Número 1 (noviembre 2026): 35-49

Tracking Ethical Principles in the Legislative Procedure of the European Artificial Intelligence Act

Angeliki Skoura[1]

Hellenic OCR Team (Greece)

Received: 05/16/2026
Accepted: 07/01/2026

https://doi.org/10.26422/REPA.2026.0100.sko

Abstract
The Artificial Intelligence (AI) Act represents a landmark regulatory framework for AI governance in the European Union. This study presents a methodology for comparing documents from the legislative procedure of the Act against seven AI ethical principles. The methodology starts with text preprocessing, and subsequently, the term frequency-inverse document frequency (TF-IDF) technique is employed to assign a weight of significance to each term in the set of documents. The similarity score between documents and ethical principles is quantified using cosine similarity. The results reveal that the proposal version shows the highest similarity scores across all principles whereas the committees/amended version has the lowest scores. The final version of the Act presents positive adjustments compared to the committees version; however, the final similarity scores do not reach those achieved by the proposal. In addition, a principle-specific analysis reveals that oversight and safety rank highest whereas well-being scores lowest across all versions.

Keywords: artificial intelligence act, legislative process, ethical principles, term-frequency analysis, cosine similarity score.

1 Angeliki Skoura has received BSc, MSc and PhD from Computer Engineering and Informatics Department, University of Patras, Greece. Her research focuses on data analysis including datasets of time-series, images and text. She is currently working at Hellenic Parliament and she is a member of the Hellenic Optical Character Recognition Team. agg.skoura@gmail.com / https://orcid.org/0000-0002-6381-8155

Seguimiento de los principios éticos en el procedimiento legislativo de la Ley de Inteligencia Artificial de la Unión Europea

Resumen

La Ley de Inteligencia Artificial (IA) constituye un marco regulador de referencia para la gobernanza de la IA en la Unión Europea. Este estudio presenta una metodología para comparar los documentos correspondientes al procedimiento legislativo de dicha ley con respecto a siete principios éticos de la IA. La metodología comienza con un proceso de preprocesamiento de textos y, posteriormente, emplea la técnica de frecuencia de términos-frecuencia inversa de documentos (TF-IDF, por sus siglas en inglés) para asignarle un peso de relevancia a cada término del conjunto de documentos. El grado de similitud entre los documentos y los principios éticos se cuantifica mediante la similitud del coseno. Los resultados revelan que la versión de la propuesta presenta los mayores índices de similitud en relación con todos los principios, mientras que la versión de las comisiones parlamentarias y sus enmiendas registran los índices más bajos. La versión definitiva de la ley incorpora ajustes positivos en comparación con la versión de las comisiones; sin embargo, los índices finales de similitud no alcanzan los valores obtenidos por la propuesta inicial. Además, el análisis específico de cada principio muestra que la supervisión y la seguridad ocupan las posiciones más destacadas, mientras que el bienestar obtiene la puntuación más baja en todas las versiones analizadas.

Palabras clave: ley de inteligencia artificial, procedimiento legislativo, principios éticos, análisis de frecuencia de términos, índice de similitud del coseno.

1. Introduction

The Artificial Intelligence Act (AI Act), formally known as Regulation (EU) 2024/1689 (European Parliament and Council of the European Union, 2024), is a regulatory framework proposed by the European Commission to regulate AI in the European Union (EU). The regulation aims to establish a legal framework for AI, ensuring its safe and ethical development and deployment, while fostering innovation and competition. The legislative procedure of the AI Act included numerous documents. The process began with the European Commission drafting a legislative proposal (European Commission, 2021). The proposal for the AI Act was presented in April 2021, as a first draft of the legal framework and was based on consultations, impact assessments and input from stakeholders such as AI researchers, industry representatives and the public. In a later stage, the Committees of the European Parliament presented their report to amend the legislation and proposed new provisions (European Parliament, 2024). These amendments were debated in the plenary sessions of the European Parliament. The AI Act was adopted in June 2024, once the Parliament and the Council reached final agreement. Following adoption, the AI Act was published in the Official Journal of the European Union. The steps of the complete parliamentary procedure are presented in (European Commission, 2025).

While the AI Act aims to align regulatory mechanisms with foundational ethical values such as transparency, fairness, accountability and human oversight, the codification of these often ambiguous and context-sensitive principles into legally binding norms remains a complex endeavor. Ethical principles are typically aspirational and interpretive, lacking the specificity required for consistent legal implementation (Stix, 2021). This creates a gap between normative ideals and enforceable obligations, complicating both compliance and evaluation. Moreover, the AI Act's risk-based classification system reflects the ethical implications of AI systems, particularly those with emergent or cumulative harms that fall outside predefined risk categories. The rapidly evolving nature of AI technologies further exacerbates these difficulties, as ethical standards and societal expectations may shift faster than regulatory frameworks can adapt. Consequently, assessing these documents in relation to fundamental AI principles poses a number of conceptual and practical challenges.

The most official and foundational document outlining the EU's AI ethical principles is the "Ethics Guidelines for Trustworthy AI", published by the European Commission's High-Level Expert Group on Artificial Intelligence in April 2019 (European Commission High-Level Expert Group on Artificial Intelligence, 2019). This document sets out seven key requirements for trustworthy AI to guide the development and deployment of AI systems in a manner that is lawful, ethical and robust: (i) human agency and oversight, (ii) technical robustness and safety, (iii) privacy and data governance, (iv) transparency, (v) diversity, non-discrimination and fairness, (vi), societal and environmental well-being, and (vii) accountability. These guidelines have been serving as the ethical foundation for subsequent policy initiatives, including the EU AI Act; however, they are non-binding.

The objective of this paper is twofold: the comparison of three documents across the legislative procedure of the AI Act against the seven AI ethical principles and the identification of the principles that are most represented in the three document versions. At a high-level, the methodology leverages Natural Language Processing (NLP) to prepare the textual data, generate quantitative document representations and compute similarity scores which capture the alignment between legislative documents and ethical principles. By systematically evaluating the AI Act's legislative documents, this study contributes to a deeper understanding of the regulatory framework and informs policymakers about inconsistencies in AI governance.

The rest of the paper is organized as follows. Section 2 reviews the state-of-the-art text mining techniques for legal documents as well as AI ethics. Section 3 presents the dataset used in this study and Section 4 describes the methodology. Section 5 presents the results, which include the document-specific analysis to capture the differences across the legislative procedure and the principle-specific analysis to detect how the documents prioritize AI principles. Finally, Section 6 discusses the findings, and Section 7 concludes with an interpretation of the results, along with future research plans to enhance the accuracy of text analysis.

2. State-of-the-art

2.1 Text mining techniques for the analysis of EU legal documents

The increasing complexity and volume of EU legal texts, mainly directives, regulations and policy papers, have prompted the adoption of computational linguistics, particularly text mining, to facilitate legal analysis (Moens, 2006). In particular, the application of text mining techniques to legal documents has gained increasing traction over the past decade, driven by the growing availability of digital legislative corpora such as EUR-Lex.

Early work in this field focused on basic information extraction aiming to identify named entities, such as institutions and legal instruments, as well as legal actions such as mandates, obligations and exemptions. Named Entity Recognition models have been trained on legal corpora, and used to extract structured information from EU regulations and directives, enabling better indexing and search functionalities (Boella et al., 2013a). Topic modeling methods, particularly Latent Dirichlet Allocation and Non-negative Matrix Factorization, have been used to identify thematic structures within large sets of EU legal texts. These techniques enable the discovery of clusters around regulatory domains such as digital policy, environmental law, or labor rights (Van Opijnen & Santos, 2017). Additionally, advanced methods of semantic analysis have enabled finer-grained analyses of legal discourse. These models capture semantic similarity between concepts, enabling cross-linguistic alignment of directives and improved retrieval across legal systems (Zhu et al., 2021).

Among the foundational methods in legal text analysis, the TF-IDF technique remains a widely employed and effective approach for information retrieval, feature extraction and

document representation. This technique is particularly suitable for handling high-dimensional sparse data, a characteristic feature of legal corpora (Salton and Buckley, 1988). Within EU legal document analysis, TF-IDF has been employed for tasks such as legal document classification, similarity search, and clustering of legal texts (Mandal et al., 2022). Furthermore, TF-IDF vectors have also served as effective input features for unsupervised learning techniques, including K-means clustering, hierarchical clustering, and topic modeling. These methods have been applied to cluster EU legal documents by thematic areas, such as digital policy, competition law, or environmental regulation (Boella et al., 2013b). Additionally, TF-IDF has been integrated with dimensionality reduction techniques such as Principal Component Analysis to visualize the semantic landscape of legislative texts (Flaounas et al., 2013). This has proven especially useful in examining the evolution of legal discourse over time or across policy areas.

In terms of the AI Act, researchers and practitioners have increasingly turned to text mining techniques to analyze the act's content, interpret its implications and assess its alignment with ethical principles. NLP techniques including tokenization and part-of-speech tagging have been used to facilitate the structural breakdown of the proposal of the AI Act, assisting in the identification of obligations, prohibitions and definitions critical to compliance and enforcement (Veale & Zuiderveen Borgesius, 2021). Researchers have also employed topic modeling techniques, notably Latent Dirichlet Allocation, to uncover underlying themes and topics within the AI Act (Cornerstone Research, 2024). Such techniques allowed for the identification of prevalent themes, such as risk classification, transparency requirements and governance structures and provided insights into the text of the Act. Furthermore, semantic similarity measures and word embeddings (such as Word2Vec and GloVe) have been utilized to assess the alignment of the provisions of the AI Act's with established ethical principles. By comparing the semantic content of the Act with existing ethical guidelines, researchers evaluated the extent to which the legislation embodies principles like fairness, accountability and transparency (Zeng et al., 2022). Sentiment analysis has been employed to process textual data from consultations, social media and public comments with regards to the AI Act (Cornerstone Research, 2024). This analysis informed policymakers about the reception of the Act and highlighted areas requiring clarification or amendments. Finally, information extraction methods coupled with the development of legal ontologies have been used to facilitate automated reasoning, compliance checking and aligning with the AI Act's documentation requirements (Golpayegani et al., 2024).

Despite the advancements in text mining applications, challenges persist. The complexity of the legal language and context-dependence of EU legislation pose difficulties for accurate automated analysis. Moreover, the dynamic evolution of AI technologies necessitates continuous updates to legal ontologies and compliance frameworks.

2.2 Text mining techniques for the analysis of AI ethical principles

AI ethics documents present a heterogeneous and often unstructured body of texts. While they generally reference common values, such as transparency, fairness, accountability, and privacy, their formulations vary significantly in terminology, granularity and legal status (Fjeld et al., 2020). This complexity makes them ideal candidates for text mining in order to uncover latent patterns, identify normative themes, and track the diffusion of ethical principles across documents and jurisdictions.

Initial studies have employed keyword frequency analysis and term co-occurrence mapping to identify dominant ethical principles and assess their prominence within and across documents (Hagendorff, 2020). For instance, Jobin et al. (2019) used frequency-based methods to show widespread agreement on core principles like beneficence and non-maleficence, while revealing less consensus on concepts such as sustainability or solidarity. Recent studies have incorporated distributional semantics and transformer-based language models to enhance semantic interpretation. These models have been used to measure similarity between principles, align ethics documents across cultural or institutional boundaries, and extract conceptually equivalent expressions of values (Zeng et al., 2022). For example, Sánchez-Monedero et al. (2020) leveraged word embeddings to map how ethical concepts are operationalized differently in corporate versus governmental guidelines. Additionally, text classification techniques, both supervised and unsupervised, have been applied to annotate ethical principles in legal and technical documents. Annotated datasets such as the AI Ethics Guidelines Global Inventory have facilitated the training of classifiers that can identify passages related to specific principles, such as fairness or transparency (Morley et al., 2020). Emerging research emphasizes the importance of multimodal data fusion (e.g., combining text mining with legal analysis or stakeholder interviews) for the text mining of AI ethics (Morley et al., 2020; Zeng et al., 2022).

A number of constraints nonetheless remain. The lack of standardized taxonomies for ethical principles complicates inter-study comparison. Additionally, ethical terms often exhibit semantic ambiguity and cultural variability, which can hinder computational modeling.

3. Dataset

This study focuses on three documents of the legislative procedure on the adoption of the AI Act, namely the proposal by the European Commission in April 2021, the landmark agreement composed by the European Parliament's Committees in February 2024 and the adopted text of the AI Act as published in the Official Journal of the European Union in July 2024.

Table 1 outlines the basic differences among the documents. The characteristics of documents are presented horizontally and include description of the legal status, purpose, main authors and publication date of the documents.

Table 1
A comparison of the basic characteristics of the documents related to this study

	Ethics Guidelines for Trustworthy AI	Proposal version	Committee/ Amended version	Final version
Legal Status	Non-binding guidelines	Draft legislative proposal	Political agreement (Trilogue outcome)	Legally binding regulation
Purpose	Promote ethical principles and trust in AI	Establish regulatory framework	Negotiate legal/ political compromises	Enforce comprehensive EU-wide AI regulation
Main Authors/ Actors	High-Level Expert Group on AI	European Commission	European Parliament, Council, Commission	European Parliament and Council
Publication Date	April 2019	April 2021	February 2024	July 2024

Source: Prepared by the author.

The seven fundamental AI ethical principles used in our study are derived from the document entitled "Ethics Guidelines for Trustworthy AI" (European Commission, 2019). Each principle is presented with a heading (for example, "1. Human agency and oversight"), followed by one or more paragraphs explaining the meaning of the principle and the ethical rationale behind it. The text used to represent each principle is the text of the corresponding section in the document. A summary of the main characteristics of the seven principles is presented in Table 2. The second column includes the name of the ethical principle according to the document and the third column presents an abbreviation of the principle used for simplicity in the rest of this paper. The key elements of each ethical principle are presented in the fourth column.

Table 2
A summarization of the seven ethical principles presented by the European Commission as guidelines for trustworthy AI

No	Ethical Principle	Abbreviation	Main considerations
1	Human Agency and Oversight	Oversight	Human-in-the-loop, human autonomy and decision making
2	Technical Robustness and Safety	Safety	General security, reliability, resilience against errors and attacks
3	Privacy and Data Governance	Privacy	GDPR compliance, data minimization
4	Transparency	Transparency	Explainability of AI decisions, documentation
5	Diversity, Non-discrimination and Fairness	Fairness	Avoidance of model bias, inclusive training datasets
6	Societal and Environmental Well-being	Well-being	Sustainability, promotion of the common good
7	Accountability	Accountability	Auditability, clear lines of responsibility

Source: Prepared by the author.

The corpus of this study consists of ten documents: the three documents from the legislative procedure of the AI act (Proposal, Committees' version and Final Text) and the seven texts, each representing one of the corresponding ethical principle identified in the said document. The corpus analyzed in this study includes publicly available textual documents obtained from official EU sources.

4. Methodology

Our methodology consists of four steps: (i) the preprocessing of the texts to eliminate common words and to perform lemmatization, (ii) the employment of the TF-IDF technique to assign a weight of significance to each term presented in the set of documents, (iii) the representation of each documents as a vector based on the TF-IDF weights of its containing terms, and (iv) the computation of similarity between each legislative document and each ethical principle using the cosine similarity metric.

The aim of text preprocessing before the application of the TF-IDF technique is to clean and normalize the original text data to ensure consistency and accuracy in term extraction. The relevant steps analytically are:

- Tokenization: Splitting the text into individual words (tokens), which are the basic units for TF-IDF calculation.
- Stopword removal: Elimination of common words (such as "the", "and", "is") as they carry little semantic value and can skew frequency counts.
- Lemmatization: Reducing words to their base forms to group variations of a word under a single term.

Additionally, all texts are converted to lowercase to avoid case-sensitive discrepancies, and non-alphanumeric characters (such as special characters and punctuation) are removed to focus on meaningful terms.

After preprocessing, the TF-IDF weight (also known as the TF-IDF score) for each term in the corpus is calculated using Equation 1. This formula quantifies how important a word is to a specific document in the collection (corpus).

$$TF - IDF\ (t, d) = TF(t, d) \times IDF(t) \quad (Equation\ 1)$$

where:

$TF\ (t, d) = f_{t,d}/N_d$ is the Term-Frequency (TF), which measures how often a term (t) appears in a document (d), normalized by the total number of words in that document (N_d), and $IDF\ (t) = \log\ (N/n_t)$ is the Inverse Document Frequency (IDF), which measures how unique a term (t) is, where N is the total number of documents and n_t is the number of documents that contain the term t.

The idea behind the TF-IDF score is to identify words that are important or relevant to a specific document in a corpus by balancing local frequency and global rarity. The combination of TF and IDF assigns greater importance to terms that appear frequently in an individual document while occurring infrequently across the corpus.

Once the TF-IDF scores have been calculated for each unique term in the corpus, each document in the corpus is represented as a vector containing the weights of the terms contained. For a corpus with N unique terms in total, each document is represented as an N-dimensional vector and each entry in the vector corresponds to the TF-IDF score of the respective term in that document. In case a term is not present in a document, its TF-IDF score, and thus the corresponding vector component, is 0.

After calculating the TF-IDF vectors for the documents in the corpus, cosine similarity is used to measure the similarity between pairs of documents. Cosine similarity score computes the cosine of the angle between two TF-IDF vectors, which reflects how similar the documents are in terms of their content, regardless of their length. The formula for the cosine similarity score between two documents in a corpus is presented below:

$$\text{cosine similarity}\,(D_i, D_j) = \frac{D_i \cdot D_j}{\|D_i\| \cdot \|D_j\|}$$

where $D_i \cdot D_j$ is the dot product of TF-IDF vectors of documents D_i and D_j, and $\|D_i\|$ and $\|D_j\|$ are the Euclidean norms of the vectors correspondingly.

The cosine similarity score of TF-IDF vectors is a non-negative value ranging from 0 to 1. Values of cosine similarity score closer to 1 indicate greater similarity between two documents as a result of similar weighted term distributions**.** On the contrary, values of cosine similarity score closer to 0 indicate two unrelated documents. Cosine similarity is particularly suitable for comparing TF-IDF vectors of documents with different lengths because it measures the angle between the vectors rather than their magnitude, allowing documents of different lengths to be compared effectively.

In terms of implementation of the methodology, a text preprocessing and analysis pipeline was implemented in Python (version 3.12) using established NLP and machine learning libraries. The documents were processed using the spaCy library, including tokenization, conversion to lowercase, and removal of common stopwords. Lemmatization was conducted using the spaCy English language model. For the transformation into TF-IDF vector representations, the Tfidf Vectorizer implementation was used from the Scikit-learn library. The similarity between legislative documents and AI ethical principles was quantified using the cosine similarity function from Scikit-learn.

5. Results

5.1 Similarity scores between documents

Table 3 presents the cosine similarity scores for all the pairs between the three document versions and the seven ethical principles. Additionally, the average similarity score of each document across all principles has also been calculated in order to understand the average similarity of each document with the full set of principles. The average similarity of the three documents across all principles is also presented.

Table 3
The cosine similarity scores between the legislative documents and the seven AI ethical principles

	Documents in the Legislative Process		
	Proposal	Committees	Final
Transparency	0.599	0.47	0.53
Accountability	0.542	0.429	0.476
Oversight	0.632	0.485	0.545
Wellbeing	0.393	0.322	0.367
Fairness	0.54	0.419	0.478
Privacy	0.432	0.344	0.39
Safety	0.622	0.51	0.579
Average Similarity Score	0.537	0.426	0.481

Source: Prepared by the author.

5.2 Document-specific comparison

Based on the similarity analysis among the documents, a key observation is that the proposal version consistently shows the highest similarity scores across all principles. On average, the proposal exhibits the highest alignment with ethical principles (Average: 0.537), suggesting a strong initial ethical foundation. The committees/amended version has the lowest scores indicating that amendments reduced alignment with the linguistic description of the ethical principles. The amended version shows a significant decline in scores (Average: 0.426). The final version improves upon the committees/amended version, showing positive adjustments and suggesting a partial restoration of ethical alignment (Average: 0.481). However, the similarity scores of the final version do not reach the level of similarity score of the proposal.

5.3 Principle-specific comparison

Comparing across principles, Oversight and Safety consistently rank highest across all versions, with score 0.632 and 0.622 respectively for the proposal. This implies that the document inherently emphasizes human governance and risk management which are retained even after amendments. Wellbeing scores the lowest across all versions (Proposal: 0.393, Final: 0.367), indicating that it is the least aligned principle in the documents. Privacy is the second-lowest scoring principle across all versions, with significant decline during amend-

ments. Additionally, to identify the most affected principles, the percentage change between the Proposal and Final versions is calculated. Oversight and Accountability present the largest declines (-13.8% and -12.2% respectively) from the proposal to the final version.

6. Discussion

The text analysis of this study reveals a pattern within the legislative procedure of the AI Act: the Proposal version exhibits the highest lexical similarity across all seven ethical principles based on TF-IDF/cosine similarity scores. This finding suggests that the Proposal may represent an "ideal" ethical standard, while later versions could reflect real-world constraints such as stakeholder requests.

The comparative analysis among the ethical principles reveals critical insights into how each ethical principle evolved from the Proposal to the Final version. The stability of Safety, the highest-scoring principle, suggests risk management remained a priority. The consistently low scores for Wellbeing across all document versions suggest a systemic issue in how this ethical dimension is defined or is addressed. One plausible explanation is the use of ambiguous, broad language or lack of explicit metrics in terms of Wellbeing, which may dilute its measurable alignment with the content of the documents. The data reveals that Privacy consistently underperforms compared to other ethical principles across all document versions. This possibly indicates that external regulatory frameworks like the General Data Protection Regulation (GDPR) (European Parliament and Council of the European Union, 2016) may influence scoring. As the GDPR sets minimum legal requirements (such as data anonymization and user consent), Privacy can be seen as a compliance checkbox rather than an aspirational ethical goal. Finally, the fact that the most pronounced declines occurred in Oversight and Accountability may reflect amendments that prioritized operational flexibility over the original ethical alignment of the proposal version.

A limitation of this study is that the proposed methodology measures lexical similarity rather than substantive ethical alignment. TF-IDF-based representations emphasize the frequency and distinctiveness of terms within documents and cosine similarity captures similarity in the corresponding vector representations. Consequently, high similarity scores indicate greater overlap in the use of ethics-related terminology but do not necessarily imply equivalent interpretation or implementation commitment to the underlying principles. The findings should therefore be considered exploratory and interpreted alongside qualitative assessments of legislative intent and legal provisions.

7. Conclusion

This study proposes a methodology for assessing the alignment between legislative documents and AI ethical principles through NLP techniques. By combining text preprocessing, TF-IDF vectorization and cosine similarity, the study provides a systematic and reprodu-

cible means of quantifying the extent to which ethical principles are reflected in successive versions of legislative texts. Applied to the legislative process of the EU AI Act, the methodology offers insights into how ethical considerations evolve throughout the drafting and negotiation processes.

The results provide an exploratory indication of the extent to which terms associated with AI ethical principles are represented across different versions of the AI Act during the legislative process. The findings indicate that, on average, the initial Proposal exhibits the highest overall alignment with the studied AI ethical principles. The Committees version presents the minimum similarity among the three legislative documents. Although the Final version shows positive adjustments towards the ethical principles compared to the committees version, it does not reach the level observed in the proposal. Furthermore, the principle-specific comparison highlights differences among ethical principles and reveals a ranking of the ethical principles for each document of the legislation process. Oversight and safety consistently receive greater representation than wellbeing throughout the legislative process.

Considering the variations in terminology that may arise across different texts or institutional authors and the limitation of the TF-IDF text mining approach to handle synonymity, our future research plans include the use of word embeddings which can capture the semantic similarity between words. An enriched set of terms, including both original words and their high-confidence synonyms, is expected to result in a more accurate capture of relevant linguistic variations by incorporating synonymous terms related to the AI ethics principles. In addition, further study of the ongoing bibliography on the AI act and especially on stakeholder feedback, would assist in a more detailed interpretation of the results and, more specifically, clarify whether the changes of ethical alignment across the parliamentary procedure reflect targeted improvements or unintended shifts in ethical emphasis.

Bibliography

Boella, G., Di Caro, L., Humphreys, L., & Robaldo, L. (2013a). *Using NLP to build legal ontologies in EU legislation. 14th International Conference on Artificial Intelligence and Law (ICAIL 2013)*. Rome, Italy, June 10-14, 2013.

Boella, G., Di Caro, L., & Robaldo, L. (2013b). *Semantic classification of legal texts. 14th International Conference on Artificial Intelligence and Law (ICAIL 2013)*. Rome, Italy, June 10-14, 2013.

Cornerstone Research. (2024). *The EU AI Act: Using AI to Analyze the Public Response*. https://www.cornerstone.com/insights/articles/eu-ai-act-analysing-public-response-with-ai/

European Commission. (2019). *Ethics Guidelines for Trustworthy AI*. High-Level Expert Group on Artificial Intelligence. https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai

European Commission. (2021). *Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021PC0206

European Commission. (2025). *Procedure 2021/106 – Artificial Intelligence Act*. EU Law Tracker. https://law-tracker.europa.eu/procedure/2021_106?lang=en

European Parliament. (2024). *Provisional Agreement from interinstitutional negotiations. Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Committee on the Internal Market and Consumer Protection. https://www.europarl.europa.eu/meetdocs/2014_2019/plmrep/COMMITTEES/CJ40/AG/2024/02-13/1296003EN.pdf

European Parliament and Council of the European Union. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*. https://eur-lex.europa.eu/eli/reg/2016/679/oj

European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). *Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI*. Berkman Klein Center Research Publication No. 2020-1. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3518482

Flaounas, I., Ali, O., Turchi, M., Shrimpton, L., & Cristianini, N. (2013). Research Methods in the Age of Digital Journalism. *Digital Journalism, 1*(1), 102-116. https://www.tandfonline.com/doi/pdf/10.1080/21670811.2012.714928

Golpayegani, D., Hupont, I., Panigutti, C., Pandit, H. J., Schade, S., O'Sullivan, D., & Lewis, D. (2024). *AI Cards: Towards an Applied Framework for Machine-Readable AI and Risk Documentation Inspired by the EU AI Act*. *arXiv preprint*. https://doi.org/10.48550/arXiv.2406.18211

Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines, 30*(1), 99-120. https://doi.org/10.1007/s11023-020-09517-8

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence, 1*(9), 389-399. https://doi.org/10.1038/s42256-019-0088-2

Mandal, K., Das, A., & Ghosh, S. (2022). Legal Document Classification using TF-IDF and Machine Learning Techniques. *International Journal of Computer Applications, 184*(11), 1-6.

Moens, M. F. (2006). *Information Extraction: Algorithms and Prospects in a Retrieval Context*. Springer. https://link.springer.com/book/10.1007/978-1-4020-4993-4

Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. *Science and Engineering Ethics, 26*, 2141-2168. https://link.springer.com/article/10.1007/s11948-019-00165-5

Van Opijnen, M., & Santos, C. (2017). On the concept of relevance in legal information retrieval. *Artificial Intelligence and Law, 25*(1), 65-87. https://doi.org/10.1007/s10506-017-9195-8

Salton, G., & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Information Processing & Management, 24*(5), 513-523. https://doi.org/10.1016/0306-4573(88)90021-0

Sánchez-Monedero, J., Dencik, L., & Edwards, L. (2020). What does it mean to "solve" the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. *ACM Conference on Fairness, Accountability, and Transparency (FAT*)*. Barcelona, Spain, January 27-30, 2020. https://doi.org/10.1145/3351095.3372849

Stix, C. (2021). Actionable Principles for Artificial Intelligence Policy: Three Pathways. *Science and Engineering Ethics, 27*(15). https://doi.org/10.1007/s11948-020-00277-3

Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the Draft EU Artificial Intelligence Act. *Computer Law Review International, 22*(4), 97-112. https://doi.org/10.9785/cri-2021-220402

Zeng, Y., Lu, E. Y. H., & Huangfu, C. (2022). *Linking Artificial Intelligence Principles*. arXiv https://doi.org/10.48550/arXiv.1812.04814

Zhu, J., Zhang, X., Zhang, Y., & Zhang, Z. (2021). A Comparative Study of Semantic Similarity Measures for Legal Text Analysis. *Journal of Information Science, 47*(2), 170-183.

Transparency in the use of artificial intelligence

The author used the free version of ChatGPT (OpenAI) in June 2025 for language refinement and for code generation in Python. The author checked and approved all outputs and retain full responsibility for the manuscript's accuracy, originality, integrity and conclusions.

Authorship roles notice

The author declares having performed all authorship roles for this article.

Conflict of interest notice

The author has no conflict of interest.