



GREEN PAPER v0.1

Quality and safety of clinical decision support systems

A draft protocol for discussion

Contributor names to be included at a later date

11 February 2002

Incorporates suggestions and amendments list - 15 April 2002

TABLE OF CONTENTS

1. Context	3
1.1. Quality and safety management in software engineering.....	4
1.2. Risk and liability assessment.....	4
2. A draft quality and safety protocol	5
2.1. Methods for assuring quality of CDSSs.....	6
2.2. Safety and hazard management techniques for CDSS development	7
2.2.1. Levels of risk.....	7
2.2.2. Safety in design.....	7
2.2.3. Operational safety	8
2.2.4. Safety case.....	8
2.2.5. Safety culture.....	9
2.3. Documentation.....	9
3. Table 1: Quality and safety requirements.....	10
4. Conclusions.....	10
5. References	11
6. Suggestions and amendments.....	11

Abstract: Developers of Clinical Decision Support Systems (CDSSs) have to date tended to be more concerned with the *efficacy* of their systems (e.g. measurable improvements in healthcare outcomes) than with *safety* (e.g. potential for hazardous side-effects). A review of quality, safety, ethical and legal liability issues suggests that CDSSs developers will be expected to comply with a “duty of care” covering all aspects of the design, development and deployment life-cycle. Borrowing from experience in the transport, power and other safety-critical industries, we identify a range of quality and safety assurance methods whose adoption may be needed before CDSSs can safely become an integral part of routine patient care, and before the trust of healthcare professionals, patients and other stakeholders can be gained. No single method will be sufficient for safe development and deployment; a range of techniques will be needed and used selectively. This OpenClinical green paper seeks to initiate a discussion of this topic. The current draft is being circulated to the OpenClinical Scientific Advisory Board and other stakeholders for comment. Later drafts will include concrete proposals for an appropriate protocol and will be published on the OpenClinical web site for open discussion prior to full publication.

1. Context

“ ... we must systematically design safety into processes of care”
[IOM Report, 2001].

There is now good evidence that clinical decision support systems (CDSSs) such as patient monitoring and reminder systems, prescribing systems, treatment management and workflow systems can make a significant contribution to quality and consistency of patient care. Interest in the use of such technologies is now growing rapidly, particularly in light of the recognition that human error in the delivery of patient care is a major source of avoidable mortality and morbidity [IOM Report, 2001].

Despite the potential of CDSSs to help improve patient care we must also anticipate possible risks in their introduction. Most new medical technologies entail new hazards (e.g. unanticipated side-effects of drugs) and even with our best efforts it will not be possible to avoid entirely the possibility that at some point in the future someone will suffer, or possibly die, in circumstances where a CDSS is involved. Software developers clearly have a responsibility to ensure that avoidable hazards are anticipated and prevented, and that unavoidable ones are properly managed should they occur.

In the context of CDSSs, there is also a further longstanding and unanswered issue concerning legal liability: if a decision support system gives bad advice, who will be held responsible? The software designers? The providers of the medical knowledge used by the system? Or the end-users - the healthcare professionals who are responsible for the final clinical decision? No-one seems to know: so far as we can establish, there is no case law to establish the relevant precedents, either in the USA, Europe or elsewhere. CDSS developers must anticipate possible legal liabilities that might result from the use of their technologies.

In this paper, we review current practices in software engineering with a view to discussing options for establishing quality methodologies that are appropriate for CDSS technology development. Further, we consider circumstances in which liability issues might come up, and propose an initial set of methods and procedures to help deal with the legal exposure that might arise should patients come to harm in situations where CDSS technology is used.

1.1. Quality and safety management in software engineering

CDSS developers have much to learn from current quality practices in software engineering, particularly in software safety engineering. Nowadays, software is increasingly developed within a “development lifecycle”, which covers the design, implementation and ongoing maintenance of software, particularly software that is intended for use in safety-critical applications [Leveson, 1995]. Quality and safety methodologies are supported by internationally accepted standards, such as the ISO 9000 quality standard [ISO 9000]. Furthermore, the software industry appears to be adopting the recently published IEC 61508 standard [IEC 61508] as a basis for establishing best practice in the design and development of safety-critical software. However, neither the International Standards Organisation nor the International Electrotechnical Commission have the authority or resources to enforce their standards (e.g. by any audit or certification process) so the current position is that industry must police itself.

In addition, no current standard can unequivocally guarantee the safety of a complex technology such as medical software; the most that one can practically achieve is to commit reasonable effort to attaining acceptable quality and safety. The problem here is that the meanings of the terms “reasonable” and “acceptable” are vague, and an organisation could commit indefinite resources in return for ever-diminishing benefit. Consequently, it is generally accepted in safety-critical industries like power and aerospace engineering that developers can only be responsible for getting the risk associated with the use of a software system to a level that is “as low as reasonably practicable” (ALARP). In short, safety is seen as a trade-off between maximising safety and investing a level of resources that is proportionate to the risk involved.

1.2. Risk and liability assessment

We have carried out an informal study of circumstances in which liability issues might arise from the use of CDSS technology. The study drew the following conclusions:

- Despite the absence of case law in this area, a supplier of CDSSs would almost certainly be viewed in the courts as having a legal *duty of care* both to patients who might be adversely affected by the technology and to health professionals who may use it in good faith in their clinical practice.
- This duty of care will probably extend to any future commercial licensees of CDSS technologies. Developers must be able to provide a high degree of assurance that the quality and safety of all components of a licensed technology have been developed according to generally agreed quality and safety standards.
- Disclaimers that attempt to limit liability by, say, restricting “proper” use of CDSSs will have limited status in law and would not by themselves be sufficient to protect a developer or supplier from legal proceedings.
- Suppliers may be able to limit their exposure by taking out insurance to cover for awards in the event of a mishap. However, this would not insulate the supplier from other costs, such as a damaged reputation.

The goal of those developing CDSSs must be to maximise the quality and safety use of this new technology, thereby minimising the risk of adverse events and exposure to legal action. All CDSS developers would wish to see their work put to effective use for the benefit of patients, but since absolute safety can never be guaranteed with any technology, they should as a minimum be able to demonstrate (in the courts, to the public or to the media) that they have fully complied with commonly accepted *best practice* during all stages of development. Currently, there appear to be no generally accepted standards of practice that developers or suppliers of CDSSs can adopt.

We have identified four primary approaches to quality and safety for CDSS technologies so far reported in the literature.

1. The use of rigorous software engineering techniques to ensure the integrity and reliability of the technology platform.

2. The adoption of systematic development life-cycles for creating and maintaining the medical content of an application and its associated scientific evidence base.
3. The application of explicit safety and hazard management techniques within the applications where this is possible.
4. The provision of comprehensive documentation to provide for quality and safety reviews by end users, technology licensees etc.

The position with regard to legal liability issues is less clear. Many CDSSs currently available (e.g. through the web) seem to depend substantially on disclaimers for their legal protection. Typical examples are “In providing this expert system, [the company] does not make any warranty, or assume any legal liability or responsibility for its accuracy, completeness, or usefulness, nor does it represent that its use would not infringe upon private rights” and “The Software is provided “AS IS”, without any warranty as to quality, fitness for any purpose, completeness, accuracy or freedom from errors”.

Given existing consumer protection legislation in many countries, legal opinion seems to be that such disclaimers offer limited protection, even when augmented with the requirement that users accept the developers’ disclaimers before access is permitted. A stronger strategy to cope with liability issues is needed.

In practice, it is likely that the degree of exposure for an organisation making CDSS technologies or applications is manageable, though this will obviously depend on common norms of litigation in different countries and markets. Nevertheless, to comply with our duty of care, it would be desirable to have an explicit protocol to guide designers and implementers in the development and distribution of CDSS systems. In the remainder of this document we consider a range of options available and propose an outline strategy for deciding when to adopt those options to comply with the ALARP principle.

2. A draft quality and safety protocol

Unfortunately, the wholesale adoption of procedures for promoting quality and safety, such as methods 1-4 above, is likely to have undesirable side-effects as well as benefits. The use of rigorous software engineering methods is difficult, particularly formal design and verification, and skills are not widely available. Many approaches also entail increased costs for the supplier, thereby reducing the commercial incentive for their development. At the other end of the process, placing strong constraints on who can have access to CDSSs and under what circumstances will inevitably curtail their use and limit their potential benefit.

This proposal therefore, abandons the idea that “one size fits all”; that a single standard of quality and safety can be adopted for the development of all applications. Rather we propose a more flexible framework based on the ALARP principle. This places a clear duty of care on CDSS developers and suppliers while permitting them to establish reasonable rules for limiting the resources to commit to system development and the restrictions they should place on its use. The approach we propose is to assess the risk of patient harm associated with a specific application and to adopt increasingly stringent quality and safety procedures proportional to increasing clinical risk.

The remainder of this document is structured as follows:

In sections 2.1 - 2.3 we consider the possible options for maximising quality and safety that are currently available, based on the four areas introduced above (2.1: software engineering techniques to maximise system quality and development life-cycles for creating and maintaining medical content; 2.2: techniques to optimise safety and hazard management in system development; 2.3: documentation.)

Section 3 is reserved for the development of an outline table to define the quality and safety techniques considered to be appropriate for CDSS applications at different levels of risk. (*This*

table will be compiled following discussions.) If a CDSS developer can demonstrate that an application has a particular level of risk and that the appropriate quality and safety requirements for that level have been complied with then the duty of care for that risk level will have been satisfied.

2.1. Methods for assuring quality of CDSSs

The quality of a decision support system needs to be considered at two levels: the level of the CDSS *technology platform* (the software which is used to build a clinical application) and/or the specific clinical application (the medical *knowledge content*). The following quality methods are applicable to both:

1. Systems should be designed, implemented, tested and documented using generally recognised methods.
2. An explicit quality plan should be developed covering all phases of implementation, testing and maintenance of the system.
3. Testing should be carried out following accepted practices, with all tests and their results recorded for review.
4. An appropriate independent individual should be nominated to sign off software and associated documentation as fit for purpose before it is made available to third parties.

Ensuring that the medical knowledge base of a CDSS is of high quality raises additional problems. Medical knowledge is subject to frequent change and research often shows that past clinical practices are ineffective, or even hazardous. Furthermore, knowledge quality will often be the professional *judgement*, either of an individual or group of experts, and efficacy and safety aspects are not necessarily always based on objective scientific evidence. Even when there is evidence, it may be limited, open to different interpretations, and subject to change as scientific knowledge advances.

A computer-based representation of medical knowledge cannot, in principle, be proved to be clinically comprehensive or objectively valid; it can only attempt to formalise the current state of professional and scientific opinion. Nevertheless, current techniques make it possible to verify formally that the medical knowledge used in a CDSS satisfies certain technical requirements like consistency and completeness, and this can be achieved, at least partially, by automatic means. In addition content will need to be endorsed by professional clinicians for the foreseeable future so it should also be a requirement that the content is humanly legible, so far as possible, and can be effectively reviewed by appropriate specialists and end-users.

The developers of decision support systems should seek to achieve at least the level of quality assurance that is associated with more traditional knowledge sources (such as medical journals and reference texts) augmented with methods that are appropriate for the new types of knowledge technology. Methods for quality control of medical knowledge bases could include:

1. Rigorous analysis to find internal inconsistencies, gaps, redundancies, ambiguities etc. (e.g. based on syntax-directed verification software).
2. Use of peer review by competent individuals. Peer review may include static assessment of content (e.g. reading the knowledge base) and dynamic assessment (e.g. testing the performance of the application against example patient data).
3. All content should be available in a legible form for review by the end users of the system, both in static form (e.g. as text) and dynamic form (e.g. as explanations of any decision or recommendation which is made for a specific patient). Provision should be made for end-users to report queries and problems to the application developers as easy as possible.

2.2. Safety and hazard management techniques for CDSS development

Safety is not the same as quality. A CDSS that is designed and implemented to high quality standards, and is working exactly as intended, can still give bad clinical advice. For example advice offered by a CDSS may not take into account atypical patient circumstances (e.g. unusual combinations of conditions; local lack of resources). In a clinical application in which there are safety considerations, therefore, an explicit protocol should be adopted which provides some assurance that the design and implementation of such systems minimises avoidable hazards to patients or others and makes provision for unavoidable hazards that are known.

2.2.1. Levels of risk

Hazards and Operability Analysis (HAZOP) is a technique which supports methodical investigation of the hazards and operational problems to which a technological system can give rise, and “is particularly effective for new systems or novel technologies” (Redmill et al, 1999). It is recommended that for all CDSS applications a limited HAZOP analysis should be carried out at the start of development in order to classify the potential level of clinical risk associated with the application. The purpose of the analysis is to classify the proposed application into one of the following risk categories.

- | | |
|---------------|--|
| Risk level 1: | There are significant, avoidable hazards that could be caused by inappropriate care based on DSS advice. |
| Risk level 2: | No hazards are expected to result from DSS use or misuse but it may be possible to neglect a pre-existing clinical condition or a situation that could warrant intervention. |
| Risk level 3: | There are no dangerous conditions that might be missed, but DSS might fail to anticipate future development of preventable conditions. |
| Risk level 4: | There are no identifiable consequences in terms of increased patient mortality or morbidity. |

2.2.2. Safety in design

If the basic HAZOP analysis suggests a risk level between 1 and 3, application development should take in a separate “safety stream” [Fox and Das, 2000] including the following activities:

- A detailed HAZOP should be carried out alongside the software requirements specification phase to identify clinical situations or events that might be associated with increased patient mortality or morbidity. Each such situation represents an *obligation* on system developers to make appropriate design changes which will prevent the anticipated hazard.
- Testing should explicitly include procedures to demonstrate that all safety obligations have been discharged.
- The application may support active safety management during operation, such as hazard monitoring and amelioration.
- A “safety case” should be prepared which documents the principle hazards, design choices and associated safety arguments, which have been considered in developing the CDSS (see section 2.2.4.).

2.2.3. Operational safety

The safety and quality techniques listed above are concerned with the responsibilities that may be imposed on designers and implementers of clinical technologies and applications. For many reasons, however, rigorous compliance with every one will not guarantee excluding the possibility of adverse events occurring in clinical operation. Systems could, for example, be misused or used in inappropriate or unforeseen situations. For these reasons, further obligations should be added to minimise, or at least monitor, the occurrence of situations that are associated with patient harm or potential harm (“near misses”). So far we have identified the following possible options:

2.2.3.1. Limiting right of access

When an application could potentially be used inappropriately, an important option would be to be able to limit access to suitably qualified or trained users. There are many possible access control options, including:

- Limiting access to licensed users (defined by an explicit contract with named individuals and/or organisations).
- Limiting access to a specific level of user whose qualifications can be verified (e.g. medical practitioners whose current professional registration has been confirmed).
- Limiting access to users who explicitly accept terms and conditions of use.
- Unrestricted open access.

2.2.3.2. “Black box”. functions

As with other technologies, notably transport, it may be desirable to keep a detailed record of the operational use of some applications, including:

- An operations audit trail. The application records all significant internal operations and external transactions (data acquisition, messages etc.) in a time-stamped format and using a mechanism that is secure.
- A clinical audit trail. The application records all medically significant information which may need to be reviewed by appropriately qualified auditors to assess the appropriateness of all recommendations, decisions taken, clinical orders issued etc.

2.2.3.3. “Guardian” functions.

For certain classes of application it may be desirable and practical to include within the application:

- The capability to monitor the use of a CDSS;
- The capability to intervene if it appears that an inappropriate decision or action is to be taken, or if a clinical hazard has not been acted upon etc.

2.2.4. Safety case

In all cases where HAZOP analysis demonstrates a significant level of patient risk, developers should document all safety-related design decisions as a “safety case”. The safety case will normally include:

- A description of the method and scope of the HAZOP analysis that has been carried out.

- All safety obligations that were identified, the design changes made to discharge them and the arguments why the design changes were necessary and/or sufficient to make the application safe.
- Revision history.

A summary of the safety case should be made accessible to end-users, ideally in electronic form from within the application. The detailed safety case will form part of the application documentation that will be available to all users on request.

2.2.5. Safety culture

The measures outlined above are to be put in place by the developing organisation, but it has become an article of faith within the safety-engineering world that safety needs to be part of the thinking of every individual in the development and support team, and possibly even those only concerned with marketing. All individuals should be exposed to clinical settings if they are to understand stakeholder needs and constraints, and all individuals should understand that their personal actions and decisions can determine patient benefit, and patient harm, resulting from the use of CDSS technology.

In short a *safety culture* should be established within the development team to which all staff will be expected to be committed. This should be supported at all levels of training and management, and even reflected in conditions of employment. For example, development organisations may require:

- Inclusion of a statement of quality and safety policy in all contracts of employment which all members of the system development and support team are required to sign.
- Discussion of the policy with collaborators, development partners and prospective clinical users, with a view to identifying safety issues.
- Ongoing assessment of compliance with this safety procedure in all staff assessments and reviews.

2.3. Documentation

An important element of a quality and safety protocol is the provision of comprehensive and clear documentation for the system to guide (i) end-users in operating it correctly and understanding its behaviour, and (ii) technical staff who are charged with support and/or maintenance.

We suggest that all CDSSs should provide the following basic documentation:

- Statement of clinical purpose of the application (e.g. diagnosis of headache, prescribing for hypertension, management of depression)
- Characteristics of the population for whom the application is intended (e.g. patient group, age, gender etc).
- Inclusion and exclusion criteria to determine the patients who should be managed using the application.
- Context of use (such as out-patient department, A&E, primary care; self-care).
- Designated users and skill levels required (e.g. qualified physician, nurse, receptionist, pharmacist, patient).
- A description of the conditions under which the application may be used (e.g. during the patient encounter; before/after the patient sees the doctor; in the patient's home etc.)

- Any associated documentation, on computer or in paper documents, which were used in the preparation of the application.
- Evidential status - one of: evidence-based (with references), consensus (who), local policy (where), individual clinical opinion (who), prototype (status), together with any further comments or significant qualifying remarks.
- Version control: a unique number for the version, build and release, together with its creation date and last revision date.
- Review date: A date at which the application should be reviewed for confirmation, revision or removal from service.
- Authors: Names, affiliations and contact details of the application authors.
- Safety case

We recommend that this documentation be accessible to the *end-user* of the application at any time from within the runtime environment.

3. Table 1: Quality and safety requirements

Risk level	Platform quality	Content quality	Safety case	Documentation	Runtime services
1					
2					
3					
4					

Table 1: Quality and safety requirements (*to be developed*)

4. Conclusions

Management of quality and safety of clinical decision support systems is an important but difficult challenge requiring technical, professional and organisational commitment. A policy that is overly lax could lead to patient harm and ethico-legal problems, while one that is overly stringent will be a disincentive to developing such technologies and achieving the potential for improved patient care that they offer. This green paper has set out a variety of options for defining quality and safety procedures to help system designers and developers demonstrate that their *duty of care* has been discharged. It is not intended that all such options should be used in all applications but that the level of investment in managing quality and safety should match the potential level of clinical risk associated with technical or operational failures. The table in section 3 (*to be completed*) will offer a tentative proposal for structuring the options that might be recommended for different risk levels.

Documented compliance with an explicit quality and safety process could provide the best practical demonstration that a developer's duty of care has been taken seriously, and that any faults, accidents or other mishaps that subsequently occur are probably unavoidable given the current state of clinical and scientific knowledge and do not represent negligence by the developer.

Comments on any aspect of this proposal or the discussion would be gratefully received.

5. References

Fox J and Das S, 2000. [Safe and Sound: Artificial Intelligence in Hazardous Situations](#). AAAI and MIT Press.

[IEC 61508](#). Standard 61508 of the International Electrotechnical Commission on the functional safety of electrical/electronic/programmable electronic safety-related systems.

IOM Report, 2001. [Crossing the Quality Chasm: A New Health System for the 21st Century](#), Report by the US [Institute of Medicine of the National Academies](#), 1 March 2001

[ISO 9000](#). ISO 9000:2000 family of international quality management standards and guidelines. (See also <http://praxiom.com/>)

Leveson N, 1995. Safeware. System safety and computers. Reading, Massachusetts, Addison-Wesley, 1995

Redmill F. and Anderson T, 1999. Towards System Safety, Springer-Verlag.

6. Suggestions and amendments

Name and date	Topic	Description	Action
Alan Rector, Manchester University arector@cs.man.ac.uk 4 March 2002	Include a consideration of decision support in the organisational context	In the final analysis, the question must be: does the overall organisation perform better with the system support or without it? On the other hand, two points need to be taken into account: the overall organisational performance might be improved by a less than perfect system but might likewise be degraded by the unintended effects of a technically (near) perfect system. The risk of system harm should be weighed against the risk of harm without the system, rather analogously to the risk of side effects in drug therapy. There is a real danger that we could place ourselves in a situation in which systems which would on balance save lives in high risk situations, cannot be deployed because they are less than perfect. This is a complex issue, particularly where, as with drugs, it may occasionally happen that a patient will suffer an adverse effect as the result of the decision support system who would not otherwise have done so, even though the balance of risk might have been very much in favour of the use of the decision support system.	Cover in the hazard assessment section.

<p>Charles Vincent Department of Psychology UCL c.vincent@ucl.ac.uk</p> <p>15 March 2002</p>	<p>The regulatory environment in UK healthcare.</p>	<p>This is changing. The Commission for Health Improvement (CHI) has recently been given new powers (or will have shortly) to 'inspect' healthcare organisations. Currently, the CHI reviews quality/safety arrangements and can investigate disasters. However the purpose of the new powers is to enable the CHI to look at any aspect of healthcare ... I see no reason why this should not extend, if problems ensue, at least to commenting on software and devices. I can't say they are anywhere near this at the moment but a major failure in, say, detection of cancer involving decision support systems might certainly bring them in. Wouldn't have any direct legal impact, but the fallout could be considerable in terms of reputation of companies etc.</p>	
<p>Charles Vincent, Department of Psychology UCL c.vincent@ucl.ac.uk</p> <p>15 March 2002</p>	<p>Operational safety</p>	<p>I appreciate the restrictions on right of access etc and other useful points made. However the general tenor of the safety case is, how can I put this, with the 'rational user' in mind. I think you might usefully add a section which considers how the software stands up to misuse, bearing in mind that people often work in a hurry, do not obey rules or follow protocols etc. What I have in mind is some rougher 'usability testing' such as is advocated for medical devices in which people (maybe expert, maybe not) are just let loose without instruction to see what happens. (There's a chapter in Vincent CA ed. <i>Clinical risk management. Enhancing Patient Safety. BMJ Publications, 2001. Second edition, 2001</i> by John Gosbee on this). The idea is to see how intuitively sensible it is, but, more importantly, to see what problems come up which the designers could never have anticipated in routine use. I realise that this is implicit in some of the safety paragraphs, but I think it would be useful to highlight it. The basic additional question is: 'Is this decision support system robust to misuse?' What are the dangers of inexpert/illegitimate use? Maybe there could even be guardian functions to close it down in certain circumstances, I don't know. Related to this I was also wondering about a kind of 'yellow card' system as they have for drugs, in which problems that are experienced could be fed back, perhaps by email, to the makers. I guess this could easily be built in.</p>	
<p>Tim Snape Abbotsbury Software Ltd. DORSET, DT3 4JT, UK tjm@wdi.co.uk</p> <p>15 April 2002</p>	<p>Data protection issues</p>	<p>There are now actually 2 data protection issues : i) confidential patient data (already discussed); ii) automatic processing.</p> <p>Individuals have a right to ensure that no decisions made as a result of automated decision making are solely based on the processing by automatic means. If an individual were to invoke this right then there would be an obligation on the controller (for the decision support system) to</p>	<p>Basically you need a new section - Data Protection and Human Rights compliance</p>

		<p>reconsider the decision. This sounds a lot worse than it is, the important points to remember are to inform the patient, and get the patients consent.</p> <p>I believe that the recommendations under Operational Safety will go a long way to satisfy DP law and Methods for assuring quality of CDSSs, in particular measure C2 (top of page 5) will give you the paper trail needed to inform the patient how decisions are made.</p>	