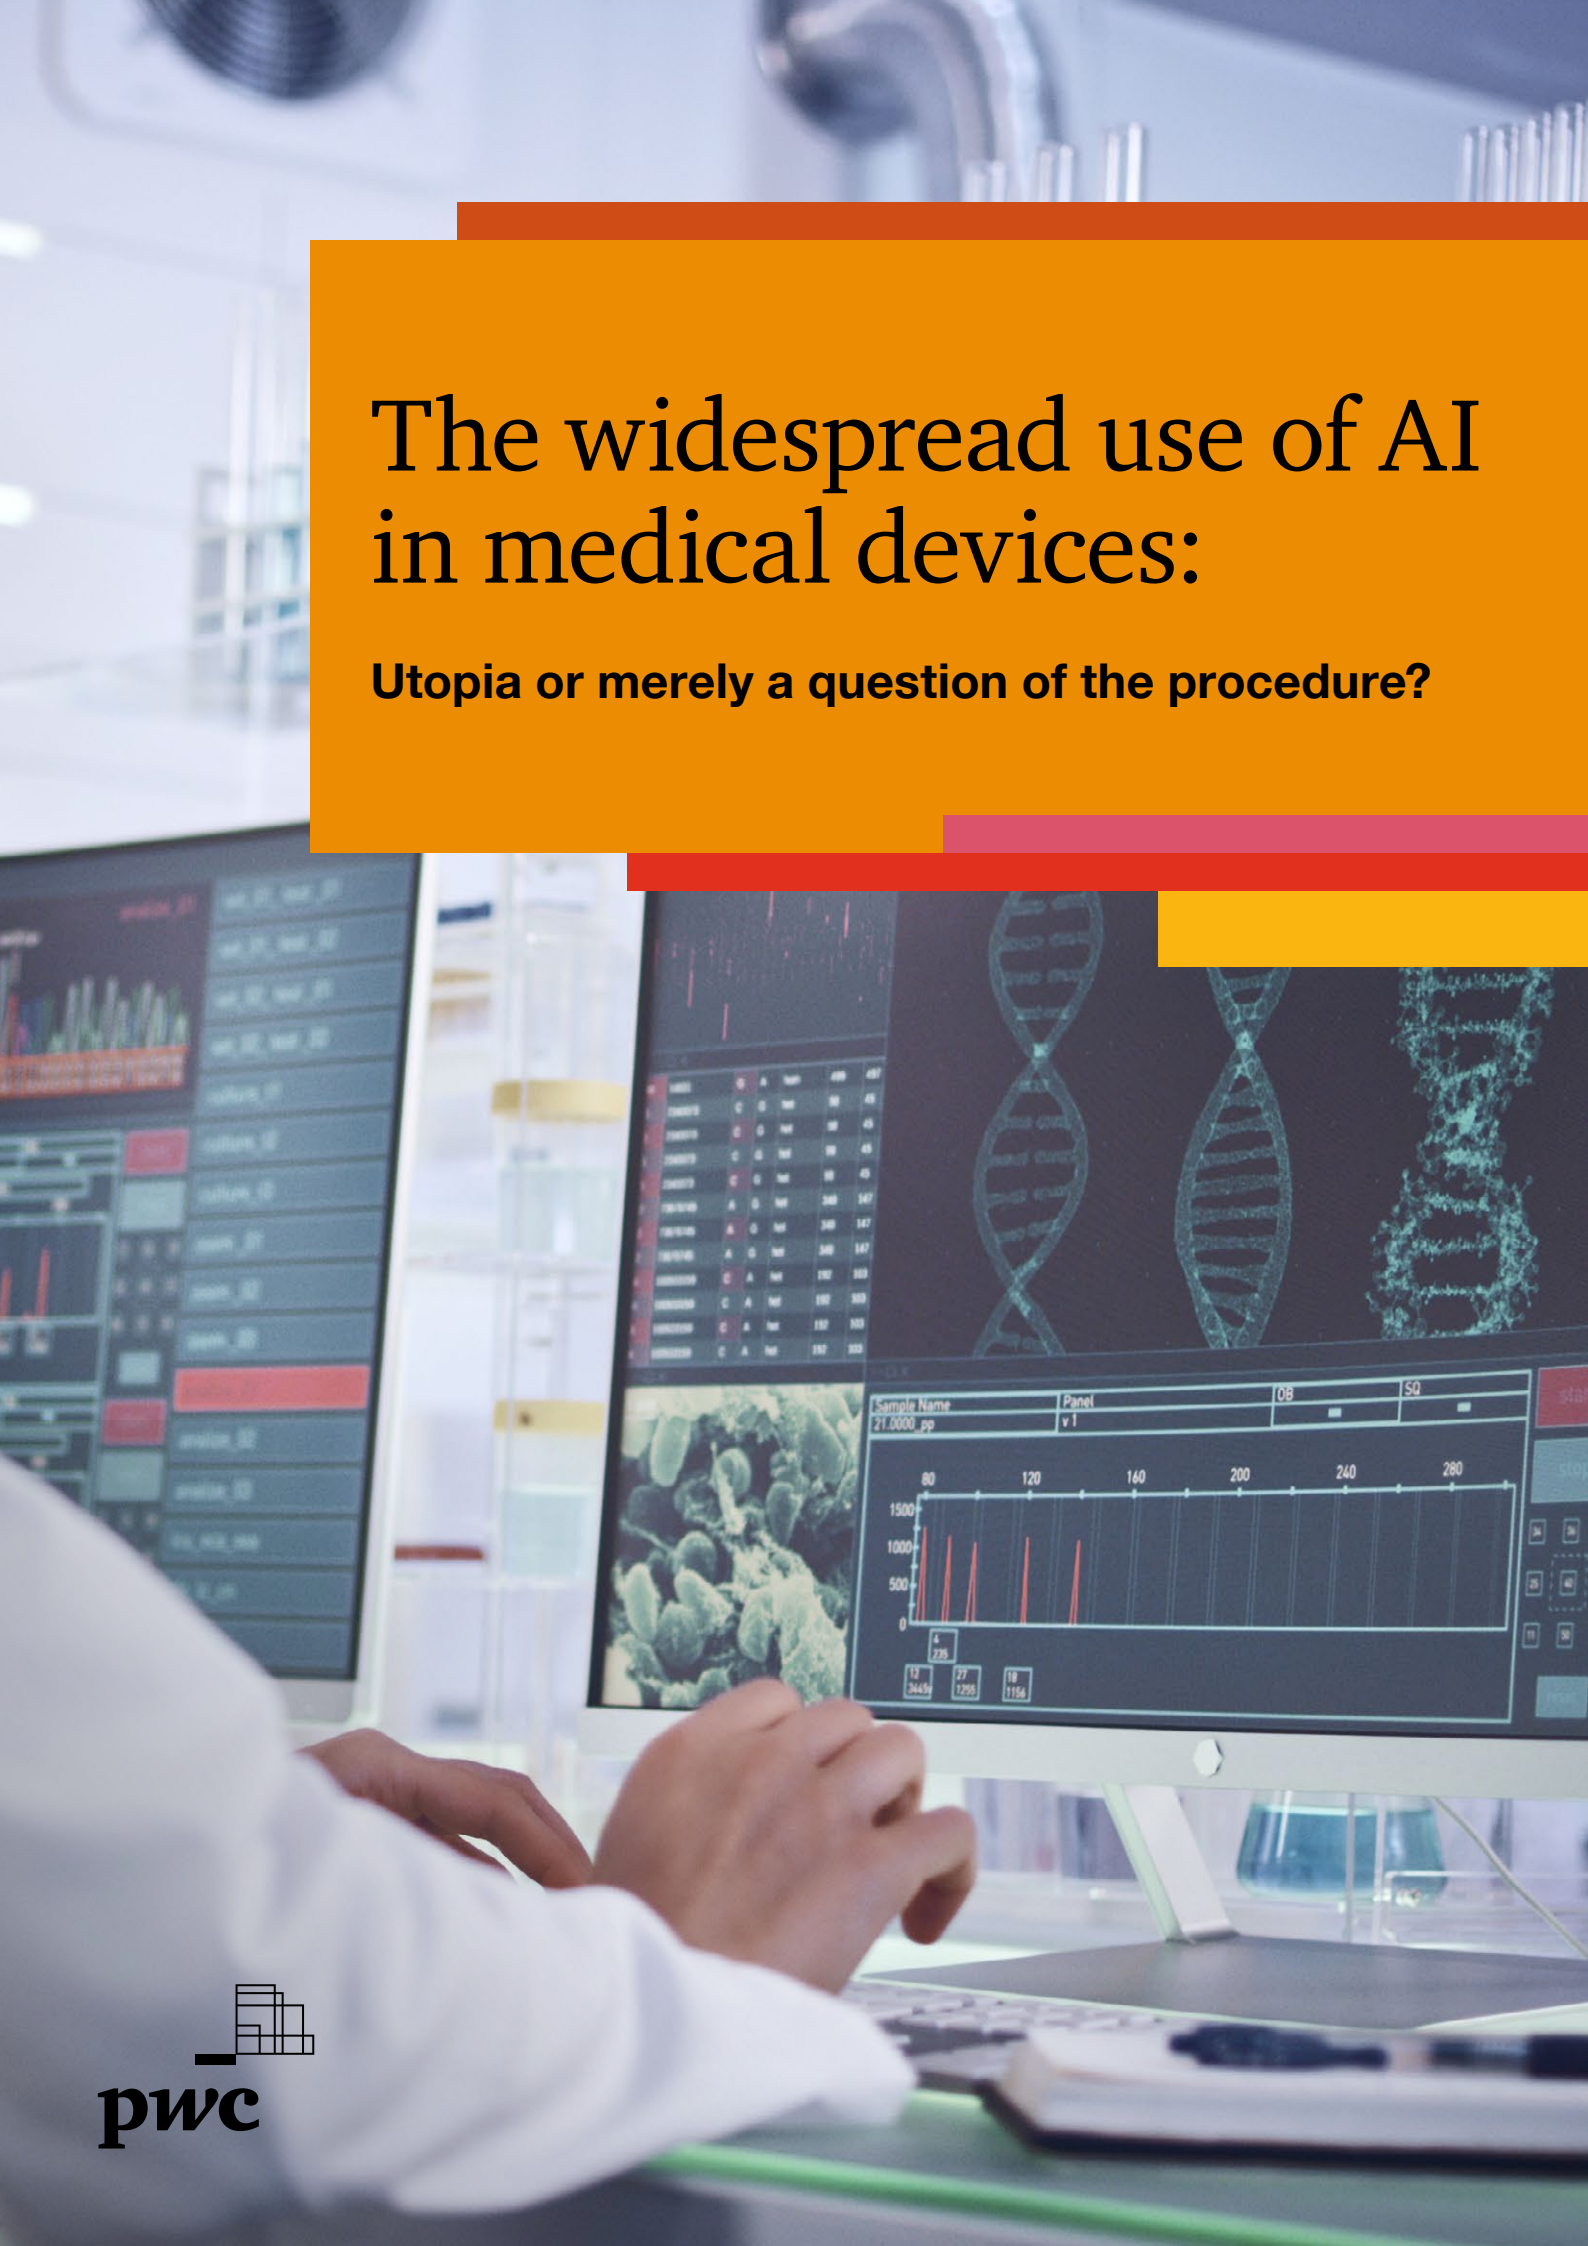

The widespread use of AI in medical devices:

Utopia or merely a question of the procedure?

pwc

Introduction

Like many other industries, the industry for medical devices is amid a digital transformation. Accelerated by the progress of processing capabilities, mobile connectivity, computer networks, and innovations, new opportunities emerge in preventing, diagnosing and treating diseases, which can help improve the healthcare system and, consequently, everyone's daily lives.

Besides fundamental technologies, artificial intelligence (AI) is a potential game changer as it can help optimise prevention, diagnosis, and therapy. Crucially, AI can also fill the gap arising from demographic changes – the shortage of physicians and therapists, which we already observe today in many areas of the healthcare system. The technical prerequisites for establishing AI in medical devices for the market exist. However, uncertainty over regulations, and physicians, therapists and patients' lack of trust are often barriers to implementation. Therefore, the full potential of AI in medical devices has yet to be unleashed.

The underrepresentation of software as a medical device (SaMD) and the associated classification in terms of the medical device regulations (MDR) contribute significantly to manufacturers' uncertainty. The MDR describes and places requirements on manufacturing, assessing and distributing medical devices in the European Union (EU). In the context of machine learning (ML) based SaMD, additional regulations focusing on ML-based software generally – e.g., the EU's Artificial Intelligence Act (AI Act) – further increase the complexity from a compliance perspective.

This complexity strongly connects to the fact that all these regulations are not yet aligned, resulting in identical or contradictory requirements for these systems. Furthermore, the increasing level of networking in terms of the Internet-of-Things (IoT) and the resulting distribution of ML-based SaMD parts between the centralised cloud and edge devices lead to a diverse solutions landscape. The combination of technology-driven changes and regulatory uncertainties with geographic dependencies leads to low trust in the technology, preventing sophisticated solutions from entering the healthcare market.

In the challenging environment of uncertainty and lack of best practices, combining existing regulations and standards from different areas is a practical approach to transferring the technology from policy to operations.

Artificial Intelligence: A sophisticated technology for reshaping the medical devices market

Artificial intelligence is estimated to be the largest commercial opportunity in today's economy[1]. In 2022, the global AI market size was estimated at US$ 56.89 billion. With a compound annual growth rate (CAGR) of 29.0 %, the market is anticipated to snowball this decade[2] (Fig. 1).

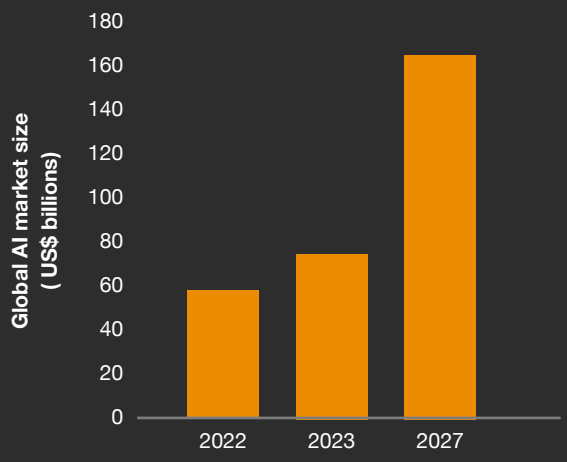

Fig. 1: Global AI market size in USD billions[2].

A survey of 360 companies conducted by Lufthansa shows that 73 % of those with more than 10,000 employees and 59 % of smaller companies used machine learning applications in 2021. The IT budget is the most relevant indicator for ML readiness and use, with a clear increase of AI in companies with budgets over 10€ million. The study also shows that these investments pay off, as 63 % of companies with AI projects profit after only three months and 84 % after one year. The improvements through AI primarily concern increased productivity and efficiency and reduced costs[3]. To unlock the potential, 49 % of executives plan to increase their IT investments by 10 %, according to PwC, as investors and stakeholders also have an increasing demand for data-driven improvements[4].

[1]PwC Artificial Intelligence everywhere | [2]The Business Research Company: Artificial Intelligence Global Market Report 2023 | [3]Lufthansa IDG study 2021 | [4]PwC study: Companies leave data potential largely unused

Aside from general AI developments, their impact on the medical device industry is often considered one of the most disruptive. AI in medical devices has the potential to disrupt all aspects of the industry and even change the whole perspective of medicine from a reactive to a preventive approach.

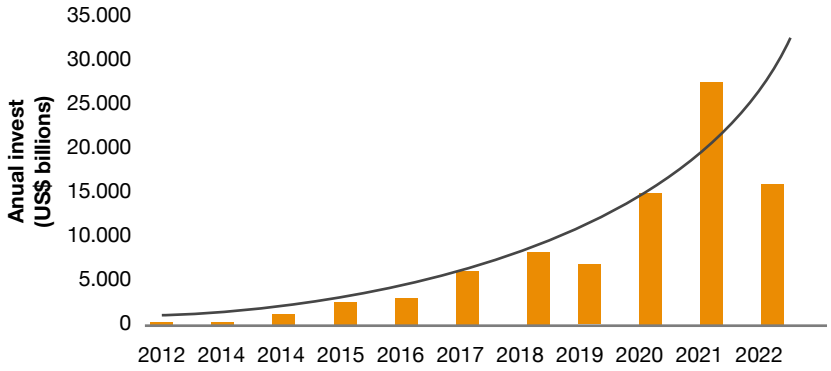

Fig. 2: Annual VC investment in healthcare, pharma and biotechnology start-ups[5].

The annual venture capital (VC) investments in AI healthcare, pharma, and biotechnology start-ups have increased significantly over the last decade (Fig. 2)[5]. Further, with a market size of US$ 167.7 billion in 2023, it's estimated that the global healthcare IT market will rise to US$ 609.1 billion by 2030[6]. The investment inflow is reflected by the number of relevant AI publications in healthcare, which have exponentially increased over the last decades. The US and Europe are two major contributors. The highly potent combination of growing investments and research resulted in a rapid increase of FDA-approved ML-based SaMD over the last ten years, as shown in Fig. 3[7]. The most active approval fields are the radiology and cardiovascular domains, which traditionally hold significant financial value[8].

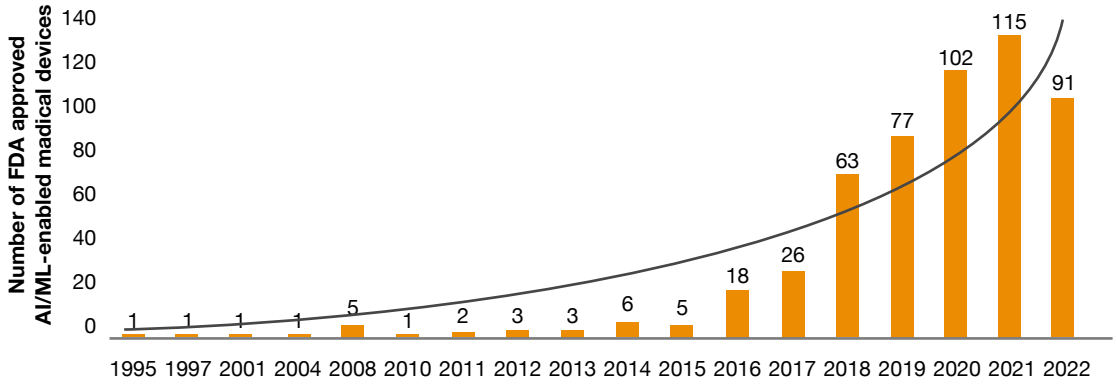

Fig. 3: FDA-approved AI/ML-enabled medical devices from 1995 to 2022[7,8].

[5]OECD.AI (2023), visualisations powered by JSI using data from Preqin, accessed on 25/1/2023 | [6]Grand View Research: Healthcare IT market size | [7]The year 2022 only includes approvals until the end of July, thus the lower number of approvals. Extrapolating the number of approvals, the overall trend should continue. | [8]US FDA: Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices

Besides the approved medical devices, promising and potentially highly disruptive AI applications are in development and already in the market, as shown in Fig. 4. In healthcare research, even end-to-end machine learning pipelines surpassing professionals' performance are introduced[9]. All these developments are estimated to create US$ 150 billion in annual healthcare savings in the US by 2026[10]. Comparing the huge capital investments and high research activity with the comparatively low number of FDA approvals raises the question: Why does this large gap exist?

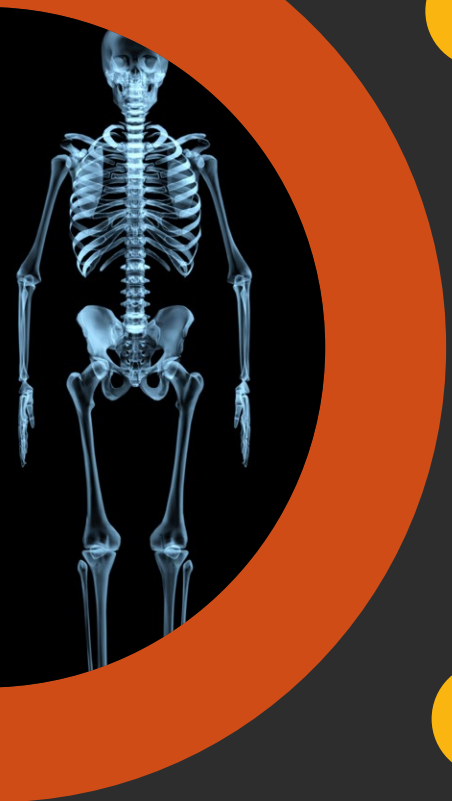

Skin Cancer Detection
A skin cancer detection app achieved comparable cancer classification to a qualified doctor in 2017[11].

AlphaFold
AlphaFold predicts 3D protein structure from the amino acid sequence[12].

Algnostics
Algnostics offers an AI as SaaS that analyzes tissue and cells for pathologies[13].

Siemens
Siemens has incorporated AI in products for the last 20 years, with 700 patents related specifically to machine learning and 275 to deep learning[14].

GE Healthcare
GE Healthcare developed the Edison information platform utilising AI applications to improve doctors' workflows and personalised medicine[15].

Fig. 4: AI-related products and services in the healthcare sector[11,12,13,14,15].

[9]Ardila, D., Kiraly, A.P., Bharadwaj, S. et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. Nat Med 25, 954–961 (2019)
[10]Accenture: ARTIFICIAL INTELLIGENCE: Healthcare's New Nervous System | [11]Stanford skin cancer | [12]AlphFold | [13]Algonstics | [14]Siemens Healthineers | [15]GE Edison Solution

What's preventing artificial intelligence from entering the healthcare market?

The highly fragmented market with high entry barriers

The healthcare market is highly fragmented; the various areas of health insurance companies, pharmaceutical groups, medical device manufacturers, hospitals and doctors typically operate in isolation and are highly specialised in their field of activity. However, ultimately all are closely intertwined with the patient. This state of affairs has evolved from the historical legal framework and increasingly specialised and extensive expertise in each area. In addition to the highly fragmented market, just a few large companies with little to no serious competition often dominate. This situation is largely due to these companies' historical devel-opment and growth in the healthcare market and the extremely high regulatory barriers to entry.

The highly fragmented healthcare market and specialist areas hinder the quick implementation of new and integrated technologies that could improve overall patient care and well-being. But with the increasingly fast technological developments in recent years and the slow innovation rate of large corporations, start-ups with long-term VC backing began pushing into the healthcare market with new innovative ideas – for example, smart devices, digital twins of patients and AI applications for low-level medical diagnoses.

Although most current medical innovations lack regulatory approval, this is only a matter of time. Consequently, large companies from all industries have begun developing data-driven and AI-based products. Now, integrating new trending technologies is creating a conflict between the traditional medical products designed for a fragmented market and the latest digital and connected devices aimed at integrated market solutions.

The inhomogeneous target group of end users

The heterogeneous target group of AI and medical products in healthcare presents a significant challenge. Potential end-users include healthcare providers and administrators, patients, and government regulators. Each of these groups has a different level of understanding of AI and medicine, different needs and expectations. This makes it challenging to develop and implement AI solutions that meet everyone's requirements.

Considering the target group of end users, especially the users' behavioural intention to use technology, is significantly influenced by their social groups and beliefs regarding the expected enhancement in performance. Therefore, the appropriate usability and user experience are key drivers for the product's market acceptance. Additionally, the general scepticism and lack of explainability associated with AI can further increase this issue. Studies have shown that while AI applications could be useful in healthcare, some stakeholders often do not perceive them as functional. This perception can be due to a lack of transparency and the complexity of the technology. However, regulatory frameworks and trustworthy AI solutions could improve user acceptance and increase their adoption of new technologies.

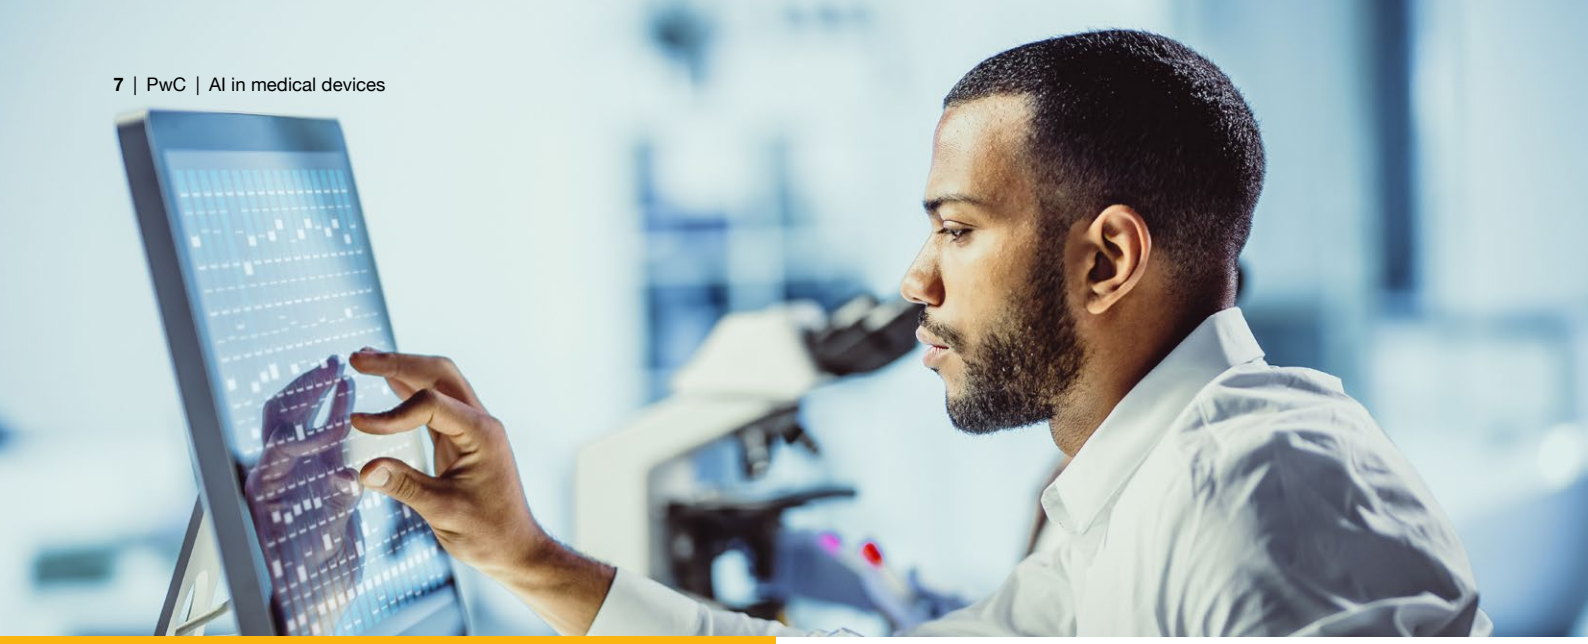

The complex regulatory landscape for ML-based medical software

As medical devices directly influence human lives with the potential to harm or save millions of people, they are subject to strict regulatory requirements. With the patient always in focus, the main goal of these regulatory frameworks is not to restrict medical device diversity but rather to give a quality framework for security and safety.

Focusing on the US market, the Food and Drug Administration (FDA) is the legal body for establishing guidelines and requirements for medical device manufacturers. In particular, the FDA clearly indicates the requirements for SaMD regarding risk categorisation, the quality management system, and clinical evaluation, which each medical device manufacturer must fulfil before entering the market.

Besides the baseline requirements for SaMD, the FDA proposed a regulatory framework for modifications to AI/ML-based soft ware for medical devices in 2019. In 2021, the FDA specified the approach into a five-step action plan, slightly adjusting its main point, the 'Predetermined Change Control Plan' after feedback. According to the FDA, evaluating SaMD that leverages AI will in the future rely on good machine learning practices developed on a global standardisation level (ISO), be patient-centred, incorporate transparency for users, focus on the scientific evaluation of bias and robustness, and be oriented to real-world performance.

Looking at the EU, the medical device regulation (MDR) has been the main quality gate across the European market since 2021. The competent authorities of each member state are responsible for implementing the MDR in a national context. Article II Section (1) of the MDR defines a medical device as a medical application and software used for diagnosing, preventing, monitoring, predicting, prognosis, treating or alleviating disease and disabilities. Depending on the device's intended purpose and inherent risk, medical devices are classified into one of three risk categories, outlining the need for a CE marking to prove they are conform with the requirements[16]. If the software influences or drives the device's use, it falls into the same risk category as the device. If the software is independent of the device, one must define its SaMD class.

[16]Medical Device Regulation (MDR)

Although of the different risk classes in Fig. 5, each SaMD should follow established principles of the development lifecycle and risk management, including information security, verification and validation (MDR 17.2). In order to meet these requirements, the MDR does not give specific recommendations but instead leaves it open to the manufacturer to choose a method that is state of the art, meaning harmonised standards.

Using AI in SaMD adds another layer of complexity to the MDR dimensions of risk classification. As the first attempt to regulate AI-based systems in the European market, the upcoming EU AI Act establishes new and different risk classes (Fig. 5). As a safety component or even as part of the critical infrastructure, most SaMD would fall under the high-risk classification, which outlines specific additional requirements in the risk management, security, robustness, and accuracy. Having a further risk classification that is not aligned with the MDR places yet another burden on introducing medical device software to the EU market, as well as the related administrative and financial costs.

Fig. 5: Comparison of risk classification in accordance with the EU AI Act and MDR.

The high requirements for data governance

In today's world, it is no surprise to most of us that every action we take on an electronic device generates a wealth of data. This connectivity and information fuels innovations worldwide and increasingly impacts the health ecosystem. With increased data quality and availability, the number of different software applications as medical devices has risen significantly. Since most medical applications try to solve specific, individual-related issues, they rely on personal and, thus, critical information about a person's condition.

The regulatory landscape regarding data changed dramatically in recent years with the EU's data strategy. Starting with the General Data Protection Regulations (GDPR) in 2018 regulating data protection and the rights of EU citizens over their data, the Data Act was added to the agenda in February 2022 and laid down rules on data sharing across economic sectors (B2B, B2C, B2G). The final piece of the general strategy, the Data Governance Act (DGA), came into force in June 2022, providing rules for the re-use of data originally created in the public sector for commercial or non-commercial purposes, for causes other than the original (government and public services, healthcare, etc.).

Looking at the introduction of the European Health Data Space (EHDS), one additional concept of the DGA plays an important role: data altruism, where data is made available without reward for purely non-commercial usage. With this concept in mind, the EHDS creates a space according to the FAIR (findable, accessible, interoperable, and reusable) principles allowing doctors, scientists, and policymakers to base their decisions on a broad, high-quality health database. In return, the EHDS enables a European digital service infrastructure and a European reference network for rare diseases, putting the patient into the centre.

All three data-specific acts add a new layer of requirements to the existing GDPR, regulated on an EU level. However, there are many nation-specific requirements to be aware of since EU law is generally reinforced and extended by national initiatives creating a jungle of regulatory requirements for using health data in the EU.

The challenges of ML-based towards traditional rule-based software

Due to the highly critical nature and relevance of the medical subject, the regulations require that a human caregiver or legislator can understand all decisions and considerations made by patient care software. This requirement is implicitly given in traditional rule-based software products. Decision and clustering rules are explicitly defined by software engineers developing the software-based products and their deterministic nature.

In contrast, this is not true for ML-based software products, which follow a probabilistic approach. This approach means decision boundaries are not formulated explicitly and cannot be easily comprehended by humans (black-box scenario). Thus, questions arise around transparency, explainability, interpretability, data protection, fairness, and bias for ML-based software products. Therefore reasoning, i.e., about the explainability of an ML-based application, requires new methodologies to enable an understanding of the ML system's underlying decision boundaries. For this, further technological developments are necessary that allow a good comprehension of the reasoning of ML-based software.

Aside from the required technological advancements, the legal framework for software-based medical products must also adapt to the new probabilistic approaches. The current regulation aims to control traditional deterministic-based software products. It thus is unsuitable for the probabilistic black-box approaches of ML-based medical products. Specifically, the regulation should be adapted so that findings from reliable probabilistic methods that explicitly formulate the implicit decision boundaries of ML-based medical software products are sufficient for successful software testing.

From policy to operations: Best practices for building trustworthy ML-based medical devices

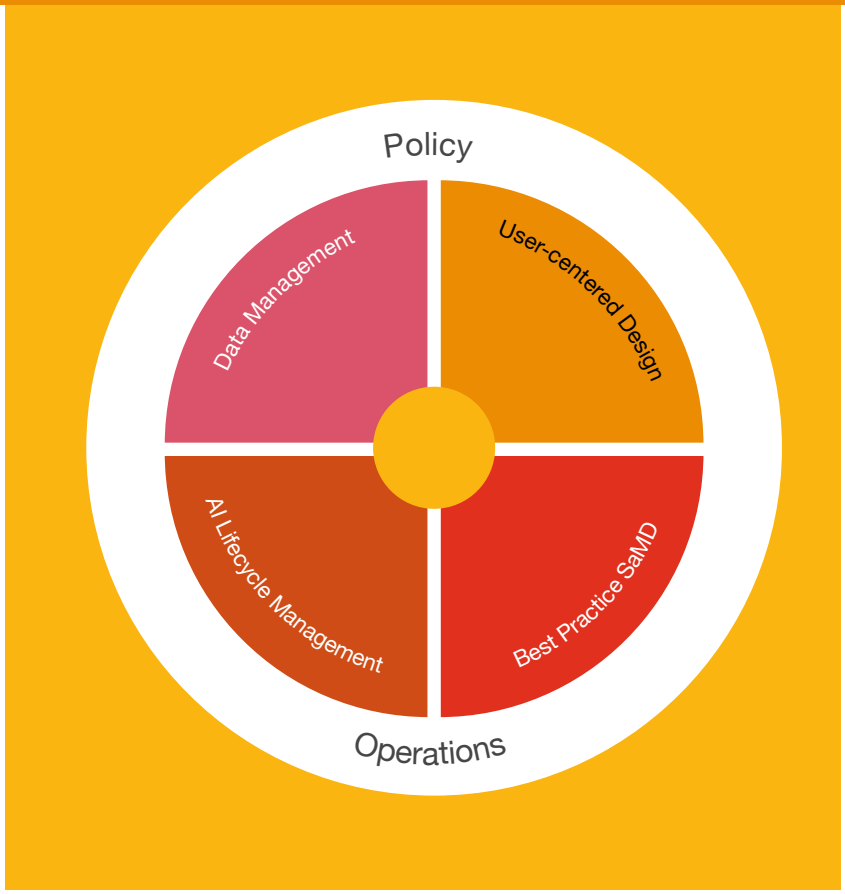

Fig. 6: Four focus areas for bridging the gap between policy and operations.

Regardless of the very specific requirements and market complexity, using ML-based software in medical technologies and solutions is not excluded. This use is demonstrated not least by ML-based SaMD, which already has approval up to class III.

Although a holistic process model cannot change the market and the regulatory framework, it can be used to develop best practices for developing and operating ML-based SaMD in conformity with the regulatory framework. The basis for such a procedural model is an iterative process that is not only standardised but also has a high degree of automation to reduce the susceptibility to errors. Furthermore, the process requires the early involvement of users and appropriate error monitoring mechanisms in the testing phase and operation. The process model focuses on data management, user-centred design, best practices for SaMD, and AI lifecycle management (Fig. 6).

Data management

High-quality data management is the foundation for developing reliable algorithms and trustworthy ML-based SaMD. Primarily, this involves creating a data governance model. The model covers the roles and responsibilities regarding the company's data and the necessary organisational processes, aiming for high-quality data. Data protection also plays an important role. It is crucial to validate which laws and regulations are in place and relevant and to know which inclusion and exclusion criteria to apply to the data, at which intervals the data are collected and available to whom, and how the data are processed.

These points apply to training, validating, and testing data and also during the pre-development phase concerning data labelling and pre-processing. Consider the lifecycle of the data holistically. Introducing a company-wide data catalogue as part of data governance is advisable to enable a single point of truth for the data's origin, access, and responsibilities. The catalogue will achieve high data quality and integrity and guarantee quick access to the data.

Finally, pipelines – managing the inflow and outflow of data, the sequential steps in between, and the ingestion of new training data – ensure efficient data management. Here, quality gates and a high degree of automation lead to higher reliability and, ultimately, higher data quality.

User-centred design

To consider the diverse requirements associated with an inhomogeneous user group, the development process for ML-based SaMD should involve users early. Their involvement brings usability to the foreground and reduces the likelihood of future complex change processes.

Introducing the user group early to the technology also builds trust from the initial phase of the development. In addition to early involvement in requirements management, quality gates through user acceptance tests built into the iterative, agile development process are an adequate tool for bringing the user into the development focus.

User acceptance tests also provide a sound basis for subsequent clinical evaluation. In this context, the standards for usability in medical devices (IEC 62366-1) provide additional requirements for consideration, which support and align with the concept of user-centred design.

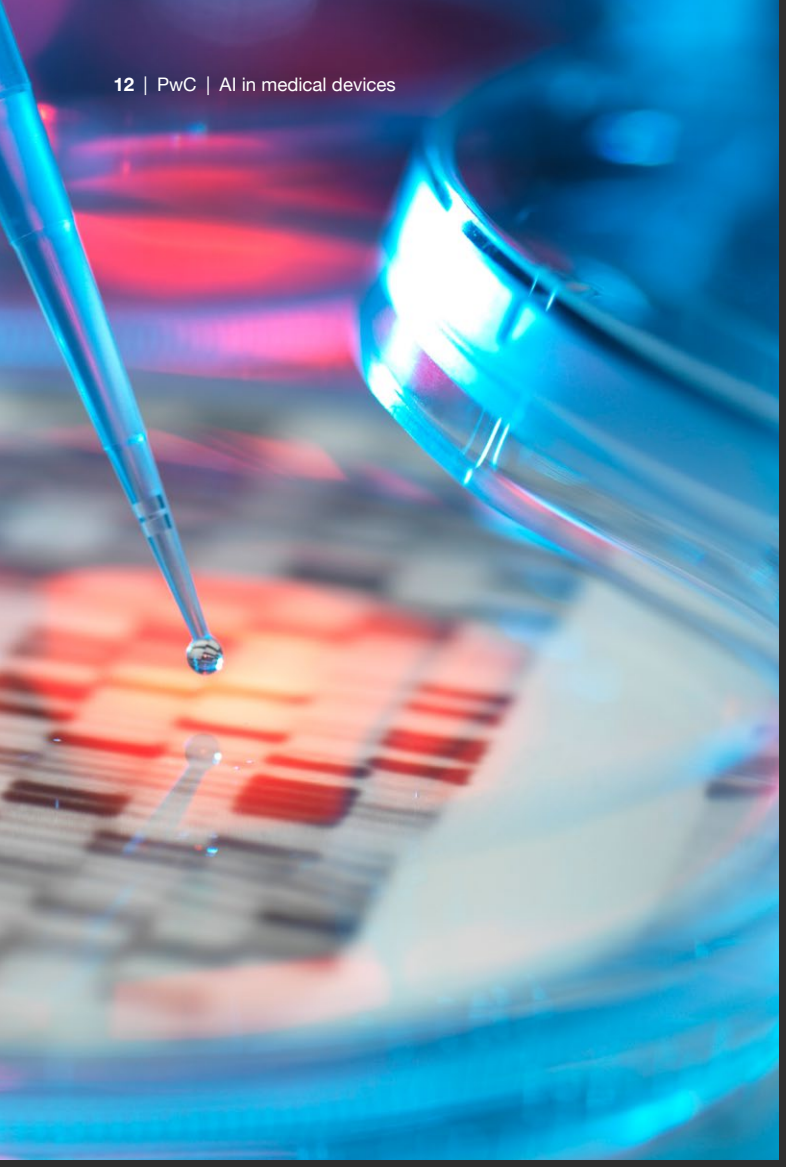

Best practices for developing software-based medical devices

A standardised software development process should be the basis for development; IEC 62304 provides the necessary requirements. As well as development, the IEC standard requires processes for maintenance and operations, issue management, and configuration management. Although IEC 62304 is based on the classic V-model, it does not contradict agile development per se. Dynamic adaptation of the individual artefacts in an iterative and agile way is still possible, but quality gates with appropriate approvals should be built in.

In parallel with the software lifecycle, the processes for risk management (ISO 14971) and quality management (ISO 13485) need to be in place, which the MDR also requires. As well as the fact that developing ML-based software requires an agile development process that differs from traditional software development, there are significantly higher hurdles in operating the software. These hurdles are caused particularly by the properties of the algorithms but also by the database.

AI lifecycle management

An agile software development process requires a high degree of standardisation and automation to guarantee consistently high quality of the solution during and after each development cycle. For this purpose, the practice of continuous integration and continuous delivery has become established, enabling a smooth transition from development to operation.

In developing ML-based SaMD, operating the software, in particular, must be considered holistically in the AI lifecycle to derive more from MLOps' best practices. Chapter 2 outlined the challenges associated with the characteristics of AI models. Measures that monitor the software's behaviour and output during operation and intervene on an event-driven basis when thresholds are exceeded or make time-driven adjustments address these challenges.

A technical safety concept at the beginning of the development is one way to identify the possible hazards that may arise from the software and to develop threshold values for metrics which trigger the system's reaction to mitigate or prevent predicted hazard events. Such strategies can be used to identify model drift or performance limits and take appropriate action.

A user-centred approach is not only a vital component during development but also for operation. User interaction logs, through which users can provide direct feedback, and regular user reports allow users to participate in the software's evaluation and optimisation cycle and thus sustainably keep the quality of the software solution high. In addition, the users should be involved in the update process.

Finally, reproducibility and explainability are essential components that can be guaranteed within the framework of MLOps. The model's decisions can be reproduced afterwards by storing model metadata, output, input and metrics. Considering explainability in model selection and parameterisation (the process of defining or choosing parameters) helps trace the reasons for the model's decision.

Focusing on these key topics can equip manufacturers with an approach or even procedures for dealing with the software lifecycle's requirements and the usability of ML-based medical devices combined with traditionally oriented standards and the challenges of applying AI.

Conclusion: Trust is an enabler that unleashes the full potential of artificial intelligence in medical devices

A combination of the healthcare market's characteristics, the challenges with ML-based software, and its lifecycle hinder many companies from bringing ML-based SaMD to the end user.

The main reasons for this are not missing regulations or standards but more the lack of established best practices for aligning MDR and ML standards. This results in uncertainty on the manufacturer's side and less user trust. End users' scepticism of the technology deepens the lack of confidence. Even if ML is already established and omnipresent in daily living, some people are warier about adopting recent technologies concerning their health. For this reason, we focus on trust as we see it as an enabler and one of the core drivers for bringing ML solutions to the user.

Consequently, our approach is a comprehensive framework for supporting manufacturers under the premise, 'don't do as much as possible, but only as much as necessary!' The framework integrates dedicated technological measures into the product lifecycle to ensure the solution's integrity, safety and security and support the entire development process from ideation to operation. These measures help unleash AI's evident and exciting potential.

Our experts

Hendrik Reese
Partner

hendrik.reese@pwc.com
+49 89 57906093

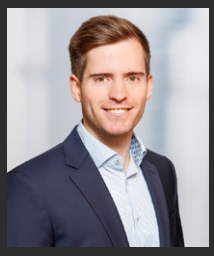

Dr. Sebastian Becker
Manager

sebastian.becker@pwc.com
+49 211 9812311

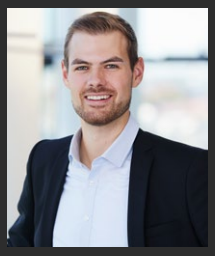

Tristan Armbruster
Associate

tristan.armbruster@pwc.com
+49 89 57906349

Jan Boysen
Associate

jan.boysen@pwc.com
+49 89 57905068