Classifying GenAI under the European Union's Medical Device Regulation

Benedikt Kolbeinsson
Askan, London, UK
K01, Reykjavík, Iceland
benedikt@askan.is

Arinbjörn Kolbeinsson
Askan, London, UK
K01, Reykjavík, Iceland
arinbjorn@askan.is

Abstract

The rapid growth of Generative Artificial Intelligence (GenAI) in healthcare has introduced novel systems and processes that require careful classification under the European Union's Medical Device Regulation (MDR). The diversity of these applications raises important questions about their regulatory categorization. Proper classification of GenAI systems is essential, as it determines the level of regulatory oversight, impacting both patient safety and the pace of healthcare innovation. This paper addresses the critical issue of whether GenAI-tools should be classified as a medical device under the MDR, focusing on Rule 11. By exploring various use cases, we provide a detailed analysis of the regulatory implications, offering insights for developers, healthcare providers and regulators on navigating this emerging field.

1 Introduction

Generative Artificial Intelligence (GenAI), a technology that creates new content based on patterns learned from existing data, is rapidly evolving with significant potential in healthcare applications. In healthcare, GenAI can generate synthetic data - artificial but realistic health information that simulates patient records, medical images, laboratory results, and other clinical data without using actual patient information (Goncalves et al., 2020). The increasing adoption of GenAI-powered tools in clinical settings raises critical questions about their regulation, particularly within the framework of the European Union's Medical Device Regulation (EU MDR).

The EU MDR (Official Journal of the European Union, 2017), implemented in 2021, provides a regulatory structure for ensuring the safety and effectiveness of medical devices marketed within the EU. Rule 11 of the MDR is particularly relevant for software applications, classifying them based on their role in clinical decision-making and the associated risk to patients. Understanding whether and how GenAI applications fall within the scope of these regulations is vital for developers, healthcare providers, and regulators alike (Muehlematter et al., 2021).

Classifying GenAI under the MDR is essential because its regulatory status will determine the level of scrutiny it faces, the compliance burdens on developers, and ultimately the safety and reliability of the healthcare applications that use synthetic data (Group et al., 2019). Misclassification could lead to under-regulation, potentially putting patients at risk if synthetic data is used improperly in clinical settings. Conversely, over-regulation could stifle innovation in a field with significant potential for advancing healthcare. This paper examines the classification of GenAI applications under the EU MDR, with particular focus on Rule 11, and discusses the implications for healthcare innovation and patient safety.

GenAI for Health Workshop @ NeurIPS 2024, Vancouver.

2 Overview of the EU Medical Device Regulation (MDR)

The EU Medical Device Regulation (MDR) came into force in May 2021, replacing the earlier Medical Device Directive (MDD). While a directive provides guidelines for member states to create their own laws, the MDR as a regulation applies directly across all EU member states, establishing stricter unified rules to ensure the safety and performance of medical devices (Official Journal of the European Union, 2017). The MDR implements a risk-based classification system, with higher-risk devices facing more stringent requirements for clinical evidence, post-market surveillance, and transparency.

Of particular relevance to GenAI is *Rule 11* of the MDR, which governs the classification of software that qualifies as a medical device. Under the MDR, software is considered a medical device if it is intended to serve one or more specific medical purposes, such as diagnosis, monitoring, or treatment of a disease or injury. The classification of GenAI applications thus depends primarily on their intended use.

For GenAI applications, the critical factor in determining medical device classification is the *intended purpose* of the system and its output. If the output is intended to inform or influence clinical decisions, such as providing data that aids in diagnosis or treatment planning, it would likely be classified as a medical device. However, GenAI applications used solely for non-clinical purposes, such as research, academic study, or model development without direct interaction with patient care, fall outside the MDR's scope.

Understanding this classification framework is essential because it determines the level of regulatory scrutiny, required documentation, and approval processes. Improper classification could either expose patients to unnecessary risks or impose excessive regulatory burdens on developers, potentially hindering innovation in healthcare.

3 GenAI Classification

Once a GenAI system is considered a medical device under the MDR, it must be classified according to its risk level. The MDR employs a risk-based classification system that assigns medical devices to four classes: I, IIa, IIb, and III, in order of increasing risk level. Class I represents minimal risk to patients, Class IIa and IIb represent moderate risk with increasing severity, and Class III represents the highest risk level for devices that may significantly impact patient safety.

All software, including GenAI, is considered an *active device* under the MDR. The classification of medical software is governed by *Rule 11*, which provides specific guidelines for categorization based on the software's intended purpose and potential impact on patient health.

According to Rule 11 of the MDR:

> *Software intended to provide information which is used to make decisions with diagnosis or therapeutic purposes is classified as class IIa, except if such decisions have an impact that may cause:*
>
> - *death or an irreversible deterioration of a person's state of health, in which case it is in class III; or*
> - *a serious deterioration of a person's state of health or a surgical intervention, in which case it is classified as class IIb.*

This means that a GenAI system providing information to support decisions for diagnostic or therapeutic purposes would generally be classified as **Class IIa**. However, if the decisions informed by the system could potentially lead to death or irreversible health deterioration, the system would be classified as **Class III**. For example, a GenAI system that generates synthetic health data used in critical treatment planning for life-threatening conditions might fall under this classification.

Similarly, Rule 11 also states that:

> *Software intended to monitor physiological processes is classified as class IIa, except if it is intended for monitoring of vital physiological parameters, where the nature of variations of those parameters is such that it could result in immediate danger to the patient, in which case it is classified as class IIb.*

If a GenAI tool is used for monitoring physiological processes, such as generating synthetic data to simulate heart rate variability or other vital signs, it would typically be classified as **Class IIa**. However, if the software is intended to monitor vital physiological parameters in a way that any significant variation could pose immediate danger to the patient, the classification would increase to **Class IIb**. For instance, a GenAI application simulating vital signs that assist in real-time monitoring during a high-risk surgery might fall under this category.

Lastly, Rule 11 concludes with:

> *All other software is classified as class I.*

Thus, if the GenAI system does not fall into the categories described above—such as software used for research, non-clinical purposes, or administrative tasks—it would be classified as **Class I**, indicating minimal risk.

In summary, the classification of GenAI under the MDR depends on its intended use and the associated risks to patient health. The classification framework progresses from Class I (minimal risk) through Class IIa and IIb (moderate risk with increasing severity) to Class III (highest risk), ensuring that more stringent controls are applied as the potential impact on patient safety increases.

4 Practical Use Cases: Classifying GenAI Applications Under MDR

This section examines a series of relevant use cases and edge cases that practitioners might encounter when attempting to classify GenAI applications under the MDR. Each scenario is analyzed to provide insight into how the MDR might be interpreted in these specific contexts. For a comprehensive classification of GenAI applications in healthcare, we applied a topology framework that evaluates both Decision-Making Relevance and Therapeutic Risk Potential, as detailed in Appendix A.

Case 1: Synthetic Training Data for Medical Image Analysis A research team develops a GenAI system that generates synthetic medical images (e.g., X-rays, MRIs) to augment training datasets for a medical image analysis algorithm.

MDR Interpretation: This case likely falls outside the scope of MDR as a medical device. The synthetic data is not directly used in patient care but is instead used to train another system. However, if the resulting medical image analysis algorithm is intended for clinical use, the GenAI system might be considered an accessory to a medical device. In this case, it could potentially be classified as Class I, with relatively lighter regulatory requirements.

Case 2: Digital Twin for Treatment Planning A hospital uses a GenAI system to create digital twins of patients based on their medical history, genetic data, and current health status. These digital twins are then used to simulate different treatment options and inform clinical decision-making.

MDR Interpretation: This application would likely be classified as a medical device under MDR, as it directly influences patient treatment decisions. Depending on the specific medical conditions and treatments involved, it could be classified as Class IIa or IIb. If used for high-risk treatments or critical care decisions, it might even be considered Class III. Compliance requirements would be substantial, including clinical evaluation, risk management, and possibly clinical trials.

Case 3: Synthetic Clinical Trial Data for Protocol Design A pharmaceutical company uses GenAI to generate synthetic patient data for designing clinical trial protocols. The synthetic data helps in power calculations and identifying potential confounding factors but is not used in the actual trial or for regulatory submissions.

MDR Interpretation: This use case would likely fall outside the scope of MDR. The synthetic data is used for research planning purposes and does not directly influence patient care or clinical decisions. However, developers should be cautious if there is any possibility that this synthetic data could indirectly influence the design of medical devices or drugs.

Case 4: GenAI-Enhanced Electronic Health Records A healthcare software company develops a GenAI system that integrates with Electronic Health Records (EHR) to generate detailed patient summaries, suggest potential diagnoses, and recommend follow-up tests.

MDR Interpretation: This system would almost certainly be classified as a medical device under MDR, as it directly influences clinical decision-making. It would likely be classified as Class IIa or IIb, depending on the level of influence it has on diagnosis and treatment decisions. If it suggests critical care decisions, it could potentially be Class III. Extensive documentation, clinical evaluation, and post-market surveillance would be required.

Case 5: Synthetic Biomarker Data for Algorithm Testing A biotech startup uses GenAI to create synthetic biomarker data to test and validate new diagnostic algorithms before using real patient data.

MDR Interpretation: If the synthetic data is used solely for initial testing and validation, and not for the final validation or training of the diagnostic algorithm intended for clinical use, it might fall outside the scope of MDR. However, if the synthetic data plays a crucial role in the development of a medical device (the diagnostic algorithm), it could be considered an accessory to a medical device. The classification would depend on the intended use and risk level of the final diagnostic algorithm.

Case 6: GenAI for Medical Education and Training A medical school develops a GenAI system that generates realistic patient scenarios for student training, including simulated patient interactions and diagnostic challenges.

MDR Interpretation: This application would likely fall outside the scope of MDR, as it is intended for educational purposes rather than direct patient care. However, if the system is later adapted for use in clinical settings (e.g., for physician decision support), it would then need to be evaluated under MDR criteria.

Case 7: Personalized Treatment Response Prediction A GenAI system is developed to predict individual patient responses to various treatments based on synthetic data generated from vast amounts of anonymized patient records.

MDR Interpretation: This system would likely be classified as a medical device under MDR, as it directly influences treatment decisions. The classification could range from Class IIa to III, depending on the types of treatments involved and the level of risk associated with the decisions it informs. If used for life-critical decisions, it would likely be Class III, requiring the highest level of scrutiny and clinical evidence.

Case 8: GenAI for Drug Discovery A pharmaceutical company uses GenAI to generate and screen potential drug candidates, including predicting their effects on synthetic patient populations.

MDR Interpretation: In the early stages of drug discovery, this application would likely fall outside the scope of MDR. However, if the GenAI system is later used to inform clinical trial designs or drug dosing decisions, it could then be considered a medical device. The classification would depend on the specific use case and the level of influence on patient care decisions.

These use cases demonstrate the complexity of classifying GenAI applications under the EU MDR. Key factors in classification include:

- The intended use of the GenAI system and its outputs
- The directness of its influence on patient care and clinical decision-making
- The risk level associated with its use in healthcare settings
- Whether it functions as a standalone system or as an accessory to other medical devices

Practitioners should carefully consider these factors when developing and deploying GenAI systems in healthcare contexts. When in doubt, it is advisable to consult with regulatory experts or the relevant authorities for guidance on classification and compliance requirements. As the field of GenAI in healthcare continues to evolve, ongoing dialogue between developers, healthcare providers, and regulators will be crucial to ensure both innovation and patient safety.

Appendix A.1 outlines the two primary dimensions—Decision-Making Relevance and Therapeutic Risk Potential—used to systematically assess each GenAI use case in terms of regulatory impact. The rationale for assigning specific scores to each use case is thoroughly explained in Appendix A.3.

For a visual overview, Figure 1 in the appendix positions each use case within the two-dimensional framework, providing a nuanced view of their regulatory implications.

5 Discussion

This paper has examined the classification of GenAI systems for synthetic health data under the European Union's MDR, with a particular focus on Rule 11. Our analysis reveals the complex nature of this classification process, which depends heavily on the intended use of the GenAI system and its potential impact on patient care and clinical decision-making.

While our examination provides a framework for understanding how GenAI might be classified under the MDR, it is important to acknowledge the limitations of this work. Firstly, the rapidly evolving nature of GenAI technology means that new applications and use cases may emerge that challenge current interpretations of the MDR (Brown, 2020; Yi et al., 2019). Secondly, our analysis is based on the current text of the MDR and existing guidance; future amendments or clarifications from regulatory bodies could alter these interpretations (Maresova et al., 2020). Additionally, the hypothetical use cases presented, while informative, do not capture the full range of potential GenAI applications in healthcare.

Furthermore, this paper does not address the technical challenges of implementing the necessary safeguards and validation processes for GenAI systems that are classified as medical devices. The unique nature of GenAI, particularly its ability to generate novel data, presents unprecedented challenges in terms of testing, validation, and ongoing monitoring that are not fully addressed by current regulatory frameworks (Doshi-Velez and Kim, 2017).

Future work in this area should focus on several key directions. First, there is a need for empirical studies on the real-world implementation of GenAI in healthcare settings, which could provide valuable insights into the practical challenges of regulatory compliance. Second, research into the development of specialized validation methodologies for GenAI systems in healthcare could help bridge the gap between current regulatory requirements and the unique characteristics of these systems.

Moreover, there is a pressing need for interdisciplinary collaboration between AI researchers, healthcare professionals, ethicists, and legal experts to develop comprehensive guidelines for the responsible development and deployment of GenAI in healthcare (Mittelstadt et al., 2016). Such guidelines should address not only regulatory compliance but also ethical considerations, data privacy, and the potential long-term impacts of synthetic health data on medical research and practice.

Lastly, as the field of GenAI continues to advance, ongoing dialogue with regulatory bodies will be crucial. This could involve the development of GenAI-specific guidance documents or even potential amendments to the MDR to more explicitly address the unique challenges posed by this technology (Pesapane et al., 2018).

6 Conclusion

The MDR provides a necessary framework for regulating GenAI applications in healthcare, though adapting it to these technologies presents unique challenges. Our analysis illustrates that the MDR's classification, particularly under Rule 11, hinges on GenAI's intended use and its impact on clinical decisions and patient safety. This structured approach, applied across various use cases, demonstrates how regulatory clarity can support responsible innovation in healthcare.

As GenAI applications continue to advance, close collaboration between developers, healthcare providers and regulators will be essential. By aligning the MDR's evolving requirements with GenAI's capabilities, we can ensure these tools enhance patient care while maintaining the highest standards of safety and ethical responsibility.

References

Brown, T. B. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Doshi-Velez, F. and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

Goncalves, A., Ray, P., Soper, B., Stevens, J., Coyle, L., and Sales, A. P. (2020). Generation and evaluation of synthetic patient data. *BMC medical research methodology*, 20:1–40.

Group, M. D. C. et al. (2019). Guidance on qualification and classification of software in regulation (eu) 2017/745–mdr and regulation (eu) 2017/746–ivdr. *European Commission*.

Maresova, P., Hajek, L., Krejcar, O., Storek, M., and Kuca, K. (2020). New regulations on medical devices in europe: are they an opportunity for growth? *Administrative Sciences*, 10(1):16.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2):2053951716679679.

Muehlematter, U. J., Daniore, P., and Vokinger, K. N. (2021). Approval of artificial intelligence and machine learning-based medical devices in the usa and europe (2015–20): a comparative analysis. *The Lancet Digital Health*, 3(3):e195–e203.

Official Journal of the European Union (2017). Regulation (eu) 2017/745 of the european parliament and of the council of 5 april 2017 on medical devices. *Official Journal of the European Union*, L117:1–175.

Pesapane, F., Volonté, C., Codari, M., and Sardanelli, F. (2018). Artificial intelligence as a medical device in radiology: ethical and regulatory issues in europe and the united states. *Insights into imaging*, 9:745–753.

Yi, X., Walia, E., and Babyn, P. (2019). Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552.

A Appendix: Topology Framework for Classifying GenAI Use Cases

To systematically classify GenAI applications under MDR, we establish a two-dimensional topology based on two critical dimensions: **Decision-Making Relevance** and **Therapeutic Risk Potential**. This framework facilitates a nuanced understanding of each use case's regulatory implications by quantitatively assessing their influence on clinical decisions and the associated risks to patient health.

A.1 Dimension Definitions

Decision-Making Relevance This dimension quantifies the extent to which a GenAI system impacts clinical or treatment decisions. It is scored on a scale from **0** to **10**, where:

- **0** indicates minimal to no influence on decision-making.
- **5** represents a moderate level of influence, providing supportive information without direct decision-making authority.
- **10** signifies a critical influence, directly guiding or altering clinical or treatment decisions.

Therapeutic Risk Potential This dimension assesses the potential risk associated with the use of a GenAI system, particularly concerning patient safety and health outcomes. It is also scored on a scale from **0** to **10**, where:

- **0** denotes negligible risk to patient health.
- **5** indicates a moderate level of risk with some implications for patient safety.
- **10** represents a high to critical risk, directly affecting patient safety and health outcomes.

A.2 Scoring of Use Cases

Each of the eight use cases is evaluated and assigned scores on both dimensions based on their intended use, impact on clinical workflows, and potential risks to patient health. The scoring aims to provide a clear, quantitative basis for positioning each case within the regulatory landscape.

Table 1: Updated Scoring of GenAI Use Cases on Decision-Making Relevance and Therapeutic Risk Potential

Use Case	**Decision-Making Relevance**	**Therapeutic Risk Potential**
Case 1: Synthetic Training Data for Medical Image Analysis	2	5
Case 2: Digital Twin for Treatment Planning	9	7
Case 3: Synthetic Clinical Trial Data for Protocol Design	3	2
Case 4: GenAI-Enhanced Electronic Health Records	9	6
Case 5: Synthetic Biomarker Data for Algorithm Testing	3	6
Case 6: GenAI for Medical Education and Training	1	1
Case 7: Personalized Treatment Response Prediction	10	8
Case 8: GenAI for Drug Discovery	5	3

A.3 Rationale for Scoring

Case 1: Synthetic Training Data for Medical Image Analysis Assigned a low Decision-Making Relevance score of **2** because the GenAI system is used for generating training data and does not directly influence clinical decisions. The Therapeutic Risk Potential is set at **5** due to the possibility that inaccuracies in the synthetic data could indirectly affect the performance of diagnostic algorithms, potentially impacting patient care.

Case 2: Digital Twin for Treatment Planning With a high Decision-Making Relevance score of **9**, this system directly influences treatment planning. The Therapeutic Risk Potential is moderately high at **7**, reflecting that while the system significantly impacts clinical decisions, the risk is somewhat mitigated by clinician oversight.

Case 3: Synthetic Clinical Trial Data for Protocol Design Scores a Decision-Making Relevance of **3**, as it informs trial design but does not directly affect patient care. The Therapeutic Risk Potential is low at **2**, given that any risks are indirect and confined to the research planning phase.

Case 4: GenAI-Enhanced Electronic Health Records Receives a high Decision-Making Relevance score of **9** due to its direct influence on diagnoses and treatment recommendations. The Therapeutic Risk Potential is set at **6**, acknowledging significant risk, though mitigated by clinician validation.

Case 5: Synthetic Biomarker Data for Algorithm Testing Assigned a Decision-Making Relevance score of **3**, reflecting its supportive role in algorithm testing. The Therapeutic Risk Potential is higher at **6** because inaccuracies in the synthetic data could lead to flawed algorithms, potentially impacting patient diagnoses if used clinically.

Case 6: GenAI for Medical Education and Training Scores a Decision-Making Relevance of **1** and a Therapeutic Risk Potential of **1**, indicating minimal impact on clinical decisions and negligible risk to patient health, as it is intended for educational purposes.

Case 7: Personalised Treatment Response Prediction Achieves the highest Decision-Making Relevance score of **10**, as it directly predicts individual treatment responses. The Therapeutic Risk Potential is slightly lower at **8**, considering that clinical judgment may mitigate some risks.

Case 8: GenAI for Drug Discovery Assigned a moderate Decision-Making Relevance score of **5**, reflecting its influence on early-stage research decisions. The Therapeutic Risk Potential is low at **3**, as risks to patients are minimal during the drug discovery phase.

A.4 Visualisation and Interpretation

Figure 1 illustrates the positioning of each use case within the two-dimensional topology, providing a visual representation of their regulatory implications under the EU MDR.

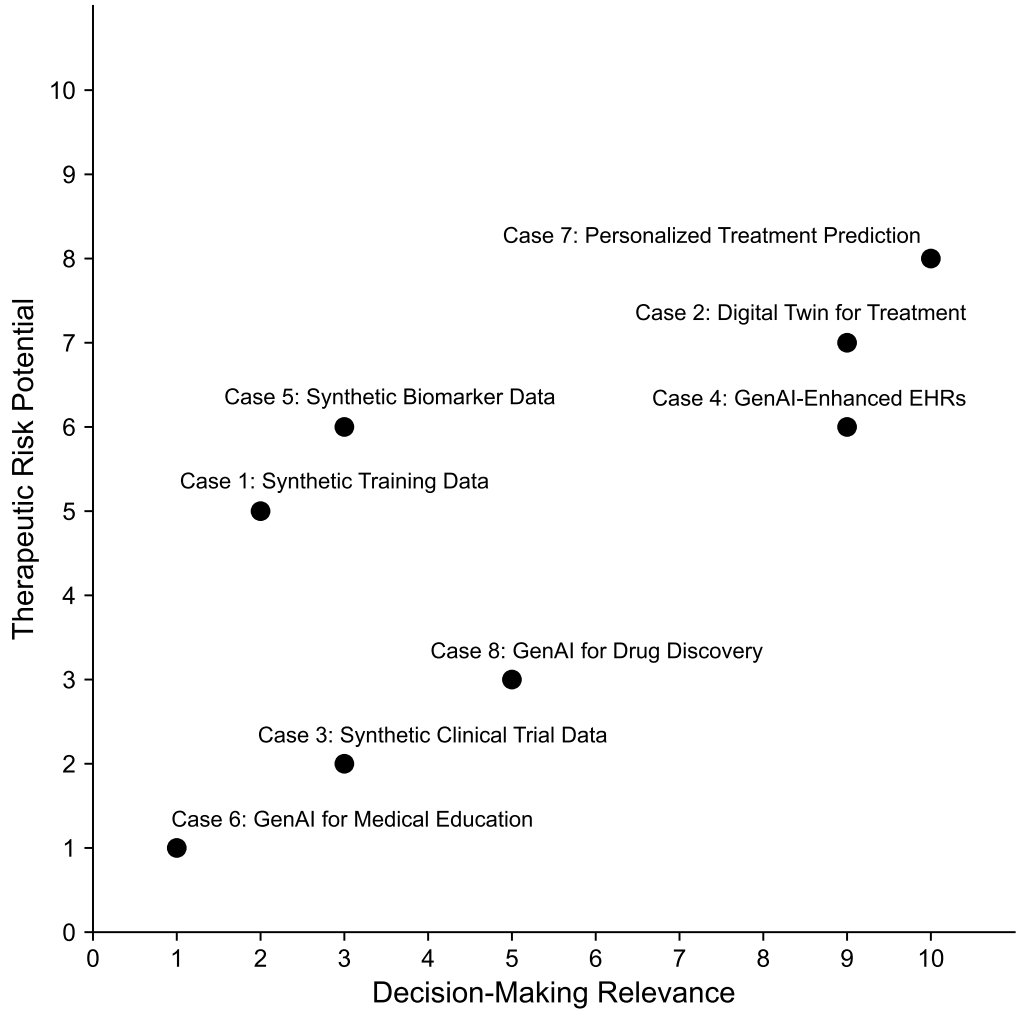

Figure 1: Topology of GenAI Use Cases based on Decision-Making Relevance and Therapeutic Risk Potential