TECHNICAL SPECIFICATION

ISO/TS 25238

First edition
2007-06-15

Health informatics — Classification of safety risks from health software

Informatique de santé — Classification des risques de sécurité à partir d'un logiciel de santé

Reference number
ISO/TS 25238:2007(E)

PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

COPYRIGHT PROTECTED DOCUMENT

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

Page

Foreword iv
Introduction v
1 Scope 1
2 Terms and definitions 1
3 Abbreviated terms 2
4 Principles of hazard and risk analysis 2
5 Assignment of a risk class to a health software product 4
5.1 Introduction 4
5.2 Assignment to consequence categories 4
5.3 Assignment of likelihood to consequences 5
5.4 Risk classes 7
5.5 Assignment of risk class to a health software product 7
5.6 Process of iteration 7
6 The analytical process 7
6.1 General 7
6.2 Involvement of stakeholders 8
6.3 Understanding the system and user environment 8
6.4 Consequence analysis 8
6.5 Likelihood analysis 9
6.6 Iteration 10
6.7 Reviews 10
6.8 Documentation 10
6.9 Incident library 11
7 Examples of assignment of risk classes to products 11
8 Relationship of risk classes to design and control of production of products 11
Annex A (informative) Health software products and medical devices: rationale 12
Annex B (informative) Examples of assignment of Risk Classes 15
Annex C (informative) Illustration of the nature of the relationship between risk classes and potential controls for risk management 20
Bibliography 23

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75 % of the member bodies casting a vote.

In other circumstances, particularly when there is an urgent market requirement for such documents, a technical committee may decide to publish other types of normative document:

— an ISO Publicly Available Specification (ISO/PAS) represents an agreement between technical experts in an ISO working group and is accepted for publication if it is approved by more than 50 % of the members of the parent committee casting a vote;

— an ISO Technical Specification (ISO/TS) represents an agreement between the members of a technical committee and is accepted for publication if it is approved by 2/3 of the members of the committee casting a vote.

An ISO/PAS or ISO/TS is reviewed after three years in order to decide whether it will be confirmed for a further three years, revised to become an International Standard, or withdrawn. If the ISO/PAS or ISO/TS is confirmed, it is reviewed again after a further three years, at which time it must either be transformed into an International Standard or be withdrawn.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights.

ISO/TS 25238 was prepared by Technical Committee ISO/TC 215, *Health informatics*.

Introduction

In the past, health-related software was primarily applied to relatively non-critical administrative functions where the potential for harm to the patient, as distinct from disruption to the organization, was low. Clinical systems were generally unsophisticated and often with a large administrative (rather than clinical) content and little in the way of decision support. Even clinical decision support systems tended to be “light touch”, relatively simple and understandable in their logic and used as a background adjunct to decisions, rather than a major influence on which to rely routinely. That has changed and will continue to change substantially. The nature of these changes will increase the potential for risks to patients.

There have been some high profile adverse incidents related to clinical software, e.g. in the area of screening and patient call and/or recall where software malfunctions have resulted in failure to “call” “at-risk” patients. Such incidents have not only caused anguish for the many patients concerned, but may also have led to premature deaths. The trust of the general public has been severely dented. The scope for screening for diseases is increasing significantly and it is in such applications involving large numbers of subjects that there will be heavy reliance, administratively and clinically, on software to detect normal and abnormal elements and to “call” or “process” those deemed to be at-risk. Such software needs to be safe for its purpose.

There is mounting concern around the world about the substantial number of avoidable clinical incidents having an adverse effect on patients, and of which a significant proportion result in avoidable death or serious disability (see References [1], [2], [3], [4], [5] and [6]). A number of such avoidable incidents involve poor or “wrong” diagnoses or other decisions. A contributing factor is often missing or incomplete information, or simply ignorance, e.g. of clinical options in difficult circumstances or cross-reactions of treatments.

It is increasingly claimed that information systems such as decision support, protocols, guidelines and pathways could markedly reduce such adverse effects. If only for this reason (quite apart from others, which do exist), this is leading to increasing utilization of decision support and disease management systems, which will inevitably increase in sophistication and complexity. It can also be anticipated that, due to pressures on time and medico-legal aspects, clinicians will increasingly rely on such systems with less questioning of their “output”. Indeed, as such systems become integrated with medical care, any failure to use standard support facilities may be criticized on legal grounds.

Increased decision support can be anticipated not only in clinical treatment, but also in areas just as important to patient safety, such as referral decision-making, where failure to make a “correct” referral or to make one “in time” can have serious consequences.

Economic pressures are also leading to more decision support systems. The area of generic and/or economic prescribing is the most obvious, but economy in number and the cost of clinical investigative tests is another.

Systems such as for decision support have considerable potential for reducing clinical errors and improving clinical practice. For example a large body of published evidence gives testimony to the reduction in errors and adverse incidents resulting from the deployment of electronic prescribing. However, all such systems also carry the potential for harm. Harm can of course result from unquestioning and/or non-professional use, even though manufacturers can mitigate such circumstances through, for example, instructions for use, training and on-screen presentation techniques, guidance or instruction. The potential for harm may lie equally in the system design, in such areas as:

- poor evidence base for design;
- failure in design logic to properly represent design intentions;
- failure in logic to represent good practice or evidence in the design phase;
- poor or confusing presentation of information or poor search facilities;
- failure to update in line with current knowledge.

Some of these system deficiencies are insidious and may be invisible to the user.

A substantial increase in spending on information management and technology is evident in many national health systems. Associated timetables are often tight and the goals ambitious. This increased spending can be expected to attract new manufacturers, some of which may be inexperienced in healthcare processes. Such circumstances could lead to an environment of increased risks to patient well-being.

Part of the foreseeable explosion in information and communications technology will be in telemedicine. Many of the health software products supporting such applications will be innovative and untried and the distance between clinicians and patients will make the scope for errors greater as well as less evident. Similarly, increasing use of innovative mobile IT devices and their application to new fields is likely to be associated with risks.

Whereas we are many years away from paperless, film-less hospitals, GP practices are heading that way. The inability to fall back on paper and film brings increased reliance on computers and databases. Corruption and loss of data can not only bring administrative chaos, but can also significantly affect patient care.

To sum up, the potential for harm to patients from the use of information and communications technology (ICT) in health applications will rise as the use of ICT in health applications rises, the sophistication of the applications increases and the reliance on ICT grows. There is evidence of increasing concern amongst professionals and the public as incidents of malfunctions of software, leading to adverse health consequences, raise public consciousness.

Consequently, a number of health organizations are increasingly focusing on “controls assurance” standards, including those on “governance” and “risk management”. An important feature of such controls is the management of risk in the context of harm to patients and deficiencies in the quality of care. These controls will often encompass the purchase and application of health software products.

Failures and deficiencies in health software products can, of course, have adverse impacts other than by causing harm to patients. They may, for example, create administrative inconvenience or even administrative chaos, with a range of impacts on the organization, including financial loss. Harm to a patient may also have a consequent impact on the organization, such as financial loss resulting from litigation. Whereas these adverse organizational impacts will be significant to an organization, they are not the subject of this Technical Specification unless they result in harm to a patient. For example, the failure of a hospital's central patient administration system will certainly cause substantial administrative inconvenience, but that adverse impact is not in itself within the scope of this Technical Specification unless it has the potential to cause harm to a patient (which is possible). It is the potential harm to the patient that is the subject of this Technical Specification.

The safety of medicines and of medical devices is assured in many countries through a variety of legal and administrative measures, e.g. in the European Union it is subject to several EU directives (see References [7], [8] and [9]). These measures are often backed by a range of safety related standards from a number of sources, both national and international, including the International Organization for Standardization (ISO), the European Committee for Standardization (CEN) and the International Electrotechnical Commission (IEC). Software necessary for the proper application or functioning of a medical device is often encompassed by these legislative controls. However, other software applied to health is not usually covered in this way. This Technical Specification is concerned with software applied to health excluding that which is necessary for the proper application or functioning of a medical device.

A necessary precursor for determining and implementing appropriate design and production controls, in order to minimize risks to patients from product malfunction or inadequate performance, is a clear understanding of the hazards that a product might present to patients if malfunction or an unintended event should occur and the likelihood of such a malfunction or event causing harm to the patient. Additionally, if guidance is to be given to manufacturers of health software products on design and production control (and corresponding standards produced), then it will need to be recognized that the controls necessary for products presenting low risks will not be the same as for those presenting high risks. Controls need to match the level of risk that a product might present to a patient. For these purposes, many standards, legislation and specifications dealing with control of risks in design and production group together products in a limited number of classes or types according to the risk they might present.

This Technical Specification presents a process for such a grouping of health software products. It proposes five risk classes and will facilitate broad screening of generic product types and of individual products to allow different levels of, or rigour in, the application of design and production controls that are matched to risk. Thus, the classification proposed may be a precursor for standards on design and production control, where the latter might require a far more detailed, in-depth and rigorous risk analysis for a particular product than that required for the broad classification process in this Technical Specification. Examples of the application of the process for assigning a risk class are given for a number of different types of health software products.

The term “health software products” refers to any health software product, whether or not it is placed on the market and whether it is for sale or free of charge. This Technical Specification therefore covers commercial products as well as, for example, open-source health software and software created for, and used in, only one health organization, such as a hospital. There is a broad range of health software products, ranging from simple research databases to call and recall systems, clinical decision support, electronic health record systems, ambulance dispatch systems, hospital clinical laboratory systems and GP systems. Annex B provides four examples of the application of this Technical Specification to different health software products. However, any software that is necessary for the proper application or functioning of a medical device is outside the scope of this Technical Specification.

Health informatics — Classification of safety risks from health software

1 Scope

This Technical Specification is concerned with the safety of patients and gives guidance on the analysis and categorization of hazards and risks to patients from health software products, in order to allow any product to be assigned to one of five risk classes. It applies to hazards and risks which could cause harm to a patient. Other risks, such as financial or organizational risks, are outside the scope of this Technical Specification unless they have the potential to harm a patient.

This Technical Specification applies to any health software product, whether or not it is placed on the market and whether it is for sale or free of charge. Examples of the application of the classification scheme are given.

This Technical Specification does not apply to any software which is necessary for the proper application or functioning of a medical device.

NOTE This Technical Specification is intended for the assignment of health software to broad risk classes, so as to aid decisions such as what controls should be applied to ensure safety. It is not intended for the application of risk analysis and risk management to the design of health software products and the mitigation of any identified risks to acceptable levels (see Annex A).

2 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

2.1
harm
death, physical injury and/or damage to the health or well-being of a patient

[adapted from ISO/IEC Guide 51:1999]

2.2
hazard
potential source of harm

[ISO/IEC Guide 51:1999]

2.3
health software product
software proffered for use in the health sector for health-related purposes, but excluding software necessary for the proper application of a medical device

2.4
manufacturer
natural or legal person with responsibility for the design, manufacture, packaging or labelling of a health informatics product, assembly of a system or adaptation of a health informatics product before it is placed on the market and/or put into service, regardless of whether these operations are carried out by that person or on that person's behalf by a third party

[adapted from ISO 14971]

2.5
patient
any person who is subject to, or who utilizes, a health software product

NOTE In this Technical Specification, this is taken to include healthy persons, where applicable (e.g. a healthy person accessing a knowledge database to obtain health-related information).

2.6
product
entire entity proffered to a user, including instructions for use and training, where applicable

2.7
risk
combination of the likelihood of occurrence of harm and the severity of that harm

[adapted from ISO/IEC Guide 51:1999]

See Clause 4.

2.8
risk analysis
systematic use of available information to identify hazards and to estimate the risk

2.9
risk class
classification of a health software product according to the underlying risk it might present to the safety of patients

2.10
safety
freedom from unacceptable risk of harm

[ISO/IEC Guide 51:1999]

2.11
tolerable risk
risk which is accepted in a given context based on the current values of society

[IEC 61508-4:1998]

3 Abbreviated terms

ICT	Information and Communication Technologies

4 Principles of hazard and risk analysis

Manufacturers of health software products should have a clear understanding of the hazards that their product might present to a patient, if it were to malfunction or to cause an unintended event, and the degree of likelihood that the hazard might be realized if it were to occur in reasonable circumstances of use. That knowledge is necessary for the extent and nature of the control measures required, and the rigour with which they need to be applied, so as to reduce the risk to patients to a tolerable level, e.g. through measures such as inherent design features, instructions for use and induction training. What is tolerable will depend on circumstances and the current views of society and regulators.

The essential precursor to this process is to undertake a hazard and risk analysis.

There are a variety of approaches to hazard and risk analysis, all of which share a set of underlying concepts. Existing standards, guidance and publications tend to focus on particular sectors of activity (e.g. electronic safety systems, aeronautics) or subject areas (e.g. financial risks, risks to property, risks to the security of personal data). As such, they need interpretation in the context of health software products. This Technical Specification draws on a variety of sources to keep in line with accepted general principles. The Bibliography provides a list of useful sources of information on the subject. In considering the approach to take for health software products, account has been taken of how medical devices are classified and controlled in terms of safety. Annex A addresses this matter.

The following presents some of the basic concepts insofar as they are utilized in this Technical Specification. This clause is not intended to cover all aspects of hazard/risk analysis.

The risk to the safety of a patient or patients from a health software product will depend on the possible consequence(s) that might result if the product malfunctioned or resulted in an adverse event or events, and the likelihood that such consequence(s) would in fact be realized. Thus, risk has two aspects: consequence and likelihood.

NOTE 1 ISO Guide 51 defines risk as the "combination of the probability of an event and its consequence", whereas this Technical Specification defines it as the "combination of the likelihood of occurrence of harm and the severity of that harm" (2.7). The probability that a hazard will be realized might, in some domains, be represented quantitatively as a probability which may be based on historical or experimental failure analysis and incident statistics. That is very unlikely to be the case with health informatics products safety, where such statistics and evidence are not available, and therefore qualitative judgements are necessary. Whereas probability can of course be qualitatively expressed, the term "likelihood" better conveys that meaning and is therefore used in this Technical Specification.

NOTE 2 ISO Guide 73:2002 defines risk as the "combination of the probability of an event and its consequence". This has the same drawback regarding the use of the term "probability" rather than "likelihood". Moreover, this Technical Specification is focussed only on events that are likely to cause harm to patients and the severity of that harm, rather than other events. Thus, the term "event" is not used.

The consequence, i.e. harm to the patient(s), may take on different forms, varying from death to minor inconvenience, for example. Consequences may be categorized. Such categories need interpretation according to their sphere of application, in this case the application of ICT to health. This Technical Specification proposes five "consequence" categories, each with a description of its scope (see 5.2).

The likelihood that a hazard will be realized in reasonably foreseeable circumstances might, in some domains, be represented quantitatively as a probability which may be based on historical or experimental failure analysis and incident statistics. That is very unlikely to be the case with health software products safety, where such statistics and evidence are not available, and therefore qualitative judgements are necessary. This Technical Specification proposes five likelihood categories, each with a description of its scope (see 5.3).

As noted earlier, the risk to the safety of a patient or patients from a health software product depends on the possible consequence(s) that might result if the product malfunctioned or resulted in an adverse event or events, and the likelihood that such consequence(s) would in fact be realized. The level of risks can be represented in a risk matrix where likelihood and consequence are its two dimensions (see Table 1).

Table 1

		Consequence				
		worst				least
Likelihood	highest	1	2			
		3				
						4
	lowest				5	6

Each cell of the matrix thereby represents a level of risk. Thus, in the risk matrix in Table 1 (above), the 25 cells represent 25 risk outcomes which reduce in severity on moving diagonally from top left to bottom right.

Such levels of risk outcomes can be grouped into classes such as the following:

- the highest risk class would be a group of cells in the top left, such as 1, 2 and 3;
- the lowest risk class would be a group of cells in the bottom right, such as 4, 5 and 6.

The cells of the risk matrix can thereby be populated with risk classes. When grouping together cells into a class, consideration needs to be taken of the circumstances within the application sector and the meanings assigned to each category of consequence and likelihood. The aim is to reduce complexity by identifying cells which broadly represent a similar degree of risk to the patient and grouping them into a class on that premise. Thus, a minor consequence with a high likelihood might broadly equate to a worse consequence but with lesser likelihood.

This Technical Specification proposes five risk classes (see 5.4).

5 Assignment of a risk class to a health software product

5.1 Introduction

This clause proposes categories for consequences arising from hazards, and categories for the likelihood of such consequences being realized, in the context of health software products. It further proposes a number of risk classes for health software products and relates those classes to the proposed categories of consequence and likelihood through a risk matrix. Annex B demonstrates the application of these proposals to different types of health software product.

5.2 Assignment to consequence categories

Hazards (potential for harm) that a health software product might present to a patient, if it were to malfunction or be the cause of an adverse event, shall be determined and the potential consequences of such hazards shall be identified. Each such consequence shall be assigned to one of the following consequence categories:

- catastrophic;
- major;
- considerable;
- significant;
- minor.

NOTE It will not be necessary to identify and categorize all possible consequences that could arise. The analysis to identify the realistic consequences and the likelihood of their occurring only needs to be undertaken to the extent required to assign with confidence the product to a risk class by means of the iterative process described in 5.6.

The consequence categories shall be interpreted as in Table 2. The descriptions have been created to suit the context of this Technical Specification, but are consistent with those in other sectors and in other complementary disciplines and approaches (see References [15], [16] and [17]).

Where there is doubt on the margins of two categories, the consequence shall be assigned to the category of worse consequence.