AI Machine Learning and Medical Devices

Methods and guidance on addressing safety

Author: Rich de la Cruz, President of Silver Lake Group, Inc.
Contributing author: Dr. Aris Tzavaras, Head of Artificial Intelligence Notified Body, BSI Group

BSI White Paper Series

By Royal Charter

Contents

3 Introduction

5 Definitions

6 Intended use/intended purpose

7 Medical Device Regulation, guidance and standards

11 Bias and trustworthiness

12 MLMD Training

15 Monitoring and mitigation

17 Regulatory views and concerns – Dr. Aris Tzavaras (BSI)

18 Conclusion

19 References

20 Authors

21 Reviewers

23 Published white papers

Introduction

"Change is the only constant in life." Heraclitus, a Greek philosopher.

If you are not running to keep abreast of the changes in computer science, regulations, and standards, you will fall behind.

If you are running to keep up, run faster! Change happens at the speed of thought. The reality is you will never know everything, but you will be closer than most.

Machine Learning for Medical Devices (MLMD) has created a level of excitement not seen since the discovery of stem cells over 60 years ago.[1] There are similarities between the excitement generated by the discovery of stem cells and Artificial Intelligence / Machine Learning and its use in medical devices.

The discovery of stem cells created quite an uproar, with the idea that the use of stem cells would quickly revolutionize the world; but the use of stem cells was also accompanied by a lot of controversy. This is very similar to what we are seeing with the advent of Artificial Intelligence (AI).

While wonderful things have been accomplished with the use of stem cells, it has not been at the pace that many would have liked due to multiple issues. Putting new technology into practice can be a long and arduous path, in particular when failure of the technology may cause harm.[2] The question we should ask is: are the training, tools, regulations and standards ready for the widespread use of MLMD?

1 https://bioinformant.com/who-discovered-stem-cells/#research; https://www.discovermagazine.com/health/is-the-dawn-of-the-stem-cell-revolution-finally-here

2 BS EN ISO 14971:2019+A11:2021, *Medical devices — Application of risk management to medical devices*

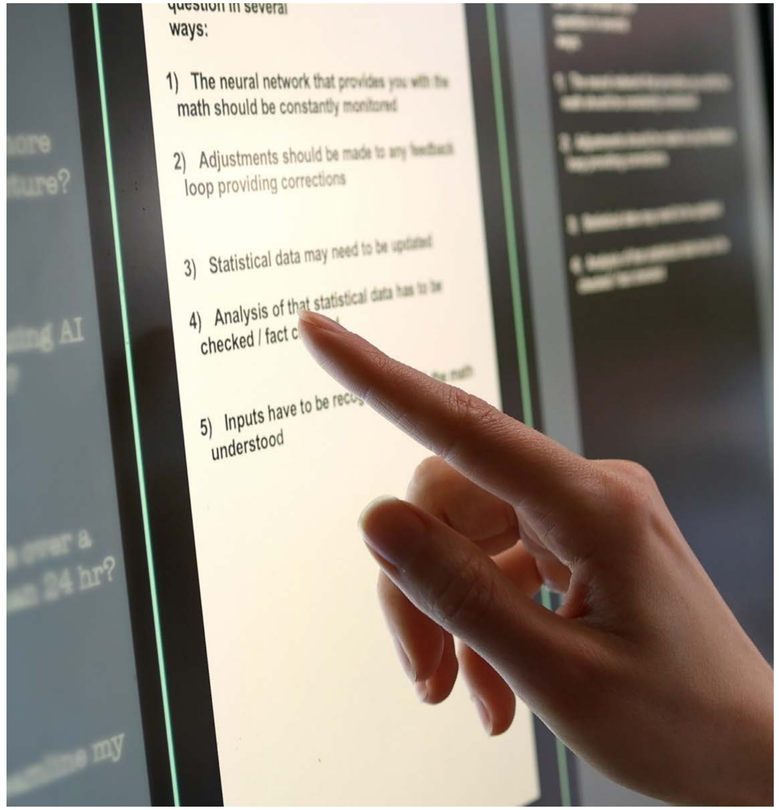

This paper will examine the following issues associated with the use of Artificial Intelligence / Machine Learning (AIML) in medical devices.

1 **Intended use / intended purpose:** The manufacturer must ensure that the output of the MLMD conforms to the verified and validated intended use of the product, its claims, its labelling and instructions for use.

2 **Regulations, guidance and standards:** MLMD must demonstrate compliance to medical device regulations, guidance and conform to standards, some of which are being created, or are changing for MLMD. MLMD must also conform to additional security and privacy concerns, as defined in the current draft of the EU Artificial Intelligence Act and other regulations as they are developed.

3 **Bias and trustworthiness:** The MLMD must be trustworthy. To be trustworthy, AI technologies must appropriately reflect characteristics such as accuracy, explainability and interpretability, privacy, reliability, robustness, safety, and security or resilience to attacks – and ensure that bias is mitigated.[3]

4 **MLMD Training:** manufacturers should create objective evidence that demonstrates that the training data is appropriate for use. Additionally, demonstrating, via objective evidence, that as learning occurs, the results stay within the approved intended use / intended purpose.
   - Note that besides the training data, test data must be used to verify that the AI produces validated output.

5 **Monitoring and mitigation:** Manufacturers must determine how to mitigate negative reactions to the use of AI in medical devices and determine how they will conduct post-market monitoring.[4]

3 https://www.Ai.Gov/strategic-pillars/advancing-trustworthy-ai/

4 WHO urges caution over use of generative AI in healthcare May 16, 2023: https://techmonitor.ai/technology/ai-and-automation/ai-in-healthcare- who#:~:text=The%20World%20Health%20Organisation%20%28WHO%29%20has%20issued%20a,taking%20appropriate%20precautions%20over%20bias%20and%20misdiagnos is%20risks.

Definitions

*Definitions used in this document with an * were created by 'the author.'*

Bias [from PD ISO/IEC TR 24027:2021]: "...the term bias is defined as a systematic difference in the treatment of certain objects, people, or groups in comparison to others..."

***Constrained AIML**: algorithm output, which is limited by design to a range of output that has been verified and validated.

***Data of Unknown Provenance (DOUP)**: data where the source, date of creation, or validity of the information cannot be confirmed.

***Dynamic-AI**: where the AIML learning occurs, both prior to release with fully validated data and in the field as the product is used.

MLMD: A medical device that uses machine learning, in part or in whole, to achieve its intended medical purpose.

***Static-AI**: where AIML occurs prior to placing on the market with fully validated data and with the output fully verified and validated. No additional learning occurs once the product is released, and the algorithm is locked or otherwise constrained.

SaMD (Software as a Medical Device) [from the International Medical Device Regulators Forum (IMDRF)]: "software intended to be used for one or more medical purposes that perform these purposes without being part of a hardware medical device."

SiMD (Software in a Medical Device) [from the FDA in *Guidance on Premarket Submissions for Device Software Functions*]: software that operates or controls a medical device and cannot operate on its own.

Verification[5]: "confirmation by examination and provision of objective evidence that specified requirements have been fulfilled."

Validation: "confirmation by examination and provision of objective evidence that the particular requirements for a specific intended use can be consistently fulfilled."
Note that validation usually is against user needs.

5 Verification and validation are from the the FDA 820.3 definitions. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfCFR/CFRSearch.cfm?fr=820.3

Intended use / intended purpose

Manufactures must ensure that their medical devices operate within the specified intended use / intended purpose after the product is released to market.

In the development of MLMD, the manufacturer must show that that the device will continue to operate within the intended use, claims and labelling.

The EU Medical Device Regulation (MDR) defines the intended purpose in MDR Article 2(12) as follows:

“Intended purpose means the use for which a device is intended according to the data supplied by the manufacturer on the label, in the instructions for use, or in promotional or sales materials or statements, and as specified by the manufacturer in the clinical evaluation.”

Medical Device Regulation, guidance and standards

Regardless of the development methodology or technology used for the medical device, all medical devices must comply to applicable regulation, and should conform to standards.[6] This can be difficult for people who are working on the forefront of AIML, as the rules for MLMD are quite dynamic.

BS EN 62304:2006+A1

When developing SiMD or SaMD, manufacturers are required to follow guidance and standards for development, which include BS EN ISO 13485:2016+A11:2021, BS EN 62304:2006+A1:2015[7], PD CEN ISO/TS 82304-2:2021 and BS EN ISO 14971:2019+A11:2021 as applicable.

Figure 1: A simplified view of the BS EN 62304:2006+A1:2015 framework for software life cycle processes and activities (Patty Krantz-Zuppan, Rich de la Cruz)

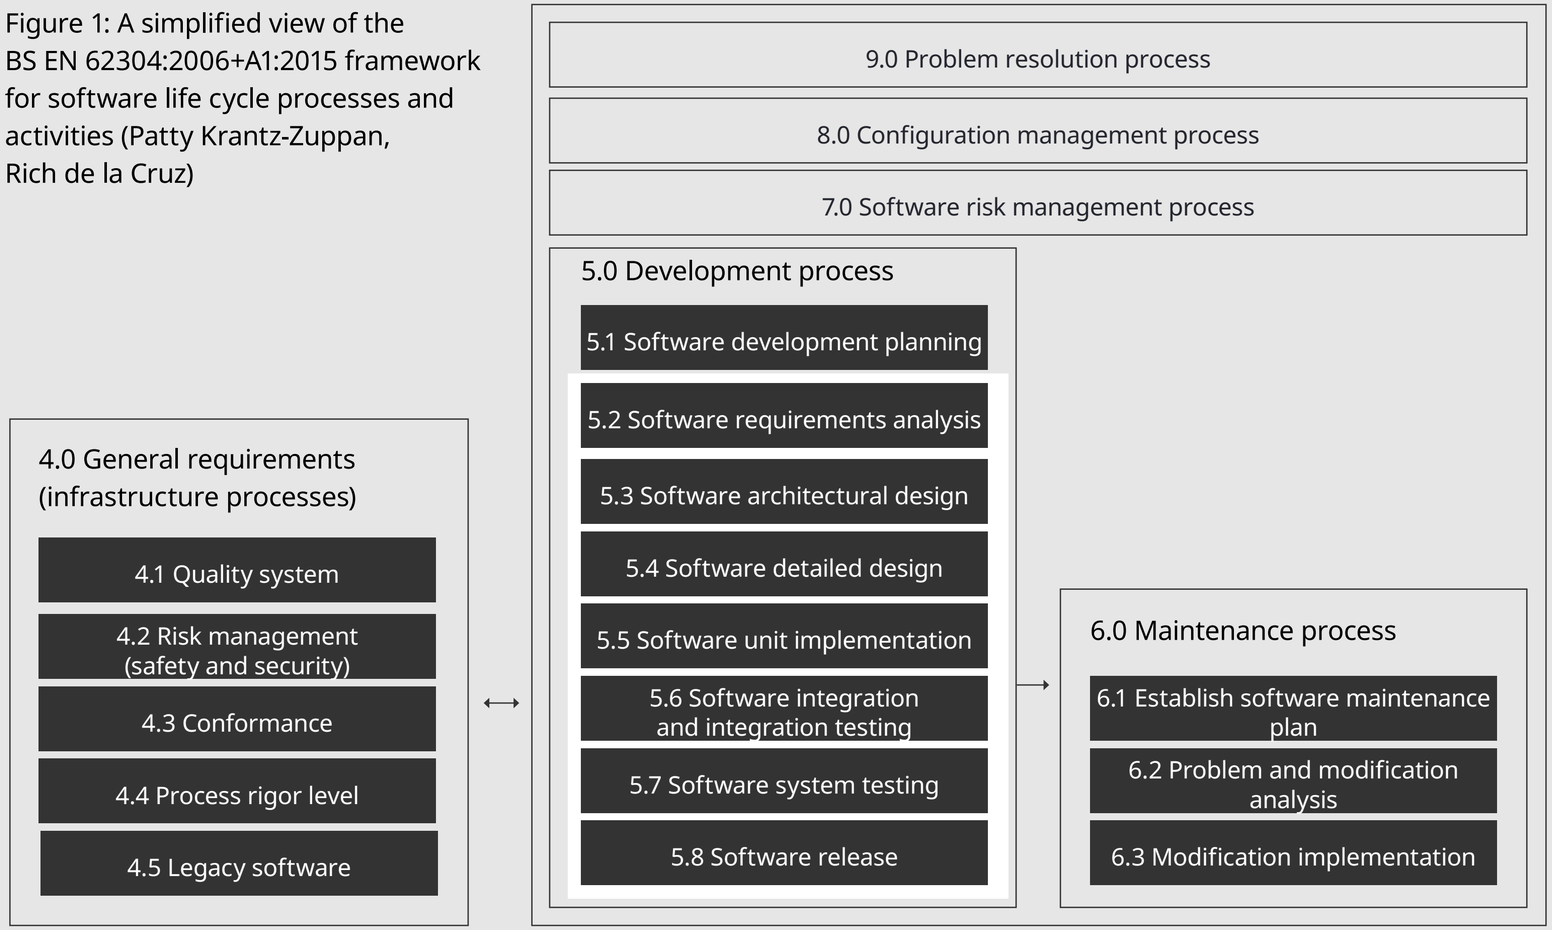

6 Technically, use of standards is voluntary in that regulated bodies will not require the use a specific standard. I have never seen a company forego the use standards because proving that your approach meets FDA or EU expectations absent conformance to an accepted standard can be very difficult.

7 https://array.aami.org/content/news/application-iec-62304-ai-and-other-technologies-s-not-rocket-science-s-computer-science

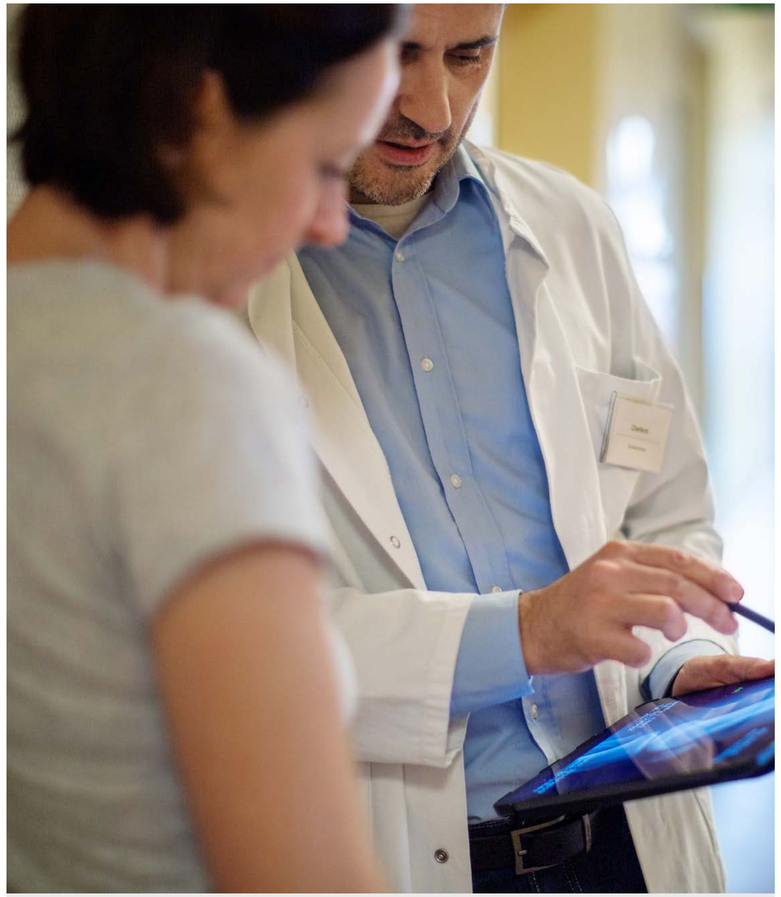

The proposed European Union AI Act (AIA)

From the proposed AIA: “The proposal sets harmonized rules for the development, placement on the market, and use of AI systems in the Union following a proportionate risk-based approach.”

For SaMD or SiMD, medical device manufacturers are already following BS EN ISO 13485:2016+A11:2021, BS EN 62304:2006+A1:2015, PD CEN ISO 82304-2:2021, BS EN ISO 14971:2019+A11:2021, and BS EN 62366-1:2015+A1:2020 as appropriate, which cover many of the new AIA’s requirements. The AIA is considered horizontal legislation, as it does not apply to a specific industry, but to multiple industries. It is not changing the requirements of the MDR, but it will add additional requirements.

The are some issues with the AIA as it is not clear how the AIA would coexist with medical device guidance and standards. How will it prevent duplication? Could it force medical device manufacturers to use less effective methods than those used now? Will its enactment require the medical device manufacturers to update existing medical device standard operating procedures to match the terminology that the AIA uses?

The AIA also indicates that the requirements for the AI systems set out in the Act will be checked as part of existing conformity assessments. While that sounds straightforward, implementation could be a challenge as regulatory bodies interpret what conformance means.

See:

- The AI Act https://artificialintelligenceact.eu/the-act/
- *capAI – A Procedure for Conducting Conformity Assessment of AI Systems in Line with the EU Artificial Intelligence Act* https://artificialintelligenceact.eu/assessment/

Application of BS EN ISO 14971:2019+A11:2021 to MLMD

Many questions have been asked about how to apply the requirement of BS EN ISO 14971:2019+A11:2021, *Medical devices — Application of risk management to medical devices* to MLMD development. AAMI has released CR34971:2022, *Guidance on the application of ISO 14971 to artificial intelligence and machine learning*,[8] which the FDA recognizes as its first AI-focused document, in the FDA list of consensus standards.[9]

MLMD is required to use BS EN ISO 14971:2019+A11:2021 for Risk Management. BS EN ISO 14971:2019+A11:2021 requires full traceability from requirements to verification and validation. In particular, the requirements for intended use and foreseeable misuse are critical to MLMD and are based on the medical device's intended use.

AAMI CR34971:2022 provides guidance on meeting the requirements of BS EN ISO 14971:2019+A11:2021, including the following:

- Risk Management File
- Risk Management Planning
- Risk Analysis – Intended Use and Misuse
- Mitigation of Safety Risks Associated with AIML
- Data Management and bias
- Identification of Hazard and Hazardous Situations for AIML
- Risk Controls and Evaluation of Overall Risk
- Post-Production Activities
- Risks Associated with Human Factors

An area of concern, related to risks associated with human factors that should be addressed for MLMD, is what I call "the computer said it was OK syndrome" or the "automation of bias."[10] Many years ago, I was developing kinetic modelling (Kt/V) SaMD, which was an accessory to a haemodialysis machine device. The application software, working with real-time haemodialysis machine readings, would suggest changes that the clinician could make to the haemodialysis machine device settings to keep the patient within the prescribed Kt/V.

I was directed to allow the software to suggest increasing either time or blood flow rate to meet the prescribed Kt/V. When conducting formative human factors testing per BS EN 62366-1:2015+A1:2020 *Medical devices – Application of usability engineering to medical devices* with a clinician, the software simulated that the Kt/V would not be met and suggested changes to increase the time of dialysis or increase the blood flow rate. The clinician elected to change the blood flow rate to meet the prescribed Kt/V. However, the blood flow rate was changed without considering whether the patient could tolerate the new increased rate, or whether the suggested blood flow rate was actually within the physician-prescribed range for the patient.

8 https://array.aami.org/content/news/u-s-fda-recognizes-first-artificial-intelligence-guidance-among-updated-list
9 https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfStandards/results.cfm

10 'Automation bias' is explained on page 11 of *Clinical Decision Support Software: Guidance for Industry and FDA Staff*

There were reasons for the clinician to change the blood flow rate. A few examples follow below, indicating reasons to increase blood flow rate instead of increasing the time of dialysis.

- Due to the pre-planned patient schedule, other patients were waiting for their appointed time to use the haemodialysis station; therefore, time could not be increased or extended.
- It was nearing the end of the daily schedule; the clinician's shift was scheduled to end and they did not have an option to work overtime.
- The patient may have refused to extend their dialysis treatment time.
- Most alarmingly, increasing the blood flow rate was considered acceptable because the computer suggested it.

After that formative test, we no longer allowed the software to suggest changing the blood flow rate. At the time, there was not enough information available for the software to suggest the change of blood flow rate as a "safe" option.

MLMD has issues understanding the full context of the overall environment and any systems associated in making clinical decisions. Presenting a "right solution in the wrong situation" is one of the concerns that must be addressed as part of Risk Management.

The following is an excerpt from a video[11] by Professor Xenophon Papademetris of Yale University:

Recent developments in AIML (primarily deep learning) have been surrounded by a large amount of hype and overpromise; a phenomenon that is common in the history of AI. One of the major problems we face is how to avoid overlearning or overtraining an algorithm from the relatively small training datasets available (as compared to what is used for non-medical applications). Researchers in the field are familiar with how an algorithm's performance can deteriorate over time as it gets applied to data from slightly different scanners (or even the same scanner after a minor software upgrade), which are both fundamental due to such overtraining. So, while there are many papers advertising exceptional performance, much of this is artificially inflated. The situation is analogous to the p-hacking (reproducibility) crisis seen in other areas of science.

"These risks make AI a uniquely challenging technology to deploy and utilize both for organizations and within society. Without proper controls, AI systems can amplify, perpetuate, or exacerbate inequitable or undesirable outcomes for individuals and communities. With proper controls, AI systems can mitigate and manage inequitable outcomes".[12]

Verification that the data used to train, tune and test the AI may be as much work as developing the product. Requirements for the makeup of the data used to train the AI should include bias definition, that is a well-defined requirement of what the developers consider positive, negative, or neutral bias which may include performance and operational characteristics. Full traceability from requirements to testing to show that the data meets the defined biases is also important.

11 Lessons from the regulatory process for medical software for image analysis and AI (Dec 12, 2022): https://www.youtube.com/watch?v=3qdrZ65DgL0

12 From the Artificial Intelligence Risk Management framework (AI RMF 1.0) by the National Institute of Standards and Technology (NIST):

Bias and trustworthiness

Discussions about bias have pervaded news outlets, college campuses, workplaces, and numerous other settings, it seems many of us were not aware we had biases. This fits nicely with the concept of “unconscious bias”.

Bias can be neutral, good or bad based, on perspective.

Let’s use the example of potato crisps in discussing biases. I have a favourite potato crisp. It is actually the only one I buy. When I see a potato crisp display, all the other crisps are deemed unacceptable. To me and the manufacturer of my favourite crisps, it’s a positive bias. To the competitors of my favourite brand, it’s a negative bias. To my partner, it is neutral because they don’t care one way or the other which crisp I like.

A bias can also be an unconscious bias, in that a person is not aware that they have it. Let’s say you meet a stranger. They seem to be a nice person, but there is the usual distance between two strangers. While in discussion with this person, they mention one of their relatives and you realize you are related to this person; a close relative, actually. Suddenly, this person is not a stranger, they are family, and your attitude towards them may change drastically.

It is critical to create an appropriate set of test data to validate the trained AI’s output. The training set of data must be independent of the test data; using the training data as part of the test data would invalidate the testing.

It must be understood how bias is used in the selection of the data used to train the AI. The AI cannot have an unconscious bias reflected by the developers, as that could result in harm.

Some issues with biases and AI:

- The average computer scientist is not trained in psychology and does not understand how to develop software that manages the many types of bias. How does someone create an AI that has to manage biases when they don’t understand it themselves? In all development, end user needs are created with input from many disciplines. To create user needs that are further elaborated to become requirements that reflect biases as per intended purpose, it is recommended to add an appropriate resource who understands bias.
- Manufacturers currently do not have consistent, industry-wide tools to be used during development to manage bias. Manufacturers develop their own approaches with potentially varied and inconsistent results.
- Manufacturers need to understand how the bias embedded into the training data for the AI impacts the output of the software, as it relates to intended use or intended purpose.

MLMD training

Common sense is a term that refers to sound, practical or reasonable judgment that is shared by, or common to, most people.

My experience is that common sense is not all that common. How many times has someone needed help because they keep making mistakes that could have been easily avoided with use of the elusive “common sense”? As we all have experienced when doing a search on the internet, a large part of the task is creating the correct filter to find appropriate information, then parsing the choices presented for acceptability. While people use common sense, experience, education and the context of the data to find what they are looking for, that is not something an AI can easily perform at this time.

Imagine some of the choices you might make without a certain degree of common sense. Think about using GPS software that states the shortest path to a destination requires a right turn into the harbour. Because the GPS software said so, the driver made a right turn and ended up in a harbour.[13]

The AI only knows what it was trained to know. Its knowledge and flexibility is only as good as the programmers of the AI (at least at this point) and the training data. How the AI is trained becomes of critical importance for medical devices.

Regulatory guidance and approval to take the product to market are dependent on a well-defined set of requirements, intended use/intended purpose, and objective evidence that the software meets the design (verification) and end user needs (validation).

With the addition of AI into SiMD or SaMD, verification activities will include that the manufacturer provide additional objective evidence of the verification of the training data and verification of the test data (which should not be the same dataset as the training data). As with all medical device testing, the data used is derived from requirements and must be traced from requirements to test. The requirements associated with the use, content and biases associated with the training data and the test data must be fully defined.

Static-AI and Dynamic-AI

Static-AI: Training of the AI that occurs during the development process. The output of the AI is fully verified and validated. Additional learning, or training of the AI, does not occur after the medical device is released for use.

“Dynamic-AI: Training of the AI which occurs both during development and after the product is released. It requires much more oversight and management of risks post-release. An issue for Dynamic-AI is creating controls that keep the output within the verified and validated intended use / intended purpose and that does not cause harm because of unvalidated output.

For both Static- and Dynamic-AI, the required verification and validation of the data used to train the AI may add additional effort to the development process. Part of the verification may indicate the outcomes are equal to, or better than, what can be accomplished using current methodologies.

13 See https://www.insider.com/tourists-hawaii-gps-drove-car-into-water-2023-5

Static-AI training

Static-AI is most similar to current development methodologies: the final outcomes produced by MLMD are fully verified and validated prior to release of the product.

See below a list of example activities that should be considered per CR34971:2022:

- Create an AI training development plan
- Develop requirements for the training data while having the ability to demonstrate (via objective evidence) any issues have been mitigated with the following as appropriate (including, but not limited to):
  - Incorrect data
  - Handling of outliers
  - Incomplete data
  - Subjective data
  - Inconsistent data
  - Atypical data
- Training data should include all types of data the AI will be processing, including the good with the bad. Training data content should be based on requirements per BS EN ISO 13485:2016+A11:2021 and FDA 21 CFR Part 820.30, Quality System regulation. The manufacturer must verify the MLMD can manage all types of information it may see, to verify the MLMD is properly trained and processes the data per requirements
- Identify the source of the data
  - If DOUP, how was the data validated as appropriate for use and where was the data sourced?
  - Synthetic data might be considered DOUP, but it would have to be proven why that data met the requirements and was all inclusive of both good and bad
- Verify the AI can manage (per requirements) outliers, incomplete data, inconsistent data, marginal data, and atypical data
- Verify the data used is not overfitting[14] or underfitting[15]
- Verify the bias in the AI conforms to intended use
- Verify that as learning occurs, the Dynamic-AI stays within its approved intended use / intended purpose
- Based on the requirements and the risks associated with use of the data, developments will include risks associated with the use of the training data in the risk management plan and create appropriate mitigation within the requirements. The verification plan for the AIML would require test data for systems verification. That test data should not be the data that was used to train the AIML.

14 Overfitting: Creating a model which fits the training data too precisely and fails to generalize on new data.
15 Underfitting: Creating a model that does not fit the training data closely enough and produces incorrect predictions on new data.

Dynamic-AI training

A concern with Dynamic-AI is that learning occurs outside of verification and validation. How can the manufacturer confirm the verified and validated performance of the medical device, when the output, or results, may change in ways not envisioned by the developers? The risks of harm associated with Dynamic-AI must be mitigated; real-time, active post-market monitoring may be required. As with stem cells, some uses of Dynamic-AI in medical devices may be ahead of the curve of safe, verified and validated use. Failure of MLMD Dynamic-AI, where that failure could result in harm, is not acceptable.

Dynamic-AI requires conforming to all Static-AI deliverables for Static-AI; and, because learning never stops, the additional deliverables might include:

- post-market monitoring;
- ongoing human feedback; and
- “AI-Watchdogs” that would perform various activities to keep the medical device from unforeseen harm and ensure the medical device is limited to the verified and validated intended use. While the medical device can stay within the intended use, that does not mean potential new learning (which could potentially change, or impact intended use) could not be reported to the manufacturer for evaluation.

Monitoring and mitigation

MLMD may be impacted by medical device hardware changes

Medical device hardware or software changes can impact MLMD. For example, if a medical device's sensors degrade or are improved, how does this impact the MLMD? How will the AI be trained to recognize whether the input has deviated from what is expected? If technological changes are made to the medical device sensors, readers, or other technology, how will that impact the AI when it encounters something deviating from what it was trained to expect?

Post-market monitoring

Static-AI: The medical device can be treated like any other as it relates to post-market monitoring, provided that monitoring includes vulnerability assessments for the MLMD, and a robust complaint management system is looking for complaints related to the AI output.

Dynamic-AI: Keeping the AI constrained to a defined intended use / intended purpose, and to the verified and validated requirements may be a challenge. However, the use of AI-Watchdogs would be critical, as well as a more stringent oversight of the activities of the AI to ensure it is staying within its intended purpose.

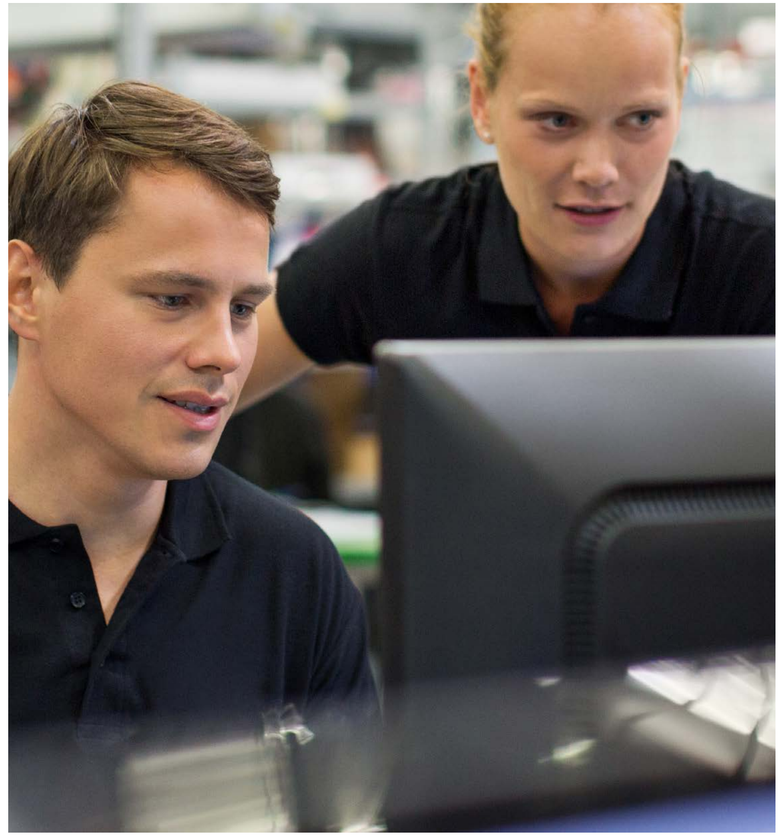

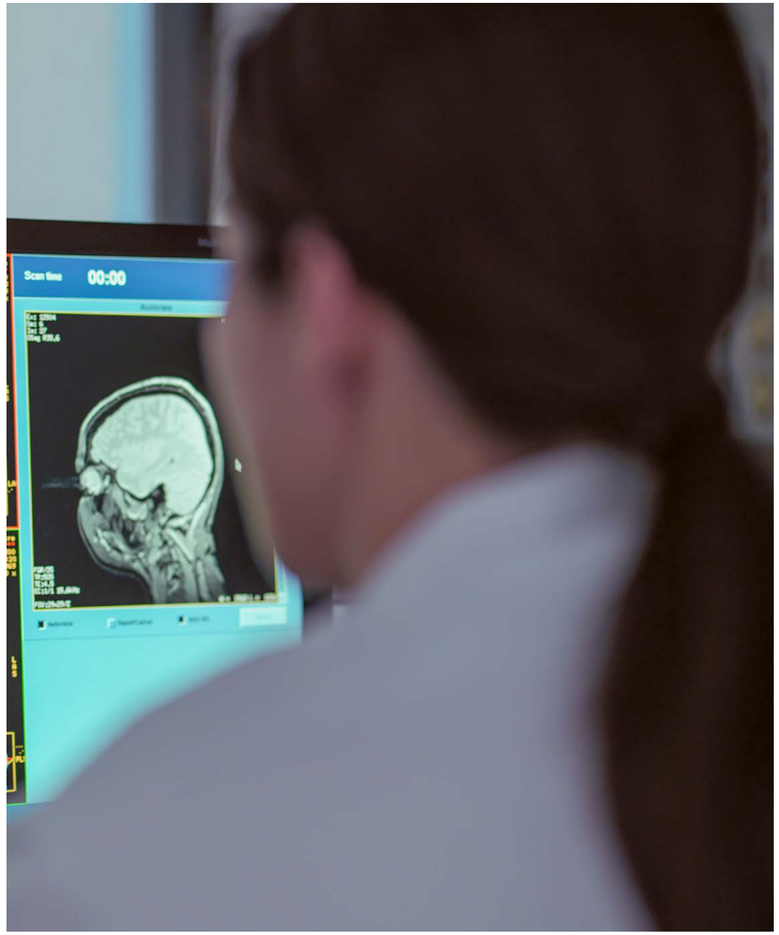

AI-Watchdogs which might be AIs

An AI-Watchdog is both an old and a new concept. It brings to mind the Master Control Unit (MCU) in the movie *TRON,* released in 1982. That movie is about a community of AIs governed by an MCU which is a powerful “virtual intelligence” that becomes “evil”. In today’s world, we would characterize the characters in the movie as AIs. The concept of all activities being monitored and managed by a system of AIs could be looked at for Dynamic-AI in particular, and it has been suggested that Static-AI may need these as well. There could be AIs that monitor data being used for learning, for monitoring and making changes to biases, AIs looking at feedback real-time, perhaps in concert with human interaction, a series of AI-Watchdogs that make sure that learning stays within intended use, remains trustworthy and does not add additional harm.

Two examples of proposed AI-Watchdog activities

MLMD Yellow-card: When the AI-Watchdog determines that the AIML finds deviations in input data or output data due to bias changes or other triggers, the AI-Watchdog throws a Yellow-card with the below results.

- Notifies the manufacturer of the issues which caused the Yellow-card.
- Solicits feedback from clinicians.

MLMD Red-card: When the AI-Watchdog determines that the AIML produces content that has deviated from the approved intended use and/or validated content, the AI-Watchdog will throw a Red-card, with the below results.

- Does not allow the output of the data that deviated from the approved intended use and/or validated content (deviated from requirements).
- Rollback learning to a prior “acceptable” state and puts further learning on hold.
- Notifies the manufacturer that the AIML needs to be evaluated and provides the data that triggered the Red-card.

Regulatory views and concerns – Dr. Aris Tzavaras (BSI)

AI-enabled medical devices have been around for a few years. Both FDA and EU Notified Bodies (NBs) have permitted access to market, assessing compliance to jurisdiction sectorial legislations and State of The Art (SoTA) which is usually expressed by voluntary compliance with relevant standards. What has changed that triggered a waterfall of actions around AI-enabled medical devices? There are several factors; however, the most important one is the understanding that current legislative frameworks and standards do not adequately address AI. AI is a different software technology that introduces new types of risks and therefore new controls, and mitigations need to be proactively in place for ensuring safety, performance and ethics.

The risks around AI are present in every life cycle phase, starting from the data collection, processing and datasets suitability for the task, AIML model explainability and AIML performance validation, to AI use and oversight, but extend further to the post-market phase as most AIML models will, ultimately, be designed to evolve. As most of the risks are linked to the intended purpose of the device, this is the starting point for every assessment. We are all familiar with the discussions around bias existing in datasets, but is bias on its own an issue? A mammography screening AI-enabled device is expected to have mainly female datasets for training and testing purposes, therefore be gender-biased in the narrow interpretation of the term.

New legislations around AI are under development, with the most advanced being the EU AI Act (AIA). The AIA approaches AI compliance by setting high-level requirements; meeting those requirements is presumed by compliance to appropriate standards. ISO/IEC JTC 1/SC 42, ISO's Artificial Intelligence committee, has already published 25 standards and 28 more are under development. Similarly, other organizations like IEEE and BSI have already published their own AI-relevant standards (CR34971:2022; BS 30440:2023, *Validation framework for the use of artificial intelligence (AI) within healthcare — Specification)*. It is obvious that the SoTA is evolving fast and both regulators and manufacturers will need to adapt.

Currently, most AI-enabled medical devices are narrow, designed for a specific task (usually with little or no autonomy) and best described as a Clinical Decision Support System (CDSS). A different type of risk is introduced by the use of such AI systems. Such risks include, but are not limited to, over- (automation bias) or under-trusting AI decisions, users' ability to interpret AI decisions, and users' ability to identify and understand the causes of AI going wrong, such as data drift and robustness issues. Some of those risks are relevant to the transparency of AI systems, in terms of sufficient information and warning to users around the capabilities and limitations of the devices use.

The thorn of evolving (dynamic) AI is addressed across jurisdictions by the use of the term "predetermined changes". Predetermined changes as described under FDA guidance documents are protocols that describe datasets' relevant processes, models' characteristics to be modified, re-training trigger criteria, methods and metrics for model evaluation and action when criteria are not met as well as a monitoring plan. The latter might be much harder to implement by manufacturers as it could be the case that AI is evolving differently in multiple settings or on a personalized patient level.

Conclusion

The inclusion of AI is exciting and there is widespread interest in being able to claim the use of AI in medical devices.

I do not believe neither regulatory bodies, nor manufacturers, currently have the development and monitoring tools, regulation, and standards to manage Dynamic-AI or, arguably, Static-AI. The difficulty will be keeping the AI constrained to a defined intended use / intended purpose, and to the verified and validated requirements.

The use of Static-AI closely aligns with current development practice, but the inclusion of AI may bring significant extra overheads to the development process. Manufacturers will have to verify training data was appropriate for use, aligned with requirements for the data (including bias) and with full traceability to requirements, and to the verification of the training data. Additionally, it has been suggested by peers that the use of Static-AI needs the type of AI-Watchdogs that I suggested for Dynamic-AI.

Using Dynamic-AI may require significant additional work by the manufacturer. In addition to demonstrating that the MLMD conforms to its intended use and purpose upon product release, the manufacturer must also verify that the MLMD continues to conform to its intended use over time as it continues to learn after the product is released.

It may be too early to integrate MLMD into some medical device applications, as the tools, regulations, and standards for MLMD development are not yet ready. Furthermore, there is no clear direction on how to manage MLMD using Dynamic-AI post-market, nor are there guidelines or standards that define the controls that must be in place to monitor the medical device, nor actions to take when the medical device no longer conforms to the approved intended use / intended purpose.

We are at a crossroads of the old and new. The application of AI in medical devices is similar to Damocles' sword: while there is potential for great benefits, there is also the risk of failure that can result in injury or death to a patient. We must choose wisely, as the wrong choice can result in harm and can set back the use of AI in medical devices for years.

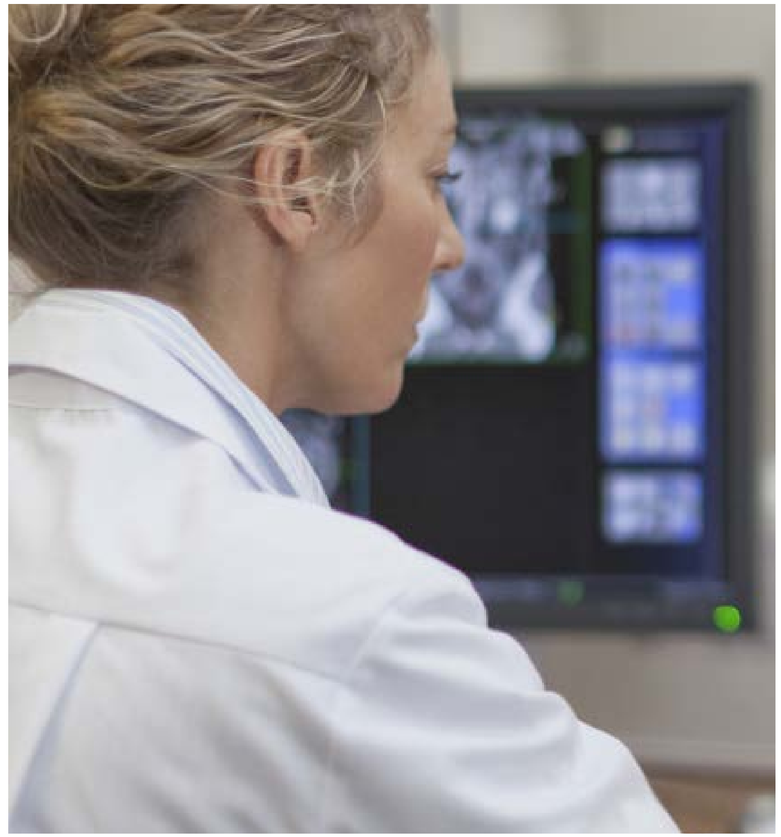

References

Draft EU Artificial Intelligence Act: https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence

IMDRF/AIMD WG/N67, *Machine Learning-enabled Medical Devices: Key Terms and Definitions*

BS EN ISO 14971:2019+A11:2021, *Medical devices — Application of risk management to medical devices*

BS EN 62304:2006+A1:2015, *Medical device software — Software life cycle processes*

CR34971:2022, *Guidance on the application of ISO 14971 to artificial intelligence and machine learning*

BS EN ISO 13485:2016+A11:2021, *Medical devices — Quality management systems — Requirements for regulatory purposes*

Clinical Decision Support Software Guidance for Industry and FDA Staff September 2022 https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software

Finding the Origination Point: Understanding Our Biases to Create a More Peaceful World, Bill de la Cruz

Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/ Machine Learning (AI/ML)-Enabled Device Software Functions

WHO urges caution with healthcare AI deployments: https://www.healthcareitnews.com/news/who-urges-caution-healthcare-ai-deployments

Rapid AI adoption could cause medical errors, patient harm, WHO warns, urging oversight: https://www.medtechdive.com/news/WHO-artificial-intelligence-AI-caution/650527/

ISO/IEC JTC 1/SC 42 Artificial intelligence: https://www.iso.org/committee/6794475.html

Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/ Machine Learning (AI/ML)-Enabled Device Software Functions, Draft Guidance for Industry and Food and Drug Administration Staff, April 2023: https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial

Authors

Author

Rich de la Cruz is President of Silver Lake Group, Inc. He has worked in the regulated medical industry for more than 35 years and has helped develop and/or manage 25 medical devices, including infusion pumps, pain pumps, haemodialysis, and peritoneal dialysis instruments, Class I and II clinical software and systems, pharmacy systems, and electronic medical record systems. Rich was the co-chair of the international working group IEC 62304 from 2015 to 2021 and actively contributes to standards and technical information report development.

Contributing Author

Dr. Aris Tzavaras is Head of Artificial Intelligence Notified Body at BSI Group, The Netherland B.V.

Reviewers

Pat Baird, Senior Regulatory Specialist, Philips

Pat Baird works at Philips as a Senior Regulatory Specialist, with a focus on the use of AI in healthcare. Pat likes to think of his job as “Policy Engineering” – understanding the unmet needs (and frustrations) of regulators and developers, and developing standards, white papers, and training to meet those needs. He co-chairs multiple committees related to artificial intelligence at AAMI, ISO, CTA, AdvaMed, and the Association of Food and Drug Officials (AFDO) and the Regulatory Affairs Professionals Society (RAPS). He is involved with other software committees regarding topics such as cloud services for a regulated environment, risk management and cybersecurity, and was a sub-team lead for the IMDRF AI for Medical Devices committee.

Jane Edwards, Head of Communications, Global Product Management, BSI

Jane holds a BSc in Chemistry and an MBA from Durham University. She has over 13 years’ experience in the medical device industry, having previously worked for Coloplast in their ostomy and continence business. Jane’s experience includes working within the pharmaceutical, chemical and telecoms industries for Glaxo Wellcome, ICI and Ericsson, allowing her to bring depth of knowledge from across many industries and technologies. Her current role in BSI allows her to work with technical reviewers across all disciplines ensuring that all BSI communications are accurate and relevant. She is a member of the European Medical Writers Association.

Paul Sim, Medical Devices Knowledge Manager, BSI Standards

Paul has worked in the healthcare industry for over 35 years, joining BSI in 2010 to lead the organization in Saudi Arabia where it had been designated as a Conformity Assessment Body. Later, he managed BSI’s Unannounced Audits programme. Since October 2015, he has been working with both the Notified Body and Standards organizations looking at how best to use the knowledge, competencies and expertise in both. Previously he held senior RA/QA leadership positions at Spacelabs Healthcare, Teleflex Medical, Smiths Medical and Ohmeda (formerly BOC Group healthcare business). Paul is a member of the Association of British Healthcare Industries (ABHI) Technical Policy Group and Convenor of the ABHI ISO TC 210 Mirror Group. He is Convenor of the BSI Committee that monitors all of the work undertaken by ISO TC 210, and Convenor of the BSI Subcommittee dealing with quality systems. As UK Delegation Leader to ISO TC 210, he is also actively involved in the work of national, European and international standards’ committees.

Eamonn Hoxey, Director, E V Hoxey Ltd

Eamonn is a technical author, trainer and consultant in a range of life science areas including regulatory compliance, quality management, sterility assurance and standards development. He worked for Johnson & Johnson for 17 years in positions of increasing responsibility for Quality and Regulatory Compliance for medical devices, pharmaceuticals and consumer products, including Vice President of Compliance, Vice President of Market Quality and leading quality implementation for the EU medical devices regulation for J&J's Medical Devices companies. Prior to joining J&J, Eamonn spent 16 years with the UK Medical Devices Agency, including six years as Head of Device Technology and Safety. Eamonn is currently chair of ISO TC 198, 'Sterilization of Healthcare products', chair of CEN TC 204, 'Sterilization of medical devices' and past chair of ISO TC 210, 'Quality management and related general aspects for medical devices'. He received the BSI Wolfe Barry medal in 2016 for his contribution to standards development.

Anette Sjögren, Consultant, Preventia

Anette has over 30 years' experience in the medical device and pharma industries, as QA, RA, QP and tasks within medical fields such as quality and risk management, clinical affairs, toxicology and biocompatibility. She has been a member of the Swedish (TK 355) and the international technical (TC 210) committees since 2010. Anette has been a consultant with Preventia since 2003. She holds an MSc in Biomedicine.

Published white papers

The growing role of human factors and usability engineering for medical devices: What's required in the new regulatory landscape? Bob North

The differences and similarities between ISO 9001:2015 and ISO 13485:2016: Can we integrate these quality management standards? Mark Swanson

Cybersecurity of medical devices: Addressing patient safety and the security of patient health information, Richard Piggin

General Safety and Performance Requirements (Annex 1) in the New Medical Device Regulation: Comparison with the Essential Requirements of the Medical Device Directive and Active Implantable Device Directive, Laurel Macomber and Alexandra Schroeder

Do you know the requirements and your responsibilities for medical device vigilance reporting? A detailed review on the requirements of MDSAP participating countries in comparison with the European Medical Device Regulation 2017/745, Cait Gatt and Suzanne Halliday

Nanotechnology: What does the future look like for the medical devices industry? Professor Peter J Dobson, with Dr Matthew O'Donnell

Recent advancements in AI – implications for medical device technology and certification, Anil Anthony Bharath

The impact and potential for 3D printing and bioprinting in the medical devices industry, Kenny Dalgarno

Sterilization – Regulatory requirements and supporting standards, Eamonn Hoxey

Medical device clinical investigations – What's new under the MDR? Maria Donawa

The convergence of the pharmaceutical and medical devices industries: Navigating the innovations and regulations, Barbara Nasto and Jonathan Sutch

Phthalates and endocrine disruptors – An overview of their safety requirements and evaluations and the standards that support them, Benjamin Seery

European Union Medical Device Regulation and In Vitro Device Regulation: unique device identification: What is required, and how to manage it, Mary Gray

Person responsible for regulatory compliance (PRRC) – MDR/IVDR Article 15: An overview of the requirements and practical considerations, Anne Jury and Maddalena Pinsi

Guidance on MDCG 2019-9: Summary of Safety and Clinical Performance, Amie Smirthwaite

Clinical evaluation under EU MDR, Amie Smirthwaite

Medical device clinical investigations — What's new under the MDR? An update, Maria Donawa

Using Standards to Demonstrate conformity with Regulations, Eamonn Hoxey

Performance Evaluation for IVD, Fiona Gould

Implementing the European Union Medical Devices Regulations – Are you prepared? (revised), Eamonn Hoxey

Developing and maintaining a quality management system for IVDs (revised), Melissa Finocchio

Ethical and trustworthy Artificial Intelligence: BSI's introduction to the European Artificial Intelligence Draft Act (AIA), Dr. Aris Tzavaras

The differences and similarities between 9001 and 13485 (revised), Mark Swanson

The European Medical Devices Regulations: What are the requirements for vigilance reporting and post-market surveillance? (revised), Eamonn Hoxey

EUDAMED – an overview of how it is being developed and deployed by the European Commission, Olga van Grol-Lawlor

Technical Documentation under the Medical Device and In Vitro Diagnostic Regulations (MDR and IVDR) (revised), Lydie Moreau

Forthcoming white papers

Requirements of EU-GDPR and PMCF studies, registries and surveys under the MDR

Symbols and information to be provided with medical devices

The UKCA mark

bsi

Read more about our services and products on our website **bsigroup.com**