

Boglárka Matúz

Implementation of AI systems in regulated environments for Life Sciences.

Master's Thesis

Supervised by Toni Manzano

Master's Degree in Biomedical Data Science



Amsterdam, 2025

Abstract

With the rapid integration of artificial intelligence (AI) into healthcare and life sciences, regulatory bodies have issued a growing number of guidelines and frameworks to govern its safe and ethical use. This thesis investigates whether these regulatory documents around AI genuinely prioritize patient well-being or are predominantly shaped by bureaucratic and procedural considerations. To explore this, a structured keyword-based content analysis was conducted on five major regulatory documents from the United States and Europe. Two analytical approaches were applied: a comprehensive full-document analysis and a refined content-focused analysis that excludes peripheral sections such as glossaries, references, and appendices.

Normalized keyword frequencies were used to calculate comparative scores reflecting either patient-centered or bureaucratic language. A benchmark was established using the average and first quartile of the scores to distinguish documents that meet acceptable levels of patient-centeredness. Findings indicate significant variation in focus across the documents, with some exhibiting a stronger bureaucratic orientation.

Keywords: Artificial Intelligence, Healthcare Regulation, Policy Analysis, Patient-Centeredness, Bureaucracy

Conflict of interest statement

The author declares no conflict of interest.

Table of contents

1. Introduction	1
1.1. Background and Motivation	1
1.2. Objectives	2
1.3. Thesis Structure	3
2. Literature Review	5
2.1. Artificial Intelligence in Healthcare	5
2.2. Regulatory Frameworks, Guidelines, and Key Milestones	7
3. Methodology	20
4. Results	23
4.1. Results from Comprehensive Document Evaluation	23
4.2. Results from Focused Main-Content Evaluation	29
4.3. Benchmark Development and Threshold Criteria	35
5. Discussion	37
5.1. Discussion of Results	37
5.2. Broader Discussion	38
6. Limitations and Future Work	40
6.1. Limitations	40
6.2. Future Work	40
7. Conclusion	42
8. Ethical-Social Impact, Sustainability, and Diversity	44
8.1. Ethical-Social Impact	44
8.2. Sustainability	44
8.3. Diversity	45
References	46

List of figures

1. Figure: Schematic representation of the roadmap for deploying effective ML systems in healthcare.....	6
2. Figure: Timeline of key guidelines and milestones related to AI implementation in life sciences.....	8
3. Figure: Workflow of AI models based on the good practices considering AI/ML as software as a medical device.	12
4. Figure: Schema of the risk-based classification of AI Systems based on the EI AI Act, including examples and requirements.	15
5. Figure: Schema describing the seven-step risk-based credibility assessment framework for AI in drug development.	17
6. Figure: Total number of patient-centered keywords	27
7. Figure: Total number of bureaucratically driven keywords	28
8. Figure: Summary of normalized keyword scores.....	29
9. Figure: Total number of patient-centered keywords in the main content,	33
10. Figure: Total number of bureaucratically driven keywords in the main content.	34
11. Figure: Summary of normalized keyword scores based on the main content analysis ...	35
12. Figure: Risk Matrix of Regulatory Abstraction and Patient Impact.....	38

List of tables

1. Table: FDA’s risk categorization regarding SaMD.....	13
2. Table: Frequency of Keywords in Proposed Regulatory Framework for SaMD.	23
3. Table: Frequency of Keywords in the Executive Order No. 14110.	24
4. Table: Frequency of Patient-centered and Bureaucratically Driven Keywords in the EU AI Act.	24
5. Table: Frequency of Keywords in the EMA Reflection paper	25
6. Table: Frequency of Keywords in the FDA Considerations for the use of AI	25
7. Table: Comparative summary of normalized keyword scores	28
8. Table: Frequency of Keywords in the main content of the Proposed Regulatory Framework for SaMD.	30
9. Table: Frequency of Keywords in the main content of the EU AI Act.	31
10. Table: Frequency of Keywords in the main content of the EMA Reflection Paper.	31
11. Table: Frequency of Keywords in the main content of the FDA Considerations for the use of AI.	32
12. Table: Comparative summary of normalized keyword scores based on the main content analysis	35
13. Table: Summary of benchmark components and threshold values	36

1. Introduction

This chapter introduces the background and motivation for the research presented in this thesis. This is followed by a statement of the research objectives. Finally, the structure of the thesis is presented to provide an overview of its organization and the logical flow of its chapters.

1.1. Background and Motivation

Artificial Intelligence (AI) is transforming the life sciences sector by enabling innovations across drug discovery, diagnostics, patient monitoring, and clinical trial optimization (Topol, 2019). These systems offer tremendous potential for accelerating research, reducing costs, and improving patient outcomes. However, the use of AI in such a highly regulated field introduces significant challenges. Regulatory bodies like the U.S. Food and Drug Administration (FDA), the European Medicines Agency (EMA), and frameworks such as Good Machine Learning Practice (GMLP) and Good Clinical Practice (GCP) require that systems be transparent, auditable, and reliable (FDA, 2021).

One of the objectives of this thesis is to analyse these existing regulations and frameworks for data governance, AI algorithm qualification, AI model validation, and AI model maintenance within regulated life sciences environments. Given the life sciences sector's commitment to patient well-being, the deployment of AI systems in this field carries profound social implications. If governance around AI applications lacks a robust approach, fairness in model development may be compromised, which can lead to incomplete, unreliable models that cannot be safely applied in patient care.

The FDA's recent draft guidance underscores both the promise and the complexity of AI's integration into regulatory processes. It highlights how AI applications are increasingly being used to support regulatory decision-making throughout the drug development lifecycle. Examples include reducing reliance on animal studies, using predictive modeling for clinical pharmacokinetics, integrating diverse datasets to understand disease progression, and analyzing large-scale real-world data for developing clinical trial endpoints. However, the document also identifies key concerns that mirror the challenges addressed in this thesis, such as variability in dataset quality and

representativeness, methodological opacity, the uncertainty of model output, and the phenomenon of model drift, whereby an AI model's performance may degrade over time due to changes in input data (FDA, 2025). These issues reaffirm the need for robust model lifecycle management and highlight the importance of transparent, explainable AI development processes.

Additionally, due to the fast-paced technological advancements, there is a significant gap in aligning AI solutions with regulatory expectations. The lack of robust model validation processes, poor data quality, and difficulties in maintaining AI systems over time can put patient safety and data integrity at risk (Wiens et al., 2019). The main motivation behind this thesis is to explore how this gap between theoretical regulatory frameworks and real-world AI deployment challenges can be bridged to ensure trustworthiness and robustness of AI systems in the life sciences' sector, meanwhile the wellbeing and safety of patients remains the center of focus. Through a critical analysis of regulatory documents, as well as industry perspectives, the research will explore methodologies that enhance the qualification, validation, and maintenance of AI models. By ensuring that AI models are developed with fairness, transparency, and inclusivity, this investigation aims to contribute to the broader discourse on responsible AI adoption in life sciences, ensuring that AI-driven solutions align with evolving compliance requirements while maintaining efficiency.

1.2. Objectives

The primary objective of this thesis is to critically examine the current landscape of regulatory documents related to AI, with a specific focus on the extent to which these documents prioritize patient well-being. This research seeks to explore whether the guidelines and frameworks issued by regulatory bodies are genuinely centered on improving patient outcomes, or whether they are predominantly shaped by bureaucratic, procedural, and administrative considerations.

To achieve this, the study undertakes a systematic keyword-based analysis of a selection of influential regulatory texts. Through this analysis, it aims to identify thematic trends and linguistic patterns that may serve as indicators of either a patient-centered or bureaucratically driven orientation.

In addition to providing insights into the substantive focus of current regulatory efforts, this thesis also proposes an initial benchmarking approach. By applying a normalization technique and setting a comparative threshold, the study introduces a preliminary metric to help distinguish between documents that demonstrate a clear emphasis on patient-centered values and those that may be more aligned with institutional or regulatory formalism. This benchmark may serve as a foundational tool for future research and policy evaluation in the domain of responsible AI governance in healthcare.

1.3. Thesis Structure

This thesis is organized into eight sections, each contributing to a comprehensive investigation of the regulatory landscape surrounding artificial intelligence (AI) in the life sciences sector.

Section 2 presents the literature review, where relevant researches and regulatory frameworks and guidelines are examined. Each document is summarized with particular emphasis on its implications for the life sciences sector, including healthcare and the pharmaceutical industry. The section concludes with a synthesis and comparative discussion that highlights overarching themes and gaps identified in the reviewed materials.

Section 3 outlines the methodology used in this study, describing the content analysis approach applied to five key documents:

- Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD): Discussion Paper and Request for Feedback (FDA, 2019),
- Executive Order No. 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Exec. Order No. 14110, 2023),
- The European Union Artificial Intelligence Act (European Union, 2024),
- Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle (European Medicines Agency, 2024),
- Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (Draft Guidance) (FDA, 2025).

This section details the rationale for document selection, the keyword frequency approach used for analysis, and the normalization technique applied to compare textual content across documents of varying lengths. In addition, a benchmark was established to provide a reference point for future regulatory guidelines.

Section 4 presents the results of the analysis, offering a comparative assessment of how each regulatory document aligns with either a patient-centered or bureaucratically driven orientation, based on keyword frequency.

Section 5 offers a discussion of the findings, interpreting the implications of the results within the broader context of AI governance in life sciences, and relating them back to the central research question.

Section 6 addresses the limitations of the study, acknowledging methodological constraints and offering suggestions for future research that could enhance or extend the current analysis.

Section 7 provides the conclusions drawn from the study as a whole, summarizing the key insights and contributions of the research.

Finally, section 8 explores broader considerations, including the ethical and social implications, as well as aspects related to sustainability and diversity.

2. Literature Review

Although artificial intelligence (AI) is not a new technology, recent advances in computing power have significantly transformed its capabilities and applications. These developments have brought renewed challenges in ensuring reliability, consistency, and regulatory compliance. Regulatory bodies recognized early on the importance of establishing guidelines to validate processes that would later become essential in AI-driven applications.

This chapter provides an overview of the existing literature relevant to the intersection of AI in healthcare, and regulatory governance. The review is organized into two main sections. The first section focuses on academic research and applied studies concerning the integration of AI within the domains of healthcare.


The second section shifts focus to the evolving regulatory landscape and governance surrounding AI technologies. This includes an overview of key policy documents, guidelines, and legislative milestones that shape the safe, ethical, and effective deployment of AI across sectors.

2.1. Artificial Intelligence in Healthcare

Wiens et al. (2019) present a comprehensive framework for responsibly translating machine learning (ML) applications into healthcare settings, addressing the persistent gap between technical innovation and clinical adoption. Despite growing enthusiasm for ML in medicine, real-world deployment remains limited due to fragmented development practices and insufficient stakeholder engagement. The authors emphasize the importance of interdisciplinary collaboration from problem formulation through to deployment, advocating early involvement of clinicians, administrators, patients, and other relevant actors to ensure clinical relevance.

It is emphasized that successful model development requires rigorous data scrutiny, assessing how and why data were collected, harmonizing datasets across systems, and ensuring representativeness to avoid reinforcing structural biases. Ethical considerations are also central, as healthcare data often reflect social inequalities that can embed unfairness into algorithms. The authors recommend involving ethicists and social

scientists to help mitigate these risks early in the design process, and stress the importance of robust evaluation practices to prevent technical issues like label leakage and to ensure the model performs reliably in its intended environment. Furthermore, the paper offers a structured roadmap aimed at promoting safe, equitable, and meaningful ML implementation in healthcare as shown in Figure 1 (Wiens et al., 2019).



1. Figure: Schematic representation of the roadmap for deploying effective ML systems in healthcare
 Extracted from Wiens, et al., 2019 (Wiens et al., 2019).

Similarly to Wiens et al. (2019), Saenz et al. (2024) present a framework outlined by a comprehensive case study for the responsible integration of AI in healthcare institutions. Recognizing the complexity of deploying AI technologies in clinical environments, the authors propose a set of guiding principles emphasizing fairness, robustness, privacy, transparency, explainability, and accountability. The framework is built around continuous monitoring and adaptation, acknowledging the challenges posed by data distribution shifts, evolving medical contexts, and demographic changes that affect model performance over time (Saenz et al., 2024).

A cross-functional team of experts, including legal, clinical, equity, and informatics professionals, collaborated to develop institution-wide guidelines. The framework prioritizes equity by requiring diverse training datasets and incorporating an equity-focused evaluation of model outcomes. Saenz et al. also emphasize the importance of assessing usability, scalability, and provider feedback to ensure real-world practicality. Transparent communication about an AI system's FDA approval status is considered essential for building trust and regulatory compliance (Saenz et al., 2024).


Importantly, the authors advocate for a risk-based oversight approach, with high-risk systems subject to stricter scrutiny and monitoring. This includes establishing an AI operations committee responsible for documentation, human oversight in clinical applications, and intensified post-deployment evaluation. A generative AI use case illustrates the operational and monitoring demands of emerging technologies. The paper concludes that sustained vendor collaboration and ongoing system reassessment are vital to ensuring the safety, performance, and ethical deployment of AI tools in healthcare (Saenz et al., 2024).

2.2. Regulatory Frameworks, Guidelines, and Key Milestones

In this thesis, the geographic focus is directed toward Europe and the United States in relation to regulatory frameworks and key milestones in the field of AI. This decision was guided by both practical and conceptual considerations. These two regions represent some of the most advanced and influential regulatory environments in the world, particularly in the domains of healthcare and technologies such as AI. Regulatory bodies such as the U.S. Food and Drug Administration (FDA) and European Medicines Agency (EMA), along with legislative instruments such as the EU AI Act, are widely recognized for setting benchmarks in AI regulation.

Additionally, the United States and Europe offer contrasting regulatory philosophies that provide a valuable foundation for comparative analysis. While the United States tends to prioritize innovation, Europe is often characterized by a more precautionary and human-centric approach. Analyzing these differing paradigms facilitates a deeper understanding of how governance culture may shape the integration of AI in the life sciences, particularly in balancing innovation, compliance, and patient-centeredness. The decision to focus on these regions was further supported by the accessibility of comprehensive policy documents in English.

This section provides an overview of the relevant documents and key milestones, as illustrated in Figure 2, which also indicates the year of issuance for each.



2. Figure: Timeline of key guidelines and milestones related to AI implementation in life sciences.

One of the earliest regulatory frameworks related to AI was the FDA's Process Validation Guidance, issued in 2011. The guidance applies to various categories of drugs, including human and veterinary drugs, biological and biotechnology products, finished products and active pharmaceutical ingredients (APIs), and drug components of combination (drug-device) products. Additionally, the guidance is relevant to processes that use automated equipment, highlighting the increasing role of automation in pharmaceutical manufacturing. While the guidance does not establish legally enforceable responsibilities, compliance with it is crucial for ensuring quality and regulatory approval (FDA, 2011).

The FDA defines process validation as the collection and evaluation of data throughout a product's lifecycle to establish scientific evidence that a process consistently produces quality products. The guidance introduces a three-stage approach to process validation, where stage 1 is the process design, stage 2 is process qualification, and stage 3 is the continued process verification. Furthermore, it emphasizes several best practices for process validation. For example, an integrated team approach, involving cross-functional collaboration among quality, manufacturing, and regulatory teams. A lifecycle approach, ensuring continuous monitoring and improvement with risk-based decision making, and viewing criticality as a spectrum rather than a binary concept. Validated analytical methods, though not always required during early development, should be scientifically sound, specific, and accurate for reliable results. Automated and electronic equipment

validation, ensuring that mechanical and software-based systems used in drug manufacturing are calibrated, inspected, and checked regularly to ensure performance. It is also highlighted that process validation is a critical component of quality assurance, ensuring that each step in the manufacturing process is controlled and produces products that meet required quality standards. The guidance identifies key principles to be followed to meet the before mentioned quality standards. First, quality, safety, and efficacy must be built into the product rather than relying solely on final testing. Secondly, quality assurance cannot be achieved by inspections alone, processes must be designed to ensure consistency. Finally, each stage of the manufacturing must be controlled to maintain quality attributes (FDA, 2011). Importantly, these principles reflect a patient-focused approach: ensuring that everything a patient receives is safe, effective, and consistent in quality is fundamental to protecting public health.

While this guidance was not explicitly designed for AI, it laid the groundwork for validating complex, data-driven processes, an approach that has since become fundamental in AI model qualification, validation, and maintenance. Its lifecycle approach, focus on data-driven decision-making, and integration of automation are directly relevant to the validation of AI algorithms, AI models, and digital manufacturing systems in today's regulatory landscape.

Similarly to the FDA, the European Medicine Agency (EMA) also published their guidance on process validation for finished products first with a draft version in 2012, followed by a final version issued in 2016. The EMA also put a great emphasis on the lifecycle approach for process validation, reinforcing that it is not a one-off activity, but an ongoing commitment for continued assurance that the process remains in a state of control during routine commercial production. Additionally, the guideline recognizes the applicability of continuous process verification (CPV) to both traditional and enhanced development approaches, enabling greater flexibility and responsiveness in modern manufacturing (EMA, 2016).

The EMA guidance specifically addresses the validation of the manufacturing process, which corresponds to Stage 2 of the product lifecycle. This aligns with the FDA's emphasis on linking development, validation, and continuous control as part of the holistic quality system. In summary, the EMA's Process Validation Guidelines promote a flexible, risk-based, and lifecycle-driven approach to manufacturing validation. By endorsing enhanced

process understanding, continuous verification, and site-specific assurance, the EMA sets the stage for incorporating modern technologies such as AI, automation, and advanced analytics within validated pharmaceutical processes.

In conclusion, both the FDA and EMA emphasize that quality must be built into pharmaceutical products from the outset, rather than relying solely on end-product testing. According to the FDA, quality, safety, and efficacy are integral to the manufacturing process itself and must be consistently maintained at every stage (FDA, 2011). Similarly, the EMA highlights that process validation is a continuous activity across the product lifecycle, supporting a model in which the manufacturing process is actively monitored and controlled to ensure reliable, reproducible outcomes (EMA, 2016). Underpinning both approaches is a strong patient focus: by embedding quality into the process, manufacturers ensure that what patients receive are not only safe and effective, but also consistently meet expectations.

This lifecycle-based, risk-driven mindset is highly applicable to the validation and governance of AI algorithms and models used in healthcare. Just as drugs must be validated throughout development and commercial use, AI models must also be monitored, tested, and updated continuously to ensure performance, accuracy, and safety. Whether it's a predictive model used in diagnostics or an algorithm managing pharmaceutical production, the ultimate goal remains the same: to protect patient safety, support therapeutic outcomes, and uphold public trust.

In 2016, the European Directorate for the Quality of Medicines & HealthCare (EDQM), a directorate of the European Council, issued the European Pharmacopoeia which sets common quality standards for the development, production, and control of medicines and their components, ensuring the safety and consistency of pharmaceuticals used in Europe (EDQM, 2016). While traditionally focused on the physicochemical and microbiological testing of pharmaceutical substances and finished products, the European Pharmacopoeia acknowledges the growing role of chemometric techniques in the analysis and interpretation of complex datasets and increasingly incorporates advanced data-driven methodologies such as chemometrics, reflecting the industry's shift toward digitalization and analytics. Specifically, chapter 5.21, titled as Chemometric methods applied to analytical data, provides guidance for data-driven modelling, a methodology closely aligned with how AI systems are trained and validated, and it also outlines expectations


for the design, implementation, and validation of these methods in pharmaceutical contexts. In the guideline chemometrics is defined as “The chemical discipline that uses mathematical and statistical methods to design or select optimal measurement procedures and experiments, and to provide maximum chemical information by analysing chemical data.” (EDQM, 2016). Chemometric methods are primarily multivariate, data-driven modelling techniques, and are empirical in nature. While powerful, these methods are not intended to replace established scientific theory when available, they complement traditional approaches by extracting meaningful insights from large and complex data.

By embedding standards for multivariate data analysis, model validation, and ongoing performance monitoring, the European Pharmacopoeia provides a structured, scientifically endorsed framework that can be extended to AI and machine learning systems used in manufacturing control, analytical method development, and quality assurance. Chemometric models must demonstrate predictive reliability, specificity, and robustness, particularly when they influence decision-making processes that may impact patient safety or regulatory compliance. In this way, the European Pharmacopoeia acts not only as a scientific benchmark but also as a regulatory anchor, ensuring that AI technologies are integrated in a manner that maintains product integrity, data reliability, and patient trust.

In 2018, the FDA approved IDx-DR, the first AI-based medical device that detects diabetic retinopathy without a clinician’s interpretation. Designed for use in primary care settings, it analyzes retinal images using an AI algorithm and provides an immediate screening result. This device aims to address the issue of under-screening among diabetic patients, nearly half of whom do not receive regular eye exams (FDA, 2018). The approval of IDx-DR marks a significant milestone in the evolving regulatory landscape of AI, by showing that autonomous AI systems can meet the stringent standards of safety, efficacy, and reliability required for medical use—especially in high-stakes, patient-facing applications.

The IDx-DR approval triggered deeper discussions across regulatory bodies about how to handle algorithm change protocols, validate so-called black-box models, and ensure transparency and explainability in AI diagnostics. A good example of this is the Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) - Discussion Paper and Request for Feedback issued by the FDA in 2019. The paper proposes a forward-looking regulatory

framework tailored to the unique challenges and opportunities presented by AI and machine learning (ML) technologies in Software as a Medical Device (SaMD). These tools offer transformative potential in healthcare by continuously learning from real-world data and improving performance over time. However, traditional medical device regulations were not designed for such adaptive and evolving systems. To address this, the FDA suggests a Total Product Lifecycle (TPLC) regulatory approach that supports continuous learning while ensuring safety and effectiveness. Under this model, certain significant software modifications, such as those affecting performance, intended use, or core algorithms would still require premarket review. However, predefined algorithm change protocols and performance objectives would allow for controlled, ongoing improvements without the need for repeated submissions. This TPLC approach emphasizes real-world performance monitoring, transparent validation processes, and a flexible oversight model that promotes innovation without compromising patient safety (FDA, 2019).



3. Figure: Workflow of AI models based on the good practices considering AI/ML as software as a medical device.

Extracted from the discussion paper "Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD)" (FDA, 2019).

The document also emphasizes the need for a risk-based approach that accommodates the unique characteristics of AI/ML systems. This risk-based approach categorizes SaMD based on two key factors: (1) the significance of the information provided by the AI system

in the context of clinical decision-making, ranging from informing clinical management to diagnosing or treating a condition, and (2) the seriousness of the healthcare condition or situation being addressed, such as whether it is critical, serious, or non-serious. These two dimensions are then used to place a given SaMD into one of four risk categories, from Category I (lowest risk) to Category IV (highest risk). This stratification not only supports proportional regulatory oversight but also helps manufacturers and regulators better align validation and lifecycle controls to the real-world clinical impact of the AI system (FDA, 2019).

State of healthcare situation or condition	Significance of information provided by SaMD to healthcare decision		
	Treat or diagnose	Drive clinical management	Inform clinical management
Critical	IV	III	II
Serious	III	II	I
Non-serious	II	I	I

1. Table: FDA's risk categorization regarding SaMD.

Adapted from the discussion paper "Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD)" (FDA, 2019).

Additionally, the paper underlines the importance of establishing good machine learning practices to ensure the quality and reliability of AI/ML algorithms. These practices would cover aspects such as data management, algorithm training, and performance evaluation. It also introduces the concept of a predetermined change control plan, which would allow manufacturers to pre-specify certain modifications to AI/ML-based SaMD. This plan would detail the types of changes anticipated post-market and the methods for implementing them, facilitating continuous learning while ensuring patient safety (FDA, 2019).


Recognizing the evolving nature of AI/ML technologies, the FDA seeks input from various stakeholders, including industry, academia, and patient groups, to refine its regulatory approach and ensure it remains relevant and effective. Ultimately, the FDA envisions a regulatory environment that enables AI/ML-based SaMD to evolve responsibly and reliably.

Similarly to the FDA, in the document Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle (EMA, 2023) the EMA outlines its views on integrating artificial intelligence and machine learning throughout the lifecycle of medicinal products and aims to stimulate dialogue with key stakeholders. The main purpose of the reflection paper is to explore opportunities, challenges, and regulatory expectations for the use of AI/ML in the development, authorization, and post-authorization phases of medicine. As a core principle, the EMA emphasizes that AI systems used in the medicinal product lifecycle must be ethically developed, transparent, robust, and fit-for-purpose to ensure that the safety of the patients remains the priority (EMA, 2023).

As for the European Union, in the summer of 2024 the EU AI Act came into force. This document aims to establish a uniform legal framework across the Union to regulate the development, marketing, deployment, and use of artificial intelligence systems, ensuring they align with Union values like health, safety, fundamental rights, democracy, the rule of law, and environmental protection (EU, 2024, Recital 1). It also aligns with the FDA's and EMA's guidelines by promoting human-centric and trustworthy artificial intelligence systems. The regulation aims to support innovation while also safeguarding against potential harms and allowing free cross-border movement of artificial intelligence products and services within the internal market (EU, 2024, Recital 1). Although the regulation acknowledges the potential, advantages and social benefits AI systems can deliver, particularly in fields like healthcare (EU, 2024, Recital 4), it never loses sight of its core principle that AI must ultimately serve as a tool designed to enhance human well-being, ensuring that human interests remain at the center of technological advancement and serve society (EU, 2024, Recital 6). As part of these core principles, the regulation clarifies that it does not interfere with existing EU data protection laws, ensuring alignment with frameworks like the GDPR (EU, 2024, Recital 10).

To foster innovation responsibly, the regulation underlines the need to respect scientific freedom and support research and development, ensuring that regulatory measures do not stifle scientific progress (EU, 2024, Recital 25). Similarly to the FDA's approach (FDA, 2019) described previously, a risk-based approach forms the backbone of the regulatory framework, tailoring obligations according to the level of risk associated with AI systems (EU, 2024, Recital 26) as shown in Figure 4. Special attention is given to high-risk AI systems, for which common, non-discriminatory rules are laid down (EU,

2024, Recital 7). The protection of public interests, particularly health, safety, and fundamental rights, is a central objective (EU, 2024, Recital 8).



4. Figure: Schema of the risk-based classification of AI Systems based on the EI AI Act, including examples and requirements.
Adapted from European Society of Radiology (European Society of Radiology, n.d.).

The EU AI Act classifies AI systems into four distinct risk categories, each with corresponding regulatory obligations. Unacceptable risk systems, which contravene fundamental EU rights and values, are outright prohibited. High-risk systems, those affecting health, safety, or fundamental rights, are subject to rigorous measures, including conformity assessments and post-market monitoring. Limited-risk applications, such as chatbots or AI-generated content, require clear information and transparency to prevent deception or manipulation. Low or minimal risk systems, such as spam filters or recommendation engines, are exempt from specific regulatory requirements (EU, 2024).

Finally, while the Act introduces binding obligations for trustworthy AI systems. These guidelines outline seven ethical principles, such as human agency and oversight, technical robustness and safety, privacy and data governance, transparency, diversity and fairness, societal and environmental well-being, and accountability, which continue to inform the vision of trustworthy, human-centric AI in Europe (EU, 2024, Recital 27).

In the United States, Executive Order 14110, titled as “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence” was also issued to establish a comprehensive framework to manage the development, deployment, and oversight of artificial intelligence. With this executive order the Biden Administration aimed to capture the opportunities and benefits of artificial intelligence while also making sure appropriate

measures are in place to mitigate risks possibly posed to national security, civil rights, the economy, and public safety (US Exec. Order No. 14110, 2023).

The executive order approaches AI from different perspectives, such as safety and security, privacy, civil rights, and equity just to mention a few. These aspects are particularly important in critical sectors like healthcare, where mistakes could harm patients. Therefore, the executive order required AI developers to share test results with the U.S. government before public release, mandated the development of rigorous standards by the National Institute of Standards and Technology for testing, evaluation, and mitigation of AI systems, and directed the Department of Homeland Security to assess the use of AI in critical infrastructure sectors. It also emphasized the need for privacy-enhancing technologies (PETs) to protect patient confidentiality and directed agencies to prevent algorithmic discrimination in hiring, housing, and the criminal justice system by ordering training and guidance for those using AI for decision making (Exec. Order No. 14110, 2023).


In section 5.2 (e) the executive order mandated the Secretary of Health and Human Services (HHS) to prioritize grant making to support the responsible development of AI tools. As part of this it aimed to accelerate the development of personalized medical care, such as immune response profiling, ensure improved data quality, and provide sufficient resources for the National Institutes of Health Artificial Intelligence/Machine Learning Consortium to Advance Health Equity and Researcher Diversity (AIM-AHEAD) program, which uses AI to address health equity and diversity in research (Exec. Order No. 14110, 2023).

As part of the revocation of several executive orders, the new Trump Administration withdrew the Executive Order 14110 in the beginning of 2025 (The White House, 2025). However, the full outcome of revoking the order remains yet unknown, the elimination of essential guardrails it established raises several pressing concerns.

Without these guardrails, AI systems could be deployed without proper accountability, and could leave communities vulnerable to unregulated decision making that can introduce systemic bias. Furthermore, the revocation of the order opens the door to unverified information, manipulation, and loss of public trust in AI systems. In healthcare particularly, the elimination of these directives might slow down innovation while posing harm to patients. In summary, the Executive Order 14110 was more than a regulatory

document, it aimed to ensure that emerging AI technologies serve the public good. By removing safeguards around AI could open the door to significant risks across multiple dimensions.

Complementing these broader concerns, the U.S. Food and Drug Administration's draft guidance, *Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products* (FDA, 2025), addresses the integration of AI within regulatory frameworks specific to drug development and evaluation. The document introduces a risk-based credibility assessment framework to evaluate AI model outputs used in support of regulatory decisions concerning drug safety, effectiveness, and quality. This includes a seven-step process for establishing credibility, guided by the specific context of use and associated risk level (FDA, 2025) as shown in Figure 5.



5. Figure: Schema describing the seven-step risk-based credibility assessment framework for AI in drug development.

Adapted from the FDA "Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (Draft Guidance)." (FDA, 2025).

Furthermore, the guidance emphasizes on lifecycle maintenance, which acknowledges that AI model performance may shift over time or across deployment contexts. To address this, the FDA advocates for proactive monitoring strategies, performance metrics, and triggers for model reassessment. These provisions aim to ensure that AI models remain fit for purpose throughout their use, particularly in high-risk scenarios. In doing so, the guidance provides a practical and forward-looking framework for balancing innovation with patient safety through structured and risk-sensitive governance (FDA, 2025).

In the context of pharmaceutical regulation, globalization necessitates harmonization and mutual recognition of standards. Countries exporting medicinal products to the United States of America must adhere to the U.S. Food and Drug Administration (FDA) regulations, while the European Union imposes its own set of requirements through the European Medicines Agency (EMA). These reciprocal expectations create an interdependent regulatory environment. As artificial intelligence (AI) becomes increasingly integrated into pharmaceutical research and manufacturing, there is a clear opportunity, and arguably an urgent need to align regulatory frameworks. Establishing standards for AI applications in the life sciences would not only streamline international compliance but also elevate the overall quality and safety of innovation (European Commission, 2024).

A compelling reason to act pre-emptively is illustrated by past regulatory shortcomings. The nitrosamine contamination crisis, which affected a range of widely used medicines, serves as a cautionary example. Regulatory responses came only after the issue had emerged, highlighting the risks of reactive governance (EMA, 2020). In the case of AI, where the pace of technological change is even faster, a post-mortem approach to regulation would be particularly dangerous. Effective oversight must be anticipatory, embedding risk mitigation strategies at the earliest stages of technological integration.

The FDA's recent draft guidance on the use of AI in drug manufacturing exemplifies the current limitations in regulatory comprehension of advanced technologies. The document mentions the term "algorithm" only once, and solely in the context of defining AI (FDA, 2023). Conversely, terms such as "model" and "data" appear frequently, indicating a focus on traditional data-driven processes rather than the broader conceptual and ethical implications of AI. This limited lexical emphasis may reflect a superficial understanding of AI's complexity, underscoring the need for regulatory bodies to actively involve domain experts in AI and machine learning. Without such expertise, guidance risks being overly simplistic and potentially inadequate. This situation poses a broader question about regulatory readiness in the face of rapidly evolving technologies. If regulators do not fully grasp the systems they are tasked with overseeing, the resulting frameworks may be insufficient, leaving public health and safety vulnerable. Therefore, proactive and informed regulation is imperative.

At this critical juncture, Europe is well-positioned to take a leadership role. The recently adopted EU AI Act (European Union, 2024) marks a pioneering effort to create a comprehensive legal framework for AI. Rather than serving solely as a regulatory instrument, it should be viewed as a visionary model that integrates ethical, scientific, and societal considerations. By setting high standards for safety, transparency, and accountability, the EU can establish itself as a global leader in the responsible use of AI in healthcare and beyond (European Commission, 2024).

3. Methodology

The central research question of this thesis is whether existing regulations and guidelines governing AI are genuinely patient-centered, or whether they are primarily bureaucratically driven. While this question is inherently qualitative and difficult to quantify directly, a proxy method has been developed to assess the patient-centeredness of these regulatory documents through a targeted content analysis.

To this end, the following key regulatory and policy documents were selected for analysis:

- Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD): Discussion paper and request for feedback. (FDA, 2019),
- Executive Order No. 14110: Safe, secure, and trustworthy development and use of artificial intelligence. (Exec. Order No. 14110, 2023),
- The European Union Artificial Intelligence Act (European Union, 2024),
- Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle (EMA, 2024),
- Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (Draft Guidance) (FDA, 2025).

These documents were selected based on their relevance, recency, and jurisdictional coverage. Collectively, they represent the most comprehensive and current efforts to regulate the application of AI systems in the United States and Europe, the two primary geographic regions of focus in this thesis. Each document addresses the governance of AI technologies either in general terms or with specific reference to medical and pharmaceutical applications.

Although Executive Order 14110 was officially revoked in January 2025, it is included in this analysis due to its significant role in shaping the United States policy discourse on AI governance during its period of effect. Its inclusion provides critical historical context for understanding regulatory shifts and priorities in the United States.

The analytical approach involved identifying and quantifying the occurrence of specific keywords indicative of either a patient-centered or bureaucratic orientation. These keywords were derived from themes commonly associated with patient advocacy, public health, and regulatory compliance.

Patient-centered Keywords:

Patient, Fairness, Bias, Equity, Safety, Security, Transparency, Disease, Rare disease, Adverse effect

Bureaucratically Driven Keywords:

Compliance, Reporting, Submission, Submit, Documentation, Timeline, Regulation, Regulatory, Quality Affairs, Regulatory bodies

The selection of keywords for this analysis was grounded in both conceptual relevance and alignment with commonly accepted themes in health policy, and biomedical ethics. The two categories, Patient-centered and Bureaucratically Driven, were designed to reflect distinct orientations within regulatory documents.

Patient-centered keywords were chosen based on terminology frequently associated with patient advocacy, clinical ethics, and public health priorities. Terms such as patient, safety, equity, and fairness are central to the goals of ensuring that technologies applied in healthcare serve the interests of patients. Words like bias and rare disease indicate efforts to protect vulnerable or underserved populations, while transparency and adverse effects are indicators of patient-focused risk communication and accountability.

Bureaucratically driven keywords were selected to capture the language typically used in administrative, procedural, and regulatory contexts. Terms such as compliance, submission, and documentation reflect a focus on procedural adherence and institutional requirements, frequently used in regulatory literature and guidance documents. Regulatory bodies, quality affairs, and timelines emphasize institutional oversight and formal processes that, while important for governance, may not directly prioritize patient outcomes.

To conduct the analysis, keyword occurrences were identified using the CTRL+F search function within each document. The number of instances for each keyword was then recorded in structured tables. The total counts for each category, patient-centered and

bureaucratically driven, were computed using the SUM() function in Microsoft Excel. Given the variation in document lengths, normalization was applied by calculating the frequency of relevant keywords per page, subsequently the total number of words for each category were divided by the number of pages in each document to obtain normalized scores. This resulted in a comparative score for each document, enabling a relative assessment of its orientation along the patient-centered versus bureaucratic spectrum.

To enhance the robustness of the findings, a two-phase analytical approach was employed. The first phase consisted of an exploratory analysis of each full document, capturing keyword frequencies across all sections, including front and back matter such as cover pages, tables of contents, glossaries, references, annexes, and appendices. Building on insights from this initial scan, the second phase aimed to provide a more nuanced and focused understanding by restricting the analysis to the core content of each document. In this refined approach, non-substantive sections were excluded from the page count used for normalization, and keyword searches were confined to the main body text, omitting footnotes and peripheral material.

Based on the content-focused analysis, a benchmark was established using the sum of the average and first quartile scores across documents as an acceptance criterion. If a guideline scores above this benchmark in terms of patient-centeredness, it is considered a positive indication of its alignment with patient-focused values. Conversely, if the bureaucratically driven score exceeds this benchmark, it may signal the need for critical review or revision of the framework to better align with practical and ethical healthcare priorities.

4. Results

This section presents the results obtained from the analysis, with findings reported separately for each regulatory document. The results are structured into two parts: the first based on a full-document analysis, and the second derived from a focused main-content analysis.

4.1. Results from Comprehensive Document Evaluation

The first document analyzed is the Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device: Discussion Paper and Request for Feedback (FDA, 2019). This document comprises 20 pages. Using the keyword-based approach described in the Methodology section, the following results were obtained:

Patient-centered		Bureaucratically Driven	
Word	Count	Word	Count
Patient	37	Compliance	1
Fairness	0	Reporting	5
Bias	0	Submission	21
Equity	0	Submit	15
Safety	17	Documentation	1
Security	0	Timeline	0
Transparency	13	Regulation	19
Disease	9	Regulatory	9
Rare disease	0	Quality Affairs	0
Adverse effect	0	Regulatory bodies	0

2. Table: Frequency of Keywords in Proposed Regulatory Framework for SaMD. ¹

Based on these figures, a normalized score of 3.80 was calculated for patient-centered keywords, while a score of 3.55 was obtained for bureaucratically driven keywords.

The next document analyzed is the Executive Order No. 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Exec. Order No. 14110, 2023). This document comprises 36 pages. Applying the keyword-based methodology, the following results were obtained:

¹ Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device: Discussion Paper and Request for Feedback (FDA, 2019).

Patient-centered	
Word	Count
Patient	7
Fairness	0
Bias	9
Equity	9
Safety	25
Security	126
Transparency	3
Disease	0
Rare disease	0
Adverse effect	0

Bureaucratically Driven	
Word	Count
Compliance	7
Reporting	10
Submission	0
Submit	19
Documentation	3
Timeline	0
Regulation	14
Regulatory	23
Quality Affairs	0
Regulatory bodies	0

3. Table: Frequency of Keywords in the Executive Order No. 14110. ²

The analysis resulted in a normalized score of 4.97 for patient-centered keywords, indicating a strong emphasis on patient-related considerations, whereas bureaucratically driven terms yielded a lower score of 2.11.

The third document included in the analysis is The European Union Artificial Intelligence Act (European Union, 2024), which spans 144 pages.

Patient-centered	
Word	Count
Patient	2
Fairness	5
Bias	25
Equity	0
Safety	113
Security	101
Transparency	41
Disease	2
Rare disease	0
Adverse effect	3

Bureaucratically Driven	
Word	Count
Compliance	150
Reporting	20
Submission	2
Submit	30
Documentation	124
Timeline	2
Regulation	770
Regulatory	90
Quality Affairs	0
Regulatory bodies	0

4. Table: Frequency of Patient-centered and Bureaucratically Driven Keywords in the EU AI Act. ³

The analysis produced a normalized score of 2.03 for patient-centered terms and 8.25 for bureaucratically driven language, suggesting a focus on regulatory and administrative aspects over direct patient-centered concerns.

² Executive Order No. 14110: Safe, secure, and trustworthy development and use of artificial intelligence (Exec. Order No. 14110, 2023).

³ The European Union Artificial Intelligence Act (European Union, 2024).

The fourth document examined is the Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle (EMA, 2024), consisting of 18 pages.

Patient-centered		Bureaucratically Driven	
Word	Count	Word	Count
Patient	20	Compliance	8
Fairness	6	Reporting	1
Bias	8	Submission	1
Equity	0	Submit	0
Safety	12	Documentation	3
Security	2	Timeline	0
Transparency	5	Regulation	31
Disease	1	Regulatory	49
Rare disease	0	Quality Affairs	0
Adverse effect	0	Regulatory bodies	1

5. Table: Frequency of Keywords in the EMA Reflection paper ⁴

Based on the keyword analysis, the document yielded a normalized score of 3.00 for patient-centered terms and 5.22 for bureaucratically driven terms, indicating a stronger orientation toward regulatory and procedural considerations than toward patient-centered priorities.

The final document analyzed is the Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products: Guidance for Industry and Other Interested Parties (Draft Guidance) (FDA, 2025), which comprises 23 pages.

Patient-centered		Bureaucratically Driven	
Word	Count	Word	Count
Patient	18	Compliance	4
Fairness	0	Reporting	2
Bias	7	Submission	12
Equity	1	Submit	10
Safety	10	Documentation	2
Security	1	Timeline	0
Transparency	1	Regulation	7
Disease	2	Regulatory	38
Rare disease	0	Quality Affairs	0
Adverse effect	0	Regulatory bodies	0

6. Table: Frequency of Keywords in the FDA Considerations for the use of AI ⁵


⁴ Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle (EMA, 2024).

⁵ Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (FDA, 2025).

The keyword frequency analysis produced a normalized score of 1.74 for patient-centered terminology and 3.26 for bureaucratically driven language. These findings suggest that the document places greater emphasis on regulatory and procedural aspects than on patient-focused considerations.


To complement the tabulated results, two bubble charts were created in Tableau to visually summarize the frequency of keywords associated with patient-centered and bureaucratically driven language across the analyzed regulatory documents. These visualizations serve to highlight the relative emphasis placed on different thematic domains within the regulatory texts.

The first chart focuses on patient-centered terminology, where each bubble represents a keyword, with its size proportional to its total frequency across the documents. Notably, “Security” and “Safety” emerged as dominant terms, indicating that regulatory language often acknowledges the importance of safe and secure AI systems. However, terms like “Fairness,” “Equity,” and “Bias” appeared significantly less frequently, suggesting that issues of inclusion and algorithmic justice may be underarticulated in current guidance documents.



6. Figure: Total number of patient-centered keywords

The second chart displays bureaucratically driven keywords across the documents included in the analysis. Among these, “Regulation” and “Regulatory” were the most frequently cited, reflecting a strong emphasis on procedural control and legal oversight, hallmarks of bureaucratic governance. However, a closer inspection of the data reveals that this dominance is largely driven by the EU AI Act, which accounts for a substantial proportion of keyword occurrences. Specifically, “Regulation” appears 770 times in the EU AI Act out of a total of 841 mentions across all documents. Similarly, “Compliance” appears 150 times out of 170, and “Documentation” occurs 124 times out of 133 in the same document. This suggests that the visual prominence of these terms in the bubble chart may overstate their presence across the broader corpus, and instead reflects the density and legalistic nature of the EU AI Act in particular. While the chart effectively communicates the bureaucratic emphasis of that document, it also highlights the need for caution when generalizing from aggregated keyword counts without considering the distribution across sources.




7. Figure: Total number of bureaucratically driven keywords

Finally, to facilitate cross-document comparison, the normalized scores for both keyword categories are summarized in Table 7 and Figure 8.

Document Title	Number of Pages	Patient-Centered Score	Bureaucratically Driven Score
Proposed Regulatory Framework for SaMD (2019)	20	3.80	3.55
Executive Order No. 14110 (2023)	36	4.97	2.11
EU AI Act (2024)	144	2.03	8.25
EMA Reflection Paper (2024)	18	3.00	5.22
FDA Considerations for the use of AI (2025)	23	1.74	3.26

7. Table: Comparative summary of normalized keyword scores

As shown, the Executive Order No. 14110 demonstrates the highest emphasis on patient-centered considerations, whereas the EU AI Act displays the strongest bureaucratic orientation.



8. Figure: Summary of normalized keyword scores

4.2. Results from Focused Main-Content Evaluation

The first document analyzed is the Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device: Discussion Paper and Request for Feedback (FDA, 2019). This document comprises 13.5 pages of main content. Using the keyword-based approach described in the Methodology section, the following results were obtained:

Patient-centered	
Word	Count
Patient	37
Fairness	0
Bias	0
Equity	0
Safety	17
Security	0
Transparency	11
Disease	9
Rare disease	0
Adverse effect	0

Bureaucratically Driven	
Word	Count
Compliance	0
Reporting	4
Submission	13
Submit	8
Documentation	1
Timeline	0
Regulation	6
Regulatory	9
Quality Affairs	0
Regulatory bodies	0

8. Table: Frequency of Keywords in the main content of the Proposed Regulatory Framework for SaMD.⁶

The adjusted analysis resulted in a patient-centered keyword frequency of 5.48, while terms aligned with bureaucratic language registered a frequency of 3.04.

The second document analyzed, Executive Order No. 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Exec. Order No. 14110, 2023), contains 36 pages of substantive content. Since this document does not include extensive supplementary sections such as a cover, glossary, or annexes, the number of pages and keyword results are identical to those reported in the full-document analysis. Therefore, the previous scores of 4.97 for patient-centered keywords and 2.11 for bureaucratically driven terms remain unchanged.

The third document in this focused assessment is The European Union Artificial Intelligence Act (European Union, 2024), with 123 pages considered for analysis after omitting ancillary sections.

The keyword-based analysis yielded a normalized score of 1.82 for patient-centered terms and 8.43 for bureaucratically driven language, indicating a strong focus on administrative and regulatory aspects.

⁶ Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device: Discussion Paper and Request for Feedback (FDA, 2019).

Patient-centered	
Word	Count
Patient	1
Fairness	5
Bias	24
Equity	0
Safety	105
Security	45
Transparency	39
Disease	2
Rare disease	0
Adverse effect	3

Bureaucratically Driven	
Word	Count
Compliance	143
Reporting	20
Submission	0
Submit	24
Documentation	98
Timeline	2
Regulation	660
Regulatory	90
Quality Affairs	0
Regulatory bodies	0

9. Table: Frequency of Keywords in the main content of the EU AI Act. ⁷

It is important to note that the high frequency of the word “regulation” (660 instances) is primarily driven by the document’s repeated internal references to itself using the phrase “this regulation”, which alone accounts for 477 of those occurrences. Similarly, the term “regulatory” appears 90 times, of which 74 references are specifically related to the concept of a regulatory sandbox. These patterns significantly inflate the score for bureaucratically driven language.

The fourth document evaluated is the Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle (EMA, 2024), comprising 10.25 pages of core content.

Patient-centered	
Word	Count
Patient	19
Fairness	6
Bias	8
Equity	0
Safety	11
Security	2
Transparency	4
Disease	1
Rare disease	0
Adverse effect	0

Bureaucratically Driven	
Word	Count
Compliance	5
Reporting	0
Submission	1
Submit	0
Documentation	3
Timeline	0
Regulation	6
Regulatory	40
Quality Affairs	0
Regulatory bodies	1

10. Table: Frequency of Keywords in the main content of the EMA Reflection Paper. ⁸

⁷ The European Union Artificial Intelligence Act (European Union, 2024).

⁸ Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle (EMA, 2024).

In this analysis, patient-centered language reached a normalized frequency of 4.98, whereas bureaucratic terminology scored slightly higher at 5.46, revealing a slight procedural emphasis.

The final document examined is the Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products: Guidance for Industry and Other Interested Parties (Draft Guidance) (FDA, 2025), with 20 pages included in the main-content analysis.

Patient-centered		Bureaucratically Driven	
Word	Count	Word	Count
Patient	15	Compliance	1
Fairness	0	Reporting	1
Bias	3	Submission	6
Equity	0	Submit	6
Safety	6	Documentation	1
Security	0	Timeline	0
Transparency	1	Regulation	2
Disease	2	Regulatory	19
Rare disease	0	Quality Affairs	0
Adverse effect	0	Regulatory bodies	0


11. Table: Frequency of Keywords in the main content of the FDA Considerations for the use of AI.⁹

Keyword counts revealed a patient-centered score of 1.35, while bureaucratic language reached 1.80.

As in the previous analysis, visualizations were created to highlight the relative emphasis placed on different thematic domains within the regulatory texts.


Notably, the only substantial decrease observed among the patient-centered keywords in the content-focused analysis pertains to the term “security.” The frequencies of all other keywords remained relatively consistent, indicating that their presence within the documents is likely contextually grounded in the main body of the text rather than being confined to peripheral sections.

⁹ Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (FDA, 2025).



9. Figure: Total number of patient-centered keywords in the main content,

Regarding the bureaucratically driven keywords, “Regulation” and “Regulatory” remained the most frequently cited terms in the content-focused analysis. Their continued prominence is primarily attributable to the EU AI Act, which contributes a significant share of the total keyword occurrences. However, in this refined analysis, a notable decrease was observed in the frequency of nearly all bureaucratic terms, indicating that narrowing the scope to the main content effectively filtered out peripheral references and provided a more focused view of the documents' substantive emphasis.




10. Figure: Total number of bureaucratically driven keywords in the main content.

The normalized scores for both keyword categories are presented in Table 12 and Figure 11. Overall, the content-based analysis provided a more nuanced understanding of the orientation of each document, with improved scores across most cases—except for the EU AI Act. This exception reinforces the impression formed during the initial exploratory analysis: the EU AI Act exhibits a stronger bureaucratic emphasis relative to the other documents. This finding contrasts with initial expectations, particularly considering Europe's traditionally strong focus on social welfare.

Document Title	Number of Pages	Patient-Centered Score	Bureaucratically Driven Score
Proposed Regulatory Framework for SaMD (2019)	13.5	5.48	3.04
Executive Order No. 14110 (2023)	36	4.97	2.11
EU AI Act (2024)	123	1.82	8.43
EMA Reflection Paper (2024)	10.25	4.98	5.46
FDA Considerations for the use of AI (2025)	20	1.35	1.80

12. Table: Comparative summary of normalized keyword scores based on the main content analysis



11. Figure: Summary of normalized keyword scores based on the main content analysis

4.3. Benchmark Development and Threshold Criteria

To establish a meaningful benchmark for evaluating the orientation of regulatory documents, the normalized scores derived from the main content analysis were utilized. Specifically, the mean and the first quartile were calculated for each category of

terminology. These two statistical measures were then summed to create a composite benchmark for each category. This approach strikes a balance between identifying documents that significantly exceed the average while still remaining grounded in the broader distribution of scores. Based on this method, the benchmark for patient-centered terminology was determined to be 5.54, whereas the benchmark for bureaucratically driven terminology was calculated as 6.28, as detailed in the table below. These benchmarks serve as preliminary thresholds or acceptance criteria, helping to distinguish documents that demonstrate a pronounced patient-oriented focus from those that may lean more heavily toward bureaucratic or administrative priorities.

Terminology	Average Score	First Quartile	Benchmark
Patient-centered	3.72	1.82	5.54
Bureaucratically Driven	4.17	2.11	6.28

13. Table: Summary of benchmark components and threshold values

5. Discussion

This section critically interprets the findings presented in the previous section and places them within the broader context of AI governance in the life sciences sector. The discussion is structured in two parts. The first part focuses on the results derived from the keyword-based evaluation of selected regulatory documents. The second part considers the wider implications of these findings, including reflections on regulatory approaches in the United States and the European Union, the clarity and applicability of guidance, and potential impacts on patient-centered care and innovation. Through this twofold structure, the chapter aims to bridge the gap between empirical results and broader theoretical and practical considerations.

5.1. Discussion of Results

Based on the results presented in the previous section, it can be observed that regulatory approaches in the United States tend to place a stronger emphasis on patient-centered considerations compared to those in the European Union, where the regulatory language appears more bureaucratically oriented. It is important to reiterate, however, that Executive Order 14110, despite achieving the highest score in terms of patient-centeredness among the documents analyzed, is no longer in force. Its inclusion in this study reflects its relevance during the time it was active and its influence on the broader discourse surrounding AI governance in healthcare.


These findings may reflect broader systemic differences between the regulatory cultures of United States and Europe. Prior to conducting the keyword analysis, it was anticipated that the European regulatory environment would reflect a strong emphasis on human well-being and societal balance, consistent with broader societal values. Additionally, Europe is often seen as approaching AI as a scientific discipline in its own right, rather than merely a technological tool. These underlying perspectives may help explain why Europe tends to emphasize precautionary principles and procedural rigor, while the United States appears to prioritize innovation, while maintaining a focus on patient outcomes. Whether these differing regulatory philosophies lead to measurable differences in patient outcomes or development of AI technologies remains an open question and a considerable direction for future empirical research. This contrast also

raises a broader consideration: will the European Union, with its more stringent regulatory framework, be perceived as a more trustworthy environment for the deployment of AI in healthcare, or will the United States emerge as a more attractive destination for investment and innovation due to its more flexible policies?

5.2. Broader Discussion

An important consideration relates to the clarity and specificity of regulatory guidance. If the language used in such documents is overly abstract or lacks concrete directives, the risk increases that stakeholders, such as healthcare providers, developers, and regulators, may rely on subjective interpretations. This can lead to inconsistent implementation and, more critically, may compromise patient safety and equity by introducing variability in how AI systems are evaluated, deployed, or monitored.

The relationship between the degree of abstraction in regulatory frameworks and the impact of AI applications on patient outcomes can be conceptualized using a risk matrix. In this model, the X-axis reflects the level of ambiguity or abstraction in regulatory guidance, while the Y-axis represents the potential severity of consequences arising from misinterpretation or inadequate implementation of such guidance.



12. Figure: Risk Matrix of Regulatory Abstraction and Patient Impact

The upper-right quadrant of the matrix denotes a high-risk zone, where highly abstract regulatory guidance intersects with high-impact AI applications, such as patient-facing systems. In such cases, vague directives may lead to inconsistent implementations and

heightened risks to patient safety and equity. Conversely, the lower-left quadrant represents a low-risk zone, where specific, actionable guidance governs low-impact applications, for example, internal workflow optimization tools, thus posing minimal risk to patients.

6. Limitations and Future Work

This section outlines the key limitations encountered during the course of this study and proposes directions for future research that could build upon and extend the findings presented.

6.1. Limitations

The study is subject to key limitations that should be acknowledged in relation to the methodology and scope of this study:

- **Lack of Contextual Analysis:** The content analysis conducted was based on the frequency of predefined keywords, without accounting for the specific context in which those terms appear. As a result, nuances in meaning or the intent behind certain language may not be fully captured.
- **Subjectivity in Keyword and Document Selection:** The selection of both the regulatory documents and the associated keywords was informed by the researcher's judgment. While guided by relevant literature and the research question, this introduces a degree of subjectivity that may influence the interpretation of findings.
- **Geographic Scope:** The analysis is limited to regulatory frameworks from Europe and the United States. Although these regions are highly influential in the development of AI policy, the exclusion of perspectives from other jurisdictions may limit the generalizability of the conclusions.

6.2. Future Work

Building on the findings of this research, several directions for future work are worth considering. One important area involves the analysis of regulatory warning letters issued by agencies such as the FDA. These documents may provide valuable insights into how concerns around data integrity, privacy, and risk assessment are operationalized in practice. Particularly, if such letters are issued in contexts closer to production or clinical application, rather than in early research phases, they may reveal how regulatory enforcement aligns with patient-centered priorities.

Another potential extension of this study involves expanding the geographic scope beyond the United States and the European Union. Including other regions would provide a more comprehensive understanding of global AI governance trends.

Moreover, the incorporation of expert interviews, such as with regulators, AI developers, or clinical practitioners, could enrich the analysis with qualitative insights that go beyond document-based content. These interviews could help to interpret the intent and impact of regulatory language and assess how it is perceived and applied in practice.

Finally, future research could benefit from adopting a more qualitative content analysis approach. While the keyword-based method used in this thesis offers a systematic and replicable way to explore thematic emphasis, qualitative methods could provide deeper insights into context, tone, and regulatory intent that are not captured by word frequency alone.

7. Conclusion

As artificial intelligence (AI) becomes increasingly embedded in the life sciences sector, the development of robust governance frameworks is essential to ensure not only compliance and performance, but also fairness and inclusivity. The findings of this thesis emphasize that AI governance mechanisms must go beyond technical specifications and regulatory compliance to also encompass ethical and social considerations, particularly in domains where AI outcomes directly influence patient health and safety.

To promote fairness, AI governance should mandate the use of diverse and representative datasets during model development and deployment. This mirrors the inclusivity standards expected in clinical trials and is essential to prevent the exclusion of underrepresented populations. Ensuring equitable model performance across demographic groups reduces the risk of algorithmic bias, which can otherwise compromise diagnostic accuracy and treatment efficacy.

The involvement of a broad range of stakeholders, including domain experts, regulators, developers, and patient representatives, is also essential in the formulation of regulatory frameworks. Such inclusive dialogue can ensure that guidelines are not only technically robust but also ethically grounded and responsive to real-world needs. As emphasized by Wiens et al. (2019), engaging stakeholders early in the development process, such as clinicians, patients, administrators, and even extended social networks like family and friends, can help prevent misguided problem formulation and lead to more contextually appropriate and effective solutions. This collaborative approach fosters public trust and supports the creation of more practical, implementable guidance (Wiens et al., 2019).

Furthermore, the implementation of lifecycle management strategies is vital to maintaining AI model integrity over time. As AI systems encounter new data inputs or are deployed across different clinical or operational settings, their performance may shift in unintended ways. Regulatory bodies have increasingly recognized this risk, emphasizing the need for ongoing oversight and model maintenance to address potential issues such as data drift and changes in model behavior over time. A lifecycle approach not only enables the early detection of such performance deviations but also facilitates timely interventions to uphold model reliability, safety, and regulatory compliance (FDA, 2025).

A risk-based approach should also be embedded within AI governance practices. Not all AI applications pose the same level of risk, therefore, regulatory scrutiny should be proportional to the potential impact of a given system on patient outcomes. High-risk applications, particularly those involved in diagnostic or therapeutic decisions, require more stringent oversight and validation procedures than lower-risk applications, such as administrative automation.

The benchmark established in this thesis serves as an indicator for assessing whether a regulatory document demonstrates a stronger emphasis on patient-centered considerations or leans more toward bureaucratic priorities. This metric can help evaluate the potential impact of such documents on patients and guide the development of more balanced, patient-oriented regulatory frameworks. Ultimately, the responsible governance of AI in healthcare must remain grounded in the commitment to protect human well-being and to ensure that technological progress serves the needs of society in a safe, ethical, and forward-looking manner.

Finally, the patient must remain at the center of all AI governance efforts. As AI technologies increasingly shape clinical workflows and therapeutic decisions, ensuring that these systems are aligned with the overarching goal of improving patient outcomes is imperative. Regulatory frameworks that prioritize patient well-being can guide the ethical and responsible integration of AI into healthcare, ultimately fostering innovation that serves both technological progress and societal good.

8. Ethical-Social Impact, Sustainability, and Diversity

As part of the research, important considerations related to ethical and social impact, diversity, and sustainability have emerged. This section seeks to provide insights into these dimensions, exploring how the study aligns with responsible research practices.

8.1. Ethical-Social Impact

The core objective of this thesis, to evaluate whether regulatory frameworks are truly patient-centered, directly aligns with ethical concerns around the human impact of AI-driven policies. By critically examining the language of regulatory documents, this research sheds light on the extent to which patient well-being and outcomes are prioritized in AI governance. This analysis promotes greater accountability in policymaking and contributes to the ethical deployment of AI technologies in clinical settings.

Importantly, the findings of this thesis may inform stakeholders, including policymakers, clinicians, and technologists, about the implications of bureaucratic bias in shaping guidelines. Highlighting such biases can help improve transparency and trust in AI systems and ensure that the deployment of such systems reflects the diverse needs of patient populations.

8.2. Sustainability

From a sustainability standpoint, this thesis employed a lightweight and resource-efficient approach to computational analysis. The methodology relied primarily on manual review, keyword searches, and spreadsheet-based quantitative analysis (e.g., Excel), without the need for large-scale machine learning models or cloud-based infrastructure. This choice aligns with low-energy, sustainable research practices, minimizing carbon emissions associated with computing.

8.3. Diversity

The analysis performed in this research relies solely on publicly available regulatory documents from major authorities such as the FDA, European Commission, and EMA. While these sources are authoritative and comprehensive, it is acknowledged that they represent a limited range of cultural, political, and healthcare perspectives, primarily from western, high-income settings. Consequently, the findings may not fully reflect regulatory considerations in low- and middle-income countries or in healthcare systems with different socio-political dynamics.

Moreover, the methodology based on keyword frequency does not capture the nuance of culturally specific terminology or implicit biases that may influence regulatory language. Future research could expand the data scope to include documents from a wider array of jurisdictions and languages, thus improving representativeness and inclusiveness.

References

U.S. Food and Drug Administration. (2011, January). *Guidance for industry: Process validation – General principles and practices (Revision 1)*. <https://www.fda.gov/media/71021/download>

European Medicines Agency. (2016, November 21). *Guideline on process validation for finished products: Information and data to be provided in regulatory submissions* https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-process-validation-finished-products-information-data-be-provided-regulatory-submissions_en.pdf

European Directorate for the Quality of Medicines & HealthCare (EDQM). (2016). *European Pharmacopoeia (9th ed.)*.

U.S. Food and Drug Administration. (2018, April 11). *FDA permits marketing of artificial intelligence-based device to detect certain diabetes-related eye problems*. U.S. Department of Health and Human Services. <https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-artificial-intelligence-based-device-detect-certain-diabetes-related-eye>

U.S. Food and Drug Administration. (2019). *Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD): Discussion paper and request for feedback*. <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>

European Commission. (2024). *Artificial Intelligence Act: Council gives final green light to the first worldwide rules on AI*. <https://www.consilium.europa.eu/en/press/press-releases/2024/05/21/artificial-intelligence-ai-act-council-gives-final-green-light-to-the-first-worldwide-rules-on-ai/>

European Medicines Agency. (2020). *Nitrosamine impurities in human medicinal products*. <https://www.ema.europa.eu/en/human-regulatory-overview/post-authorisation/pharmacovigilance-post-authorisation/referral-procedures-human-medicines/nitrosamine-impurities>

U.S. Food and Drug Administration. (2023). *Artificial Intelligence in Drug Manufacturing: Draft Guidance for Industry*. <https://www.fda.gov/media/165743/download>

European Union. (2024). *Regulation (EU) 2024 of the European Parliament and of the Council of 13 March 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts*. Official Journal of the European Union, L 206, 1–84.

<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>

Topol, E. (2019). *High-performance medicine: the convergence of human and artificial intelligence*. *Nature Medicine*, 25(1), 44–56. <https://www.nature.com/articles/s41591-018-0300-7>

Wiens, J., Saria, S., Sendak, M., et al. (2019). *Do no harm: a roadmap for responsible machine learning for health care*. *Nature Medicine*, 25(9), 1337–1340. <https://doi.org/10.1038/s41591-019-0548-6>

European Medicines Agency. (2024). *Reflection paper on the use of artificial intelligence (AI) in the medicinal product lifecycle* (EMA/CHMP/CVMP/83833/2023). https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf

Biden, J. R. (2023). *Executive Order No. 14110: Safe, secure, and trustworthy development and use of artificial intelligence*. *Federal Register*, 88(210), 75191–75210. <https://www.govinfo.gov/content/pkg/FR-2023-11-01/pdf/2023-24283.pdf>

The White House. (2025). *Initial rescissions of harmful executive orders and actions*. <https://www.whitehouse.gov/presidential-actions/2025/01/initial-rescissions-of-harmful-executive-orders-and-actions/>

U.S. Food and Drug Administration. (2025). *Considerations for the use of artificial intelligence to support regulatory decision-making for drug and biological products: Guidance for industry and other interested parties (Draft Guidance)*. <https://www.fda.gov/media/184830/download>

European Society of Radiology. (n.d.). *Risk-based classification of AI systems under the EU AI Act, with examples and regulatory requirements*. In *EU AI Act – Policy Topics*. Retrieved May 25, 2025, from <https://www.myesr.org/eu-international-affairs/policy-topics/>

Saenz, A. D., Mass General Brigham AI Governance Committee, Centi, A., Ting, D., You, J. G., Landman, A., & Mishuris, R. G. (2024). *Establishing responsible use of AI guidelines: A comprehensive case study for healthcare institutions*. npj Digital Medicine, 7, Article 55. <https://doi.org/10.1038/s41746-024-01300-8>