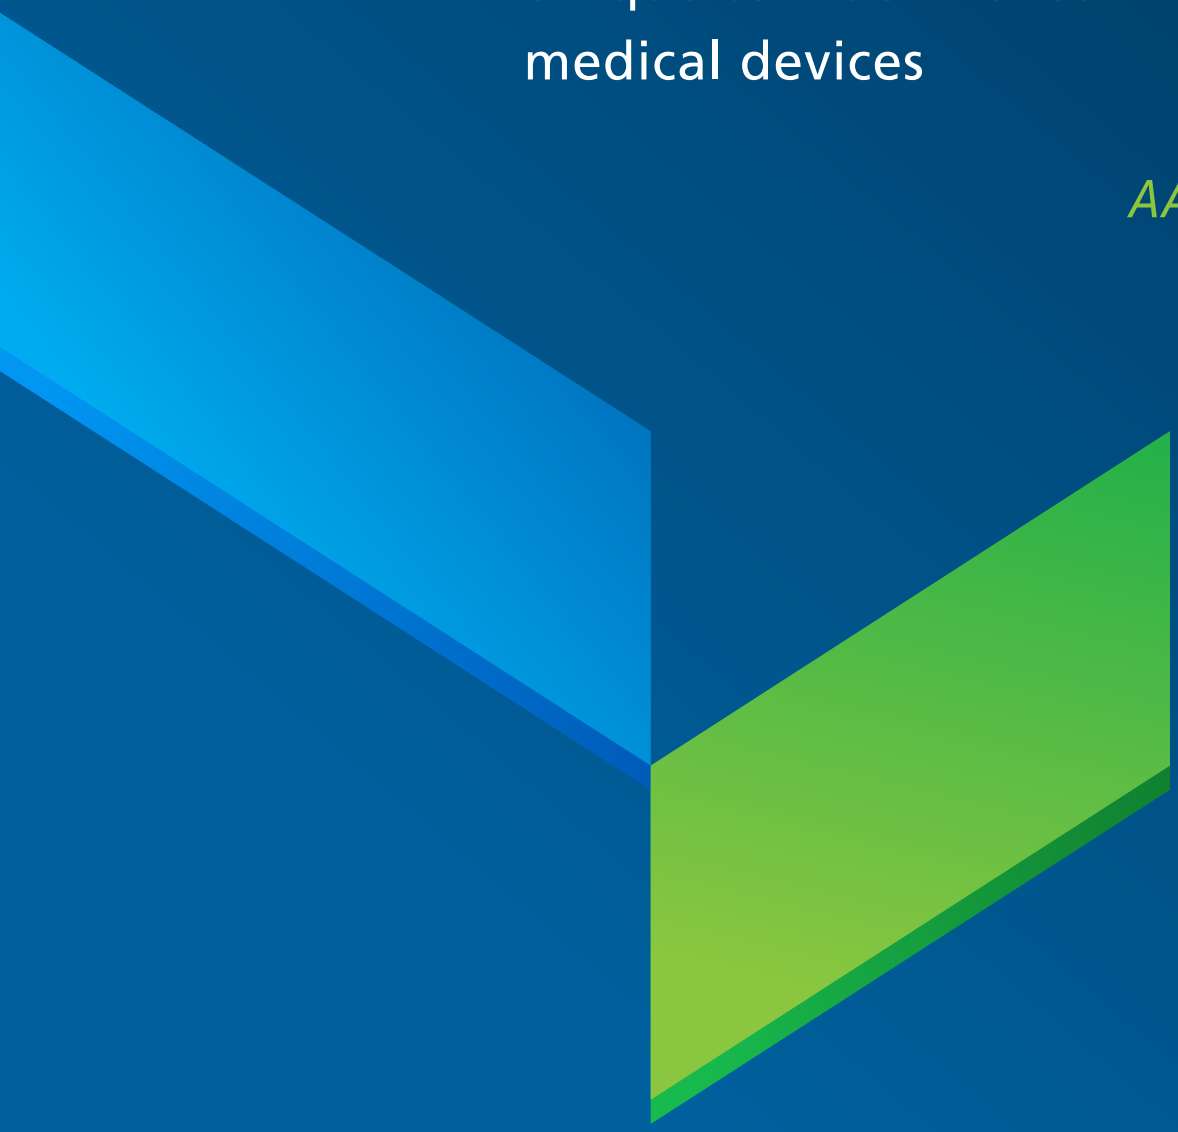
AAMI
Advancing Safety in Health Technology

AAMI Consensus Report

Cybersecurity considerations
unique to machine learning–enabled medical devices

AAMI CR515:2025



Cybersecurity considerations unique to machine learning–enabled medical devices

Approved 23 June 2025 by
AAMI

Abstract: This consensus report specifically addresses the unique cybersecurity threats associated with developing and deploying machine learning–enabled medical device software, as distinct from the broader cybersecurity risks applicable to all medical device life cycle phases. This includes potential cybersecurity threats that can arise from or during data collection, product design, product deployment, product use, and maintenance.

Keywords: machine learning, cybersecurity threats, medical devices

AAMI Consensus Report

A consensus report (CR) is a publication of the Association for the Advancement of Medical Instrumentation (AAMI) developed to provide concise, prompt, and practical guidance on narrowly focused topics of high importance to the health technology community. A consensus report is intended to provide initial consensus guidance in response to an urgent/immediate need for guidance in the following instances:

- While more robust data/information develops on emergent areas
- When variation in the development, implementation or use of a product or process exists
- When existing standards or other documents require additional context/clarification

A consensus report is not subject to the same formal process as a standard and while similar in nature to a technical information report (TIR), a CR is based on the collective knowledge and experience of a selected group of stakeholders and has not undergone the wider reviews of a TIR or standard and offers an even greater response time.

CAUTION NOTICE: This AAMI CR may be revised or withdrawn at any time. Because it addresses a rapidly evolving field or technology, readers are cautioned to ensure that they have also considered information that may be more recent than this document.

All standards, technical information reports, consensus reports and other types of technical documents developed by AAMI are *voluntary*, and their application is solely within the discretion and professional judgment of the user of the document. Occasionally, voluntary technical documents are adopted by government regulatory agencies or procurement authorities, in which case the adopting agency is responsible for enforcement of its rules and regulations.

Comments on this technical information report are invited and should be sent to AAMI, Attn: Standards Department, 901 N. Glebe Rd, Suite 300, Arlington, VA 22203.

Published by

AAMI
901 N. Glebe Rd., Suite 300
Arlington, VA 22203
www.aami.org

© 2025 by the Association for the Advancement of Medical Instrumentation

All Rights Reserved

Publication, reproduction, photocopying, storage, or transmission, electronically or otherwise, of all or any part of this document without the prior written permission of the Association for the Advancement of Medical Instrumentation is strictly prohibited by law. It is illegal under federal law (17 U.S.C. § 101, *et seq.*) to make copies of all or any part of this document (whether internally or externally) without the prior written permission of the Association for the Advancement of Medical Instrumentation. Violators risk legal action, including civil and criminal penalties, and damages of $100,000 per offense. For permission regarding the use of all or any part of this document, visit the Copyright Clearance Center.

Printed in the United States of America

ISBN 978-1-57020-912-3

Contents

Page

Task Group representation iv
Introduction vi
1 Scope 1
2 References and resources 1
3 Terms and definitions 2
4 Threat modeling for MLMDs 2
5 ML algorithm life cycle description 3
6 Unique threats, vulnerabilities, and mitigations for MLMDs 7
Annex A (informative) Data validation 15
Bibliography 16

Task Group representation

Association for the Advancement of Medical Instrumentation

Artificial Intelligence Committee Members

This AAMI consensus report (CR) was developed by a task group under the auspices of the AAMI Artificial Intelligence (AAMI/AI) Committee.

At the time this document was published, the AAMI Artificial Intelligence committee had the following members:

Cochairs:
Pat Baird
Jesse Ehrenfeld

Members:
Pat Baird, Philips
Prithul Bom, Regulatory Technology Services (RTS)
Moe Chan, Alcon Laboratories Inc
Yu Chen, Draeger Inc.
Tatiana Correia, Solventum (fka 3M Health Care)
Richard De La Cruz, Silver Lake Group Inc
Jesse Ehrenfeld, Medical College of Wisconsin
Chase Elms, (Individual)
Anura Fernando, UL LLC
Ajay Gopinath, Abbott Laboratories
Jason Gorman, ResMed Pty Ltd
James Hamlyn, Skin Analytics Ltd.
Jeanetta Harris, (Individual)
Tim Jackson, Arthrex Inc
Patricia Krantz-Zuppan, Medtronic Inc Campus
Christine (Tina) Krenc, KTA Compliance Consulting
Li Li, (Individual)
Larry Lugo, Lugo & Associates LLC
Karana M Pierre, Karana M Pierre Person
Ian Marsden, Pfizer Parenteral Center of Excellence
Magnus Miller-Wilson, Boston Scientific Corporation
Haseeb Mohammed, AbbVie
Bharath Motamarri, Cordis US Corp
Marc Neubauer, Baxter Healthcare Corporation
Susumu Nozawa, Siemens Healthineers
Andrew O'Keeffe, Olympus America Inc
Alexander Paczynski, Cardinal Health
Siva Paramanandam, Amgen Inc
Deborah Perkins, Johnson & Johnson
Mike Powers, St Lukes Hospital and Health Network
Luke Ralston, US FDA/CDRH
Manjunath Bisalehalli, Intuitive Surgical Inc
Leigh Ruiz, (Individual)
Felicia Sheriff, (Individual)
Dina Sifri, Meddevsoft
Dylan Sinks, Eli Lilly & Company
Dr. Jeffrey Smoot, (Individual)
Chaitanya Srinivasamurthy, ICU Medical Inc
Daniel Stoller, SVQC Professional Services
Mark Supanich, (Individual)
Arleen Thukral, Veterans Administration (VA) Central Office
Jerica Vie, Stryker
Monica Walker, Insulet Corporation
Hollis Whittington, Micro Systems Engineering Inc
Andrew Wingen, Bayer Medical Care, Inc.
Fubin Wu, GessNet
Ying Xiao, (Individual)

Alternates:

George Abouzei, Insulet Corporation
AmirHessam Aminfar, Amgen Inc
Amen Arfi, AbbVie
Morgan Beaver, Cardinal Health
Jan Bom, Regulatory Technology Services (RTS)
Jason Fuller, Stryker
Shelly Leacock, Veterans Administration (VA) Central Office
Bill Liao, US FDA/CDRH
Bill Prucka, Eli Lilly & Company
Elena Raja, Werfen USA
Ravi Reddy, Micro Systems Engineering Inc
Chris Reed, Medtronic Inc Campus
Mark Rohlwing, ICU Medical Inc
Williams Sammons, Baxter Healthcare Corporation
Nisha Shah, Pfizer Parenteral Center of Excellence
Jason Smith, UL LLC
Jacob Sunol, Abbott Laboratories
Vinay Tikka, Arthrex Inc
Jason Tjia, ResMed Pty Ltd
Eric Wespi, Boston Scientific Corporation
Michael Yramategui, Intuitive Surgical Inc

The committee would like to extend its appreciation to the project leaders and the task group for their outstanding guidance, dedication, and leadership in the development of this consensus report. Their expertise and commitment were instrumental in shaping the quality and depth of the final work:

Pat Baird, Philips/Project lead
Arvin Eskardania, FDA CDRH/Project lead
Steven M Christey, MITRE
Zack Hornberger, AdvaMed
Randy Horton, Orthogonal
Michelle Jump, MedSec
Patricia Krantz-Zuppan, Medtronic
Christine “Tina” Krenc, KTA Compliance Consulting
Catherine Lowe, LoweRisk QA Consulting
Robert Turpin, BSI
Varun Verma, Philips
Jessica Wilkerson, FDA CDRH
Axle Wirth, MedCrypt

NOTE—Participation by federal agency representatives in the development of this document does not constitute endorsement by the federal government or any of its agencies.

Introduction

Cybersecurity incidents in healthcare have become increasingly widespread and disruptive in recent years. The healthcare sector is a prime target for cyber threats due to the criticality of its services and the vast amount of valuable patient data involved. The importance of cybersecurity for medical device technologies cannot be overstated, as it directly impacts patient safety, data security, and the overall integrity and availability of healthcare systems. Consequently, securing medical devices from cybersecurity incidents is a key priority for healthcare organizations, medical device manufacturers, and regulatory agencies.

Mature medical device technology manufacturers and developers may be well versed in securing traditional software functions in devices; however, increasing adoption of new technologies and the associated cybersecurity complications deserves more attention, particularly in the healthcare space. From the medical device manufacturer or software developer perspective, several measures are being taken to enhance the cybersecurity of medical device technologies. Manufacturers are increasingly adopting secure product development life cycles and incorporating security-by-design principles into the development of medical device technologies. At the fundamental level, this involves, but is not limited to, identifying potential cybersecurity risks and vulnerabilities across the total product life cycle, and implementing security controls accordingly. The U.S. Food and Drug Administration (FDA) has been rigorously evaluating devices for cybersecurity through its premarket review process, now further strengthened by the official authority granted under section 524B of the FD&C Act. Additionally, the FDA requires manufacturers' adherence to cybersecurity practices in the postmarket, as recommended in the FDA's postmarket policies, and as required by section 524B. This comprehensive approach ensures robust cybersecurity measures are in place across the total product life cycle (TPLC).[1,2]

A field that has grown rapidly in the past several years is artificial intelligence (AI) and, more specifically, its subset machine learning (ML). Developers of medical technologies are increasingly harnessing the promise that AI/ML holds, utilizing and incorporating it into software used in medical devices and healthcare applications. In the recent past, the U.S. Food and Drug Administration has cleared a large number of AI/ML–enabled medical devices[3] and this number is only expected to grow along with the need to increase awareness of the distinctive associated cybersecurity challenges.

While ML–enabled medical devices (MLMD) are likely susceptible to similar cybersecurity risks as traditional medical devices, there are certain unique security risks that accompany the development and utilization of MLMD that require more attention from medical device manufacturers (MDM) and developers. Examples include the following:

1. Development of MLMD is unique in that it is data-centric, and instead of an explicit algorithm that directly solves a problem, ML models learn patterns and make predictions based on data. This means that collecting, cleaning, and preparing high-quality data is a significant part of the development process. Handling sensitive data in ML applications requires additional considerations to protect against cybersecurity incidents.
2. In ML, a crucial step is training the ML model. This involves feeding it annotated data to learn patterns and relationships. The model adapts and improves its performance over time through this training process, unlike traditional software development that may leverage preprogrammed algorithms.
3. ML models are also susceptible to unique risks due to data leakage. ML Data Leakage (Target Leakage) refers to the introduction of information from outside the training dataset or from future information that wouldn't

[1] See *FDA guidance, Cybersecurity in medical devices: Quality system considerations and content of premarket submissions,* available at www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-medical-devices-quality-system-considerations-and-content-premarket-submissions.

[2] See the *FDA guidance, Postmarket management of cybersecurity in medical devices,* available at www.fda.gov/regulatory-information/search-fda-guidance-documents/postmarket-management-cybersecurity-medical-devices.

[3] See the the FDA's updated list, *Artificial intelligence-enabled medical devices,* available at www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices.

© 2025 Association for the Advancement of Medical Instrumentation ■ AAMI CR515:2025

be available during real-world predictions. This causes the model to learn patterns that will not generalize to new, unseen data, leading to inflated accuracy during evaluation and poor performance in real applications.

4. ML models can produce unexpected results, especially when dealing with complex, real-world data. These unexpected results may happen due to a cyber threat or misconfigurations during the training phase.

MLMD present unique cybersecurity concerns and threat vectors. These concerns stem from the complexity of ML models, the reliance on data, and the potential for unique adversarial attacks.

This consensus report (CR) describes cybersecurity considerations related to development of MLMD, including concerns, such as privacy, integrity (input and output), model security, and reduced performance due to potential attacks, among others. Information on associated vulnerabilities in ML models and potential mitigation techniques are also identified where possible.

Addressing unique cybersecurity concerns related to MLMD software development requires a multidisciplinary approach that combines medical device development and cybersecurity expertise with domain-specific knowledge in data science and ML. It also necessitates a proactive plan for data quality, model robustness, and ongoing cybersecurity measures to protect against potential cyber threats and vulnerabilities. Securing medical device technologies from cyber threats is an ongoing effort, and the healthcare industry should continuously work to improve the security of these critical technologies given that the threat landscape is continually evolving. Cybersecurity measures must adapt accordingly to address new and emerging threats, including those that manifest from adopting and incorporating MLMD in the healthcare space.

ML is a rapidly evolving field, and multiple standards and guidance documents are continually being developed. While, at the time of writing, this document is aligned with other publications, things can change over time, resulting in slightly different advice. For example, multiple standards projects are developing an algorithm life cycle similar to the one used in section 5, but there is no current consensus on exactly what the life cycle should look like. Other organizations have published related documents**;** for example, the Cybersecurity & Infrastructure Security Agency (CISA) has created a roadmap for AI,[4] and NIST has developed a taxonomy of attacks and mitigations for ML systems.[5]

Although this team could have waited until there was more agreement across various stakeholders, we felt that it was important to publish this CR as quickly as possible. This committee plans to develop a TIR that will explore the cybersecurity challenge more deeply, and we hope there will be more harmonization across the various stakeholders at that time.

[4] See the website of the America's Cyber Defense Agency: National Coordinator for Critical Infrastructure Security And Resilience, at www.CISA.gov/ai.

[5] See NIST AI 100-2 E2023, *Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations,* at https://csrc.nist.gov/pubs/ai/100/2/e2023/final.

Cybersecurity considerations unique to machine learning–based medical devices

1 Scope

This consensus report specifically addresses the unique cybersecurity threats associated with developing and deploying machine learning (ML)–enabled medical device (MLMD) software, as distinct from the broader cybersecurity risks applicable to all medical device life cycle phases. This includes potential cybersecurity threats that can arise from or during data collection, product design, product deployment, product use, and maintenance. While MLMD software, like any software-based technology, is susceptible to traditional cybersecurity threats and vulnerabilities, this document focuses only on those threats that are unique to ML–based technologies or that may differently challenge such technologies.

This consensus report covers various cybersecurity threats that can manifest across the data life cycle and the total product life cycle of MLMD software, including associated vulnerabilities and potential mitigation strategies (see Figure 1 and Table 1).

This document is intended to supplement the information on security risk management of medical devices provided in AAMI SW96. It does not alter the security risk management process described in AAMI SW96, but rather, it provides information and guidance to supplement the application of SW96 to MLMD software.

It is assumed that the reader is familiar with the basic concepts and terminology associated with medical device software development, cybersecurity, and machine learning. A list of references that the reader may find helpful in gaining a high-level understanding of cybersecurity concepts for MLMDs is provided in the Bibliography (Appendix B) along with all documents referenced in this consensus report.

2 References and resources

AAMI SW96:2023, *Standard for medical device security — Security risk management for device manufacturers*

AAMI TIR66:2017/(R)2020, *Guidance for the creation of physiologic data and waveform databases to demonstrate reasonable assurance of the safety and effectiveness of alarm system algorithms*

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

ISO/IEC 22989: *Information technology — Artificial intelligence — Artificial intelligence concepts and terminology*.

ISO/IEC 2382:2015, *Information technology — Vocabulary*

ISO/IEC 5469:2023, *Artificial intelligence — Functional safety and AI systems*

© 2025 Association for the Advancement of Medical Instrumentation ■ AAMI CR515:2025