

**T. Beale**

Ocean Informatics, Australia  
Centre for Health Informatics and  
Multiprofessional Education  
University College London  
London, UK

## Synopsis

### *The Health Record – why is it so hard?*

It is often asked: what is the difference between health IT and IT in other domains? One well-known answer is “the patient”. Systems in other domains such as banking and airline reservation have “customers” or “travelers” but these are grossly simplified abstract versions of a person. “Patients” in clinical systems are anything but: their biological and social complexity is manifested directly in clinical information, posing a far greater challenge than in other domains.

Consider the complexity of the human organism measured in terms of its DNA: 30,000 genes containing 3 billion base pairs which specify the proteins, enzymes and other functional products of the human body. The total amount of information has been estimated as being equivalent to a medical school library containing thousands of books [1]. By contrast, a Boeing 747 only has 6,000,000 parts (half of which are fasteners), while 75,000 engineering drawings were used to produce the first 747 in 1970 [2]. If we allow 500 pages per book, this is the equivalent of 150 books, or 300 if we double it to allow for the most recent version of the 747. But the comparison is not so simple. In both cases, the true complexity relates to the dynamic system *defined* by the specifications, and particularly the interaction of the parts of the built sys-

tem with each other and their environment. While a 747 is a vastly complex (and successful) work of engineering, human complexity outstrips it by several orders of magnitude. The complexity of the human organism is manifested in proteins (and diseases of their genetic errors); the immune system, the body’s amazing multi-layered, multi-strategy communication and logistic defence force (implicated in diseases as diverse as asthma, hyperthyroidism and AIDS); exquisite works of engineering such as the nephrons of the kidney (whose failure may lead to the medical and social complications of dialysis and kidney transplantation); interaction between the human organism and pathogens (malaria, dengue and common cold), and the functioning of whole organs and systems such as the gastro-intestinal system (problems occur at every level from the genetic to poor lifestyle).

Of all the manifestations of human biological complexity, the brain is the most important. The expressed complexity of the brain is not primarily encoded in the relevant genes, but in what it can do due to its immense connectivity. Pattern recognition, language, emotions, planning, humour and moral thought are just some of its abilities. The brain is implicated in health issues including schizophrenia,

the emotional reaction to the death of a loved one, and the emotional ability of a patient to deal with a condition such as cystic fibrosis. In populations, human mental capabilities lead to complex emergent social structures and group behaviours, many having a direct bearing on health and its recorded information. Lack of nutrition in early childhood and subsequent poor diet and inactivity (often the combined result of poverty and aggressive marketing of fast food) may lead to diabetes; overwork can cause mental health problems; lack of economic opportunity may lead to substance abuse and all its concomitant problems. Another societal human behaviour which directly impacts on the design of clinical information systems is travel: people who are on the move, whether for holidays, while in the military, or as refugees all access the health system at different locations and for very different reasons. Other societal phenomena impacting on personal health include war and privatised health provision.

Miller [3] identifies eight levels of organisational complexity in living systems from the cellular to the super-societal; all are relevant in clinical information systems. This innate complexity of life at these levels of organisation – concretised in the “pa-

“patient” - is a key difference between clinical information and the information of other domains. In health, the “patient” cannot be abstracted away in the way a “customer” can be abstracted to a mortgage account in a banking system; on the contrary, he must be considered and recorded in all his detail.

The concrete informational complexity in clinical systems can easily be understood in the common case of an insulin-dependent diabetic patient. The diabetic patient’s GP records numerous things: observations of initial symptoms (glucose tolerance test result) leading to a diagnosis, the search for an appropriate specialist to manage the disease, and ongoing management of the patient over her lifetime. The GP and/or specialist may record counselling of the patient and/or parents, suggested lifestyle changes, and advice on how to recognise and manage hypoglycaemia. With other carers such as nurses, ophthalmologists and podiatrists, they will potentially record evidence and interventions relating to the complications of peripheral neuropathy, retinal disease, heart disease and stroke. Any one of these problems will balloon into numerous sub-problems lasting for years – each generating significant amounts of data in the patient record. The diabetic patient record will thus be an ongoing story of managed medication, laboratory test results, lifestyle management and intervention. Each step on the way is documented, and preferably shared among the care team. And like any non-diabetic, every other infection, pregnancy and mental health issue of the patient is of interest to one or more carers; most issues will interact with the others in some way. Diabetic patient care requires recording of information which ranges from the technically precise to subjective social narrative. If the items in a diabetic patient’s record were represented as linked nodes, the result would be more like the myriad fila-

ments of a large, messy spiderweb, than a neat ledger of simple entries. This common example shows why the computerised health record is so difficult to implement: it documents happenings at many of Miller’s eight levels, and it needs to work for all of them.

One of the symptoms of the difficulty is that graphical user interfaces (GUIs) for health applications are notoriously difficult to get right: how can a single GUI adequately represent an HbA1c result but also a family counselling session? Another is the design of an information model. How can one database schema accommodate thousands of qualitatively different kinds of medical information from “apgar score recording” to “cardiology examination”? In contrast, a database containing the design of a jumbo jet will be larger than a single patient record, but does not span such a range of phenomena at so many levels of organisational hierarchy.

The Electronic Health Record (“EHR”; ISO TC215 term [4]) is where the story of the patient comes together. A shared-care community-based patient-centric EHR is the desire of many health carers and their patients; the reality today is that most health records are still location-based, and hence more or less episodic. Nevertheless, apart from sharing among the wider care team, such records (usually called EMRs – electronic medical records, or EPRs – electronic patient records) exhibit many of the complexities of the full EHR, and have the potential to become integrated in a shared EHR computing framework.

Conceptually, the EHR seems simple: it is a shared repository where any carer can record and view observations, decisions, and intended actions relating to the patient. At any moment, a user can find out what the patient’s situation is. Decision support applications can interrogate it, and make diagnostic or therapeutic suggestions based on existing data. Sec-

ondary epidemiological analysis may interrogate thousands of de-identified records, searching for patterns relating to specific diseases. However if we compare its requirements to other types of systems, the level of difficulty of the EHR quickly becomes apparent. Consider the following requirements of the EHR:

- information and efficient user interface reflecting multiple levels of hierarchical biological and social organisation;
- mobile patients;
- longevity of information (e.g. 100 years);
- multi-lingual;
- data shared and authored by multiple users simultaneously;
- integrated with knowledge bases such as terminology and clinical guidelines;
- wide geographical availability of a given record to multiple carers and applications;
- consent-based, potentially fine-grained privacy rules on information use (with exceptions for emergency access);
- multiple sources of constant change to requirements including medical technology, clinical procedures and guidelines, genomic/proteomic medicine;
- reliable medico-legal support for all users.

For any one of these categories there is undoubtedly at least one application more exigent than the EHR. The problem that the EHR poses is that it has significant needs in every one of these categories, and as such stretches the boundaries of multiple areas of ICT at once. To make things worse, some of the requirements appear to be in direct conflict. Requirements relating to privacy and “need-to-know” access to information sit uneasily with those for open availability not only within and among care delivery enterprises and ultimately over whole countries, but also across the

software applications of multiple vendors. Controlled access to sensitive information for multiple users over a wide area will clearly remain a challenge for the EHR for some time to come.

Such needs are part of the second major difficulty for the EHR. The optimal deployment is as an interoperable, shared access resource within a distributed security infrastructure, not only within enterprises but between them. However, the prevailing paradigm of ICT delivery is intra-enterprise deployment and vendor “lock-in”. As a consequence, almost all attempts at deploying the EHR outside the boundaries of a single enterprise have met with economic, political and logistic barriers, on top of the not insignificant technical challenges. The truth is, health provision cannot function solely as a market.

In response to these challenges, progress in the world of the EHR is occurring in ways not always familiar to the typical corporate or government development. Standards, open source, and knowledge development are three areas where alternative paradigms are being brought to bear on the EHR. In terms of standardisation, requirements for the EHR information architecture have been described in ISO TS18308 [5], while ISO TC215 committees are making progress on various aspects of interoperability and security. The European standards organisation committee, CEN TC/251 [6] is near completing a major revision of its EN13606 EHR communication standard, which will find use not only in Europe but also in Australia. EN13606 will also be fast-tracked into ISO. The OMG’s Health Domain Taskforce [7] standards (which were arguably a decade ahead of their time) continue to find use, including as inspiration for more recent standards work. HL7, the US-based health standards organisation is making fast progress with its EHR-S (EHR System) functional specifica-

tion standard which is gaining wide acceptance across the US; its Clinical Document Architecture (CDA) XML-schema specification is being used in a number of emerging clinical document interoperability projects. Openly published specifications which are not of official standards but are nevertheless finding use in government and private development sectors include *openEHR* [8] and OASIS [9].

While standards provide a definitional basis for interoperability, the open source software development sector is providing real solutions which break the single-vendor, closed-data paradigm of commercial development. The most famous of these systems, the US Veteran’s Health Administration’s VistA system is enjoying renewed success. Newer projects which have found operational deployment include LANL’s openEMed [10], TORCH [11], gnumed [12], and a large number of EU-funded systems, including PICONIC [13] and HARP [14].

Another area experiencing significant research and development activity is that of knowledge. Clinical terminologies and ontologies continue to evolve, with SNOMED-CT [15] and LOINC [16] being well-known examples. Development of computerised clinical guidelines is growing, and is finding success with systems such as Prodigy [17] and Proforma [18]. A new area has emerged, loosely characterised as “templates and archetypes” which provides a formal way to link ontologies (which can for example define meanings for breathing and reflexes) and captured information (such as an Apgar score, a 5-way assessment of a newborn including a 0-2 score for breathing, reflexes, muscle tone etc). The technology of archetypes and templates is being developed by the *openEHR* Foundation, has been adopted by CEN TC/251 and is under consideration by HL7. Numerous research groups are studying the combined use of terminologies,

archetypes, and EHR systems, including at University College London, the University of Manchester, the University of Seville, the Middle East Technical University (Turkey), the University of Aalborg (Denmark), and the Mayo Clinic in the US.

Private enterprise participates in much of the afore-mentioned activity by way of membership in standards organisations and joint development with large government initiatives, such as the UK’s National Program for IT in Health (NPfIT), Canada Infoway, Australian HealthConnect and the emerging US National Health Information Infrastructure (NHII) programme.

As with any area of research and development, quality research relies heavily on validated empirical evidence. Unlike the activities at the highly innovative and creative leading edge, validation requires real, deployed systems and proper statistical studies much as for a new drug. Both implementation and validation exercises are expensive and time-consuming, and it is inevitable that such activity is some way behind the front line of work in the domain. The studies in this chapter constitute a prime example of much needed validation at the simpler end of the EHR spectrum, namely intra-enterprise ambulatory patient record systems. Although in terms of new ideas, such studies may not be ground-breaking, they are an essential part of the evidence base for future work. Such studies allow today’s researchers to confidently make assumptions about future approaches, rather than having to prove everything from scratch each time. They also turn up surprising facts about workflow and costs.

In the first paper, Hippisley-Cox *et al* provide solid statistical evidence of the most basic proposition of the EHR: that electronic patient records are of better overall quality than the equivalent paper records, despite the authors’ default assumption that weak

keyboard and computer skills might in fact result in worse recording. Nilsson, Åhlfeldt and Strender provide evidence of informational quality in the EHR, via their study of Swedish electronic GP records in which extensive narrative and high levels of coding (ICD10) were found. Those records organised using Weed's well-known problem-oriented [19] approach were judged to be of the best clinical quality. Rotich *et al* describe their experience with a patient-centric electronic record in rural Kenya, which replaces 30 years' of clinic-centric paper records. Basic innovations such as an encounter form, a unique patient identifier, and simple backup and security measures underpin a wholly electronic, if relatively modest health record. Their time-motion study shows how, over 10,000 encounters, the workflow efficiency has been significantly improved by the advent of the electronic record along with appropriate procedures. The last study, by Wang, Middleton *et al*, provides evidence of the significant financial benefits of electronic patient records in ambulatory settings in the US. Their analysis indicates the sources of savings include reduced transcrip-

tion, averted costs due to decreased utilisation, and fewer adverse drug events due to basic medication decision support. Such proof is essential ammunition both for obtaining funds and resources for the further development of the EHR, and for guiding future design approaches for the computerisation of what may be the most challenging domain of all.

#### References

1. Ross DW. The Human Genome: Information Content and Structure. Available at <http://www.hosprract.com/genetics/9906mmc.htm>. Accessed 10 November 2004.
2. Boeing website. [http://www.boeing.com/commercial/747family/pf/pf\\_facts.html](http://www.boeing.com/commercial/747family/pf/pf_facts.html). Accessed 10 November 2004.
3. Miller J G. Living Systems. University Press of Colorado; 1995.
4. ISO/TR 20514. Health Informatics - Electronic Health Record Definition, Scope and Context. 2005.
5. ISO/TS 18308:2004. Health informatics - Requirements for an electronic health record architecture. Available from <http://www.iso.org>.
6. CEN TC/251 prEN13606. <http://www.cen251.org>. Accessed 10 November 2004.
7. OMG HDTF standards. See <http://healthcare.omg.org/>, [http://www.omg.org/technology/documents/domain\\_spec\\_catalog.htm](http://www.omg.org/technology/documents/domain_spec_catalog.htm). Accessed 10 November 2004.
8. <http://www.openEHR.org>
9. OASIS. <http://www.oasis-open.org/home/index.php>. Accessed 10 November 2004.
10. OpenEMed. <http://openemed.org/>. Accessed 10 November 2004.
11. TORCH. <http://www.openparadigms.com/>. Accessed 10 November 2004.
12. Gnumed. <http://www.gnumed.org/>. Accessed 10 November 2004.
13. PICNIC. <http://picnic.euspirit.org/>. Accessed 10 November 2004.
14. HARP. <http://www.telecom.ntua.gr/~HARP/HARP/HARP.htm>. Accessed 10 November 2004.
15. SNOMED-CT. <http://www.snomed.org/>. Accessed 10 November 2004.
16. LOINC. <http://www.regenstrief.org/loinc>. Accessed 10 November 2004.
17. Prodigy decision support system. <http://www.prodigy.nhs.uk/>. Accessed 16 November 2004.
18. Proforma. <http://www.openclinical.org/>. Accessed 16 November 2004.
19. Weed LL. Medical Records that guide and teach. *N Engl J Med* 1968;278:593-600; 652-7.

Address of the author:  
 Thomas Beale  
 CTO Ocean Informatics, Australia  
 Senior Research Fellow  
 Centre for Health Informatics and  
 Multiprofessional Education  
 University College London.  
 E-mail: [Thomas@DeepThought.com.au](mailto:Thomas@DeepThought.com.au)