Ethical considerations for the development of Decision Support Systems

Alan J Green

Author's e-mail: *ethics@alan-j-green.me.uk*

This paper is extracted from a dissertation submitted as part of the MA in Medical Ethics at Keele University 2004.

The author thanks Dr Monique Jonas of Keele University for her advice and comments concerning both the MA dissertation and this extract.

Introduction

Decision making in the doctor-patient relationship will inevitably be a balance between the patient's right to control his or her own life ("autonomy") and the doctor's duty to choose and prescribe the best course of treatment for the patient ("paternalism")[1]. How much information is given to the patient and what options he is offered will always be, at least partially, decided by someone else, generally acting to secure the patient's best interest[2].

There is much emphasis in current Western healthcare systems on respect for autonomy; this is not to say that patients should be forced to make autonomous decisions (even if that were possible), but that they should be offered the tools needed for making such decisions. For the purpose of this paper, I will be concerned with the provision of knowledge and understanding that form part of the conditions needed to fulfil this. A doctor cannot expect all her patients to have wide ranging clinical knowledge, and she must decide how much knowledge is required for a particular patient to be equipped to make a particular decision. As the complexity of the medical condition increases the difficulty of giving the necessary information will increase; the range of treatments available, and the pros and cons of each may confuse patients. To avoid this, the doctor might decide that beneficence requires her to offer only a restricted range of treatment options that she deems most appropriate in the prevailing circumstances. This decision inevitably contains an element of paternalism.

Whatever balance is chosen between autonomy and paternalism the patient trusts his doctor to decide what information he should receive, and to provide that information accurately. Where a Decision Support System (DSS) is involved in the process it will have an effect on the working of the doctor-patient relationship. Design and implementation of DSSs should take into account the effect of the system on overall patient care. In this paper I will look at the need for trust and the requirements of a duty of care in the design and maintenance of DSSs.

Open Clinical has stated "A review of quality, safety, ethical and legal liability issues suggests that [Clinical] DSS developers will be expected to comply with a "duty of care" covering all aspects of the design, development and deployment life-cycle." (Open Clinical Taskgroup, 2002).

For the thought experiments below, I will consider the ethical issues raised by two different theoretical types of DSS:

> A static DSS, distributed on CD ROM. This can be seen as analogous to a textbook. It contains a fixed body of knowledge and the editors and authors will be established experts in the field. The publication medium has no bearing here (a "conventional" text book could as easily be published on CD ROM as paper) so that the only relevant difference is the use of a "knowledge engine" or differential diagnosis algorithm. It is reasonable to assume that a GP will turn to such a system in situations where she currently turns to textbooks or colleagues.
>
> A dynamic DSS, which is regularly updated. This is more analogous to consulting current findings in research journals. It could be a subscription internet site where current research and clinical trial data are posted.

I will suggest ways of addressing the issues raised by these systems.

---

[1] Both Autonomy and Paternalism are more complex concepts than is implied by the usage here; however this will suffice for this paper.

[2] I will not be concerned with the restriction of choice by health care economics and rationing.

Trust and trustworthiness

The use of a DSS in the process of medical diagnosis and treatment requires that the DDS be brought into the system of trust that is integral in the medical process. It is important to understand the nature of trust in order to understand how this can happen.

It will be important for the following discussion to state that trust is placed in specific other agents and only in them. I may say that I trust the braking system on my car, I believe this to be "shorthand" for the trust placed in a combination of individuals; the designer, the manufacturer (and his employees) and the mechanic who services the brakes are all being trusted when a child runs into the road in front of me. It is nonsensical to say I place trust in the machine itself. The machine behaves according to physical laws, if it fails it is not because the laws of mechanics have let me down, it is because one of the individuals above did not do a good job or because I am using the machine outside of the designer's specifications for its operating environment. When I come to consider the trust placed in decision support systems, it is the individuals who design and maintain them who are being trusted.

Being trustworthy and being trusted are not the same; the most obvious disparity being when trust is placed in untrustworthy individuals (as in the case of Dr Harold Shipman). There are also many examples of trustworthy individuals or groups failing to gain the trust of the populace; for example the National Radiological Protection Board has done much highly detailed investigation into the supposed dangers of living near mobile telephone transmission masts, no health implications have been detected, yet there is huge public concern over the placement of these masts (see for example Derbyshire, D., 2000).

It is not possible to enforce trust, we can only ensure and publicise trustworthiness in the hope that recognition of such will engender trust.

The duty of care

The special relationship between doctor (or other health-care professional) and patient is based on trust: the doctor is declaring herself to be competent and knowledgeable and, by accepting an individual as a patient, is accepting a duty of care towards that patient. This acceptance of a duty of care towards an individual or a group may be seen as a declaration of trustworthiness. The doctor-patient relationship is a complex one but one of its mainstays is the duty of care accepted by the doctor towards the patient and the consequent expectation of trustworthiness.

Chains of Trust

When I trust my doctor I am doing several things, I am accepting that she is a reasonable person who is acting in my best interests, I am believing that her knowledge is accurate and up to date, and I am believing she has the technical competence to operate any clinical instrument she uses (be it a computer or an otoscope). In doing this I am acting reasonably because, when she accepted me onto her list we entered into the type of relationship wherein she has accepted a duty of care towards me. She can only fulfil this duty if, within the scope of our relationship, she acts in my best interest and if she strives to keep her knowledge and clinical skills up date and critically assesses any new knowledge presented to her.

The knowledge a doctor uses to reach a diagnosis is of two types:

> Knowledge of the patient's symptoms and of clinical signs; these are acquired from interviewing the patient (and sometimes those close to the patient) and from diagnostic investigations. Such investigations will often require the use of equipment (from a simple stethoscope to complex, computer-controlled blood-chemistry assay machines); such equipment is expected to perform repeatedly and reliably and, if it fails, to fail to a safe condition (e.g. report itself as having failed or produce no result rather than report an inaccurate, but believable result).

> Knowledge of medicine: the doctor is expected to know which conditions can cause which signs and symptoms, and which treatments are available and useful for those conditions. She should understand the primary effect and the side effects of the possible treatments and how they interact with other conditions the patient may have and treatments that he may be undergoing. She must also know which conditions should be referred to more specialised units.

Trustworthy equipment

From the earliest days of modern medicine doctors have relied on the skill of equipment manufacturers and have trusted them to make equipment that behaves predictably. As the equipment used in medicine has become more complex, trust in it (and thus, I believe, trust in the designers and manufacturers), by members of the medical profession, is engendered by clinical trials and audit procedures, and by the publication of reports in the medical press of dangerous failures or unexpected behaviour. Patients, in general, in the UK NHS, place their trust in their doctors and rely on them to use or supply trustworthy equipment.

Trustworthy knowledge

Doctors acquire their medical knowledge from a number of sources, each of which they believe to be trustworthy. The knowledge sources, and the balance between them, will depend on the doctor's position and seniority; medical training, subsequent medical experience; peer reviewed scientific journals, literature from drug companies and advice from senior colleagues all feature as sources. More junior doctors will refer to more senior who will, in general, rely more on experience and literature.

If a GP, consulted by a patient, thinks she has insufficient knowledge to advise the patient on the best treatment she might consult a senior colleague within the same practice or a specialist at a hospital. In order to do this she must adjudge these sources of knowledge to be trustworthy. The patient's trust in her encompasses trust in her judgement concerning these senior colleagues and a chain of trust is established. If the hospital specialist decides to consult another colleague a further link is placed in this chain. Two further possible links in this chain are important here, the use of established knowledge contained in medical textbooks and the use of new knowledge gathered from medical journals.

The "peer review process" is designed to establish (among other things) trust in the published results. Trust is placed in the editor of a journal to identify trustworthy peer reviewers, and in those reviewers to do their job professionally. Much of this chain of trust is obscure to the patient; he must put his trust in the established procedures of the medical and scientific professions.

The chain of trust brings with it a chain of duty of care; this is not the weak, general duty which some philosophers think we owe to all moral agents, but a (possibly implicit) contractual duty taken on when the provider (of equipment or knowledge) accepts that he is in a chain leading to patient care. Although the end of the chain is the patient, it is the doctor who plays a pivotal role in deciding which of the possible chains should be considered. Here again there will be a balance between the patient's autonomy ("look at all these places with information and advice about your condition") and paternalism ("I want you to see Professor Smith, he's the country's leading expert in your condition"); the doctor must decide for each individual patient how much information the patient wants, needs and can comprehend.

Extension to Decision Support Systems

Decision Support Systems (DSS) bridge the categories I discussed above; they are items of equipment that are used to support the doctor's knowledge of medicine rather than her knowledge of the patient's signs and symptoms (although they might indicate further signs to investigate).

DSS in the chain of trustworthiness

The mechanisms of trustworthiness seen in the above examples form a base from which to investigate how chains of trust should work with the introduction of DSSs

Ethical duty of care

In taking a patient onto her list a doctor takes on an "agent-relative" duty of care for that patient, this is an interaction between two people known to each other. Chains of trustworthiness develop as medicine adopts technology which is relied on by doctors but the technical details of which are outside of medical training and expertise. The duty of care of a manufacturer of equipment is more abstract, but he is entering into a specific relationship with (unknown) other agents and is undertaking a form of agent-relative duty.

To see what this duty may be, consider a product designed for a specific market: a mobile telephone manufacturer may intend to serve a fashion-conscious teenage market; the phones he produces will serve the market by being a fashionable size and shape, connecting to the picture-messaging networks, and being cheap; it is acceptable (to his market) that they wear out in a couple of years as the owners will wish to "upgrade" by then. If a paramedical service adopts these cheap phones to control the distribution of its staff, I do not believe it can blame the phone manufacturer if patients are harmed due to failure of communication. The paramedical service has a duty to ensure its communication network is trustworthy. On the other hand, if they go to the phone manufacturer and discuss their needs and the requirement to trust their phones, a different (possibly more expensive) phone may be recommended; by making a recommendation in the knowledge of the "safety critical requirements" the manufacturer is making a statement of trustworthiness, and has accepted a duty to supply phones of the required quality. The manufacturer is entering an implicit contract with those who call on the paramedical service in the future. With this contract comes a duty of care. The specific agent-relative duties, then, depend on the relationship entered into; where equipment is expected to be used in safety critical applications, the corresponding duties will be more stringent.

It is, thus, this position of knowledge of use in patient safety critical situations that, I believe, brings a duty of care. The contract (formal or informal) to supply or work in such a situation is indivisible from the duty of care towards the patients at the end of the chain. If in the face of that knowledge a person declines the duty of care s/he must decline the whole contract.

Legal duty of care

At the time of writing (2004) Open Clinical were not aware of any case law concerning clinical DSSs (Open Clinical Taskgroup, 2002 §1.2). It is however reasonable to expect the courts to hold that foreseeable harm is tied to knowledge of intended use.

Precedent can be taken from similar areas where liability has been linked to failure in design, bearing in mind expected use. In a case concerned with an a new building developing cracks in its structure, the design engineers Baynham Meikle & Partners were employed by a building contractor, Greaves & Co. (Contractors) Ltd, to design a warehouse for the storage of lubricating oils in barrels which would be moved around using fork-lift trucks. When the building was finished and put to its intended use, it developed cracks due to the floor not being strong enough to withstand the vibration from the fork-lift trucks. Greaves & Co sued the design engineers for breach of contract and a declaration of liability; finding for the plaintiff Kilner Brown J held[3]

[3] *Greaves & Co (Contractors) Ltd v Baynham Meikle & Partners* [1974] 3 All E.R. 666, (1974) 118 S.J. 595

> The design was inadequate for the purpose. … In the special circumstances of this case, by his knowledge of the requirement …, it can be said that there was a higher duty imposed upon him than the law in general imposes … [on an]other professional man. In my judgment, the duty to a client or patient or employer is not necessarily a matter of general principle stated in objective terms. … It seems to me that it is a different situation where an engineer fails to design properly what he is specifically engaged to design. … If it be necessary to allot a category for the breach of duty which I find to have occurred I would say that, not only was there here a breach of duty, as I have described, but there was a breach of an implied term that the design should be fit for [its intended use]. There was an implied term or warranty here and it was broken.

The judgement was upheld on appeal[4]. In this case the law imposes a "higher duty" because the designers knew the intended use of the building and the design must be fit for its intended use. It is reasonable to expect the law to impose an equivalent "higher duty" on the designers, manufacturers and maintainers of clinical DSS.

A couple of thought experiments will clarify the arguments:

A Static Decision Support System

Without DSS support

Consider two patients (Alfred and Bernard) who refer themselves to their GP (Deirdre) with a common set of symptoms, consistent with a viral infection (V); the GP looks for a few signs to confirm her initial diagnosis and prescribes anti-inflammatory drugs, rest and increased fluid intake. A and B both take the medicine as prescribed. In ten days A is fit and well, and returns to work; B is still unwell and returns to D who refers him to the local hospital for further tests. It transpires that B has a rare disorder (Z) that is normally quiescent but is provoked by V; for the sake of argument suppose that 1 in 20,000 patients with V have Z. By the time B is diagnosed as having Z he is beyond cure, he dies a month later. Had he received prompt treatment for Z when first seen by D he would, in all probability, have been cured, D however did not consider Z when making her diagnosis and prescribing treatment.

D may feel that she failed B, but the rarity of Z means that no blame can be attached to her; other GPs would have acted as she did; we can see that she acted in accordance with normal practice. We must simply say B is unlucky to have had Z.

With a static DSS

Now consider D's practice has installed a new static DSS which knows about Z. D is now alerted to the possibility of Z for both A and B. The DSS will indicate how a differential diagnosis can be made, consider the following possibilities:

T_0 – there is no known test for Z;

T_1 – a simple test can be carried out there and then;

T_2 – most cases of V show reduced symptoms within 5 days, whereas most with Z show no change;

T_3 – a blood test must be sent for analysis;

T_4 – the patient must attend hospital for a scan; and

T_5 – a liver biopsy[5] is required.

In the case T_0 the availability of the DSS does not change D's prescribed treatment for A or B, however D is now aware of the possibility of Z. Is D justified in telling A and B of the possibility of Z? It is difficult to see how A or B have their autonomy compromised by lack of this knowledge. There seems to be no choice of action for A or B for which this knowledge is relevant. We may say D is acting paternalistically by keeping this knowledge from them; D

---

[4] *Greaves & Co (Contractors) Ltd v Baynham Meikle & Partners* [1975] 3 All E.R. 99, (1975) 119 S.J. 372

[5] Percutanious liver biopsy carries a 1 in 10,000 chance of mortality (Maxwell, J. H., 2004)

cannot fully know A's or B's goals in life. However the worry caused to 20,000 patients, one of whom (who may or may not be B) might decide to spend the next 10 days engaged in fulfilling one of his life goals (be it reading War and Peace or learning to knit) in case he doesn't survive, seems unjustified.

Case T_1 is straightforward, a simple test, at negligible cost, can save B's life, A can be told that D has checked for a rare complication which is not present in his case.

In case T_2 should D tell A and B of the possibility of Z? If she does so she will cause some concern to both but will provoke B to return if he has not responded to treatment. In order to provoke this she will have to be prepared to cause concern to 20,000 patients without Z. Should she not tell A or B about Z but suggest they return if they have seen no improvement in 5 days? This approach increases the level of paternalism in D's practice (she is protecting them from knowledge of Z) and, inevitably, more that 1 in 20,000 patients suffering from V will not feel sufficiently better after 5 days and will return to the surgery. Can D justify the cost (in terms of finances and workload) of more patients in the waiting room? If there is "a lot of V about" a lot of patients suffering from V will spend longer in the waiting room, or waiting for an appointment. Either of these approaches would lead to increased discomfort for many patients for the slight chance of saving one. D may decide to dismiss the possibility of Z as too unlikely to bother with; this may be a paternalistic decision by D if she thinks it is A and B's best interest not to worry about Z, however it is a decision which condemns B to die. D may decide not to tell A or B of Z because she believes that practice resources would be wasted by a large number of V sufferers returning after five days; she may feel that the consequentialist calculation involved in this decision is the only way to act with justice for all her patients; again B dies as a result of her analysis. In order for A and B to act with full autonomy D must tell them of the possibility of Z; here B will be saved but many patients will be caused to worry and many resources will be used on patients without Z.

T_3 and T_4 bring further problems of resource allocation into the decision process.

If T_3 is a simple blood test, where blood can be taken in D's surgery and sent off for an automated assay, few resources will be required, D need not worry A or B by telling them what the test is for (unless they ask) and the case is little different from T_1.

In case T_4, attendance at hospital will always cause worry for many patients. If the scan is expensive, or the resource required is of limited availability, then the resources used may be thought disproportionate to the number of lives saved. A consequentialist analysis may show that a scanner (having a limited throughput capacity) would save more lives if used for other examinations. Even if resources were unlimited it would be necessary to balance the saving of B against the worry caused to, and possibly the radiation dose given by the scan[6] to 20,000 patients without Z.

There is, then, a continuum of increasing resource requirement; at one end it is easy to see that the test should be done at the other it is easy to see that it should not. Some analysis will be required to indicate the balance point where cost outweighs benefit. The balance may be shifted if D considers each patient as an individual; a very small risk of a fatal complication may be negligible for an elderly patient suffering from Alzheimer's disease, but not for a young adult patient who is the sole breadwinner for a family of six.

In the case of T_5 the chance of dying as a result of complication from the diagnostic test is greater than the chance of having the disease. There is, then, no *meaningful* method for differential diagnosis between the case of A (having V) and B (having V and Z), so T_5 is not significantly different from T_0.

[6] In this case it might be difficult to justify the use of PET-CT or CT alone under the current UK regulations (Statutory Instrument 2000 No. 1059 The Ionising Radiation (Medical Exposures) Regulations 2000)

Is D negligent if B's Z is undiagnosed in these cases? I think there is a strong case for saying D is negligent in case T_1 if the DSS informs her of the position and she does not perform a simple negligible-cost test; ethically she has failed in her duty of care, the duty was not onerous (a simple there and then test) and there were no competing duties (a negligible cost test). The strength of the case against D will decrease as the costs (physical and emotional to the patient or resources to the NHS) increase.

If we reanalyse these cases changing the rarity of Z, the cost-benefit balance, inherent in the resource allocation consideration, will change and a consequentialist calculus may indicate different actions to be right. Consider that 1 in 200 patients with V have Z, much higher-cost diagnostic tests now seem justified: there are more potential lives to save. If the proportion is 1 in 2 then failure to organise a liver biopsy seems a failure in the duty of care (it is reasonable to expect D to know about Z without a DSS in this case).

This analysis indicates the importance of the accuracy on knowledge (how rare is Z, how can Z be diagnosed). When this knowledge comes from the DSS database then a chain of trust operates bringing a chain of duty of care. If the knowledge gained from the DSS is inaccurate, and B dies because D was not informed about the simple diagnostic test for Z, or was misinformed about the probability of occurrence of Z, then failure in duty of care cannot be ascribed to D; she acted reasonably according to her knowledge. To locate the failure in the chain of duty the origin of the error must be found, and the role of the people at that link in the chain must be examined:

Company programmer

> The programmer (or the programming team leader) who is designing and implementing the DSS, with knowledge of eventual use, has a duty of care towards patients whose diagnosis or treatment will be affected by the DSS.

General software house programmer writing general database tool

> If the database tools written by a general software company are employed by a specialist software company there is no reason to think the general programmer has entered into a special relationship with patients. The programmer may have a general duty of competence but not have a "higher duty" as described above. S/he has no agent-relative duty of care towards patients.
>
> Where, then, is the duty of care associated with these database tools? Here again I think the crucial point is knowledge of eventual use. When the database tool is brought into the DSS system it must undergo a set of validation tests to show that it is fit for its purpose. The people who design and run these tests know the intended use of the tools and, thus, acquire a duty of care towards the patients.

Data entry operator

> Data entry operators will have knowledge of eventual use and are thus incorporated in the chain of trust. A useful analogy here is the duty of staff administering drugs in a hospital ward; a second competent person must check the drugs before they are given to the patient. This analogy indicates that data verification procedures must be implemented to guard against any errors in data input.

As well as identifying responsibility in the case of failure, there is a positive role for QC and QA procedures. A patient critical system can be made the subject of a continuing audit programme. This audit process will identify weaknesses in the program, which can then be addressed, and confirm strengths in the program, thereby helping to build trust in the system.

In some respects the static DSS is like an electronic textbook in that it is written by identifiable authors and contains an established body of knowledge. It is unlike a textbook in certain key ways which impact on ethical medical practice.

Diagnosis algorithm

The most immediate and obvious difference between a DSS and a textbook is that the DSS incorporates a diagnosis algorithm. Once given a set of signs and symptoms this will offer a set of possible diagnoses and suggest further tests to distinguish between them. This is supplementing the skill and knowledge of the doctor; this brings both the providers of information (presumably senior medical experts) and the programmers of the system (often with little or no medical training) into the chain of trustworthiness involved in the patients' treatment.

Immediacy and completeness

It will be possible to have a DSS available on the computer in any consulting room, which will make them much easier to access than a textbook. The process of regular update tied with continuous audit should keep the knowledge encapsulated in the DSS up to date and complete.

Whether this will make them more likely to be consulted will be a matter for empirical investigation once the systems have been established. Their adoption may depend partly on the attitude of the courts if either use of, or failure to use, available DSS is cited in actions for negligence. For instance, in the test case above, in case T_1 if D had a DSS available but did not use it, and was subsequently accused of negligence, she might offer the defence that a responsible body of GPs would not have used a DSS when consulted by B. However if reports in the medical press were showing that patients fared better if their GPs were using DSS, it would be difficult to maintain that the body of GPs not doing so counted as "reasonable". If it is verified that patients were advantaged by their GPs using DSS then the duty of care entails the use of the systems.

Sole source

A problem with any system such as this is that, once established, it can be seen as the only source required. Why should a GP spend her time reading BMJ or Lancet when the DSS acquires all the knowledge she requires and presents it to her in easy pieces?

The professional autonomy of the family doctor can be seen as threatened if she is expected to use such a system, and act upon its pronouncements, without the latitude to decide what is in each patient's individual interest. In this the medical profession would not be unlike other professions where the growth in information technology and processing has brought major changes to working practices. It is to be hoped that the introduction of DSSs is seen as an opportunity to use technology for the advantage of patients and doctors, and to strengthen the role of the GP in caring for patients and enabling their autonomy.

A Dynamic Knowledge Model Based Decision Support System

To investigate this type of system I will consider a patient (P), who is referred to an oncology specialist (O); O confirms than P has cancer (C). Standard treatment options (T) for C show a 10% chance of survival for five years. The oncology specialist consults a dynamic, web-based DSS; it indicates that a different treatment (T') is showing 60% chance of survival for five years for patients with a particular subclass of this type of cancer (C'), indicated by a specific transcriptome pattern (R') found in a biopsy of the tumour. T' gives only a 1% chance of 5 year survival for patients with C. P undergoes biopsy and is found to have R' indicating C'.

What information should the ethical clinician give P, and how should she advise him?

It seems to me that both parts of this question depend on O's assessment of the trustworthiness of the information found by the DSS. The central questions relate to who supplies the data, how the supply of data is regulated and how the data is transformed into knowledge.

The simplest form of DSS is a database that gathers the data on a specific disease into a single place for the physician's convenience. Obviously a DSS that incorporates all and only peer-reviewed results from the printed scientific press will carry the same level of trustworthiness

as the current system of publication; this DSS is similar to the regularly updated static one in terms if its impact.

A more useful system will edit the data to be presented to aid the physician, for example by eliminating some of the more preliminary data. I will take the National Cancer Institute in the USA as an example of how this might work in practice. NCI has a WWW site (National Cancer Institute, 2004c) that carries an up to date information system (called "PDQ") aimed primarily at physicians but open to the general public:

> The PDQ cancer information summaries are peer reviewed and updated monthly by six editorial boards comprised of specialists in adult treatment, pediatric treatment, supportive care, screening and prevention, genetics, and complementary and alternative medicine. The Boards review current literature from more than 70 biomedical journals, evaluate its relevance, and synthesize it into clear summaries.
>
> National Cancer Institute, 2004c

The PDQ database includes an assessment of the level of evidence of the summaries "where appropriate" based on the "strength" of the study design and on the importance of the end-points; randomised, double-blinded controlled trials being ranked as containing the strongest evidence and non-consecutive case studies the weakest; total mortality is rated as the most important end point with other measures of success rated lower.

PDQ, thus, incorporates two types of assurance of trustworthiness, studies must be published in one of the favoured biomedical journals, and an assessment of the quality of the data is made. While either of these may impact on the patient's autonomy, it is the latter which is the more controversial. PDQ rates quality of life as fourth in importance of its five possible end-points, and whilst the editorial boards of PDQ state scientific justification for this ranking, it might not square with a patient's preferences. In the above example P may be less interested in total mortality of a treatment option, than in the quality of his remaining life. For example, he might consider life not worthwhile if he has to stop undertaking long treks through uninhabited waste lands, and thus not be willing to undergo a therapy with a high risk of renal toxicity, leading to the need for renal dialysis for his remaining years. For PDQ to be the DSS in the above example, the basic scientific work (linking T' with C' via R') would have been done prior to a clinical trial of T' in C' patients, and the results published in a top rated journal, a cycle which takes several years.

More dynamic systems giving current developments and linking data from different groups have been discussed and are under development. Recent initiatives by the research community and by the research funding bodies will lead to an increase in the data available for meta-analysis[7], for example the Medical Research Council in the UK have published a draft policy on data sharing which says:

[7]Meta-analysis is defined by NCI as:

> Meta-analyses of randomized studies offer a quantitative synthesis of previously conducted studies. The strength of evidence from a meta-analysis is based on the quality of the conduct of individual studies.
>
> National Cancer Institute, 2004b

> MRC promotes the creation of a diverse range of datasets, many of which are rich sources of information, unique, and cannot be readily replicated. Data sharing allows scientists to increase the value of these datasets through new high-quality ethical research and exploitation. It also reduces unnecessary duplication of data collection. Building preservation systematically into routine data management is good research practice: it strengthens quality, enables replication and audit, and provides a sound basis for data sharing.
>
> The MRC expects investigators supported by MRC funding to make their research data available in a timely and responsible manner to the scientific community for subsequent research with as few restrictions as possible.
>
> MRC, 2004

Other funding bodies are instituting similar policies; at the time of writing there do not seem to be tight definitions of "the scientific community" or what restrictions should be put in place.

If this data is used in a DSS, where is the element of duty of care introduced? Bench scientists, undertaking basic research into cancer, may establish the link between R' and C', but might not expect their work to have a direct impact on patient care. Just by working in cancer they have some increased duty of care since they hope the work will eventually have a clinical impact, but when they started the project they had no expectation of becoming part of the chain of duty leading directly to patients. On the other hand, if a bio-informatics programmer develops a model of radiation therapy of cancer into a program which can be used to predict which patients will benefit from a specific treatment, I believe the onus must be on that person to ensure that the data being relied on is sound (that it comes from a reputable source, where possible that it is independently verified or, at least, that it has not been challenged or refuted). So, it is again with knowledge of intended use that stronger duties are introduced.

Turning to the link between data and knowledge we see the computer taking a new role in the process. Not only can current theories be incorporated, but also statistical techniques can indicate previously unseen links. This moves the system into the field of artificial intelligence (AI). For the concerns of this paper it is only necessary here to note that the AI programmer is analogous to the database programmer discussed above; duty of care will follow knowledge of intended use and validation and audit techniques will be the tools for ensuring trustworthiness.

Many aspects of this type of system have the same ethical considerations as the static system; there are, however, aspects which are not seen in static systems.

The immediacy of availability of research data puts considerable duty on the suppliers of data and on the verification procedures. The definition of "suppliers of data" needs some clarification. I indicated above that the pre-clinical scientists who established a link between R' and C' did not expect an immediate clinical impact for their work; their duty in putting the data into a research database was not the "higher duty" entailed by knowledge of use. The higher duty lies with the meta-analysis scientists who aim to link the basic research with clinical research. It is the process of taking data from the research database and supplying it to the meta-analysis system that aims at immediate clinical use and it is here that the duty of care requires strict verification of imported data.

Having looked at the systems and how their trustworthiness can be established it is reasonable to ask *Quis custiodet ipsos custiodes?*

The systems presented here will generally have a committee of the great and the good to decide what is incorporated in the DSS and what is not. The immediate question arises "Who are *they* to decide?" Taking the example of PDQ again: NCI publish a list of the editorial board for each topic covered (adult, paediatric etc) so that experts in the field can see whom

they are trusting and, if unsure of their credentials, search academic reference databases for their academic output. The Editorial Boards are overseen by Advisory Boards:

> Each PDQ Editorial Board is supported by a corresponding Editorial Advisory Board. The Editorial Advisory Boards are similarly comprised of experts in cancer and related specialties. These Boards review the PDQ cancer information summaries on a regular basis and make recommendations for changes to the corresponding Editorial Boards.
>
> National Cancer Institute, 2004a

Although a search of the NCI web site does not reveal the membership of these Advisory Boards, NCI is a government institute and this information could presumably be obtained through the US government. The question is not about who oversees the work of the Editorial Board but how that oversight is managed.

PDQ boards make an assessment of the level of evidence present in studies they summarise, that assessment is open and public and we can decide whether to accept or reject their findings on that basis. For more rapidly reacting DSS trustworthiness can be increased by an open system of risk assessment for the incorporation of data or knowledge models, and procedures introduced to minimise the possibility of poor or positively dangerous advice being given.

Summary

Decision Support Systems can have both positive and negative effects on ethical medical practice. The increase in the doctor's and the patient's knowledge have a positive effect on autonomy; in contrast where knowledge is withheld, either by not being presented by the DSS, or by the doctor not passing on all the knowledge imparted by the DSS autonomy may be reduced; where this is thought to be in the patient's best interest there will be an increase in paternalism.

The incorporation of such systems into the healthcare process brings ethical considerations that extend the features of the doctor-patient relationship to non-medical professionals. I have introduced the concept of chains of duty and shown how trustworthiness and duty of care propagate along the chain of people involved in the development and maintenance of the systems. I have indicated that knowledge of eventual use provides a good end point for these chains.

The eventual incorporation of DSSs into normal practice will depend on many factors. Education of doctors in the use (and abuse) of the systems will play a part in this process (the Royal College of General Practitioners run many continuing professional development and other courses and publish a number of information guides for GPs; these could offer a conduit for knowledge of the use of DSSs to be given to GPs). Availability of funding and the recommendations of National Institute for Clinical Excellence will play a part as will the attitude of the courts if, for example, the use of DSSs forms part of any case of negligence. However the attitude of the public will also have a strong influence on the acceptance of such systems; in order that the public see the systems as offering a benefit the systems must be seen as trustworthy; the use of open methods of control of the systems will play a crucial role in this.

References

Derbyshire, D. (2000) *Living near phone masts "is safe"* [online]. Telegraph Group Limited. Available from: http://www.telegraph.co.uk/news/main.jhtml?xml=/news/2000/06/30/nmast30.xml [downloaded: 12 September 2004]

Maxwell, J.H. (2004) *What you need to know: Liver Biopsy* [online]. HCV Advocate. Available from: http://www.hcvadvocate.org/hepatitis/About_Hepatitis_pdf/1.1_Hepatits_C/biopsy.pdf [downloaded: 17 August 2004]

MRC (2004) *Draft MRC statement on Data Sharing and Preservation policy* [online]. Medical Research Council. Available from: http://www.mrc.ac.uk/index/strategy-strategy/strategy-science_strategy/strategy-strategy_implementation/strategy-other_initiatives/strategy-data_sharing/strategy-data_sharing_policy.htm [downloaded: 17 August 2004]

National Cancer Institute (2004a) *About the PDQ® Editorial Boards* [online]. US National Institutes of Health. Available from: http://www.cancer.gov/cancerinfo/pdq/about-editorial-boards [downloaded: 17 August 2004a]

National Cancer Institute (2004b) *Levels of Evidence for Adult Cancer Treatment Studies (PDQ®)* [online]. US National Institutes of Health. Available from: http://www.cancer.gov/cancerinfo/pdq/levels-evidence-adult-treatment [downloaded: 17 August 2004b]

National Cancer Institute (2004c) *PDQ® - NCI's Comprehensive Cancer Database* [online]. US National Institutes of Health. Available from: http://www.cancer.gov/cancerinfo/pdq/cancerdatabase#summaries [downloaded: 17 August 2004c]

Open Clinical Taskgroup (2002) *Quality and safety of clinical decision support systems: A draft protocol for discussion.* [online]. Open Clinical. Available from: http://www.openclinical.org/docs/int/docs/qands011.pdf [downloaded: 21 April 2004]