Chapter 2
Why Interoperability is Hard

This chapter takes a broad look at some of the reasons why health interoperability and healthcare computing in general are hard. The focus is on interoperability, but we also cover aspects of electronic patient records.

The benefits of joined-up healthcare, to provide the right information at the right time and place, depend on computer systems being able to exchange information in a way that is safe, secure, and reliable.

The term interoperability means different thing to different people. *HIMSS Dictionary of Healthcare Information Technology Terms, Acronyms and Organizations* lists 17 definitions from the strictly technical to include social political and organizational factors [1].

A widely used definition is:

> Interoperability is ability of two or more systems or components to exchange information and to use the information that has been exchanged. IEEE [2]

This includes two separate ideas: first, the exchange of information, which is technical interoperability, and second, the ability of the recipient to use that information, which is semantic interoperability.

A third concept, pertaining to the actual use of the information, is process interoperability. The HL7 EHR Interoperability Work Group has set out a framework, which covers these different points of view [3]:

Technical interoperability moves data from system A to system B, neutralizing the effects of distance. Technical interoperability is domain independent. It does not know or care about the meaning of what is exchanged. Information theory, which shows how it is possible to achieve 100% reliable communication over a noisy channel, is foundation stone of technical interoperability [4]. Technical interoperability is now taken for granted.

Semantic interoperability: Dolin and Alschuler [5] define semantic interoperability as

> The ability to import utterances from another computer without prior negotiation and have your decision support, data queries and business rules continue to work reliably against these utterances.

T. Benson, *Principles of Health Interoperability HL7 and SNOMED*,
Health Information Technology Standards,
DOI 10.1007/978-1-4471-2801-4_2,

This is the core of what we mean by healthcare interoperability, ensuring that sender and recipient understand the same data in the same way. Semantic interoperability allows computers to share, understand, interpret, and use data without ambiguity. However, semantic interoperability is almost always specific to domain and context, and usually involves the use of codes and identifiers.

Process interoperability is achieved when human beings share a common understanding across a network, business systems interoperate, and work processes are coordinated. People only obtain benefits when they use information originating elsewhere in their day-to-day work. The importance of reengineering work processes to take full advantage of electronic systems has long been recognized, but the lessons have not yet been well learnt.

The more we understand about the three types of interoperability, the less likely we are to underestimate the work required to make health systems interoperable. These types of interoperability are interdependent, and all three are needed to deliver significant business benefits.

Why Standards are Needed

The number of interfaces needed to connect N systems increases exponentially using the formula $(N^2 - N)/2$. Linking two nodes needs only a single interface, which can easily be agreed by people sitting round a table; linking 6 nodes requires 15 interfaces; but linking 100 nodes requires 4,950 interfaces.

The center of the star at the right of the figure below (Fig. 2.1) indicates a single specification being used for linking the six domains. This replaces the 15 separate specifications shown on the left hand side.

The problem with standards is not that there are so many to choose from, but that we have not deployed those we have and that there has been no one with the power to make deployment happen. Standards that are not deployed are a waste of time and effort.

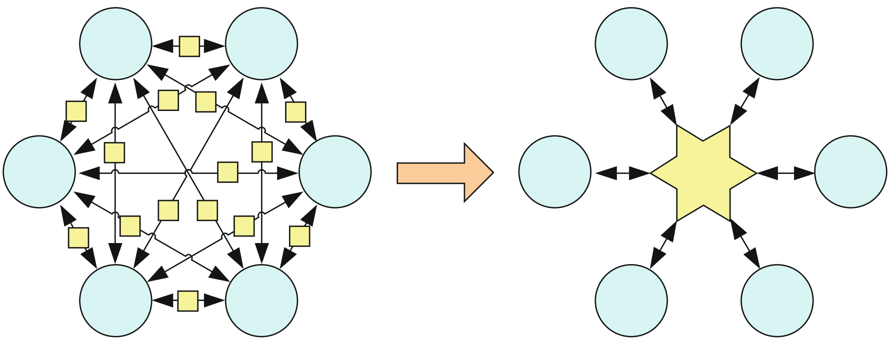

Fig. 2.1 The benefits of one standard

The number of transactions in healthcare systems can be vast. For example, a single system (the EHR system) at one large hospital (the Mayo Clinic in Rochester, Minnesota) processes over 660 million HL7 messages a year, or about 2 million messages a day (Anthony J, 2008, personal communication).

Examples of functions needing interoperability cover a wide range including:

- Requests for investigations such as laboratory tests and radiology
- Prescriptions for medication and other therapy
- Orders for nursing care, equipment, meals, and patient transport
- Investigation reports from laboratories, imaging, and other diagnostic departments
- Administrative data such as patient registration and identification; admissions, discharges, and transfers (ADT); and appointments
- Letters and memos from one clinician to another, including referral, clinic, and discharge letters
- Transfer and merging of electronic medical records
- Information used for management, audit, and monitoring
- Commissioning, billing, and accountancy

Any interchange typically involves two translations: first from the native language of the sender to the wire format; second from the wire format to the native language of the recipient. Problems begin because every computer system stores data internally in a different way. This means that to communicate, data has to be translated from one format or internal language into another. The solution involves translating to a standard wire format (a *lingua franca*) such as a version of HL7 that is understood by each party.

The Rosetta Stone from ancient Egypt, now in the British Museum, provides an analogy. The Rosetta Stone contains the same proclamation in three languages, used by the priests (Hieroglyphic), the court (Greek), and the people (Demotic). In our context, the three languages could be those used by a sending system, the receiving system, and a common wire format used for information interchange, such as HL7. The meaning of a message is precisely the same in each language, but the notation is quite different. The inscribers of the Rosetta Stone only needed to perform their translation once only, but in computer interoperability, each and every message has to be translated from one format to another without error. The choice of interchange language is not sufficient to ensure interoperability. Each transaction needs to be defined in unambiguous detail as part of a complete, consistent, coherent, and computer-readable set of specifications for that transaction to ensure interoperability between the machines and minimize any possibility of error.

You may ask why have health interoperability standards worked in some places and not in others? One explanation is to consider the individual self-interest. It is usually in each vendor's financial self-interest to provide a proprietary nonstandard interface to a customer, even though they know well that this is ultimately creating an interoperability nightmare.

There is an analogy with *The Tragedy of the Commons*, where it is each farmer's interest to add an extra cow to the common grazing land, even though that degrades the pasture as a whole [6]. The selfish farmer gains 100% of the benefit from his extra cow, but the downside is shared between everyone. This is because the division of costs and benefits is unequal. It pays each farmer to add additional cows to the pasture even while it continues to degrade.

The solution to this conundrum is to appoint a *regulator* with teeth, who everyone agrees shall specify what standards shall be used within their geographical area, ideally in full consultation with all concerned interests. You have to do what the regulator says, assuming he has appropriate means to enforce what he says. The ONC's meaningful use regulations are one example.

Electronic Health Records

An electronic health record (EHR) is best thought of as a collection of statements, which are a faithful record of what clinicians have heard, seen, thought, and done [7]. The EHR is not really a collection of facts, but rather a set of observations about a particular patient, which have been made by clinicians, each at a specific time and place for some purpose.

Each clinical statement is an observation, and so it is quite possible for two statements about the same event to disagree with each other. Such disagreements can often be resolved if the context or provenance of each statement (who stated it, when, and where) is recorded. As with a work of art, a statement without provenance is of doubtful validity.

The ISO 13606 reference model for electronic health record communication sets out a useful hierarchical structure for clinical information in the context of exchanging clinical information between parties [8]:

- *Composition*: The EHR is made up of compositions. The composition is a key concept; it is the set of information committed to one EHR by a clinician relating to a specific clinical encounter. Each composition shares common metadata such as the author, subject (patient), date/time, and location. Progress notes, laboratory test reports, discharge summaries, clinical assessments, and referral letters are all examples of compositions.
- *Folder*: Compositions may be grouped together into folders and subfolders. Folders may be used as containers for various purposes, grouping together the records by episode, care team, clinical specialty, condition, or time period.
- *Entry*: Each composition comprises a number of entries, also known as clinical statements. An entry is the information recorded in the EHR as a result of a single clinical action, observation, interpretation, or intention. It may be thought of as a line in the record. Examples include the entries about a symptom, a laboratory result, a diagnosis, or a prescribed drug.
- *Section*: Entries may be grouped together in sections. A section is a grouping of related data within a composition usually under a heading such as Presenting History, Allergies, Examination, Diagnosis, Medication, and Plans. Sections may have subsections.

- *Element*: The leaf node of the EHR hierarchy is an element, which is a single data value, such as systolic blood pressure, a drug name, or body weight.
- *Cluster*: Related elements may sometimes be grouped into clusters. For example, systolic and diastolic blood pressures are separate elements, but may be grouped into a cluster (e.g., 140/90), which represents one item in an entry (Fig. 2.2).

Every clinical specialty has its own way of working. The grand vision of joined-up healthcare is predicated on the notion that patient records can be shared electronically between clinicians from different specialties. Historically, this has been largely wishful thinking. There have been few successes and many failures. Yet, this is the promise of President Obama's HITECH Act and

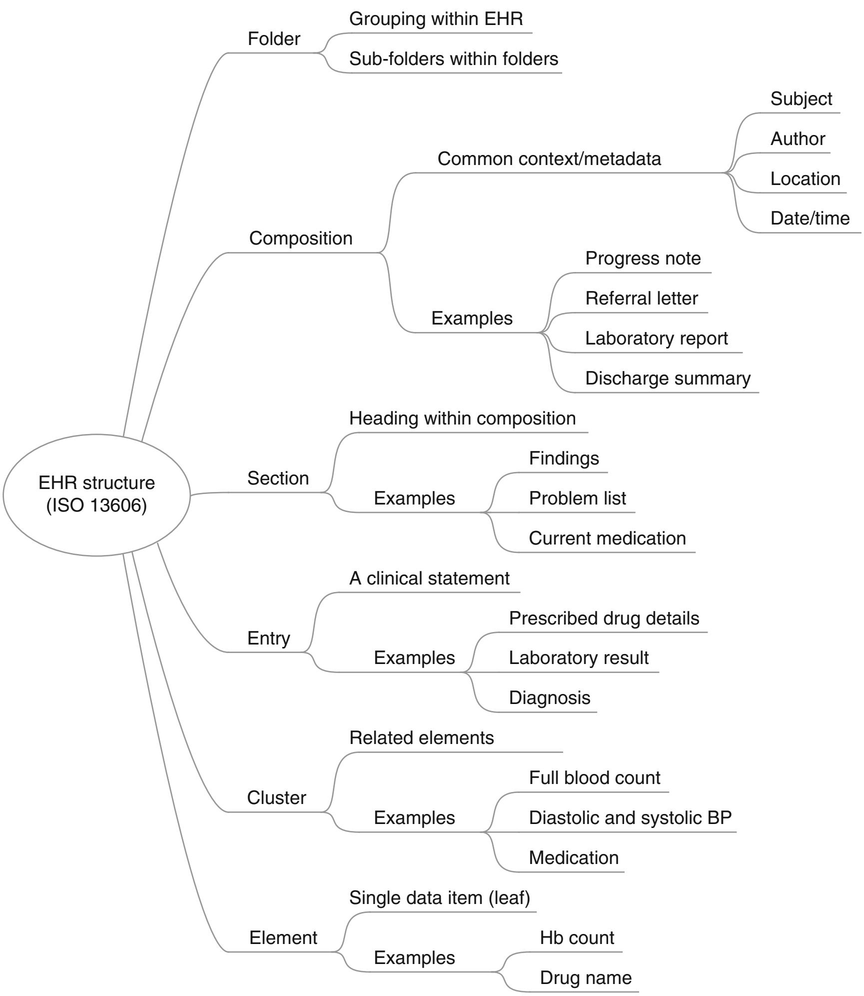

Fig. 2.2 Structural aspects of the EHR

was a specific but unrealized objective of the ill-fated NHS National Programme for IT.

It is difficult to share information between different computer applications even within the same specialty. This is because each computer application stores data in a different way and may use different internal codes. Furthermore, even within the same specialty, the information about an outpatient visit can differ greatly from that about an elective surgical operation or an emergency admission.

The GP2GP project in England illustrates the point. Patients have a lifelong medical record, which follows them when they move from one GP to another. In an ideal world, each patient's records would be sent electronically from their old practice to the new in a manner that avoids the need to reenter information. The GP2GP project set out to do just that, although the project's leaders recognized that it could a poisoned chalice [9]. The work has been every bit as difficult as predicted, and each record has to be carefully quality-checked before transmission and on receipt; after 10 years work it is a qualified success.

Now, consider some of the variety by thinking about some examples of a "discharge letter" [10]. Each of these is a composition and relates to a discharge event, but the information contents are diverse:

- An elderly patient discharged home after recovering from a fractured femur after a fall
- Mother and baby following birth
- A family after a course of counseling by a clinical psychologist
- Initial consultation report from an ophthalmologist notifying a proposed operation for cataract
- Letter to GP notifying that a patient has been diagnosed with cancer and outlining the treatment plan
- Discharge from hospital following coronary artery bypass grafts
- Postmortem report

Similarly, the content of different types of clinical laboratory report are quite different, as is the work done in each type of laboratory. Histopathology examines cells with a microscope; microbiology grows bacteria to identify them; hematologists count blood cells, and clinical chemistry measures chemical concentrations by measuring the intensity of color changes when chemicals are added. The only commonality is that they all work with specimens extracted from patients.

Yet, even the specimen workflow is not fixed; sometimes the requester supplies the sample, sometimes the sample is taken by the laboratory, and sometimes the patient is required to be present in person.

The Devil is in the Detail

Much of the hard work involved in health interoperability and healthcare computing lies in teasing out the detail of hundreds and thousands of different use cases. Information technology analysts and those who pay them tend to focus on the high-volume transactions, which are common across all specialties, and pass over the specific needs of smaller specialties. Yet, the common stuff is not usually the most important clinically. In medicine, as in art, the value of any piece of information is often related to its rarity.

One way to simplify the problem is to distinguish between information that needs to be processed by computer and that which needs to be read and understood by human users. Humans are good at judging the significance of small discrepancies, but digital computers are unforgiving of a single unexpected bit.

Computer processing is essential when data has to be identified, matched, retrieved, or counted. For computer processing, the information needs to be structured, complete, unambiguous, and validated.

Human readers need information in a format that they can understand. This does not have to be highly structured, although it needs to be easy to read and accompanied by supporting contextual data such as who wrote it, when and where, and for what purpose.

Even apparently simple concepts such as name and address are surprisingly complex on close inspection. In spite of years of effort, there are still no satisfactory international standards for name and address. Even the order in which names and addresses are written varies substantially between countries.

For example, one person may have several names and several addresses, which they can change at will. A woman may use her maiden name in one context and her married name in another. One person may use several addresses (home, work, previous, holiday, etc.), and each address is likely to be associated with different sets of people, such as family members, friends, or colleagues.

Addresses

We all know what an address is, or do we? Some of the difficulties were outlined at an ISO workshop on address standards [11]. An address may be defined as a label used to reference a geographical object such as a property through the use of identifiable real-world objects.

The most common form of address is the postal address, used for the delivery of mail, where the address is essentially a routing instruction leading to the property. However, addresses are not restricted to properties that receive deliveries of mail. They need to be created for a wide range of geographic objects such as:

- Domestic properties
- Commercial properties
- Industrial premises
- Public buildings (schools, hospitals, prisons, halls, leisure facilities, public toilets, etc.)
- Other buildings (churches, monuments, etc.)
- Places where events take place (sports fields, parks, etc.)

While some of these have postal addresses, many do not, if they do not receive deliveries of mail. However, they need to be identified and accessed for a range of purposes, such as:

- To identify delivery points for goods or services
- To uniquely identify people via their place of residence
- To identify where people live and work, for planning public services
- To levy taxes on people and organizations
- For deployment and contingency planning by emergency services
- For property registration and transactions
- To identify customers and potential customers

Addresses generally follow a simple structure incorporating the names or numbers of a nested set of spatial units:

- Subunit within a building or property
- Building or property within a street
- Street but some rural areas do not have street names
- One or more geographic areas (locality, town, county, etc.)
- Country

Part of such an address is often abbreviated by a code (e.g., a postcode or area code). The exact definition of each of these levels in an address varies from country to country.

The granularity of property units is an issue. Many buildings large and small can incorporate several addressable objects within a single physical structure. Such multiple-occupancy properties include bed-sits with shared bathrooms and/or kitchen facilities, shared houses, student and worker accommodation, residential care homes for the aged and disabled, flats with third-party access to the inside of the property for delivery purposes, flats where there is a single point of delivery for all residents, business premises with residential owners, managers or staff, shared business properties with no particular differentiation (normally companies that are associated companies), businesses, each with their own private area, but shared reception and toilet facilities, and self-contained businesses with one shared entrance. There are no clear rules to define at what level of granularity these types of premises should be recorded.

The life cycle of an address is complex. Addresses often need to be created before the building itself is created. For example, temporary addresses are often allocated during the planning or construction phase of new developments. Changes to addresses can occur due to merging of two or more properties, extension, subdivision or demolition of a property, change of property number or name, occupancy or use and the names of areas used in the address (e.g., due to administrative area reorganization).

Complexity Creates Errors

Building a single link to exchange data between two computers is relatively straightforward. Everyone sits round a table and works out what they are going to do. This approach works for very small projects, where each person is colocated, but it does not scale. An alternative approach is to provide rigorous implementation guidelines, but these are often complex and voluminous. For example, in the NHS Pathology Message Implementation Project (PMIP), a successful national project to send clinical chemistry and hematology laboratory test reports to all GPs in England, the implementation guidelines comprise almost a million words (more than ten times the length of this book). However, the endeavor to be rigorous can create errors caused by the sheer length and complexity of the specifications.

A related problem arises when the domain experts (such as doctors, nurses, and managers) are unable to fully understand these specifications due to the complexity of language or simply the time it takes to read them. As a consequence, these specifications may not be reviewed at the specification stage as thoroughly as is required.

Errors multiply according to:

1. The probability of misunderstanding any part of the specification. This depends on difficulty of language and domain and technical knowledge of participants (people with high levels of both technical and domain knowledge are rare).
2. The length of specification. In a long specification, exactly the same idea may be presented in different ways in two places, but each may be understood differently. If large blocks of information are replicated in different sections, with small but important differences, these differences may be missed.
3. The number of options permitted. Optionality greatly increases the chance of error.
4. The number of times different implementations to be made. Each implementation involves mapping or translating the specification into the local implementation language.

Misunderstanding inevitably leads to error [12]. Errors increase costs and reduce quality, create delays, and hit profits and reputation. Successful specifications avoid errors by limiting scope, being easy to understand, relatively short and simple, with few if any options. Many problems could be avoided by adequate thought and preparation by both users and suppliers. If time is running out, it is all too easy to be vague in a specification or offer the implementer a choice of options depending on the local context.

Often, both users and suppliers genuinely believe that they are in full agreement until the moment when users try to use the final product. Problems lie on both sides.

Users do not fully understand what they want, let alone what other parties can or cannot provide; they do not commit enough time or effort up front to fully review written requirements specifications; they then will not commit to these and insist on new features after the schedule and budget have been fixed. Most end users are technically unsophisticated, do not understand the development life cycle, and are simply unable to perform the sort of scrutiny that is demanded of them. This is why users need a much higher level of education in health informatics than has been provided in the past.

Suppliers are equally guilty. They often try to shoehorn the users' requirements to fit their existing systems or patterns, believing that it will be quicker, cheaper, and lower risk to reuse what already exists while failing to grasp that the user really needs something else and will never be happy without it. Suppliers often lack specialized domain knowledge and fail to understand the user's business processes at the required level of detail. Many suppliers focus attention on the high-volume aspects of health informatics, in part because they lack the domain knowledge to deal with the idiosyncrasies of every specialty.

Shared meaning between computers requires shared understanding between all of the human participants. As an analogy, consider the purchase of a new kitchen. The kitchen designer prepares a plan of the new kitchen. This plan is checked, reviewed, and signed off by the customer and is the basis of the contract. This plan uses a precise technical notation, which provides a means of communicating precisely the user's needs to the implementer (manufacturer), in a form that can be understood by both. Manufacture only begins after the customer has agreed the specification. A different but related set of plans are used in the manufacturing process to specify the exact details of every last screw.

The challenge in interoperability is similar but harder; it is to ensure understanding both horizontally across business processes and vertically between users and developers (Fig. 2.3).

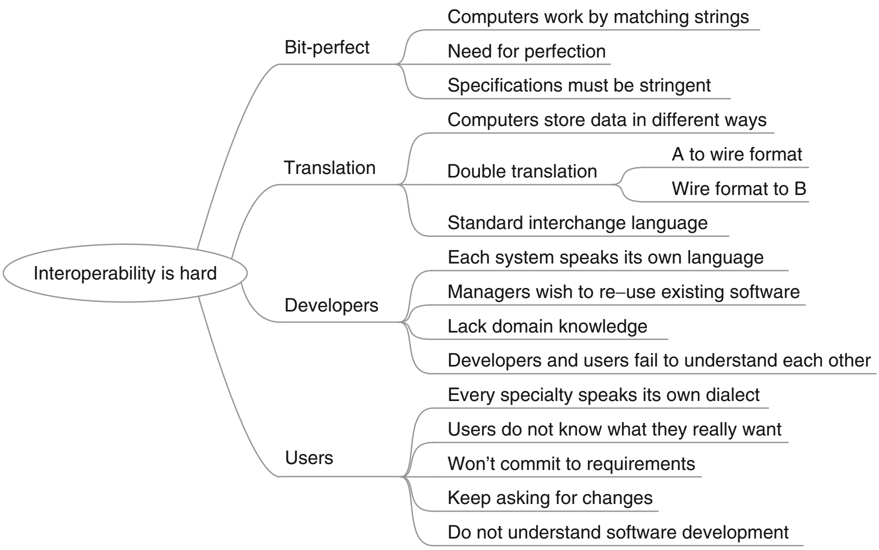

Fig. 2.3 Why interoperability is hard

Change Management

It is a sad truth that many large health IT projects fail to achieve their objectives. British MP Maria Eagle was addressing the Chief Executive of the NHS about a succession of disastrous Health IT projects when she said:

> If you saw someone walk into a brick wall, pick themselves up and walk into that same brick wall again, and pick themselves up and walk into the brick wall again, would the thought maybe cross your mind that they might be drunk or of unsound mind or would you perhaps conclude that they were incapable of avoiding brick walls? PAC [13]

John Kotter [14] in his book *Leading Change* has identified eight common errors in managing innovation and change (Table 2.1).

Table 2.1 Errors and remedies in management of change

Common errors	Proposed remedies
1. Not establish a great enough sense of urgency	Establish a sense of urgency that something really has to be done (it even helps to have a crisis)
2. Fail to create a sufficiently powerful guiding coalition	Form a powerful guiding coalition of key stakeholders with position power, expertise, credibility, and leadership
3. Underestimate the power of vision	Create a vision and strategy which is desirable, feasible, focused, flexible, and can be communicated in less than 5 min
4. Undercommunicate the vision by a factor of 10 (or 100 or even 1,000)	Communicate the change vision simply with examples. This needs to be repeated over again in multiple forums; address apparent inconsistencies, listen and be listened to
5. Permit obstacles to block the new vision	Empower employees to modify structures and systems to bring about the changes required. Changes will be needed in process, workflow, and information systems
6. Fail to create short-term wins	Generate short-term wins which are visible, clearly related to the change effort and build momentum
7. Declare victory too soon	Consolidate gains, reduce interdependencies, and produce more change
8. Neglect to anchor changes firmly in the corporate culture	Anchor new approaches in the culture. This comes last, not first

References

1. HIMSS. HIMSS dictionary of healthcare information technology terms, acronyms and organizations. Chicago: HIMSS; 2006.
2. IEEE. IEEE standard computer dictionary: a compilation of IEEE standard computer glossaries. New York: Institute of Electrical and Electronics Engineers; 1990.
3. Gibbons P et al. Coming to terms: scoping interoperability in healthcare. Final. HL7 EHR Interoperability Work Group; Feb 2007.

4. Shannon CE. A mathematical theory of communication. Bell Syst Tech J. 1948;27:379–423, 623–656.
5. Dolin R, Alschuler L. Approaching semantic interoperability in Health Level Seven. J Am Med Inform Assoc. 2011;18:99–103.
6. Hardin G. The tragedy of the commons. Science. 1968;162(3859):1243–8.
7. Rector A, Nowlan W, Kay S. Foundations for an electronic medical record. Methods Inf Med. 1991;30:179–86.
8. Health Informatics – Electronic health record communication – Part 1: Reference Model, International standard, ISO 13606–1. Geneva: ISO; 2008.
9. Purves I, Fogarty L, Markwell D. The Holy Grail or poisoned chalice: the GP-GP record transfer project. Newcastle: HIRI; 2001.
10. Benson T. Why industry is not embracing standards. Int J Med Inform. 1998;48:133–6.
11. Walker R. A general approach to addressing. In: ISO Workshop on address standards: considering the issues related to an international address standard. Copenhagen; 2008. p. 23–7.
12. Benson T. Prevention of errors and user alienation in healthcare IT integration programmes. Inform Prim Care. 2007;15(1):1–7.
13. PAC. Public accounts committee sixty-second report. London: House of Commons; 1998.
14. Kotter JP. Leading change. Boston: Harvard Business School Press; 1996.