

UDK: 004.738.5 : 34 (094.5)
341.24 : 615 (4-67
EU)(094.5)
Preliminary communication

Lisa Marie Breyer, PhD student
Faculty of Law, University of Vienna, Austria
E-mail: Lisa.marie.breyer@univie.ac.at
ORCID-ID: 0009-0004-0072-3302

THE PRESUMPTION OF CONFORMITY FOR MEDICAL AI UNDER THE AI ACT: STRUCTURAL LIMITS OF MDR CONFORMITY ASSESSMENT FOR CONTINUOUS-LEARNING GENETIC DIAGNOSTIC SOFTWARE

ABSTRACT

Continuous-learning deep-learning (DL) systems used in genetic diagnostics sit at the intersection of the EU Artificial Intelligence Act (AI Act), the Medical Device Regulation (MDR), and data protection law. Under Article 6(1) AI Act and Annex I, Section A, such systems qualify as high-risk AI systems where they constitute medical device software within the meaning of Article 2 MDR and are subject to third-party conformity assessment under the MDR. For Annex I systems, Article 43(3) channels conformity assessment through the MDR, while compliance with harmonised standards may give rise to a presumption of conformity under Article 42 AI Act. This article assesses whether MDR conformity assessment—predominantly ex ante and premised on stable device identities—provides an equivalent basis for that presumption when applied to continuously learning genetic diagnostic software. Doctrinal analysis of MDR classification (Annex VIII, Rule 11) and the conformity assessment routes (Annexes IX–XI) reveals a structural mismatch. MDR conformity assessment relies on snapshot-based technical documentation, verification/validation testing, and quality management systems that presuppose a fixed intended purpose and stable performance characteristics. While Annex IX offers procedural completeness through controlled change governance, it does not enable continuous substantive reassessment of evolving behaviour; Annex X and Annex XI reinforce type-based and unit-based verification logics that are difficult to reconcile with model-state evolution. The article demonstrates that, for continuous-learning genetic diagnostics, the MDR's conformity architecture may not be functionally equivalent to the AI Act's lifecycle-oriented risk logic, thereby limiting the reliability of the presumption of conformity in this domain.

Keywords: Artificial Intelligence Act, Medical Device Regulation, Presumption of Conformity, Continuous Learning Systems, Genetic Diagnostics

1. INTRODUCTION

Deep-learning (hereinafter: DL) systems in genetic diagnostics raise unprecedented legal, ethical, and technical challenges. Artificial Intelligence (hereinafter: AI) used for genetic analysis is governed by several regulatory frameworks at both European and national levels, including the Artificial Intelligence-Act (hereinafter: AI Act),¹ the Medical Device Regulation (hereinafter: MDR),² the General Data Protection Regulation (hereinafter: GDPR),³ and, for example, in Austria, the Genetic Technology Act (Gentechnikgesetz, hereinafter: GTG).⁴ These frameworks pursue distinct objectives⁵ and are structured around different regulatory logics,⁶ risk concepts,⁷ and enforcement mechanisms.⁸

The AI Act entered into force on 1 August 2024.⁹ Legal scholarship has established that a deep-learning model used for diagnosis qualifies as software, and accordingly, as a medical device within the meaning of the MDR.¹⁰ Consequently,

¹ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689 (hereinafter: AI Act).

² Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices (Medical Device Regulation), OJ L 2017/117 (hereinafter: MDR); Königshofer, W., Medical Device Law, in: Resch, R.; Wallner, F. (eds.), Handbook of Medical Law, 3rd ed., Manz, Vienna 2020, para. 23; European Commission, *Medical devices: Guidance document – Qualification and Classification of stand alone software (MEDDEV 2.1/6)*, July 2016, p. 9 *et seq.*, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026. This document is based on Directive 93/42/EEC; Annex II of the Council Resolution of 7 May 1985 on a new approach to technical harmonisation and standardisation (85/C 136/01), OJ C 1985/136, 1; Schneeberger, D., *Intelligente Medizinprodukte: Rechtsfragen am Schnittpunkt von DSGVO, MPVO und AI Act*, DaKo, 1/2024, p. 5 *et seq.*

³ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation), OJ L 2016/119 (hereinafter: GDPR); Herbst, T., *Regulierung intelligenter Medizinprodukte*, ZfPC, 2023, p. 246.

⁴ Genetic Technology Act (Gentechnikgesetz – GTG), Federal Law Gazette (BGBl) No 510/1994, as amended by BGBl I No 8/2022 (hereinafter: GTG).

⁵ Recitals 1, 8, 176 AI Act; Recital 1, 2, 101 MDR; Recital 170 GDPR.

⁶ Recital 26 AI Act; *European Commission*, Commission Notice – The ‘Blue Guide’ on the implementation of EU product rules 2022, OJ C 247, 29.6.2022, 9-13, 70-80.

⁷ Recitals 26, 27, 64-66, Article 4, 9 AI Act; Recitals 31-33, 58, 59, Article 1, 51, Annex VIII MDR; Recital 75, Article 35 GDPR.

⁸ Chapter IX Section 3 AI Act; Chapter VII MDR; Article 51 and 83 GDPR.

⁹ European Commission, AI Act enters into force, 1 August 2024, https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en, accessed 21 February 2026.

¹⁰ Article 2(1) MDR; Königshofer, W., *op. cit.* note 2, para. 23; European Commission, *Medical devices: Guidance document – Qualification and Classification of*

such software constitutes a high-risk AI system within the meaning of Article 6(1) of the AI Act,¹¹ while, at the same time, remaining subject to the applicable provisions of the GDPR.¹²

Deep-learning models introduce specific risks, including misdiagnosis, bias,¹³ performance uncertainty¹⁴ or unpredictable model behaviour.¹⁵ Their feedback loops¹⁶ and ability for continuous learning¹⁷ result in evolving neuronal network configurations¹⁸ that may interact with patient data in unforeseen ways. The AI Act seeks to address certain of these risks, *inter alia*, through the establishment of a risk management system¹⁹, specific technical requirements relating to feedback loops,²⁰ and conformity assessment obligations.²¹

stand alone software (MEDDEV 2.1/6), July 2016, p. 9 *et seq.*, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026. This document is based on Directive 93/42/EEC; Annex II of the Council Resolution of 7 May 1985 on a new approach to technical harmonisation and standardisation (85/C 136/01), OJ C 1985/136, 1.

¹¹ This analysis is based on Regulation (EU) 2024/1689 as adopted. At the time of writing, the European Commission has proposed amendments that would move medical devices governed by Regulation (EU) 2017/745 (MDR) from Annex I, Section A to Section B of the AI Act. Should such amendments be adopted, the direct applicability of certain high-risk AI system requirements may be modified pursuant to Article 2(2) AI Act. The structural question examined in this article—namely whether the MDR conformity architecture is conceptually capable of governing continuously learning systems—remains relevant irrespective of the final legislative configuration. See: <https://www.osborneclarke.com/insights/eu-mdr-and-ivdr-poised-remain-main-framework-medical-ai-draft-ai-interface-reforms>; <https://www.taylorwessing.com/en/insights-and-events/insights/2026/02/reformder-mdr-und-ivdr>; <https://www.emergobyul.com/news/ec-requests-feedback-proposed-regulation-simplify-mdr-ivdr>.

¹² Article 2(7) AI Act; Herbst, T., *op. cit.* note 3, p. 246.

¹³ Zhang, A. *et al.*, Shifting Machine Learning for Healthcare from Development to Deployment and from Models to Data, *Nature Biomedical Engineering*, 6(12), 2022, p. 1343-1344. doi:10.1038/s41551-022-00898-y; Wang, X. *et al.*, Inconsistent Performance of Deep Learning Models on Mammogram Classification, *Journal of the American College of Radiology*, 17(6), 2020, p. 796-803. doi: 10.1016/j.jacr.2020.01.006.

¹⁴ Zhang *et al.*, *op. cit.* note 13, p. 1343; Gama, J. *et al.*, A Survey on Concept Drift Adaptation, *ACM Computing Surveys*, 46(4), 2014, p. 3-4, doi:10.1145/2523813.

¹⁵ Gama *et al.*, *op. cit.* note 14, p. 3-4.

¹⁶ Sculley, D. *et al.*, Hidden Technical Debt in Machine Learning Systems, *Advances in Neural Information Processing Systems*, 2015, p. 2506.

¹⁷ Goodfellow, I.; Bengio, Y.; Courville, A., *Deep Learning*, MIT Press, Cambridge (MA) 2016, p. 164-325; for adaptive learning see: Gama *et al.*, *op. cit.* note 14, p. 3-4.

¹⁸ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17 p. 164-325; Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343.

¹⁹ Article 9 AI Act.

²⁰ Article 15(4) AI Act.

²¹ Article 43(3) AI Act.

By contrast, the MDR is primarily designed to ensure that only safe and effective medical devices are placed on the market.²² Its risk focus is centred on patient safety,²³ side effects,²⁴ risk-benefit considerations²⁵ and clinical efficacy.²⁶ While these elements are fundamental to product safety law, it remains questionable whether the MDR adequately captures the technological complexity and dynamic behaviour of continuously learning deep-learning modules.

Effective compliance with the AI Act requires careful consideration of the challenges posed by feedback loops²⁷ and continuous learning,²⁸ because these characteristics may lead to changes in predictive behaviour over time.²⁹ Despite several recitals addressing system modification, it remains unclear when a "continuous learning system" within the meaning of the AI Act must be regarded as having undergone a "substantial modification." Recital 128 AI Act provides only a general framework and does not offer conclusive criteria. As a result, the practical interpretation of "substantial modification" requires further legal clarification, particularly in the context of genetic diagnostics. This includes the question of whether retraining on new datasets necessitates a renewed conformity assessment and how such an obligation is to be operationalised in practice. Under Article 3(20) AI Act, a conformity assessment is a procedure designed to demonstrate whether a high-risk AI system complies with the requirements laid down in Title III, Section 2. With regard to high-risk AI systems, Article 43 AI Act establishes the applicable conformity assessment framework. The AI Act further provides for a presumption of conformity where systems comply with harmonised standards whose references have been published in the Official Journal of the European Union in accordance with Regulation (EU) No 1025/2012.³⁰ This raises a central question: whether the conformity assessment procedures under the MDR provide an equivalent and sufficiently comprehensive basis for such a presumption in the case of continuously learning genetic diagnostic systems. A core difficulty in applying the AI Act to genetic diagnostic systems lies in the adaptive nature of deep-learning algorithms, which may evolve

²² Recital 1, 2 MDR; *European Commission*, Commission Notice – The 'Blue Guide' on the implementation of EU product rules 2022, OJ C 247, 29.6.2022, 5, 12.

²³ Article 59, 80, 87, 89 MDR.

²⁴ Recital 8 MDR.

²⁵ Recital 81 MDR.

²⁶ For example, the general obligations of manufacturers in Article 10(9) in conjunction with Annex IX, Chapter I MDR.

²⁷ Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

²⁸ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325.

²⁹ *Ibid*; for adaptive learning see: Gama *et al.*, *op. cit.* note 14, p. 3-4.

³⁰ Feiler, L.; Forgó, N., Kommentar KI-VO, EU-Verordnung über künstliche Intelligenz, Verlag Österreich, Vienna 2024, p. 6.

over time.³¹ At the same time, Article 43(4), second sentence, clarifies that changes already taken into account in the original conformity assessment do not constitute significant changes. Thus, changes that occur within a previously defined and assessed framework do not constitute a "substantial modification" within the meaning of Article 43(4) AI Act and, accordingly, do not trigger a renewed conformity assessment.

However, in the context of genetic analysis, it remains uncertain what precisely constitutes a "substantial modification,"³² particularly where models are regularly retrained.³³ A deep-learning system performing genetic analyses may access new data post-deployment and detect correlations not covered by the initial conformity assessment. Such evolving behaviour may generate unreliable outputs, including erroneous diagnoses or risk predictions,³⁴ with potentially far-reaching consequences for patients and, indirectly, for biological relatives.

Finally, the concurrent application of the AI Act, MDR, and GDPR seeks to harmonise obligations for AI-based medical software.³⁵ Yet this regulatory layering also creates the risk of fragmentation, overregulation, and potential gaps in protection. Without a coherent doctrinal analysis of the interaction between these frameworks, there is a risk of both unnecessary regulatory burden³⁶ and insufficient safeguards for affected individuals. This is particularly relevant in the field of genetic diagnostics, where technical performance, clinical reliability, data protection, fundamental rights protection, and the interests of biological relatives are closely interlinked and cannot be assessed in isolation.

1.1. Methodology

This article adopts a doctrinal legal methodology. It undertakes a systematic and teleological interpretation of the AI Act and the Medical Device Regulation, with

³¹ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for adaptive learning see: Gama *et al.*, *op. cit.* note 14, p. 3-4.

³² Medical Device Coordination Group (MDCG), *MDCG 2020-3 Rev.1 – Guidance on significant changes regarding the transitional provision under Article 120 MDR*, 2023, p. 3-5, https://health.ec.europa.eu/document/download/5fb74e56-fd4c-4dae-9d24-78149edfe56b_en?filename=md_mdcg_2020_3_rev1_en.pdf, accessed 21 February 2026. While the MDCG guidance is non-binding and drafted for Article 120 MDR devices; it is cited here solely as evidence of the regulator's event-based understanding of 'significant change' in practice.

³³ Zhang *et al.*, *op. cit.* note 13, p. 1336-1345.

³⁴ Zhang *et al.*, *op. cit.* note 13, p. 1343-1344; Wang, X. *et al.*, *op. cit.* note 13, p. 796–803.

³⁵ Recitals 8-10, 27, Article 43 AI Act.

³⁶ Herbst, T., *op. cit.* note 3, p. 246.

particular focus on their conformity assessment architectures and underlying regulatory assumptions. The analysis is based on the consolidated versions of Regulation (EU) 2024/1689 and Regulation (EU) 2017/745 as retrieved from EUR-Lex and the Official Journal of the European Union. Secondary literature was identified and reviewed through legal and interdisciplinary databases, including LexisNexis, Rechtsdatenbank Manz (hereinafter: RDB), and Beck-Online. The study further draws on guidance documents issued by the European Commission and the Medical Device Coordination Group (hereinafter: MDCG). Relevant interdisciplinary scholarship from medical and technical journals was reviewed where necessary to contextualise the regulatory analysis.

1.2. Structure of the article

Section 2 sets out the Annex I logic of the AI Act and explains how Article 43(3) links MDR conformity assessment to the AI Act's high-risk framework. Section 3 analyses MDR classification under Annex VIII, with particular focus on Rule 11 and its stability-oriented assumptions. Section 4 examines the ex ante structure of MDR conformity assessment under Annexes IX–XI, highlighting the implications for continuously learning genetic diagnostic software. Section 5 concludes by assessing the limits of functional equivalence underlying the presumption of conformity in this domain.

2. AI ACT ANNEX I LOGIC AND THE PRESUMPTION OF CONFORMITY

It is widely recognised in the literature that a deep-learning model used for disease diagnosis, prediction, or prognosis constitutes stand-alone software and is therefore subject to the MDR in addition to the AI Act.³⁷ Accordingly, a deep-learning model performing genetic analysis qualifies as a medical device within the meaning of Article 2(1) MDR and thus as a high-risk AI system pursuant to Article 6(1) AI Act. Continuous-learning genetic diagnostic systems therefore fall

³⁷ Königshofer, W., *op. cit.* note 2, para. 23; European Commission, *Medical devices: Guidance document – Qualification and Classification of stand-alone software (MEDDEV 2.1/6)*, July 2016, p. 9 *et seq.*, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026. This document is based on Directive 93/42/EEC; Annex II of the Council Resolution of 7 May 1985 on a new approach to technical harmonisation and standardisation (85/C 136/01), OJ C 1985/136, 1.; Schneeberger, D., *op. cit.* note 2, p. 5 *et seq.*

within Annex I, Section A AI Act³⁸ and are subject to the obligations applicable to high-risk AI systems, including requirements relating to risk management, bias mitigation, transparency, and conformity assessment.³⁹

In this regulatory constellation, the AI Act establishes a presumption of conformity for systems that comply with harmonised standards or parts thereof, provided that the references to those standards have been published in the Official Journal of the European Union in accordance with Regulation (EU) No 1025/2012.⁴⁰ Recital 46 AI Act confirms this classification by explicitly referring to the MDR as a Union harmonisation instrument. This confirms a sector-specific regulatory approach for high-risk AI systems listed in Annex I.⁴¹ Consequently, where an AI system falls within the scope of the MDR, it benefits from a rebuttable presumption of conformity with the corresponding requirements of the AI Act, in particular Articles 8 to 15.⁴² Subsumption under a harmonisation act such as the MDR entails that conformity must be demonstrated prior to placing the system on the market, in order to verify compliance with the applicable substantive requirements of that sectoral legislation.⁴³ In particular, Article 40 (1) AI Act—referring to conformity “to the extent that those standards cover those requirements or obligations”—underscores the need for individualised analysis and evaluation.

Chapter III, Section 2, Article 8(1) AI Act provides that the requirements set out in that section must be fulfilled, taking into account both the intended purpose and the state of the art in relation to AI. Article 8(2) further clarifies that providers of high-risk AI systems listed in Annex I, Section A bear responsibility for full compliance with all applicable requirements. Although formulated concisely, this provision establishes far-reaching obligations. The AI Act therefore applies in parallel to sector-specific legislation such as the MDR and complements, rather than replaces, the latter.⁴⁴ However, a separate AI Act risk management system is not required where the MDR already provides an equivalent framework, in which

³⁸ Feiler, L.; Forgó, N., *op. cit.* note 30, p. 11-12.

³⁹ Article 9 of the AI Act requires the implementation of a risk management system; in addition, a regular conformity assessment and associated monitoring as defined in Article 43 of the AI Act is required.

⁴⁰ Feiler, L.; Forgó, N., *op. cit.* note 30, p. 6.6.

⁴¹ Forgó, N.; Prainsack, B., New AI regulation in the EU seeks to reduce risk without assessing public benefit, *Nature Medicine*, 30(5), 2024, p. 1235, doi: 10.1038/s41591-024-02874-2; Feiler, L.; Forgó, N., *op. cit.* note 30, p. 14-15.

⁴² Article 40 AI Act, Feiler, L.; Forgó, N., *op. cit.* note 30, p. 10.

⁴³ Feiler, L.; Forgó, N., *op. cit.* note 30, p. 10.

⁴⁴ Recital 64 AI Act; Feiler, L.; Forgó, N., *op. cit.* note 30, p. 11-12.

case compliance under the MDR may give rise to a presumption of conformity under the AI Act.⁴⁵

Pursuant to Article 16(f) in conjunction with Article 43 AI Act, providers must carry out a conformity assessment prior to placing a high-risk AI system on the market or putting it into service. While Article 42 AI Act establishes a presumption of conformity for compliance with harmonised standards, this presumption does not automatically extend to the conformity assessment procedure itself. Article 43(3) AI Act addresses this interaction by providing that, for high-risk AI systems covered by Union harmonisation acts listed in Annex I, Section A – including the MDR – the conformity assessment procedures under the relevant sectoral legislation apply in lieu of the stand-alone AI Act procedure. However, the substantive requirements of Chapter III, Section 2 AI Act continue to apply. As established above, a continuous deep-learning model used for genetic diagnosis qualifies as stand-alone software and thus as a medical device under Article 2(1) MDR.⁴⁶ Article 2(4) MDR further clarifies that software constitutes an active medical device. Consequently, pursuant to Article 43(3) AI Act, such systems must undergo the conformity assessment procedure established under the MDR as the applicable sectoral harmonisation act.

The presumption of conformity rests on the assumption that the relevant harmonised standards and sectoral conformity procedures sufficiently cover the substantive requirements imposed by the AI Act. This assumption warrants critical scrutiny. The MDR was adopted prior to the emergence of continuously learning AI systems as a central regulatory object and was not specifically designed to address the technical and epistemic characteristics of adaptive deep-learning architectures.⁴⁷ It therefore remains questionable whether conformity assessment

⁴⁵ Recital 22 MDR, which emphasises the importance of harmonised standards for medical devices.

⁴⁶ Article 2(1) first sentence *Regulation (EU) 2017/745* of the European Parliament and of the Council of 5 April 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC, OJ L 2017/117; Königshofer, W., *op. cit.* note 2, para. 23; European Commission, *Medical devices: Guidance document – Qualification and Classification of stand alone software (MEDDEV 2.1/6)*, July 2016, p. 9 *et seq.*, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026. This document is based on Directive 93/42/EEC; Annex II of the Council Resolution of 7 May 1985 on a new approach to technical harmonisation and standardisation (85/C 136/01), OJ C 1985/136, 1; Schneeberger, D., *op. cit.* note 2, p. 5 *et seq.*

⁴⁷ U.S. Food and Drug Administration (FDA), 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, 2019, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial->

under the MDR adequately captures the technical complexities and adaptive characteristics of continuously learning high-risk AI systems. Although Article 43(3) AI Act seeks to integrate the requirements of Chapter III, Section 2 into the sectoral conformity assessment, the structural design of the MDR may not be fully equivalent for systems characterised by ongoing behavioural evolution. Accordingly, this structural tension forms the starting point for the subsequent analysis.

3. MDR CLASSIFICATION: RULE 11 AND INSTABILITY FOR CONTINUOUS LEARNING

As set out in Chapter V MDR, the conformity assessment procedure serves to demonstrate whether a device complies with the requirements laid down in the Regulation.⁴⁸ Before determining the applicable conformity assessment route, a deep-learning system conducting genetic analysis must first be qualified and classified under the MDR. Pursuant to Article 51(1) MDR, medical devices are divided into classes I, IIa, IIb and III, taking into account their intended purpose and the inherent risks associated with their use. The determination of the intended purpose therefore operates as a decisive gatekeeper for the subsequent classification exercise. However, the classification of continuously learning deep-learning systems under Annex VIII MDR raises significant conceptual and practical challenges.

3.1 Duration of Use and the Static Product Paradigm

At an early stage of the classification exercise, Chapter I, Section 1 of Annex VIII MDR introduces a distinction between transient, short-term and long-term use. Duration of use is defined according to whether a device is intended for continuous use of less than 60 minutes, between 60 minutes and 30 days, or for long-term use exceeding 30 days. In the context of a continuously learning deep-learning model, the meaning of "duration of use" becomes ambiguous, because it is unclear whether this notion captures only the moment at which the system performs a specific genetic analysis, or whether it extends to the model's full lifecycle once deployed and continuously updated. Continuous-learning systems do not exhibit clearly bounded "use phases" in the manner presupposed by the MDR. Instead,

Intelligence-and-Machine-Learning-Discussion-Paper.pdf, p. 1, accessed 21 February 2026.

⁴⁸ Recital 4 MDR.

they may remain operational indefinitely, while their internal parameters evolve over time.⁴⁹ This ambiguity reveals an underlying structural assumption of the MDR: First, the Regulation is predicated on the notion of a fixed, self-contained product with a definable beginning and end of its lifecycle. Second, it suggests that the classification framework was primarily designed for stable and static devices, rather than for continuously evolving, data-driven systems whose performance characteristics may change after market placement.

Annex VIII Section 3.6(a) further clarifies that "continuous use" refers to the entire duration of use of the same device, irrespective of temporary interruptions (for example, for cleaning or disinfection). Such examples reflect a paradigm tailored to tangible, physical devices. By contrast, updating, monitoring, or retraining a deep-learning system constitutes a fundamentally different process from cleaning or temporarily removing a physical device. Continuous-learning architectures therefore strain the conceptual categories embedded in the MDR classification framework.

Additional complexity arises from Implementing Rule 3.5(b), which addresses accumulated use where a device is intended to be replaced immediately by another device of the same type. In the context of continuously retrained models, each updated algorithmic version could arguably be interpreted as a functional replacement of the previous iteration.⁵⁰ This raises a doctrinally sensitive question: what constitutes the relevant "device" for classification purposes—the initial approved model architecture, each retrained version, or the evolving system as a whole? The temporal cycle of continuous learning is therefore difficult to reconcile with the static product-based logic underlying Annex VIII.

Moving further, Rule 11 of Annex VIII MDR has been identified in both guidance and literature as the central classification rule for software performing diagnostic or therapeutic functions.⁵¹ Under Rule 11, software intended to provide

⁴⁹ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; for adaptive learning see: Gama *et al.*, *op. cit.* note 14, p. 3-4.

⁵⁰ Medical Device Coordination Group (MDCG), *MDCG 2020-3 Rev.1 – Guidance on significant changes regarding the transitional provision under Article 120 MDR*, 2023, p. 12, https://health.ec.europa.eu/document/download/5fb74e56-fd4c-4dae-9d24-78149edfe56b_en?filename=md_mdcg_2020_3_rev1_en.pdf, accessed 21 February 2026.

⁵¹ Königshofer, W., *op. cit.* note 2, para. 23; European Commission, *Medical devices: Guidance document – Qualification and Classification of stand alone software (MEDDEV 2.1/6)*, July 2016, p. 9 *et seq.*, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026. This document is based on Directive 93/42/EEC; Annex II of the Council Resolution of 7 May 1985 on a new approach to technical harmonisation and standardisation (85/C 136/01), OJ C 1985/136, 1; Schneeberger, D., *op. cit.* note 2, p. 5 *et seq.*

information used to make decisions for diagnostic or therapeutic purposes is classified as class IIa. However, Rule 11 introduces an escalation mechanism: where decisions supported by the software may cause death or irreversible deterioration of a person's state of health, the software is classified as class III. Where such decisions may cause serious deterioration of health or require surgical intervention, the software is classified as class IIb. This differentiation is particularly challenging in the context of AI-based systems, as it raises the question of where the relevant threshold is to be drawn, especially given the distinction between foreseen and unforeseen risks.

In the field of genetic analysis, incorrect assessments—especially if combined with ineffective human oversight—may lead to undetected hereditary diseases, misidentified diagnostic risks, or inappropriate treatment decisions.⁵² Such failures may have severe consequences not only for the patient but also, indirectly, for biological relatives. In these scenarios, it becomes difficult to determine whether the system should be classified as class IIb or class III, and which legal criteria should govern this determination. Rule 11 also provides that software intended to monitor physiological processes is classified as class IIa, except where monitoring of vital physiological parameters may result in immediate danger to the patient, in which case classification shifts to class IIb. All other software falls into class I. This residual category is particularly problematic in the AI context. A deep-learning system may escape higher classification if neither a clear causal link nor a sufficiently concrete likelihood of severe harm can be established. This outcome is reinforced by the MDR's structural assumption of effective human oversight. The classification logic presumes that clinicians retain ultimate decision-making authority and can identify and correct erroneous outputs. In the context of opaque ("black-box") models and continuous learning,⁵³ however, this assumption may function as a normative fiction. If oversight is limited in practice, the risk profile of the system may exceed what is reflected in its formal classification.

An intermediate approach may be to classify the deep-learning module as class Iia, where the likelihood of death, irreversible deterioration⁵⁴, or serious deterioration of a person's state of health⁵⁵ cannot be sufficiently established at the time of assessment. Yet this solution forces a continuously evolving system into a fixed regulatory category. This tension becomes particularly evident when considering the technical dynamics of adaptive models. Continuous-learning

⁵² Zhang *et al.*, *op. cit.* note 13, p. 1343-1344; Wang, X. *et al.*, *op. cit.* note 13, p. 796–803.

⁵³ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for adaptive learning see: Gama *et al.*, *op. cit.* note 14, p. 3-4, p. 30.

⁵⁴ Annex VIII, Chapter III, Rule 11, Sentence 1 MDR.

⁵⁵ Annex VIII, Chapter III, Rule 11, Sentence 2 MDR.

systems may change their performance characteristics over time. As new data are incorporated and internal parameters are updated, the system's predictive accuracy, error patterns, and clinical impact may shift.⁵⁶ Consequently, a model initially classified as class IIa could, in practice, begin to support decisions that entail a higher risk profile, potentially falling within the scope of class IIb or even class III.

Rule 11 provides no explicit mechanism to accommodate such dynamic risk evolution. Classification is performed *ex ante*, based on the intended purpose⁵⁷ and foreseeable risks⁵⁸ at the time of conformity assessment.⁵⁹ It does not contain a structured framework for reassessing classification where risk profiles evolve due to adaptive learning. Consequently, classification under Rule 11 is conducted on a case-by-case basis, often relying on interpretative guidance and the assessment of notified bodies, without clearly defined criteria for continuous-learning systems. The wording "may cause"⁶⁰ reinforces this tension. It presupposes a relatively stable and foreseeable risk scenario, whereas continuously learning systems operate within dynamic and probabilistic environments.⁶¹ The combination of *ex ante* classification logic and evolving model behaviour creates a structural mismatch between regulatory design and technological reality.

In sum, the classification regime under Annex VIII—particularly Rule 11—reflects a stability-oriented regulatory paradigm. When applied to continuously learning deep-learning systems in genetic diagnostics, this paradigm produces instability, uncertainty, and potential under- or over-classification. The static classification logic of the MDR therefore, sits uneasily with adaptive AI technologies whose risk profiles may evolve after market placement.

4. MDR CONFORMITY ASSESSMENT IS STRUCTURALLY EX ANTE

The analysis above demonstrates that the foundation for conformity assessment procedures and routes provided under Article 52 MDR and Annexes IX to XI are inherently *ex ante* in nature. Their design is explicitly oriented toward the period

⁵⁶ Gama *et al.*, *op. cit.* note 14, p. 3-4; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

⁵⁷ Article 2(12) MDR; Annex VIII, Chapter II, Section 3.1 MDR

⁵⁸ Annex I, Chapter I, Section 8 MDR.

⁵⁹ Annex IX-XI MDR.

⁶⁰ Annex VIII, Chapter III, Rule 11, Sentence 1 MDR.

⁶¹ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

prior to placing the device on the market or putting it into service. This temporal orientation reflects an implicit structural assumption: that the device assessed during conformity evaluation will remain materially unchanged thereafter.

This assumption entails that the medical device is characterised by a defined technical configuration, a fixed and declared intended purpose, and a stable risk profile.⁶² It further presupposes that no additional functionalities, risk dynamics, or inferential capacities emerge beyond those foreseeable at the time of assessment. In particular, it assumes that system behaviour will remain within the scope of what was documented, tested, and certified, thereby excluding the dynamic effects of continuous learning, feedback loops⁶³ or post-deployment adaptation. Feedback loops may alter the distribution of training data, resulting in parameter updates and shifts in decision boundaries,⁶⁴ which in turn may lead to different predictive outputs over time. In genetic diagnostic contexts, such shifts may surface previously non-salient probabilistic associations, including risk correlations that were not reflected in the originally assessed configuration and that may have implications for biological relatives.

By contrast, the technical architecture of continuously learning deep-learning systems departs fundamentally from this stability paradigm. Not all AI- and machine-learning-based software medical devices are "locked".⁶⁵ In this context, "locked" describes an algorithm that provides the same result each time the same input is provided, whereas a continuous learning algorithm adapts its behaviour using a specified learning process. Whereas locked algorithms produce identical outputs for identical inputs,⁶⁶ continuously learning systems are capable of modifying internal parameters during operation. Through retraining, dataset updates, or parameter optimisation, such systems may alter their performance characteristics over time.⁶⁷ Their defining feature is precisely the capacity for

⁶² See, e.g., Art 2(12) MDR.

⁶³ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley *et al.*, *op. cit.* note 16, p. 2506.

⁶⁴ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

⁶⁵ U.S. Food and Drug Administration (FDA), 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, 2019, p. 3, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>, accessed 21 February 2026.

⁶⁶ *Ibid.*, p. 3.

⁶⁷ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

behavioural adaptation based on real-world use.⁶⁸ It can be argued that regulatory frameworks originally developed for physical or functionally stable devices did not anticipate such ongoing model evolution.⁶⁹ The core tension therefore, lies in a conformity assessment that evaluates a specific historical model state, while the operational system continues to evolve after certification. Conformity assessment is temporally anchored; system behaviour is temporally dynamic.

The MDR was not originally conceived with continuously learning software as a central regulatory object.⁷⁰ Although the Regulation formally encompasses software, its structural architecture reflects a product-based paradigm centred on identifiable configurations and foreseeable risks. It remains, therefore, ambiguous whether the MDR's obligations are sufficiently tailored to address continuously adaptive deep-learning systems conducting genetic analysis, particularly where changes in data environments and post-deployment updating can alter outputs, which may facilitate secondary (relational) inferences concerning biological relatives.⁷¹ Having identified the *ex ante* character of conformity assessment under the MDR, it becomes necessary to examine the substantive mechanisms through which conformity is verified. In functional terms, conformity assessment under the MDR rests on three core components: (1) auditing the quality management system (hereinafter: QMS), (2) assessing technical documentation, and (3)

⁶⁸ U.S. Food and Drug Administration (FDA), 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, 2019, p. 1, 3-4, 7, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>, accessed 21 February 2026.

⁶⁹ Ludvigsen, K.R.; Nagaraja, S.; Daly, A., The Intention; Requirements for Software as a Medical Device in EU Law, SSRN Paper, 2021, p. 1; see further: U.S. Food and Drug Administration (FDA), 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, 2019, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>, p. 2, accessed 21 February 2026.

⁷⁰ Ludvigsen, K.R.; Nagaraja, S.; Daly, A., *op. cit.* note 69, p. 1; see further: U.S. Food and Drug Administration (FDA), 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, 2019, p. 2, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>, accessed 21 February 2026.

⁷¹ Erlich, Y.; Narayanan, A., Routes for Breaching and Protecting Genetic Privacy, *Nature Reviews Genetics*, 15, 2014, p. 417. doi:10.1038/nrg3723; on system evolution and feedback loops see: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; see also: Spiliopoulou A. *et al.*, Genomic Prediction of Complex Human Traits: Relatedness, Trait Architecture and Predictive Meta-Models, *Human Molecular Genetics*, 24(14), 2015, p. 4168-4169, doi:10.1093/hmg/ddv145.

performing verification and validation testing.⁷² The QMS framework requires documented procedures and predefined control mechanisms, with technical documentation capturing the device's characteristics at a specific point in time. Verification and validation testing assess the performance of a defined system configuration before market placement. This assumes that the tested and documented configuration remains representative of the device's ongoing operational behaviour. Consequently, the conformity assessment regime assumes stability of device identity over time.

For continuously learning systems, however, this assumption becomes fragile. As internal parameters evolve and performance characteristics shift, the relationship between the originally certified configuration and the deployed system may weaken. A deep-learning module initially assessed under a particular risk classification may, over time, exhibit altered error patterns or enhanced inferential capacities that were neither fully foreseeable nor fully representable at the time of assessment. This structural tension is further amplified by the technical complexity of genetic analysis systems. In the context of genetic analysis, deep-learning systems may facilitate indirect or relational inferences concerning biological relatives.⁷³ Where such systems undergo retraining or adaptation, the scope or precision of such inferences may shift over time⁷⁴ or enable emergent correlations not explicitly programmed.⁷⁵ As a result, manufacturers may seek to limit the scope of outputs through predefined design constraints or specification of intended use.⁷⁶ However, such modifications may influence performance metrics and error distributions, thereby affecting the system's effective risk profile.⁷⁷ These dynamics sit uneasily with a conformity logic that relies on *ex ante* assumptions regarding device behaviour and foreseeable risks, while largely excluding systematic reassessment of evolving characteristics during deployment. This results in a temporal asymmetry; while conformity assessment is completed at a fixed moment, the risk evolution may continue throughout the lifecycle of the device. In this respect, the *ex ante* evidentiary structure of the MDR does not result

⁷² Stange, A.F.; Julian, E., *Conformity Assessment of Medical Devices*, *Journal of Market Access & Health Policy*, 14(1):4, 2026, p. 2. doi:10.3390/jmahp14010004.

⁷³ Erlich, Y.; Narayanan, A., *op. cit.* note 71, p. 417; on system evolution and feedback loops see: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; see also: Spiliopoulou A. *et al.*, *op. cit.* note 71, p. 4168-4169.

⁷⁴ Erlich, Y.; Narayanan, A., *op. cit.* note 71, p. 417; on system evolution and feedback loops see: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; see also: Spiliopoulou A. *et al.*, *op. cit.* note 71, p. 4168-4169.

⁷⁵ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

⁷⁶ See: Article 2(12), Annex II MDR.

⁷⁷ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

in an enforcement challenge; rather, it reveals a design limitation when applied to continuously evolving AI systems. Consequently, even with full formal compliance, the conformity assessment may fail to capture behavioural developments that occur post-certification. In summary, the structural architecture of the MDR presupposes stability, whereas continuous-learning systems are structured by evolution.

4.1 Annex IX MDR: Quality Management System and Controlled Change

Annex IX MDR establishes a conformity assessment procedure centred on a quality management system combined with technical documentation, thereby grounding conformity assessment in a predominantly process-oriented framework designed to preserve effectiveness throughout the device's lifecycle.⁷⁸ Annex IX constitutes the standard conformity assessment route for devices classified as class IIa, IIb, and III. Chapter I, Section 2 specifies the QMS requirements, complemented by the technical documentation obligations in Chapter II. The QMS framework itself is structurally divided into an *ex ante* component (Section 2) and a post-market surveillance and monitoring component (Section 3).

Proportionate to the applicable risk class and the nature of the device,⁷⁹ this architecture reflects a regulatory model premised on predefined system controls and documented procedures. It assumes that compliance can be secured through organisational safeguards, documented responsibilities, traceability mechanisms, and monitored corrective processes.⁸⁰ The focus lies on process integrity and control structures rather than on continuous substantive reassessment of evolving outputs or adaptive learning dynamics. Although Annex IX incorporates post-market surveillance, its structure channels post-certification changes primarily through documentation, monitoring, and — where necessary — renewed conformity assessment. It does not establish a mechanism for ongoing substantive re-evaluation of device behaviour as such. The regulatory emphasis remains procedural rather than behavioural.

The MDR recognises that medical devices may undergo design modifications after conformity assessment, compliance may be preserved where the manufacturer operates a QMS capable of managing such modifications within appropriate frames and ensuring continued conformity with the applicable procedure.⁸¹ In this

⁷⁸ Stange, A.F.; Julian, E., *op. cit.* note 72, p. 3.

⁷⁹ Recital 32 MDR.

⁸⁰ Annex IX, Chapter I, Section 2.1 and 2.2 MDR.

⁸¹ Article 10(9) MDR.

context, software devices may incorporate recalibration or reconfiguration functionalities, provided that such modifications occur within predefined boundaries that were covered by the initial conformity assessment.⁸² The QMS therefore operates as a governance structure for controlled change, not as a substitute for substantive reassessment. Crucially, these modifications must remain within the originally declared intended purpose prior to market placement,⁸³ which entails manufacturers are prohibited from recommending uses deviating from the intended purpose underlying the conformity assessment.⁸⁴ Furthermore, planned substantial changes must be communicated to the notified body.⁸⁵ The conformity architecture thus presupposes that relevant modifications are identifiable, intentional, and assessable within a controlled change framework.

This logic aligns comfortably with "locked" AI systems, which generate identical outputs for identical inputs.⁸⁶ It also accommodates adaptive systems operating within predefined change envelopes covered by the original conformity assessment.⁸⁷ Therefore, it presupposes that adaptive capacity is ex ante limited, foreseeable, and governable through predefined tolerance ranges.

However, deep-learning architectures may, by virtue of their technological design, evolve beyond such predefined boundaries. As a consequence of their adaptive capabilities, these systems may exhibit behaviour that exceeds the limits initially assessed and thereby fall outside the scope of the original conformity assessment. Accordingly, deep-learning modules challenge the regulatory assumptions underlying the MDR, which are premised on behavioural stability and ex ante foreseeable developments. In this respect, Annex IX does not provide a dedicated mechanism to accommodate continuous behavioural evolution within the same conformity framework. Its regulatory logic permits adaptability only insofar as it remains ex ante bounded and controllable through predetermined organisational

⁸² Predefined modification during runtime refers to adjustments, optimisation, or recalibration of a medical software device within these predefined boundaries, in relation to an individual patient or a specific healthcare environment. See: COCIR – European Coordination Committee of the Radiological, Electromedical and Healthcare IT Industry, *Artificial intelligence in EU medical device legislation*, 2021, p. 10, https://cocir.org/uploads/2021/10/COCIR_AI_in_EU_Medical_Device_Legislation_2021.pdf, accessed 21 February 2026.

⁸³ *Ibid.*, p. 9.

⁸⁴ Article 7(d) MDR.

⁸⁵ Annex IX, Chapter I, Section 2.4 MDR.

⁸⁶ COCIR – European Coordination Committee of the Radiological, Electromedical and Healthcare IT Industry, *Artificial intelligence in EU medical device legislation*, 2021, p. 10, https://cocir.org/uploads/2021/10/COCIR_AI_in_EU_Medical_Device_Legislation_2021.pdf, accessed 21 February 2026.

⁸⁷ *Ibid.*, p. 12.

safeguards.⁸⁸ Once adaptive behaviour extends beyond those predefined limits, the architecture reverts to the dichotomy of either triggering a new conformity assessment or operating outside the assessed scope.

This reveals the structural tension underlying Annex IX: procedurally, it is comprehensive; substantively, it presupposes behavioural stability. The framework offers extensive documentation, monitoring, and quality control mechanisms. Yet these mechanisms remain anchored in the assumption that relevant behavioural change is either planned or foreseeable within predefined boundaries.




Figure 1: Input-output-based illustration of algorithmic change scenarios in AI-based genetic diagnostic systems under Annex IX MDR. The figure contrasts locked AI, adaptive AI operating within predefined change boundaries covered by conformity assessment, and deep-learning systems evolving beyond the assessed scope, highlighting the regulatory implications of behavioural evolution at output level.⁸⁹

⁸⁸ COCIR – European Coordination Committee of the Radiological, Electromedical and Healthcare IT Industry, *Artificial intelligence in EU medical device legislation*, 2021, p. 12, https://cocir.org/uploads/2021/10/COCIR_AI_in_EU_Medical_Device_Legislation_2021.pdf, accessed 21 February 2026.

⁸⁹ The conceptual design of this figure was developed by the author with the assistance of an AI-based language model (ChatGPT) and subsequently refined for academic use. The substantive content, legal analysis, and interpretation remain the sole responsibility of the author.

For continuously learning systems in genetic diagnostics, where model behaviour may evolve in response to real-world data and generate new inferential patterns, this assumption becomes increasingly fragile. The regulatory completeness of Annex IX at the procedural level thus coexists with a substantive blind spot at the behavioural level.

4.2. Technical documentation and testing under the MDR

Manufacturers of medical devices are required to establish technical documentation for the purpose of demonstrating conformity with the requirements of the MDR. Annex II MDR sets out the substantive requirements for such documentation and primarily operates as an *ex ante* regulatory instrument, whereas Annex III complements this framework by governing post-market surveillance documentation.⁹⁰

The structure of technical documentation under Annex II reflects its function as the substantive evidentiary foundation of conformity assessment. Conformity is verified through a documentation-based assessment of the device's technical characteristics, performance claims, and risk profile at the time of submission. By its very nature, this assessment captures the device at a defined point in time. Technical documentation, therefore reflects a regulatory snapshot rather than the ongoing behavioural development of a system during operation.

For continuously learning deep-learning systems used in genetic diagnostics, this snapshot logic becomes problematic. Where system behaviour evolves through retraining, feedback loops, or changing data environments,⁹¹ a documentation-based assessment may not fully represent the device's future operational characteristics. The evidentiary model presupposes temporal stability; adaptive architectures introduce temporal dynamism. Annex II further requires that technical documentation be presented in a "clear, organised, readily searchable and unambiguous manner" and accurately reflect the characteristics of the medical software device at the time of submission. Central to this documentation are the device description and specification, including the intended purpose and intended users.⁹² These elements function as fixed regulatory reference points and are defined *ex ante* by the manufacturer within the conformity assessment dossier.

⁹⁰ Article 10(4) MDR.

⁹¹ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

⁹² Annex II, Section 1.1(a) MDR.

This ex ante fixation becomes structurally fragile where feedback-driven updates⁹³ or evolving deployment conditions⁹⁴ alter model behaviour over time. In systems processing inherently relational genetic data, behavioural evolution may affect the scope of inferential outputs, potentially extending beyond the initially declared intended purpose and indirectly implicating biological relatives. Outputs may generate secondary inferences concerning biological relatives⁹⁵ and thereby expand the practical scope of the device's impact, even in the absence of formal modifications to the model architecture or algorithm.⁹⁶ The conformity assessment remains formally intact; yet the behavioural reach of the system may expand beyond what was initially documented. The result is not necessarily non-compliance, but rather a structural risk of regulatory incompleteness after submission of the technical documentation⁹⁷ While the conformity framework presupposes that the documented configuration remains representative of the deployed system, continuous learning challenges that presupposition.

The ex ante logic is further reinforced in Annex II, Point 6, which requires the inclusion of "results and critical analyses of all verifications and validation tests and/or studies undertaken to demonstrate conformity of the device with the requirements of the MDR." Verification and validation are conducted on a defined system configuration prior to market placement. In continuously learning systems, however, operational behaviour may evolve post-deployment through parameter updates or changing data distributions, even where the underlying architecture remains unchanged.⁹⁸

Testing and validation under the MDR therefore focus on the demonstrated performance of a particular model state, rather than on the adaptive trajectory of the system over time. The regulatory assumption underlying this framework is that the validated configuration remains sufficiently representative of the device's

⁹³ Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

⁹⁴ Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

⁹⁵ Erlich, Y.; Narayanan, A., *op. cit.* note 71, p. 417; Spiliopoulou A. *et al.*, *op. cit.* note 71, p. 4168-4169; on system evolution and feedback loops see: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

⁹⁶ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

⁹⁷ Additionally, the feedback loops and model performance may improve unevenly among different sub-populations groups. For example, see: Zhang *et al.*, *op. cit.* note 13, p. 1331, 1342; Seyyed-Kalantari, L. *et al.*, CheXclusion: Fairness Gaps in Deep Chest X-Ray Classifiers, Pacific Symposium Biocomputing, 26, 2021, p. 232-243. However, this thesis does not focus on bias and ethical issues in this context, therefore it remains mentioned but not analysed in detail due to the limitation of the objective of the paper.

⁹⁸ Gama *et al.*, *op. cit.* note 14, p. 4.

future operational behaviour. However continuous-learning architectures render that assumption contestable.

In this respect, the conformity architecture of the MDR does not merely impose extensive documentation obligations; it embodies a specific epistemic model of regulatory control. That model is oriented toward identifiable, stable configurations capable of being documented, tested, and certified. Where system behaviour evolves beyond that configuration through ongoing learning processes, the correspondence between certified documentation and operational reality may progressively weaken.

4.3 Annex X MDR: Type Examination and Static Device Assumptions

Annex X establishes the EU type-examination procedure, under which a notified body assesses and certifies that a medical device complies with the applicable provisions of the MDR. In this procedure, the notified body examines the technical documentation, evaluates relevant lifecycle processes, and requires a representative sample of the envisaged device production. Conformity is established through the assessment of a representative specimen, which presupposes that the essential characteristics of the device — including its performance parameters — remain sufficiently stable over time. This reflects a type-examination logic that is structurally *ex ante*. The assessment framework under Annex X presupposes a medical device whose relevant characteristics are sufficiently stable at the time of market placement to allow certification on the basis of a defined set of technical parameters.⁹⁹ Conformity assessment is therefore conducted at a specific moment in time, with the implicit expectation that this evaluation remains representative of the device's subsequent performance.

In the context of continuously learning software systems, this expectation becomes fragile again, because behavioural parameters may evolve during deployment through retraining, dataset shifts, or internal parameter updates,¹⁰⁰ potentially altering performance characteristics after certification. Annex X does not expressly provide a mechanism to evaluate prospective behavioural evolution at the time of assessment. The notified body verifies conformity on the basis of documented and tested characteristics of the submitted configuration,¹⁰¹ rather

⁹⁹ Article 56, Annex II, Annex X Section 4 MDR.

¹⁰⁰ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; for feedback loops: Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

¹⁰¹ Article 52 MDR, Annex IX, Annex II MDR.

than on the basis of dynamically unfolding operational behaviour. As a result, a temporal asymmetry arises between the certification snapshot and the subsequent evolution of the system during real-world use. The rigidity of the type-examination model is further reinforced by Annex II, Section 5, which governs changes to the approved type. Manufacturers must inform the notified body of any planned change to the approved medical device or to its intended purpose and conditions of use.¹⁰² The emphasis on "planned" modifications reflects an assumption that relevant changes are identifiable, intentional, and temporally bounded. Continuously learning systems challenge this assumption, because adaptive behaviour may expand functional scope, alter inferential capacities, or modify performance characteristics without a discrete, intentionally planned intervention by the manufacturer. Where such evolution affects the intended purpose or conditions of use, renewed approval may be required. Continuous learning could therefore repeatedly trigger the obligation to conduct a new conformity assessment.

Annex X thus reinforces the stability-oriented paradigm underlying the MDR. Type-examination is designed for devices whose essential characteristics can be fixed, documented, and certified as a stable reference configuration. Consequently, continuously learning architectures destabilise the notion of a single identifiable "type", thereby exposing the structural limits of the type-examination model.

4.4 Annex XI MDR: Product Conformity Verification and Output-Oriented Control

Conformity assessment under Annex XI constitutes the third conformity assessment pathway available under Article 52 MDR. Annex XI operates as a verification mechanism and builds upon prior conformity assessment steps, particularly the European Union type-examination (hereinafter: EU type-examination) under Annex X.¹⁰³ In line with Annexes IX and X, Annex XI presupposes a previously defined and approved device type derived from the prior EU type-examination. Part A of Annex XI focuses on production quality assurance through the implementation and surveillance of an approved quality management system. Verification of such a system presupposes the continued implementation of an established compliance structure. It does not, however, entail an independent reassessment of evolving behavioural characteristics inherent in a continuously learning software device. The conformity logic remains anchored in the approved

¹⁰² Annex II, Section 5.1 MDR.

¹⁰³ Article 52 MDR; Annex XI, Section 1 MDR.

type and the manufacturer's obligation to ensure that production consistently corresponds to that certified configuration. The manufacturer's obligation to certify and declare conformity with the type described in the EU type-examination certificate presupposes a stable and identifiable device identity. This assumption becomes fragile where behavioural evolution occurs in a continuously learning module, as the operational system may progressively diverge from the originally approved type without a discrete, formally identifiable modification event. Part B establishes a product verification procedure requiring the examination of each manufactured device to ensure and declare its conformity with the type described in the EU type-examination certificate and with the applicable requirements of the MDR.¹⁰⁴ This examination-based model presupposes identifiable units characterised by defined parameters, stable testing conditions, and reproducible characteristics. In the context of physical devices, this unit-by-unit verification logic is conceptually coherent; in contrast, centrally deployed or cloud-based software architectures may involve ongoing model update cycles and continuously learning features at the level of neural network parameters.¹⁰⁵ Parameter adjustments, retraining processes, and evolving data distributions may lead to version modifications and shifting operational states over time.¹⁰⁶ Such evolution challenges the representativeness of the examined configuration and raises the question of which model state constitutes the relevant "device" for purposes of verification.

Continuous learning architectures are characterised by parameter updates and evolving model states across training and update cycles.¹⁰⁷ Consequently, continuously learning systems introduce evolving parameter states during operation. This creates a structural tension between the unit-by-unit verification logic embedded in Annex XI and the fluid, model-state-based nature of adaptive AI systems. Where the certified configuration is no longer identical to the deployed configuration, the conceptual basis of product verification — namely, conformity with a stable approved type — becomes increasingly difficult to maintain.

Annex XI therefore, does not depart from the stability-oriented paradigm underlying Annex X; rather, it reinforces it. Product verification and production quality assurance remain anchored in the assumption of a fixed and reproducible device identity. For continuously learning genetic diagnostic systems, this

¹⁰⁴ Annex XI, Part B, Section 11 MDR.

¹⁰⁵ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325; Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

¹⁰⁶ Sculley, D. *et al.*, *op. cit.* note 16, p. 2506; Zhang *et al.*, *op. cit.* note 13, p. 1333-1335.

¹⁰⁷ Goodfellow, I.; Bengio, Y.; Courville, A., *op. cit.* note 17, p. 164-325.

assumption exposes a structural mismatch between the conformity assessment architecture and the adaptive reality of post-deployment learning.

5. CONCLUSION

This article examined the presumption of conformity for medical AI under the AI Act through the lens of continuously learning deep-learning systems in genetic diagnostics. Because diagnostic DL software qualifies as a medical device under the MDR and therefore as a high-risk AI system under Article 6(1) AI Act and Annex I, Section A, the AI Act's enforcement architecture for such systems is structurally intertwined with sectoral medical device conformity assessment. Article 43(3) AI Act channels conformity assessment for Annex I systems through the applicable Union harmonisation act—in this case, the MDR—while maintaining the continued applicability of the substantive requirements in Title III, Chapter III, Section 2 AI Act. The practical significance of this design lies in the presumption of conformity: it implicitly assumes that sectoral conformity procedures, in combination with harmonised standards, can function as an equivalent and sufficiently comprehensive mechanism for demonstrating compliance with AI Act requirements.

The analysis shows that this equivalence assumption becomes contestable for continuous-learning genetic diagnostic systems. A first structural tension emerges at the level of MDR classification. Annex VIII and, in particular, Rule 11 are built around an *ex ante* assessment of intended purpose and foreseeable risks. Yet continuously learning systems are characterised by adaptive behaviour, probabilistic inference, and the potential for evolving performance characteristics after deployment. The phrase "may cause" in Rule 11 is calibrated to stable risk scenarios; it provides limited doctrinal traction for capturing dynamic risk evolution driven by retraining, dataset shifts,¹⁰⁸ or feedback loops.¹⁰⁹ As a consequence, classification can become unstable across time: a system may be assessed under one risk profile at market entry while exhibiting materially different error patterns or clinical impacts during real-world use. This instability matters because classification functions as a gateway to the applicable conformity assessment route and the intensity of regulatory controls.

Second, more fundamental tension concerns the temporal architecture of MDR conformity assessment. Across the routes provided under Article 52 MDR and

¹⁰⁸ Zhang *et al.*, *op. cit.* note 13, p. 1343-1345.

¹⁰⁹ Sculley, D. *et al.*, *op. cit.* note 16, p. 2506.

Annexes IX to XI, conformity assessment remains anchored in a snapshot logic. Technical documentation under Annex II is designed to evidence conformity at the time of submission; verification and validation tests relate to a defined system configuration; and quality management systems govern compliance through documented procedures and controlled change processes. These mechanisms presuppose that the certified configuration remains representative of the device's operational behaviour. By contrast, continuous-learning systems are defined by the possibility of model-state evolution after market placement. In genetic diagnostics, this divergence is amplified by the relational nature of genomic data and by the potential for changing inferential patterns over time. Even where the declared intended purpose remains formally unchanged, the practical scope and impact of a system's outputs may evolve, including through the emergence of secondary inferences that were not observable at the time of assessment.

Third, the comparative analysis of conformity assessment routes underscores that procedural completeness does not resolve the substantive blind spot for adaptive behaviour. Annex IX offers a process-oriented lifecycle framework and channels post-market changes through documentation, monitoring, and, where required, renewed conformity assessment. However, it does not constitute a mechanism for continuous substantive reassessment of evolving model behaviour as such. The regulatory logic is designed for manufacturer-initiated and ex ante bounded modifications; autonomous behavioural evolution sits uneasily within this paradigm. Annex X intensifies the static premise by tying compliance to an EU type-examination of a representative specimen, with an implicit expectation that certified characteristics remain stable. Annex XI, in turn, reinforces reliance on the certified type and, particularly in Part B, relies on unit-by-unit verification logic that is conceptually coherent for tangible devices but strained for centrally deployed software architectures characterised by ongoing update cycles and shifting operational states. Together, these routes reflect a coherent MDR conformity architecture—yet one that is structurally premised on stable device identity and point-in-time verification.

These findings have direct implications for the presumption of conformity under the AI Act. The presumption is functional only to the extent that harmonised standards and sectoral conformity procedures actually cover the relevant AI Act requirements for the system in question. For continuous-learning genetic diagnostics, the MDR's ex ante conformity logic risks producing a temporal and epistemic misalignment: extensive procedural obligations are maintained, yet the framework may not adequately capture the evolution of behaviour and risk that the AI Act seeks to govern through lifecycle-oriented requirements, including change management and post-deployment control of performance-relevant

modifications. This does not imply that the MDR is ineffective as product safety regulation, nor that continuous-learning systems are necessarily incompatible with the MDR. Rather, it clarifies that, in this domain, MDR conformity assessment may not be fully equivalent as an evidentiary basis for AI Act compliance where system behaviour is expected to evolve beyond what can be represented in a certification snapshot. This is particularly relevant because the presumption of conformity operates not merely as a procedural simplification, but as a mechanism that reallocates evidentiary confidence from direct regulatory scrutiny to standardised and sectoral assessment structures. Where those structures are not designed to observe or reassess behavioural evolution over time, the presumption risks becoming under-inclusive in relation to the very risks that distinguish adaptive AI systems from conventional medical software.

A forward-looking implication is that legal certainty will depend on clarifying how continuous learning is operationalised within the AI Act’s “substantial modification” logic and how this interacts with MDR change governance. Article 43(4) AI Act distinguishes between performance-relevant changes that require renewed conformity assessment and changes already accounted for within the original assessment. For genetic diagnostic software, the key regulatory task is therefore to specify what it means, in practice, to “take into account” adaptive behaviour *ex ante*, and how to define and evidence permissible change envelopes over time. While MDCG guidance and harmonised standards can provide procedural support, the central challenge is conceptual: aligning a stability-oriented conformity model with systems whose defining feature is behavioural evolution. Addressing this alignment is a precondition for ensuring that the presumption of conformity remains credible—both as a legal shortcut and as a substantive guarantee of safety and compliance for continuously learning medical AI.

REFERENCES

BOOKS AND ARTICLES

1. Erlich, Y.; Narayanan, A., Routes for Breaching and Protecting Genetic Privacy, *Nature Reviews Genetics*, 15, 2014, 409–421, doi:10.1038/nrg3723.
2. Feiler, L.; Forgó, N., *Kommentar KI-VO, EU-Verordnung über künstliche Intelligenz*, Verlag Österreich, Vienna 2024.
3. Forgó, N.; Prainsack, B., New AI regulation in the EU seeks to reduce risk without assessing public benefit, *Nature Medicine*, 30(5), 2024, 1235–1237, doi: 10.1038/s41591-024-02874-2.

4. Gama, J. *et al.*, A Survey on Concept Drift Adaptation, *ACM Computing Surveys*, 46(4), 2014, 1–37, doi:10.1145/2523813.
5. Goodfellow, I.; Bengio, Y.; Courville, A., *Deep Learning*, MIT Press, Cambridge (MA) 2016.
6. Herbst, T., Regulierung intelligenter Medizinprodukte, *ZfPC*, 2023.
7. Königshofer, W., Medical Device Law, in: Resch, R.; Wallner, F. (eds.), *Handbook of Medical Law*, 3rd ed., Manz, Vienna 2020, paras 137–151.
8. Ludvigsen, K.R.; Nagaraja, S.; Daly, A., The Intention; Requirements for Software as a Medical Device in EU Law, *SSRN Paper*, 2021, 1–38.
9. Schneeberger, D., Intelligente Medizinprodukte: Rechtsfragen am Schnittpunkt von DSGVO, MPVO und AI Act, *DaKo*, 1/2024, 4–7.
10. Sculley, D. *et al.*, Hidden Technical Debt in Machine Learning Systems, *Advances in Neural Information Processing Systems*, 2015, 2503–2511.
11. Seyyed-Kalantari, L. *et al.*, CheXclusion: Fairness Gaps in Deep Chest X-Ray Classifiers, *Pacific Symposium Biocomputing*, 26, 2021, 232–243.
12. Spiliopoulou, A. *et al.*, Genomic Prediction of Complex Human Traits: Relatedness, Trait Architecture and Predictive Meta-Models, *Human Molecular Genetics*, 24(14), 2015, 4167–4182, doi:10.1093/hmg/ddv145.
13. Stange, A.F.; Julian, E., Conformity Assessment of Medical Devices, *Journal of Market Access & Health Policy*, 14(1):4, 2026, doi:10.3390/jmahp14010004.
14. Wang, X. *et al.*, Inconsistent Performance of Deep Learning Models on Mammogram Classification, *Journal of the American College of Radiology*, 17(6), 2020, 796–803, doi:10.1016/j.jacr.2020.01.006.
15. Zhang, A. *et al.*, Shifting Machine Learning for Healthcare from Development to Deployment and from Models to Data, *Nature Biomedical Engineering*, 6(12), 2022, 1330–1345, doi:10.1038/s41551-022-00898-y.

LEGAL RULES

1. European Commission, Commission Notice – The Blue Guide on the implementation of EU product rules 2022, OJ C 247, 29 June 2022.
2. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689.
3. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation), OJ L 2016/119.

4. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on medical devices (Medical Device Regulation), OJ L 2017/117.
5. Genetic Technology Act (Gentechnikgesetz – GTG), Federal Law Gazette (BGBl) No 510/1994, as amended by BGBl I No 8/2022.

NETWORK RESOURCES

1. COCIR – European Coordination Committee of the Radiological, Electromedical and Healthcare IT Industry, *Artificial intelligence in EU medical device legislation*, 2021, https://cocir.org/uploads/2021/10/COCIR_AI_in_EU_Medical_Device_Legislation_2021.pdf, accessed 21 February 2026.
2. European Commission, AI Act enters into force, 1 August 2024, https://commission.europa.eu/news-and-media/news/ai-act-enters-force-2024-08-01_en, accessed 21 February 2026.
3. European Commission, *Medical devices: Guidance document – Qualification and Classification of stand alone software (MEDDEV 2.1/6)*, July 2016, <https://ec.europa.eu/docsroom/documents/17921>, accessed 21 February 2026.
4. Medical Device Coordination Group (MDCG), *MDCG 2020-3 Rev.1 – Guidance on significant changes regarding the transitional provision under Article 120 MDR*, 2023, https://health.ec.europa.eu/document/download/5fb74e56-fd4c-4dae-9d24-78149edfe56b_en?filename=md_mdcg_2020_3_rev1_en.pdf, accessed 21 February 2026.
5. U.S. Food and Drug Administration (FDA), 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, 2019, <https://www.fda.gov/files/medical%20devices/published/US-FDA-Artificial-Intelligence-and-Machine-Learning-Discussion-Paper.pdf>, accessed 21 February 2026.

TABLE OF FIGURES

Figure 1: Input–output-based illustration of algorithmic change scenarios in AI-based genetic diagnostic systems under Annex IX MDR.