — Artificial intelligence — Artificial intelligence concepts and terminology.* International Organization for Standardization; 2022.

Juric M: **Legal regulation on the use of artificial intelligence for national security purposes in Europe.** *Eur Integr Stud.* 2024; **20**(2): 107–136.
Publisher Full Text

Khan AA, Alsufyani A, Alsufyani N, *et al.*: **BAML: A decentralized approach to secure, privacy-preserving financial compliance for enhancing anti-money laundering with blockchain hyperledger and federated learning.** *Peer-to-Peer Netw Appl.* 2025; **18**(5): 270.
Publisher Full Text

Kilian R, Jäck L, Ebel D: **European AI Standards – Technical Standardisation and Implementation Challenges under the EU AI Act.** *European Journal of Risk Regulation.* 2025; **16**(3): 1038–1062.
Publisher Full Text

Kosenkov O, Elahidoost P, Gorschek T, *et al.*: **Systematic mapping study on requirements engineering for regulatory compliance of software systems.** *Inf. Softw. Technol.* 2025; **178**: 107622.
Publisher Full Text

Kremidas-Courtney C: *From Dual-Use to Omni-Use: Securing Europe's Future.* European Policy Centre; 2025, October 1.
Reference Source

Lietuvos Respublikos Seimas: *Lietuvos Respublikos asmens duomenų, tvarkomų nusikalstamų veikų prevencijos, tyrimo, atskleidimo ar baudžiamojo persekiojimo už jas, bausmių vykdymo arba nacionalinio saugumo ar gynybos tikslais, teisinės apsaugos įstatymas. 2011 April 21. Nr. XI-1336.* Vilnius: Lietuvos Respublikos Seimas; 2011.
Reference Source

Lietuvos Respublikos Seimas: *Lietuvos Respublikos Kibernetinio Sugumo Istatymas. 2024 December 11. Nr. XII-1428.* Vilnius: Lietuvos Respublikos Seimas; 2024.
Reference Source

NATO: *NATO Core Metadata Specification (NCMS) - ADatP-5636 Edition A.* NATO Standardization Office; 2022.
Reference Source

NATO: *Summary of NATO's revised Artificial Intelligence (AI) strategy.* North Atlantic Treaty Organization; 2024a, July 10.
Reference Source

NATO: *NATO'S Digital Transformation Implementation Strategy.* Brussels: North Atlantic Treaty Organization; 2024b, October 17.
Reference Source

NATO: *Data Strategy for the Alliance.* Brussels: North Atlantic Treaty Organization; 2025, May 5.
Reference Source

NATO's Strategic Warfare Development Command: *Digital Transformation.* Brussels: NATO Allied Command Transformation; n.d.
Reference Source

Ndiaye NG, Waedt K, Dejon N, *et al.*: **Safety and Cybersecurity Under Emerging EU Legislation for Industry: A Use-Case Driven Perspective.** Coppens B, Volckaert B, Naessens V, *et al.*, editors. *Availability, Reliability and Security, Ares 2025, Pt I.* Springer International Publishing Ag; 2026; **15994**: 304–321.
Publisher Full Text

OECD: *Intellectual property issues in artificial intelligence trained on scraped data.* OECD Artificial Intelligence Papers; 2025. February 2025, No. 33.
Reference Source

Orlando A, Santoro M: **A semantic approach to understanding GDPR fines: From text to compliance insights.** *Computer Law & Security Review.* 2025; **59**: 106187.
Publisher Full Text

Paskauskas RA: **A Preliminary Ontology for 5G Network Resilience: Hybrid Threats, Risk Reduction, Compliance.** *2025 Ieee International Conference on Cyber Security and Resilience.* Csr; 2025; 496–503.
Publisher Full Text

Paskauskas RA: SecOntologyLab/regulatory-compliance-ontology: Initial Public Release (Zenodo Archival for ORE Submission) (v1.0.0). [Data set]. *Zenodo.* 2026.
Publisher Full Text

Pn AR, Ramachandra R, Rao KS, *et al.*: **Straight Through Gumbel Softmax Estimator based Bimodal Neural Architecture Search for Audio-Visual Deepfake Detection.** *2024 IEEE International Joint Conference on Biometrics (IJCB).* 2024; 1–9.
Publisher Full Text

Powell R: *The EU AI Act: National Security Implications.* CETaS Explainers; 2024. August 2024.
Reference Source

RAND: *General-Purpose Artificial Intelligence (GPAI) Models and GPAI Models with Systemic Risk: Classification and Requirements for Providers.* RAND Corporation; 2024.
Publisher Full Text

Topalli N, Badii A: **A User-Centric Context-Aware Framework for Real-Time Optimisation of Multimedia Data Privacy Protection, and Information Retention Within Multimodal AI Systems.** *Sensors.* 2025; **25**(19): 6105.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Vogiatzoglou P: **The AI Act National Security Exception.** 2024.
Publisher Full Text

Wang L: **Can machines think beyond words? A critique of AI's meaning-production process.** *AI & Society.* 2025; **41**.
Publisher Full Text

Xia L: **Enabling University Mental Health Monitoring Through 6G Edge-Cloud IoT Environmental Perception Frameworks.** *IEEE Communications Standards Magazine.* 2025; **1–8**: 1–8.
Publisher Full Text

Open Peer Review

Current Peer Review Status: ? ? ✓ ?

Version 1

Reviewer Report 27 May 2026

https://doi.org/10.21956/openreseurope.25029.r73927

© **2026 Dalal A.** This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Aryendra Dalal
Middle Georgia State University, Macon, Georgia, USA

This paper proposes a six-step ontology-based methodology for specifying regulatory compliance requirements for dual-use AI systems, implemented as an RDF/Turtle knowledge graph and validated through a case study involving an AI-powered espionage detection system. The regulatory scope spans EU instruments (AI Act, GDPR, NIS2) and NATO frameworks, with SPARQL queries enabling cross-domain compliance comparison.

Is the work clearly and accurately presented and does it cite the current literature? — Partly The paper is well-written overall and references recent literature through 2026. However, the academic positioning needs tightening. The related work section lists prior ontological frameworks as bullet points without meaningfully engaging with how they fall short or how this work specifically improves upon them. That gap weakens the novelty claim considerably.

Is the study design appropriate and does the work have academic merit? — Yes Using ontology-based requirements engineering for dual-use compliance is a reasonable and timely choice. The layered civilian-plus-defence compliance model captures something genuinely important that simpler approaches miss. The six-step structure is coherent, and the espionage detection system is a credible test case for the methodology.

Are sufficient details of methods and analysis provided to allow replication by others? — Partly The SPARQL queries are included and the GitHub/Zenodo links are provided, which is commendable. That said, the ontology class hierarchy is not visually presented in a way that would allow an independent team to reconstruct it without difficulty. The process by which regulatory text is translated into ontological representations is described only at a high level — this expert-driven abstraction step is arguably the most consequential part of the methodology and deserves considerably more transparency. A worked example mapping a specific regulatory article to an ontological class would help significantly.

Statistical Analysis — Not Applicable

Are all source data underlying the results available? — Partly The repository exists and is cited, but the full query set used to generate the tables in Section 4 is not completely documented within the manuscript itself. Readers should not have to navigate externally to verify core results.

Are the conclusions drawn adequately supported by the results? — Partly The four validated capabilities (dual-use coverage, lifecycle organization, cross-domain comparison, hierarchical decomposition) are all reported as successfully confirmed, which reads somewhat circular — the framework is introduced, populated with requirements the authors themselves defined, then queried to confirm those same requirements are present. This is structural completeness testing, not genuine validation against an independent benchmark or real deployment scenario. Formal engagement with legal or compliance experts would substantially strengthen the validation claim. The paper also acknowledges static representation as a limitation of existing approaches but then lists it again as a limitation of its own system, without addressing that tension.

Minor Points Two typographical errors need correcting: "focuses number of" on page 9 and "as included in. the table" on page 21. Several sections would read better as continuous prose rather than bulleted lists, particularly the future work section.

Is the work clearly and accurately presented and does it cite the current literature?
Partly

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Partly

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Cybersecurity

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Reviewer Report 25 May 2026

https://doi.org/10.21956/openreseurope.25029.r72142

© 2026 Dragomir-Constantin F. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Florentina-Loredana Dragomir-Constantin
National Defense University Carol I, Bucharest, Romania

The article “Automated regulatory compliance for AI systems in the security domain: The case of dual-use deployment” addresses an important and timely research topic situated at the intersection of artificial intelligence, information systems, national security, and dual-use technologies. The paper proposes an ontology-based methodology for specifying requirements applicable to AI systems that may be developed or deployed across both civilian and defence-related contexts. The material evaluated presents a coherent conceptual and methodological framework, supported by a case study focused on an AI-powered espionage detection system. The main contribution of the article consists in its attempt to formalize complex legal, technical, and security-related requirements into machine-readable representations through knowledge graphs, RDF/Turtle modelling, SPARQL-based analysis, and semantic validation.
The paper is relevant both academically and practically, as it addresses the growing need for structured and traceable approaches to AI systems operating in sensitive and security-relevant environments.
Q1: Is the work clearly and accurately presented and does it cite the current literature?
Response: Yes.
The work is clearly and accurately presented. The article is logically structured and introduces the reader to the complexity of AI systems that may operate across civilian and defence contexts. The discussion of overlapping obligations, system lifecycle phases, and dual-use deployment scenarios is generally well organized and easy to follow. The paper cites current and relevant literature, particularly in relation to the EU AI Act, digital regulation, AI governance, dual-use technologies, and ontology-based approaches. The references support the overall argument and demonstrate that the authors are engaging with a contemporary and evolving research area. For future development, the paper could further strengthen its academic framing by positioning the proposed methodology more explicitly in relation to existing debates on dual-use AI, national security applications, and ontology-based requirements engineering.
Q2: Is the study design appropriate and does the work have academic merit?
Response: Yes.
The study design is appropriate for the objectives of the article. The ontology-based methodology is suitable for organizing complex and overlapping requirements and for making them more transparent, traceable, and reusable. The six-step methodological structure provides a clear pathway from regulatory scoping to ontology development, knowledge graph deployment, and SPARQL-based validation. The work has academic merit because it addresses a relevant interdisciplinary problem that connects artificial intelligence, information systems, national security, security governance, and dual-use technologies. The article contributes to an emerging area of research by proposing a structured approach for translating complex requirements into system-level specifications. The case study also adds value by illustrating how the proposed framework may be applied to a concrete AI system operating in a sensitive security-related environment.
Q3: Are sufficient details of methods and analysis provided to allow replication by others?
Response: Partly.

The article provides a clear overview of the proposed methodology and presents the main technical elements used in the framework, including RDF/Turtle, knowledge graphs, SPARQL queries, and semantic validation. These elements provide a useful basis for understanding how the approach may be implemented. However, full replication would benefit from additional methodological detail. In particular, the authors could further explain how legal, policy, and security-related texts are interpreted and transformed into ontology classes, properties, relations, and constraints. More examples of mappings between requirements and ontology elements would also improve transparency. For future versions of the work, it would be useful to include a more detailed explanation of the role of experts in the validation process, the criteria used for assessing the ontology, and the steps required for another researcher to reproduce the knowledge graph and validation results.
Q4: If applicable, is the statistical analysis and its interpretation appropriate?
Response: Not applicable.
The article does not rely on statistical analysis. It is primarily a conceptual, methodological, and technical contribution based on ontology modelling and case-study validation. Therefore, this criterion is not applicable.
Q5: Are all the source data underlying the results available to ensure full reproducibility?
Response: Yes.
The article provides sufficient indications regarding the availability of the underlying resources, including ontology-related files and SPARQL queries. This supports transparency and contributes positively to the reproducibility of the study. For future improvement, the authors may consider offering more detailed documentation of the ontology files, query library, validation process, and knowledge graph structure. This would make the proposed framework more accessible to researchers and practitioners interested in applying or extending the methodology.
Q6: Are the conclusions drawn adequately supported by the results?
Response: Yes.
The conclusions are adequately supported by the methodology and the case study presented in the article. The authors show that an ontology-based approach can support the structuring, mapping, and analysis of requirements for AI systems deployed in dual-use and security-related contexts. The conclusions are consistent with the results of the study and reflect the potential value of the proposed framework. For future development, the article could further strengthen its conclusions by discussing in more detail the practical limits of the approach, including scalability, ontology maintenance, updates following legal or policy changes, and possible conflicts between civilian and defence-related requirements. The article is timely, relevant, and well aligned with current challenges in artificial intelligence, information systems, national security, and dual-use technology governance. The proposed ontology-based methodology is a valuable contribution and has clear potential for further academic and practical development.
For future improvement, the authors may consider the following points:
The academic contribution could be clarified further by explaining more explicitly how the proposed framework differs from and improves upon existing ontology-based or requirements engineering approaches.
The discussion of dual-use AI could be expanded by engaging more deeply with the academic literature on national security, defence applications, and the relationship between civilian and military uses of AI systems.
The methodology could be strengthened by providing more detail on how textual requirements are interpreted, selected, categorized, and translated into ontology components.
The validation section could be improved by specifying clearer criteria, thresholds, or indicators for evaluating the functionality and reliability of the proposed framework.

The article could include a more detailed discussion of scalability, especially regarding the future updating of the ontology when regulations, standards, or security policies change.
The presentation could be refined by replacing some bullet-point sections with continuous analytical discussion and by correcting minor wording issues.
These recommendations are offered as constructive suggestions for future refinement. They do not diminish the overall quality, originality, or relevance of the work.
Final Recommendation
Approved.

The manuscript represents an important and timely contribution to the study of AI systems in the security domain, with particular relevance for information systems, national security, and dual-use technologies. The material evaluated demonstrates methodological coherence, academic relevance, and practical potential. Although some aspects could be developed further in future versions, particularly the methodological details, validation process, and academic positioning, these are recommendations for improvement rather than reasons to withhold approval. Overall, the article is well conceived, relevant, and suitable for approval.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Information systems, artificial intelligence, national security, dual-use technologies, machine learning, decision theory.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Author Response 26 May 2026

Giedre Sabaliauskaite

We thank the reviewer for the careful assessment of our manuscript and for the constructive recommendations provided. We appreciate the approval of the article and acknowledge the suggestions regarding the academic positioning, dual-use AI framing, methodological detail, validation process, scalability, ontology maintenance, and presentation. We will take these points into account in preparing the revised version of the manuscript, particularly by clarifying the methodology, strengthening the discussion of validation and limitations, and improving the presentation where appropriate.

Competing Interests: No competing interests were disclosed.

Reviewer Report 14 May 2026

https://doi.org/10.21956/openreseurope.25029.r72145

© 2026 Schmid S. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Stefka Schmid
Aalto University, Helsinki, Finland

The work titled “Automated regulatory compliance for AI systems in the security domain: The case of dual-use deployment” contributes to an important and timely empirical topic by presenting an ontology for automated compliance of dual-use AI systems. At the same time, the work should connect more to relevant academic debates and situate their case better in the real-life environment of dual-use AI automated compliance. While the development of the ontology is overall understandable, it should be ensured that the approach is systematic and also entails a reflection on the boundaries of the proposed artifact.

Introduction and Framing
The introduction clearly lays out the empirical relevance of the author’s work. In this regard, it would be useful to open with a concrete example of a dual-use system that may need to comply not only with the AI Act but also with defence standards such as NIAP and NSP, as this would immediately ground the paper's subject matter. The academic contribution also needs to be clarified early on, as it is currently not clear what other approaches look like or how this work positions itself against them. While the research gap is laid out, the EU case needs to be better plausibilized; the reader should understand from the outset why it is necessary to focus specifically on this issue. Finally, the conclusion should circle back to this framing with a short note on the critical nature of automated regulatory compliance, referencing the state of the art and explaining why it is useful and promising to focus on this issue despite challenges such as those identified in the paper.

Academic Contribution: Situating in Scholarship
While the empirical relevance of the work is clear, the paper needs to more explicitly address questions of feasibility, potential pitfalls, and the current state of the scholarly debate. The

academic contribution requires stronger clarification, particularly regarding how this work differs from and builds upon existing approaches. A significant gap currently lies in the treatment of ontological frameworks for automating regulatory compliance: these are presently only listed as bullet points, with no substantive engagement with how the authors assess this body of work or how it connects to automated compliance tools for both civilian and military contexts. This is central to the paper's academic relevance and must be situated more carefully within the relevant scholarships. Related to this, the paper should address how well existing frameworks have performed, what their level of technological readiness is, and how differences between civilian, military, and dual-use automated compliance systems are reflected in the literature. The current limitations of other approaches should be explained explicitly, with a clear account of how this work specifically addresses them. The claim that existing AI governance compliance approaches rely on static documentation needs to be supported with references, and the added value of automated validation should be illustrated accordingly. Further, static representation is criticized as a limitation of existing approaches, yet it is subsequently named as a limitation of the presented work.

Dual-Use Concept and AI Classification
The general definition of dual use is acceptable, but should be accompanied by references to the scholarly debate on this concept. A strength of the paper is its emphasis on AI being inherently dual-use. This is an important intervention, given that much existing scholarship foregrounds the "general purpose" character of AI in a way that treats it as neutral and thereby overlooks the different regulatory regimes that may apply to security-critical domains. However, when the paper criticizes the "strict binary logic of dual use" as something that slows down innovation, this requires clarification: given that the paper has already argued that AI is inherently dual-use, the critique here seems to be directed more at the institutional separation of civilian and military spheres than at the fact that systems can serve both purposes. Even though the link between the high-risk classification and dual-use systems is understandable, the relationship between dual-use and high-risk AI systems also needs to be elaborated, particularly to justify why high-risk AI is introduced in section 2.3 and how it connects to the dual-use issue. The paper should additionally engage with national legislation that mandates the separation of civilian and military spheres (e.g. in Germany). When the espionage detection system is introduced as an example, its dual-use character should be justified immediately by explaining the range of contexts in which such a system might be deployed. Comparable systems that exist should be referenced, given that the example system apparently does not.

EU Regulatory Framework
The EU regulatory debate and its dynamics should also be situated in relation to national decision-making processes and dynamics. In discussing national data protection, examples beyond the mere national implementation of GDPR should be included. The claim that "in the legal scholarship, it is emerging that real applicability of the exemptions included under certain EU regulations such as the GDPR or AI Act is uncommon in practice, one of the reasons being that governments and militaries often tend to rely on private companies to develop AI" needs to be unpacked and explained more clearly. The list of obligations for high-risk AI systems should be properly referenced, as it appears not to be the authors' own formulation. The use of CJEU rulings as examples should also be better justified through written explanation rather than bullet points. Finally, the paper raises the important point that it requires "extreme cautiousness to capture all possible scenarios so that no AI system component slips through the regulatory gaps" but does not explain how this is typically achieved, i.e., whether humans are successfully doing this

currently, or whether this has become a more pressing concern in light of the EU's new dual-use strategy and what specifically has changed in that context.

International Humanitarian Law
The claim that the Geneva Conventions of 1949 and their Additional Protocols lay down which actions can be taken during warfare, including when AI is used, requires more careful framing. While this is a reasonable choice, it should be acknowledged that AI is not specifically addressed in the Geneva Conventions; their categories are broad and, where they do not deal with specific weapon systems, they are technology-agnostic. The paper should also engage with the UN Group of Governmental Experts (GGE) and perhaps the REAIM summits. While the ICRC is a very relevant organization, it is a non-governmental organization. More broadly, the paper should consider that not all the regulatory instruments it lists carry equal normative weight. Finally, the factors identified as relevant, such as intended end users, data categories, primary purpose, and geographic location, should be more clearly connected to what has been established in the preceding discussion of IHL and related frameworks, so that readers understand how these factors were derived and why they matter.

Defence Standards and Compliance
When referencing the NATO Information Assurance Policy and similar instruments in section 2.4, proper citations should be included. The claim that "compliance in defence deployments is, therefore, mainly achieved through accreditation processes and security authorization rather than conformity assessment under EU market legislation" also needs to be explained more clearly, as it is not immediately obvious what this distinction entails in practice and why it matters for the paper's argument.

AI Life Cycle
With regard to different life cycle phases, the paper could engage with Bode et al.'s recommendations regarding human interaction with decision support systems across different phases. The lifecycle phase organization is a notable strength of the paper, and the issue of dual regulatory coverage across life cycle phases is also important but life cycle phases should be clarified specifically regarding AI in section 2.1. It should also be clarified whether the activities and outputs described in the steps are meant to be followed in chronological order; if so, they should be numbered accordingly.

Methodology: Ontology Development
The paper would benefit considerably from grounding its approach in a recognized methodological framework, for example from information systems research or an established ontology development method. Without this, it is difficult to assess how the work proceeds and to what extent it builds on existing methods. In particular, it is unclear whether steps 1 to 6 are derived from such a method. It should also be clarified whether the first step of defining the technology is intended solely for dual-use contexts, or whether the artifact is meant to be applicable more broadly. The inclusion of domain experts is welcome, but it is currently unclear what kind of participatory approach is being used, how experts are to be included specifically, and the artifact should be introduced as such from the outset. The description of "regulatory texts being initially translated into ontological representations through expert-driven abstraction" is a key methodological step that warrants considerably more explanation. More generally, since the data consists of written legal text that requires interpretive skill and involves variation in formulation, the paper needs to provide concrete examples of how text is analyzed in practice.

This includes a comparison with human interpretation to make a convincing case for an automated approach.

Automated Compliance and Validation
The authors acknowledge that “we cannot rely on AI models yet” but then proceed to outline steps for automating compliance without adequately plausibilizing why this endeavor is worthwhile given that caveat. The validation of the artifact needs to be made more transparent: Readers should understand how the validation was conducted and any references to “recent advancements in semantic modelling and automated compliance” that support the validation results should be properly cited. The current validation is problematic in that all capabilities are reported as validated, which risks appearing tautological, as if the system is introduced and then simply confirmed to work. It is essential that validation thresholds, metrics, and criteria are made explicit and intersubjectively verifiable, and that nuances in performance are acknowledged and discussed. The paper should explain not only whether the artifact can be validated but how, drawing on the existing knowledge base to justify the approach. It should also be explained how the automated system ensures that the legal basis, as reflected in the right column of Table 4, is checked adequately. Finally, while expert validation is a valid element, the work could be further strengthened by comparing the proposed artifact against existing approaches. Building on the work, implications are derived. However, currently, it is not clear how the implications presented in the paper are derived, whether from the validation process or from comparison with existing approaches and from which specific finding a specific implication can be formulated.

Presentation, Minor Points
There are two instances of missing words that should be corrected: on page 9, “focuses [on a] number of” and on page 21, “as included in. the table.” The use of CJEU rulings would be better handled through written paragraphs rather than bullet points. Throughout the paper, bullet points should be revised into continuous written text, particularly in the future work section, where it should be made clear which directions of future work follow from which limitations or contributions. The ordering of the final sections should be revised so that findings precede implications, as implications should build on findings. The reference to Lithuania in Table 3 requires explanation, though it could potentially be incorporated into the case study. The identification of applicable defense policy within the case study should also include national legislation. The case study summary in section 5.7 should be removed. Section 6.1 should drop its summary in favor of a focused reflection on the paper’s contribution. The statement that dual-use systems should not implement prohibited functionality may come across as trivial and should either be better justified or reframed. Finally, when maritime and critical infrastructure systems are mentioned, the reader should be clearly informed that these are application fields for dual-use systems.

Is the work clearly and accurately presented and does it cite the current literature?
Partly

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?

Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?
Yes

Are the conclusions drawn adequately supported by the results?
Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: dual-use AI, innovation research, technology and international security

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 15 May 2026

Giedre Sabaliauskaite

We thank the reviewer for the detailed assessment of our manuscript and for the constructive comments provided. We acknowledge the concerns raised regarding the framing of the academic contribution, the treatment of dual-use AI, the positioning of the EU regulatory context, the discussion of international humanitarian law and defence standards, and the methodological grounding of the ontology-development and validation process. We will address these points promptly in the revised manuscript by strengthening the introduction and scholarly positioning, clarifying the scope and boundaries of the proposed artefact, expanding the methodological explanation, and revising the discussion of validation, findings, implications, and presentation. We will provide a detailed point-by-point response with the resubmission.

Competing Interests: No competing interests were disclosed.

Reviewer Report 12 May 2026

https://doi.org/10.21956/openreseurope.25029.r73924

© 2026 Zangana H. This is an open access peer review report distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

?

Hewa Majeed Zangana

Duhok Technical College, Duhok Polytechnic University, Duhok, Iraq

1. Article Summary

The paper addresses the "compliance gap" created by the overlapping regulatory landscapes for AI systems, specifically focusing on **dual-use deployment**. It explores the intersection of civilian regulations (e.g., EU AI Act, GDPR, NIS2) and defense-specific frameworks (e.g., NATO's Principles of Responsible Use).

The authors propose an **Ontology-based Requirements Engineering (ObRE)** approach to automate the identification and mapping of these requirements. By developing a specialized ontology and utilizing semantic reasoning (SPARQL queries), the study demonstrates how to transform abstract legal norms into a machine-readable knowledge graph. The framework is validated through a case study involving an AI-driven espionage detection system, concluding that defense compliance acts as an "additive" layer to civilian requirements rather than a replacement.

2. Detailed Responses to Review Questions

Q1: Is the work clearly and accurately presented and does it cite the current literature?

- **Response: Yes.**
- **Assessment:** The paper is exceptionally well-referenced, citing 2025 and 2026 literature. It effectively contextualizes the EU AI Act within the broader security domain. The narrative flow from the "regulatory thicket" problem to the ontological solution is logical and accessible to both legal and technical audiences.

Q2: Is the study design appropriate and does the work have academic merit?

- **Response: Yes.**
- **Assessment:** Using ObRE for legal compliance is a scientifically sound choice. The methodology bridges the gap between high-level legal principles and low-level system requirements. The academic merit lies in the specific application to "dual-use" scenarios—a niche that is rapidly becoming critical as defense sectors increasingly adopt commercial AI off-the-shelf (COTS) solutions.

Q3: Are sufficient details of methods and analysis provided to allow replication by others?

- **Response: Partly.**
- **Critique:** While the six-step methodology is clear, the paper lacks a detailed visualization or schema of the developed ontology. Without seeing the class hierarchy (e.g., how LegalRequirement relates to SystemConstraint), a third party cannot replicate the semantic model accurately.
- **Constructive Feedback:** (See Section 3 for specific actions).

Q4: If applicable, is the statistical analysis and its interpretation appropriate?

- **Response: Not applicable.**
- **Assessment:** The work is a technical framework and case study; it does not rely on statistical inference.

Q5: Are all the source data underlying the results available to ensure full reproducibility?

- **Response: Partly.**
- **Critique:** The paper includes snippets of SPARQL queries, but the full RDF/Turtle file (the "source data" for an ontology paper) is not included in the manuscript or a persistent repository link within the text.
- **Constructive Feedback:** (See Section 3 for specific actions).

Q6: Are the conclusions drawn adequately supported by the results?

- **Response: Yes.**

- **Assessment:** The case study on espionage detection provides a concrete illustration of the framework. The conclusion that compliance is "omni-use" (incorporating both civilian and defense layers) is a direct result of the mapping exercise performed in Step 5.

3. Constructive Criticism & Requirements for Scientific Soundness
To ensure the article is scientifically sound and fully reproducible, the authors must address the following points:
A. Methodology & Replication (Crucial for "Partly" to "Yes")

1. **Ontology Schema:** Provide a visual diagram of the **Regulatory-Compliance Ontology**. This should show the core classes, object properties, and data properties. Without this, the "automated" nature of the work remains a "black box" to the reader.
2. **Mapping Logic:** Step 4 describes "Mapping and Intersection." The authors should provide a more detailed table or logic flow showing exactly how a requirement from the AI Act (e.g., Transparency) was technically mapped to a NATO principle (e.g., Explainability).

B. Data Availability & Reproducibility (Crucial for "Partly" to "Yes")

1. **Repository Access:** The authors must provide a link to a public repository (e.g., GitHub or Zenodo) containing the .ttl (Turtle) or .owl files for the ontology and the Knowledge Graph used in the case study.
2. **Query Documentation:** While two SPARQL queries are provided, the paper should include the full set of queries used to generate the results in Table 1 to allow for verification of the automated compliance check.

C. Technical Discussion (Enhancement)

1. **Scalability:** The authors should briefly discuss the effort required to update the ontology when regulations change (e.g., an amendment to the AI Act). This is a known challenge in RegTech and addressing it would strengthen the "Academic Merit."
2. **Conflict Resolution:** If a defense requirement (NATO) contradicts a civilian requirement (GDPR) regarding data retention, how does the ontology handle this conflict? A paragraph on conflict-handling logic in semantic reasoning would be highly valuable.

Final Recommendation
The manuscript is **technically strong and timely**. However, for it to meet the standards of full scientific reproducibility, the **ontological source files and visual schema must be made available**. Addressing the lack of a shared knowledge graph is the primary requirement for making the article scientifically sound.

Is the work clearly and accurately presented and does it cite the current literature?
Yes

Is the study design appropriate and does the work have academic merit?
Yes

Are sufficient details of methods and analysis provided to allow replication by others?
Partly

If applicable, is the statistical analysis and its interpretation appropriate?
Not applicable

Are all the source data underlying the results available to ensure full reproducibility?

Partly

Are the conclusions drawn adequately supported by the results?
Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: AI, Cybersecurity, Computer Vision, Machine Learning

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 15 May 2026

Giedre Sabaliauskaite

We thank the reviewer for their careful assessment of the manuscript and for identifying specific areas where the article can be strengthened. We acknowledge the points raised regarding methodology, reproducibility, ontology schema presentation, repository access, SPARQL query documentation, scalability, and conflict-handling logic. We will address these issues in the revised version of the manuscript by clarifying the ontology structure and mapping logic, improving the reproducibility materials, and expanding the technical discussion where appropriate. We appreciate the reviewer's constructive guidance and will respond to each point in detail in the revised submission.

Competing Interests: No competing interests were disclosed.