

Organising medical and health-related information

An account by Conrad Taylor of the 7 June 2018 meeting of the Network for Information and Knowledge Exchange, held in Leeds.

The speakers were **Conrad Taylor** and **Ewan Davis**. Sadly, the event was not well attended, but the subject matter and discussion were of high quality, and therefore worth sharing. The text below is a page-formatted version of the Web report at <http://www.conradiator.com/kidmm/netikx-ehr-leeds.html>. I shall be linking this print-oriented version (as PDF) to that page.

Background to meeting

When the NetIKX Committee asked members for ideas for future development, some asked for meetings to be held outside London.

This seminar was co-sponsored by the UK chapter of the **International Society for Knowledge Organization** (ISKO UK), and was held in Leeds, where we were kindly hosted by the **Medical Protection Society** (MPS). MPS is an organisation with nearly 300,000 members worldwide, medical and dental practitioners, to whom it provides professional indemnity insurance, expert advice, and representation at tribunals.

Because of MPS involvement, and also because Leeds and neighbouring Bradford are centres of medical training, we chose the topic of medical information. Leeds is also home to two leading suppliers of GP practice medical records management software – EMIS and TPP.

Our main speaker, Ewan Davis, talked about a particular aspect of the electronic medical records system – namely, why future development in this area should be based on open standards. Conrad Taylor, who has collaborated with Ewan and others on health informatics projects within the BCS and the HANDI project, opened the meaning with a short overview, and after Ewan's talk, we continued with informal discussion, some highlights of which are also reported below.

An overview of medical information and knowledge

I spoke about this topic. As a warm-up, I spoke about the role that information and knowledge have played for centuries in the practice of medicine; what we mean by 'the medical record'; and what are some of the key issues, difficulties, and approaches to solutions.

On medical knowledge: Influential individual medical teachers and scholars of the classical period include Hippocrates, Galen, and the less well known Dioscorides, whose *De Materia Medica* was a valued pharmacopeia and catalogue of herbal treatments.

The first 'university teaching hospital' was the Academy at Gondishapur in pre-Islamic Persia, around 520 CE, which blended Greek and Syriac, Iranian and Indian medical

traditions, translated existing literatures into Syriac, and laid the basis for later Muslim medical scholarship.¹ An interesting innovation is that junior doctors were trained practically in a hospital environment, and were supervised by the whole medical faculty, not just a single master.

Modern scientific medical knowledge has now developed into a vast and hugely complex field: no one person could know it all. Knowledge development in medical science involves observation, research and experimentation, and produces huge volumes of texts and data. How this is managed, classified, indexed and accessed would be a topic for many days of conference!

Until the 20th century, knowledge about diseases and how to prevent or treat them was based on observing correlation rather than discovering causation. Often, people jumped to wrong conclusions. The word ‘malaria’, for example, is from the Italian for ‘bad air’ – the fevers were thought to be caused by poor air quality around swampy land. The first clues that it was caused by a micro-organism, and that mosquitos could be a vector for disease transmission, only arrived in the 1880s. Modern biomedical science has put medical knowledge on a better footing.

Information in the service of health and medical practice

If we turn from research to the kind of practical knowledge and information that medical folk need in order to do their jobs, we could divide it into four types:

- ◆ information about diseases and ailments, to assist in diagnosis and expected prognosis;
- ◆ information about available treatments and other interventions;
- ◆ information about service availability, for example when your GP sends you to the hospital for an echocardiogram;
- ◆ information about the individual patient, which is held in medical records.

There is arguably another kind of health information – the kind we need to better look after ourselves, such as understanding consequences of alcohol intake, what to do if you cut yourself, come down with ‘flu, etc.

Introducing ‘the medical record’

Wikipedia gives an adequate definition of what we mean by the Medical Record: it is the systematic documentation of a single patient’s medical history and care, across time, within one particular care provider’s jurisdiction.²

1 See Wikipedia article https://en.wikipedia.org/wiki/Academy_of_Gondishapur

2 From the 1960s, the structure of medical records was influenced by the writing of the innovative American physician and teacher, **Lawrence Weed** (1923–2017). In opposition to a tendency for records to be structured by source (e.g. x-rays, prescriptions...) his **Problem-Oriented Medical Record** system grouped all documentation around the patient’s medical problems. In the late 60s, Weed also helped develop an early computerised medical record system, and he helped found the American College of Medical Informatics.

The pre-digital record was a paper system, in which a physician wrote notes by hand so that at a later date he could recall what the patient told him, what he himself observed, what his assessment was, and what course of action he decided on, including medication. A problem was that doctors tended to write in a 'shorthand code' of jargon and abbreviations which would not make sense to the patient, and perhaps not even to another physician.

Even before medical records could be computerised, there were attempts made to standardise nomenclatures and abbreviations, and to reduce the possibility of dangerous communication errors with the use of pre-printed forms (e.g. when ordering blood tests).

Also note the 'one care provider's jurisdiction' issue – with an increasing trend towards medical specialisation, the patient may be referred to screening and scanning services, maternity services, community clinics etc. The patient does not have only one medical record, but a collection of records held by different care providers. This is the root of many problems.

The electronic record

Use of computers to create and maintain medical records started in the late 1960s and strengthened in the 1980s. Today every GP practice uses electronic health records, integrated in systems which also handle appointments, prescriptions, and routine office automation tasks. In the UK, a handful of suppliers created these systems; the largest shares of the current market are held by EMIS and TPP SystemOne. Both, as it happens, are headquartered in Leeds. GP systems development in the UK has often been led by people who were themselves doctors in general practice, and knew what was needed.

In parallel, systems of codes were developed which could be entered into the record as a compact standardised shorthand. Tim Benson recalls working with Dr James Read on the first iteration of what became the **Read Codes**, commonly used across General Practice in the UK and New Zealand:

“We recognised that like all computer users, GPs are fundamentally lazy and mildly computer-phobic. We wanted a coding scheme that would allow a one-fingered typist to enter data in the consulting room, by typing in a few letters and the computer doing the rest... We also wanted a system that would be quick to use, so that GPs could use it themselves in the consulting room, and could generate reports almost instantly.”

The Read codes were first published in May 1986 with a total of 15,124 codes, identifying diseases, procedures, patient history, examinations, preventative actions, plus the patient's occupation and some administrative codes. The Read codes are constructed hierarchically, becoming more specific from left to right. Thus an initial **G** indicates all circulatory system diseases, **G7** limits this to cerebrovascular diseases, and **G712** signals an intra-cerebral haemorrhage. In the context of the limited data storage available on the micro-computers of the mid-eighties (data for 3,000 patients on a single floppy disk?) this conciseness was very valuable.

Read codes were designed for general practice and made little sense in a hospital context. But further developments in the UK between 1990 and 1998 led to the scheme called **Clinical Terms Version 3 (CTV3)**, which were a few years after merged with the American Systematized Nomenclature for Medicine Reference Terms, to produce **SNOMED-CT**, around which the NHS is now standardising.

But note that coding, and terminologies, despite their many advantages for ‘computable healthcare’, can never make a complete medical record. To understand and track progress in some conditions, ‘imaging’ in its broadest sense needs to belong to the record (ultrasound, ECG, X-ray, CT, MRI etc). To follow the patient journey in a holistic way, a strong narrative component is also helpful, especially in mental health, and social care.

Terms: EHR and PHR

Throughout this report for convenience I am using the term **EHR**, Electronic Health Record, to signify a record of a patient’s health maintained by the clinician(s). Alternatively the word ‘Medical’ or ‘Care’ is used in the middle (thus, EMR, ECR)...

In the discussion after Ewan’s talk the **PHR** (personal health record) also came up, and the convention is that this is a record managed by the patient. We also discussed what is sometimes called the co-PHR, a record co-produced by patient and clinicians.

Crossing the paper bridge

While almost every GP uses a computer in the consulting room, most hospital doctors do not. Some GP practices are paperless, or at least ‘paper-light’ (paper documents are scanned and stored as page images); meanwhile in hospitals it is a common and depressing sight to see big trolleys laden with fat folders of paper records being wheeled around the corridors.

Hospitals are complicated organisations, which makes them more difficult to computerise successfully. They are also Balkanised into specialist departments which often adopt their own records systems, as Ewan would later describe, making data integration difficult across the site.

The current state of play is a state of slow transition. Most textual or coded data pertinent to the health record today is captured in a computer system, usually via some kind of on-screen templated form, perhaps with contextual drop-down menus and pick-lists. But often the means of communication of that information with other care providers is via a printed letter or report, such as a pathology lab report or hospital discharge summary, and all the ‘paper-light’ GP practice can do to capture it is to scan it as an image, which is not computer-processable.

What we need is better interoperability between the computer systems. It is beginning to arrive. My blood test results now come to my GP practice electronically, and the repeat prescription requests which I make via the surgery’s Web portal are sent to my chosen pharmacy via the Electronic Prescription Service.

Dreams of a monolithic universal health records system, which arose in the early 1990s, are now generally realised to be unattainable, and perhaps undesirable. This has

strengthened the case for **interoperability**, so that patient information can transfer between different healthcare providers' computer systems as smoothly and securely and efficiently as possible.

Standards and Openness

Shared standards make interoperability possible – after all, without them, we wouldn't have the Internet! In the health field, one important international family of standards is **Health Level 7**, which got started in the late 1980s and is beginning to get some real traction. HL7 standards have the purpose of helping both clinical and administrative data to move between the software applications used by various healthcare providers, without everyone having to use the same software.

For example, the CDA standard (**Clinical Document Architecture**) defines an XML-based markup scheme for the electronic transfer of documents such as hospital discharge summaries or specialist reports, helping to define the document's structure and semantics (the meaning and purpose of its parts). CDA can cope with structured or unstructured text, images, and links to other files; the structured bits rely on standardised coding system vocabularies such as SNOMED. In the UK, the NHS Interoperability Toolkit (ITK) recommends using CDA.

Another more recent HL7 standard which NHS England is getting enthusiastic about is **FHIR** (Fast Healthcare Interoperability Resources, pronounced as 'fire'). While CDA takes a document-centric approach to clinical communication, FHIR is a way of defining data formats and data elements as 'resources' which can be provided as a lightweight information messaging service from one computer system to another, perhaps directly populating the database on the receiving end.

It has been argued that because FHIR is based on modern web-based technologies such as Representational State Transfer (REST), Cascading Style Sheets to determine screen representation, and data encoding with XML or JavaScript Object Notation (JSON), it's easier to find programmers who can spin up such services quickly, write 'transform' processes for getting database content into FHIR packages, and even share healthcare data to web-oriented devices such as tablet and cellphone apps.

One overall barrier to integration is that most healthcare software systems are proprietary. Getting vendors to reveal the data structures they use, or make needed changes, is difficult. There are severe issues with 'vendor lock-in'. This is where the argument comes in for Openness – using an open standard for the healthcare record, such as OpenEHR, or even Open Source software tools, as our main speaker, Ewan Davis, would explain.

Open standards and systems for healthcare records

Introducing Ewan Davis: Since 1981, Ewan Davis has worked in computerised health informatics – what is now called 'digital health'. He founded AAH Meditel, in its day one of the leading GP systems. Since that was sold he has been an independent consultant/contractor, both with NHS bodies such as NHS England, NHS Digital and Connecting for Health, as well as on the industrial provider side.

In terms of voluntary engagements, Ewan has been the chair of the British Computer Society Primary Healthcare Group. Also, about six years ago he was a prime mover in setting up HANDI Health, a not-for-profit which encouraged the trend towards using apps on hand-held devices for health purposes.

(Conrad adds: ‘I got involved in the documentation aspects of HANDI, filming and recording a number of events.’)

Ewan is now CEO of Inidus, a new company committed to delivering a secure cloud-based platform for health and social care applications, building on open and well-documented standards such as HL7 FHIR, SNOMED CT, and IHE Cross Document Exchange (IHE-XDS).

Defining an open platform

What is an ‘open platform’? Ewan flourished a copy of a publication he recently wrote for the Apperta Foundation, ‘Defining an Open Platform’. (Apperta is a not-for-profit created with grant aid from NHS England – in our discussions later he told us more about where it came from.) That publication is the ‘manifesto’ for what he would be talking about – an approach that seeks to end vendor lock-in and liberate data. He had brought a few copies to give out.

You can get your own copy from <https://apperta.org/news/open-platform-rfc>

Problems and opportunities for Digital Health

The healthcare system is a carer in crisis – not only the NHS, but right across the developed world, due to growing demand from an ageing population, the workforce problem, and growing costs. Faced with this, there is a desire to adopt digital technologies to deliver ‘transformational change’ in healthcare, similar to changes in banking and finance, retail, travel etc. But in healthcare, that hasn’t happened. Why?

Approaches so far have resulted in data silos. The GP practice has a system; the hospital site could have several thousand systems, storing data in proprietary formats, making it difficult to share data between them; bits of an individual patient’s care might be represented in a number of these, so it is difficult to get a holistic picture of the patient’s health and social care record.

That leads to data lock-in, and vendor lock-in. Anybody who has had to go through the process of changing from one GP system or hospital system to another is unlikely to want to go through the same experience again. In London, several hospital-level systems from Cerner were installed under the NHS National Programme for IT (NPfIT). It was a painful process and didn’t go as well as it might. But when the five-year contract came to an end, those hospitals re-procured Cerner – not because they loved it, but because the thought of changing to an alternative was unbearable.

There are only four vendors of big hospital systems, all American: Epic, Cerner, Allscripts, Meditech. In the UK GP market, for all practical purposes there are now only two vendors, EMIS and TPP. Likewise, in the pharmacy sector, there are a couple of vendors; in maternity systems, a couple of vendors, so there isn’t a lot of choice.

Indeed, there has been no new entrant to the UK digital health market for 25 years that ever got to a scale larger than an SME.

This means that there has been very little innovation; and without innovation, you don't get transformational change. It's not that the incumbent vendors don't understand the sector – they have smart people who are very clued-up about what healthcare IT can do. But the companies are locked into technology and business models which belong to the last century, and have little motivation to change. The market is broken. We need a new paradigm.

Three kinds of information to enable medical practice

Ewan is just one among many across the world who believe the answer is to move to Open Platforms. That is about making the data that healthcare applications need available in an open, computable, shareable format. That data, for the clinical practitioner, has three elements:

- ◆ Information about the individual patient – the electronic health record;
- ◆ Medical knowledge;
- ◆ Information about resources available to call on.

What any clinician does – and therefore, what any application supporting that clinician must do – is to combine these kinds of information so that the patient's health issues can be diagnosed and a course of action chosen, within the constraints of the resources available. It would be ideal to have these three kinds of information available in an open, computable format. We have made some progress as far as information about patients is concerned.

Computer systems in healthcare don't support clinical practice well – they can even get in the way. Consider the Epic system, which in 2014 went into Cambridge University Hospitals NHS Foundation Trust at a cost of £200 million. After 2.1 million records were transferred to it, the system became unstable, information exchange was poor, there were delays in emergency care, problems with clinical letters and pathology test results. There were costs in suffering, delays, and possibly lives.

But it is hard for a new supplier to break into this market and do better. Healthcare is complex; data in healthcare is complex. There are significant regulatory barriers, with good reason – there is a lot of sensitive information about individuals, posing many governance issues. There are issues of clinical safety. Finally, the commercial environment is difficult: selling to the NHS is hard (in general, public procurement is a nightmare). And when your potential customer says, 'show us examples of your system in use', it's difficult for newcomers.

What we need instead is progressive development of health informatics. This is where you have data in open and portable formats, so a new developer doesn't need to work out from scratch how to represent blood pressure in their system. As for the methods for devising standards around this area, it needs to become more democratic and more agile – currently it's too slow. These are two things Ewan later addressed.

Anatomy of a healthcare application

A healthcare application always has three components – even though this division may not be apparent to the user, or even the developer:

- ◆ somewhere that you store the data – a database;
- ◆ an information model, which defines how the data in which we are interested will be represented;
- ◆ a user interface, which deals with the display of the data and provides a means of interaction with the system.

With most applications, those three things are all tangled up together in the software code, without clear separation.

What do we mean by the information model? An information model may be explicitly defined; it could be represented as a UML diagram. It may be embodied in a paper or on-screen form. It defines the structure of information, and what the various data ‘fields’ are permitted to contain. It may define messages which can pass within the system, or between systems, or which will be represented to the user via the interface.

If we want to share information between systems, we need information models which those systems also share. Ewan showed us two mismatched models for administering medication. One, a hospital system, defined substance, route of administration, dose, and frequency, because hospitals tend to do dose-based prescribing. But GPs tend to look at supply into the patient’s hands, where the relevant fields are product name, strengths, and timing. Between these two information models, a satisfactory transformation isn’t possible for lack of a shared model.

Introducing the potential of OpenEHR

What Ewan and others are looking to for help is **OpenEHR** (sometimes pronounced like ‘open air’). This is not software: it’s a specification for how to represent clinical content, defined in a vendor-neutral and well documented way. OpenEHR is maintained by a charitable foundation based at UCL.

OpenEHR was developed from work in the UK, mostly in London at UCL/UCH, but like many British inventions is implemented more widely elsewhere. Brazil has adopted it as the basis for their national electronic health record, Moscow’s health system runs on it, it’s in use across Norway. And in Leeds, it is the basis of the Personal Health Record.

OpenEHR allows you to specify an element of clinical content, called an **archetype**. So, for example, a ‘blood pressure’ could be defined as an archetype; as might an allergy. An archetype can be linked across into a standard terminology – for use in the NHS, and indeed generally across the world, that would be SNOMED CT, which arose from an NHS & US collaboration, and is now managed by an independent organisation based in the UK. Ewan later returned to the detail of how archetypes are developed.

Megasuites, best-of-breed, feral systems

Ewan used the term '**megasuite**' to refer to a software system that aims to handle the bulk of the information management requirements of a healthcare provider. Epic aims to do most of what an academic medical hospital does; TPP or EMIS aim to do most of what a GP practice requires. In a hospital setting, such a suite might cover about 70% of that need; in the GP sector, it is more like 90%.

The alternative approach is '**best of breed**', where systems are procured that handle information for specialist areas rather better than the megasuites can; then the challenge is to try to get them to work together. Even in megasuite deployments, other more targeted systems get added on to the side, for example for cardiology or maternity services. But these become data silos; so for example, if a maternity patient needs to see a cardiologist, the data about her pre-eclampsia may not be available to the cardiologist.

Then there are what Ewan called '**feral systems**', cooked up on the quiet to help where larger systems fail. If you were to investigate a large hospital, there could be a thousand information systems that the CIO doesn't know about! Often these are simple databases knocked up in Microsoft Access.

(While contracting for BT, Ewan did an audit of Guy's and St Thomas' and found 1,200 such feral systems. Because most of these kinds of data system breach governance and data protection rules, there were probably a lot more kept hidden from the audit team – maybe 7,000?)

Most megasuite vendors see the writing on the wall, and are trying to re-invent their systems as 'platforms' – sort of. This means, breaking through the sealed interface, and exposing an Application Programming Interface (API), through which third-party apps can access the database, and maybe write to it. But the megasuites still hold data in a proprietary format. Any serious challenge to make their data properly open, they would probably see as an existential threat.

Opening the architecture: benefits

Moving to an open platform architecture separates elements further. For a start, you work with a vendor-neutral information model. You can turn to a range of suppliers who can provide you with a suitable data store; anybody can produce an application that is compliant with that standard, and any application will work with any data store.

This gives you substitutability. If your data repository causes you grief, you can move your data to another run by a competing vendor, and it doesn't disturb your applications. Or if you don't like the prescribing software module you've got, you can switch to a better, without the need to transform the underlying data.

Moscow City Council did a pilot of OpenEHR. The city authorities run GP practice and quite a bit of outpatient care too. At first they patriotically chose a Russian vendor to supply the clinical data repository (CDR). They were happy with the pilot of OpenEHR, but they didn't like the performance of the CDR, so they brought in another from a Slovenian company (Marand – <http://www.marand.com/>). Marand, which runs the national health database for Slovenia, had originally used a back-end product from

an Australian company (Ocean), then decided they could do a better job themselves; when the time came to change, it took just four hours, and the data was not disturbed.

Similarly, when Leeds started its experiment with OpenEHR, they started using the Australian system, then moved to the Slovenian alternative, and are now moving to an Open Source CDR.

Defining clinical content

If want to model your clinical information, you can do so as OpenEHR ‘archetypes’, or FHIR ‘resources’, or InterOPEN ‘clinical data elements’. (They are very similar, with somewhat different approaches to implementation.) These things might be a blood pressure, an allergy, a prescription.

Building those models is a ‘long tail’ kind of problem. You could define a small number of such things and they would cover a lot of what a clinician would use on a day-to-day basis. Then there are less frequently encountered concepts – within clinical medicine, about 3,000 of them. If you want to extend the scope to Personalised Medicine, and the Internet of Things (devices and sensors around and on patients), then maybe beyond clinical medicine to general health and social care – well, Ewan guesses that there could be about 10,000 concepts you would want to define and maintain. Which is a massive task, unless you take a piecemeal and progressive approach.

The typical standards development process is ponderous and committee-based. In the health sector, usually there is late engagement of the vendors, who then have to cope with the expectations of the clinicians. There are fixed review cycles. If you want a new feature in the standard, it might take three years to add. So it’s no surprise that for lack of a fast standards-making process, vendors cook up their own internal standards, to get the job done. But there is an alternative: switch to OpenEHR.

(David Penfold asked which standards organisations are relevant in this field. In the NHS there is the Data Coordination Board, and the Professional Record Standards Body (PRSB) which was spun out from the Royal College of Physicians – its remit extends to social care, too. There are the formal standards bodies such as BSI and ISO; more specifically to healthcare, there are HL7 and SNOMED.)

Who’s adopting OpenEHR?

There are a number of examples of OpenEHR adoption: a Slovenian children’s hospital, a Norwegian GP system. Salford NHS trust uses AllScripts as its hospital megasuite, and isn’t likely to replace it, but they added an OpenEHR data repository beside it as an innovation-testing system: this is known as the ‘bimodal’ approach.

Plymouth have decided to go down a Moscow-like route: they have put an OpenEHR clinical data repository into their big teaching hospital. They started small with an e-prescribing system and a clinical portal. Plymouth has 190 applications within their overall architecture, and their intention is that over the next four years they will all move onto an open platform. Either existing application vendors will have to re-platform, or Plymouth will get new vendors in, who can comply.

Genomics England have chosen OpenEHR to store phenotypic data to get to grips with rare diseases (that's data about individuals, the 'phenotypes' – you wouldn't store genome sequence data that way). There are eight such data stores across London, with West Midlands and Manchester probably going to follow on.

In Scotland, OpenEHR is used to provide a common interface between existing systems, plus a new clinical decision support system from Scandinavian supplier Cambio. Leeds City Council has committed to going with OpenEHR for what they are calling the 'Person-Held Record', a patient-centred personal health record, potentially going even beyond health and social care concerns, for the citizens of Leeds.

The progressive standards development process

How can we democratise and speed up the standards development process?

At the moment, these standards tend to be defined by technicians within healthcare, fairly knowledgeable about its needs. But the people who really understand clinical content are the clinicians, and really they should be driving the process. (One reason why GP computing in the UK took off so early and so well, was because GPs led the development of systems: EMIS was founded by two GPs; AAH Meditel had about a dozen clinicians on the staff.)

To get clinicians involved in standards development, you've got to be practical – not many are interested in becoming technicians. OpenEHR provides a 'two-level modelling' approach, very accessible to a clinician.

Underpinning OpenEHR is a Reference Model, which expresses clinical content in fairly abstract terms such as an 'action', an 'order'. On top of that, you build these models called archetypes. For example, into a 'blood pressure' archetype you put everything you might possibly want to know about a person's blood pressure. Possibly no single clinician would recognise every part of that archetype – there are elements only of interest to a paediatrician, for example. But the archetype has been put together by a group of clinicians with interests in that subject.

Or, take the example of 'visual acuity', important to ophthalmologists. There are probably half a dozen people who know everything about visual acuity – one at Moorfields in London, one at Mayo in Minnesota, etc. They can work together and decide what should go into that archetype.

Archetypes are then assembled into what's called a Template. This does two things: it allows you to combine archetypes, and also to constrain them. If you are a GP, there are just a couple of things that interest you about blood pressure: systolic and diastolic pressure. A blood pressure measurement template for a GP could present only these parts of the underlying archetype for use. But for a diabetes assessment you want to add in information about blood glucose; someone else might want to record heart rate. By building templates you create re-usable components, shared with the commons, and linked into relevant terminologies such as SNOMED, ICD-10, LOINC. Effectively, those archetypes and templates are Open Source, and shared.

These components can then be used as a message format, or can inform an API definition for an application, or a GUI interface component such as an input form or

screen report. There are tools that will take these things as input and automatically generate forms from them, or messages which you can send using HL7 standards.

The role of CKM for collaborative development

The OpenEHR community has a kind of social media platform for standards evolution, called the **Clinical Knowledge Manager (CKM)**. It's an online hub which is used to document each element being developed or already published, or updated, and in all the languages into which it has been translated (including in non-latin scripts such as Arabic or Chinese).

If someone wants to represent clinical content in an application, the first thing to do is to see if the model has been worked out already; if so, you just take it and use it, saving a lot of time. If the concept hasn't yet been incorporated, you can find out if anyone is interested in working with you to develop that, which helps spread the cost and widens the pool of expertise.

So you assemble a group, a coalition of the willing, experts in that particular field. Rather than go through a committee-based approval process, the group beavers away together until they reckon, 'well, this is good enough'. The definition then goes into 'under review' status where comments and suggestions are made. After revisions, the editing group decides it should be published. Publishing is a significant step in OpenEHR; once a resource has been published, it is subject to fairly rigorous change control, because the user community needs to rely on the stability of that concept definition.

Standards definitions may be extended or revised, but they change slowly. Technology implementations can change fast – shall we encode in XML or JSON, for example? – but OpenEHR works at the information level. Again, the concept of 'blood pressure' provides an example – physiologically it hasn't changed, and the archetype definition has only been modified once, when technology evolved to make 24-hour monitoring possible.

The Clinical Knowledge Manager runs as a 'do-ocracy' – decisions are made by those who bother to get involved. The granularity of archetypes means that the process can function as a collection of virtual sub-committees, able to move at speed, and with no need to ever sit in a room together. As for the participation of formal standards bodies, the OpenEHR community sees that as a kind of secondary endorsement.

This is like the way that the development of Internet standards works. To propose a new Internet standard, you write a 'Request For Comments' (RFC), which gets debated and bashed into shape. Eventually the standard gets adopted by the IEEE or Internet Engineering Task Force.

In the case of OpenEHR concepts, one would hope to get endorsement from the likes of the Royal College of Physicians or the Professional Record Standards Body (PRSB).

Closing thoughts from Ewan

Some people have grand ideas about semantic interoperability of clinical data, and computable ontologies, but clinical data is messy. With OpenEHR, ambition is limited to getting enough sense out of the data to make it useful.

Ewan had a slide which mentioned FHIR, the HL7 standard for clinical messaging, but only to say that OpenEHR plays nicely with FHIR. The point which Ewan and others have been making to PRSB is that OpenEHR has been around for ten years and compared to FHIR is a more mature way to develop models of clinical content, but it is then relatively trivial to convert the concepts into FHIR resources, and there are tools around which semi-automate the process.

What is more difficult, is how to represent **clinical knowledge** in a standard format – there are approaches, such as the process description language PROforma, and Arden Syntax (a procedural language for representing medical algorithms for clinical decision support), which deal with some parts of it. Quite a lot of clinical knowledge is actually about workflows, so standard workflow modelling can help, at least for workflows which are deterministic – but many aren't.

Snippets from our discussion

When we arrived at our refreshment break, all but one of the MPS attendees had slipped away to their desks to get on with work, so the seven of us who were left gathered round a table to continue the conversation informally. Below are a few interesting points which came up.

Electronic observations ('e-obs')

In the hospital setting, on wards, logging observations is an activity which can consume a lot of nursing and care assistant time. For example, '**NEWS**' is the National Early Warning Score, a structured set of observations of patients in whom a risk of sepsis is suspected. Respiration, blood pressure and pulse, oxygen saturation, temperature and mentation are noted, and each measure converted to score. A higher score indicates greater departure from normal physiology.

Recording these on paper and manually doing score conversion calculations is messy, and fraught with the possibility of mistakes. It would save time if the data were entered into an app: machine algorithms can then generate the scores. Also, having NEWS records stored electronically and centrally would allow for automated triggering of alerts, and aggregated data allows supervisors to manage the hospital better. They have now implemented such 'e-obs' for NEWS at Guy's Hospital in South London.

Co-production, patient access, and patient data input

We discussed the idea of the 'co-produced record' – created and curated jointly by a patient and relevant clinicians. **There's increasingly a belief that a patient should have access to at least some of their clinical data.** These days, a patient could collect data (say, blood glucose measurements) to share with clinicians; being able to do so via an electronic portal, perhaps interfacing with a smartphone app, would help.

An SME active in this area is '**Patients Know Best**', founded by Dr Mohammed al-Ubaydli. In the PKB system, the patient has a facility rather like a social media account. Physicians and other care providers are invited to share data and messaging with the patient and other providers – and perhaps concerned family members and care-givers. This has worked well for patients with complex and multiple chronic conditions and tangled care pathways.

We also discussed where patient data should be stored, who decides that, and who should have access. Perhaps the patient should decide, and there are several business models for that. In the era of the global Internet, information storage doesn't have to be on site (EMIS Web systems already store all data on EMIS central servers, not at the GP practice). You might want your data stored securely on a server in Switzerland, or you may be happy to have the NHS look after it for free.

'Anonymised' data for research etc

Anonymised data can have great value in medical research. Unfortunately the NHS undermined confidence with its inept 'care.data' scheme, which planned to harvest patient data wholesale from GP records, and made opting-out difficult. This worried patients and professionals alike and was scrapped in 2016. (See Wired story at <http://www.wired.co.uk/article/care-data-nhs-england-closed>.) Prior to that, fewer than one percent objected to sharing health data – now it's around 8%. Yet many would be happy to donate data on a voluntary basis, just as people donate blood.

HES, the Hospital Episode Statistics, is a data warehouse of all admissions, A&E attendances and outpatient appointments at hospitals. Its primary function is to ensure that hospitals are paid for the work they do, but it is also made available for research purposes, and for planning healthcare provision.

Having clinical and personal data available in a computable format could be massively life-saving and health-enhancing – for example, in risk stratification. Risk stratification calculations have been reliant on HES data, but that data is limited to episodes. GP data is the closest one can get to a holistic picture of a person's health, and it's all computerised – and coded. If that data could be mined, better predictive algorithms could be created.

One of the concerns, however, is that you can't really anonymise a rich healthcare record, and many would be concerned about that. Perhaps by means of privacy and governance around datasets, the risks could be mitigated – and people be encouraged to donate.

The road to open platforms

Ewan described the journey towards Openness in the NHS. When Tim Kelsey was NHS England's National Director for Patients and Information, he went to the USA and saw **VistA**, the open-source health records and business information management system of the Veterans Health Administration (VHA). VHA is the largest integrated healthcare delivery system in the USA, caring for 8 million military veterans. VistA had played a key role in turning around what had been for years a badly-run, failing agency, now highly regarded.

On his return, Kelsey promoted the idea of the NHS adopting the VistA software, but it was soon apparent that the cost of localising it would be huge. The initiative changed into the Open Source Programme; later, that was thought too restrictive a brief, and the project morphed again towards Open Platforms.

The NHS Open Source Foundation was funded with a grant from NHS England, and later took the name of Apperta, which published the pamphlet Ewan showed us. Code4Health built an example open platform, to allow people to experiment with it. All this was done with half a million pounds. Compare that to £200 million for the Epic deployment in Cambridge, or the £700 million compensation paid to Fujitsu when NPfIT failed – just think what could be done with a serious budget!

The safety benefits of electronic prescribing

We also discussed how electronic prescribing could help avoid accidents. If people are given the wrong drug, or the wrong dose, or in the wrong way, or not given the right drug – these can all cause serious harm.

A 2001 American study by the Institute of Medicine called ‘Crossing the Quality Chasm: a new health system for the 21st century’ looked at the impact of medication errors on hospital admissions, deaths and overstay. In the same year the Audit Commission for Local Authorities and the NHS in England and Wales released its own report, ‘A Spoonful of Sugar: Medicines Management in NHS Hospitals’. (One figure put forward: medication errors were costing the NHS half a billion pounds a year in longer stays in hospital.)

Moving to electronic management of prescribing would help. Computers are good at spotting that a nominated dose falls outside the range of what makes sense. The system can also test for drug interactions, contra-indications, and cross-sensitivities – if given the data to enable these checks.

But using computers can also introduce errors. Where drugs have similar names, a doctor would be unlikely to write the wrong one – but if working in haste with an on-screen pick-list, it would be easy to click on the wrong one. But if the system is told which medications are commonly mis-prescribed, a further ‘are you sure?’ layer could be inserted in the process.

Electronic prescribing is universal in GP practice: it is built into GP systems.

In hospitals, it is not. Forty percent of hospital encounters include some level of error in prescribing, so there is huge scope for improvement. (Note we don’t mean the electronic transmission of prescriptions to pharmacies, but the preparation of the prescription request, including when it is printed onto a form and given to the patient.)

But in hospital informatics, implementing e-prescribing is not easy to do, because to check for contra-indications, the system needs access to data about a patient’s other conditions, and lab results – such as liver function, which affects the rate of drug metabolism. You need a whole joined-up electronic health record to do e-prescribing properly.

Starting the journey

In concluding, Ewan says that he is sure that the open platform approach is right. He has no problem ‘selling’ the idea inside the health service: it makes such sense. The problem is getting people to start on that journey, when for a large hospital trust at the moment the safe answer is almost certainly to give Epic some hundreds of millions of pounds. In fact, what Ewan said in recent consultancy was: look, go and procure a megasuite system from one of the three vendors who can meet your needs – but, in the negotiation, get the vendor to commit to progressively opening up the data. Because that’s the moment when you, as commissioner, have some leverage over the vendor.

– Conrad Taylor, June 2018