



D2.3

Success criteria set by ACM and IEEE

Victoria Wiegand, Delaram
Golpayegani, Marta Lasek-Markey,
Arjumand Younus, Monique
Munarini, Yuening Li, Aphra Kerr,
and Dave Lewis

FORSEE

| Forging Successful AI Applications
| for European Economy and Society



Forging Successful AI Applications
for European Economy and Society

D2.3. Success criteria set by ACM and IEEE

Project Information	
Project Number:	101177579
Project Title:	Forging Successful AI Applications for European Economy and Society - FORSEE
Funding Scheme:	HORIZON-CL2-2024-TRANSFORMATIONS-01-06
Project Start Date:	1 February 2025

Deliverable Information	
Title	Success criteria set by ACM and IEEE
Work Package	WP 2 - Mapping of social expectations: institutional/formal approaches to success
Lead Beneficiary	Trinity College Dublin
Due Date	31/01/2026
Revision Number:	V2.0
Authors	Victoria Wiegand ¹ , Trinity College Dublin Delaram Golpayegani ² , Trinity College Dublin

¹ <https://orcid.org/0009-0002-7211-152X>

² <https://orcid.org/0000-0002-1208-186X>



	<p>Marta Lasek-Markey³, Trinity College Dublin Arjumand Younus⁴, University College Dublin Monique Munarini⁵, Trinity College Dublin Yuening Li⁶, University College Dublin Aphra Kerr⁷, University College Dublin Dave Lewis⁸, Trinity College Dublin</p>
Dissemination Level	PU - Public
Deliverable Type	Report

³ <https://orcid.org/0000-0002-0183-3982>
⁴ <https://orcid.org/0000-0001-7748-2050>
⁵ <https://orcid.org/0009-0003-3966-623X>
⁶ <https://orcid.org/0000-0002-0458-693X>
⁷ <https://orcid.org/0000-0001-5445-7805>
⁸ <https://orcid.org/0000-0002-3503-4644>

Revision History				
Version	Changes implemented by	Revision Date	Reviewed by:	Brief description of changes
V1.0	Victoria Wiegand, Delaram Golpayegani, Marta Lasek-Markey, Arjumand Younus, Aphra Kerr, Dave Lewis	19/12/2025	Abeba Birhane, Eugenia Siapera	Initial draft
V2.0	Victoria Wiegand, Monique Munarini, Dave Lewis, Delaram Golpayegani, Yuening Li, Aphra Kerr	28/01/2026		Revised based on reviewers' feedback, with major revisions including: refining the terminology, deeper explanation of methodological process, and alignment with deliverable D2.2.

Executive Summary

This document is part of a larger study aiming to distil success criteria that are present in documents developed by different classes of formal bodies that play a role in AI policy development, namely supranational bodies, standardisation bodies, health and legal professional bodies, and, in this document, ICT professional bodies. The ICT professional bodies selected for this study are the ACM and IEEE, chosen due to their large global membership and their role in setting professional and research guidelines for responsible, trustworthy and ethical AI adoption. A total of 11 documents issued by ACM and IEEE are selected on the basis of their relevance to these issues. These documents are a mixture of policy documents, codes of conduct and standards.

The methodology adopted combines thematic analysis using the NVivo tool and unsupervised BERTopic modelling. The resulting codes and topics are the basis for a structured analysis intended to identify the themes representing expectations in these documents for criteria of success. These themes are set in the frame of meta-themes derived from the parallel study of supranational organisation policies (FORSEE Deliverable D2.2) and a micro-meso-macro assessment of where expectations are directed and who is responsible.

Key findings are:

- 1) There is a reasonable correspondence between the codes developed using NVivo and the topics developed using BERTopic, indicating that unsupervised topic modelling can play a supportive role in professional document analysis when combined with thematic analysis;
- 2) The themes identified from the combined NVivo and BERTopic analyses align well with the framing of meta-themes related to AI success (identified in a parallel study in D2.2). However, the ICT professional bodies do not discriminate between meso and macro levels and these are therefore combined when placing themes representing expectation to be addressed by external organisations;
- 3) Assessed within the current institutional context of the ICT professional bodies, the themes developed on the bodies' expectation of success seem to have evolved largely into ethical practice for AI researchers and training for individual professionals, rather than representing a normative framework that in some cases was expected to be adopted at the organisational (micro) or policy (meso/macro) level.
- 4) The persistence of ethics, ethical AI and ethics principles across the years in the documents from these organisations is noteworthy when compared to supranational documents, while human wellbeing is a distinctive focus.
- 5) A focus on individual professionals is core to these documents and a sense that it is the individual's responsibility to be a competent and ethical leader.

Acronyms

ACM	Association for Computing Machinery
AI	Artificial Intelligence
EU	European Union
ICT	Information and Communications Technology
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organization for Standardisation
LDA	Latent Dirichlet Allocation

Terminology

The following provides the definitions of the key concepts used in this deliverable.

Code	A concept that is manually identified in the thematic analysis process as being both present and important within the document.
Document	A formally issued textual artefact, as a whole, issued by supranational bodies, technical standardisation organisations, academic institutions, or professional bodies, and intended to articulate principles, requirements, or procedures relevant to the governance, development, or deployment of artificial intelligence.
Thematic Analysis	A popular method of manual qualitative data analysis that systematically analyses datasets and identifies patterns of meaning.
Theme	A broad interpretable semantic concept derived through expert interpretation, either from the topics generated by a topic modeller or from the codes generated during the thematic analysis process.
Topic Modelling	A statistical method used to identify prevalent topics in large corpora of text using unsupervised machine learning techniques.
Topic	A concept that has been identified by a computational model (e.g., LDA or BERTopic) as a latent concept associated with a subset of documents (in our work a subset of sentences within a document). Topic is typically summarised by salient words that tend to appear when that concept is present.

Table of Contents

Executive Summary	4
Acronyms	5
Terminology	6
Table of Contents	7
1. Introduction	8
2. Sociology of Expectations and Institutional Approaches to AI Governance	10
3. ICT Professional Bodies: Institutional Context	12
4. Methodology	15
4.1. Selecting Documents and Compiling the Corpus	15
4.2. Combining Thematic Analysis and Topic Modelling	16
4.2.1. Thematic Analysis	17
4.2.2. Topic Modelling	18
4.2.3. Interpreting Codes and Topics to Identify Themes	21
5. Findings	22
5.1. Codes Identified in NVivo-Supported Manual Thematic Analysis	22
5.2. Topics Identified using BERTopic	26
5.3. Comparison of Thematic Analysis Codes and Topic Modelling Labels	30
5.4. Limitations of BERTopic-based Topic Identification	32
6. Key Themes, Insights, and Discussion	33
7. Conclusion	36
Positionality statement	37
Acknowledgements	38
References	39
Appendix A: Codes from Supranational AI Documents Analysis	42
Appendix B: NVivo Code Results	45
Appendix C: Document-Specific Comparisons of NVivo Codes and BERTopic Topic Labels	48

1. Introduction

This deliverable (D2.3) is part of a four-part study in Work Package 2 (WP2) of the FORSEE project, as shown in Figure 1. Together, these studies investigate the viewpoints and criteria set forth by a range of supranational bodies, technical organisations, academic institutions, and professional bodies in defining successful AI applications. Supranational bodies, technical organisations, academic institutions, and professional bodies collectively play a pivotal role in shaping perceptions of success in AI by establishing standards, guidelines, and ethical frameworks that influence the development, deployment, and evaluation of AI applications. These bodies serve as crucial actors therefore in shaping formal approaches, criteria and social expectations as to AI success.

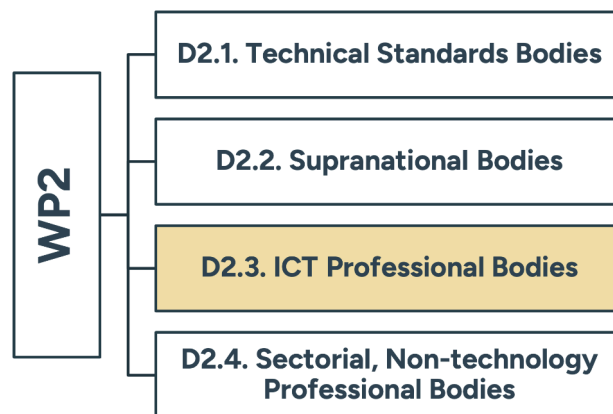


Figure 1: The structure of workpackage 2 deliverables.

For context, the other three parts of WP2 address respectively AI governance outputs from:

- **Technical Standards Bodies (D2.1):** representing international and national bodies developing technical and process standards for AI.
- **Supranational Bodies (D2.2):** Supranational institutions that represent, or who are selected to be representative of, interests of multiple different countries in developing guidelines or rules for the governance of AI.
- **Sectorial, Non-technology Professional Bodies (D2.4):** Bodies representing professionals internationally in sectors that are not primarily ICT-based but which are sufficiently impacted by AI technology to engage in sector-specific AI governance.

This deliverable (D2.3) aims to uncover common themes, both shared and divergent, within the **relevant guidelines issued by ICT (Information and Communication Technology) professional bodies, i.e. ACM and IEEE**. The research question we explore in D2.3 is: **What**

are the prevalent themes in relevant policies, standards, and guidelines published by ICT Professional Bodies, particularly IEEE and ACM, between 2018 and 2025?

D2.3 is focussed on the expression of AI success criteria offered in documents developed by **ICT professional bodies** addressing issues of AI governance. These include specifically **ACM, the Association for Computing Machinery, and IEEE, the Institute of Electrical and Electronics Engineers**. The ACM is an international professional society aiming to advance the theory, practice, and application of computing through research publications, conferences, and educational initiatives. Similarly, the IEEE is a professional association with a focus on electrical, electronics, communications and computer engineering. Both are non-profit organisations with a long history, established in the United States.

Overall, D2.3 aims to discern the set of baseline criteria for AI success produced by these ICT professional bodies and their respective committees. We seek to identify and categorise recurring themes, priorities, and nuanced dimensions of success criteria outlined in the documents issued by **ICT professional bodies**. D2.3. therefore aims to shed light on the evolving perspectives and priorities that shape ICT professional bodies' discourse on the responsible development and deployment of AI. The ACM and IEEE were selected as these are the two pre-eminent international professional bodies in the ICT domain globally in terms of international membership, participation across industry, government and academia and leadership in peer-review ICT research publication.

To do so, this study employs a combined methodological approach that integrates thematic analysis, originating from social sciences (Braun & Clarke, 2021b, 2021a; Terry et al., 2017), and unsupervised topic modelling approaches to identify key themes that signal success criteria for AI within a corpus of policies, standards, and guidelines issued by ICT professional bodies. Combining thematic analysis and topic modelling ensures that a contextual understanding of concepts within documents are captured through human-based analysis and can be used to frame and interpret patterns gleaned from across a large and diverse document set using an unsupervised machine learning algorithm.

This study selects documents from 2018 to 2025, as 2018 followed a period of significant negative discourse about AI and whistleblowing about practices in major AI companies, including the Cambridge Analytica scandal (see Kerr et al., 2020). From a sociology of expectations perspective one could discern that these negative expectations were impacting upon public and citizen willingness to accept AI in their everyday lives. Thus from 2018 there has been a period of rapid growth in multistakeholder AI policy making during which a number of supranational institutions have reached agreement on some basic criteria for AI governance. In this period both the ACM and IEEE have also developed or revised guidelines that have a bearing on the professional conduct of their members in relation to AI and to the conduct of organisations of which they are part.

Addressing our research question, we first create a corpus of the AI policies, standards, and guidelines issued by prominent ICT professional bodies, including the ACM and IEEE in the time frame identified. Second, we perform thematic analysis and unsupervised topic modelling and consolidate the results to identify themes. Following, we analyse the AI success factors articulated by ICT professional bodies through the lens of sociology of expectations.

The remainder of this report is structured as follows. Section 2 provides some background on the theories accompanying sociology of expectations, which is the lens through which we aim to integrate the results of this study with other elements of the broader study described above. Section 3 provides an overview of the institutional context of the ICT professional bodies studies, namely the ACM and IEEE. Section 4 details the methodology taken, including criteria for the selection of documents for analysis, the parallel approaches of thematic analysis to identify codes using NVivo and topic analysis using BERTopic to identify topics. Section 5 presents the findings from these two strands of analysis and their integration. Section 6 presents the identification of themes representing expectations set in the frame of meta-themes derived from the parallel study of supranational organisation policies and a micro-meso-macro assessment of where expectations are directed. Section 7 offers conclusions and plans for further research.

2. Sociology of Expectations and Institutional Approaches to AI Governance

WP2 is informed by the Sociology of Expectations (SoE). This is a strand of the Social Construction of Technology Paradigm (SCOT) (Pinch & Bijker, 1984) that posits technologies are shaped by social forces, including the values, beliefs, interests, and power dynamics of the actors involved in their development and use. Specifically, SoE highlights the prominent role that expectations play in shaping technological processes. Expectations are perceived as socio-technical visions of the future that shape the present. As (Brown et al., 2003), note: Within this process, discourses of hope often serve as vectors, embedding promissory futures with real world effects (Brown, 2015). While expectations are primarily articulated through narratives, once they gain traction they “may materialise in experiments and prototypes” (van Lente et al., 2013). By focusing on the potential material impact of promises about the future, scholars working within the SoE paradigm shift attention away from verifying the truth of specific claims towards examining their meaningfulness and their

capacity to mobilise resources, organisations, and people to act.

Expectations can be categorised according to their level—micro (e.g. research groups), meso (technological fields), or macro (societal contexts) (van Lente et al., 2013). They can also be distinguished by their content, since “expectations may concern technical, commercial, or societal aspects, and probably a mix of these” (van Lente, 2012). In practice, particular technological artefacts and areas of scientific research are rarely associated with a single, neatly bounded set of expectations. Rather, they tend to attract multiple and overlapping expectations that vary in scope, content, and degree of internal coherence.

In summarising the potential force that expectations yield, Van Lente (van Lente, 2012) notes the following key aspects:

- “First, what expectations do is to raise attention and legitimise investments: a project or programme can be defended by referring to a promising future.
- Second, expectations provide direction to the search processes of science and technology...Typically there are many possible paths while choices have to be made. The optimal direction cannot simply be calculated.
- Finally, there is a coordination effect of expectations...Technical development is not solitary work, but the work of networks of companies and research institutions. When a central control is lacking, as is usually the case, expectations indicate pieces of work and stipulate roles”.

AI has a long and well-documented history, characterised by recurrent cycles of hype and subsequent disappointment, often followed by recalibration of expectations in order to sustain funding and public attention (Galanos, 2023). These cycles, commonly referred to as “AI winters” and periods of renewed optimism, reflect the complex interplay between technological capabilities, societal expectations, and industrial interests. In the past three years, the commercial applications of large language models (LLMs) have reignited this pattern, situating AI at the centre of a hype cycle. Hype is further reinforced by the technical characteristics of AI; its adaptability and the promise of applications across multiple industries strengthen initial expectations and increase the chances of sustaining them beyond phases of disappointment (van Lente et al., 2013).

Media narratives and investor discourse frequently emphasise the transformative potential of these technologies, fostering a techno-solutionist ethos in which AI is presented as an inevitable fix for a wide range of social, economic, and scientific challenges (Lindgren & Dignum, 2023). In a similar vein, AI is treated as an “inevitability”, an unstoppable force that is not subject to debate; thusly the techno-deterministic narrative of “permissionless innovation” (Dotson, 2015) is echoed in the discourse around AI, fuelling aversion to regulation and democratic governance of emerging technologies. The same effect is created by positive-negative ideal types of expectations that associate AI either with utopian or dystopian visions of the future, obscuring limitations and a grounded debate on

AI's actual impact (Vicsek, 2020).

By contrast, the SoE approach carries a democratisation potential by identifying different expectations, associated with different social positions and perspectives. Indeed prior to 2018 a range of professional whistleblowers, including around the Cambridge Analytica scandal, shared insights into how AI tools were being deployed in companies and by public sector bodies and their potential individual and social harms. In order to address these concerns there has been a move away from a performative understanding of ethics (Kerr et al., 2020), that tends to embed ethics into the discursive strategies of AI developers (Steinhoff, 2024) and towards more robust professional standards. This prior work indicates potential conflicts between organisational-professional expectations surrounding AI and broader societal demand for better regulations, standards and professional training (Kerr et al., 2020). Indeed the social expectation in many European countries has been that public authorities and professional organisations must find a way to govern the design and use of AI in line with an agreed set of principles and values.

The Sociology of Expectations, in common with SCOT, identifies that different social groups or stakeholders may have competing narratives or expectations (Jasanoff & Kim, 2009). These narratives and expectations may stabilise over time, and some narratives, concepts and meanings may become dominant, and win out at a particular moment in time over alternative views and perspectives. Social expectations can be positive or negative, and they can be merely performative (discursive) or have real impact and shape actions in the world (Kerr et al., 2020). Additionally, the concept of **social imaginaries** provides a complementary theoretical lens to the sociology of expectations for interpreting how different social groups imagine and understand AI, as well as how it is used, governed, and made meaningful. Emerging in STS in relation to technologies such as nuclear power (Jasanoff and Kim, 2009) and the internet (Mansell, 2012) social imaginaries refer to collectively held, negotiated, and stabilised visions of how technologies are expected to function in society, including how responsibilities and benefits should be distributed.

Professional documents guide and shape professional training and practice while also reflecting how ICT professionals understand and address issues such as algorithmic bias, human well-being in autonomous and intelligent systems, and responsible system design. These understandings are embedded in, and made visible through, the phrasing, conceptual framing, and terminological choices used across documents produced by the ACM and the IEEE, and can be contextualised as the social imaginaries proposed by these professional bodies.

3. ICT Professional Bodies: Institutional Context

The ACM and IEEE represent the two largest professional bodies in the ICT domain. Both have global membership and a mix of individual member types, primarily from industry but

also including academic, governmental, civil society and student members. In common with ICT professional practice internationally, membership of these or similar professional bodies are not a requirement to practice as is common in medicine or law, though this may be sought by employers especially in safety critical ICT sectors. Therefore, professional codes in ICT can be characterised as voluntary, though professional codes are enforced by these bodies on members against whom complaints are made⁹.

In considering ACM and IEEE documents their institutional context and goals are key. Both bodies offer voluntary guidelines and standards. The ACM guidelines address the ethics code of practice to its members as individual professionals without specifying guidelines that are specific to AI. However, the widespread adoption since the early 2020's of the ACM code by peer-review conference programme committees in the AI domain means that a body of self-declared ethics assessments against the principles of the code is developing within the AI peer review corpora.

The IEEE took an approach that was much more specific to AI (scoped as 'intelligent and autonomous systems') with its Ethically Aligned Design guidelines. This took a similar approach to contemporary guidelines, such as those from the EC HLEGAI and OECD, in providing principles that can be adopted by organisations developing or deploying AI rather than structuring these as principles for ethical conduct by individual professionals as taken by the ACM. This institutional choice underpinned a broader standards development activity by the IEEE Standards Association under the P.7000 series. This was originally positioned as a set of voluntary standards that could support the development of AI systems with trustworthy and ethical characteristics and potentially also support regulatory compliance. However, there is little evidence of coordination between the development of the IEEE P.7000 series and the AI standards developed by organisations representing national bodies, i.e. ISO/IEC JTC21 SC42 internationally and CEN/CENELEC JTC21 in Europe (see deliverable 2.1).

Specifically, no standards from the IEEE P.7000 series have been admitted into the standards adopted by CEN/CENELEC JTC21 in relation to officially sanctioned standards development related to the EU AI Act. In parallel the IEEE Standards Association opted to make the P.7000 available for free (whereas IEEE standards are typically available for a fee) and has developed corresponding training and certification schemes for individual professionals¹⁰. This could therefore be viewed as a pivot by the IEEE from developing ethical AI standards for the purpose of regulatory compliance to prioritising professional standards for training and certification activities.

Therefore, despite some similarities between the IEEE P.7000 standards development activities and the AI standards development in ISO/IEC and CEN/CENELEC, we characterise the IEEE documents as now positioned similarly to the ACM code in focussing

⁹ <https://www.acm.org/about-acm/code-complaints-report>

¹⁰ <https://standards.ieee.org/about/training/ethics-for-ai-system-design/>

on voluntary professional code of practice. We therefore can summarise the expectation represented by the documents analysed in this document as being that voluntary professional codes and standardised measures can contribute to the ethical development of AI.

In this deliverable, we focus on the ACM Code of Ethics and Professional Conduct (*ACM Code of Ethics and Professional Conduct*, 2018) and the IEEE Guidelines on Ethically aligned design for intelligent and autonomous systems (*Ethically Aligned Design: A Vision for Prioritizing. Human Well-Being with Autonomous and Intelligent Systems, Version II*, 2019) and the associated IEEE standards that were inspired by this.

The ACM Code has become significant in the field of AI in that it has underpinned the code of ethics that the ACM and other learned societies such as the Association of Computational Linguistics (ACL)¹¹ and the Association for the Advancement of AI (AAAI)¹² expect works submitted to peer review venues to consider. This is in part a recognition that much of the advances in AI come from training and evaluation using secondary data and therefore such research avoids institutional ethics review that is typically triggered by research with human subjects. Researchers are therefore encouraged to comment in submission on the risk of harms arising from their research results (Srikumar et al., 2022).

The IEEE approach has taken a different path. With its Guidelines on Ethically aligned design, the IEEE was a detailed source of guidance for organisations developing and deploying AI contemporary to efforts from the European Commission's High Level Expert Group and the OECD (Lewis et al., 2020), rather than restricting itself to the conduct of individual researchers. This inspired the wider P.7000 programme of standards development programme by the IEEE Standards Association addressing ethical autonomous and intelligent systems (Adamson et al., 2019).

¹¹ <https://www.aclweb.org/portal/content/acl-code-ethics>

¹² <https://aaai.org/about-aaai/ethics-and-plurality/>

4. Methodology

In this section, we introduce how we selected our corpus of documents and how our methodology for analysing these documents combined manual thematic analysis and automated topic modelling.

4.1. Selecting Documents and Compiling the Corpus

Documents to be examined in this deliverable are selected based on the following criteria:

- **Scope:** Policies, standards, or guidelines with an explicit focus on AI and issues of trustworthy, ethical and responsible practice.
- **Timeline:** Documents published during the period of 2018 - 2025, corresponding to the period in which AI governance emerged as a concern globally.
- **Issuing body:** Prominent ICT professional bodies who have issued AI related governance documents in the specified time period. We selected documents from the ACM and IEEE.
- **Length of document:** No restrictions were applied in respect to the length of the documents.

The corpus of 11 selected documents is detailed in Table 1. The ACM Code of Ethics and Professional Conduct is openly available online¹³. IEEE-issued documents are available under the IEEE GET Program™ for AI Ethics and Governance Standards as free resources promoting trustworthy AI¹⁴.

Table 1. Documents issued by ACM and IEEE (ICT professional bodies) included in the corpus

ID	Document	Issuer	Type	Year
ICT_01	<i>(ACM Code of Ethics and Professional Conduct, 2018)</i>	ACM	Code of conduct	2018
ICT_02	<i>(Ethically Aligned Design: A Vision for Prioritizing. Human Well-Being with Autonomous and Intelligent Systems, Version II, 2019)</i>	IEEE	Policy document	2019
ICT_03	<i>(IEEE Standard Model Process for Addressing Ethical Concerns during System Design 7000, 2021)</i>	IEEE	Standard	2021
ICT_04	<i>(IEEE Standard Model Process for Addressing Ethical Concerns during System Design 7000, 2021)</i>	IEEE	Standard	2021
ICT_05	<i>(IEEE Standard for Data Privacy Process 7002, 2022)</i>	IEEE	Standard	2022

¹³ <https://www.acm.org/code-of-ethics>

¹⁴ See <https://standards.ieee.org/news/get-program-ai-ethics/>

ID	Document	Issuer	Type	Year
ICT_06	(IEEE Standard for Algorithmic Bias Considerations 7003, 2024)	IEEE	Standard	2024
ICT_07	(IEEE Standard for Transparent Employer Data Governance 7005, 2021)	IEEE	Standard	2021
ICT_08	(IEEE Ontological Standard for Ethically Driven Robotics and Automation Systems 7007, 2021)	IEEE	Standard	2021
ICT_09	(IEEE Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems 7009, 2024)	IEEE	Standard	2024
ICT_10	(IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being 7010, 2020)	IEEE	Guideline	2020
ICT_11	(IEEE Standard for Ethical Considerations in Emulated Empathy in Autonomous and Intelligent Systems 7014, 2024)	IEEE	Standard	2024

4.2. Combining Thematic Analysis and Topic Modelling

The terms “**code**”, “**topic**”, “**theme**” are used extensively throughout this deliverable. Although they may appear as synonyms, they are treated as distinct concepts and are not used interchangeably. To clarify this distinction, in the following, we provide background information about how we combine **thematic analysis** and **topic modelling** in this deliverable. These distinctions are further illustrated in Figure 2.

In this deliverable, we consider a **document** as a formally issued textual artefact published by an ICT professional body regarding AI. Such documents can range from standards to codes of conduct and are collectively compiled within a corpus.

In our methodology, we first apply **thematic analysis**, which is a popular method of qualitative data analysis that systematically organises datasets and helps to identify patterns of meaning. First described in the 1970s by Houlton – albeit the term itself had been in use even earlier – it became more prominent in the late 1990s with researchers like Boyatzis and Hayes (Braun & Clarke, 2021; Terry et al., 2017). In recent times, thematic analysis has been understood as an umbrella term for different approaches. While popular in qualitative interview data analysis, computer-assisted expert thematic analysis of legal texts, such as legislation or policy documents, appears to be less commonly employed.

In our work, we use manual thematic analysis to code the documents. The output of this process is a set of **codes**. **Code refers to a concept that is manually identified as being both present and important within the domain.** However, thematic analysis relies on the

individual’s judgements and expertise and therefore has limited reproducibility. In addition, it is a time-consuming and resource-intensive process which makes it difficult to conduct thematic analysis of a large corpus of documents.

Addressing these limitations, we also employ **topic modelling**, which is a statistical method used to identify prevalent topics in a large corpora of text. Topic modelling uses unsupervised machine learning techniques for uncovering topics within large sets of documents. There are different unsupervised topic models (see Churchill & Singh, 2022) for evolution of such models). Currently, BERTopic (Grootendorst, 2022) and Latent Dirichlet Allocation (LDA) (Blei et al., 2003) are the most popular topic models. The output of these models are topics. **Topic** is a concept that has been identified by a computational model (e.g., LDA or BERTopic) as being present in the corpus consisting of multiple documents. In both LDA and BERTopic, a topic can be understood as a latent concept associated with a subset of documents (in our work subset of sentences within a document as explained later) and typically summarised by salient words that tend to appear when that concept is present. It should, however, be noted that the underlying mathematical object defining a topic differs across LDA and BERTopic (a probability distribution over words in LDA versus a cluster of document embeddings in BERTopic).

Though topic modelling is usually applied to a set of documents, we apply it to each document separately to gain insights into each single document. Therefore, in this work topics are associated with a subset of sentences within a document.

Finally, we manually analyse the codes and topics through the lens of sociology of expectations to identify **themes**. **Theme** is a broad interpretable semantic concept derived through expert interpretation, integrating codes and topics into a set of themes that are then manually grouped into meta-themes within the framing of high level themes and expected levels where responsibility for addressing themes will lie.




Figure 2: Diagram of methodological process to generate themes from identified topics and codes

4.2.1. Thematic Analysis

The qualitative analysis of relevant ICT professional bodies documents was conducted with the use of NVivo software. As explained by (Dhakai, 2022), NVivo is a CAQDAS (Computer-assisted qualitative data analysis software) programme that assists, rather than replaces, a human researcher. NVivo analysis is, thus, researcher-led and in this project it was performed by a computer science researcher. The **documents** were imported into NVivo and coding was performed by selecting a section of text, usually a paragraph or bulleted list, and then tagging it with a node corresponding to a **code**. Note that figures or tables in the **documents** that are applicable to the **codes** are encoded with a singular code.

Throughout the encoding process, the computer science researcher (Researcher #1) referenced a list of **codes** created previously by a legal expert (Researcher #2, see their profile in the positionality statement) performing NVivo thematic analysis on supranational institutional AI policy documents¹⁵. This approach was conducted in order to align the analysis of these documents with the analysis of the governmental and intergovernmental policies already conducted on supranational documents and to identify any divergences. The list of **codes** and their corresponding **documents** is available in Appendix A. However, additional **codes** were created by the computer science researcher (Researcher #1) performing the thematic analysis as deemed necessary. As a result, the ICT professional bodies documents analysis resulted in **90 codes**.

4.2.2. Topic Modelling

Building upon our previous work on supranational bodies in FORSEE Deliverable D2.2¹⁶ (also see (Golpayegani et al., 2025)), we use BERTopic as the suitable topic modeller for quantitative analysis of the documents. The methodology for determining **topics** from the corpus of **documents** consists of the following:

1. **Text preprocessing:** Before application of BERTopic, we employed common natural language processing (NLP) approaches to apply minimal preprocessing using the `nltk` library for:
 - Lowercasing,
 - Removing non-alphabetic tokens (e.g. punctuation marks),
 - Removing tokens that consist solely of numbers, and
 - Removing tokens shorter than two characters.

Note that we did not remove stop words or perform lemmatisation prior to embedding, given that BERTopic uses transformer-based embeddings, which require the text to be persevered for better accuracy. As explained later, we removed the stop words, after generating embeddings, using the `CountVectorizer` when initialising BERTopic.

2. **Applying BERTopic:** Since this work uses BERTopic at the sentence-level, the description of how BERTopic works is made consistent with the wording used

¹⁵ <https://zenodo.org/records/18352574>

¹⁶ <https://zenodo.org/records/18352574>

throughout this deliverable. The following describes how BERTopic generates topics:

- a. Each sentence is converted to its embedding representation (numerical representation) using a pre-trained language model. We used `all-MiniLM-L6-v2`¹⁷ which is a pre-trained model that converts sentences and paragraphs into vector-based representations.

```
Python
embedding_model = SentenceTransformer("all-MiniLM-L6-v2")
```

- b. To group sentences into semantically similar clusters, the dimensionality of the resulting embeddings is reduced using UMAP (Uniform Manifold Approximation and Projection) and HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise). Since applying BERTopic default parameters on short documents did not generate any coherent topics, we identified two hyperparameter configurations: one for longer texts and another for shorter documents. These configurations only differ in the HDBSCAN clustering model, with a lower minimum number of samples per word cluster is set for shorter documents. These hyperparameter configurations were used uniformly (based on the length of document) to ensure consistency and reproducibility of results.

```
Python
map_model = umap.UMAP(n_neighbors=15, n_components=5, min_dist=0.0,
metric='cosine', random_state=42)
hdbscan_model = HDBSCAN(min_cluster_size=10, min_samples=5)
```

- c. For topic representations, we used `CountVectorizer` to enable frequency analysis, without down-weighting the common words in the document. Instead, we removed some common words, such as "AI", "Artificial Intelligence", and latin numerals in this step.

```
Python
vectorizer_model = CountVectorizer(ngram_range=(1, 1), stop_words=
stop_words)
```

¹⁷ <https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

The snippets used for preprocessing and topic modelling are available on GitHub under the MIT licence at https://github.com/DelaramGlp/forsee_topicmodelling. The source code and the results are available in an archived release with a persistent digital object identifier (DOI): <https://doi.org/10.5281/zenodo.18303112>.

3. **Topic labelling:** An example of BERTopic output for IEEE 7003-2024 Standard for Algorithmic Bias Considerations is in Figure 3. In this, “topic” denotes a unique numeric identifier for each topic generated by the model. The “Count” column shows the frequency of each topic, i.e. the number of sentences assigned to the topic. “Name” shows a short topic label, however, as shown in the figure, BERTopic does not automatically assign a semantic label and only shows top keywords. “Representation” provides a list of keywords representing the topic and “Representative_Docs” is a set of sentences that are associated with the topic.

As noted earlier, the *Name* value generated by BERTopic does not constitute meaningful semantic labels for topics. Consequently, assigning interpretable labels as topics was conducted manually. In this deliverable, the computer science researcher (see positionality statement) labelled the topics based on the codes previously generated in the NVivo thematic analysis. Note that the computer science researcher worked document by document, first conducting NVivo thematic analysis then immediately following, labelled the topics generated by BERTopic. An example of label assignments is shown in Figure 4.

Despite text pre-processing and hyperparameter optimisation of BERTopic, some of the keywords associated with a topic exhibited insufficient coherence (high noise) to allow meaningful topic labelling. Additionally, while some keywords were coherent, multiple topic labels could be assigned to them. We adopted a minimalistic approach when assigning labels, i.e. attributing the smallest set of labels to each topic identified by BERTopic. In cases where comprehensive representation required multiple labels, we established a constraint of assigning a maximum of 3 labels to each topic. This is illustrated in Figure 4 where Topic #7 is deemed not relevant, and Topic #3 has two topic labels.

Topic	Count	Name	Representation	Representative_Docs
0	52	0_bias_biases_mitigate_assess	['bias', 'biases', 'mitigate', 'assess', 'foresight', 'caught', 'hindsight', 'inform', 'categorizing', 'policy']	[intentionally by dating specific policy to mitigate unwanted bias for example gender discrimination unintentionally through', 'the interaction of data and model in both of which unwanted bias has been addressed does not preclude unwanted bias in the outputs', 'policy intended to mitigate discrimination but can actually increase unwanted bias for example mandating can not be used to assess loan rates but then unintentionally the model data of stakeholders categorized cording to sex can not be assessed for unwanted bias']
1	51	1_algorithmic_bias_model_criteria	['algorithmic', 'bias', 'model', 'criteria', 'automated', 'detect', 'normative', 'processing', 'metrics', 'consideration']	[the bias profile purpose the bias profile is defined see clause as repository of information created and maintained through the activities of algorithmic bias consideration defined in this its purpose is to provide record of how bias has been considered in relation to an ais', 'understanding algorithmic bias and aspects of human bias', 'understanding algorithmic bias and aspects of human bias the complex and multifaceted nature of algorithmic bias means there have been number of definitions and elaborations associated with this phenomenon']
2	34	2_risk_risks_ria_impacts	['risk', 'risks', 'ria', 'impacts', 'mitigation', 'process', 'mission', 'guidance', 'assessments', 'inventory']	[risk and impact assessment purpose the purpose of the risk and impact assessment ria section is to set out series of actions for the cation and analysis of risks and impacts arising from the ais under consideration', 'risk and impact assessment the goal of risk and impact assessment ria is the identification and analysis of risks arising from the ais', 'risk and impact assessment']

Figure 3: An excerpt of the output generated by BERTopic for IEEE 7003-2024 Standard for Algorithmic Bias Considerations.

Topic	Count	Representation	Researcher1_Topic
0	52	['bias', 'biases', 'mitigate', 'assess', 'foresight', 'caught', 'hindsight', 'inform', 'categorizing', 'policy']	bias
1	51	['algorithmic', 'bias', 'model', 'criteria', 'automated', 'detect', 'normative', 'processing', 'metrics', 'consideration']	bias
2	34	['risk', 'risks', 'ria', 'impacts', 'mitigation', 'process', 'mission', 'guidance', 'assessments', 'inventory']	risk (assessment and management)
3	34	['cultural', 'cultures', 'social', 'technologies', 'family', 'biases', 'gender', 'stereotypes', 'organisational', 'ethnic']	social context, diversity and inclusion
4	33	['bias', 'ais', 'requirements', 'consideration', 'bug', 'objectives', 'testing', 'methods', 'scenarios', 'process']	bias
5	28	['stakeholder', 'stakeholders', 'ais', 'influencing', 'governance', 'business', 'map', 'io', 'responsible', 'drivers']	stakeholders
6	27	['attributes', 'stakeholder', 'identification', 'stakeholders', 'annex', 'debt', 'traits', 'guide', 'diverse', 'dependent']	stakeholders
7	25	['definitions', 'standard', 'equals', 'mandatory', 'abbreviations', 'conformance', 'requirements', 'footnote', 'statements', 'accordance']	NOT RELEVANT

Figure 4: An example of assigning labels to the outputs generated by BERTopic by the computer science researcher (note positionality statement) for the IEEE 7003-2024 Standard for Algorithmic Bias Considerations.

4.2.3. Interpreting Codes and Topics to Identify Themes

Following inductive manual coding and automated topic analysis, the research team (see positionality statement at the end of the document) went through a process of iterative interpretation aimed at generating, refining and naming themes. This process of meaning making involved both colour coding of topics and codes, clustering, referring back to the original texts and analysis. In line with prior work on the sociology of expectations, we also categorised the themes according to whether the themes primarily referred to actions that might be taken by actors at a micro or meso/macro levels. The multidisciplinary team then reviewed, discussed and refined the themes and clustered them into meta-themes guided by the overall research question, our disciplinary backgrounds and our grounding in the corpus and the prior literature. All anomalies and queries were answered in the joint discussion meeting and did not require external consultation.

5. Findings

5.1. Codes Identified in NVivo-Supported Manual Thematic Analysis

The prominent **codes** from the NVivo thematic analysis of the ICT professional bodies documents are represented in Figure 5 and Figure 6. **Data, stakeholders, ethics, transparency, risk (assessment and management), and human well-being** are the most common codes. These reflect the nature of the documents as an implementable standard, policy, or guidelines for individual ICT professionals and organisations developing and/or employing AI technologies.




Figure 5. Word cloud overview of NVivo results (**codes**) for ICT professional bodies documents

As the documents differ in application areas, such as data governance, emulated empathy and robotics, the most common codes showcase the overarching themes present. **Data** is encoded in reference to any large amount of data that is to be collected, stored, and used for some purpose. *Note:* this differs from **personal information** which is synonymous with personal data or personally identifiable information (PII). **Stakeholders** are commonly referenced throughout the documents to suggest that AI developers and deployers consider all affected parties of AI technologies. The documents employ **ethics** in various ways: evaluation of deontological, utilitarian, or virtue ethical frameworks, request to broadly consider ethical implications of AI technologies, or suggestion to consult ethicists in the development process. **Transparency** denotes transparency of AI technologies through mechanisms of explainability or traceability, transparency with all forms of data handling (when possible), or transparency of development processes. A frequent suggestion is to include **risk (assessment and management)** within all stages of the AI development process. **Human well-being** is emphasized as a high priority consideration when assessing the (potential) impact of AI technologies.



Figure 6. Tree map overview of NVivo results (**codes**) for all ICT professional bodies documents

Lesser encoded codes are more applicable to specific documents. Table 2 shows an overview of key codes identified manually through NVivo supported thematic analysis of AI documents issued by ICT professional bodies, such as ACM and IEEE. Document level encoding results are available in Appendix B.

Table 2. Overview of the key codes identified in NVivo-supported manual thematic analysis.

ID	Document	No. of codes	Key codes identified through NVivo-supported manual thematic analysis
----	----------	--------------	---

ICT_01	ACM Code of Ethics and Professional Conduct	43	Common good, Harm Reduction, Social Responsibility, Transparency, Stakeholders, Codes of Practice, Fairness, Impact Assessment, Privacy, Responsibility
ICT_02	Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems (Version 2)	84	Ethics, Human Well-being, Data, Transparency, Personal Information, Norms, Metrics, Privacy, Accountability, Mixed Reality
ICT_03	IEEE 7000-2021 - IEEE Standard Model Process for Addressing Ethical Concerns during System Design	60	Ethics, Stakeholders, Standards, Risk (Assessment and Management), Value-Based Design, Life Cycle, Transparency, Controllability, Impact Assessment, Principles
ICT_04	IEEE 7001-2021 - IEEE Standard for Transparency of Autonomous Systems	47	Transparency, Standards, Stakeholders, Autonomy, Explainability, Audit, Data, Risk (Assessment and Management), Ethics, Frameworks
ICT_05	IEEE 7002-2022 - IEEE Standard for Data Privacy Process	37	Privacy, Risk (Assessment and Management), Data, Frameworks, Standards, Personal Information, Life Cycle, Policy, Responsibility, Governance
ICT_06	IEEE 7003-2024 - IEEE Standard for Algorithmic Bias Considerations	42	Bias, Stakeholders, Data, Social Context, Life Cycle, Risk (Assessment and Management), Impact, Impact Assessment, Diversity and Inclusion, Standards
ICT_07	IEEE 7005-2021 - IEEE Standard for Transparent Employer Data Governance	50	Data, Personal Information, Consent, Governance, Risk (Assessment and Management), Safety and Security, Responsibility, Transparency, Policy, Accountability

ICT_08	IEEE 7007-2021 - IEEE Ontological Standard for Ethically Driven Robotics and Automation Systems	40	Ethics, Personal Information, Legal Responsibility, Norms, Governance, Data, Ontology, Responsibility, Privacy, Accountability
ICT_09	IEEE 7009-2024 - IEEE Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems	21	Stakeholders, Health and Safety, Harm Reduction, Risk (Assessment and Management), Standards, Life Cycle, Conformance, Legal Responsibility, Autonomy, Dependability
ICT_10	IEEE 7010-2020 - IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being	58	Human Well-being, Impact Assessment, Stakeholders, Data, Impact, Standards, Risk (Assessment and Management), Harm Reduction, Monitoring, Bias
ICT_11	IEEE 7014-2024 - IEEE Standard for Ethical Considerations in Emulated Empathy in Autonomous and Intelligent Systems	61	Stakeholders, Ethical AI, Risk (Assessment and Management), Data, Standards, Codes of Practice, Guidelines, Privacy, Bias, Harm Reduction

5.2. Topics Identified using BERTopic

In this section, we provide the key topics identified from analysis of the output generated automatically by BERTopic, following our topic modelling methodology. As a reminder, in this step, BERTopic is applied to a pre-processed text of each document. The output generated by BERTopic is stored in a tabular format (CSV) and top keywords related to each topic are visualised using a bar chart as well as a word cloud. The use of multiple visualisations enhanced assignment of meaningful semantic topic labels to BERTopic-generated topics. An example bar chart and world cloud illustrating topics identified from the ACM’s Code of Ethics and Professional Conduct are illustrated in Figure 7 and Figure 8, respectively. Note that in the word cloud (Figure 8), Topic -1 represents the outlier topic (i.e. sentences that are not assigned to any meaningful topic).

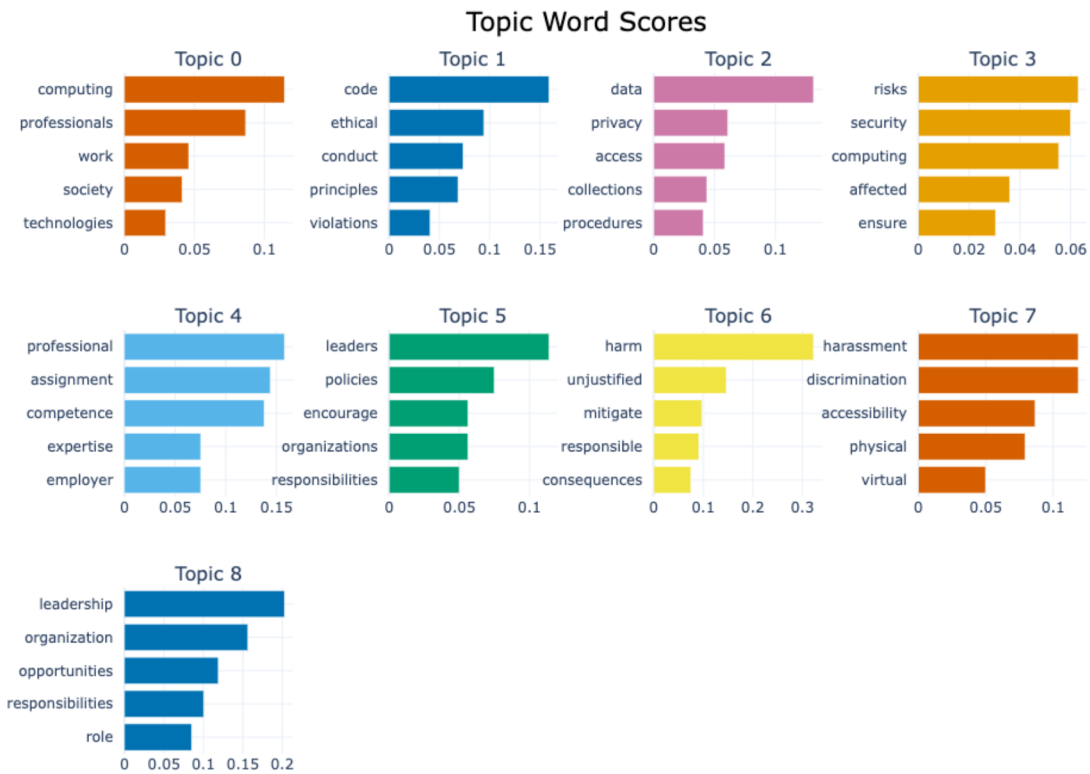


Figure 7. Bar charts representing the **topics** identified by BERTopic for ACM Code of Ethics and Professional Conduct.



Figure 8. World clouds demonstrating the **topics** identified by BERTopic for ACM Code of Ethics and Professional Conduct.

Applying the BERTopic model to each ICT professional body document produced 407 topics in total. Table 3 shows the number of topics identified by the BERTopic and the

number of labels we assigned to these topics (148 of labels in total). As noted earlier, the two figures differ because a single topic may receive multiple labels (up to three), and some topics were judged to be too noisy to warrant any meaningful label assignment, shown in Figure 4. The complete set of results for all the selected documents are published online at <https://doi.org/10.5281/zenodo.18303112>¹⁸.

To create the aggregated topic overview across the documents, we used the Count value associated with each topic, which we labelled in a manual process. Since we applied BERTopic at sentence level, the Count value represents the number of sentences assigned to each topic. Consequently, Count is an indicator of topic occurrence and prevalence within each document. To assess cross-document topic prevalence, counts from the documents were aggregated by summing the number of sentence assignments corresponding to equivalent topics across all documents. Figure 9 provides an aggregated topic overview using a word cloud visualisation, wherein relative size of each topic reflects the cumulative number of sentences associated with that topic across the corpus of ICT professional bodies documents.

Table 3. Overview of the BERTopic-generated topics and the assigned topic labels associated with each document.

ID	Document	No. of topics from BERTopic	No. of assigned topic labels	Key topics (from manually assigned labels)
ICT_01	ACM Code of Ethics and Professional Conduct	9	12	Responsibility, Ethics, Codes of Practice, Privacy, Personal Information
ICT_02	Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems (Version 2)	63	33	Ethics, Data, Privacy, Ethical AI, Metrics

¹⁸ Also available on GitHub at https://github.com/DelaramGlp/forsee_topicmodelling/tree/main/results/ICT.

ICT_03	IEEE 7000-2021 - IEEE Standard Model Process for Addressing Ethical Concerns during System Design	46	25	Ethics, Standards, Stakeholders, Value-Based Design, Impact Assessment
ICT_04	IEEE 7001-2021 - IEEE Standard for Transparency of Autonomous Systems	36	18	Transparency, Stakeholders, Audit, Data, Autonomy
ICT_05	IEEE 7002-2022 - IEEE Standard for Data Privacy Process	21	13	Privacy, Risk (Assessment and Management), Data, Life Cycle, Guidelines
ICT_06	IEEE 7003-2024 - IEEE Standard for Algorithmic Bias Considerations	30	12	Bias, Stakeholders, Social Context, Diversity and Inclusion, Metrics
ICT_07	IEEE 7005-2021 - IEEE Standard for Transparent Employer Data Governance	46	19	Personal Information, Consent, Data, Safety and Security, Accountability
ICT_08	IEEE 7007-2021 - IEEE Ontological Standard for Ethically Driven Robotics and Automation Systems	46	16	Personal Information, Ethics, Privacy, Legal Responsibility, Ontology
ICT_09	IEEE 7009-2024 - IEEE Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems	36	11	Harm Reduction, Stakeholders, Risk Assessment and Management), Legal Responsibility, Conformance
ICT_10	IEEE 7010-2020 - IEEE Recommended Practice for Assessing the Impact of Autonomous and	38	17	Impact Assessment, Human Well-being, Data, Metrics, Stakeholders

	Intelligent Systems on Human Well-Being			
ICT_11	IEEE 7014-2024 - IEEE Standard for Ethical Considerations in Emulated Empathy in Autonomous and Intelligent Systems	35	27	Affective Computing, Standards, Data, Impact Assessment, Consent

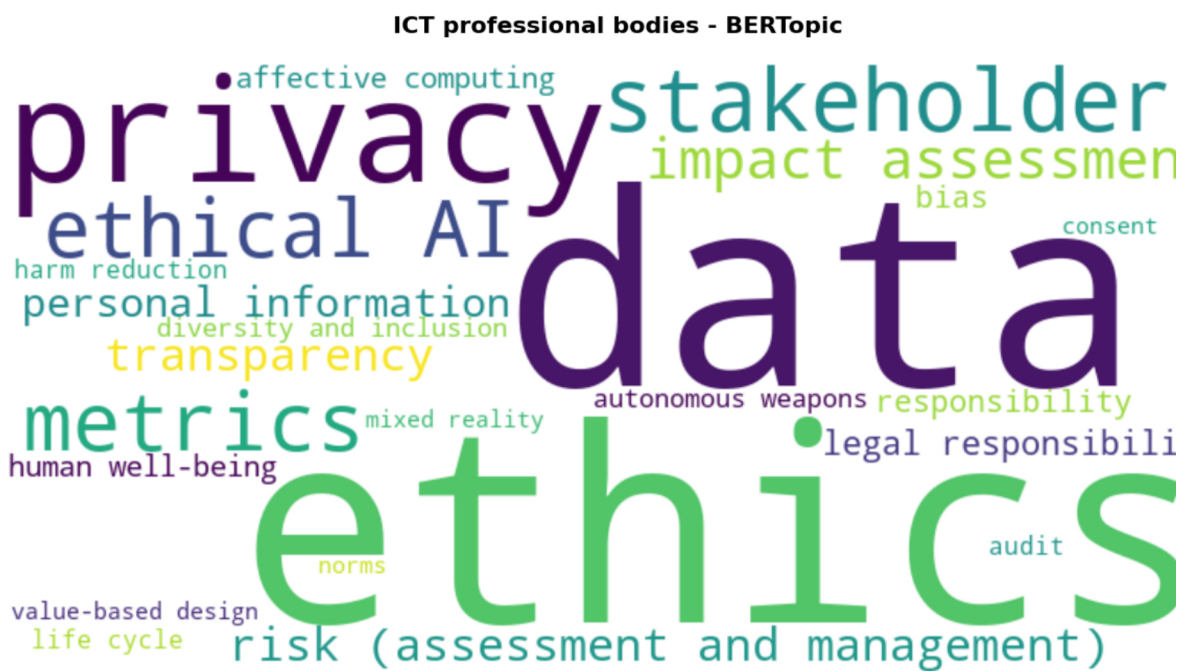



Figure 9. Word cloud of the **topics** from ICT professional bodies documents generated by BERTopic.

5.3. Comparison of Thematic Analysis Codes and Topic Modelling Labels

The BERTopic-based topic labels generated largely supported the NVivo analysis codes but fell short in identifying nuances. NVivo thematic analysis produced a larger number of codes for all ICT documents. Document specific comparisons can be found in Appendix C. Figure 10 represents a side by side visualisation of word clouds to contrast the results of the NVivo thematic analysis and BERTopic.



(a) Word cloud overview of NVivo results (codes) for ICT professional bodies documents.

(b) Word cloud of the topics from ICT professional bodies documents generated by BERTopic.

Figure 10. Side by side comparison of codes and topics identified from ICT professional bodies documents

The size of the word differs when comparing codes (Figure 5) and topics (Figure 9); for example: *privacy*, *metrics*, and *transparency*. *Ethics*, *data*, and *stakeholders* remain central for both forms of analysis. Table 4 compares the contents of the word clouds for NVivo and BERTopic. Note that this table does not take into account the size, or frequency of the themes, as displayed in the word cloud. As illustrated in the Divergences and Unique to BERTopic section of Table 4, BERTopic fixated on the case studies presented in the documents, such as affective computing, autonomous weapons, and mixed reality. Researcher-led thematic analysis in NVivo enabled a nuanced examination of the documents for finding synonymous phrases or words or overarching themes of discussed examples.

Table 4. Comparison of codes and topics identified from ICT professional bodies documents (specific to the prevalent codes and topics displayed in the word clouds, i.e. Figure 5 and Figure 9).

Comparison of prevalent codes and topics for ICT Professional Bodies Documents	
Quantitative comparison	24 (BERTopic - topics) < 27 (NVivo - codes)

<p>Overlaps</p>	<p>18 codes and topics:</p> <ul style="list-style-type: none"> ● Stakeholders ● Impact assessment ● Privacy ● Life cycle ● Data ● Ethics ● Harm reduction ● Transparency ● Metrics ● Responsibility ● Ethical AI ● Legal responsibility ● Risk (assessment and management) ● Personal information ● Consent ● Bias ● Human well-being ● Standards 	
<p>Divergences</p>	<p>Unique to BERTopic</p>	<p>Unique to NVivo</p>
	<ul style="list-style-type: none"> ● Value-based design ● Audit ● Mixed reality ● Affective computing ● Diversity and inclusion ● Autonomous weapons 	<ul style="list-style-type: none"> ● Impact ● Policy ● Autonomy ● Governance ● Social context ● Norms ● Safety and security ● Accountability ● Frameworks

5.4. Limitations of BERTopic-based Topic Identification

Sensitivity to hyperparameter settings: BERTopic results heavily rely on several key hyperparameters, including those determining how the text converts into numerical representations (embedding model), and those controlling the way these representations group into clusters with semantic similarity (parameters such as UMAP and HDBSCAN). In

this work, the hyperparameter tuning was conducted at an initial stage using a subset of representative documents, with the objective of yielding a sufficient number of relevant topics from each document while limiting the stochastic variability. As mentioned earlier, two hyperparameter configurations were used in this work: one for longer documents and another for shorter ones which only differ in the number of minimum samples in each cluster of keywords. These hyperparameter configurations were used uniformly (based on the length of document) to ensure consistency and reproducibility of results. However, a caveat in following this approach is that the performance of BERTopic might not be equally optimal for each document, for instance it might have resulted in a higher **proportion of outliers or topic omission or merging in one document compared to another**. Further, we applied BERTopic to individual documents by segmenting each into sentences, though BERTopic is typically used to identify latent topics across a large corpora of documents. This is mainly done to deal with the issue of small-scale data that BERTopic struggles with; and the approach in this case essentially enables BERTopic to overcome sparsity in data issues.

Processing of figures and tables: The documents in the corpus often feature tables or figures. As mentioned earlier, we use the PyMuPDF library to extract text from PDF files. While PyMuPDF extracts the text within tables, it does not process pictures. This can also be influenced by how the tables or figures are embedded in the document, whether in a text format or as a file attachment.

Subjectivity of the topic labelling process: the manual interpretation and labelling of BERTopic results to identify topics can be to some extent subjective (Cisek & Kelleher, 2023).

6. Key Themes, Insights, and Discussion

A thematic analysis was undertaken based on an assessment of the concepts represented by the combined NVivo-supported codes and topics identified from BERTopic-generated outputs presented in Table 5. To facilitate comparison with the thematic analysis conducted for supranational AI policy documents in FORSEE deliverable D2.2¹⁹, this analysis adopted the same meta-themes identified in the thematic analysis of supranational AI policy documents, namely: AI Technical Issues; AI Uses; AI Risks and Harms; and AI Governance.

Themes are then categorised within these four meta-themes, using the micro-meso-macro framing already used in themes in the supranational document set. This framing aims to offer a sense of the expectations of bodies as to where the concepts, i.e. aggregation of

¹⁹ <https://zenodo.org/records/18352574>

codes and topics, collected in each theme would be best addressed. The micro categorisation indicates that the concepts would be best addressed by individual organisations, but with some degree of freedom in how they do so and therefore accommodating of diversity and innovation in approaches taken. The meso categorisation indicates that concepts would be best addressed at the level of specific sectors or within a particular government policy area, while the macro categorisation indicates that concepts would be best addressed at a societal level.

The ICT professional bodies documents, however, largely addresses their expectation to individual professionals and the organisations developing and deploying AI within which they work. There is little focus within these documents on the distinction between concepts where the expectation that they could be addressed at a sectoral or government policy level (meso) and those that were expected to be best addressed at a societal level (macro). *Note that the documents suggest adhering to all legislation applicable to organisations' jurisdiction.* Therefore, this analysis does not distinguish between themes addressing expectations best suited to be handled at the meso and macro levels. Instead it treats the meso and macro level as a **single external level** to distinguish from the micro level, which we view as an **internal level**.

Following the reflexive process in Braun & Clarke's thematic analysis, the codes and topics are categorised into themes and are subjective and based on the research team's interpretation and discussion (as noted in the positionality statement). Different themes may have different meanings in different subject matter areas. Thus, the researchers based the conceptual understanding of the codes, topics, themes, and meta-themes as they are discussed in the corpus.

Under the **AI Technical Issues** meta-theme, two themes addressing expectation at the internal level are identified. The *Transparency* internal theme collects concepts that can be used by organisations and professionals including explainability, open source systems and traceability. The *Operationalising Values* theme brings together concepts such as ethical AI, interdisciplinary, common good that both relate strongly to individual professional expectations but can also be integrated into organisational behaviour and policies. External level expectations under the AI Technical meta-theme are grouped under a theme of *Standards* which express the desire for concepts such as harmonisation, conformance and metrics to be addressed at this extra-, inter-organisational level (in which the bodies themselves are active).

Under the **AI Uses meta-theme**, an internal theme of *Use Cases* is identified which is focussed on a limited set of specific applications of AI, such as autonomous weapons or affective computing which present particular concerns for responsible practice. At the external level a theme of *Societal and Personal Impact* is identified which assembles areas where there is expectation that AI use has particular sectoral or societal concerns, including digital sovereignty, sustainability and consumer protection.

Within the **AI Risks and Harms** meta-theme, two internal level themes are identified. One collects concepts related to *AI safety* including expectations of health and safety as well as more speculative risk related to organisations pursuing artificial general intelligence. The other internal theme collects concepts related to *Data Concerns*, which is a dominant term in the codes and topics extracted and includes issues such as consent, personal data and accuracy. At the external level there are expectations on how sectorial or societal actors should support responsible behaviour at the internal level via concepts collected under an Accountability theme, including expectations for normative guidance related to concepts including audits, independent oversight, impact assessment and compliance monitoring.

Under the fourth **AI Governance** meta-theme, an internal theme is identified around *Stakeholder Participation*, including concepts related to social responsibility, vulnerable groups and public education, highlighting this role on responsible AI adoption by organisations. At the external level two themes are identified. One theme *Human Rights Centered AI Governance* collects concepts aligned with notions of rule of law on justice, democracy and fundamental/human/inalienable rights and the other is focused on the articulation of *AI Governance Instruments* drawn from more general governance mechanisms such as codes of practice, guidelines and norms. These themes therefore position the ICT bodies as framing AI responsibilities as shared between organisations and the sectors, jurisdiction and societies in which they operate.

Table 5. Themes identified from ICT professional bodies

Meta Themes	Internal (Micro) themes where expectations are attended to by individual organisations, but they may do so in ways different to other organisations	External (Meso/Macro) themes where expectations are attended beyond the level of individual organisational, but without distinguishing between sectoral, policy domain or societal level responsibilities
AI Technical Issues	<i>Theme 1: Transparency</i> (n=7) <i>Theme 2: Operationalising Values</i> (n=10)	<i>Theme 3: Standards</i> (n=8)
AI Uses	<i>Theme 4: Use Cases</i> (n=3)	<i>Theme 5: Societal and Personal Impact</i> (n=11)
AI Risks and Harms	<i>Theme 6: AI Safety</i> (n= 7) <i>Theme 7: Data Concerns</i> (n=6)	<i>Theme 8: Accountability</i> (n=11)
AI	<i>Theme 9: Stakeholder Participation</i>	<i>Theme 10: Human Rights Centered AI</i>

<p>Governance</p>	<p>(n=10)</p>	<p>Governance (n=12)</p> <p><i>Theme 11: AI Governance Instruments</i> (n=6)</p>
--------------------------	---------------	---

7. Conclusion

This deliverable (D2.3) examined how success in AI is articulated by key international professional bodies through an analysis of 11 major AI policies, guidelines, and standards published between 2018 to 2025. The examined corpus comprised documents issued by ACM and IEEE, as detailed in Table 1.

Utilising a mix-method approach that combined thematic analysis and topic modelling, we identified prevalent codes and topics within documents and further interpreted them to define “successful AI” through the lens of sociology of expectations. Drawing on thematic analysis and topic modelling of these documents, a range of themes has been identified that combine normative principles for AI systems, such as privacy, ethics and ethical AI, harm reduction, and human well-being, with governance mechanisms intended to operationalise these principles. These mechanisms include stakeholder engagement, impact assessment, lifecycle approaches, data governance, transparency, metrics, responsibility and legal responsibility, risk assessment and management, personal information and consent, bias, and standards. Together, these themes constitute professionally situated imaginaries of how AI systems ought to function in society, including underlying assumptions about responsibility, risk, and acceptable forms of governance. Across the corpus, four meta-themes which emerged in FORSEE Deliverable D2.2²⁰ were used to guide the analysis in this report: AI Technical Issues, AI Uses, AI Risks and Harms and AI Governance. Within these, 11 distinct themes were identified.

In terms of limitations, the selection criteria for the corpus excludes relevant national-level policies that might support professionals in this field, considering cultural and regional aspects. Second, topic modelling was conducted at the level of individual documents, which may limit the capture of cross-document topics. Finally, while the combination of thematic analysis and topic modelling strengthened the analysis through inclusion of both qualitative and quantitative methods, further methodological refinement and triangulation

²⁰ <https://zenodo.org/records/18352574>

could enhance the outcomes.

In the remainder of the broader study undertaken by the FORSEE project, we aim to explore how the EU’s “innovation-with-protections” expectation supports the professionals behind the technology: to raise attention and legitimise investments more broadly that upholds fundamental rights, e.g. in media discourse and attitudes of stakeholder classes such as professional bodies, SMEs and CSOs; in providing direction to the search processes of science and technology, e.g. in stakeholder expectations for further investments to realise this narrative; and in its coordination effect, both in the related coordination expectations of stakeholders and the forms of coordination observed in related consensus-based outputs from standards bodies and different professional bodies. Further, we will compare the set of themes identified in this document with the themes identified in document sets selected from other types of bodies considering the alignment in meta-themes already established from FORSEE Deliverable D2.2.

Positionality statement

The researchers directly involved with identification of codes, topics, and themes are:

1. Researcher #1: First year PhD candidate in computer science
 2. Researcher #2: Holds a PhD in EU law and has 3 years of postdoc experience
 3. Researcher #3: Legal scholar, senior AI policy analyst, and PhD candidate in AI
 4. Researcher #4: Holds a PhD in Communication Studies, a Full Professor of Digital Media & Communication with 25 years experience.
- The thematic analysis to identify codes was conducted by Researcher #1.
 - The list of codes referenced by Researcher #1 in the thematic analysis process was created by Researcher #2.
 - The analysis of the BERTopic outputs to label topics was conducted by Researcher #1.
 - Researcher #1 and Researcher #3 worked together to interpret codes and topics to identify themes. Researcher #3 first organised the topics and codes into initial theme groups. This was followed by a joint discussions of identified themes and potential alternate suggestions by Researcher #1. Within all themes, all disagreements were solved in the joint discussion meeting and did not require external consultation. The analysis of the identified themes through the lens of sociology of expectation were conducted using the framing provided by Researcher #4.

Acknowledgements

This work has received funding from the European Commission's Horizon Europe Research and Innovation Programme under grant agreement No. 101177579 (FORSEE) and from Research Ireland at ADAPT, the Research Ireland Centre for AI-Driven Digital Content Technology at TCD and UCD #13/RC/2106_P2. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

We are grateful to Prof. John D. Kelleher and Dr. Vasudevan Nedumpozhihana for providing valuable advice on topic modelling. We would like to extend our thanks to Prof. Eugenia Siapera and Prof. Abeba Birhane for reviewing this manuscript.



References

ACM Code of Ethics and Professional Conduct. (2018). <https://www.acm.org/code-of-ethics>

Adamson, G., Havens, J. C., & Chatila, R. (2019). Designing a Value-Driven Future for Ethical Autonomous and Intelligent Systems. *Proceedings of the IEEE*, 107(3), 518–525. <https://doi.org/10.1109/JPROC.2018.2884923>

Braun, V., & Clarke, V. (2021). *Thematic Analysis: A Practical Guide*. SAGE.

Churchill, R., & Singh, L. (2022). The Evolution of Topic Modeling. *ACM Computing Surveys*, 54(10s), 1–35. <https://doi.org/10.1145/3507900>

Cisek, K., & Kelleher, J. D. (2023). Current Topics in Technology-Enabled Stroke Rehabilitation and Reintegration: A Scoping Review and Content Analysis. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31, 3341–3352. <https://doi.org/10.1109/TNSRE.2023.3304758>

Dhakar, K. (2022). NVivo. *Journal of the Medical Library Association: JMLA*, 110(2), 270.

Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, Version II (Intelligent Systems, Control and Automation: Science and Engineering). (2019). IEEE. https://doi.org/10.1007/978-3-030-12524-0_2

Golpayegani, D., Lasek-Markey, M., Younus, A., Kerr, A., & Lewis, D. (2025). Uncovering AI Governance Themes in EU Policies Using BERTopic and Thematic Analysis. In R. Markovich, L. Di Caro, A. Rapp, & C. Schifanella (Eds.), *Frontiers in Artificial Intelligence and Applications*. IOS Press. <https://doi.org/10.3233/FAIA251621>

IEEE Ontological Standard for Ethically Driven Robotics and Automation Systems 7007. (2021). <https://doi.org/10.1109/IEEESTD.2021.9611206>

IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being 7010. (2020, May). <https://ieeexplore.ieee.org/document/9084219>

IEEE Standard for Algorithmic Bias Considerations 7003. (2024). <https://doi.org/10.1109/IEEESTD.2025.10851955>

IEEE Standard for Data Privacy Process 7002. (2022). <https://doi.org/10.1109/IEEESTD.2022.9760247>

IEEE Standard for Ethical Considerations in Emulated Empathy in Autonomous and Intelligent Systems 7014. (2024). <https://doi.org/10.1109/IEEESTD.2024.10576666>

IEEE Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems 7009. (2024). <https://doi.org/10.1109/IEEESTD.2024.10582898>

IEEE Standard for Transparent Employer Data Governance 7005. (2021). <https://doi.org/10.1109/IEEESTD.2021.9618905>

IEEE Standard Model Process for Addressing Ethical Concerns during System Design 7000 (Nos. 7000–2021). (2021). <https://doi.org/10.1109/IEEESTD.2021.9536679>

Kerr, A., Barry, M., & Kelleher, J. D. (2020). Expectations of artificial intelligence and the performativity of ethics: Implications for communication governance. *Big Data & Society*, 7(1), 2053951720915939. <https://doi.org/10.1177/2053951720915939>

Lewis, D., Hogan, L., Filip, D., & Wall, P. J. (2020). Global Challenges in the Standardization of Ethics for Trustworthy AI. *Journal of ICT Standardization*, 123–150. <https://doi.org/10.13052/jicts2245-800X.823>

Srikumar, M., Finlay, R., Abuhamad, G., Ashurst, C., Campbell, R., Campbell-Ratcliffe, E., Hongo, H., Jordan, S. R., Lindley, J., Ovadya, A., & Pineau, J. (2022). Advancing ethics review practices in AI research. *Nature Machine Intelligence*, 4(12), 1061–1064. <https://doi.org/10.1038/s42256-022-00585-2>

Terry, G., Hayfield, N., Clarke, V., & Braun, V. (2017). Thematic analysis. *The SAGE*

*Handbook of Qualitative Research in Psychology, 2(17–37), 25. es from
Supranational AI Documents Analysis*



Appendix A: Codes from Supranational AI Documents Analysis

Table A.1. List of codes created from NVivo thematic analysis on a specified subset of supranational AI policy documents, i.e. subset of corpus from FORSEE Deliverable D2.2.


Codes from Thematic Analysis	Supranational AI policy Documents
<ul style="list-style-type: none"> • (criminal) justice • (incident) reporting • (technical) documentation • academic institutions • accountability • accuracy • AI evaluation • AI governance • AI literacy and education • AI management • audit • autonomy • Benefits • bias • children's rights • codes of practice • compliance • conflict resolution • conformity assessment • consumer protection • controllability • cybersecurity • data science • democracy • digital sovereignty • discrimination • diversity and inclusion • economy • employment and future of work • empowerment • environmental protection • ethical AI • EU values • explainability • fairness • freedom of expression • fundamental rights • gender equality 	<ul style="list-style-type: none"> • Ethics guidelines for trustworthy AI - EU HLEG, 2019 • EU AI Act - EU Council and Parliament, 2024 • Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law - Council of Europe, 2024 • OECD, Recommendation of the Council on Artificial Intelligence - OECD, 2019 (amended in 2024) • UNESCO Recommendation on the Ethics of Artificial Intelligence - UNESCO, 2022 • United Nations Governing AI for Humanity: Final Report - United Nations, 2024 • G20 Ministerial Statement on Trade and Digital Economy - G20, 2019

<ul style="list-style-type: none"> ● good administration ● harm prevention ● harmonisation ● health and safety ● human agency ● human autonomy ● human dignity ● human oversight ● human rights ● human-friendly technology ● impact assessment ● innovation ● institutional actors ● intellectual property ● internal market ● knowledge ● labour ● legal certainty ● level playing field ● life cycle ● media and misinformation ● monitoring ● policy ● predictability ● privacy and data protection ● professional bodies ● proportionality ● reliability ● resilience ● risk (assessment and management) ● robustness ● rule of law ● safety and security ● sandboxes ● SME ● social expectations ● social responsibility ● society and societal impact ● stakeholders ● standards ● state of the art ● successful AI ● supranational bodies ● sustainable development ● technical organisations ● testing ● transparency 	
---	--


<ul style="list-style-type: none">• trustworthy AI• value chain• vulnerable groups	
--	--

Appendix B: NVivo Code Results


Figure B.1. NVivo results (codes) for each ACM and IEEE document.




ICT_05: IEEE-7002-2022 Standard for Data Privacy Process




ICT_06: IEEE-7003-2024 Standard for Algorithmic Bias Considerations




ICT_07: IEEE-7005-2021 Standard for Transparent Employer Data Governan




ICT_08: IEEE-7007-2021 Ontological Standard for Ethically Driven Robotics




ICT_09: IEEE-7009-2024 Standard for Fail-Safe Design



ICT_10: IEEE-7010-2020 Impact of A/I Systems on Human Well-Being



ICT_11: IEEE-7014-2024 Standard for Ethical Emulated Empathy



Appendix C: Document-Specific Comparisons of NVivo Codes and BERTopic Topic Labels

Table C.1. Comparison of codes and topics for ACM code of ethics and professional conduct.

ICT_01. ACM Code of Ethics and Professional Conduct	
Quantitative comparison	12 (BERTopic - topics) < 43 (NVivo - codes)
Unique to NVivo	<p>31 codes:</p> <ul style="list-style-type: none"> ● Fairness ● Common good ● Diversity and inclusion ● Compliance ● Consent ● Trustworthy ● Human dignity ● Health and safety ● Accountability ● Data rights ● Creator rights ● Fundamental rights ● Society and societal impact ● Stakeholders ● Social context ● Quality work ● Conflict resolution ● Cybersecurity ● Impact ● Open source ● Sustainability ● Accessibility ● Autonomy ● Principles ● Vulnerable groups ● Justice ● Transparency ● Ethical AI ● Impact assessment ● Public education

	<ul style="list-style-type: none"> • Rule of law
--	---

Table C.2. Comparison of codes and topics for IEEE Ethically Aligned Design.

ICT_02. IEEE Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems (Version 2)	
Quantitative comparison	33 (BERTopic - topics) < 84 (NVivo - codes)
Unique to NVivo	51 codes: <ul style="list-style-type: none"> • Codes of practice • Fairness • Common good • Harm reduction • Responsibility • Empowerment • Compliance • Human well-being • Risk (assessment and management) • Consent • Explainability • Consumer protection • Human dignity • Inalienable rights • Professional bodies • Health and safety • Life cycle • Accountability • Data rights • Creator rights • Accuracy • Fundamental rights • Society and societal impact • Ai misuse • Intellectual property • Human-centric • Social context • Social responsibility • Cybersecurity • Impact • Open source

	<ul style="list-style-type: none"> ● Accessibility ● Harmonisation ● Dynamic ● Controllability ● Digital sovereignty ● Surveillance ● Frameworks ● Autonomy ● Principles ● Interdisciplinary ● Vulnerable groups ● Monitoring ● Leadership ● Societal progress ● Justice ● Traceability ● Employment and future of work ● Predictability ● Bias ● Rule of law
--	---

Table C.3. Comparison of codes and topics for IEEE standard for ethical concerns during system design.

ICT_03. IEEE 7000-2021 - Standard Model Process for Addressing Ethical Concerns during System Design	
Quantitative comparison	24 (BERTopic - topics) < 60 (NVivo - codes)
Unique to NVivo	36 codes: <ul style="list-style-type: none"> ● Codes of practice ● Fairness ● Common good ● Harm reduction ● Empowerment

	<ul style="list-style-type: none"> ● Norms ● Explainability ● Human dignity ● Policy ● Professional bodies ● Health and safety ● Accountability ● Data rights ● Conformance ● Governance ● Accuracy ● Fundamental rights ● Society and societal impact ● Personal information ● Discrimination ● Social context ● Social responsibility ● Impact ● Sustainability ● Harmonisation ● Surveillance ● Frameworks ● Vulnerable groups ● Monitoring ● Societal progress ● Human agency ● Freedom of expression ● Bias ● Public education ● Human rights ● Legal responsibility
--	--

Table C.4. Comparison of codes and topics for IEEE standard for transparency of autonomous systems

ICT_04. IEEE 7001-2021 - Standard for Transparency of Autonomous Systems	
Quantitative comparison	18 (BERTopic - topics) < 47 (NVivo - codes)

<p>Unique to NVivo</p>	<p>29 codes:</p> <ul style="list-style-type: none"> ● Codes of practice ● Fairness ● Independent oversight ● Diversity and inclusion ● Public trust ● Trustworthy ● Policy ● Life cycle ● Data rights ● Conformance ● Accuracy ● Society and societal impact ● Ai misuse ● Intellectual property ● Impact ● Accessibility ● Controllability ● Dynamic ● Frameworks ● Principles ● Monitoring ● Traceability ● Ethical AI ● Guidelines ● Bias ● Impact assessment ● Privacy ● Public education ● Rule of law
-------------------------------	--

Table C.5. Comparison of codes and topics for IEEE standard for data privacy process

<p>ICT_05. IEEE 7002-2022 - Standard for Data Privacy Process</p>	
<p>Quantitative comparison</p>	<p>13 (BERTopic - topics) < 37 (NVivo - codes)</p>

<p>Unique to NVivo</p>	<p>24 codes:</p> <ul style="list-style-type: none"> ● Value-based design ● Codes of practice ● Fairness ● Responsibility ● Diversity and inclusion ● Compliance ● Ethics ● Accountability ● Data rights ● Conformance ● Governance ● Ai misuse ● Social context ● Impact ● Accessibility ● Dynamic ● Safety and security ● Audit ● Monitoring ● Leadership ● Traceability ● Transparency ● Predictability ● Impact assessment
-------------------------------	---

Table C.6. Comparison of codes and topics for IEEE standard for algorithmic bias considerations

<p>ICT_06. IEEE 7003-2024 - Standard for Algorithmic Bias Considerations</p>	
<p>Quantitative comparison</p>	<p>12 (BERTopic - topics) < 42 (NVivo - codes)</p>
<p>Unique to NVivo</p>	<p>31 codes:</p> <ul style="list-style-type: none"> ● Fairness ● Independent oversight ● Responsibility ● Empowerment ● Norms

	<ul style="list-style-type: none"> ● Ethics ● Trustworthy ● Policy ● Accountability ● Conformance ● Governance ● Society and societal impact ● Ai misuse ● Personal information ● Discrimination ● Impact ● Accessibility ● Dynamic ● Surveillance ● Frameworks ● Principles ● Monitoring ● Justice ● Human oversight ● Transparency ● Human agency ● Ethical ai ● Guidelines ● Impact assessment ● Legal responsibility ● Rule of law
<p>Unique BERTopic</p>	<p>to 1 topic:</p> <ul style="list-style-type: none"> ● Accuracy <p><i>Note: ICT_06 is the only document to have a topic unique to BERTopic but not identified as a code via the NVivo analysis</i></p>

Table C.7. Comparison of codes and topics for IEEE standard for transparent employer data governance

ICT_07. IEEE 7005-2021 - Standard for Transparent Employer Data Governance	
Quantitative comparison	19 (BERTopic - topics) < 50 (NVivo - codes)
Unique to NVivo	<p>31 codes:</p> <ul style="list-style-type: none"> ● Fairness ● Harm reduction ● Diversity and inclusion ● Empowerment ● Compliance ● Norms ● Explainability ● Ethics ● Trustworthy ● Data rights ● Conformance ● Accuracy ● Fundamental rights ● Discrimination ● Social context ● Impact ● Accessibility ● Controllability ● Surveillance ● Frameworks ● Principles ● Leadership ● Human oversight ● Traceability ● Human agency ● Guidelines ● Bias ● Impact assessment ● Predictability ● Public education ● Human rights

Table C.8. Comparison of codes and topics for IEEE ontological standard for ethical robots and automation systems

ICT_08. IEEE 7007-2021 - Ontological Standard for Ethically Driven Robotics and Automation Systems	
Quantitative comparison	16 (BERTopic - topics) < 40 (NVivo - codes)
Unique to NVivo	24 codes: <ul style="list-style-type: none"> ● Value-based design ● Fairness ● Harm reduction ● Human well-being ● Consent ● Explainability ● Trustworthy ● Health and safety ● Accountability ● Life cycle ● Stakeholders ● Impact ● Dynamic ● Frameworks ● Interdisciplinary ● Safety and security ● Audit ● Monitoring ● Justice ● Guidelines ● Predictability ● Standards ● Human rights ● Rule of law

Table C.9. Comparison of codes and topics for IEEE standard for fail-safe design of (semi) autonomous systems

ICT_09. IEEE 7009-2024 - Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems	
--	--

Quantitative comparison	11 (BERTopic - topics) < 21 (NVivo - codes)
Unique to NVivo	10 codes: <ul style="list-style-type: none"> ● Transparency ● Responsibility ● Accountability ● Governance ● Autonomy ● Principles ● Safety and security ● Standards ● Social context ● Monitoring

Table C.10. Comparison of codes and topics for IEEE recommended practice on impact assessment of AI/S on human well-being

ICT_10. IEEE 7010-2020 - Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being	
No. of themes	16 (BERTopic - topics) < 58 (NVivo - codes)
Unique to NVivo	42 codes: <ul style="list-style-type: none"> ● Value-based design ● Codes of practice ● Fairness ● Common good ● Diversity and inclusion ● Norms ● Explainability ● Trustworthy

	<ul style="list-style-type: none"> ● Human dignity ● Policy ● Life cycle ● Accountability ● Data rights ● Governance ● Society and societal impact ● Personal information ● Discrimination ● Social context ● Social responsibility ● Cybersecurity ● Accessibility ● Surveillance ● Frameworks ● Ontology ● Autonomy ● Interdisciplinary ● Principles ● Safety and security ● Vulnerable groups ● Monitoring ● Leadership ● Transparency ● Employment and future of work ● Guidelines ● Bias ● Standards ● Privacy ● Public education ● Human rights ● Autonomous weapons ● Legal responsibility ● Rule of law
--	---

Table C.11. Comparison of codes and topics for IEEE standard for ethics in emulated empathy

ICT_11. IEEE 7014-2024 - Standard for Ethical Considerations in Emulated Empathy in Autonomous and Intelligent Systems	
Quantitative comparison	27 (BERTopic - topics) < 61 (NVivo - codes)

<p>Unique to NVivo</p>	<p>35 codes:</p> <ul style="list-style-type: none"> ● Codes of practice ● Fairness ● Common good ● Metrics ● Independent oversight ● Diversity and inclusion ● Compliance ● Norms ● Public trust ● Trustworthy ● Policy ● Accountability ● Data rights ● Governance ● Fundamental rights ● Society and societal impact ● Ai misuse ● Human-centric ● Social responsibility ● Impact ● Sustainability ● Accessibility ● Controllability ● Dynamic ● Autonomy ● Interdisciplinary ● Audit ● Principles ● Vulnerable groups ● Human agency ● Freedom of expression ● Guidelines ● Predictability ● Human rights ● Legal responsibility
-------------------------------	--