Technical Specification

ISO/TS 24971-2

First edition
2026-06

# Medical devices — Guidance on the application of ISO 14971 —

## Part 2: Machine learning in artificial intelligence

*Dispositifs médicaux — Recommandations relatives à l'application de l'ISO 14971 —*

*Partie 2: Apprentissage automatique en intelligence artificielle*



Reference number
ISO/TS 24971-2:2026(en)

**COPYRIGHT PROTECTED DOCUMENT**



ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

# Contents

Page

# Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

ISO draws attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO takes no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents. ISO shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html.



This document was prepared by Technical Committee ISO/TC 210, *Quality management and corresponding general aspects for products with a health purpose including medical devices*, in collaboration with Technical Committee IEC/TC 62, *Medical equipment, software, and systems*, Subcommittee SC 62A, *Common aspects of medical equipment, software, and systems*, and with the European Committee for Standardization (CEN) Technical Committee CEN/CLC/JTC 3, *Quality management and corresponding general aspects for medical devices*, in accordance with the Agreement on technical cooperation between ISO and CEN (Vienna Agreement).

A list of all parts in the ISO 24971 series can be found on the ISO website.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

# Introduction

Artificial intelligence (AI) is rapidly advancing and offers transformative potential for the healthcare sector. These advantages can be related to improved *benefits* for the patient, increased efficiency in clinical workflows and more effective management of healthcare overall. However, the implementation of new technologies such as AI can also present new *risks* and can, for example, jeopardize patient *safety*, affect privacy and security, influence user actions, undermine trust in healthcare or adversely affect the management of healthcare.

The *safety* and effectiveness of AI in *medical devices* was explored in an AAMI-BSI document[16], which identified three ways in which AI-based *medical devices* differed from "traditional" (non-AI) *medical devices*:

a) *Training*. These *medical devices* can process large amounts of data and learn from these data to improve their results. Thus, they can have positive effects on patient health within the scope of the *intended use* of the *medical device*.

b) Level of autonomy. These *medical devices* can generate different treatment options, select the best option based on a trained model and execute the selected option (see for example IEC/TR 60601-4-1[8]). These steps can be performed with reduced or even without direct user action. Therefore, human oversight is critical for their safe application.

c) *Explainability*. These *medical devices* often rely on complex algorithms and large datasets to generate output. However, the inherent opacity of these algorithms makes it challenging to interpret how specific conclusions or recommendations are derived. This aspect is sometimes referred to as the "black-box problem". This can make it difficult for clinicians, other healthcare personnel and individuals without specialist knowledge to understand the rationale for the output.

Many different AI-based technologies and algorithms exist today, which are differentiated by data topology and model architecture, including decision trees, genetic algorithms and deep learning-based technologies such as generative AI and neural networks. ISO/IEC 22989[4] and ISO/IEC 23894[6] provide general guidance on AI concepts, terminology and *risk management*, but they do not specifically address the application of AI to *medical devices*. It is noted that "risk" is defined in these documents as the effect of uncertainties on objectives (see also ISO 31000[3]). This definition is useful for organizational or business *risk management*. The term "*risk*" used in the healthcare sector is different and is defined in ISO 14971:2019 as the combination of the probability of occurrence of *harm* and the *severity* of that *harm*.

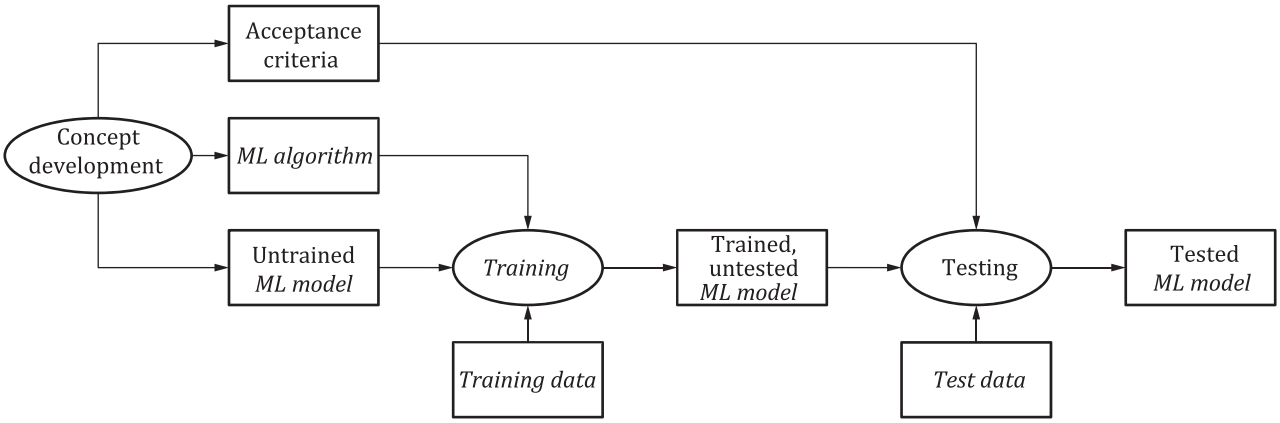


**Figure 1 — Concept development, *training* and testing stages of the *ML model* and its relationship with the *ML algorithm*, the *training data* and the *test data***

This document focuses on *machine learning* (*ML*) techniques and is restricted to *ML-enabled medical devices* (*MLMD*). *Machine learning* is considered a subset of AI that involves an *ML model* and an *ML algorithm*. See Figure 1. The *ML model* and the *ML algorithm* are the results of the concept development for a new *MLMD*, together with acceptance criteria for the eventual *MLMD*. It is noted that the supporting infrastructure (computing framework, hardware, network and other IT components) can be an important aspect in concept development. The *MLMD* acceptance criteria are different from the criteria for *risk* acceptability. It

is important to establish the *MLMD* acceptance criteria at the start, as part of concept development, and not at the end of the *MLMD* development. Otherwise, the results of *MLMD* testing could influence the decisions when establishing those criteria.

After concept development, the *ML model* is trained by using an *ML algorithm* enabling it to learn patterns from *training data* without being explicitly programmed. Next, the trained *ML model* is applied to *test data* to verify its performance. The *training data* and the *test data* are different (disjoint) sets. They can be actual patient data or synthetic data, i.e. data created to simulate a patient for *training* or testing purposes. The tested *ML model* can then be applied to new patient data in a clinical setting. More information on *MLMD* can be found in IMDRF documents N67[21] and N88[23] and in guidance documents[17][18] from FDA, Health Canada and MHRA.

It is recognized that the *ML model* can require retraining after a period of use to redefine its parameters and to ensure its continued performance. This can be achieved by planned retraining with collected patient data or on a continuous basis with each new patient data. The latter is referred to as “continuous learning” throughout this document and sometimes as “adaptive” in other documents.

All *medical devices* come with inherent *risks*. *Manufacturers* are required to demonstrate that their *medical devices* do not pose unacceptable *risks*, and that the *benefits* of the *intended use* outweigh the overall *residual risk*. ISO 14971 details how *manufacturers* can identify, assess and control *risks* to protect the patients, the users and other persons as well as property (for example objects, data, other equipment) and the environment. This includes *risks* related to data and systems security and cybersecurity. Guidance on the application of ISO 14971 is provided in document ISO/TR 24971[2]. Additionally, IEC 80001-1[11] and IEC/TR 80002-1[12] address software internal to a *medical device* that can support AI or *ML*.

This document was developed to provide specific guidance on the application of ISO 14971 to *MLMD*. It does not provide a new *risk management process*, nor does it expand the requirements of ISO 14971. This document addresses *risks* related to *machine learning* and topics such as data management, feature extraction, unwanted *bias*, information security, *training* the *ML model* by an *ML algorithm*, evaluation and testing of the trained *ML model*. See Annex A for an explanation of *bias*. The report AAMI TIR34971[14] provided valuable input for this document.

# Medical devices — Guidance on the application of ISO 14971 —

# Part 2: Machine learning in artificial intelligence

## 1 Scope

This document provides guidance on *risks* specific to artificial intelligence (AI) and *machine learning* (*ML*) and how to apply the *risk management process* of ISO 14971 to *ML-enabled medical devices* (*MLMD*). This document is intended to be used in conjunction with ISO 14971 and ISO/TR 24971[2].

This document does not apply to *MLMD* employing large language models (LLM) or generative AI.

## 2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO 14971:2019, *Medical devices — Application of risk management to medical devices*



## 3 Terms and definitions



For the purposes of this document, the terms and definitions given in ISO 14971:2019 and the following apply.

ISO and IEC maintain terminology databases for use in standardization at the following addresses:

— ISO Online browsing platform: available at https://www.iso.org/obp

— IEC Electropedia: available at https://www.electropedia.org/

**3.1**
**bias**
systematic difference in treatment of certain objects, people, or groups in comparison to others

Note 1 to entry: Treatment is any kind of action, including perception, observation, representation, prediction or decision.

[SOURCE: ISO/IEC TR 24027:2021, 3.2.2[7]]

**3.2**
**explainability**
property of a system to express important factors influencing the system's results in a way that humans can understand

Note 1 to entry: It is intended to answer the question "Why?" without actually attempting to argue that the course of action that was taken was necessarily optimal.

[SOURCE: ISO/IEC 22989:2022, 3.5.7[4], modified — "AI system" was changed to "system".]

**3.3**
**machine learning**
**ML**
function of a system that can learn from input data instead of strictly following a set of specific instructions

Note 1 to entry: *Machine learning* focuses on prediction based on known properties learned from the input data.

[SOURCE: AAMI TIR66:2017[13], 3.19]

**3.4**
**machine learning algorithm**
**ML algorithm**
algorithm to determine parameters of a *machine learning model* (3.5) from data according to given criteria

[SOURCE: ISO/IEC 22989:2022, 3.3.6[4], modified — The example was removed.]

**3.5**
**machine learning model**
**ML model**
logical representation of a system that generates predictions based on input data

Note 1 to entry: A *machine learning model* results from *training* (3.9) based on a *machine learning algorithm* (3.4).

Note 2 to entry: Adapted from ISO/IEC 22989:2022, 3.1.23 and 3.3.7[4].

**3.6**
**ML-enabled medical device**
**MLMD**
*medical device* that utilizes *machine learning* (3.3)

Note 1 to entry: *MLMD* can involve multiple *ML models* and multiple *ML algorithms* (3.4).



**3.7**
**overfitting**
creating a model that fits the *training data* (3.10) too precisely and fails to generalize on new data

Note 1 to entry: *Overfitting* can occur because the trained model has learned from non-essential features in the *training data* (i.e. features that do not generalize to useful outputs), because of excessive noise in the *training data* (e.g. excessive number of outliers) or because the model is too complex for the *training data*.

Note 2 to entry: *Overfitting* can be identified when there is a significant difference between errors measured on *training data* and on separate *test data*. The performance of overfitted models is especially impacted when there is a significant mismatch between *training data* and new data.

Note 3 to entry: See Figure 2 for a graphical comparison with *underfitting* (3.12).

[SOURCE: ISO/IEC 23053:2022, 3.1.4[5], modified — In the definition, “which” was changed to “that”; in Note 1 to entry, “because of” was added; in Note 2 to entry, “separate test and validation data” was changed to “separate test data”, and “production data” was changed to “new data”; Note 3 to entry was added.]

**3.8**
**test data**
data used to assess the performance of the *machine learning model* (3.5)

Note 1 to entry: *Test data* is disjoint from *training data* (3.10).

[SOURCE: ISO/IEC 22989:2022, 3.2.14[4], modified — The preferred term “evaluation data” was removed; in the definition, “a final model” was changed to “the *machine learning model*”; “validation data” was deleted from Note 1 to entry.]

**3.9**
**training**
*process* to determine or to improve the parameters of a *machine learning model* (3.5), based on a *machine learning algorithm* (3.4), by using *training data* (3.10)

[SOURCE: ISO/IEC 22989:2022, 3.3.15[4], modified — The admitted term "model training" was removed.]

**3.10**
**training data**
data used to train a *machine learning model* (3.5)

[SOURCE: ISO/IEC 22989:2022, 3.3.16[4]]

**3.11**
**transparency**
property of a system that appropriate information about the system is made available to relevant stakeholders

Note 1 to entry: Appropriate information for system *transparency* can include aspects such as features, performance, *residual risks*, limitations, components, *procedures*, measures, design goals, design choices and assumptions, data sources and labelling protocols.

Note 2 to entry: Inappropriate disclosure of some aspects of a system can violate security, privacy or confidentiality requirements.

[SOURCE: ISO/IEC 22989:2022, 3.5.15[4], modified — "*residual risks*" was added in Note 1 to entry.]

**3.12**
**underfitting**
creating a model that does not fit the *training data* (3.10) closely enough and produces incorrect predictions on new data



Note 1 to entry: *Underfitting* can occur when features are poorly selected, when there is insufficient *training* time or when the model is too simple to learn from large *training data* due to limited model capacity (i.e. expressive power).

Note 2 to entry: See Figure 2 for a graphical comparison with *overfitting* (3.7).

[SOURCE: ISO/IEC 23053:2022, 3.1.5[5], modified — In Note 1 to entry, "when there is" was added; Note 2 to entry was added.]

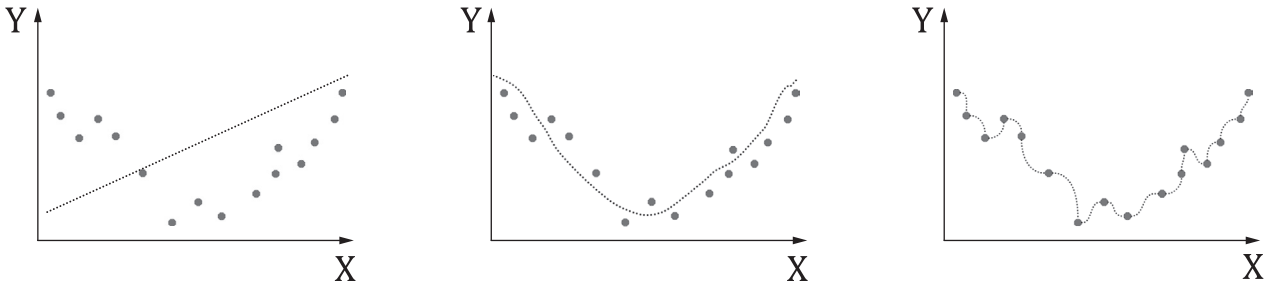


**a) *Underfitting*** **b) Good fit** **c) *Overfitting***

**Key**

| | |
|---|---|
| ● | *training data* |
| ..... | model prediction |
| X | input variable |
| Y | output variable |

**Figure 2 — Fit of the *ML model* on the *training data***

# 4 General requirements for *risk management* system

## 4.1 *Risk management process*

ISO 14971 specifies a *process* for *risk management* of *medical devices*. This *process* also applies to *ML-enabled medical devices*. The *process* covers the identification of *hazards* and *hazardous situations*, the estimation and evaluation of the associated *risks*, the need for *risk control*, the evaluation of overall *residual risk* and the production and *post-production* activities for monitoring the effectiveness of the *risk control* measures.

## 4.2 Management responsibilities

ISO 14971 requires *top management* to:

a) provide adequate resources;

b) assign competent personnel;

c) define and document a policy for establishing the criteria for *risk* acceptability; and

d) review the suitability of the *risk management process* to ensure its continuing effectiveness.

There is no specific guidance for *MLMD* in addition to what is provided in ISO/TR 24971[2].

## 4.3 Competence of personnel

ISO 14971 requires that persons performing *risk management* tasks are competent and, where appropriate, have knowledge of and experience with the particular *medical device*, its use, the technologies involved, and the *risk management* techniques employed.



For *MLMD* specifically, this knowledge should cover the relevant stages of the *MLMD life cycle*. This can include general knowledge about the *training data* and *test data* used for the *MLMD*, awareness of the data quality, understanding of the context of the data and knowledge of what the data actually mean in clinical practice.



The team performing the *risk management* tasks for the *MLMD* should include:

a) individuals with an understanding of good *machine learning* practices (see IMDRF document N88[23]), who are trained on *ML model* and *ML algorithm* development to identify *hazardous situations* that can occur for the *MLMD*;

b) individuals with knowledge of information technology (IT) to identify *risks* and *risk control* measures related to the computing platforms (when applicable) where the *MLMD* is to be installed and used, including the security *risk control* measures for *MLMD* connected to the internet or to other equipment;

c) individuals trained on simulation and algorithm testing to identify *risk control* measures for the *MLMD*, either in software or in hardware (see IEC 62304[9]);

d) individuals with knowledge of software validation to develop *verification* methods for the *risk control* measures;

e) individuals trained on usability engineering (human factors engineering) in the context of the *MLMD* use in accordance with IEC 62366-1[10], including how *transparency* and *explainability* can affect the *MLMD* use (e.g. the potential for *bias* in interpreting the output);

f) individuals with knowledge of the clinical workflow for the *MLMD*, understanding of the environment in which the *MLMD* is to be used, understanding of the interactions of the *MLMD* with medical practitioners and patients, and the capability to evaluate the relevance of the results produced by the *MLMD* in clinical practice; and

g) individuals with an understanding of good data management practices to identify *risks* related to data management, including data integrity, quality of *training data* and *test data* (e.g. data robustness, data fitness to the *ML model* and free of unintended *bias*), and data privacy.

NOTE An explanation of *bias* is given in Annex A.

### 4.4 *Risk management* plan

ISO 14971 requires the *manufacturer* to establish a *risk management* plan that includes:

a) the scope of the *risk management* activities;

b) the assignment of responsibilities and authorities;

c) requirements for the review of *risk management* activities;

d) the criteria for *risk* acceptability;

e) the method to evaluate overall *residual risk* and the criteria for acceptability of the overall *residual risk*;

f) activities for *verification* of the implementation and effectiveness of *risk control* measures; and

g) activities to collect and review production and *post-production* information.

Specifically for *MLMD*, the *manufacturer* should consider whether the *post-production* activities need to include monitoring the *MLMD* performance, updating the *MLMD* or retraining the *ML model*.

If monitoring the *MLMD* performance is considered necessary for *safety*, the *risk management* plan shall define the methods and *processes* for such monitoring and which data are to be collected and reviewed to support this activity. If updating the *MLMD* or retraining the *ML model* is considered necessary for the purpose of managing *risk*, the *risk management plan* should include the criteria for initiating an update or retraining, the frequency of applying those criteria (occasionally, after a set period or continuously) and the planned activities in the update or retraining *process* itself. These activities may further include software maintenance and provisions for returning the *MLMD* to a previous version.



NOTE Retraining the *ML model* and returning the *MLMD* to a previous version can be subject to regulatory requirements (see ISO 13485:2016, 7.3.9 and 8.2.3[1]).

EXAMPLE 1 Methods for monitoring the *MLMD* performance, including periodic testing with new *test data*, to verify that the *MLMD* continues to perform as intended.

EXAMPLE 2 *Processes* for collecting specific data from reported complaints, *processes* for collecting necessary input data (e.g. required data input fields to be filled by users before the *MLMD* can operate).

### 4.5 *Risk management file*

ISO 14971 requires the *manufacturer* to establish and maintain a *risk management file*. The *risk management file* contains the results produced by the *risk management process* as well as traceability from each identified *hazard* to the *risk analysis*, the *risk evaluation*, the implementation and *verification* of *risk control* measures and the results of evaluation of the *residual risk*.

There is no specific guidance for *MLMD* in addition to what is provided in ISO/TR 24971[2].

## 5 *Risk analysis*

### 5.1 *Risk analysis process*

ISO 14971 requires the *manufacturer* to perform *risk analysis* for the particular *medical device* and to record the results of the planned *risk analysis* activities in the *risk management file*. The techniques that support *risk analysis* (see ISO/TR 24971:2020, Annex B[2]) can also be applied to *ML-enabled medical devices*.

## 5.2 *Intended use* and *reasonably foreseeable misuse*

ISO 14971 requires the *manufacturer* to document a clear description of the *medical device*, its *intended use* and the *reasonably foreseeable misuse*. It is emphasized that *reasonably foreseeable misuse* includes *use error* (for example related to lack of *transparency*), use outside the intended patient population, use with inadequate input data, overreliance on *MLMD* output. Questions related to *MLMD* use that can assist in identifying *hazards* and characteristics related to *safety* are given in Annex C.

Differences in *MLMD* performance can relate to the selection of *training data* from specific patient groups, leading to differences in performance for other patient groups. The *manufacturer* can decide to restrict the *intended use* based on such differences in performance.

## 5.3 Identification of characteristics related to *safety*

ISO 14971 requires the *manufacturer* to identify the characteristics of the *medical device* that can affect *safety*. ISO/TR 24971:2020, Annex A[2] provides an extensive list of questions that can assist the *manufacturer* in identifying those characteristics. Additional questions that are specific to properties of *MLMD* are listed in Annex C of this document. Further guidance on the characteristics of software related to the *safety* of *medical devices* is given in IEC/TR 80002-1[12].

NOTE ASME report[15] provides considerations that can be useful for *risk management* of *MLMD*. For example, a definition of the context of use including the role and scope of the *ML model*, an assessment of the *risks* associated with the *MLMD*, the influence that the *MLMD* results can have on clinical decisions.

## 5.4 Identification of *hazards* and *hazardous situations*

ISO 14971 requires the *manufacturer* to identify and document known and foreseeable *hazards* associated with the *medical device* based on the *intended use*, *reasonably foreseeable misuse* and the characteristics related to *safety* in both normal and fault conditions as well as the *hazardous situations* that can result from each identified *hazard*.

The *hazards* and *hazardous situations* that can occur for *medical devices* in general can also occur for *MLMD*. The questions in ISO/TR 24971:2020, Annex A[2] and in Annex C of this document can help the *manufacturer* to understand the factors affecting the sequence of events and contributing to *hazardous situations*. Annex B of this document provides several examples of characteristics related to *safety*, reasonably foreseeable sequences or combinations of events, *hazardous situations* and possible *harm* in addition to those given in ISO 14971:2019, Annex C.

## 5.5 *Risk estimation*

ISO 14971 requires the *manufacturer* to perform *risk estimation*. Factors to consider in the *risk estimation* for *MLMD* include, but are not limited to, the following.

a) In cases where the probability of occurrence of *harm* cannot be estimated (which can be the case for some *ML*-related *risks*), the *risk* should be estimated based on the *severity* of possible *harm* alone.

b) The *risks* can be influenced by factors such as the reliability and usability of the hardware and other IT components. These factors can complicate the estimation of the probability of occurrence of *harm*.

c) Additional activities such as usability evaluations (including identification and analysis of *use error*) and quality assessment of the *training data* and the *test data* (e.g. for being representative of the expected patient data) can be needed to improve the *risk estimation*.

# 6 *Risk evaluation*

ISO 14971 requires the *manufacturer* to evaluate all *risks* associated with the *medical device* and to determine if each *risk* is acceptable or not. There is no specific guidance for *ML-enabled medical devices* in addition to what is provided in ISO/TR 24971[2].