

Supporting Cybersecurity Risk Management for Medical Devices via the SECUMAN Ontology and Shapes

Martin DILLER^{a,1}, Anne ESSLINGER^b, Piotr GORCZYCA^a, Evi HARTIG^b,
Lia KACHOLDT^c, and Hannes STRASS^a

^a*Faculty of Computer Science, TU Dresden, Germany*

^b*Else Kröner Fresenius Center for Digital Health, Dresden, Germany*

^c*secunet Security Networks AG*

Abstract. We propose the SECUMAN ontology and shapes for representing and analysing cybersecurity risk-management documentation for medical devices. Cybersecurity risks are increasingly relevant for connected medical devices and may have direct consequences for patient safety. Current risk-management files are often maintained as semi-structured natural language text, which makes consistency checking, certification review, and reuse difficult. SECUMAN provides a formal OWL-based vocabulary for modelling security-risk context, assessment, control measures, and residual-risk evaluation, and uses SHACL constraints to check structural completeness and conformity with the intended documentation model. The ontology is aligned with VDE Spec 90025 and the related RISKMAN ontology and shapes, while extending their safety-oriented approach to concepts relevant to cybersecurity risk documentation such as threat scenarios, protection goals, attacker profiles, exposure levels, assets, and secure design arguments. SECUMAN is intended to support automated first-pass validation, traceability, and integration of cybersecurity and safety risk-management documentation.

Keywords. Risk management, Cybersecurity, OWL EL, SHACL, Medical devices

1. Introduction

Medical devices are safety-critical systems whose failures may harm patients, users, or the environment. Manufacturers must therefore provide documented justification that device-related risks have been systematically identified, evaluated, and controlled, as required in the EU by the Medical Device Regulation [1]. In practice, however, risk-management documentation is still often exchanged as semi-structured natural language. This complicates certification and maintenance: notified bodies must inspect large volumes of text, manufacturers must keep risk files consistent across device versions, and reusable parts of risk analyses and risk controls remain difficult to identify and transfer.

To address these problems for safety risk management, the RISKMAN ontology and shapes were proposed in [2]. Building on ISO 14971 and VDE SPEC 90025 [3,

¹Corresponding Author: Martin Diller, martin.diller@tu-dresden.de

4], RISKMAN encodes central concepts and structures used in safety-oriented risk-management documentation for medical devices, while its SHACL shapes express structural and semantic constraints relevant to such documentation. The approach is not intended to determine whether a concrete control measure is technically adequate; rather, it supports automated first-pass validation, improved navigation, and more systematic reuse of risk-management documentation.

For connected medical devices, including wearable monitors, implantable systems, and networked clinical equipment, safety-oriented documentation alone is insufficient. Such devices are now an integral part of modern healthcare: they enable continuous patient monitoring, real-time physiological data acquisition, and remote management, thereby improving diagnosis, treatment, and healthcare efficiency [5]. However, increasing connectivity also introduces cybersecurity risks that may affect patient safety [6]. Cybersecurity risks in healthcare are therefore not merely theoretical. In Europe, ransomware attacks on hospitals have led to cancelled operations and reduced treatment capacity, as observed in incidents in Spain and the United Kingdom [7,8], while attacks on national health systems have disrupted access to patient data and critical services [9]. In the United States, healthcare data breaches have remained at a sustained high level, with more than 700 large breaches reported annually and over 275 million patient records exposed in 2024 alone [10,11]. At the same time, the cybersecurity threat landscape continues to intensify, with ransomware accounting for more than half of attacks on the health sector [12,13]. These developments show that cybersecurity has become a critical aspect of healthcare systems, with direct implications for patient safety.

Regulatory and guidance documents explicitly reflect this connection. The EU MDR requires software-based systems to be developed in accordance with state-of-the-art principles, including information security principles (MDR Annex I, Sections 3 and 17.2). MDCG 2019-16 likewise treats safety and security risks as interdependent and requires cybersecurity risks and risk controls with potential safety impact to be considered within safety risk-management processes, and vice versa [1,14]. Similar expectations are reflected in international guidance and standards, including IMDRF principles on medical device cybersecurity, IEC 81001-5-1, and AAMI TIR57, which emphasize the need to manage cybersecurity risks in conjunction with patient-safety considerations across the device lifecycle [15,16,17].

A separate ontology for cybersecurity risk management is needed because security risks are conceptualized differently from safety risks. Safety risk management typically starts from device-related hazards and analyzes how they may lead to hazardous situations and harms. Cybersecurity risk management must instead account for threat scenarios shaped by the operating environment, product security properties, attack types, protection goals, vulnerabilities, attacker capabilities, affected assets, and impacts. These concepts are often handled by different expert groups and documented using different terminology, evaluation methods, and data structures [18]. A security ontology must therefore preserve cybersecurity-specific concepts while remaining compatible with safety-oriented documentation, so that security–safety dependencies can be made explicit where needed.

In this work, we provide such an ontology, based on a model for risk-management documentation inspired by VDE SPEC 90025 and RISKMAN, but specialised to cybersecurity using relevant standards and guidance such as AAMI TIR57, AAMI SW96, NIST SP 800-30, and IEC 81001-5-1 [17,19,20,15]. Specifically, our contributions are:

(i) an ontology for representing central concepts and structural elements of cybersecurity risk-management documentation for medical devices; (ii) SHACL shapes expressing constraints on such documentation; and (iii) a RISKMAN-compatible modelling approach that supports traceability from threat context through risk assessment to risk control and residual-risk evaluation. SECUMAN thereby complements RISKMAN by covering the security-risk domain and provides a basis for future integration between cybersecurity and safety risk-management documentation.

2. Related Work

Risk-management ontologies for medical devices. Several ontologies have been proposed for medical-device regulation, risk-related concepts, or adjacent compliance tasks, but only few target risk-management documentation itself – see also the discussion in [2]. Uciteli et al. [21] defined a Risk Identification Ontology embedded in the General Formal Ontology (GFO) [22,23], but their focus was on identifying risks in time periods surrounding surgical procedures. Kim et al. [24] integrated concepts from IEC 60601-1, IEC 62304, and ISO 14971 in an ontology for medical-software development processes; however, their model is primarily intended to support process integration and standards compliance, rather than automated reasoning over risk-management files. Schütz et al. [25] developed an ontology for medical devices in Germany with a broader focus on manufacturers, operators, and legal procedures, yet with the aim of general semantic interoperability, and based on a legal framework that has since been superseded by the EU MDR.

More recent work has explored semantic technologies for neighboring medical-device regulatory tasks: Chatteraj et al. [26] proposed an OWL/RDF knowledge-graph architecture for representing and querying FDA medical-device regulatory rules, while Zhu et al. [27] introduced a methodology for constructing a medical-device risk knowledge graph that combines standards, adverse-event data, and large language models. These works illustrate a continued interest in machine-processable representations for medical-device regulation and risk-related information, but they do not provide an ontology that supports the generation and validation of medical-device risk-management documentation itself. As already mentioned, this work builds on the RISKMAN ontology [2], which formalizes safety-oriented risk-management documentation for medical devices using OWL and SHACL. SECUMAN adopts this documentation-centered approach and is intended to integrate with RISKMAN, for example by mirroring its separation into “context”, “assessment”, and “control” aspects. However, SECUMAN targets the security domain: it models concepts required for medical-device security-risk management, including security-specific mitigation strategies represented as secure-design arguments, in contrast to the safe-design arguments used in RISKMAN.

Cybersecurity and security-risk ontologies in the medical domain. Ontology-based security modelling covers a wide range of concerns, including security taxonomies, vulnerability and attack representation, threat intelligence, operational technology and industrial Internet of Things security, semantic log analysis, and ontology-supported large-language-model applications, as summarized in surveys on these topics, e.g. [28,29,30, 31,32]. As to work closely related to our setting, Sills et al. [33] constructed a cybersecurity knowledge graph for medical-device cyber threat intelligence by integrating man-

manufacturer bulletins, CISA Industrial Control Systems Cyber Emergency Response Team alerts, Wikidata, and FDA AccessGUDID data; their focus, however, is on cyber-threat-intelligence aggregation and embedding-based retrieval rather than on a documentation ontology for security-risk management. Hannou et al. [34] introduced SafecareOnto, a cyber-physical security ontology for healthcare infrastructures that models hospital assets, dependencies, impacts, and protections in order to reason about incident propagation, whereas SECUMAN targets medical-device security-risk management files and their validation. Dart and Ahmed [35] propose CYBER-AIDD, a Unified Modeling Language-based ontology and governance framework for cybersecurity resilience in large healthcare providers, but this is a high-level organizational governance model rather than a formal OWL and SHACL framework for machine-checking manufacturer risk-management documentation. Liu et al. [36] presented an ontology-based framework for classifying forensic evidence from Internet of Medical Things devices; this work is security-adjacent and medical-domain-specific, but it concerns post-incident forensic evidence prioritization rather than prospective security-risk analysis and control documentation.

Work on Internet of Medical Things vulnerability modelling and ontology-assisted automation is closer to our setting, but still differs in its target artefacts. Bughio et al. [37] developed an Internet of Medical Things cybersecurity ontology for remote patient monitoring that represents devices, vulnerabilities, exploits, attacks, services, vendors, and affected assets, and in related work proposed a knowledge-graph-based framework for Internet of Medical Things vulnerability detection [38]; but they are primarily oriented toward infrastructure, vulnerabilities, and threat intelligence rather than toward documentation and validation. Ghosh et al. [39] presented CVE-LLM, an ontology-assisted large language model system for vulnerability evaluation in post-market medical-device cybersecurity, using existing cybersecurity knowledge such as CWE and CAPEC via the SEPSES knowledge graph to support artefacts such as affected-status classifications, justifications, comments, and environmental CVSS vectors; this complements SECUMAN by showing the usefulness of ontology-enriched automation, but it does not define a formal ontology and SHACL constraint layer for security-risk management files themselves. Finally, Chander et al. [40] proposed an Intelligent Cybersecurity Ontology Framework for Internet of Medical Things-enabled remote patient monitoring that combines ontology-based modelling of devices, dependencies, vulnerabilities, and threats with graph learning, temporal updates, risk scoring, and adaptive mitigation, thereby addressing operational threat detection and response rather than regulatory or engineering documentation.

Overall, the closest related works are those on medical-device cyber threat intelligence, Internet of Medical Things vulnerability modelling, and ontology-assisted vulnerability evaluation [33,37,38,39]; these approaches are complementary to SECUMAN, since they structure vulnerability, threat, asset, or operational-security knowledge, while SECUMAN focuses on a lightweight OWL ontology and SHACL shapes for representing, reasoning over, and validating medical-device security-risk management documentation.

3. Background

The SECUMAN ontology follows the RISKMAN ontology in being based on the recent VDE Spec 90025 [4], which proposes a structured format for digitalizing risk manage-

Security Risk 1																			
Context <table border="1"> <tr> <td>OR ENV</td> <td>NETWORK SECURITY No network connectivity</td> <td>PHYSICAL SECURITY Accessible to all medical staff</td> <td>SYSTEM/PLATFORM SECURITY Standard Windows, no additional</td> </tr> <tr> <td>PRODUCT</td> <td>IDENTITY/ACCESS MANAGEMENT No existing measures</td> <td>DATA SECURITY/PRIVACY No existing measures</td> <td>OPERATIONS/MAINTENANCE No existing measures</td> </tr> <tr> <td>ATTACK TYPE</td> <td colspan="3">Physical access</td> </tr> <tr> <td>PROTECTION GOAL</td> <td colspan="3">Confidentiality</td> </tr> </table>				OR ENV	NETWORK SECURITY No network connectivity	PHYSICAL SECURITY Accessible to all medical staff	SYSTEM/PLATFORM SECURITY Standard Windows, no additional	PRODUCT	IDENTITY/ACCESS MANAGEMENT No existing measures	DATA SECURITY/PRIVACY No existing measures	OPERATIONS/MAINTENANCE No existing measures	ATTACK TYPE	Physical access			PROTECTION GOAL	Confidentiality		
OR ENV	NETWORK SECURITY No network connectivity	PHYSICAL SECURITY Accessible to all medical staff	SYSTEM/PLATFORM SECURITY Standard Windows, no additional																
PRODUCT	IDENTITY/ACCESS MANAGEMENT No existing measures	DATA SECURITY/PRIVACY No existing measures	OPERATIONS/MAINTENANCE No existing measures																
ATTACK TYPE	Physical access																		
PROTECTION GOAL	Confidentiality																		
Assessment <table border="1"> <tr> <td>INITIAL RISK</td> <td>IMPACT LEVEL Medium</td> <td>LIKELIHOOD FACTOR Attacker Profile: Script Kiddie</td> <td>EXPOSURE LEVEL Known / easily detectable / limited access = 3</td> </tr> <tr> <td>IMPACT</td> <td colspan="3">Patient data read out (theft), GDPR violation SAFETY IMPACT: Affected: false</td> </tr> <tr> <td>COMPONENT</td> <td colspan="3">User Interface</td> </tr> <tr> <td>ASSET</td> <td colspan="3">Patient data</td> </tr> </table>				INITIAL RISK	IMPACT LEVEL Medium	LIKELIHOOD FACTOR Attacker Profile: Script Kiddie	EXPOSURE LEVEL Known / easily detectable / limited access = 3	IMPACT	Patient data read out (theft), GDPR violation SAFETY IMPACT: Affected: false			COMPONENT	User Interface			ASSET	Patient data		
INITIAL RISK	IMPACT LEVEL Medium	LIKELIHOOD FACTOR Attacker Profile: Script Kiddie	EXPOSURE LEVEL Known / easily detectable / limited access = 3																
IMPACT	Patient data read out (theft), GDPR violation SAFETY IMPACT: Affected: false																		
COMPONENT	User Interface																		
ASSET	Patient data																		
Control <table border="1"> <tr> <td>RESIDUAL RISK</td> <td>IMPACT LEVEL Medium</td> <td>LIKELIHOOD FACTOR Attacker Profile: Pen Tester</td> <td>EXPOSURE LEVEL Known / easily detectable / limited access = 3</td> </tr> </table>				RESIDUAL RISK	IMPACT LEVEL Medium	LIKELIHOOD FACTOR Attacker Profile: Pen Tester	EXPOSURE LEVEL Known / easily detectable / limited access = 3												
RESIDUAL RISK	IMPACT LEVEL Medium	LIKELIHOOD FACTOR Attacker Profile: Pen Tester	EXPOSURE LEVEL Known / easily detectable / limited access = 3																
Secure Design Argument <table border="1"> <tr> <td colspan="2"> SecDA sda0 TEST REPORT Reference to test result (only strong passwords are accepted, login is only possible with a second factor), reference to penetration test STRATEGY Through authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access </td> <td colspan="2"> TOP-LEVEL MITIGATION </td> </tr> <tr> <td colspan="2"> SecDA sda1 STRATEGY Through password security, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION Identities are established; passwords are not written on a Post-it attached to the device IMPACT Reference to configuration showing that encryption is enabled SAFETY IMPACT Affected: true SECURITY IMPACT Affected: false </td> <td colspan="2"> SecDA sda2 STRATEGY Through two-factor authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION The authenticator app is not used on a permanently installed device right next to it IMPACT Reference to configuration showing that the second factor is enabled SAFETY IMPACT Affected: false SECURITY IMPACT Affected: false </td> </tr> </table>				SecDA sda0 TEST REPORT Reference to test result (only strong passwords are accepted, login is only possible with a second factor), reference to penetration test STRATEGY Through authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access		TOP-LEVEL MITIGATION		SecDA sda1 STRATEGY Through password security, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION Identities are established; passwords are not written on a Post-it attached to the device IMPACT Reference to configuration showing that encryption is enabled SAFETY IMPACT Affected: true SECURITY IMPACT Affected: false		SecDA sda2 STRATEGY Through two-factor authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION The authenticator app is not used on a permanently installed device right next to it IMPACT Reference to configuration showing that the second factor is enabled SAFETY IMPACT Affected: false SECURITY IMPACT Affected: false									
SecDA sda0 TEST REPORT Reference to test result (only strong passwords are accepted, login is only possible with a second factor), reference to penetration test STRATEGY Through authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access		TOP-LEVEL MITIGATION																	
SecDA sda1 STRATEGY Through password security, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION Identities are established; passwords are not written on a Post-it attached to the device IMPACT Reference to configuration showing that encryption is enabled SAFETY IMPACT Affected: true SECURITY IMPACT Affected: false		SecDA sda2 STRATEGY Through two-factor authentication, reduce likelihood, preventive, so that only authorized/trusted persons have access LOCATION Location: Login service of the internal employee system ("intranet.company.com/login") CONTEXT Context: Corporate policy "Access Control 101". Responsible: IT security team. Lifecycle phase: Operation ASSUMPTION The authenticator app is not used on a permanently installed device right next to it IMPACT Reference to configuration showing that the second factor is enabled SAFETY IMPACT Affected: false SECURITY IMPACT Affected: false																	

Figure 1. HTML rendering of a security risk.

ment files, as well as a machine-readable exchange format using HTML with RDFa [41], which annotates selected HTML tags with Resource Description Framework (RDF) triples. However, whereas VDE Spec 90025 and RISKMAN focus on safety aspects of medical devices, primarily following ISO 14971 [3], SECUMAN formalises a model for cybersecurity risk management documentation drawing on cybersecurity-relevant standards and guidance, such as AAMI TIR57, AAMI SW96, NIST SP 800-30, and IEC 81001-5-1 [17,19,20,15]. At the same time, SECUMAN follows the same philosophy as VDE Spec 90025 and RISKMAN with respect to the digitalisation and validation of risk management files, and is intended to support integration with safety-focused risk management documentation.

In contrast to vulnerability-centric approaches, which assess isolated technical weaknesses, the model formalised by SECUMAN is built around contextualized threat scenarios, where risk emerges from the interaction of the operating environment, device-specific properties, attack types, and attacker capabilities. The model is structured into three sections: Context, Assessment, and Control, mirroring the conceptual structure of safety risk management as formalised in VDE Spec 90025. This separation enables consistent reuse of context attributes across multiple risk entries, thereby reducing the workload involved in creating a security risk management file, and supports traceability from threat description to risk control. The model is designed to support structured documentation of security risks and to enable alignment and systematic integration with safety risk management.

Figure 1 shows the rendering of an HTML file containing a cybersecurity risk structured according to the model formalised by SECUMAN. The context of the risk describes the conditions under which a threat can occur. It is defined by the combination of information pertaining to the operating environment, product, attack type, and protection goal. These elements represent reusable attributes that are not specific to a single risk, but can be applied across multiple risk scenarios and even across multiple devices employed in similar contexts.

Specifically, the operating environment captures external conditions relevant to cybersecurity. It includes physical security, network security, and system and platform security or hardening. The product describes security-relevant properties of the device itself, including identity and access management, data security and privacy, and operations and maintenance. The protection goal is one of confidentiality, integrity or availability.

The assessment aspects describe how the threat affects the system, its operational environment, and relevant stakeholders. In particular, the assessment refers to the context and extends it with the component through which the attack occurs, the affected asset, the impact, and the initial security risk level. The component represents the high-level system element that serves as the entry point of the attack, while the asset represents the physical or digital entity that has value to an individual, an organization, or a government and that is affected by the attack. The impact may include different types of consequences, such as harm to people, financial loss, legal implications, or operational disruption. It also explicitly records whether an impact on safety can occur. The initial security risk level, i.e., the security risk level before any mitigation has been proposed, is composed of an exposure level and an attacker profile, which together form the likelihood factor, as well as an impact level.

The control section represents the risk after the application of control measures. It refers to the assessment aspects and includes a secure design argument and the residual security risk level. The latter has the same structure as the initial security risk level, but represents the risk level after mitigation measures have been applied.

Cybersecurity risk control in the proposed model is represented through secure design arguments (SecDAs). SecDAs are a variant of the safety design arguments proposed in VDE Spec 90025 and formalised in RISKMAN – which in turn are a simplified form of assurance cases as defined in standards such as ISO/IEC/IEEE 15026 and ISO/TS 81001-2-1 [42,43] – specialised for mitigation strategies addressing security risks.

SecDAs are structured as trees, with the top-level SecDA representing the claim that a given cybersecurity risk is adequately controlled. Sub-arguments, or sub-SecDAs, contribute to this claim by refining the control strategy into more specific measures. The tree structure enforces a clear separation of abstraction levels: the top-level SecDA defines the overall strategy and includes a test report for effectiveness verification; intermediate SecDAs refine the strategy without implementation detail; and leaf SecDAs define concrete control measures. The strategy defines the overarching approach to risk mitigation, including the control type, i.e., preventive, detective, or corrective, the rationale for the control, and its intended effect on risk parameters, e.g., the likelihood or impact factor. In addition to the strategy, each leaf SecDA contains attributes describing the intervention method, the location of intervention, the organisational context of intervention, assumptions, and an assessment of the impact on safety and security.

Specifically, the intervention method describes the mechanism by which the control is applied. It is the most concrete description of the control and specifies what must be implemented. It may include technical, operational, or administrative measures. The location of intervention specifies where the control is applied within the system, including device internals, connected infrastructure, or, where relevant, the deployment environment. The organisational context of intervention describes the operational setting of the control, including the device lifecycle phase, responsible stakeholders, and relevant policies or practices. The assumptions capture conditions that must hold for the control to be effective, such as the availability of infrastructure or expected user behaviour. Finally,

the assessment of the impact on safety and security indicates whether the control measure may itself introduce a safety impact or an additional security impact that requires mitigation.

SecDAs inherit the subcategorisation of safe design arguments from VDE Spec 90025 and RISKMAN into implementation-specific arguments and assurance arguments. Implementation-specific SecDAs include an implementation manifest, which provides evidence of how the control measure has been realised within the medical device; in particular, leaf SecDAs must always include an implementation manifest. Assurance SecDAs, on the other hand, additionally reference established state-of-the-art practices, such as standards, regulatory guidance, or accepted security frameworks. These references support the justification and traceability of control strategies. In contrast to safety, where standards often define the state of the art more comprehensively, cybersecurity assurance may also rely on evolving best practices and accepted security baselines.

4. Design & Overview

4.1. Design

The SECUMAN ontology was developed based on requirements and conceptual design decisions defined by a consortium of experts in medical devices, regulatory affairs, security risk management, ontology engineering, and software development. We followed the Linked Open Terms (LOT) methodology [44], which builds on the NeOn methodology [45]. The development process comprised three steps: (1) Requirement specification: We studied relevant norms and standards and examined real medical device security risk management files to establish a shared understanding of the required terms. This was done through regular consortium meetings, workshops, and meetings with manufacturers and security risk experts. The result was a conceptual model of security risk documentation that formed the basis of the SECUMAN ontology. (2) Implementation: We encoded the ontology and shapes in OWL and SHACL. During development, we primarily used the HermiT reasoner [46] and the Ontology Pitfalls Scanner (OOPS!) [47] to support validity and consistency checks. (3) Publication and maintenance: The SECUMAN ontology and shapes² as well as the associated validation pipeline³ are maintained in separate GitHub repositories for version control, issue tracking, and open online access.

4.2. Formal Background

The SECUMAN ontology is encoded in the $\mathcal{EL}$ profile of OWL [48], while the shapes are specified in SHACL [49]. For ease of presentation, we describe the ontology here using the description logic $\mathcal{EL}^{++}$ [50,51] that underlies the $\mathcal{EL}$ profile of OWL. For shapes, we use the abstract SHACL syntax introduced by Corman et al. [52,53]. Before giving an overview of SECUMAN, we therefore briefly recall the syntax, semantics, and aspects of reasoning relevant to our approach of $\mathcal{EL}^{++}$ as well as SHACL in the sense of Corman et al.

²<https://github.com/cl-tud/secuman>

³<https://github.com/cl-tud/semeco-q2-validation-pipeline>

The Description Logic $\mathcal{EL}^{++}$ $\mathcal{EL}^{++}$'s concept constructors and their semantics are recalled in the upper part of Table 1; the middle (lower) part shows the constructs allowed in a TBox (ABox). As usual for description logics, the semantics of $\mathcal{EL}^{++}$ is defined

Table 1. Syntax and semantics of concept, TBox, and ABox expressions of $\mathcal{EL}^{++}$.

Name	Syntax	Semantics
individual name	$a \in N_I$	$a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$
concept name	$A \in N_C$	$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
role name	$R \in N_R$	$R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
top	$\top$	$\Delta^{\mathcal{I}}$
bottom	$\perp$	$\emptyset$
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists R.C$	$\{x \in \Delta^{\mathcal{I}} \mid \exists y \in \Delta^{\mathcal{I}} : (x, y) \in R^{\mathcal{I}} \ \& \ y \in C^{\mathcal{I}}\}$
domain restriction	$dom(R) \sqsubseteq A$	$R^{\mathcal{I}} \subseteq A^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
range restriction	$ran(R) \sqsubseteq A$	$R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times A^{\mathcal{I}}$
general concept inclusion	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
role inclusion axiom	$R_1 \circ \dots \circ R_k \sqsubseteq R$	$R_1^{\mathcal{I}} \circ \dots \circ R_k^{\mathcal{I}} \subseteq R^{\mathcal{I}}$
concept assertion	$A(a)$	$a^{\mathcal{I}} \in A^{\mathcal{I}}$
role assertion	$R(a, b)$	$(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$

via *interpretations* $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ with a non-empty *domain* $\Delta^{\mathcal{I}}$ and an *interpretation function* $\cdot^{\mathcal{I}}$. An interpretation $\mathcal{I}$ is a *model* of an ABox $\mathcal{A}$ (TBox $\mathcal{T}$) if it satisfies all elements of $\mathcal{A}$ ($\mathcal{T}$) as per Table 1. An assertion α is *entailed* by $\mathcal{T} \cup \mathcal{A}$, written $\mathcal{T} \cup \mathcal{A} \models \alpha$, if every model of $\mathcal{T} \cup \mathcal{A}$ is a model of α .

SHACL To define SHACL, we conveniently re-use description logic vocabulary, viz., pairwise disjoint sets N_C of *classes*, N_R of *properties*, and N_I of *individuals*. A finite set $\mathcal{A}$ of assertions (an ABox) can then be seen as representing a labelled *graph*, with individuals acting as nodes, classes labelling nodes, and properties labelling edges. The syntax of *shape expressions* ϕ and *path expressions* E is shown in Table 2 (top). For the semantics, a given graph (ABox) $\mathcal{A}$ with nodes (individuals) $N_I(\mathcal{A})$ defines an evaluation function $\llbracket \cdot \rrbracket^{\mathcal{A}}$ that assigns to each path expression E a binary relation $\llbracket E \rrbracket^{\mathcal{A}} \subseteq N_I(\mathcal{A}) \times N_I(\mathcal{A})$, and to each shape expression ϕ a set $\llbracket \phi \rrbracket^{\mathcal{A}} \subseteq N_I(\mathcal{A})$ via induction as shown in Table 2 (bottom).

A *shape constraint* is an expression of the form $A \leftarrow \phi$, with $A \in N_C$ and ϕ a shape expression. A *shape schema* is a pair $(\mathcal{C}, \mathcal{B})$ where $\mathcal{C}$ is a set of shape constraints and $\mathcal{B}$ is a set of *target* concept assertions. Intuitively, a target $A(a)$ expresses the requirement that a be labelled by A . Formally, an ABox $\mathcal{A}$ is a *model* for a set $\mathcal{C}$ of constraints if $\llbracket \phi \rrbracket^{\mathcal{A}} \subseteq \llbracket A \rrbracket^{\mathcal{A}}$ for all $A \leftarrow \phi \in \mathcal{C}$. An ABox $\mathcal{A}$ is *validated* against a schema $(\mathcal{C}, \mathcal{B})$ if there exists a set $\mathcal{B}'$ of concept assertions such that (1) $\mathcal{B} \subseteq \mathcal{B}'$, (2) $N_I(\mathcal{B}') \subseteq N_I(\mathcal{A})$, and (3) $\mathcal{A} \cup \mathcal{B}'$ is a model for $\mathcal{C}$. For brevity, in the following we also use the syntactic abbreviations $\leq_n E.\phi$ for $\neg(\geq_{n+1} E.\phi)$ and $=_1 E.\phi$ for $\geq_1 E.\phi \wedge \leq_1 E.\phi$.

Reasoning For reasoning with respect to the SECUMAN ontology and shapes, we adopt a materialisation-based approach. First, all relevant logical consequences of a given ABox together with the $\mathcal{EL}^{++}$ ontology are computed; more precisely, we derive all entailed concept and role assertions involving concept, role, and individual names occur-

Table 2. Syntax and semantics of path and shape expressions.

The syntax of path expressions E and shape expressions ϕ is given by the grammars		
$E ::= R \mid R^- \mid E \cup E \mid E \cdot E \mid E^*$ and $\phi ::= \top \mid A \mid a \mid \phi_1 \wedge \phi_2 \mid \neg \phi \mid \geq_n E.\phi \mid \forall E.\phi \mid E = E$		
where $n \in \mathbb{N}^+$, $A \in \mathbf{N}_C$, $a \in \mathbf{N}_I$, and $R \in \mathbf{N}_R$ with R^- indicating the inverse of R .		
$\llbracket R \rrbracket^{\mathcal{A}} = \{(a, b) \mid R(a, b) \in \mathcal{A}\}$	$\llbracket E_1 \cdot E_2 \rrbracket^{\mathcal{A}} = \llbracket E_1 \rrbracket^{\mathcal{A}} \circ \llbracket E_2 \rrbracket^{\mathcal{A}}$	
$\llbracket R^- \rrbracket^{\mathcal{A}} = \{(b, a) \mid R(a, b) \in \mathcal{A}\}$	$\llbracket E_1 \cup E_2 \rrbracket^{\mathcal{A}} = \llbracket E_1 \rrbracket^{\mathcal{A}} \cup \llbracket E_2 \rrbracket^{\mathcal{A}}$	$\llbracket E^* \rrbracket^{\mathcal{A}} = \left(\llbracket E \rrbracket^{\mathcal{A}} \right)^*$
$\llbracket \top \rrbracket^{\mathcal{A}} = \mathbf{N}_I(\mathcal{A})$	$\llbracket A \rrbracket^{\mathcal{A}} = \{a \mid A(a) \in \mathcal{A}\}$	$\llbracket a \rrbracket^{\mathcal{A}} = \{a\}$
$\llbracket \phi_1 \wedge \phi_2 \rrbracket^{\mathcal{A}} = \llbracket \phi_1 \rrbracket^{\mathcal{A}} \cap \llbracket \phi_2 \rrbracket^{\mathcal{A}}$	$\llbracket \neg \phi \rrbracket^{\mathcal{A}} = \mathbf{N}_I(\mathcal{A}) \setminus \llbracket \phi \rrbracket^{\mathcal{A}}$	
$\llbracket \forall E.\phi \rrbracket^{\mathcal{A}} = \{a \mid \forall b : (a, b) \in \llbracket E \rrbracket^{\mathcal{A}} \text{ implies } b \in \llbracket \phi \rrbracket^{\mathcal{A}}\}$		
$\llbracket \geq_n E.\phi \rrbracket^{\mathcal{A}} = \{a \mid \{(a, b) \in \llbracket E \rrbracket^{\mathcal{A}} \text{ and } b \in \llbracket \phi \rrbracket^{\mathcal{A}}\} \geq n\}$		
$\llbracket E_1 = E_2 \rrbracket^{\mathcal{A}} = \{a \mid \forall b : (a, b) \in \llbracket E_1 \rrbracket^{\mathcal{A}} \text{ iff } (a, b) \in \llbracket E_2 \rrbracket^{\mathcal{A}}\}$		

ring in the input ABox. Second, the resulting completed ABox is validated against the SHACL constraints.

The feasibility of this approach in our setting is ensured by three observations. First, the role inclusion axioms of SECUMAN satisfy the syntactic restriction imposed by Baader et al. [51, Section 3], and hence entailment of assertions can be decided in polynomial time. Second, SECUMAN uses existential quantification only on the left-hand side of general concept inclusion axioms. As a consequence, materialisation does not introduce fresh individuals, and thus remains finite. Finally, under the formal semantics of Corman et al., SHACL validation is NP-complete in general, whereas certain positive fragments are PTIME-complete [52].

4.3. The SECUMAN Ontology & Shapes

Common Aspects Figure 3 presents the GCIs for classes shared across the assessment and control aspects of security risk management documentation. In particular, the central class *SecurityRisk* is characterized compositionally: every individual with an attacker profile and an exposure level is a *LikelihoodFactor*, every individual with an impact level and a likelihood factor is a *SecurityRiskLevel*, and every individual with a security risk level is a *SecurityRisk*. Finally, there is the GCI that introduces the auxiliary class *ImpactAssessment*, which captures an assessment of whether a security risk or a control measure may induce further risks, in particular in the safety or security dimension. Concretely, it states that every individual with an *isAffected* value is an *ImpactAssessment*.

$$\begin{aligned}
 \exists \text{isAffected}. \top &\sqsubseteq \text{ImpactAssessment} & (1) \\
 \exists \text{hasAttackerProfile}. \top \sqcap \exists \text{hasExposureLevel}. \top &\sqsubseteq \text{LikelihoodFactor} & (2) \\
 \exists \text{hasSecurityRiskLevel}. \top &\sqsubseteq \text{SecurityRisk} & (3) \\
 \exists \text{hasImpactLevel}. \top \sqcap \exists \text{hasLikelihoodFactor}. \top &\sqsubseteq \text{SecurityRiskLevel} & (4)
 \end{aligned}$$

Figure 3. General concept inclusions for common aspects.




Figure 2. Schema diagram of the SECUMAN classes and properties divided into context, assessment, and control aspects. Any arrow representing a property without an explicit name means that the property has the name *hasX* where *X* is the class in the range, e.g. *hasAsset* for the property with domain *AnalyzedSecurityRisk* and range *Asset*. Property arrows are also labelled with information pertaining to the Shape constraints, e.g. an *AnalyzedSecurityRisk* can have 1 or more assets or a *SecDA* only has a *TestReport* if it is a top *SecDA* (i.e. is not the child of any other *SecDA*), and then it must have exactly one *TestReport*. If there are no such labels, then each instance in the domain must be related to exactly one instance in the range. Moreover, any two classes without subclass relationship are disjoint.

The domain and range axioms in Figures 4 and 5 complement these GCI by fixing the intended source and target classes of the relevant properties, including *hasAttackerProfile*, *hasExposureLevel*, *hasImpactLevel*, *hasLikelihoodFactor*, *hasSecurityRiskLevel*, *hasRationale*, *hasSafetyImpactAssessment*, *hasSecurityImpactAssessment*, and *isAffected*. In particular, both *hasSafetyImpactAssessment* and *hasSecurityImpactAssessment* have domain *ImpactAssessmentDomain*. As shown by further GCI introduced below, this makes it possible to associate impact assessments both with risks and control measures (security design arguments). Moreover, the range of *isAffected* is a *Boolean*, while an *ImpactAssessment* individual may also have associated a *Rationale*, which captures, as its name suggests, the rationale underlying a particular impact assessment.

$$\text{dom}(\text{hasAttackerProfile}) \sqsubseteq \text{LikelihoodFactor} \quad (5)$$

$$\text{ran}(\text{hasAttackerProfile}) \sqsubseteq \text{AttackerProfile} \quad (6)$$

$$\text{dom}(\text{hasImpactLevel}) \sqsubseteq \text{SecurityRiskLevel} \quad (7)$$

$$\text{ran}(\text{hasImpactLevel}) \sqsubseteq \text{ImpactLevel} \quad (8)$$

Figure 4. Range and domain restrictions for common aspects, part I.

$dom(hasLikelihoodFactor) \sqsubseteq SecurityRiskLevel$	(9)
$ran(hasLikelihoodFactor) \sqsubseteq LikelihoodFactor$	(10)
$dom(hasRationale) \sqsubseteq ImpactAssessment$	(11)
$ran(hasRationale) \sqsubseteq Rationale$	(12)
$dom(hasSafetyImpactAssessment) \sqsubseteq ImpactAssessmentDomain$	(13)
$ran(hasSafetyImpactAssessment) \sqsubseteq ImpactAssessment$	(14)
$dom(hasSecurityImpactAssessment) \sqsubseteq ImpactAssessmentDomain$	(15)
$ran(hasSecurityImpactAssessment) \sqsubseteq ImpactAssessment$	(16)
$dom(hasSecurityRiskLevel) \sqsubseteq SecurityRisk$	(17)
$ran(hasSecurityRiskLevel) \sqsubseteq SecurityRiskLevel$	(18)
$dom(hasExposureLevel) \sqsubseteq LikelihoodFactor$	(19)
$ran(hasExposureLevel) \sqsubseteq ExposureLevel$	(20)
$dom(isAffected) \sqsubseteq ImpactAssessment$	(21)
$ran(isAffected) \sqsubseteq Boolean$	(22)

Figure 5. Range and domain restrictions for common aspects, part II.

At the constraint level, the SHACL shapes in Figure 6 make these dependencies explicit for instance data. In particular, each `LikelihoodFactor` must have exactly one `AttackerProfile` and exactly one `ExposureLevel`, each `SecurityRisk` must have exactly one `SecurityRiskLevel`, and each `SecurityRiskLevel` must have exactly one `ImpactLevel` and exactly one `LikelihoodFactor`. Moreover, each `ImpactAssessment` must specify exactly one `isAffected` value and may have at most one `Rationale`.

$$ImpactAssessment \leftarrow =_I isAffected.T \wedge \leq_I hasRationale.T \quad (23)$$

$$LikelihoodFactor \leftarrow =_I hasAttackerProfile.T \wedge =_I hasExposureLevel.T \quad (24)$$

$$SecurityRisk \leftarrow =_I hasSecurityRiskLevel.T \quad (25)$$

$$SecurityRiskLevel \leftarrow =_I hasImpactLevel.T \wedge =_I hasLikelihoodFactor.T \quad (26)$$

Figure 6. Shape constraints for common aspects.

Context The GCIs in Figure 7, together with the domain and range restrictions in Figure 8 capture the context-specific information used to characterize a security concern more precisely. The central class `ContextSpecificThreat` brings together four dimensions: the operating environment, the product, the attack type, and the protection goal. The two classes `OperatingEnvironment` and `Product` further structure this information: `OperatingEnvironment` groups physical, network, and system/platform security aspects, whereas `Product` groups identity and access management, data security and privacy, and operations and maintenance. The protection goals are given by the distinguished values `availability`, `confidentiality`, and `integrity`.

$$\begin{aligned}
& \exists \text{hasOperatingEnvironment}. \top \sqcap \exists \text{hasProduct}. \top & (27) \\
& \sqcap \exists \text{hasAttackType}. \top \sqcap \exists \text{hasProtectionGoal}. \top \sqsubseteq \text{ContextSpecificThreat} \\
& \exists \text{hasPhysicalSecurity}. \top \sqcap \exists \text{hasNetworkSecurity}. \top & (28) \\
& \sqcap \exists \text{hasSystemAndPlatformSecurity}. \top \sqsubseteq \text{OperatingEnvironment} \\
& \exists \text{hasIdentityAndAccessManagement}. \top & (29) \\
& \sqcap \exists \text{hasDataSecurityAndPrivacy}. \top \\
& \sqcap \exists \text{hasOperationsAndMaintenance}. \top \sqsubseteq \text{Product} \\
& \text{ProtectionGoal(availability)} \sqsubseteq \text{ProtectionGoal(confidentiality)} & (30) \\
& \text{ProtectionGoal(integrity)}
\end{aligned}$$

Figure 7. General concept inclusion axioms and concept assertions for context aspects.

$$\begin{aligned}
& \text{dom}(\text{hasAttackType}) \sqsubseteq \text{ContextSpecificThreat} & (31) \\
& \text{ran}(\text{hasAttackType}) \sqsubseteq \text{AttackType} & (32) \\
& \text{dom}(\text{hasDataSecurityAndPrivacy}) \sqsubseteq \text{Product} & (33) \\
& \text{ran}(\text{hasDataSecurityAndPrivacy}) \sqsubseteq \text{DataSecurityAndPrivacy} & (34) \\
& \text{dom}(\text{hasIdentityAndAccessManagement}) \sqsubseteq \text{Product} & (35) \\
& \text{ran}(\text{hasIdentityAndAccessManagement}) \sqsubseteq \text{IdentityAndAccessManagement} & (36) \\
& \text{dom}(\text{hasNetworkSecurity}) \sqsubseteq \text{OperatingEnvironment} & (37) \\
& \text{ran}(\text{hasNetworkSecurity}) \sqsubseteq \text{NetworkSecurity} & (38) \\
& \text{dom}(\text{hasOperationsAndMaintenance}) \sqsubseteq \text{Product} & (39) \\
& \text{ran}(\text{hasOperationsAndMaintenance}) \sqsubseteq \text{OperationsAndMaintenance} & (40) \\
& \text{dom}(\text{hasOperatingEnvironment}) \sqsubseteq \text{ContextSpecificThreat} & (41) \\
& \text{ran}(\text{hasOperatingEnvironment}) \sqsubseteq \text{OperatingEnvironment} & (42) \\
& \text{dom}(\text{hasPhysicalSecurity}) \sqsubseteq \text{OperatingEnvironment} & (43) \\
& \text{ran}(\text{hasPhysicalSecurity}) \sqsubseteq \text{PhysicalSecurity} & (44) \\
& \text{dom}(\text{hasProduct}) \sqsubseteq \text{ContextSpecificThreat} & (45) \\
& \text{ran}(\text{hasProduct}) \sqsubseteq \text{Product} & (46) \\
& \text{dom}(\text{hasProtectionGoal}) \sqsubseteq \text{ContextSpecificThreat} & (47) \\
& \text{ran}(\text{hasProtectionGoal}) \sqsubseteq \text{ProtectionGoal} & (48) \\
& \text{dom}(\text{hasSystemAndPlatformSecurity}) \sqsubseteq \text{OperatingEnvironment} & (49) \\
& \text{ran}(\text{hasSystemAndPlatformSecurity}) \sqsubseteq \text{SystemAndPlatformSecurity} & (50)
\end{aligned}$$

Figure 8. Range and domain restrictions for context aspects.

The corresponding SHACL shapes in Figure 9 require these structures to be present explicitly in the data. Thus, each `ContextSpecificThreat` must specify exactly one value for each of its four dimensions, while each `OperatingEnvironment` and each `Product` must provide exactly one value for each of their respective subdimensions. In addition, every `ProtectionGoal` value must belong to `{availability,confidentiality,integrity}`.

$$\text{ContextSpecificThreat} \leftarrow =_I \text{hasOperatingEnvironment}.T \wedge =_I \text{hasProduct}.T \wedge =_I \text{hasAttackType}.T \wedge =_I \text{hasProtectionGoal}.T \quad (51)$$

$$\text{OperatingEnvironment} \leftarrow =_I \text{hasPhysicalSecurity}.T \wedge =_I \text{hasNetworkSecurity}.T \wedge =_I \text{hasSystemAndPlatformSecurity}.T \quad (52)$$

$$\text{Product} \leftarrow =_I \text{hasIdentityAndAccessManagement}.T \wedge =_I \text{hasDataSecurityAndPrivacy}.T \wedge =_I \text{hasOperationsAndMaintenance}.T \quad (53)$$

$$\text{ProtectionGoal} \leftarrow \{\text{availability, confidentiality, integrity}\} \quad (54)$$

Figure 9. Shape constraints for context aspects.

Assessment The GCI in Figure 10 capture the essential structure of analyzed security risks. In particular, every individual that is related to a context-specific threat, an impact, an initial security risk level, a component, and at least one asset is characterized as an *AnalyzedSecurityRisk*. Moreover, the property *hasInitialSecurityRiskLevel* is specified as a subproperty of *hasSecurityRiskLevel*, which implies that every *AnalyzedSecurityRisk* is also a *SecurityRisk* (see Figure 3), while *Impact* is placed under *ImpactAssessmentDomain*, thereby linking impacts – and indirectly also analyzed risks – to the common impact-assessment pattern introduced above.

$$\text{hasInitialSecurityRiskLevel} \sqsubseteq \text{hasSecurityRiskLevel} \quad (55)$$

$$\text{Impact} \sqsubseteq \text{ImpactAssessmentDomain} \quad (56)$$

$$\begin{aligned} & \exists \text{hasContextSpecificThreat}.T \sqcap \exists \text{hasImpact}.T \\ & \sqcap \exists \text{hasInitialSecurityRiskLevel}.T \sqcap \exists \text{hasComponent}.T \\ & \sqcap \exists \text{hasAsset}.T \sqsubseteq \text{AnalyzedSecurityRisk} \end{aligned} \quad (57)$$

Figure 10. General concept and role inclusion axioms for assessment aspects.

Complementing this conceptual characterization, the domain and range restrictions in Figure 11 determine the intended typing of these relations by assigning *AnalyzedSecurityRisk* as their domain and *ContextSpecificThreat*, *Impact*, *SecurityRiskLevel*, *Component*, and *Asset* as the corresponding ranges.

$$\begin{aligned}
\text{dom}(\text{hasAsset}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (58) \\
\text{ran}(\text{hasAsset}) &\sqsubseteq \text{Asset} & (59) \\
\text{dom}(\text{hasComponent}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (60) \\
\text{ran}(\text{hasComponent}) &\sqsubseteq \text{Component} & (61) \\
\text{dom}(\text{hasContextSpecificThreat}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (62) \\
\text{ran}(\text{hasContextSpecificThreat}) &\sqsubseteq \text{ContextSpecificThreat} & (63) \\
\text{dom}(\text{hasImpact}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (64) \\
\text{ran}(\text{hasImpact}) &\sqsubseteq \text{Impact} & (65) \\
\text{dom}(\text{hasInitialSecurityRiskLevel}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (66) \\
\text{ran}(\text{hasInitialSecurityRiskLevel}) &\sqsubseteq \text{SecurityRiskLevel} & (67)
\end{aligned}$$

Figure 11. Range and domain restrictions for assessment aspects.

The ontological axioms are mirrored at the constraint level by the SHACL shape in Figure 12, which makes the intended dependencies explicit for data: each *AnalyzedSecurityRisk* must specify exactly one context-specific threat, exactly one initial security risk level, and exactly one component, while requiring at least one impact and at least one asset.

$$\begin{aligned}
\text{AnalyzedSecurityRisk} \leftarrow \text{hasContextSpecificThreat.T} \wedge \geq_1 \text{hasImpact.T} \\
\wedge \text{hasInitialSecurityRiskLevel.T} \wedge \text{hasComponent.T} \\
\wedge \geq_1 \text{hasAsset.T}
\end{aligned} \quad (68)$$

Figure 12. Shape constraints for assessment aspects.

Control The main GCI in Figure 13 defines controlled security risks: every individual that is related to an analyzed security risk, a residual security risk level, and a secure design argument is characterized as a *ControlledSecurityRisk*. In addition, and analogously to *hasInitialSecurityRiskLevel* for analyzed risks, *hasResidualSecurityRiskLevel* is introduced as a subproperty of *hasSecurityRiskLevel*, from which it follows that controlled security risks are also security risks.

$$\begin{aligned}
\text{hasResidualSecurityRiskLevel} &\sqsubseteq \text{hasSecurityRiskLevel} & (69) \\
\exists \text{hasAnalyzedSecurityRisk.T} \sqcap \exists \text{hasResidualSecurityRiskLevel.T} \\
\sqcap \exists \text{isMitigatedBy.T} &\sqsubseteq \text{ControlledSecurityRisk} & (70)
\end{aligned}$$

Figure 13. General concept and role inclusion for control aspects.

The domain and range restrictions in Figure 14 again fix the expected typing of the associated relations, with *ControlledSecurityRisk* as domain and *AnalyzedSecurityRisk*, *SecurityRiskLevel*, and *SecDA* as the corresponding ranges.

$$\begin{aligned}
\text{dom}(\text{hasAnalyzedSecurityRisk}) &\sqsubseteq \text{ControlledSecurityRisk} & (71) \\
\text{ran}(\text{hasAnalyzedSecurityRisk}) &\sqsubseteq \text{AnalyzedSecurityRisk} & (72) \\
\text{dom}(\text{hasResidualSecurityRiskLevel}) &\sqsubseteq \text{ControlledSecurityRisk} & (73) \\
\text{ran}(\text{hasResidualSecurityRiskLevel}) &\sqsubseteq \text{SecurityRiskLevel} & (74) \\
\text{dom}(\text{isMitigatedBy}) &\sqsubseteq \text{ControlledSecurityRisk} & (75) \\
\text{ran}(\text{isMitigatedBy}) &\sqsubseteq \text{SecDA} & (76)
\end{aligned}$$

Figure 14. Range and domain restrictions for control aspects.

Finally, the SHACL shape in Figure 15 formulates the corresponding constraints at the data level by requiring each `ControlledSecurityRisk` to specify exactly one analyzed security risk, exactly one residual security risk level, and exactly one secure design argument.

$$\begin{aligned}
\text{ControlledSecurityRisk} \leftarrow \text{hasAnalyzedSecurityRisk}.\top \wedge \text{hasResidualSecurityRiskLevel}.\top \\
\wedge \text{isMitigatedBy}.\top
\end{aligned} \quad (77)$$

Figure 15. Shape constraints for control aspects.

Secure Design Arguments The GCI in Figure 16 define the central structure of secure design arguments and their refinements. In particular, every individual that is related to a strategy is characterized as a `SecDA`, while every `SecDA` with an implementation manifest is classified as a `SecDAI`. The remaining GCIs introduce the assurance-specific specializations `AssuranceSecDA` and `AssuranceSecDAI`. In addition, the axiom $\text{SecDA} \sqsubseteq \text{ImpactAssessmentDomain}$ places secure design arguments into the common impact-assessment pattern introduced earlier.

$$\begin{aligned}
\text{SecDA} \sqcap \exists \text{hasAssurance}.\top &\sqsubseteq \text{AssuranceSecDA} & (78) \\
\text{SecDAI} \sqcap \text{AssuranceSecDA} &\sqsubseteq \text{AssuranceSecDAI} & (79) \\
\text{SecDA} &\sqsubseteq \text{ImpactAssessmentDomain} & (80) \\
\exists \text{hasStrategy}.\top &\sqsubseteq \text{SecDA} & (81) \\
\text{SecDA} \sqcap \exists \text{hasImplementationManifest}.\top &\sqsubseteq \text{SecDAI} & (82)
\end{aligned}$$

Figure 16. General concept and role inclusion axioms for secure design argument aspects.

The domain and range restrictions in Figure 17 fix the expected typing of the relations involving secure design arguments, with `SecDA` as domain and, depending on the property, assumptions, assurances, intervention methods, implementation manifests, locations and organisational contexts of intervention, strategies, subordinate secure design arguments, and test reports as corresponding ranges.

$dom(hasAssumption) \sqsubseteq SecDA$	(83)
$ran(hasAssumption) \sqsubseteq Assumption$	(84)
$dom(hasAssurance) \sqsubseteq SecDA$	(85)
$ran(hasAssurance) \sqsubseteq Assurance$	(86)
$dom(hasInterventionMethod) \sqsubseteq SecDA$	(87)
$ran(hasInterventionMethod) \sqsubseteq InterventionMethod$	(88)
$dom(hasImplementationManifest) \sqsubseteq SecDA$	(89)
$ran(hasImplementationManifest) \sqsubseteq ImplementationManifest$	(90)
$dom(hasLocationOfIntervention) \sqsubseteq SecDA$	(91)
$ran(hasLocationOfIntervention) \sqsubseteq LocationOfIntervention$	(92)
$dom(hasOrganisationalContextOfIntervention) \sqsubseteq SecDA$	(93)
$ran(hasOrganisationalContextOfIntervention) \sqsubseteq OrganisationalContextOfIntervention$	(94)
$dom(hasStrategy) \sqsubseteq SecDA$	(95)
$ran(hasStrategy) \sqsubseteq Strategy$	(96)
$dom(hasSubSecDA) \sqsubseteq SecDA$	(97)
$ran(hasSubSecDA) \sqsubseteq SecDA$	(98)
$dom(hasTestReport) \sqsubseteq SecDA$	(99)
$ran(hasTestReport) \sqsubseteq TestReport$	(100)

Figure 17. Range and domain restrictions for secure design argument aspects.

Once more, the SHACL constraints in Figure 18, which include several helper constraints, specify the structural constraints of secure design arguments at the data level. They require each SecDA to specify exactly one strategy, while allowing at most one location of intervention, at most one assumption, at most one implementation manifest, and at most one assurance. Moreover, the recursive organization of secure design arguments is constrained explicitly: every SecDA must have at least one implementation-level argument reachable via the hasSubSecDA relation, which implies in particular that every leaf secure design argument must have an implementation manifest. In addition, top-level arguments must have exactly one test report, whereas non-top-level arguments must have none, and only leaf arguments may carry impact and intervention related information as well as assumptions. For assurance-specific arguments, the corresponding SHACL shape further requires exactly one assurance and propagates the assurance character recursively to subordinate arguments. Moreover, the shapes for SecDAI and AssuranceSecDAI require exactly one implementation manifest. In this respect, the subdivision of secure design arguments into assurance-specific and implementation-specific arguments, as well as the corresponding constraints, parallels that of safe design arguments from [2]; for convenience, however, these structures are duplicated explicitly in the SECUMAN shapes.

$$\begin{aligned}
\text{TopSecDA} &\leftarrow \leq_0 \text{hasSubSecDA}^{\neg}. \top & (101) \\
\text{NotTopSecDA} &\leftarrow \geq_1 \text{hasSubSecDA}^{\neg}. \top & (102) \\
\text{LeafSecDA} &\leftarrow \leq_0 \text{hasSubSecDA}. \top & (103) \\
\text{NotLeafSecDA} &\leftarrow \geq_1 \text{hasSubSecDA}. \top & (104) \\
\text{TopSecDAAndOneTestReport} &\leftarrow \text{TopSecDA} \wedge =_1 \text{hasTestReport}. \top & (105) \\
\text{NotTopSecDAAndNoTestReport} &\leftarrow \text{NotTopSecDA} \wedge \leq_0 \text{hasTestReport}. \top & (106) \\
\text{LeafSecDAAndOneA} &\leftarrow \text{LeafSecDA} \wedge =_1 h. \top & \\
&\quad \text{for all } (A, h) \text{ in} & \\
&\quad \{(\text{InterventionMethod}, \text{hasInterventionMethod}), & \\
&\quad (\text{OrganisationalContextOfIntervention}, & \\
&\quad \text{hasOrganisationalContextOfIntervention}), & \\
&\quad (\text{SafetyImpactAssessment}, \text{hasSafetyImpactAssessment}), & \\
&\quad (\text{SecurityImpactAssessment}, \text{hasSecurityImpactAssessment})\} & (107) \\
\text{NotLeafSecDAAndNoA} &\leftarrow \text{NotLeafSecDA} \wedge \leq_0 h. \top & \\
&\quad \text{for all } (A, h) \text{ in} & \\
&\quad \{(\text{Assumption}, \text{hasAssumption}), & \\
&\quad (\text{InterventionMethod}, \text{hasInterventionMethod}), & \\
&\quad (\text{LocationOfIntervention}, \text{hasLocationOfIntervention}), & \\
&\quad (\text{OrganisationalContextOfIntervention}, & \\
&\quad \text{hasOrganisationalContextOfIntervention}), & \\
&\quad (\text{SafetyImpactAssessment}, \text{hasSafetyImpactAssessment}), & \\
&\quad (\text{SecurityImpactAssessment}, \text{hasSecurityImpactAssessment})\} & (108) \\
\text{AssuranceSecDA} &\leftarrow \forall \text{hasSubSecDA}. \text{AssuranceSecDA} & \\
&\quad \wedge =_1 \text{hasAssurance}. \top & (109) \\
\text{AssuranceSecDAI} &\leftarrow =_1 \text{hasAssurance}. \top & \\
&\quad \wedge =_1 \text{hasImplementationManifest}. \top & (110) \\
\text{SecDA} &\leftarrow =_1 \text{hasStrategy}. \top & \\
&\quad \wedge \leq_1 h. \top \text{ for all } h \text{ in} & \\
&\quad \{\text{hasLocationOfIntervention}, \text{hasAssumption}, & \\
&\quad \text{hasImplementationManifest}, \text{hasAssurance}\} & \\
&\quad \wedge \exists \text{hasSubSecDA}^*. \text{SecDAI} & \\
&\quad \wedge (\text{NotTopSecDA} \vee \text{TopSecDAAndOneTestReport}) & \\
&\quad \wedge (\text{TopSecDA} \vee \text{NotTopSecDAAndNoTestReport}) & \\
&\quad \wedge (\text{NotLeafSecDA} \vee \text{LeafSecDAAndOneA}) & \\
&\quad \text{for all } A \text{ in} & \\
&\quad \{\text{InterventionMethod}, & \\
&\quad \text{OrganisationalContextOfIntervention}, & \\
&\quad \text{SafetyImpactAssessment}, & \\
&\quad \text{SecurityImpactAssessment}\} & \\
&\quad \wedge (\text{LeafSecDA} \vee \text{NotLeafSecDAAndNoA}) & \\
&\quad \text{for all } A \text{ in} & \\
&\quad \{\text{Assumption}, & \\
&\quad \text{InterventionMethod}, & \\
&\quad \text{LocationOfIntervention}, & \\
&\quad \text{OrganisationalContextOfIntervention}, & \\
&\quad \text{SafetyImpactAssessment}, & \\
&\quad \text{SecurityImpactAssessment}\} & (111) \\
\text{SecDAI} &\leftarrow =_1 \text{hasImplementationManifest}. \top & (112)
\end{aligned}$$

Figure 18. Shape constraints for secure design argument aspects.

Disjointness Finally, the disjointness axioms in Figure 19 operate at two conceptual levels. First, pairwise disjointness is imposed on the ontology’s primitive classes, i.e., classes that are not introduced as subclasses of other domain classes and therefore serve as basic modeling categories. Second, further disjointness axioms are stated for specialized classes that refine more general concepts. In particular, *AnalyzedSecurityRisk* and *ControlledSecurityRisk*, both subsumed by *SecurityRisk*, are required to be disjoint, and *Impact* and *SecDA*, both subsumed by *ImpactAssessmentDomain*, are likewise required to be disjoint. Together, these axioms exclude unintended class overlap both among primitive categories and among their specialized refinements.

$$\begin{aligned}
& A_i \sqcap A_j \sqsubseteq \perp \text{ for all distinct } A_i, A_j \text{ in} \\
& \quad \{ \text{Assumption, Assurance, Asset,} \\
& \quad \text{AttackType, AttackerProfile, Component,} \\
& \quad \text{ContextSpecificThreat, DataSecurityAndPrivacy,} \\
& \quad \text{IdentityAndAccessManagement, ImpactAssessment,} \\
& \quad \text{ImpactAssessmentDomain, ImpactLevel,} \\
& \quad \text{ImplementationManifest, InterventionMethod,} \\
& \quad \text{LikelihoodFactor, LocationOfIntervention,} \\
& \quad \text{NetworkSecurity, OperatingEnvironment,} \\
& \quad \text{OperationsAndMaintenance,} \\
& \quad \text{OrganisationalContextOfIntervention,} \\
& \quad \text{PhysicalSecurity, Product, ProtectionGoal, Rationale,} \\
& \quad \text{SecurityRisk, SecurityRiskLevel, Strategy,} \\
& \quad \text{SystemAndPlatformSecurity, TestReport,} \\
& \quad \text{ExposureLevel} \} \tag{113} \\
& \text{AnalyzedSecurityRisk} \sqcap \text{ControlledSecurityRisk} \sqsubseteq \perp \tag{114} \\
& \text{Impact} \sqcap \text{SecDA} \sqsubseteq \perp \tag{115}
\end{aligned}$$

Figure 19. Disjointness axioms.

5. Usage & Extensibility




Figure 20. Overview of the SECUMAN validation pipeline.

Figure 20 shows the validation pipeline for risk management documentation of medical devices proposed in [2], which we adapt here for risk management documentation concerning security aspects. As noted above, the approach assumes that risk documentation is submitted as HTML files containing risk management data encoded as RDF triples and expressed using the terminology of the SECUMAN ontology. An *RDF distiller* extracts




Figure 21. Partial graphical representation of the RDF for the security risk from Figure 1.

the RDF graph from the HTML by removing the markup and isolating the RDF triples – see Figure 21 for a graphical depiction of the RDF encoding the security risk rendered via HTML in Figure 1. The RDF triples are then used as input to an OWL ($\mathcal{EL}$ profile) *reasoner* together with SECUMAN and, where needed, additional ontologies. As a result, a materialized knowledge base is obtained and subsequently validated against the SECUMAN SHACL constraints, and potentially additional constraints, using a *SHACL validator*. The result is then communicated in the form of a human-readable *validation report*.

In our extension of the validation pipeline for RISKMAN (see the URL in Section 4.1), we in particular use the HermiT reasoner [46] for materialization, although additional options are provided through a realization wrapper that we also make available online⁴. For SHACL validation, we use PySHACL [54].

Like RISKMAN, the SECUMAN ontology and shapes are intentionally lightweight so that the validation pipeline can be adapted easily to different modelling needs by incorporating additional ontologies and shapes, as already indicated above. One area in which the ontology will clearly require extension is the modelling of security risk levels, for example attacker profiles, exposure levels, and impact levels, as well as the way in which these are combined to determine whether a risk level is acceptable or unacceptable. A minimalist “plugin” for adding such levels to the SECUMAN ontology is also provided in our validation pipeline. In this setting, the user may enrich the SECUMAN ontology with an ontology $\mathcal{K}_{\pi,\sigma,\tau}^{a-v-i} := \mathcal{T}_{gt} \cup \mathcal{A}_{\pi,\sigma,\tau}$ for π attacker profiles, σ exposure levels, and τ impact levels ($\pi, \sigma, \tau \in \mathbb{N}$). This ontology not only introduces individual profiles and levels represented by nominals, but also an ordering gt over these profiles and levels:

⁴<https://github.com/cl-tud/realization-wrapper>

$$\begin{aligned}
\mathcal{A}_{\pi, \sigma, \tau} = & \{ \text{AttackerProfile}(a_i) \mid 1 \leq i \leq \pi \} \cup \\
& \{ \text{ExposureLevel}(e_i) \mid 1 \leq i \leq \sigma \} \cup \\
& \{ \text{ImpactLevel}(s_i) \mid 1 \leq i \leq \tau \} \cup \\
& \{ \text{gt}(a_{i+1}, a_i) \mid 1 \leq i < \pi \} \cup \\
& \{ \text{gt}(e_{i+1}, e_i) \mid 1 \leq i < \sigma \} \cup \\
& \{ \text{gt}(i_{i+1}, i_i) \mid 1 \leq i < \tau \} \\
\mathcal{T}_{\text{gt}} = & \{ \text{tra}(\text{gt}) \}
\end{aligned}$$

To encode a concrete security risk matrix – that is, the way in which risk levels are derived from combinations of attacker profile, exposure level, and impact level – further axioms would need to be added to $\mathcal{T}_{\text{gt}}$, since such matrices are typically manufacturer-specific and therefore do not form part of SECUMAN. As an example, the following axioms first identify critical attacker profiles, exposure levels, and impact levels, and then use these classifications in the last two GCIs to determine critical likelihood factors and, subsequently, critical security levels:

$$\begin{aligned}
& \exists \text{gt}.\{a_2\} \sqsubseteq \text{CriticalAP} \\
& \exists \text{gt}.\{e_3\} \sqsubseteq \text{CriticalEL} \\
& \exists \text{gt}.\{i_2\} \sqsubseteq \text{CriticalIL} \\
& \exists \text{hasAttackerProfile.CriticalAP} \sqcap \exists \text{hasExposureLevel.CriticalEL} \sqsubseteq \text{CriticalLF} \\
& \exists \text{hasLikelihoodFactor.CriticalLF} \sqcap \exists \text{hasImpactLevel.CriticalIL} \sqsubseteq \text{CriticalRiskLevel}
\end{aligned}$$

Checking for controlled security risks with critical residual risk levels can then be achieved by adding the SHACL constraint

$$\text{ControlledSecurityRisk} \leftarrow \neg(\exists \text{hasResidualSecurityRiskLevel.CriticalRiskLevel})$$

to the SECUMAN shapes. Of course, more complex conditions for deriving and detecting critical security risk levels can likewise be expressed by adding further axioms and shapes.

6. Conclusion

We presented the SECUMAN ontology and accompanying SHACL shapes for representing and analysing cybersecurity risk-management documentation for medical devices. Analysed risks and their mitigations are represented as an $\mathcal{EL}^{++}$ ABox; the SECUMAN ontology is used together with an $\mathcal{EL}$ reasoner to infer implicit knowledge; and SHACL constraints are used to check whether the resulting data conform to the intended documentation model. The ontology, shapes, and a reference implementation of the complete validation pipeline are freely available.

The SECUMAN ontology complements the existing RISKMAN ontology by targeting security rather than safety risks of medical devices. It was designed with the ultimate objective of enabling the combined analysis of safety and security risk-management documentation. This objective is reflected not only in the similar structure of the RISKMAN and SECUMAN ontologies, but also, for instance, in the explicit safety-

impact assessments associated with both analysed security risks and SecDAs in the SECUMAN ontology. The next step is therefore to develop a bridging ontology for safety and security aspects. Such an ontology should make it possible, for example, to propagate a safety impact identified for a security control measure into the safety-risk-management domain, where the resulting risk can then be analysed and controlled.

Both RISKMAN and SECUMAN are lightweight ontologies intended to support first-pass validation and, in particular, to facilitate the analysis, navigation, and reuse of medical-device risk-management documentation. More substantive automated analysis will require integration with additional domain ontologies and structured cybersecurity resources; see, e.g., [55,56,57,58,59,60,61], as well as the existing medical-device-specific cybersecurity ontologies and knowledge graphs discussed in Section 2. Quality assurance is a particularly significant challenge [62,63]; possible techniques include dialogue-inspired explication methods from computational argumentation [64,65], as well as other approaches for providing justifications in logic-based knowledge-representation formalisms [66].

Finally, the systematic reuse of risk-management documentation remains an important open challenge. We believe that a shared, Semantic Web-based repository of risk-management documentation, including reusable secure and safe design arguments, could substantially accelerate both the creation and the evaluation of such documentation. This potential is especially promising when combined with large language models in retrieval-augmented generation setups [67,68]. Realising this vision, however, will require appropriate incentives for sharing risk-management knowledge, as well as infrastructure that supports controlled access and reuse without exposing manufacturer secrets.

Acknowledgements.

This work was supported by funding from BMFTR (Federal Ministry of Research, Technology and Space) within the project SEMECO (grant no. 03ZU1210BG).

References

- [1] Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices; 2017. Available from: <http://data.europa.eu/eli/reg/2017/745/oj>.
- [2] Gorczyca P, Arndt D, Diller M, Hampe J, Heidenreich G, Kettmann P, et al. Supporting risk management for medical devices via the RISKMAN ontology and shapes. In: SEMANTiCS. Studies on the Semantic Web. IOS Press; 2025. p. 226-46.
- [3] International Organization for Standardization. ISO 14971:2019: Medical devices – application of risk management to medical devices; 2019. Published by ISO, Geneva, Switzerland. Available from: <https://www.iso.org/standard/72704.html>.
- [4] Arndt D, Bank P, Gorczyca P, Heidenreich G, Kettmann P, Kobel I, et al.. VDE SPEC 90025 V1.0 (en): Medical devices – risk management: Framework of a computerized risk analysis format for transmission and submission (MD-CRAFTS); 2024. VDE Verband der Elektrotechnik Elektronik Informationstechnik e.V. Available from: <https://www.vde.com/resource/blob/2323050/6df908c47e69cd01126530ab07e78a52/vde-spec-90025-v1-0--en--data.pdf>.
- [5] Kang M, Park E, Cho BH, Lee KS. Recent patient health monitoring platforms incorporating internet of things-enabled smart devices. International neurology journal. 2018;22(Suppl 2):S76.
- [6] Schwartz S, Ross A, Carmody S, Chase P, Coley SC, Connolly J, et al. The evolving state of medical device cybersecurity. Biomedical instrumentation & technology. 2018;52(2):103-11.
- [7] AP News. Cyberattack hits major hospital in Spanish city of Barcelona; 2023. Accessed: April 27, 2026. Available from: <https://apnews.com/article/barcelona-hospital-cyberattack-ransomware-37e0fee33798c56459e63866ca8b449f>.

- [8] Campbell D, Hern A. Services disrupted as London hospitals hit by cyber-attack; 2024. The Guardian. Available from: <https://www.theguardian.com/society/article/2024/jun/04/cyber-attack-london-hospitals>.
- [9] Mashinchi MI, Acton T, Datta PM. When healthcare becomes sick: Recovering from ransomware. *Journal of Information Technology Teaching Cases*. 2024;20438869241279443.
- [10] Alder S. 2024 healthcare data breach report; 2025. The HIPAA Journal. Available from: <https://www.hipaajournal.com/2024-healthcare-data-breach-report/>.
- [11] Alder S. Healthcare data breach statistics; 2026. The HIPAA Journal. Available from: <https://www.hipaajournal.com/healthcare-data-breach-statistics/>.
- [12] ENISA. Health; 2025. Accessed: April 27, 2026. Available from: <https://www.enisa.europa.eu/topics/cybersecurity-of-critical-sectors/health>.
- [13] European Commission. Cybersecurity in healthcare; 2026. Accessed: April 27, 2026. Available from: https://commission.europa.eu/topics/digital-economy-and-society/cybersecurity-healthcare_en.
- [14] Medical Device Coordination Group. MDCG 2019-16: Guidance on cybersecurity for medical devices; 2020. MDCG 2019-16.
- [15] International Electrotechnical Commission. IEC 81001-5-1:2021: Health software and health IT systems safety, effectiveness and security – part 5-1: Security – activities in the product life cycle; 2021. IEC 81001-5-1:2021. Available from: <https://www.iso.org/standard/76097.html>.
- [16] International Medical Device Regulators Forum. IMDRF/CYBER WG/N60FINAL:2020: Principles and practices for medical device cybersecurity; 2020. IMDRF/CYBER WG/N60FINAL:2020.
- [17] Association for the Advancement of Medical Instrumentation. AAMI TIR57:2016/(R)2023: Principles for medical device security – risk management; 2016. Reaffirmed 2023.
- [18] Martinez J, Godot J, Ruiz A, Balbis A, Ruiz Nolasco R. Safety and security interference analysis in the design stage. In: *International Conference on Computer Safety, Reliability, and Security*. Springer; 2020. p. 54-68.
- [19] Association for the Advancement of Medical Instrumentation. ANSI/AAMI SW96:2023: Standard for medical device security – security risk management for device manufacturers; 2023. ANSI/AAMI SW96:2023.
- [20] Joint Task Force Transformation Initiative. NIST Special Publication 800-30 Revision 1: Guide for conducting risk assessments. National Institute of Standards and Technology; 2012. NIST SP 800-30 Rev. 1.
- [21] Uciteli A, Neumann J, Tahar K, Saleh K, Stucke S, Faulbrück-Röhr S, et al. Ontology-based specification, identification and analysis of perioperative risks. *J Biomed Semant*. 2017;8(1):36:1-36:14.
- [22] Herre H. General formal ontology (GFO): A foundational ontology for conceptual modelling. In: Poli R, Healy M, Kameas A, editors. *Theory and Applications of Ontology: Computer Applications*. Dordrecht: Springer Netherlands; 2010. p. 297-345.
- [23] Herre H, Heller B, Burek P, Hoehndorf R, Loebe F, Michalek H. General formal ontology (GFO): A foundational ontology integrating objects and processes. Part I: Basic principles (version 1.0). University of Leipzig: Research Group Ontologies in Medicine (Onto-Med); 2006. Onto-Med Report.
- [24] Kim DY, Park YS, Lee B, Lee JW. Ontology-based process integration incorporating reference associations between medical standards from the perspective of medical software developers. *Journal of Ambient Intelligence and Humanized Computing*. 2019.
- [25] Schütz AE, Fertig T, Weber K. Defining a core ontology for medical devices in Germany to ensure semantic interoperability. In: Simian D, Stoica LF, editors. *Modelling and Development of Intelligent Systems – 7th International Conference, MDIS 2020, Sibiu, Romania, October 22-24, 2020, Revised Selected Papers*. vol. 1341 of *Communications in Computer and Information Science*. Springer; 2020. p. 394-410.
- [26] Chatteraj S, Walid R, Joshi KP. Semantically rich approach to automating regulations of medical devices. In: *ICDH*. IEEE; 2024. p. 132-7.
- [27] Zhu W, Zhang P, Xia W, Gao Z, Li W, Tian R, et al. AI-driven medical device risk management: A new paradigm integrating large language models and prompt engineering for standard-risk knowledge graph construction and application. *Risk Management and Healthcare Policy*. 2026;19:571156. Available from: <https://www.tandfonline.com/doi/abs/10.2147/RMHP.S571156>.
- [28] Rosa FdF, Bonacin R, Jino M. The security assessment domain: A survey of taxonomies and ontologies. *arXiv preprint arXiv:170609772*. 2017.

- [29] Adach M, Hänninen K, Lundqvist K. Security ontologies: A systematic literature review. In: International Conference on Enterprise Design, Operations, and Computing. Springer; 2022. p. 36-53.
- [30] Hollerer S, Sauter T, Kastner W. A survey of ontologies considering general safety, security, and operation aspects in OT. *IEEE Open Journal of the Industrial Electronics Society*. 2024;5:861-85.
- [31] Jarwar MA, Freng JWC, Ali S. Modelling industrial IoT security using ontologies: A systematic review. *IEEE Open Journal of the Communications Society*. 2025.
- [32] Lourenço B, Adão P, Ferreira JF, Marques MM, Vaz C. Structuring security: A survey of cybersecurity ontologies, semantic log processing, and LLMs application. *CoRR*. 2025;abs/2510.16610.
- [33] Sills M, Ranade P, Mittal S. Cybersecurity threat intelligence augmentation and embedding improvement – a healthcare use case. In: *IEEE International Conference on Intelligence and Security Informatics (ISI)*. IEEE; 2020. p. 1-6.
- [34] Hannou F, Atigui F, Lammari N, Cherfi SS. SafecareOnto: A cyber-physical security ontology for healthcare systems. In: *DEXA (2)*. Lecture Notes in Computer Science. Springer; 2021. p. 22-34.
- [35] Dart M, Ahmed M. CYBER-AIDD: A novel approach to implementing improved cybersecurity resilience for large Australian healthcare providers using a Unified Modelling Language ontology. *Digital Health*. 2023;9:20552076231191095.
- [36] Liu J, Sasaki R, Uehara T. An ontology-based framework for medical IoT forensic evidence. In: *QRS Companion*. IEEE; 2023. p. 863-4.
- [37] Bughio KS, Cook DM, Shah SAA. Developing a novel ontology for cybersecurity in Internet of Medical Things-enabled remote patient monitoring. *Sensors*. 2024;24(9):2804.
- [38] Bughio KS, Cook DM, Shah SAA. Novel knowledge graph-based modeling for vulnerability detection in the Internet of Medical Things. In: *Recent Challenges in Intelligent Information and Database Systems*. Communications in Computer and Information Science. Springer; 2024. p. 314-25.
- [39] Ghosh R, von Stockhausen H, Schmitt M, Marica VG, Karn SK, Farri O. CVE-LLM: Ontology-assisted automatic vulnerability evaluation using large language models. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39. AAAI Press; 2025. p. 28757-65.
- [40] Chander AS, Harshavardhana M, Sathyanarayana N, Hemalatha K, Gagana B. Intelligent cybersecurity ontology framework for Internet of Medical Things-enabled remote patient monitoring. In: *ITM Web of Conferences*. vol. 79. EDP Sciences; 2025. p. 01030.
- [41] Adida B, Birbeck M, McCarron S, Herman I. RDFa core 1.1 – third edition, 17 March 2015; 2015. W3C Recommendation. Available from: <http://www.w3.org/TR/rdfa-core/>.
- [42] International Organization for Standardization. ISO/IEC/IEEE 15026-1:2019: Systems and software engineering – systems and software assurance – part 1: Concepts and vocabulary; 2019. ISO/IEC/IEEE 15026-1:2019.
- [43] International Organization for Standardization. ISO/TS 81001-2-1:2025: Health software and health IT systems safety, effectiveness and security – part 2-1: Coordination – guidance and requirements for the use of assurance cases for safety and security; 2025. ISO/TS 81001-2-1:2025. Available from: <https://www.iso.org/standard/82210.html>.
- [44] Poveda-Villalón M, Fernández-Izquierdo A, Fernández-López M, García-Castro R. LOT: An industrial-oriented ontology engineering framework. *Engineering Applications of Artificial Intelligence*. 2022 may;111:104755.
- [45] Suárez-Figueroa MC, Gómez-Pérez A, Fernández-López M. The NeOn methodology framework: A scenario-based methodology for ontology development. *Appl Ontology*. 2015;10:107-45. Available from: <https://api.semanticscholar.org/CorpusID:5094899>.
- [46] Shearer RD, Motik B, Horrocks I. HermiT: A highly efficient OWL reasoner. In: *Owled*. vol. 432; 2008. p. 91.
- [47] Poveda-Villalón M, Gómez-Pérez A, Suárez-Figueroa MC. OOPS! (OntOlogy Pitfall Scanner!): An online tool for ontology evaluation. *International Journal on Semantic Web and Information Systems (IJSWIS)*. 2014;10(2):7-34.
- [48] Motik B, Grau BC, Horrocks I, Wu Z, Fokoue A, Lutz C. OWL 2 web ontology language profiles (second edition); 2012. W3C Recommendation. Available from: <https://www.w3.org/TR/owl2-profiles/>.
- [49] W3C. Shapes constraint language (SHACL): Technical report; 2017. Available from: <https://www.w3.org/TR/shacl>.
- [50] Baader F, Brandt S, Lutz C. Pushing the EL envelope. In: Kaelbling LP, Saffiotti A, editors. *IJCAI-05*, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh,

Scotland, UK, July 30 – August 5, 2005. Professional Book Center; 2005. p. 364-9. Available from: <http://ijcai.org/Proceedings/05/Papers/0372.pdf>.

- [51] Baader F, Lutz C, Brandt S. Pushing the EL envelope further. In: Clark K, Patel-Schneider PF, editors. Proceedings of the Fourth OWLED Workshop on OWL: Experiences and Directions, Washington, DC, USA, 1–2 April 2008. vol. 496 of CEUR Workshop Proceedings. CEUR-WS.org; 2008. Available from: https://ceur-ws.org/Vol-496/owled2008dc_paper_3.pdf.
- [52] Corman J, Reutter J, Savkovic O. Semantics and validation of recursive SHACL. In: Vrandečić D, Bontcheva K, Suárez-Figueroa MC, Presutti V, Celino I, Sabou M, et al., editors. The Semantic Web – ISWC 2018 – 17th International Semantic Web Conference, Proceedings, Part I. vol. 11136 of Lecture Notes in Computer Science. Springer; 2018. p. 318-36.
- [53] Andresel M, Corman J, Ortiz M, Reutter J, Savkovic O, Šimkus M. Stable model semantics for recursive SHACL. In: Huang Y, King I, Liu T, van Steen M, editors. Proceedings of The Web Conference 2020. WWW’20. New York, NY, USA: ACM / IW3C2; 2020. p. 1570-80.
- [54] Sommer A, Car N. pySHACL. Zenodo; 2024.
- [55] Syed Z, Padia A, Mathews ML, Finin T, Joshi A. UCO: A Unified Cybersecurity Ontology. In: Proceedings of the AAAI Workshop on Artificial Intelligence for Cyber Security. AAAI Press; 2016. p. 195-202.
- [56] Iannacone M, Bohn S, Nakamura G, Gerth J, Huffer K, Bridges R, et al. Developing an Ontology for Cyber Security Knowledge Graphs. In: Proceedings of the 10th Annual Cyber and Information Security Research Conference. CISRC ’15. Association for Computing Machinery; 2015. .
- [57] Kiesling E, Ekelhart A, Kurniawan K, Ekaputra F. The SEPSES Knowledge Graph: An Integrated Resource for Cybersecurity. In: The Semantic Web – ISWC 2019. vol. 11779 of Lecture Notes in Computer Science. Springer; 2019. p. 198-214.
- [58] Hemberg E, Kelly J, Shlapentokh-Rothman M, Reinstadler B, Xu K, Rutar N, et al.. Linking Threat Tactics, Techniques, and Patterns with Defensive Weaknesses, Vulnerabilities and Affected Platform Configurations for Cyber Hunting; 2021.
- [59] MITRE. CAPEC: Common Attack Pattern Enumeration and Classification;. Available from: <https://capec.mitre.org/>.
- [60] MITRE. CWE: Common Weakness Enumeration;. Available from: <https://cwe.mitre.org/>.
- [61] National Institute of Standards and Technology. National Vulnerability Database;. Available from: <https://nvd.nist.gov/>.
- [62] de la Vara J, Jiménez G, Mendieta R, Parra E. Assessment of the quality of safety cases: A research preview. In: Knauss E, Goedicke M, editors. Requirements Engineering: Foundation for Software Quality – 25th International Working Conference, REFSQ 2019, Essen, Germany, March 18-21, 2019, Proceedings. vol. 11412 of Lecture Notes in Computer Science. Springer; 2019. p. 124-31.
- [63] Foster S, Nemouchi Y, Gleirscher M, Wei R, Kelly T. Integration of formal proof into unified assurance cases with Isabelle/SACM. Formal Aspects Comput. 2021;33(6):855-84.
- [64] Diller M, Gaggli S, Gorczyca P. Flexible dispute derivations with forward and backward arguments for assumption-based argumentation. In: Baroni P, Benz Müller C, Wang YN, editors. Logic and Argumentation – 4th International Conference, CLAR 2021, Hangzhou, China, October 20-22, 2021, Proceedings. vol. 13040 of Lecture Notes in Computer Science. Springer; 2021. p. 147-68.
- [65] Diller M, Gorczyca P. ABA Disputes in ASP: Advancing Argument Games Through Multi-shot Solving. In: PRIMA. vol. 16366 of Lecture Notes in Computer Science. Springer; 2025. p. 566-84.
- [66] Denecker M, Brewka G, Strass H. A formal theory of justifications. In: Calimeri F, Ianni G, Truszczyński M, editors. Logic Programming and Nonmonotonic Reasoning – 13th International Conference, LP-NMR 2015, Lexington, KY, USA, September 27-30, 2015. Proceedings. vol. 9345 of Lecture Notes in Computer Science. Springer; 2015. p. 250-64.
- [67] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In: Advances in Neural Information Processing Systems. vol. 33; 2020. p. 9459-74.
- [68] Fan W, Ding Y, Ning L, Wang S, Li H, Yin D, et al. A Survey on RAG Meeting LLMs: Towards Retrieval-Augmented Large Language Models. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Association for Computing Machinery; 2024. p. 6491-501.