Theme Article: Special Issue on AI Ethics and Trust: From Principles to Practice

Open Requirements Modelling for Compliance and Conformity of Trustworthy AI

Julio Hernandez, ADAPT Centre, School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland.

David Lewis, ADAPT Centre, School of Computer Science and Statistics, Trinity College Dublin, Dublin, Ireland.

Abstract—The many initiatives on trustworthy AI result in a confusing and multipolar landscape that organizations are operating within the fluid and complex international value chains must navigate in pursuing trustworthiness AI. The EU's proposed Draft AI Act will now shift the focus of such organizations toward the normative requirements for regulatory compliance. Understanding the degree to which standards compliance will deliver regulatory compliance for AI remains a complex challenge. This paper offers a simple and repeatable mechanism for extracting and sharing the terms and concepts relevant to normative statements in the legal and standards texts into open knowledge graphs. This representation is used to assess the adequacy of standards conformance to regulatory compliance and thereby provide a basis for identifying areas where further technical consensus development in trustworthy AI value chains will be required to achieve regulatory compliance.

The global interest in AI's societal and ethical risks has grown rapidly in recent years. Academic, governmental, and commercial initiatives have proposed a large array of guidelines for the responsible development of AI[1]. These are typically presented as structured statements of principles that organizations can opt to adopt with the goal of demonstrating some degree of ethical and trustworthy characteristics in their development and use of AI technology. There is, however, increasing recognition by public authorities that there is a wide range of applications through which AI can impact on people lives and that are currently developed and deployed with little external oversight.

Several jurisdictions are now developing legislations that introduces some level of regulatory oversight over AI that offers protection for people and groups in their jurisdiction from potentially harmful impacts of AI applications. Regulatory proposals must balance such protections to develop more open and competitive markets for AI-based products and services, supporting the associated value chains of datasets and AI models to realize AI innovation's economic and societal benefits.

One of the more detailed attempts to develop such a regulatory balance is the Draft AI Act[a], AI Act hereafter, proposed by the European Commission (EC) in April 2021 for scrutiny and enactment by the European Parliament[2]. This proposal specifies a tiered system of risk in different AI application areas. Some areas are proscribed, and others are identified as a sufficiently low risk that only consumer labels or voluntary codes of

[a] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206

practice are required. However, the AI Act also identifies a range of high-risk AI application areas and requires structure-focused risk and quality management measures to comply with regulations. These measures follow the regulatory mechanism, called the New Legislative Framework, developed to provide a single health and safety mechanism for products across the European Single Market. The AI Act extends this mechanism to products and services containing AI and extends the scope of protection beyond health and safety to include all fundamental rights. In this way, the legislation aims to support the development and deployment of Trustworthy AI. However, while the EU has separately provided guidelines for developing trustworthy AI[3], these do not form part of the AI Act, which instead delegates the detailed rulemaking on how risks are assessed, managed, and monitored to technical standards. These can be in the form of standards that have been harmonized with the requirement of the AI Act by a European Standardization Organization (ESO), namely CEN[b], CENELEC[c], or ETSI[d]. Relevant standards are already being addressed internationally by standards development organizations such as ISO/IEC JTC1 Subcommittee 42 on AI (SC42) and the IEEE P7000 series on ethical autonomous and intelligent systems[4]. However, these standards development initiatives involve complex sets of interrelated standards, many of which are still under development[5] and will be evolving in parallel to the AI Act and similar legislations being considered in other jurisdictions.

Following the mechanisms established in the New Legislative Framework, providers of high-risk AI applications must demonstrate their compliance with the relevant requirements of the AI Act through a conformity assessment process that is either self-certified or certified by a recognized authority, known as a Notified Body. The conformity assessment process must address AI Act requirements related to risk management, data governance, and technical documentation under a quality management system for the compliance of the product to be certified. This mechanism is well aligned with the standards developed by the ISO committee on conformity assessment (CASCO) through the ISO 17000 series of standards that provides guidance on the terminology and concepts, requirements, processes, and competencies that regulators can use in establishing certification rules. These standards are then complemented by standards that an organization can follow to implement the risk and quality systems compatible with CASCO-defined certification, known as Management System Standards (MSS). ISO/IEC JTC1 SC42 is developing a MSS for AI, ISO/IEC 42001[e]. Therefore, this alignment with an existing standardized conformity framework means that this AI MSS and other SC42 standards references form a strong candidate for adoption by ESO in response to a harmonized standards request from the EC for the AI Act. The SC42 international standards will therefore provide some guidance to the developers aiming to place AI systems onto the European Market in compliance with the AI Act.

However, several challenges remain when considering the vertical nature of the AI Act's high-risk categorization, the potentially complex value chains involved, and the international nature of AI innovation. Firstly, the AI Act focuses its provisions for high-risk AI based on a specific set of applications of AI systems, categorized into two groups of AI applications. One is those already subject to specific European product health and safety regulations, e.g., in products such as machinery, toys, medical devices, agricultural vehicles, rail systems, etc. The other is AI applications that are not yet regulated but are identified by the EC as presenting high risks to health, safety, or fundamental rights. However, the technical requirements for compliance with the AI Act and the potential harmonized standards from SC42 are horizontal, i.e., they are specified in terms that apply to any form of AI system. For instance, if we consider the risk of a voice recognition system misunderstanding the same utterance in different accents, the level of acceptable risk when used in ambulance dispatch may involve different considerations from use in primary school student assessment.

Secondly, many AI providers may already be undertaking some form of proprietary trustworthy AI risk assessment and quality process, e.g., Microsoft[6]. Such AI providers will need to undertake a mapping to assess whether the proprietary approach fully satisfies the requirements of the AI Act. They may also wish to establish a transition mapping from the proprietary standard to the relevant harmonized standard to reduce the cost of demonstrating compliance with the AI Act and improving the potential for establishing such compliance, and thereby its trustworthy AI competencies to its customers and to affected societal stakeholder more broadly.

Thirdly, there may be populations of AI providers that have invested in undertaking a trustworthy AI risk and quality assessment based on standards from national

[b] Comité Européen de Normalisation.

[c] Comité Européen de Normalisation Electrotechnique.

[d] European Telecommunications Standards Institute

[e] https://www.iso.org/standard/81230.html

bodies, e.g., NIST[f], DIN[g], BSI[h], or other international standards, e.g., P7000. Mapping between such standards and AI Act's harmonized standards may be important for AI providers to manage the cost of maintaining compliance with regulations in multiple jurisdictions. Providing such mappings could also support future equivalence agreements for trustworthy AI compliance between the EU and other jurisdictions regulating AI.

The evolving nature of international standards for trustworthy AI exacerbates these requirements and compliance mapping challenges. The need for the harmonization request for the AI Act to be satisfied by European SDOs[i] that are not current driving those standards and the proliferation of other proprietary, international, and national guidelines and standards for trustworthy AI. This paper presents an open approach to capturing requirements from different regulations and associated standards documents so that the sufficiency of the local management process and resulting artifact exchanges between value chain actors can be compared and compliance with different regulatory and policy requirements can be assessed and tracked.

SEMANTIC MODELLING OF TRUSTWORTHY AI REQUIREMENTS

Any mapping between regulatory compliance requirements for trustworthy AI and technical standards that enable conformance and certification functions that satisfy those requirements will require flexible, extensible, transparent to third parties, and auditable solutions to satisfy regulatory and organizational rules on governance process integrity. Open standards should be used as far as possible to increase third-party inspection and, therefore, confidence in the completeness and accuracy of mapping. We take an approach based on Open Knowledge Graphs (OKG) specified using standards from the W3C[j], which have been provide successful in promoting interoperability between approaches satisfying requirements of the EU Data Protection Regulation (GDPR)[7] and expressing high-risk information through an AI risk ontology based on the requirements of the AI Act and ISO 31000 series of standards[8]. Such OKG are grounded in the Resource Description Framework (RDF)[9], which allows an unlimited knowledge graph of nodes and links to existing online resources on the web, thus lending themselves to third-party scrutiny. Nodes and associations in this knowledge graph are typed according to ontologies, also known as data vocabularies, that can be developed independently and published to a distinct namespace on the web. This namespace typing allows the free combination of types and associated conceptual knowledge from any published vocabulary. This highly decentralized approach aligns well with the goal of promoting the participation of those generating standards, organizational policies, and regulations, as well as those with interest in how these documents develop and map to each other. OKGs also offer predictable and controlled upgrade paths for expressing compliance rule as new regulation or regulatory guidance and case law emerges, allowing regulatory compliance for trustworthy AI to remain robust and cost-controlled amidst rapid evolution in the relevant regulation.

In developing a semantic model for any specific domain, different levels of semantic commitment can be employed to express semantic relationships between possible information elements. The Web Ontology Language (OWL)[10] allows information elements to be modeled as classes or instances, like object-oriented software engineering models. OWL classes can be structured in hierarchies such that one class can be declared a subclass of another class. Properties can be declared between classes and literal types that allow facts or axioms about the world to be asserted and inferred.

However, Trustworthy AI is a domain with a wide range of competing conceptual models but a relative paucity of concrete instances where trustworthy characteristics have been modeled, tested, and subject to third-party scrutiny. It is, therefore, more appropriate to capture some structure of knowledge without a full understanding of the instances that define the conceptual classes, the relationships between them, and the nature of any hierarchical structures, or we may not necessarily have the goal of checking data model consistency. In such scenarios, the Simple Knowledge Organization System (SKOS)[11] can be used to organize concepts into concept sets and establish hierarchical relationships that are useful to build taxonomies. In SKOS, hierarchical associations are defined as a 'narrower' or 'broader' relationship between concepts, which makes no semantic commitment about these concepts being classes of instances and therefore makes no claims about the relationships of instances. The existence of concept relationships that do not have a hierarchical characteristic can also be captured by a 'related' association between those concepts. SKOS concepts and their associations can

[f] National Institute of Standards and Technology

[g] Deutsches Institut für Normung

[h] British Standards Institution

[i] Standard Development Organizations

[j] World Wide Web Consortium

be grouped into concept sets that can represent the consensus developed on a domain by a group at a particular time.

CONCEPTUAL REQUIREMENTS CAPTURE FROM AI ACT AND PROSPECTIVE HARMONIZED STANDARDS

To address the challenges of mapping normative statements from regulations such as the AI Act against those in standards from different SDOs, we need to be able to catalog the normative statements from these different source documents. We do this in a way that mirrors the granularity of authority and the revision cycles of these separate documents. Specifically, we have analyzed the sections of the AI Act, specifically the compliance requirements for AI Providers for high-risk AI systems, and the terms and concepts defined by SC42 in foundational standards ISO/IEC 22989, as well as the template for ISO MSS, which forms the basis for the development of the AI MSS, ISO/IEC 42001 (Figure 1).

We aim to enable the capture of terms and concepts related to regulatory requirements and standards to which organizations in the AI value chain can conform to demonstrate their compliance with their regulatory obligations. The approach specifically aims to enable the interlinking of requirements between regulatory text and texts specifying such international standards and thereby check the extent to which prospective harmonized standards requirements will deliver regulatory compliance. This requires an analysis of the normative scope of requirements of both the relevant compliance clauses of the AI Act and the MSS template.

Our semantic modeling leverages the core commonality of the harmonized structure for MSS to provide a minimal and reusable approach, determining the extent to the requirements present in normative statements specified in a regulatory text for trustworthy AI are satisfied by normative statements in technical standards documents used in conformance, specifically those stemming from AI MSS. This is taken as a specific assessment of the more general goal to assess whether this approach allows machine-readable mapping for specific proposed trustworthy AI guidelines or standards to be mapped against requirements of specific regulatory text. The target forms of mapping consider:

- Whether all captured regulatory requirements are addressed by available management system requirements or other technical requirements.
- Do regulatory requirements have mappings to specific technical activities or entities/artifacts defined in the technical standards.
- Whether some requirement mapping is partial in that they use a different definition of concepts or different levels of normative strictness, i.e., the requirement (must/shall) compares to a recommendation (should), permission (may), or possibility (can).
- Are there terms in the regulatory requirements for which mapping to technical standard requirements, activities, or entities cannot be fully determined.

The conceptual extraction and mapping process first involves extracting explicitly defined terms as SKOS concepts, using different concept sets for different source documents. The structure of terminological lists (for example, subsection in the terminology section of

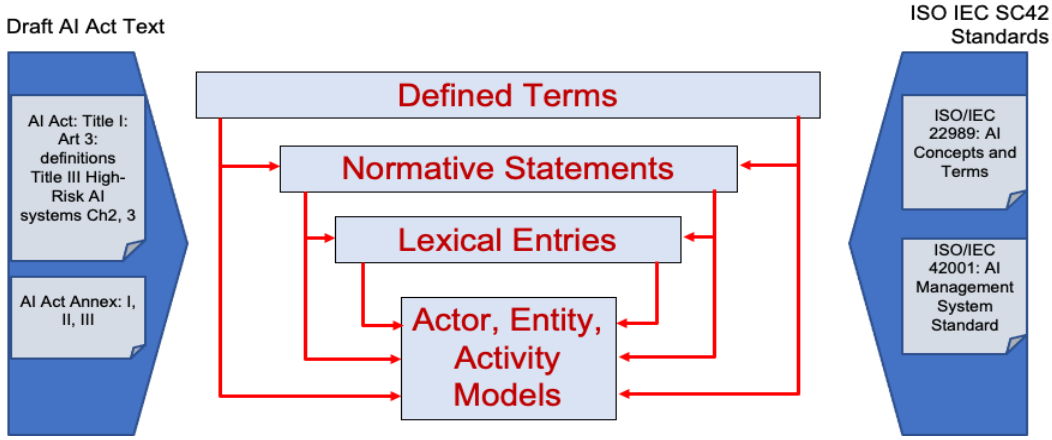

FIGURE 1. Mapping concepts between the AI Act and the ISO/IEC SC42 standards.

ISO/IEC standards), the text of the definitions, and cross references between these are used to capture taxonomical structures, using ‘narrower’, ‘broader’, and ‘related’ relationships. Then, normative clauses of the source documents are converted to individual normative concepts. These are single requirement statements with a specific irreducible requirement on a subject to provide for a specific situation involving named actors, activities, or entities. For example, a clause that requires a subject to undertake a list of activities would be converted to a single normative concept for each, enabling links for previously extracted terminological concepts. Where the requirement or required situation from an individual normative statement conceptualization does not correspond to a term from that same source document, the terms used is captured as a lexical entry, indicating that it is a concept that may require further definition to support future compliance checking. Such lexical entries are, therefore, candidate for alignment with definitions from another document, e.g., from another referenced legislative document or technical standard.

The use of SKOS concepts set for terms and requirements in this approach differs from existing approaches which aim to allow formal expression of a requirement that can be subject to deontic reasoning as part of a requirement management process. Our approach is focused instead on facilitating the mapping between separate developed sets of definitions and compliance standards in a flexible, extensible, and repeatable manner appropriate to the evolving trustworthy AI landscape where compliance rules and associated standards are still under development.

Next, we present the semantic web ontology used to allow this term and requirements conceptualization and mapping to be undertaken.

THE TAIR ONTOLOGY

The Trustworthy AI Requirements (TAIR) ontology[k] provides the elements to describe requirements and concepts associated to a specific ISO standard. Additionally, provides the resources to related requirements or concepts from different standards but with similar purposes, e.g., domain-specific ISO MSS. Figure 2 depicts the TAIR ontology, where

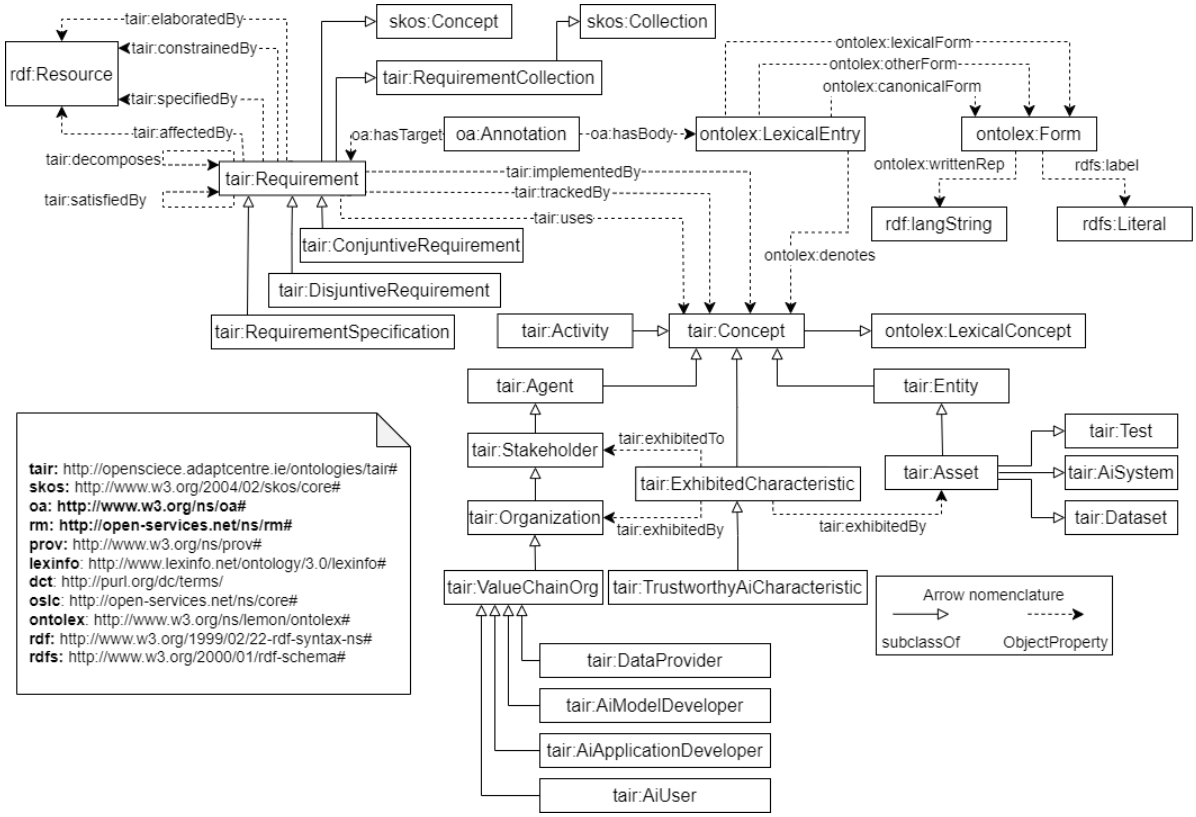

FIGURE 2. The TAIR ontology.

[k] The technical implementation is available at https://tair.adaptcentre.ie

tair:Requirement and tair:Concept are the main classes in the ontology.

The tair:Concept class is a subclass of the ontolex[l] vocabulary, which describes linguistic resources such as the representation of dictionaries or annotations commonly found in lexicography. This class is the parent of any other resource in the TAIR ontology, which means a tair:Agent class or an ISO standard concept will be described by the tair:Concept class.

ISO STANDARDS MAPPING REQUIREMENTS AND CONCEPTS

ISO standards are accessible through a text document that difficulties their implementation for two or more standards from a different, or the same, domain for an organization, e.g., ISO MSS. TAIR ontology aims to map the ISO standards requirements into linked data resources, making them available to consult and query.

The mapping requirements into linked data resources will help to create systems capable of defining the requirements needed to comply with a domain-specific standard, e.g., information security, quality management, etc. Additionally, identify and represent the concepts related to a standard, i.e., the words or phrases defined in the document with a specific meaning.

Figure 3 provides an overview of the mapping process. The mapping process considers the ISO/IEC document structure divided into clauses, expressing requirements to comply with the standard based on the verbal forms of shall and shall not[m]. The mapping steps followed are described next:

- Text processing (Figure 3, S1). Transforms the original standard document, usually in a PDF format, into a set of requirements sentences (considering the verbal form of shall) and concept definitions.
- Text mapping (Figure 3, S2). Describes each requirement sentence and concept definition into a linked data element, defining the relationship between requirements and between concepts (if it exists). The TAIR ontology is used to define the corresponding classes and properties for each case. The tair:RequirementCollection defines the clause of an ISO/IEC standard, e.g., the Context of the organization clause from the ISO/IEC 27001. The tair:Requirement defines a particular requirement from the clause. The tair:Concept describes a particular concept from an ISO/IEC standard, e.g., the concept of “top management” from the ISO/IEC 27001.

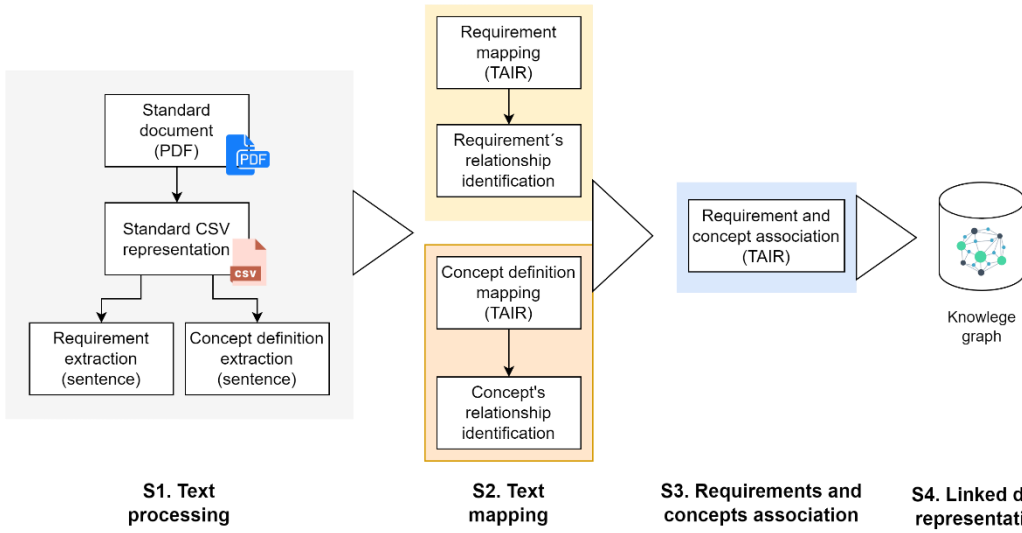

FIGURE 3. The mapping process description divided in four steps (S1-S4)

[l] https://www.w3.org/2019/09/lexicog/

[m] ISO/IEC Directives, Part 2 - https://www.iso.org/sites/directives/current/part2/index.xhtml

- Requirement and concepts association (Figure 3, S3). Determines if a concept is associated with a specific requirement. The properties of tair:uses, tair:constrainedBy, tair:implementedBy defines the relationship between a requirement and concept(s).
- Linked data representation (Figure 3, S4). Provides the mechanisms to consult and query the mapped requirements and concepts.

CONCEPTUAL MAPPING BETWEEN AI ACT AND ISO/IEC AI MSS

The TAIR ontology could be used to represent semantic relationships between standards. In this case, the mapping process considers the semantic relationship between the AI Act and the ISO/IEC 42001 standard (AI MSS).

The mapping process considers the provider's requirements from the AI Act, semantically relating them to the AI MSS. Figure 4 illustrates the semantic relationship between AI Act articles and AI MSS clauses. These relationships are drawn through the skos:related predicate.

The extraction of requirements from the AI Act related to compliance obligations on AI providers, specifically, resulted in 118 separate requirements, each one captured as a SKOS concept per the TAIR ontology. Where relevant, these are linked to the 46 explicitly defined concepts from Article 3 of the draft. Additionally, a further 23 concepts are used in those requirements that are not explicitly defined in the AI Act (Table 1).

The defined concepts were extracted from potential harmonized standards from SC42, specifically ISO/IEC 22989, AI Terms and concepts, and ISO/IEC 42001 AI MSS. The former yielded 106 defined concepts, grouped under collection categories of Artificial Intelligence; Machine Learning; Neural Networks; Trustworthiness, and Natural Language Processing. Notably, the only terms from this set that coincides with terms defined or extracted from the AI Act are AI System; Risk; Training Data; Validation Data; Testing Data. This reflects the more technical focus on the SC42 terminology, compared to the AI Act, which is couched more in terms related to the compliance and conformance mechanism and the application rather than the technical form of AI. A further 20 concepts was extracted from the MSS HS with a further five added in the AI MSS, of which only the concept of Risk coincided with the AI Act terms, though Documented Information is close to the AI Act term of Technical Documentation.

This relative lack of coincident terminology indicates that careful terminological mapping will be needed to ensure the coherent use of these SC42 standards for the conformance aspects of AI Act compliance if adopted as harmonized standards by the ESO. However, it also points to an appropriate distinction of scope between legal compliance concerns and standardization of the two concept sets, which may minimize the likelihood of conflicting or incompatible concepts. Within the two sets of requirements extracted separately from the AI Act and the MSS HS, we have been able to capture the conceptual relationship between requirements within each set which is helpful for their organization and navigation. However, mapping requirements between the two sets to assess the satisfaction of legal compliance by standards conformity will require analysis of further specifications. This will be needed in terms of the sectorial compliance knowledge relevant to the AI application area defined as high risk in the AI Act, which provides concrete details that horizontal normative statements are unable to achieve alone. It will also need to be supported through mapping to more detailed normative concepts and requirements from SC42, both in terms of the management system control requirements being finalized in the annexes of ISO/IEC 42001 and the further SC42 normative specifications being developed and referenced by those control statements.

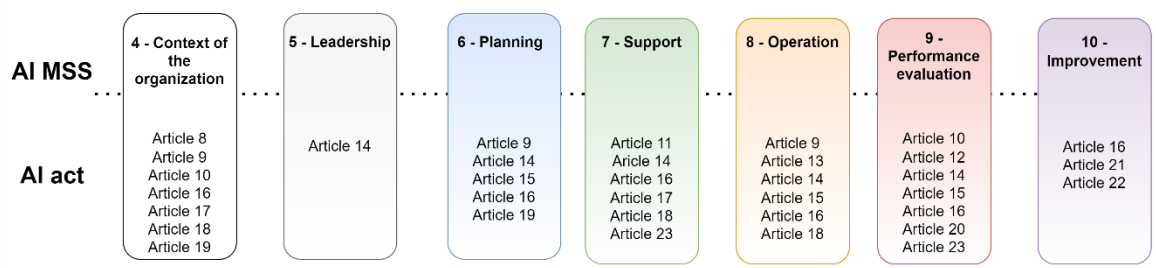

FIGURE 4. Provider´s requirements relationship from AI Act and the ISO 42001 (AI MSS)

CONCLUSIONS AND FUTURE WORK

The Trustworthy AI Requirements (TAIR) ontology provides a basis for capturing and analyzing terms and requirements as concept sets taken from normative statements respectively from the AI Act and from the conformance focused international standard on AI from SC42. This is made partially available as an OKG resource that allows the links between defined terms, other relevant concepts, and the requirements themselves to be established. In the future, we will invite subject matter experts in specific domains, such as healthcare, to explore whether these links can be resolved to establish AI Act compliance through harmonized standards conformance when applied to those domains. Further horizontal requirements mapping will be explored, especially as the SC42 AI MSS is supported by further standards in risk management (ISO/IEC 23894), ML robustness testing (ISO/IEC 24029-2), AI quality model (ISO/IEC 25059) and ML data quality (ISO/IEC 5259 series). Longer term, this approach, and resources could be used for comparing existing proprietary or national trustworthy AI mechanisms to the conformance and compliance system offered by the AI Act and its harmonized standards. Such mapping resources could also assist civil society organizations to monitor the future implementation and enforcement of the AI Act, especially in relation to fundamental rights protections. Fully realizing this potential would, however, require agreement between the EC and ESO on how harmonized standards can be made publicly available without the current paywall fees.

ACKNOWLEDGMENTS

This project has received funding as a research gift from Meta and is supported by the Science Foundation Ireland under Grant Agreement No 13/RC/2106_P2 at the ADAPT SFI Research Centre and the European Union's Horizon 2020 Marie Skłodowska-Curie grant agreement No 813497 for the PROTECT ITN.

REFERENCES

1. A. Jobin, M. Ienca, and E. Vayena,, "The global landscape of AI ethics guidelines". Nat Mach Intell 1, 389-399 (2019). DOI: 10.1038/s42256-019-0088-2
2. Proposal for a Regulation of the European Parliament and of the Council laying down harmonised rules on Artificial intelligence (Artificial Intelligence Act) and amending certain union legislative actions, COM(2021) 206 final,

TABLE 1. Concepts defined in the AI Act or extracted from its compliance requirements statement.

Actor Concepts	**Artefact Concepts**	**Process Concepts**
Concepts from Ai act definitions (Article 3)		
Provider; Small Scale Provider; User, Authorised Representative; Importer; Distributor; Operator; Notifying Authority; Conformity Assessment Body; Notified Body; Market Surveillance Authority; Law Enforcement Authority; National Supervisory Authority; National Competent Authority	AI System; Intended Purpose; Reasonably Foreseeable Misuse; Safety Component; Instructions For Use; Performance; Substantial Change; CE Marking; Harmonized Standard; Common Specification; Training Data; Validation Data; Testing Data; Input Data; Biometric Data; EmotionRecognitionSystem; Biometric Categorization System; Remote Biometric Identification System; Real Time Remote Biometric Identification System; Post Remote Biometric Identification System; Publicly Accessible Space; Serious Incident Infrastructure	Placing On The Market; Making Available On The Market; Putting Into Service; Recall; Withdrawal; Conformity Assessment; Post Market Monitoring; Law Enforcement
Concepts from extraction of AI Provider compliance (chapter 2, on requirements for high-risk Ai systems, articles 8 to 13, and 15; and chapter 3, on obligations for AI providers, articles 16 to 23)		
children	high risk AI system; risk management system; risk management measures; harmonized standards; residual risk; preliminarily defined metrics; preliminarily defined probabilistic thresholds; technical documentation; logs; quality management system; quality assurance system; risk	continuous iterative process; eliminating or reducing risks; perform consistently for their intended purpose; testing of high-risk AI systems; testing procedures; development process; placing on the market; putting into service; children; human oversight; quality control; and resource management

Brussels, 21.4.2021
3. HLEG (2019) European Commission's High Level Expert Group, "Ethics Guidelines for Trustworthy AI", April 2019, https://ec.europa.eu/futurium/en/ai-alliance-consultation/guidelines
4. ISO/IEC JTC 1 (2022) ISO/IEC JTC 1 Information Technology, December 2022. https://www.iso.org/committee/45020.html.
5. AI Watch: AI Standardization Landscape (v. 2.0) state of play and link to the EC proposal for an AI regulatory framework. (2021)
6. Microsoft Responsible AI Standard, v2 GENERAL REQUIREMENTS, June 2022, https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-Responsible-AI-Standard-v2-General-Requirements-3.pdf
7. H. J. Pandit, D. O'Sullivan, and D. Lewis, "Queryable Provenance Metadata For GDPR Compliance", in Procedia Computer Science, ser. Proceedings of the 14th International Conference on Semantic Systems 10th – 13th of September 2018 Vienna, Austria, vol. 137, Jan. 1, 2018, pp. 262–268. DOI: 10/gfdc6r. Available: http://www.sciencedirect.com/science/article/pii/S1877050918316314 (visited on 16/12/2022).
8. G. Delaram, H. J. Pandit, D. Lewis, "AIRO: an Ontology for Representing AI Risks based on the Proposed EU AI Act and ISO Risk Management Standards", International Conference on Semantic Systems (SEMANTiCS), Vienna, Austria, 13-15 September 2022, 2022
9. RDF 1.1 Primer,. Available: https://www.w3.org/TR/rdf11-primer/ (visited on 16/12/2022).
10. OWL 2 Web Ontology Language Document Overview (Second Edition), W3C Recommendation 11 December 2012, http://www.w3.org/TR/owl2-overview/
11. SKOS Simple Knowledge Organization System Reference, W3C Recommendation 18 August 2009, https://www.w3.org/TR/skos-reference/
12. International Organization for Standardization. (2009) ISO/IEC Directives, Part 1 - Consolidated ISO Supplement - Procedure for the technical work - Procedures specific to ISO Retrieved from https://www.iso.org/sites/directives/current/consolidated/index.xhtml
13. Normative Requirements as Linked Data Fabien Gandon, Guido Governatori, Serena Villata doi:10.3233/978-1-61499-838-9-1SeriesFrontiers in Artificial Intelligence and Applications EbookVolume 302: Legal Knowledge and Information Systems
14. Standard, O. A. S. I. S. (2021). OSLC Requirements Management Version 2.1. Part 1: Specification. https://docs.oasis-open-projects.org/oslc-op/rm/v2.1/os/requirements-management-spec.html

Julio Hernandez is a research fellow from Trinity College Dublin. His research interests are semantic web technologies, machine learning, and text mining. He is collaborating in the mapping requirement process from AI Act and ISO 42001. E-mail: julio.hernandez@adatpcentre.ie

Dave Lewis is an Associate Professor at the School of Computer Science and Statistics at Trinity College Dublin, where he served as the head of its Artificial Intelligence Discipline. He is the Interim Director of Ireland's ADAPT Centre for human centric AI and digital content technology research. He investigates open semantic models for trustworthy AI and data governance and contributes to international standards in digital content processing and trustworthy AI at ISO/IEC JTC1/SC42 and CEN/CENELEC JTC21. E-mail: dave.lewis@adaptcentre.ie