Experiences from the Use Cases in SemanticHealthNet

Ian McNicoll, Charlie McCay
Tony Austin, Shanghua Sun,
Gerard Freriks, James Cunningham

Experts involved in the use cases of SHN WP1 (heart failure) and WP2 (public health)

Annex 7
of SHN Work Package 3
Deliverable D3.3

Final version, May 29, 2015

Document description

Deliverable:	Annex 7 to SHN WP3 Deliverable D3.3
Publishable summary:	In this annex, we collected the experiences of key players in the use cases of the SHN project. The first Use Case was focused on the development of a clinical model for the outpatient letter for heart failure. The second Use Case dealt with the secondary use of clinical data for public health research in cardiovascular disease.
Status:	Final draft
Version:	1.5
Public:	x Yes
Deadline:	May 29, 2015
Contact:	Ian McNicoll ian@freshehr.com, Charlie McCay charlie@ramseysystems.co.uk ; Tony Austin tonyaustin@heliconhealth.co.uk ; Shanghua Sun postmaster.electronsea@gmail.com ; Gerard Freriks postmaster.electronsea@gmail.com ;, James Cunningham j.a.cunningham@gmail.com
Editors:	Robert Vander Stichele robert.vanderstichele@ugent.be

Table of content

Introduction 2

Reflexions of a clinical modeller 4

Experiences from mapping HL7 to the Heart Failure Patient Summary Requirements....7

Experiences of interacting with clinicians to finalize clinical specifications of the Heart Failure Patient Summary 9

Experiences of converting clinical requirements of a Heart Failure Patient Summary into a functional application 12

Experiences of a representative of EN13606 Association : From Semantic Interoperability to Semantic Interpretability 16

Experiences of a public health researcher....20

Note:
The experts involved in the use cases of SHN were commissioned by Prof. Dr. R. Vander Stichele, Workpackage Leader of SemanticHealth Net WP3, to write up their experiences.

Introduction

In this annex, we collected the experiences of key players in the use cases of the SHN project. The first Use Case was focused on the development of a clinical model for the outpatient letter for heart failure.
The second Use Case dealt with the secondary use of clinical data for public health research in cardiovascular disease.

Reflexions of a clinical modeller

Ian McNicoll
Co-chair OpenEHR

The key challenge at the heart of eHealth interoperability might be seen as trying to maximise the value of information captured as part of routine care in one part of the health and social care system, so that it can be re-used and re-purposed in other parts of the health system.

Where information can be captured and re-used in coded and structured formats, the processability and re-usability should be maximised to the benefit of direct care, quality improvement and secondary analysis and research. This is the hope and promise of semantic interoperability but a goal which has remained elusive over many years of considerable research effort.

In truth, most health data, even when highly structured and coded, remains highly siloed and difficult to re-use outside of its original context of collection.

In the SHN Heart failure example one might imagine that information collected and coded by a GP surgery might be easily re-used in the context of a Heart Failure Nurse-led clinic, re-purposed again by a consultant review clinic, through into quality register and research secondary uses. In practice, the patterns of information collected, seemingly for the same clinical requirement are often orthogonal, and rarely transferrable without complex and expensive transformation exercises.

Efforts to resolve this issue seem largely to regard this problem as an issue of logics and engineering. The challenge is seen to be one of technical and linguistic differences between the different stakeholders. Semantic disjoints are due to the artificial constraints applied by differing technical solutions or modelling formalisms. These differences can be resolved by the use of more sophisticated and logically correct formalisms after increasingly detailed ontological analysis of the problem space. This is a scientific endeavour where more analysis and better understanding of the problem space will allow a largely ‘hands-free’ translation between differing representations of the same clinical concept, the mythical ‘Babel-fish’ of Douglas Adams.

Underlying this approach is an assumption that the differences between these siloed representations of clinical concepts are simply artefacts of the technology used i.e. that these different clinical groups are trying to ‘say the same thing’ but that technology (and human linguistics) is getting in the way.

The possibility that these different clinical groups actually have real, though subtle, differences in their information needs and uses seems rarely to be considered. In other words, part of the problem may be that different clinical groups are not actually ‘trying to

say the same thing', or at least, require to re-process similar information to maximise its value within their working environment.

These inconsistencies in the use of clinical information are often not apparent to clinicians themselves, or indeed to their system developers, both of whom are highly focused on maximising the fit of information structures to local requirements. This goal often conflicts directly with an expressed desire to interoperate.

The EHR seems often to be thought of a record of physical phenomena experienced by the patient when in truth it is a set of documents used by care professionals to inform and guide the care of the patient. This will, of course, includes records of various physical phenomena but always couched in a manner which best supports those professionals who author the record. The EHR is a document of diverse working practices centred around a specific patient, not per-se a biological record of the patient themselves. The 'siloed' nature of clinical documentation is therefore inherent and often 'by-design'.

Lack of interoperability is not an artefact of the technology used to collect or transmit the information, and indeed it is likely that as the potential for information exchange improves, controls will need to be put in place to prevent local data quality being compromised by the automatic import of external data which is not locally fit-for-purpose.

We therefore have a clinical community used to expressing requirements for local information capture and use, but when those requirements expand to include the need to share patient information with other clinical groups, there is no capacity to allow conversations directly between those groups to develop a shared understanding of their disparate information needs and constraints. Such discourse is typically mediated through slow and unrepresentative 'standards' authorities, or expressed in complex and arcane technological formalisms. Interoperability therefore comes to be seen as an exercise in technical plumbing, rather than the development of negotiated, shared clinical practice.

The openEHR archetype-based methodology is designed to expose aspects of clinical information models in a way that allows ordinary clinicians to directly participate in such discourse. It allows clinicians to directly explore the potential for shared information use and understand the limitations. It gives clinicians the power to directly influence the development of shared information models but also the responsibility to define when the effective limits of such sharing have been reached i.e the points at which the need for siloed or localised models outweighs any benefit of shared models.

One of the frustrations of the SemanticHealthNet project was that while openEHR tools and methods were used to create candidate semantic artefacts, the opportunity was lost to engage the various clinical groups involved in gaining a better understanding of how their information uses and needs differ, and of how negotiation and consensus building can lead to useful interoperability as a pragmatic and practical exercise.

My experience of building a number of openEHR-based technical artefacts, representing different stages in the Heart Failure patient information journey, is that a high degree of re-use of basic clinical information components is possible, as long as clinicians and local

system developers are made aware of some of the difficult issues that are not easily resolvable, or of the local assumptions that may cause confusion. A good example is that while the concept of 'primary diagnosis' has clear utility within the context of condition or episode-based management, it makes no sense within a longitudinal care record such as a GP record. This does not mean that such information cannot be transferred between disparate care settings, just that it needs understanding and careful handling.

There will, of course, always be a need for condition-specific models e.g. NYA Heart Failure Score but these are not necessarily of sole interest to secondary / specialist services, as more traditionally specialist monitoring is carried out in the community or by patient self-assessment.

For me it was unfortunate that SemanticHealthNet seemed quickly to become dominated by efforts to promote a highly theoretical approach based on arcane ontological principles that were impenetrable to the majority of participants, including those like myself who are well-versed in the complexity of health informatics.

It was certainly completely opaque to contributing clinicians, who already seemed confused by the goals of the project and the process seemed to relegate other technical participants to the role of providing test material for what had become the 'primary thesis' - the ontological approach.

The opportunity was lost to allow these other groups to explore how a best-of-breed approach to a combination of these technologies might allow real, immediate and practical progress to be made and to identify gaps in provision.

With regard to Terminology binding, the openEHR-based Heart Failure attempted to maximise the use of SNOMED CT where appropriate e.g to provide values for diagnoses, procedures or lab investigation names. We deliberately did not attempt to provide bindings or mappings for other node names or values where the relative semantic value is low or coverage in SNOMED CT is poor. In short, make use of SNOMED CT where it is strong and the value is high but recognise the inherent limitations and challenges of aligning the SNOMED CT concept models with structural models such as openEHR.

This limited approach tends to cause much frustration in the SNOMED CT community who regard it as under-using the potential of SCT-based inferencing but it reflects both the relative immaturity of significant aspects of SNOMED CT, and more significantly, the limited capacity of the majority of developers to harness this potential within current system architectures. It is interesting that the HL7 FHIR community have taken a largely identical approach to the use of SNOMED CT in their modelling efforts.

While openEHR does allow multiple terms and terminologies to be mapped to specific openEHR nodes, this was not performed as part of the SHN project, the emphasis being on use of terminology as nomenclature, rather than to provide classification.

Experiences from mapping HL7 to the Heart Failure Patient Summary Requirements

Charles McCay
HL7 Technical Lead for SemanticHealthNet

The requirements were initially provided as a patient story from which an example CDA document was created, based on the epSOS CDA Patient Summary templates. This example driven approach was very quick and effective in identifying issues that would need to be addressed for a complete specification to be created. The same story was also expressed using the openEHR, 13606 and ontological modeling frameworks

The objective for the HL7 CDA modeling work was to use an existing specification methodology that would scale. After evaluated various CDA template authoring methods, the pivot table approach used in the epSOS specification was selected to provide a definition of the set of constraints on the information items to be included in the summary. The rationale was that no special tools were needed to extend and repurpose the epSOS pivot tables, and also there was a substantial pan-european investment in materials to support implementation using the epSOS Patient Summary.

The pivot tables made it easy for someone familiar with CDA to see how the information items to be included in the patient summary could be represented by using and extending either the epSOS CDA patient summary, or the Consolidated CDA patient summary. It also facilitated the use of a common collection of SNOMED-CT value sets, and a basis for mapping into the openEHR representation.

The limited testing that was done pointed to the possibility of a set of patterns being defined that would allow information items needed for the patient summary to be represented at the same level of granularity across the different information models, and in the set of ontological patterns that were the focus of the technical development work in the project.

Indeed it was apparent that differing assumptions and understanding of the information requirements generally caused the variations between examples and authored constraints, rather than these differences being a consequence of the modeling approach used.

The ontology, openEHR, 13606, and HL7 CDA representations could be could each be constrained to provide ‘slots’ for the same information items, and hence to support bi-directional conversions, with two significant limitations:

1. While this appeared practical for the set of sections in a patient summary, and for specific entry-level information items such as smoking status, the equivalence did not extend to every possible representation in each modeling paradigm. Thus a specific set of equivalent models would need to be authored and validated in each paradigm in the context of the use case.

2. All of the modeling frameworks seemed relatively complex and arcane to those experts who had not arrived at the project with prior knowledge of the framework. This perception persisted throughout the project despite email and face-to-face collaboration.

Each modeling approach has a technical communicating community around it. This is the group the individuals who are comfortable using the modeling approach, and discussing issues using "their" models. Moving to understanding and using another modeling paradigm required significant effort and motivation. This motivation was provided to some extent by the collaborative nature of the SemanticHealthNet Network of Excellence, however it seems improbable that this would be replicated at scale outside of the project unless the effort required is significantly reduced.

While the experts did not "go native" and seek to use the alternate representations, producing mappings in the context of a well-defined use case appeared possible. The use of "slots" and examples, and the knowledge that the information items were being used for the same clinical purpose were important factors in reducing the anticipated issues with creating these mappings. The crucial factor here was that the information was being mapped between structures that were defined to achieve the same purpose for the same clinical communicating community.

The detailed work on modeling smoking related information showed how there were different clinical communicating communities who were collaborating for specific reasons, and each communicating community had subtly different and incompatible information needs that could not be bridged by logical or semantic transformations. Getting a clear enough understanding of the respective requirements to establish this incompatibility with respect to information about a common subject took time and effort.

The way that clinical and technical communicating communities interact and intersect at a point in time and over time would be a fruitful area to explore further. This would help to establish the extent to which there is a real need for interoperable sets of models to be maintained. For while work within SemanticHealthNet indicated that specifications and mappings could be created and maintained, it has not answered the question as to when and for how long such efforts deliver sufficient value to justify the cost.

Given the importance of interoperability and the appropriate sharing of information to the healthcare community, there is also a need to reduce the costs of becoming an effective member of one or more of the technical communication communities. As well as reducing the barriers associated with using the associated modeling paradigm, this will reduce the effort needed to build bridges into and out of the paradigm. HL7 has been focusing substantial effort on the FHIR suite of standards that have ease of implementation and adoption as a driving design goal.

Experiences of interacting with clinicians to finalize clinical specifications of the Heart Failure Patient Summary

Tony Austin
Imperical College London

Our team was commissioned to deliver a rendering of the SemanticHealthNet (SHN) deliverable 1.2 (for heart failure) as a running application. The team possesses mature tools for the modelling of records in a clinical domain and the turnkey emission of user application equivalents of these models.

At the outset to the development project we were in possession of the finalised written artefact (the deliverable) that we could operationalise. Unfortunately this is often not the case. Past experience suggests that just the generation of this initial document is hard because it requires up-front engagement by the user community — they are the domain experts after all.

It is often easier to build from scratch than it is to build over an earlier attempt. This is an engineering principle rather than something unique to software or to healthcare. Think for example of laying high-speed rail, vs. updating existing standard rail to high-speed rail. As well as the written deliverable we were also presented with a technical artefact created using a different representation formalism. However, it is actually very hard to convert one technical format into another. Representations of data types, depth of aggregate container nesting, different fundamental approaches to representation such as relational vs. object-oriented, all sum to a significant scale. SemanticHealthNet itself planned a "round-trip" test where candidate data was converted to and from three pre-selected formalisms. Of course, this is exactly the same problem as using other technical formalisms as the basis for representation, which we would then have had to solve by ourselves. Assuming that all the chosen formalisms are comprehensive representations (which is by no means guaranteed), that change is possible, but it remains an effort. We have considerable experience at preparing technical artefacts within the constraints of our own formalism. We chose therefore to return to the ground truth clinical statement of the requirement and avoid the overhead of parsing a different artefact with different potential divergences from it.

The visualisation suggested different things to different participants: an Electronic Healthcare Record (EHR) for Heart Failure specialists; an outpatient letter; even a minimal data set for an international registry of heart failure patients. However, in general it's hard to scope a development at the outset. This again isn't limited to healthcare by any means, it applies to software engineering generally and is one of the reasons why an iterative model of agile development is increasingly seen as more realistic than a "waterfall" model. During the period in which the heart failure application was developed, the scope expanded from a simple programmatic representation of a deliverable in operational form, through to an EHR for heart failure specialists that includes the export of an audit dataset that might populate a physical letter as well. Only in specifically SHN project meetings was the idea of a registry communicated. It's possible that this doesn't excite clinicians quite as much as an operational tool for their daily practice, but more likely that it would indeed have been reached but at a moment now after the project's end.

Personally, I am confused by the difference in visualisation potential between a picture of a screen and a screen that you can type into. They seem equivalent to me. Consequently I don't personally believe that a real application is necessary to visualise the heart failure summary. Nevertheless it can't be denied that a real application engages clinicians in a way that pictures do not. Perhaps something in the medical experience encourages clinicians to strive towards a target and the belief that feedback will lead to a better application amounts to a target worthy of their time and energy. So I must finally conclude that to really engage healthcare professionals they have to imagine themselves using the written artefact in an operational way. Since our toolset presents them at every iteration with a real application they could take away and use immediately, it did indeed have the motivating effect.

Once in progress, discussions between clinicians and software engineers are sometimes "robust" (fortunately to the later amusement of both groups) but they are usually conducted amicably. They are somewhat systematic. A clinician spends some time alone with the application and then while issues are still fresh in the mind meets the development team. The application is studied together screen by screen before finally new functionality is proposed and discussed. Shortly following the meeting the rough notes are converted into specific actions on all parties by the senior developer which is then shared widely to avoid mistakes.

The number of iterations needed is relatively few. By the conclusion of SemanticHealthNet there will only have been four for this activity and every stage has tended more obviously towards an approved dataset. Undoubtedly the biggest issue we have faced is the amount of wait between iterations. The difficulty of marshalling the time of very senior clinicians is easily the biggest brake on progress towards final ratification. A strong recommendation before embarking on such a project is to impress upon the commissioning party the importance of having enough protected time to engage personally and stay engaged throughout.

Of lesser difficulty but still notable are the different approaches taken by clinicians when visualising the same problem. This has been less of an issue here where discussion has mostly centred around the relative importance of particular values — sometimes the benefit of one diagnostic value over another, or the real likelihood of obtaining a value from a specific test. However, on previous occasions clinicians have clashed more fundamentally. On one project a healthcare professional routinely obtained test results in a document which was sufficient for him. He therefore questioned why detailed data capture was needed when he could post a test result into a text box on the screen (effectively creating a convenient document repository for himself). Others in the group pointed out the potential value of analysing more differentiated results and the use in audit. The discussion circuited for some time before both possibilities were catered for in the data capture, thereby serving neither party.

As non-clinicians we are not really in a position to declare what is and is not important in the proposed dataset and instead parse input in a robot-like fashion. For the purposes of shared care we find it instructive to note how multiple teams working to different briefs describe the same clinical concept differently. Although this project proposes a heart failure summary and would therefore expect to omit certain aspects of, for example an echocardiogram, other teams with whom we have worked have a much more expansive set. Since we would like to encourage shared care we would prefer everyone to use the same definition (even if in some

cases not all the data is provided). Having two such definitions means that a conversion between them is needed if either is to serve the other's purpose.

Experience also flags domain modelling that looks likely to be the result of one clinician's or one group's "pet hates". In contrast, fussiness resulting in careful modelling is to be encouraged!

Not for any project in all our twenty-plus years of domain modelling have we ever been asked to include terminology bindings except by vendors of terminologies. This is partly because the international standard on which our work is based includes the exposition of a knowledge component of its own, that enables us to precisely model the structure of a healthcare record. The two technologies are potentially complimentary in the sense that record structure doesn't strive to include comprehensive taxonomic or mereologic relationships forming a medical knowledge base (like "JVP is a pressure" or "the toe is part of the foot") whereas terminologies don't tend to describe best practice in recording (like "smoking history should be established in a first prenatal visit"). From our experience it is the record structure statement whose value is most keenly felt by clinicians and knowledge acquisition professionals.

If a systematic evaluation of similar use cases for clinical model development in other disciplines were to be set up, it would be sensible to ask of each project:

1. Did it agree a publicly accessible structured documentary artefact for the data capture?
a. Did the process conclude (that is, did it finish with what the participants believed to be a definitive outcome)?

2. Did this result in publicly accessible technical artefacts for the data capture?
a. How readily usable are the artefacts (for example, do they depend exclusively on a specific underlying model)?
b. Have the artefacts been shown to be consistent (for example, have they been used in an application)?

3. How engaged has the community become (for example, is there evidence of widespread agreement, endorsement by professional societies, or commercialised implementations)?

The application resulting from this work has continued to expand, incorporating new requests and new scope. To have engaged the user community to such an extent it must have fired the imaginations and the expectations of those who have seen it. As a developer, this is the reward we always hope for!

Experiences of converting clinical requirements of a Heart Failure Patient Summary into a functional application

Shanghua Sun
Imperical College London

I was part of the team tasked with implementing a visualisation of the heart failure summary specification (SemanticHealthNet deliverable 1.2) as a functioning application. Specifically, my role was to convert the clinical requirements into a knowledge model specification from which the application could be derived. To do this, I used a representation formalism known as the "Pattern". Patterns are used in regular expressions to mean the constraints on an input that define those which are acceptable and the word Pattern is reused in this formalism to mean the same for instances of clinical data. The representation borrows heavily from class annotations used in programming languages. An example from the heart failure summary might be: http://aruchi.chime.ucl.ac.uk/pattern/describe?id=432

As an implementer I am tasked with implementing what I am presented with and the clinical correctness or scope isn't generally mine to resolve. The deliverable was well written, clear, seemingly comprehensive and computationally logical. However, on some occasions the document was not structured in a way that was directly representable using our formalism and in those cases I guessed at the real underlying need. There were also occasions where I needed to look up additional information that was not needed at the user requirements level, for example measurement units used in a test, or the expansion of an abbreviation.

All Patterns must carry a description and clinicians very familiar with the meaning and use of a field do not pass this information along in a way that enables me to teach medicine to a computer! This is the gap of understanding from a clinician's point of view to that of a software engineer. Over time of course a body of code tends to grow and lend understanding to that which comes later. Fortunately we have developed other applications with similar data items and the formalism is deliberately designed to enable sharing between domain models to occur. Only for values that we had not come across before did I need to check the meaning for this use case.

The overall procedure is as follows:

1. Trim the original document to its hierarchical data values and containers, including any descriptive information for those such as the measurement units;

2. Translate it into a documentary technical plan, and thence into computational artefacts using the Pattern approach;

3. Finally prepare and deploy the application.

Step 2 is the task for which I was chiefly responsible. I try to approach the definition of the computational artefacts keeping the following steps in mind:

2a. First, rewrite the requirement in terms of what can be displayed

The range of things that can be displayed in an application is deliberately few. This keeps application behaviours consistent and intuitive, and prevents users getting "lost" in deep record structure hierarchies. However, this does limit the full expressiveness with which requirements are sometimes presented and requires me to refactor them into simpler alternatives.

2b. Save time and promote sharing by reusing existing definitions

I am looking as much as possible to facilitate reuse of definitions to encourage shared care. For example, if I'm to model an Echocardiogram I would like the model to represent a definitive statement,

or as near as I can make it at the time (since clinical knowledge is always evolving). I want to know what an Echocardiogram *is*, so that everyone who uses one has the same model. Where different types of Echocardiogram specification exist, I must manage the definitions separately and any use of the data for record instances must be converted.

2c. Rewrite the requirement to facilitate reuse

Where there is no suitable model to reuse I will try to make the one I build generic so that it can be reused later. For example, if asked to create a field for "Assessment of Systolic Function" with values NORMAL, MILD_IMPAIRMENT, MODERATE_IMPAIRMENT and SEVERE_IMPAIRMENT, I would rather create a Pattern for "Impairment" that would be useful elsewhere. There is nothing Systolic about the values given and the label only need be changed in the container. We consider this different to a sub-Pattern where a more restricted set of non-curatorship constraints would apply (for example, without NORMAL).

2d. Repeat

After using our in-house tool set to create the Pattern resulting from the above, I then continue for each field value the same way. With the elemental data values completed I then look for appropriate groupings that will look attractive and clinically sensible on screen. I try to avoid including fields that look unlikely to survive appraisal by a broader clinical team. If a field needs to be removed, deploying the database for the application becomes much more difficult as the data has to be migrated. In contrast, adding new fields that turn out to be needed after all is much easier.

In SemanticHealthNet these rounds of cyclic revision and feedback will have reached the 4th iteration by the close and this is consistent with our previous clinical system development experience in other domains. The methodology of communication worked nicely. Once we got the chance to communicate with clinicians and other developers, it was very efficient and productive. The only problem is the “busy” nature of healthcare professionals, so waiting for a communication opportunity did involve a very long period of time.

Meetings benefit from including a broad range of skills. We have in contrast certainly known projects where a very engaged commissioning person has driven the requirement from the clinical side without including representatives from the team destined to use the system, and unsurprisingly it turns out that the system doesn't meet their needs. However the technical and knowledge acquisition teams also need to marshal different skills and personalities. The sparking brilliance of senior clinicians has to be funnelled into a logical flow that if successful, will have an explanatory tutorial quality about it — the unstated feature of passing on their experientially-gained but apparently diffuse approach to other practitioners that do not have their hard-won background. At the same time, the technical team needs to defend the logical flow against clinical “requirements” that actually reduce rather than enhance it. For example, repeating identical values on multiple screens, or using headings to group values that have no explicable unity. Still, although I am not a domain expert, this use case does indicate that with a good routine of development and feedback, together a healthcare professional and technical team can achieve a high level of granularity and precision in the end result application.

SemanticHealthNet had what we understand is an unusual provision for "expert advice" that enabled funding to be secured for us to perform this exercise. Even so, the level of funding was strongly capped in terms of time taken rather than results. One recommendation for future projects of this sort is to ensure that similar or better funding provision is secured. Costing this is of course a non-trivial task, since projects are often planned financially at the outset before any documentary statement of the complexity of the models to be created and supported exists. However it’s clear that clinicians need to see how a system looks in order to validate their data capture requirements. Also, some financial support for initial commercial deployment beyond the end of the research-led part of a project is also worth considering. Arguably, there is a commercial disincentive for application vendors to share data freely where once access might have been revenue-generating and the least amount of effort required to make this marketable is obviously the better.

Artefacts that document clinical models must be made public in order to enable discussion and to increase buy-in, at a national level and from professional societies, but especially from vendors who are the real users of the models and whose data is to be shared. It takes a great deal of time and

effort to create and iterate a clinical model and it is probably inappropriate to expect vendors to perform this function. They might choose to recoup their costs by restricting access to the end result. A publicly-funded body such as a university is better placed to commission clinical models and then apply their intellectual property to facilitate broad use.

Overall, this use case was a very useful exercise and it has become a valuable pilot application. There is still a lot of future work on the application that can potentially be done if further funding and support can be secured. A pilot application like that of the use case described here can certainly be extended to a fully working system used in clinics. It is desirable to create such applications for use across nation-states as the only viable way of ensuring transnational validity of the associated model.

Experiences of a representative of EN13606 Association : From Semantic Interoperability to Semantic Interpretability

Gerard Freriks

EN 13606 Association

In the context of the Semantic Health Net an position paper[1] was written. This is an introduction to the Annex that explains what is needed for the next level of semantic interoperability named: semantic interpretability. What is needed mimics the way humans exchange data using syntax, words from shared dictionaries, shared encyclopedias and shared phrases.

Common problems:

- EHR-systems are not really interoperable.
- It takes too much time create a patient safe exchange.
- And when the exchange is realised new, adaptations to the systems are resource intensive (time and money).
- Data inside EHR-systems can not or only partially be made available to the HcProvider and its organisation.
- Re-use of data for reporting and research is complex. Too much implicit human knowledge is necessary to interpret the data fully. The 'idea' that is in the mind of the sender needs to be made explicit in order to create the same 'idea' in the mind of the receiver.

Semantic Interpretability

At this level of interoperability almost all contextual data is captured and made available for analysis. In present Semantic Interoperability the complete context is not captured. Communicating parties need to interpret the data received for which they need a lot of implicit knowledge. In Semantic Interpretability all implicit knowledge can be exchanged as well.

The next set of questions makes clear why attention for the data and its full context is important.

- Is a diagnosis used in a statistical report the same as a new inference documented in the EHR as diagnosis?

[1] Annex8 of SHN D3.3

- Is a diagnosis documented by a trainee the same as one provided by the patient or one provided by a specialist?
- Is a diagnosis of the same value when documented 50 years ago compared with one documented yesterday
- Is a body temperature measurement interpreted the same when it was measured: after a sprint of 100 meters, or in an high or low ambient temperature, using calibrated or non-calibrated method, etc?
- Is a measurement of the same value when it is labelled as uncertain or labelled as certain?

Semantic Interoperability and the Messaging paradigm

Semantic Interoperability for many is the ability to transport data from one data base to an other database. Messages are used as intermediate format. Software is used to retrieve data from the sending database and is transformed into the message specification. Subsequently the receiving system translates the data from the message into the local database.

Sender and receivers systems have to develop software. Both sender and receiver have to interpret the message specification and their local database schema.

Because a lot of the context is not transmitted via the message humans with shared implicit knowledge are necessary to interpret the data safely and a documented profile in an implementation guide is necessary. It is a resource intensive process (time and money) that needs the Integrating the Healthcare Enterprise (IHE) profiles and Connectathons.

In addition this method needs a lot of human implicit knowledge and is not very flexible, or agile.

Semantic Interoperability and the Two Level Modeling paradigm

Two models define on one hand a Reference Model that can hold any data in an EHR and, on the other hand, allows flexibly the extension of that Reference Model such that each artefact (archetype) created defines what can be documented about a topic. (e.g. demographics, or diagnosis, or finding, etc.) These archetypes are defined using the Archetype Object Model.

With this system EHR-Extracts (messages) can flexibly be designed and used. When EHR-systems know how to deal with the Reference Model, and the Archetype Object Model, EHR-systems can exchange data easily and flexibly in an agile way.

Since each user or user group can define in its own way the EHR_Extract is (and its archetypes are) using a lot of implicit knowledge that is needed to interpret the data in those EHR-systems or EHR-Extracts.

Semantic Interpretability

The next level of sophistication of Semantic Interoperability is Semantic Interpretability.

When data is normalised such that it can be interpreted safely and fully it means that the context of the data must be as complete as possible. The goal is to create a set of methods and use standards such

that a person yet to be born, using a system, yet to be designed, is capable of understanding fully the data in its full context.

Data in context

In order to record data in its full context the next set of methods must be defined and used:

- agreement on modeling styles
- agreement on the use of codes from reference coding systems
- ability to capture data in its full context (what, where, why, who, and how)
- ability to capture the presence and non-presence, or certainty or uncertainty
- part of the context are the processes the documented data is used in
- all archetypes must carry in an explicit way the meaning of the nodes in the archetype
- resulting in a set of shared Reference Archetypes that are deployed to store and retrieve and exchange data

Modeling styles secure that the same thing is modeled the same.

Codes can be used in many ways in archetypes. E.g. '*Fracture in left femur*' or problem= '*fracture*', Location= '*Femur*', Laterality= '*left*'. Both versions mean the same. In order to reduce misunderstandings and increase patient safety one way of using codes need to be chosen.

Data (e.g. Diagnosis = Diabetes) depends on its correct interpretation on many aspects: when was it diagnosed, on what evidence, by whom, when and where plus what method was used and what confounding factors were present? The clinical weight depends on these context factors. E.g. a diagnosis by a patient is not the same as when a specialist is the author, or a blood glucose test is interpreted differently depending on the kind of test or whether the measurement was after fasting or immediately post prandial or after a fixed amount of time after a fixed challenge.

Processes are part of the context

We must be aware that data is used in processes. When data is observed in the patient system and documented in an EHR it is different from the same Observation that is re-used in the ordering of an administrative process (reporting for instance). All kinds of organisational processes and clinical diagnostic/treatment cycles influence the interpretation of data in an EHR-system.

Ubiquitously sharing health data

The more healthcare actors (patients, GP's. specialists, nurses, devices and services) are cooperating they need to share and exchange data plus the ideas they represent.

In addition re-using data for research and reporting necessitates that data is stored, queried and exchanged with as much as possible implicit data made explicit. The moment all health data is available in this normalised way constitutes a ubiquitous way a data platform for eHealth and mobile health applications can make use of.

Gerard Freriks
EN 13606 Association

Experiences of a public health researcher

James Cunningham
University of Manchester, UK

Deliverables 2.1 and 2.2 outlined a series of vignettes based around the example of Cardiovascular disease illustrating the challenges relating to summarising healthcare records relating to CVD on a population level and looked at the informatics required at to address these population level challenges. It has been noted that in these cases CVD serves very much as an exemplar and that both the lessons learnt and the challenges outlined in these deliverables can be mapped to other disease areas and healthcare in general. In a sense we can view the message expressed in these deliverables as a vision of how we would want healthcare systems positioned in future so that they would be best placed to address the needs of population-level care and research.

Whilst outlining this vision of where we would like to be in terms of semantically rich and harmonised healthcare systems and approaches the deliverable in a sense begged the question as to what we need to do to get ourselves into this situation. The vision itself does not give us a map of how to get there. This note outlines the two endpoints of the path that we want to be on -- on the one side the current state of EU wide healthcare systems and operations in terms of their applicability to semantically rich and harmonised research and practice and on the other the ideal state of such systems. Further we then expand on a use case and methodology that will give us a potential insight into the path that we can take to reach our goal -- the eLab, a system (or in the general case a type of research environment) that takes un-harmonised silos of data and the research and care that exist at present and draw together an ecosystem of people, methods and data that can enable the type of scenarios we saw in D2.2. The notion of eLabs have been used as the underlying informatics solution across a range of population-level healthcare research projects of the sort that match to the scenarios addressed in the workpackage 2 deliverables. The eLab example comes from within the UK research community, but the lessons learnt and the insight into an approach to enabling the 'de-silo'ing' of data can be applied and scaled across the EU healthcare landscape.

The issue of obtaining semantic interoperability often refers to the (highly non-trivial) aim of relating items of coded data to each other in a semantically rich and meaningful way. One of the points that we can draw from Deliverable 2.2 is that the interoperation and integration of healthcare systems, particularly as they relate to enabling they key factors of population-level healthcare research, should encompass not only the semantic interoperability of healthcare data, but should also provide an ecosystem that encompasses the people (healthcare practitioners, researchers and patients) and the methodologies (the approaches, formal or otherwise, to enacting research) as well as the data and provides interoperability between these entities in a broader sense. A healthcare system that could act as the enabler of such an ecosystem would provide a basis for enacting population-level care in a way that would address the types of challenge outlined in the vignettes of the workpackage 2 deliverables.

The CVD scenarios that were looked at as a part of workpackage 2 looked a range of outcomes related to tackling the problem of Cardiovascular disease that covered broader

aspects of healthcare than the immediate treatment of a given patient. These include the development of social interventions, aspects of clinical audit and commissioning, consumer healthcare applications as well as research using the EHR. It is clear that purely harmonising data alone, whilst perhaps a pre-requisite, would not in its own right enable such goals. It is the existence of a wider environment, one that enables the interoperability of not just items of data, but also enables the social aspects of communication between the practitioners of healthcare and healthcare research and the more technical aspects of coordinating the methods applied to the data, that is needed to address these goals.
This then gives us a broad idea of where we want to be, that is, in a situation where the interoperability between healthcare systems extends not only to the data being recorded and used, but to people and methods surrounding the use of that data. Deliverable 2.2 developed in depth some of the ideas that have driven this vision. One of the key conclusions of D2.2 was summarized as follows:

- Many of the technologies required to deliver the systems outlined in this report are already in place. They may even be in widespread use, however take up is often piecemeal, driven by forces outside direct healthcare and often they do not provide a complete, end-to-end solution.

It is not primarily the technical challenges that stand in the way of us getting to where we want to be on the population-level, rather the coherent adoption of these solutions and it could be added a willingness or agreement amongst key stakeholders to push in the same directions. D2.2 then makes explicit one of the aspects that is holding us back from integrated healthcare environments, namely human factors. Again from D2.2:

- Human factors, including motivation to change, may be more limiting to the provision of semantically interoperable population summaries of EHR data than are technical factors.

What we have is then is a reticence to change amongst those actors in the healthcare environment that is hindering the adoption of (often pre-existing) technological systems. If we can foster an environment that brings these actors together, and encourages a socially agreed upon push to change, then we will be on the path towards our vision of an environment in which informatics solutions can be deployed to address population level health care challenges. Another point brought forward in D2.2 concerned the need for a coherent economic case for change to be made:

- The “business cases” that encourage people and organisations to change existing ways of working can be viable on large and small scale however it is important for those business cases to be formally developed in order to provide the evidence.

One of the key aspects here is that it is the provision of evidence that is needed to develop these business cases. At the moment there is a lack of evidence that would drive change - The evidence for backing up such economic arguments necessarily comes from a range of different stakeholder and cannot be made purely from the data alone. Once again it is an

environment that enables the collaboration of different stakeholders that is lacking - the joined-up thinking required for formulating these economic arguments needs to be encouraged.

Coming back to a purely data-centric view of where we are now, the current lack of interoperability on a data level has lead to data being 'siloed'. Unlinked, standalone pools of data exist and whilst clear medical use can be derived from them, much richer benefits would come from the data being integrated (or at least integrate-able) with the other silos of data out there. Further, the siloing of data often encourages notions of ownership to develop around it. Where data is siloed and not immediately compatible with other data sources the onus on the custodians of that data to protect its use in an ethico-legal sense will often get formulated in terms of ownership of data. Once data is notionally owned by its custodians economic value tends to be assigned to it and there is again a reticence to share it, despite potential economic benefits that would come from taking a long-term view of the population-level benefits of integrated data. Once again an healthcare ecosystem that encourages a social aspect can mitigate the creation and protection of these data silos.

Clearly where we are now is some way away from where we want to be. There exists a gap, both in the UK system and EU wide, between the vision of integrated population level systems that will enable the type of healthcare outlined by D2.2 on the one hand and the state and existance of semantically rich and integrated health systems on the other. We now move on to looking at a solution that has been deployed across a number of population-level health research scenarios and goes some way to bridging the gap that has been highlighted between where we are and where we want to be in terms of systems for addressing the various population-level raised in D2.1.

Clearly where we are now is some way away from where we want to be. There exists a gap, both in the UK system and EU wide, between the vision of integrated population level systems that will enable the type of healthcare outlined by D2.2 on the one hand and the state and existance of semantically rich and integrated health systems on the other. We now move on to looking at a solution that has been deployed across a number of population-level health research scenarios and goes some way to bridging the gap that has been highlighted between where we are and where we want to be in terms of systems for addressing the various population-level raised in D2.1.

An eLab is an information system for bringing together people, data and analytical methods at the point of investigation or decision-making.[2] We can view an eLab as either a specific instantiation of an information system aimed at addressing a particular problem or class of problems, or as an approach or methodology for addressing such problems. As an application an eLab aims to provide an online collaboration space where data can be shared (within existing ethico-legal boundaries) and integrated. Beyond simply providing shared access to data an eLab should provide a collaborative environment for developing methods that can be applied to data. On top of this, and perhaps more importantly, an eLab should provide an environment that enables and enhances social interaction between users of the data held within the eLab .[3]

[2] Ainsworth, John D., and Iain E. Buchan. "e-Labs and work objects: towards digital health economies." Communications Infrastructure. Systems and Applications in Europe. Springer Berlin Heidelberg, 2009. 205-216.

[3] Couch P, O'Flaherty M, Sperrin M, Green B, Balatsoukas P, Lloyd S, McGrath J, Soiland-Reyes C, Ainsworth J, Capewell S, Buchan I. e-Labs and the Stock of Health Method for Simulating Health Policies. Studies in health technology and informatics. 2013 Jan; (192):288–92.

eLabs have been used to address a range of population-level research questions ranging from providing information on the prevalence of long term conditions in the geographic area covered by primary care practices to enhancing the quality of information driving the commissioning of Bariatric Surgery in the North of England.[4] The key idea underlying all the use cases of eLab-type systems was that the primary focus of the informatics platform in question was in creating a social space for bringing together the actors who worked on data rather than having the primary focus be solely on the integration or harmonization of the data itself.

This annex has outlined the view from workpackage 2 as to where we are and where we need to be in terms of creating an environment that will enable the full potential use of healthcare data integrated on a population level. Further it has outlined as an exemplar some research pertaining to eLabs, their development and deployment. eLab type environments (i.e. those that) are the type of informatics solution needed to bridge the gap between the current state of siloed and over-protected pools of data and a more integrated (both semantically and socially) state of healthcare system which would foster informatics solutions that can successfully address population-level challenges. It has been argued that it is not novel technological solutions to interoperability that are needed but rather the creation of social environments that can bring together the people and methods surrounding data, and from these integrated environments novel informatics approaches to population-level healthcare can emerge.

[4] John AINSWORTH, James CUNNINGHAM, and Iain BUCHAN. "eLab: Bringing Together People, Data and Methods to Enhance Knowledge Discovery in Healthcare Settings." HealthGrid Applications and Technologies Meet Science Gateways for Life Sciences 175 (2012): 39.