Internet based repository of medical records that retains patient confidentiality

Roy Schoenberg, Charles Safran

A patient's medical record has always been a dispersed entity. Literally defined, it is the accumulation of medical information concerning the patient. Ideally, this information is bundled in a single folder with the patient's identification data on the cover. In real life, this information is scattered between several archives (computerised and paper based) in various locations, often under different identifier numbers. Much of the information in the records is obsolete, redundant, duplicated, or indecipherable to the extent that it does not benefit the patient at the point of care.¹

Ownership of the data is also a limiting issue. Many hospitals consider the records in their systems to be their property, whereas many patients argue that their medical information is their own.^{2,3} Consequently, a distinction is made between ownership of the physical record and the right to access (or duplicate) data that are stored in it. Policies on this issue differ substantially between delivery networks, states, and countries. That said, it is typically agreed that patients have the right to be informed of the general content of their medical record and that patients' care providers must be allowed access to any information that is relevant to a patient's treatment. This approach endorses locking sensitive information (such as psychiatric evaluation or various serological findings) from some care providers but promoting access to what is "needed to know" for the provision of appropriate care. It is thus reasonable to assume that, between the patient and his or her primary healthcare coordinator (such as the family doctor), most of the "critical information" is within reach. For the purpose of our argument, we will assume that patients have rights of access to their medical information and are entitled to decide which parts of their record can be exposed or electronically published.

In this article we describe a patient controlled, "granularly secured," cross sectional medical record that is accessible via the world wide web. Unlike existing information systems, the patient initiates the service. The patient's primary healthcare coordinator suggests which clinical content is worth "risking" for the benefit of making it available when needed. The patient secures his or her identity and each data element in the medical record by specifying which identifying data anyone requesting the information must supply in order to gain access.

Summary points

Using the internet to transmit medical information could allow providers access to medical information at the point of care, but it might violate patient confidentiality

Obstacles that have prevented such implementation include patient and provider identification, security requirements, content issues, format, and language

A patient controlled, "granularly secured," cross sectional medical record that is accessible via the world wide web may be simple enough to implement and practical enough to show benefit

Patient and doctor agree which clinical content is worth "risking" for the benefit of making it available when needed

The patient determines the level of security for each data element

Center for Clinical Computing, Beth Israel Deaconess Medical Center, Harvard Medical School, 21 Autumn Street, Boston, MA 02155, USA

Roy Schoenberg
fellow

Charles Safran
director

Correspondence to:
R Schoenberg
rschoenb@caregroup.harvard.edu

BMJ 2000;321:1199-203

What is relevant in the record?

The content of the medical record is extremely heterogeneous. Information is gathered over time from various sources that use different formats and standards (which also change over time), making the record difficult to follow.⁴ In addition, medical terminology and diagnostic techniques change over time, so that the record becomes an aggregation of loosely related documents.

It can be argued that most of the information relevant to a patient at the point of care is in the most recent entries of the record or, if one is produced, its abbreviated summary. Thus, it is a patient's allergy to penicillin that is critical to future care, not previous hospitalisations for wrongful administration of the drug. The latest electrocardiographic results of a patient with coronary artery disease would be useful for emergency care, but not the numerous archived results typically found in the record. This vital information is not cumulative but cross sectional rather in nature. It explicitly does not cover a patient's medical

history but reflects information that is pertinent for future (emergency) care.

The question as to what information falls into this category is controversial. Ironically, this is one place where the shortage in time allocated for doctor-patient encounters has a positive side effect. In the United States managed care is openly promoting reductions in the length of such encounters (in some cases down to 7 minutes) in order to make best use of resources. With time so short, many healthcare providers now keep an accessible summary to avoid re-reading the whole of a patient's medical record at each encounter. This summary holds brief, concise, and relevant information, exactly the type of information our system is built to deliver. We therefore believe that a patient's principal healthcare coordinator is the appropriate authority to determine which medical information is relevant, beneficial, and "worth risking exposure" at the point of care.

Where should medical records be available?

Health administration organisations (such as managed care organisations in the United States) would like to have medical information available for questions about eligibility for treatment. Delivery networks (such as hospitals) would like it to be available to the services within the system (clinical, administrative, and financial). Individual healthcare providers would like the record to be available to them as the patient enters their practice. Patients will want their records to be available wherever they present themselves to get care.⁵ This scope extends far beyond the walls of the facility where a patient is usually seen. For example, a patient with haemophilia who is involved in a car accident and is taken by ambulance to the nearest hospital would need his records to be available there.

Evidently, any system that attempts to provide the "correct" scope of access should be able to cover all of the above simultaneously.⁶ In consequence, such a system should focus on a flexible delivery mechanism rather than on specific delivery end points. As the location of the point of care cannot be predetermined, global availability is needed. The world wide web could fulfil this requirement.⁷

The medium for data transfer

Medical data have always been exchanged between care providers. Traditional methods include the telephone, fax, and post, but these are inferior to computerised communication methods in ease of use, speed of access, cost, and reliability. The development of email has made medical data exchange simple and quick,⁸ but it is still considered insecure and operates only between users who know each other's address. Depending on the parties involved, email correspondence may be slow or unreliable.

As concern about breaches in internet security occupies more of the public and legislative agenda, "smart cards" are being considered as a possible solution.⁹ The low cost of the cards, their increasing storage capacity, and the fact that patients carry their card to the point of care make the smart card an attractive option either as an identification token or as a data container. However, people's tendency to lose small

objects, the need for a standard format, and the problem of compatible hardware at the point of care are some of the reasons why smart cards are not gaining popularity. In addition, the cards must be physically present at the time information is reviewed, which makes remote consultation impossible.

Using the web might solve some of these problems (speed, scope, security, and cost) but would pose new problems.^{10 11} One would be the location of the service. For any web data service, the requestor would probably need to know a web address (URL) before he or she could initiate the transaction. Using search engines or agents would be a possible solution. Another possibility would be putting the URL on a smart card that could be used to facilitate communication but would not be essential for using the system. Although some aspects of using the web are not resolved, it is still the most promising platform for rapidly delivering data in multiple forms to care providers whose identity and location cannot be anticipated in advance.

Identifying the requestor

During a consultation, a healthcare provider would engage the information system to acquire the patient's data. Before releasing the data, the system should be able either to recognise a trusted requestor¹² or confirm the requestor's "need to know."¹³ The degree of certainty regarding the requestor's identity and the need to validate his or her motives are key security issues.¹⁴

Identification of the requestor by means of trusted hardware tokens (such as SecureID cards) is a reliable but essentially local solution. Limiting data access to trusted subnets by means of IP addresses would identify the machine, not the requestor, and even a machine's identity could be "hacked." Furthermore, in order to allow "global" access, the reference list of trusted requestors could become too large to maintain regardless of the technique used.

Instead, we advocate a methodology that focuses on requestors' need to know rather than on their identity to approve data transaction. The definition of need to know criteria must also not be fixed. Such criteria may change over time and may differ by the content of each transaction or by patient preference. We suggest a flexible architecture that allows patient and healthcare coordinator to decide how familiar a requestor should be with a patient and his or her condition before being allowed access to data. In other words, what degree of authentication is required to satisfy need to know criteria for each clinical data item?

Identifying the patient

The difficulty of patient identification will vary in proportion to the scope of the information system. In the United States the traditional use of patients' names and dates of birth is error prone, and the likelihood of multiple matches makes queries with just these data items impractical. However, in the absence of a national patient identifier most US systems still use such queries.

Some companies have even tried to market software that shows the probability of accurate identification for any given patient database. Master patient indexes are equivalent to a medical record number that

spans several medical facilities—typically part of a distributed delivery network (a chain of hospitals). Such index systems allow the consolidation of scattered medical entries that relate to the same patient. The yield of the index is nevertheless limited once the patient attends a facility outside the network, as there is no way of knowing to which directory the index belongs. Master patient indexes are useful as an internal aggregator but not as a way to identify and get access to the patient record externally.

National indexes such as the UK NHS number greatly facilitate locating patient information because they permit unique identification with a single data item. Our proposed system takes advantage of such an identifier by treating it as yet another identifying attribute of a patient. With such a unique identifier our query algorithm would identify the patient on its first attempt; when such an identifier is not available or does not exist (as in the United States) our algorithm attempts identification using any other available attributes and may take longer to complete its task. This flexibility becomes important when patients move beyond the scope of “their” index (for example, outside their hospital network in the United States or take a business trip outside Britain). The benefit from not assuming the availability of a unique identifier, though taking advantage of it when present, is most apparent when an information system needs to scale up to a wider scope. Our scalable identification algorithm attempts to match a patient using any identification data available to the requestor at the point of care. Although the algorithm requires a unique match for the transaction to continue, it is capable of establishing that match in numerous ways, unlike traditional systems.

What medical data should be available?

There are no rules that automatically determine the optimal content or “granularity” (degree of separation of associated data items) of “vital” medical data. For each patient, the principal healthcare coordinator would have to tailor an appropriate cross section of data as discussed above. Major events in a patient’s health would require updates of the data, much as a problem list in a patient’s file needs to be updated. Data would be entered into designated categories (containers) that are general enough to prevent misunderstanding (such as cardiology). Some subcategorisation would also be implemented, mainly to allow efficient, direct data access. The ability to specify different access restrictions¹⁵ for each data item (and therefore separate them) is mandatory if “global scope” is considered. In such a system cardiac patients could allow access to, and thus risk exposing, their latest electrocardiographic results but could keep the results of a CD4 cell count (and the fact that it was performed) in a deeper, more secure layer of data. For that purpose, the system would have to be able to distinguish between cardiac and serological data, as well as a more granular distinction between different serological studies.

In what form should the information be delivered?

Information should be delivered in a standard manner to allow wide use, but it should also be conveyed in an

Fig 1 Work flow in patient controlled, medical information system

understandable manner. Unfortunately, these two approaches don’t always overlap, since the definition of “understandable” and the issue of “understandable to whom” remain open.¹⁶ However, a major incentive for using standards is the need to make data reusable in computer systems (that is, to ensure that data fit into exactly the same field in both the transmitting and accepting databases). This applies to the communication layer (such as HTTP), the visual representation of the data (such as HTML), the categorisation method of data objects (HL7 or XML), and the uniformity of the data items themselves (ICD, CPT, UMLS, SnoMed, etc).

When humans are the receptors of information many of these problems of understandability disappear. The plain textual description of a patient’s penicillin allergy is sufficient to prevent penicillin administration. Moreover, “penicillin allergy” is globally better understood than “ICD code 721.08.” When coding is required, however, a pair of tags identifying the coding system and the data element can be attached. This would allow use of the system for research purposes such as identifying patient eligibility for clinical trials.

How our system would be used

Our suggested online system would be maintained and provided by the healthcare service to which each patient belongs. This could be a centralised entity like the NHS in Britain or a discrete health plan like those in the United States. The system would enable access to patients’ vital medical information when their medical records were inaccessible (within or outside the delivery network).

Service initialisation and data upload

A patient and his or her healthcare coordinator initialise the system by using a secure internet connection (such as Secure Socket Layer) to construct a list of identifiers that can later be used to uniquely identify the patient (fig 1). The list contains demographic data (such as first and last name, social security number, NHS identifier, postal code, area code, and telephone number), non-demographic data (such as passport

number, native language, etc), and physical attributes (such as eye colour, hair colour, appendectomy scar). Finally, a list of user definable fields (such as the patient's secret code, the doctor's key, the hospital medical record number, or the patient's dog's nickname) is entered. The resulting list of identifiers includes both "easy" identifiers like the patient's first name and more cryptic data items like a password. The list is checked against the database to ensure it creates a unique record. Although a unique identification could probably be established with a fraction of these identifiers, the large number of identifiers allows security flexibility (see below).

The patient's medical information is entered next. The doctor suggests the content on the basis of the patient's clinical state and potential benefit. The patient, considering the benefit of having the data available to future care providers and the risks of exposure, decides which data are recorded into the system. With the patient's consent, data are uploaded into pre-defined medical data containers. Data are presented to users (the patient's care provider) in a hierarchical manner—mitral stenosis observations are a branch of valvar disease, which in turn is a branch of the cardiology stem. The data repository, on the other hand, follows the relational convention—where all observations are similarly stored in one table while another table contains indexes that define the temporal relations between the observations

Security structure and data retrieval

The patient can define two tiers of security for access. The first tier is the patient's identity. In order to identify a patient, the person requesting the information has to provide enough of the patient's attributes to establish uniqueness. The system has a query algorithm that checks for a unique match using any

Fig 2 Organisation of data items in a sample patient record on medical information system

Fig 3 A sample client interface for the medical information system

data items that are provided. If the scope of the system is a single hospital identification will typically require one or two items, whereas if the system covers a network identification may require three or four items. The transition from one scope to the other requires no change in the system, only a larger set of identifiers. For example, an attempt to uniquely identify a "John Smith" in a 900 000 patient database required only four identifiers. Less common names required three. The algorithm is designed so that supplied data for which no reference has been entered in the patient file will not prevent a match.

Once uniqueness is established, the system can enforce patient identification constraints (data the requestor must supply to gain access, as set by the patient). The data items required by the patient may be among the items used to establish the patient's identity but may also be complementary. Using this structure, the patient can control identification and make it as easy ("no complementary requirements, any combination that uniquely identifies me is sufficient") or as difficult ("unique identification and three passwords and the hospital medical record number") as he or she desires.

The second tier of patient customisation controls access to individual items of the medical record. This tier addresses both security and data granularity requirements. Every medical data item entered into the system is linked to a series of required "authorisers." The authorisers are any combination of medical and patient identification data, and the requestor must supply these data to gain access to medical data. Different combinations of authorisers are possible for different categories of medical information, so that a dynamic matrix of authorisation requirements can be tailored for the granular content of the record. Naturally, the requestor is unaware of the process of identification and authorisation for data access and is simply returned with data items for which he or she has satisfied the security requirements (figs 2 and 3).

Justifying the extra workload

The major obstacle for implementing any information system is the extra work required, especially in the hectic healthcare setting.¹⁷ The uploading of information

into the system requires patients' healthcare providers to invest a precious resource—time. Such an investment can be justified only if it yields a tangible return for providers as well as patients.

Patient referrals to other healthcare providers or studies usually require the writing of referral notes, which contain the same information that would be available through our system. The system enhances the patient-provider relationship by designating a patient's principal care provider as the custodian and administrator of the patient's record in the system. The system identifies the provider as the health coordinator for the patient to any other care provider the patient encounters. In addition, internal use of the system for reference purposes within a non-computerised clinic is likely to be less time consuming than the retrieval of the physical record. In practices where the medical records are computerised, automated updates of the system (such as replacement of an old electrocardiogram with a new one) can reduce interference in the provider's work. Although any diversion from medical care is in fact interference, we believe that our system's functionality has a chance of paying back the provider, delivery network, and patient for the time taken to enable it.

Discussion

Our system is intended to prevent unnecessary or erroneous medical care from being delivered by making relevant information available when and where it is needed.¹⁸ Our premise is based on the notion that such information usually exists in patients' records but is not practically attainable. Lack of medical data can lead to inefficient or inappropriate practice¹⁹ or to necessary care being delayed or withheld. An intervention that addresses even a fraction of this problem will have many financial and clinical benefits.

The introduction of the electronic medical record was, in part, an attempt to solve this issue by making the record available on all hospital terminals.²⁰ By now, the plethora of "legacy systems" and "platforms" within the same delivery network has again made the medical record dispersed. The intervention we suggest uses a globally accepted standard platform (HTTP, HTML, and the internet) to reach a much more ambitious scope (one that will not need expansion). Its architecture allows deployment both within and outside the delivery network simultaneously (a solution for delivery networks that do not have a fully integrated information system). It is inexpensive to deploy, with evident potential benefit. The system complies with the security requirements for the confidentiality of electronic health data published by the US National Library of Medicine.²¹ Thus, it is both legal to implement and can be endorsed by large delivery networks.

The success of such a system depends mostly on its endorsement. If providers will not look for the information in this system, or if numerous parallel systems coexist, the location of patients' medical data will once again be ambiguous. It is nevertheless possible to use search agents (like the ones used in web portals) to obtain unique matches within parallel services. By allowing patients to control the level of security of their medical data while permitting access on a "need to know"

basis, our proposed system may be simple enough to implement and practical enough to show benefit.

RS can be contacted by email (rschoenb@caregroup.harvard.edu) or by telephone (00 1 617 290 1678)

Competing interests: None declared.

- 1 Kuilboer MM, van der Lei J, Bohnen AM, van Bommel JH. The availability of unavailable information. *Proc AMIA Annu Fall Symp* 1997;749-53. (Annual volume.)
- 2 Annas GJ. A national bill of patients' rights. *N Engl J Med* 1998;338:695-9.
- 3 Stanberry B. The legal and ethical aspects of telemedicine. 1: Confidentiality and the patient's rights of access. *J Telemed Telecare* 1997;3:4, 179-87.
- 4 Dudeck J. Aspects of implementing and harmonizing healthcare communication standards. *Int J Med Inf* 1998;48:1-3, 163-71.
- 5 Gibby GL, Schwab WK. Availability of records in an outpatient preanesthetic evaluation clinic. *J Clin Monit Comput* 1998;14:6, 385-91.
- 6 Espinosa AL. Availability of health data: requirements and solutions. *Int J Med Inf* 1998;49:1, 97-104.
- 7 Luxenberg SN, DuBois DD, Fraley CG, Hamburg RR, Huang XL, Clayton PD. Electronic forms: benefits drawbacks of a world wide web-based approach to data entry. *Proc AMIA Annu Fall Symp* 1997;804-8. (Annual volume.)
- 8 Mandl KD, Kohane IS, Brandt AM. Electronic patient-physician communication: problems and promise. *Ann Intern Med* 1998;129:6, 495-500.
- 9 Neame R. Smart cards—the key to trustworthy health information systems. *BMJ* 1997;314:573-7.
- 10 Woodward B. The computer-based patient record and confidentiality. *N Engl J Med* 1995;333:1419-22.
- 11 Masys DR, Baker DB. Patient-centered access to secure systems online (PCASSO): a secure approach to clinical data access via the world wide web. *Proc AMIA Annu Fall Symp* 1997;340-3. (Annual volume.)
- 12 Bakker A. Security in perspective; luxury or must? *Int J Med Inf* 1998;49:1, 31-7.
- 13 Epstein MA, Pasieka MS, Lord WP, Wong ST, Mankovich NJ. Security for the digital information age of medicine: issues, applications, and implementation. *J Digit Imaging* 1998;11:1, 33-44.
- 14 Rind DM, Kohane IS, Szolovits P, Safran C, Chueh HC, Barnett GO. Maintaining the confidentiality of medical records shared over the Internet and the world wide web. *Ann Intern Med* 1997;127:2, 138-41.
- 15 Toyoda K. Standardization and security for the EMR. *Int J Med Inf* 1998;48:1-3, 57-60.
- 16 Cimino JJ, Sengupta S, Clayton PD, Patel VL, Kushniruk A, Huang X. Architecture for a Web-based clinical information system that keeps the design open and the access closed. *Proc AMIA Symp* 1998;21-5. (Annual volume.)
- 17 Slack WV, Bleich HL. The CCC system in two teaching hospitals: a progress report. *Int J Med Inf* 1999;54:183-96.
- 18 Leape LL. Error in medicine. *JAMA* 1994;272:1851-7.
- 19 Leape LL, Woods DD, Hatlie MJ, Kizer KW, Schroeder SA, Lundberg GD. Promoting patient safety by preventing medical error. *JAMA* 1998;280:1444-7.
- 20 Barrows RC Jr, Clayton PD. Privacy, confidentiality, and electronic medical records. *J Am Med Inform Assoc* 1996;3:139-48.
- 21 National Library of Medicine. *Confidentiality of electronic health data: methods for protecting personally identifiable information*. Bethesda, MD: National Library of Medicine, 1996. (CBM 95-10.)

(Accepted 7 March 2000)

Endpiece

General impressions

General impressions are never to be trusted. Unfortunately when they are of long standing they become fixed rules of life, and assume a prescriptive right not to be questioned. Consequently those who are not accustomed to original inquiry entertain a hatred and a horror of statistics. They cannot endure the idea of submitting their sacred impressions to cold-blooded verification. But it is the triumph of scientific men to rise superior to such superstitions, to desire tests by which the value of beliefs may be ascertained, and to feel sufficiently masters of themselves to discard contemptuously whatever may be found untrue.

Galton F. General impressions are never to be trusted. *Ann Eugenics* 1925;1:i.

Submitted by J H Baron, honorary professorial lecturer, Mount Sinai School of Medicine, New York