Foundations for an Electronic Medical Record

A.L. Rector, W.A. Nolan, S. Kay

Medical Informatics Group
Department of Computer Science
University of Manchester
Manchester M13 9PL
Tel +44-161-275-6188/7183
FAX: +44-161-275-6204
email {arector,skay}@cs.man.ac.uk
http://www.cs.man.ac.uk

Published in Methods of Information in Medicine 30: 179-86, 1991

Foundations for an Electronic Medical Record

A.L. Rector, W.A. Nolan, S. Kay

Abstract: Given the many efforts currently under way to develop standards for medical records, it is important to step back and re-examine the fundamental principles which should underlie a model of the electronic record. This paper presents an analysis based on the experience in developing the PEN & PAD prototype clinical workstation. The fundamental contention is that the requirements for a medical record must be grounded in its use for patient care. The basic requirement is that it must be a faithful record of what clinicians have heard, seen thought, and done. The other requirements for a medical record, e.g., that it be attributable and permanent, follow naturally from this view. We use the criteria developed to re-examine Weed's Problem-Orientated Medical Record and also relate the criteria to secondary uses of the medical record for population data, communications and decision support.

***Key-Words*:** Medical Record, Information Models, Standards, Problem-Orientated Medical Record.

Introduction

Developing the PEN & PAD prototype clinical workstation (1-3) has required a re-examination of the fundamental principles and assumptions underlying the medical record. The user-centred design process used in PEN & PAD has emphasized the many different functions performed by clinical records in patient care and underlined the limited focus of many existing electronic medical record systems. Our long term goal is a comprehensive patient care information system. Our first step is a clear understanding of the medical record and its role in such a comprehensive system.

A reassessment of the foundations of the medical record is timely for two reasons. Firstly, many of the technical constraints on storage and computing power which have conditioned the design of existing electronic medical record systems are disappearing. Secondly, a number of standards bodies are now looking at the medical record. Standards for Electronic Medical Records are being mooted by, among others, the IEEE P1157 (Medix) Committee (4), the Institute of Medicine (5), and the European Workshop on Open Systems (EWOS) (6), Key Projects under the European Community initiative on Advanced Informatics in Medicine (AIM) (7) are involved in standards (8, 9) or have made medical records and models of the medical record major parts of their work (10). Work on terminology and nomenclatures such as the Read Clinical Classification (11) and the unified Medical Language System (12) are also rapidly moving towards extensions which will effect models of the medical record. Before introducing standards, it is essential to examine again the fundamental assumptions about the functions and requirements of the medical record.

This paper describes the underlying assumptions and analysis of requirements for an electronic medical record which have emerged from the PEN & PAD project. Separate papers describe the framework for the PEN & PAD model itself and the semantics of Structured Meta Knowledge representation in which it is formulated (13, 14).

Basic Assumptions

A Faithful and Structured Record of Patient Care

Our view, presented in this paper, is based on two fundamental assumptions. Firstly, we maintain that the principle purpose of the medical record is to support individual patient care. The design of many existing electronic medical record derives, implicitly or explicitly, from support for the use of aggregated data for research, audit, finance, or planning. We maintain that such designs are inappropriate for a record for clinical use, and, ultimately, inadequate. While the use of aggregated data presents important requirements to any medical record system, clinical information, as it is generated and used during patient care, is the only sound basis for a model of the medical record.

Our second assumption is that all clinical information will be held in a structured representation which can be manipulated by the system. Structure is essential if the system is to be active in organising information for the clinician as well as to support valid aggregation of data. By "clinical information" we mean that information which falls, strictly within the medical domain. For example, that the pain is aggregated by cold is clinical; that the pain comes on when walking past the freezing compartments at the local supermarket is anecdotal. Knowledge of supermarkets and freezing compartments lies outside the domain of clinical medicine.

The goal of developing analyses and models for the medical record is, therefore, to create an architecture for structured information which is both faithful to the process of clinical care and adequate for the other uses of the information collected. As the number of information systems increases, the demands to use the same information in many different ways also increases. The greater the demands to make multiple use of information, the more important it is that the underlying models reflect accurately the nature and structure of that information. Many of the difficulties experienced in attempting to generalize existing systems stem from the fact that they have pre-selected and distorted information in order to fit into particular applications, usually clinical research and epidemiology. The models emit much of the information actually used in clinical care and do not accurately reflect the real status of the data they record. For example, many systems require that a single correct value be entered for each sign or symptom. If there is a disagreement it must be resolved outside of the system, either before data entry or through correction at a later date. Such distortion of the information, so that it is more suitable for use in one application, will almost certainly make it less suitable for another.

We therefore make *faithfulness* to the clinical history and care of the patient the fundamental criterion for the record. The contention of this paper is that the medical record consists of what clinicians have said about what they have heard, seen, thought and done.

The first consequence of our view of *faithfulness* is that the information in the medical record itself is not about what was "true" of the patient but what was observed and believed by clinicians. We can make inferences about what was "true" on the basis of these observations with greater or lesser confidence depending on the circumstances, but these are only inferences.

A second consequence of this view of *faithfulness* is that the model of the medical record should be *descriptive* rather than *prescriptive*. The function of the record is to record what was observed; the description of what was actually observed cannot be constrained to fit within a predefined view of what ought to have happened.

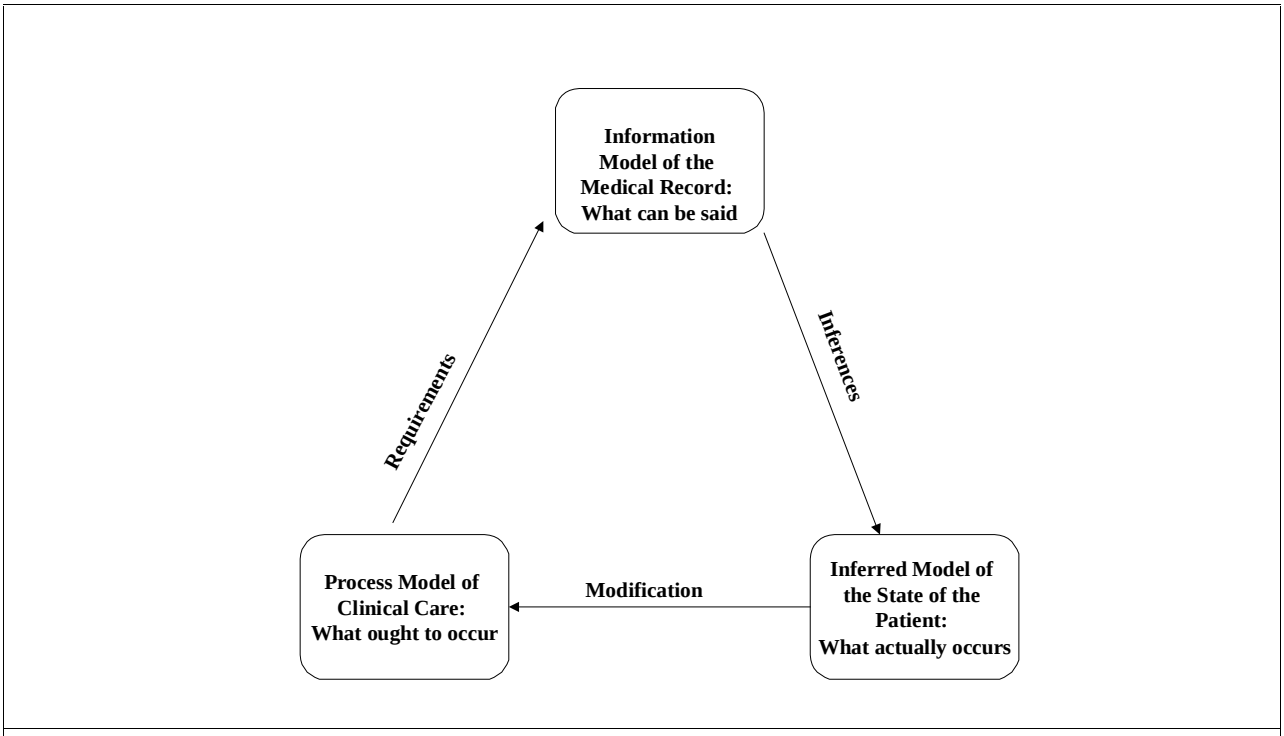

Figure 1: Relationship of the three related models: the information model of the medical record, the process model of clinical care, and the inferred model of the state of the patient.

Three models for a Comprehensive Patient Care Information System

If the statement that the medical record is about what was said, rather than what was true seems strange, this is partially because the medical record is only one component of a comprehensive patient care information system. Such a system will require at least three types of information each with its own model:

1. The information model of the medical record,
2. The process model of clinical care,
3. The inferred model of the state of the patient.

The relationship shown between these three models is shown in figure 1.

The information model of the medical record is of what *can sensibly be said* about observations of the patient and their care, and will form the main focus of this paper.

The process model of clinical care is about what *ought to be done or recorded*. The process of clinical care imposes requirements on the information model of the medical record, since the information model must be capable of expressing all of the concepts needed in the process model. For example, if the process model of clinical care requires that all therapies be related to a problem, then the information model of the medical record must be capable of representing the links between therapies and problems. However, it would be inappropriate to include the constraint that therapies linked to problems within the information model. It *makes sense* to describe a therapy independent of the reason for which it has been prescribed, even if it is considered poor medical practice. To build the constraint into the information model would make it impossible to express information from other sources or to deal with new situations in which the constraint turned out to be inappropriate.

The relationship between the information model of the medical record and the inferred model of the true state of the patient is similar. The information model must be able to express the

concepts needed to describe the true state of the patient which do not apply to observations of the patient. For example, it is perfectly reasonable for two observers to disagree as to whether or not the patient suffers from a diabetic retinopathy, but it is not possible for the patient in truth to have *and* not have the retinopathy simultaneously. In addition, the medical record may contain speculations and suggestions related to the process of decision making which are not "about" the patient in any simple way. For example, to say that a doctor wishes to rule out a diagnosis of cancer is different to saying that the patient may be suffering from cancer.

Recognition of the importance of distinctions between types of "knowledge" is one of the lessons we draw from artificial intelligence research, where making distinctions such as those between control and domain knowledge (15) and shallow and deep *knowledge* (16) have been important breakthroughs in understanding.

The remainder of this paper concerns itself only with the information model of the medical record itself.

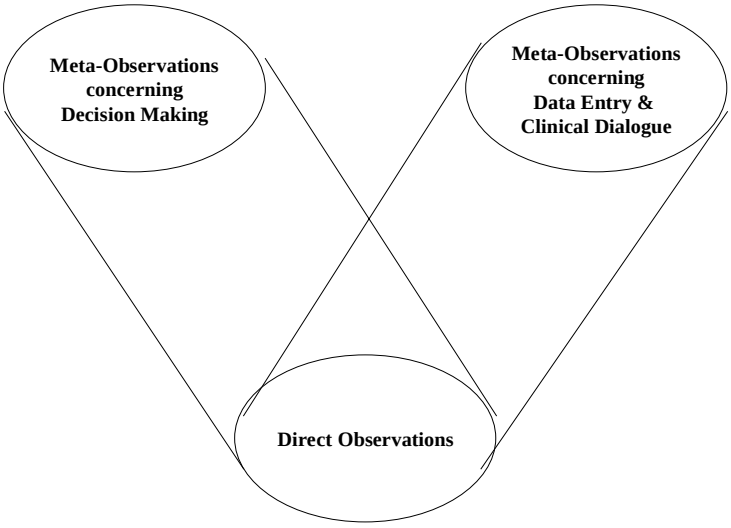

Figure 2: Two Levels of observations in the medical record

Aspects and Levels of Information

We have stated that the first criterion for the Information Model of the Medical Record is *faithfulness* to the clinical history and process of care. Within the information model of the medical record we can consider three different aspects *faithfulness*:

a) Faithfulness to the clinicians' *observations* of the patient,
b) Faithfulness to the *decision making process,*
c) Faithfulness to the *clinical dialogue.*

The first aspect deals with the *direct observations* themselves. The second and third aspects deal with the way the direct observations are related together in the decision making process and dialogue. We, therefore, divide the analysis of the medical record into two distinct levels as shown in figure 2:

Level 1. The direct observations: what was heard, seen , thought, and done concerning the patient, corresponding to *aspect a* above.

Level 2. The meta-observations about those direct observations; meta-observations concern the decision making process and the clinical dialogue, corresponding to *aspects b* and *c* above.

General requirements which apply to both of these levels will be discussed first, then each of the two levels will be considered in turn.

General Requirements Applying to both Levels

The medical record is comprised only of observations and meta-observations. Two general requirements follow immediately from this point of view.

1. *Attributability*

Every observation must be made by an agent at a particular place and time. In manual records, attribution is signified by the signature at the bottom of a note, and all entries are required, in theory, to be signed. Attributability is a major concern with regards to the security and data protection in electronic medical records, but usually regarded as a separate feature of the system. In the analysis presented here, attributability follows as a natural consequence of the understanding of the nature of the record.

2. *Permanence*

The fact that an observation was made at a particular time and place is not affected by the fact that it was later found to be incorrect. In most manual record systems the record is not supposed to be altered or removed once it is signed, albeit that there are limited mechanisms to enforce this restriction. The desirability of a deletionless electronic medical record and the necessity for adequate audit trials has been put forward (e.g. (17), (18)) in order to satisfy the requirements for legal documentation placed on clinical records. Viewing the record as a series of observations and meta-observations, makes it natural to say that once made, an observation cannot be retracted.

Requirements for Level 1: Authenticity

The first level of the medical record consists of agents' direct observations concerning the patient, as shown in figure 2. The prime requirement is that the information model of the medical record must allow an authentic account of the clinicians' understanding and not force clinicians to formulate statements in ways which they find unnatural. This requirement leads to a series of more technical requirements for the information model of the medical record, that it should:

1. *Allow conflicting statements*

Disagreements among observers and changes of opinion reflect the reality of clinical practice. To collapse disagreements into a single representation of the "truth" about a patient distorts the record of the process of care. There is no conceptual difficulty in conflicting statements, provided they are attributed to different agents or to different times.

2. *Accept uncertain and negative statements*

Observations can be negative or uncertain as well as positive. Pertinent negatives and uncertain observations can be critical to clinical decisions and can have important legal consequences. For example, it is entirely reasonable to say that the patient was asked about

"discharge from the ear" and said that they were not sure. Likewise a clinician might say that he or she has examined the tympanum and was uncertain as to whether or not it was distended. Similarly clinicians like to record their doubt about diagnoses or other inferences. Medical records contain many statements such as "hypertension".

3. Allow expression at clinicians' natural level of abstraction

The clinicians' observations must be recorded in the form in which they are made, if the record is to be faithful to the process of care. Clinicians freely mix statements about symptoms, signs and diagnoses at various levels of abstraction and detail. Our experience is that clinicians find it natural to find terms of different levels of diagnostic precision occurring within the same list of choices, for example "cough" and "acute bronchitis". In different situations, clinicians appear to perceive the presenting problem at either level of abstraction. For this reason, we draw no sharp distinction between "diagnoses", "symptoms" and "signs" and treat them all as "observations". There is strong evidence (19,20) that expert clinicians perform much of their diagnosis by pattern matching. Forcing clinicians to generate a *post hoc* rationalisation distorts the record of the clinical process and fails to communicate an authentic model of the clinicians' observations to other clinicians.

4. Allow descriptions to an arbitrary level of detail

For clinical care, the medical record should contain information at the level of detail at which it was observed. The basic concepts must be capable of being described to the required level, for example, that "diabetes has been poorly controlled over the past three weeks" or that the "fracture was severe, painful and comminuted".

5. Include the context of observations

Many items in the medical record cannot be judged without their context. For example, laboratory values require normal ranges or reference distributions if they are to be interpreted later at a different time and place from where they originated, and imaging procedures may require details of the exact techniques used before their validity can be assessed.

6. Accept multiple simultaneous measurements and instances of count nouns

Measurements, such as blood pressure, are often repeated at a single session by the same observer. For all the practical purposes such multiple measurements should be regarded as simultaneous. While a person can have only one blood pressure at any given time, it is perfectly reasonable to observe and record several independent measurements of that blood pressure. There is a related distinction between what linguists refer to as "count nouns" and "mass nouns", e.g., it is perfectly possible (if rather unfortunate) to have simultaneously two distinct "fractures of the right leg", but it is not reasonable to have simultaneously two different "hypertensions". To be faithful to the clinical history, the record should be capable of representing multiple measurements and multiple instances of count nouns.

Requirements for Level 2: Faithful Recording of Decision Making, Dialogue and Data Entry

The second level of the medical record consists of meta-observations about the observations at the first level. The characteristic feature of meta-observations at the second level is that they can be altered without changing the underlying observations. Whether the observations of "chest pain" and "swelling of the ankles" relate to one problem or two separate problems does not alter the status of the initial observations. Whether two visits for "bronchitis" should be grouped together as a single episode or separated into two episodes does not alter the status of the observations of "bronchitis" on two separate occasions. Whether a haemoglobin

measurement was made in response to a request or not, does not alter the value of measurement.

The meta-observations on the second level concern two broad areas, decision making and the clinical dialogue, which will be discussed separately.

Decision making

Problems, plans and episodes

Most meta-observations are used to group observations together and order them with respect to time, causality, or importance. The concept of a problem resides within the layer of meta-observations. The evolution of the definition of a problem, the grouping of observations pertinent to that problem, the recording of its activity over time, and the subsumption of one problem by another, are all meta-observations. The problem must be able to develop and change without altering the direct observations.

Related concepts in other record systems include *plans*, which link managements to problems or observations, and *episodes*, which group visits or periods of care together in ways which may reflect either the clinicians' understanding or administrative rules.

Justifications

The links between the evidence in the record and the conclusions which the clinician reached based on that evidence constitute a particularly important form of meta-observation. In many cases the evidence on which a decision is based is as important as the decision itself to the understanding of the course of events. While it should be possible to record events based on simple pattern matching without further justification, it should also be possible to record a detailed pattern of reasoning leading to a conclusion. Detailed arguments in support of conclusions are common in the manual records of patients with complex problems.

The ability to record justifications will take on added significance as intelligent inference systems are developed and attached to record systems. The evidence used by an inference process must be open to scrutiny by the clinician.

The Clinical Dialogue and Process of Recording

Requests and responses

The medical record also describes the complex dialogue amongst clinicians consisting of requests for tests and procedures, referrals, and opinions, and the corresponding responses. A diagnostic procedure, such as an X-ray examination, may be ordered at one time and one place by one clinician, performed at a second, interpreted at a third, and finally accepted and integrated into the medical record at a fourth. The information involved may have been entered into several different information systems at still other points in time. Each event involves separate observations, but can only be properly understood if they are linked together.

The information model of the medical record must therefore include meta-observations which record the clinical dialogue. In its most general form it must be able to record that one statement is a *response to a request* expressed in another statement. Only if this information is available is it possible to answer questions such as "Have all of the test results outstanding on Mrs Smith for more than one week been received and acted on ?"Note that, as with observations about decision making, whether or not an observation is a response to an earlier request has no effect on the validity of that observation *per se*, e.g., the patients' chest X-ray shows an enlarged heart regardless as whether it was taken as part of a routine admission procedure or in response to a specific request. However, the total process of care can only be understood and audited if the links between the request and response are available.

Data entry

It is also necessary to record who actually entered information in a system, when and where. Ultimately, everything in the medical record must be entered by some agent at a particular time and place. We assume that the record will normally by the clinician at the time of the patient encounter. The intervention of a third party in the recording process should, eventually, become an anomaly. For this reason we take the observation as the basic unit of the medical record. Observations are assumed to be entered by the observer unless stated otherwise. Where information is entered by a third party, the model must allow this to be indicated by a suitable meta-observation.

The Problem Orientated Medical Record Revisited

Weeds' Problem Orientated Medical Record (21) has been the most influential work on the structure of medical records over the past three decades. Weed was instrumental in focusing the attention of the medical profession on the medical record and in establishing the idea of "problems" and the "problem list" as central features of many medical record systems. A major goal of Weeds' work was to facilitate computerization of medical records, and many computer-based systems claim to implement all or part of Weeds' system. Nevertheless, Weeds' Problem Orientated Medical Record has not been widely accepted in the form in which it was originally conceived.

The analysis in this paper owes much to Weeds' work. However, there are three key issues on which it differs from Weeds' original formulation. We contend that misunderstandings relating to these issues are responsible for much of the difficulty in implementing systems based on Weeds' work, whether manual or computer based.

1. Weeds' Problem Orientated Medical Record is essentially a prescriptive model of how care should be delivered rather than a descriptive model of the medical record. Weed is entirely explicit about this and sometimes speaks of a "problem orientated system of medicine". However, its prescriptive nature makes the Problem Orientated Medical Record inappropriate, in itself, as an information model for a medical record, since the constraints which it imposes on what *should be done* cannot be assumed to hold for what *can be said*. We contend that, if required, the prescriptive constraints in Weeds' work should form a separate process model of clinical care. Questions of the desirability or otherwise of these constraints should be divorced from the information model needed to support them.

2. We treat problems as fundamentally different from *observations* of diagnoses, signs and symptoms. *Problems* belong to the second level of the medical record consisting of meta-observations about the underlying observations. As Weed indicated, a *problem* can be anything – a symptom, an abnormality, a treatment, or anything else which has been observed of the patient. However, this is not because "problem" is a superordinate concept *subsuming* all other medical statements; rather it is because problems are based on meta-observations *about* other medical statements.

In other words, it is not correct to say that signs, symptoms, diagnoses, treatments, social complaints, etc. are all *kinds* of problems in the sense that pneumonia is a *kind* of infection or that signs and symptoms are both *kinds* of indicants. Problems behave entirely different from signs, symptoms, and diagnoses. For example, it is reasonable to say that the problem was originally thought to be chest pain, then diagnosed as pneumonia and later found to be a myocardial infarction. The same problem is said to persist through all three definitions.

Rather than say that signs, symptoms, diagnoses, etc. are *kinds* of problems, we represent problems using special meta-relationships which can hold between an observation and a problem. One such meta-relationship is used to say that an observation *defines* that problem;

a second meta-relationship is used to say that other observations are *pertinent* to that problem. In this, we do not so much disagree with Weed as with the interpretations of Weed made in the course of implementing many record systems.

3. We insist that information be recorded at the level of abstraction at which the doctor perceived it, qualified where appropriate by a marker for uncertainty. Weed drew attention to the importance of recording information at the level of symptoms, signs and abnormalities where no more precise diagnosis was available. However, he constantly argued against the converse – recording questions such as "?hypertension". Doctors routinely use rubrics such as "?hypertension" and an important part of their observation and decision making process (18). To argue for a prescriptive process model which forbids expressions of uncertainty is legitimate, but not to provide for them in an information model that cripples the model, rendering it incapable of recording information which many doctors believe to be important.

Other Functions of the Medical Record

Although it is not the task of this paper to enumerate or analyse all the possible functions of the medical record, there are several important functions which any proposed model of the medical record must be shown to serve.

Population Data for Research, Clinical Audit and Health Service Management

Historically, the most important motivation for computerized medical records has been the need for data on populations to support research, clinical audit and health service management. Many early record systems such as ARAMIS (22) were developed primarily to support a single research project. When data on individuals are used in population analyses, the two prime requirements are comparability and flexibility. We contend that the most effective means for routinely collecting data which meet these requirement is through medical records which faithfully record the clinical care of individual patients.

Comparability can be achieved in two broad ways. The first approach is to collect data in standardised ways according to standard definitions. Essentially, this requires a process model of clinical care. The second approach is to establish comparability *post hoc*. This always involves difficulties, but to have any chance of success, establishing comparability *post hoc* demands detailed clinical information. Without detailed information, it is unlikely that it will be possible to determine whether or not a given patient meets the required case definition, let alone whether or not they possess other characteristics of interest. For example when analysing data from a large set of routine sources it would be unwise to assume the presence of a code for diabetes in a clinical record always indicated a case of diabetes in the context of the analysis. It would normally be necessary to examine each case in detail against pertinent criteria to ascertain its eligibility for inclusion in the analysis.

Furthermore, a simple record of unconnected observations often provides an inadequate basis for analysis. Meaningful medical audit, in particular, requires knowledge of the reasons for an action as well as details of the action itself (23). The proposed two level model, which records decision making and justifications as meta observations, meets this requirement. Existing systems based on simpler models can provide no such facilities.

Flexibility, in the context of the analysis of population data, means the ability to aggregate the data in many different ways along many alternative axes. Traditional classification systems are goal-orientated and pre-define the data aggregations according to that goal. When attempts are made to use the information for another purpose, the pre-defined aggregates are often found to be inappropriate or too coarse-grained. Details which have vital discriminatory significance for the new purpose have been eliminated. For example, data classified according to major aetiological factors may be an unsuitable basis for drawing conclusions on health service

provision. If, however, the detailed relationships between the diseases of individual patients and their treatments are maintained, then appropriate inferences may be possible.

Communication, Data Interchange, and Cooperative Working

An important function of the medical record is to provide means of communication amongst clinicians. The development of models of medical records is essential in order to facilitate interchange of clinical information. Communication imposes many requirements on the technical model of the record, but the key conceptual requirements are the ability to link statements together into sequences of requests and responses.

Decision Support and Access to other information

A major motivation for developing new electronic medical records is the hope that they can be used as part of a decision support system. Typically decision support systems such as the Oxford System of Medicine (24, 25) or QMR (26) require much more information than it is practical to collect at the time they are used. The natural solution is to link them to the patient record system. However, this requires that the record system represents the information in an appropriate form. Fundamentally, this requires that the model of the medical record be capable of supporting a rich description of the patient and process of care – precisely those requirements which we have outlined in our criteria of *faithfulness*.

Legal Documentation

Medical records are legal documents, and there are legal constraints on their handling in almost all countries. The fundamental requirements for attributability and permanence follow naturally from the view expounded here of the record as a collection of observations. Most of the requirements for control of access and transfer amount to demands that it be possible to make additional meta-observations about the authorisation required to gain access to statements in the medical record, and the system guarantees that these statements are respected.

Conclusion

This paper presents an approach to the medical record based on its use in clinical care. It is an attempt to step back from assumptions which originated in the limitations of previous hardware and software and in the origins of many medical records systems as support for research, epidemiology, or financial management. As a result, it sees the medical record as much more than a series of date-stamped codes. It sees it as the record of clinicians' observations and dialogue – a highly organised structure of observations which record what clinicians have heard, seen, thought and done; meta-observations which record their reasoning; and further meta-observations which record the structure of their dialogue of requests and responses and the process by which the record was recorded.

The approach which emerges from this analysis not only satisfies the requirements of clinical practice but also makes it easier to satisfy the requirements of other uses of the medical record. Attributability and permanence are seen as natural consequences of the nature of the record rather than as special features to satisfy requirements for confidentiality, security, and legal documentation. Problem orientation becomes a special case of more general techniques for recording the decision making process. Issues remain concerning the comparability of data in different medical records which need to be resolve, but these issues reflect fundamental problems in collecting data for clinical research. The approach advocated here separates the clinical discussion of the comparability of the data from the technical discussion of the structure of the information system. It accepts that the medical record consists of observations rather than facts and distinguishes the descriptive information model of what can be said from the prescriptive process model of what it is correct to say or do.

We took as our starting point the view that the medical record must be structured. Workers such as Gell (27) and Scherrer (28) have developed extensive medical record systems based on free text and natural language recording. Despite the success of these systems for information retrieval, we do not believe that free text represents a suitable basis for general clinical information systems for two reasons. Firstly, since the information is not "understood" as it is entered, the system cannot provide information management and decision support functions for direct patient care during the encounter. Secondly, comparability and information interchange for aggregating data are extremely difficult to achieve without a pre-defined structure. Eventually natural language systems will be able to convert free text and speech to a structured form in real time. When this occurs, whether natural language or structured input is preferred will be a matter of convenience, ease of use and empirical determination of the quality of the resulting data. In the short term, however, structured data entry is computationally more tractable. It is also proving to be a valuable tool for studying the underlying semantics of medical language (29).

More generally, increased detail and a more complex structure might potentially be seen as being in conflict with ease of use for the clinician and flexibility when entering data. With respect to ease of use, our experience strongly suggests that structured data entry can be made quick enough and easy enough for routine clinical use (1,30). With regards to flexibility, we would contend that the models advocated here will result in systems which are more flexible than existing systems, because these models more faithfully represent the structure of the information recorded. Many of the restrictions and much of the complexity of existing systems results from their use of models which do not accurately reflect the true structure of the data being represented, e.g., the mechanisms for data audit and corrections of erroneous entries.

The analysis has emerged in the course of the development of the PEN & PAD clinical workstation, and has resulted in a serious attempt to capture the full richness of the narrative record within a formal structured framework. The details of the models underlying PEN & PAD are reported separately. Although we have so far concentrated on primary care, the analysis is intended to be general, and preliminary investigations of hospital environments have produced no serious indications to the contrary. The success of PEN & PAD in satisfying clinicians' demands gives us considerable confidence in the soundness and feasibility of the approach.

Acknowledgements

This research was supported in part by the United Kingdom Medical Research Council, grant number SPG8800091, and the Department of Health.

References

1. Howkins TJ, Kay S, Rector AL, Goble CA, Horan B, Nowlan WA, Wilson A. An overview of the PEN & PAD project. In: O'Moore R, Bengtsson S, Bryant JR, Bryden JS, eds. MIE 90. Lecture notes in Medical Informatics No.40. Berlin: Springer-Verlag, 1990: 73-8

2. Horan B, Rector AL, Sneath EL, Goble CA, Howkins TJ, Kay S, Nowlan WA, Wilson A. Supporting a humanely impossible task: the clinical human-computer environment. In Diaper D, ed. Proceedings of Interact 90. Amsterdam: Elsevier science Publ. 1990:247-52.

3. Rector AL, Goble CA, Horan B, Howkins TJ, Kay S, Nowlan WA, Wilson A. Shedding light on patients' problems : integrating knowledge-based systems into medical practice. In: Aiello L. ed. Proceedings of the Ninth European Conference on Artificial Intelligence. ECAI 90. Pitman Publ, 1990: 531-4.

4. Kay S. SKRAM: Semantics and Knowledge Representation of the Medical Record. Subcommittee of IEEE P1157 (MEDIX) 1989.

5. National Academy of Sciences. Institute of Medicine. Medical Record Project (proposal march 21st 1988). Final report (to appear).

6. European Workshop for Open Systems (EWOS PT007 Medical Informatics). Study and Investigation of Problems Related to Standardisation and the Relevance of Open Systems Standards in Medical Informatics Final report (Second Draft v 3.12).

7. Commission of the European Communities. Directorate of Telecommunications. Information Industries and Innovation. Operation 1992: Investigation of Requirements and Options in the Field of Advanced Informatics in Medicine: Consolidation of Results (Ref X111/F/GE0392). Brussels: EEC DG X111, 1990.

8. Englebrecht R, Fitter M, Newton P, Rector AL, Robinson D, Sneath L, PRECISE: Prospects for an Extra Mural and Clinical Information System Environment (AIM Project A1043, Subproject CRS: Requirements for Consulting Room Systems). Report on User Workshop,1990.

9. SESAME. Standardization in Europe on Semantical Aspects of Medicine (AIM Project A1031. Deliverable 9. Structure for Classifications on Primary Health Care) 1990.

10. Harrison J, Kain B, Harvey F, Framework Model: Level 2 AIM Eurodiabeta Project (Deliverable 2, Eurodiabeta). London: Department of Chemical Pathology, St.Thomas's Hospital 1990.

11. Read JD. Computerising Medical Language. In: De Glanville H, Roberts J, eds. Current Perspectives in Health Computing, 1990: 203-8.

12. Evans DA, Miller RA. Unified Medical Language System (UMLS) Project. Initial phase in developing representations for mapping knowledge: INTERNEST-I/QMR, HELP and MESH (Technical Report No.CMU-LCL-87-1). Pittsburgh, PA: Laboratory for Computational Linguistics, Carnegie Mellon University, 1987.

13. Rector AL, Nowlan WA, Kay S. Unifying a medical information using an architecture based on descriptions. In: Miller RA, ed. Proceedings of the Symposium on Computer Applications in Medical Care. SCAMC 90. Washington, 1990:190-4.

14. Rector AL, Kay s. Descriptive models for medical records and data interchange. In: Barber B, Cao D, Qin D, Wagner G, eds. MEDINFO 89. Amsterdam: North-Holland Publ Comp, 1989: 230-4.

15. Clancey W. The epistemology of a rule based expert system: a framework for explanation. Artif Intell 1983; 20: 215-51.

16. Steels L. Second Generation Expert Systems. In Bramer MA. ed. Research and Development in Expert Systems 111. Cambridge: Cambridge University Press, 1986: 175-83.

17. Barber B, Jensen OA, Lamberts H, Roger-France R, De Schouwer P, Zollner H. The six safety first principles of health information systems: a programme of implementation, part 1: Safety and Security. In: O'Moore R, Bengtsson S, Bryant JR, Bryden JS, eds. MIE 90. Lecture Notes in Medical Informatics no.40. Berlin: Springer-Verlag, 1990: 608-13.

18. Barber b, Jensen OA, Lamberts H, Roger-France R, De Schouwer, Zollner H. The six safety first principles of health information systems: A programme of implementation. Part 2: Convenience and Legal Issues. In O'Moore R, Bengtsson S, Bryant JR, Bryden JS eds. MIE 90. Lecture Notes in Medical Informatics no.40. Berlin: Springer-Verlag, 1990: 614-9.

19. Gale J, Marsden P. Diagnosis: Process not Product. In: Sheldon MG, Brooke JB, Rector AL, eds. Decision Making In General Practice. London: MacMillan Press Ltd, 1985:59-61.

20. Patel VL, Groen GJ. Knowledge-based solution strategies in medical reasoning. Cognitive Science 1986; 10: 91-116.

21. Weed LL. The Problem Orientated Record as a basic tool in medical education, patient care and clinical research. Ann Clin Res 1971; 3: 131-4.

22. Fries JF. Time-orientated patient records and a computer data bank. Amer Med Ass 1972; 222: 1536-42.

23. Van der Lie J, Westerman RF, Mosseveld BMTh. Critiquing based on automated medical records: An evaluation of HYPERCRITIC. In O'Moore R, Bengtsson S, Bryant JR, Bryden JS, eds. MIE 90. Lecture Notes in Medical Informatics no.40. Berlin: Springer-Verlag, 1990: 369-74.

24. Gordon C, Fox J, Glowinski A, O'Neill M. The design of the Oxford System of Medicine: An overview. In: O'Moore R, Bengtsson S, Bryant JR, Bryden JS, eds. MIE 90. Lecture Notes in Medical Informatics no.40. Berlin: Springer-Verlag, 1990:225-70.

25. O'Neill M, Glowinski A, Fox J. A symbolic theory of decision-making applied to several medical tasks. In: Hunter J, Cookson J, Wyatt J, eds. AIME 89. Lecture Notes in Medical Informatics No.38. Berlin: Springer-Verlag, 1989:62-71.

26. Miller RA, Masaric FE. Quick Medical Reference (QMR): A microcomputer based diagnostic decision support system for general internal medicine. In: Miller RA, ed. Washington DC: IEEE Computer Society Press, 1990: 986-8.

27. Gell G. AURA: Experience and Future. In: Scherrer JR, Côte RA, Mandil SH, eds. Computerized Natural Medical Language Processing for Knowledge Representation. Proceedings of the IFIP-IMIA WG6 International Working Conference on Computerized Natural Medical Language Processing for Knowledge Representation. Amsterdam: North-Holland Publ Comp, 1988: 253-61.

28. Scherrer J-R, Baud RH, Hochstrasser D, Osman R. An integrated hospital information system in general. MD Computing 1990; 7:81-9.

29. Nowlan WA, Rector AL. Predictive Data Entry: A new paradigm for medical knowledge representation. In: Proceedings of AIME-91. Berlin: Springer-Verlag (in press).

30. Nowlan WA, Rector AL, Kay S, Horan B, Wilson A. A patient care workstation based on user-centred design and a formal theory of medical terminology: PEN & PAD and the SMK formalism (submitted to SCAM 91).