The implications of electronic health records for personalized medicine

Amnon Shabo (Shvo)

IBM Research Lab in Haifa,
Haifa 31905, Israel
Tel.: +972 544 714070;
Fax: +972 4 8296116;
E-mail: shabo@il.ibm.com

The emerging concept of an electronic health record (EHR) targeted at a patient centric, cross-institutional and longitudinal information entity (possibly spanning the individuals lifetime) has great promise for personalized medicine. In fact, it is probably the only vehicle through which we may truly realize the personalization of medicine beyond population-based genetic profiles that are expected to become part of medication and treatment indications in the near future. The new EHR standards include mechanisms that integrate clinical data with genomic testing results obtained through applying research-type procedures, such as full DNA sequencing, to an individual patient. Although the most optimal process for the utilization of integrated clinical-genomic data in the EHR framework is still unclear, the new Health Level Seven (HL7) Clinical Genomics Draft Standard for Trial Use suggests using the 'encapsulate & bubble-up' approach, which includes two main phases: the encapsulation of raw genomic data and bubbling-up the most clinically significant portions of that data, while associating it with clinical phenotypes residing in the individual's EHR.

The vision of personalized medicine relies on the ability to effectively associate personalized genomic data with clinical data to support the clinical decision at the point of care for the individual patient [1–3]. It is evident that the healthcare information world is extremely fragmented [4,101]. This situation has resulted in the low availability of patient data and more importantly, data that is incomplete, typically episodic, and created in the proximity of the point of care [5,6]. Achieving complete patient data has the following three main barriers, due to institutional, temporal and topical boundaries.

The institutional boundary is perhaps the most apparent. Every healthcare provider is required by current law to maintain the medical records it created in a long-term repository, which also means that medico-legal copies of those records are those residing in the archives of that provider. Since each provider is running different information systems using proprietary data formats, and executing different security and privacy rules, it is an extremely difficult challenge to integrate data for a specific patient from disparate sources (e.g., various hospitals where the patient was hospitalized, various lab testing facilities, and various consultants/specialists).

The temporal barrier is simply the average current lifespan. This poses a question: who could possibly sustain patient data for about 70 years? Current laws require providers to maintain the data for a much shorter time (approximately 25 years or much less in many places).

The topical barrier is the semantic challenge, in that given a pile of medical records (time-stamped data items such as documents, lab results and medications), how could we integrate patient data so that it could be presented by topic, and, for example, as a summary of a patient's diabetes, based on data from all temporal records?

Medical records are attested data created by healthcare providers and handled by healthcare information systems often called electronic medical record (EMR), electronic patient record (EPR), computerized patient record (CPR), and so forth. The information systems and information entities (i.e., the medical records) are all centred on the healthcare enterprise needs and workflows. The needs of the patient and the secondary uses of medical records, such as clinical trials and medical research, are not adequately addressed by the current constellation of information in healthcare [7]. Into this turmoil we now add genomics, and as genomic research and testing become increasingly reliable, it has become a growing challenge to integrate the new data into the reality of dispersed and dissimilar medical records that belong to the same patient.

Rise of the electronic health record

The aforementioned challenges have led to the emergence of the electronic health record (EHR) concept. Health records differ from medical records in that they are longitudinal and cross-institutional, and may include individual health data not necessarily created

future medicine

during medical practice (e.g., self documentation, lifestyle, environmental details, workplace hazards, personal preferences, and so on). The primary goal of the EHR is to support continuity of care by providing a complete history of the patient with topical summarizations, based on ongoing secondary processing of the raw data by intelligent software applications linked to ontologies of the most updated medical knowledge, as well as the specifications of all recognized data standards.

A major question is how to compile and sustain a coherent EHR across the lifetime of an individual. Shabo and colleagues suggested a new vision arguing that lifetime EHRs should be sustained by new players in the healthcare arena who will function as Independent Health Record Banks (IHRBs) [8]. These multiple competing organizations should be established following pre-emptive legislation, and regulated by that legislation. They should also be independent in the sense that they are neither owned by healthcare providers nor by health insurer/payers or government agencies. The fundamental principle of the new legislation is that the medico-legal copies of all individual's records are stored in these banks, and healthcare providers would no longer serve as the legal record keepers. This revolutionary constellation brings two main benefits to providers as well as to consumers:

- Providers will be able to cut the costs of long-term record keeping, which is mandatory according to current law.
- Better care will be provided based on the availability of new patient life-long health records.

This model is not centered on any of the current players in the field. It is not government-centric, as it does not suggest having national repositories of citizens EHRs. Additionally, it is not provider-centric, as it suggests moving records created by the providers to an objective custody. It is not even consumer-centric, as it does not suggest allowing patients to own and maintain copies of their records. This model is non-centric, as in essence it is focused on the objective and independent service of sustaining individual EHRs, much like financial banks are perceived with regard to certain financial assets.

The new IHRBs would be able to better serve the vision of personalized medicine, as they are solely focused on maintaining a coherent clinical history, which could include personalized genomic data. Such a complete history of the patient makes the utilization of genomic data more effective.

Electronic health record standard – standardization efforts

Any sustainability model of longitudinal EHRs requires the use of EHR standards. However, the active standard development organizations (SDOs) in healthcare have not yet produced an agreed-upon standard for EHRs. The Health Level Seven (HL7) SDO [102] has created a specification for a medical record message, which essentially only relays a clinical document to the medical records unit of the healthcare enterprise. The clinical document format is defined in a newly created standard from HL7 called the Clinical Document Architecture (CDA) [9]. A collection of patient CDA documents might serve as a good basis for an EHR. However, that idea is beyond the scope of the HL7 spec. A major problem with this approach is the redundancy and lack of coherency in such an aggregation. In Europe and Australia, attempts in the 1990s to standardize the EHR resulted in a few 'pre-standard' specifications, such as prENV 13606 by TC-251 of the European standardization institute Comité Europeén de Normalisation (CEN) [103]. As opposed to the document-centric approach, the European standard views the EHR as a standardized information entity to which clinicians should enter the data directly during attested transactions. Another interesting division arises when we examine proposed definitions of EHR. In a USA Institute of Medicine (IOM) document, EHR is proposed to be "components that form the mechanism by which patient records are created, used, stored, and retrieved", thus focusing on the processing aspect of EHR. In contrast, the ISO TC-215 specification draft (influenced by CEN and OpenEHR [104]) has the following definition proposal [106] which focuses on the informational aspect of EHR: "a longitudinal collection of personal health information concerning a single individual, entered or accepted by healthcare providers, and stored electronically."

The electronic health record system functional model and standard

In 2003, the HL7 SDO was asked by several US agencies to produce an 'EHR functional model and standard' aimed at improving the care provided to their patients and to all other health consumers [107]. The functional model focuses on the requirements of EHR systems and the respective

functionalities needed to fulfill those requirements. It includes functions organized by either:

- A hierarchy within the broad headings of care delivery and infrastructure functions.
- A list of functions deemed essential or desirable within four common care settings: hospital care, ambulatory care, nursing home care, and care in the community and personal healthcare.

The last setting indicates the intention of this model to be comprehensive and go beyond the requirements of enterprise clinical records. There are no functions in this model that relate to genomic data but efforts are being made to introduce such functions in later versions of that standard. It should be noted that semantic interoperability is not emphasized in this effort, as it leaves the implementation of each set of functions to the vendor community. Since the information involved in each implementation could differ substantially, it does not guarantee semantic interoperability of a patients EHR between healthcare providers.

Clinical trials and clinical genomics data standards

Standardization efforts are also undertaken in the intersection of medical informatics and bioinformatics (for example, clinical trials and clinical genomics). In clinical trials, the Clinical Data Interchange Standards Consortium (CDISC) [105] is active in creating standards for describing all data involved in the clinical trials Operational Data Model (ODM), Submissions Data Standards (SDS), Laboratory Data (LAB), and Analysis Dataset Model (ADaM). In recent years, CDISC activists have established a technical committee in HL7, which expresses the CDISC standards in the HL7 'language'. Furthermore, several HL7 organizational members (e.g., IBM, Mayo Clinic, Cap Gemini, and Ernst & Young) have initiated a new clinical genomics special interest group whose scope is the actual use of genomic data in healthcare practice [108]. Recently, more clinical trials have begun using genomic data, and thus the two groups in HL7 are collaborating towards a shared model for genomic data. In the Clinical Genomics specifications, bioinformatics mark-up formats such as microarray gene expression markup language (MAGE-ML) [109] for gene expression, and bioinformatic sequence markup language (BSML) [110] for biological sequencing, are being utilized much like controlled medical vocabularies (such as the Systematic Nomenclature of Medicine

[SNOMED] [111]) and are being referenced from the current clinical standards.
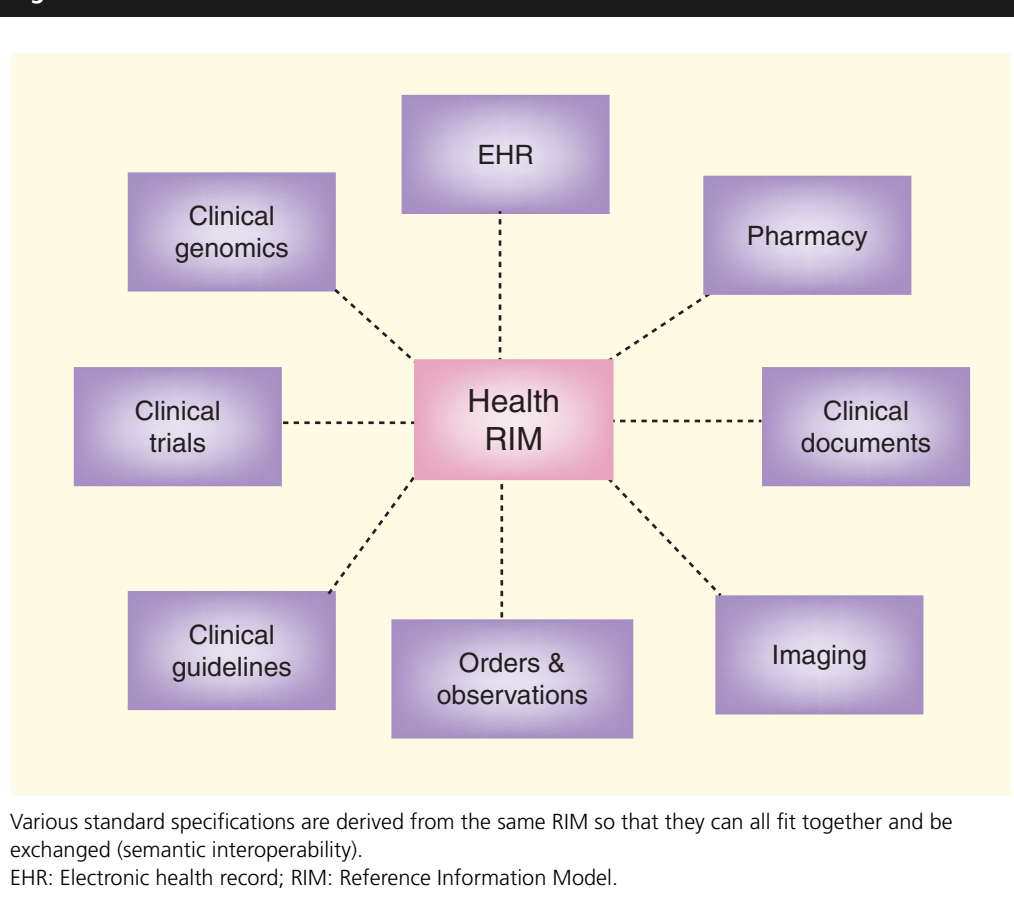
How could the standards fit together?

How could these bioinformatics-oriented efforts be integrated into an individual's EHR? The answer lies in a recent 'methodology shift' that major SDOs are undergoing, including the HL7 SDO. Since 1987, HL7 has been developing messages. In that time, it produced a number of American National Standards Institute (ANSI) standards for healthcare messaging. However, these standards lack a common 'language', and their development methodology has not been consistent. The new version 3 represents a quantum leap with the emergence of a central Reference Information Model (RIM). The RIM consists of common building blocks such that each committee can derive its specifications from the RIM and base it on those building blocks. These include core classes such as acts (observations, procedures etc.), entities (patients, organizations etc.) and associations (act relationships, participations etc.), organized in a coherent model called the Unified Service Action Model (USAM). This forms a common 'language' that all HL7 standards use: EHR and clinical trials as well as other specifications. It is anticipated that all health SDOs will eventually agree on a common 'Health RIM', and this agreement will boost the usability of all specifications derived from that RIM – enabling the desired goal of semantic interoperability (**Figure 1**).

HL7 clinical genomics

As previously noted, the Clinical Genomics Special Interest Group (SIG) develops HL7 standards derived from the HL7 RIM, ensuring that they fit all other standards and in particular EHR standards under development. The standards enable the exchange of clinical and personalized genomic data between interested parties. In many cases, the exchange is carried out between disparate organizations (e.g., healthcare providers, genetic labs, research facilities etc.) and acceptable standards are crucial for the usefulness of the data in healthcare practice.

The Genotype model

The main specification developed by the HL7 Clinical Genomics SIG is the Genotype model. It describes various types of genomic data relating to a chromosomal locus, and is proposed to be the basic unit of genomic information exchange in healthcare. This model is not intended to be a comprehensive biologic model;

Figure 1. Central Health Reference Information Model.

Various standard specifications are derived from the same RIM so that they can all fit together and be exchanged (semantic interoperability).
EHR: Electronic health record; RIM: Reference Information Model.

rather, it is aimed at the needs of healthcare with the vision of personalized medicine in mind. Nevertheless, it could facilitate the needs of clinical research conducted within healthcare enterprises when it involves a closer association with clinical data. The Genotype model is the result of the SIG's efforts to find commonalities in genomic-oriented storyboards that have been explored (e.g., tissue typing, cystic fibrosis, BRCA, and pharmacogenomics). To date, it is a generic model, but in future versions it might be refined to create specialized models by subject (e.g., human, animal and viral genotypes) or by type of data (e.g., DNA, expression and proteomics). It has been balloted in HL7 as a Draft Standard for Trial Use (DSTU), and thus the trial use of this standard will guide the derivation process in a way that is most useful to its early adopters and use cases.

The entry point to the Genotype model is a specific chromosomal locus that could be optionally associated with a number of alleles or directly associated with gene expression data if no allelic information is available. Each of the alleles could

be associated with sequence, sequence variation and expression data. These core classes in the model are also those that encapsulate raw genomic data. In addition to the core classes, other classes hold extracted and derived data such as certain types of variations, main expression results, haplotype references and proteomics data (for example, determinant peptides). Both the genotype and allele classes could recurse to represent relationships to genes/alleles from other loci. The sequence class could also recurse to allow the representation of the translational path, (i.e., DNA, RNA and protein sequences) derived from each other. Finally, all genomic classes can be associated with clinical phenotypes residing either internally within a genotype instance, or externally outside of the genotype instance (for example, in the patient's EHR). **Figure 2** shows a bird's eye view of the Genotype model.

'Encapsulate & bubble-up'
When exploring the Genotype model, it is possible to identify the utilization of bioinformatics markups, such as BSML, in a way that allows a

Figure 2. A bird's eye view of the Genotype model and its focal areas, all associated with clinical phenotypes (observed or known).
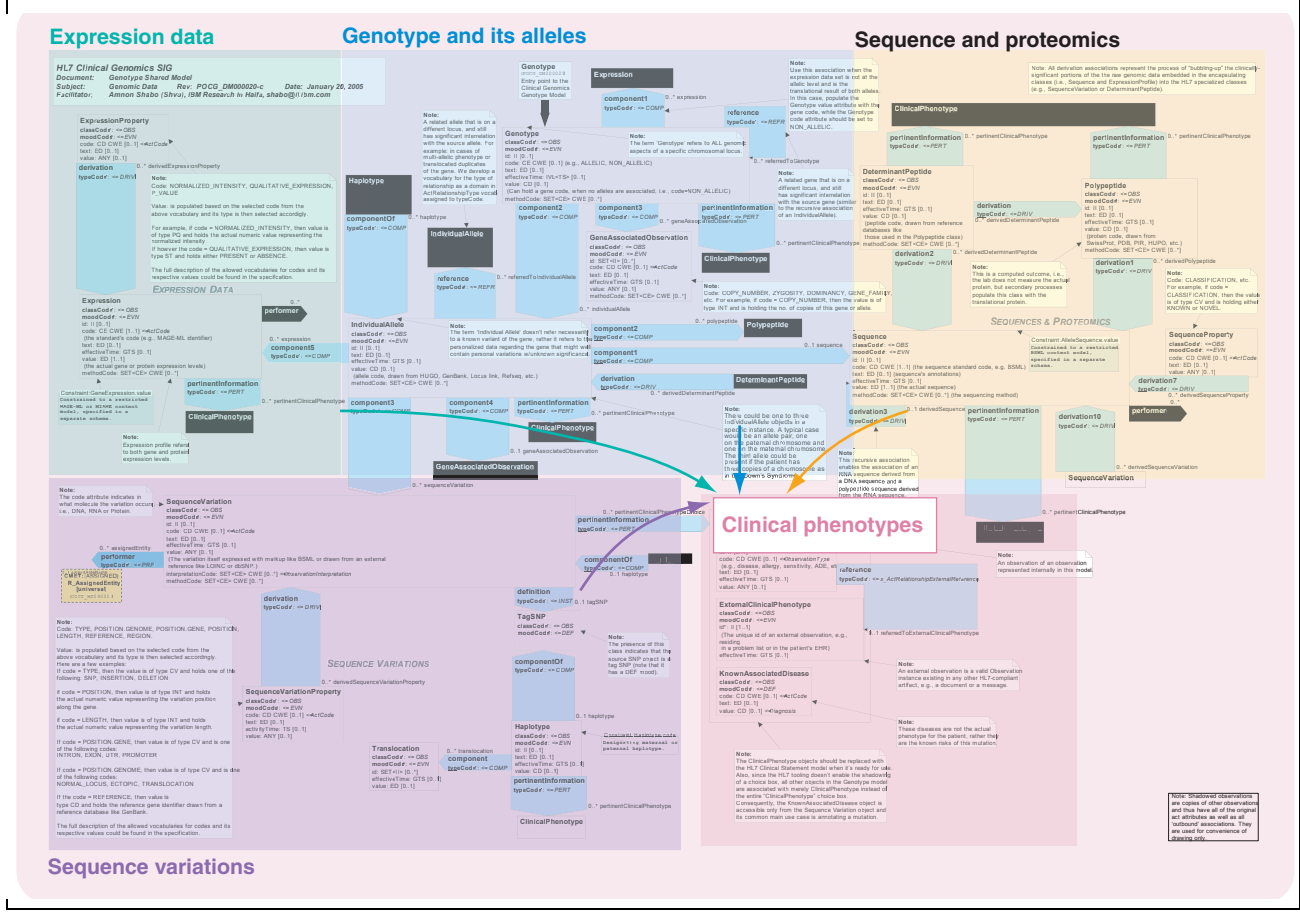

few of the HL7 classes, such as the sequence variation class, to overlap with elements of the bioinformatics markup. For example, it is possible to find a single nucleotide polymorphism (SNP) represented in both the sequence encapsulating class (using BSML) as well as in the sequence variation class (using HL7 attributes). The question then arises: what is the relationship of the two representations and how do they coexist? Data representations, such as MAGE and BSML, are considered the 'raw genomic data' in this context, and the assumption is that much of this raw data might have relevance in the near future as new discoveries are made through biomedical research.

The main paradigm underlying the design of this standard can be described as 'encapsulate & bubble-up', and aims at addressing the coexistence mentioned above. The 'encapsulation' phase is a static process in which certain 'encapsulating objects' in the Genotype model are populated with portions of raw genomic data based on predefined constrained bioinformatics markups. The constraining process is part of the standardization effort and is designed to exclude portions that seem irrelevant to clinical practice, for example, the display elements in the BSML markup. The constraining process also ensures that the data refers to only one patient and one gene, in order to fit the Genotype scope. To represent multiple gene/allele/locus data, higher level models utilizing the Genotype model are being applied (e.g., the Genetic Profile model, the Tissue Typing Observation and the Family History specification).

The 'bubbling-up' phase is a dynamic process where genomic-oriented decision–support applications will parse the raw genomic data already encapsulated in the HL7 instance, and surface up those portions that seem to be clinically significant to the patient's clinical history and treatment goals based on the most updated scientific knowledge. The results of this

Figure 3. Bubbling-up the most clinically-significant portions of the raw genomic data embedded in the encapsulating objects.

DS: Decision support; EHR: Electronic health record; HL7: Health Level Seven; KB: Knowledge base.
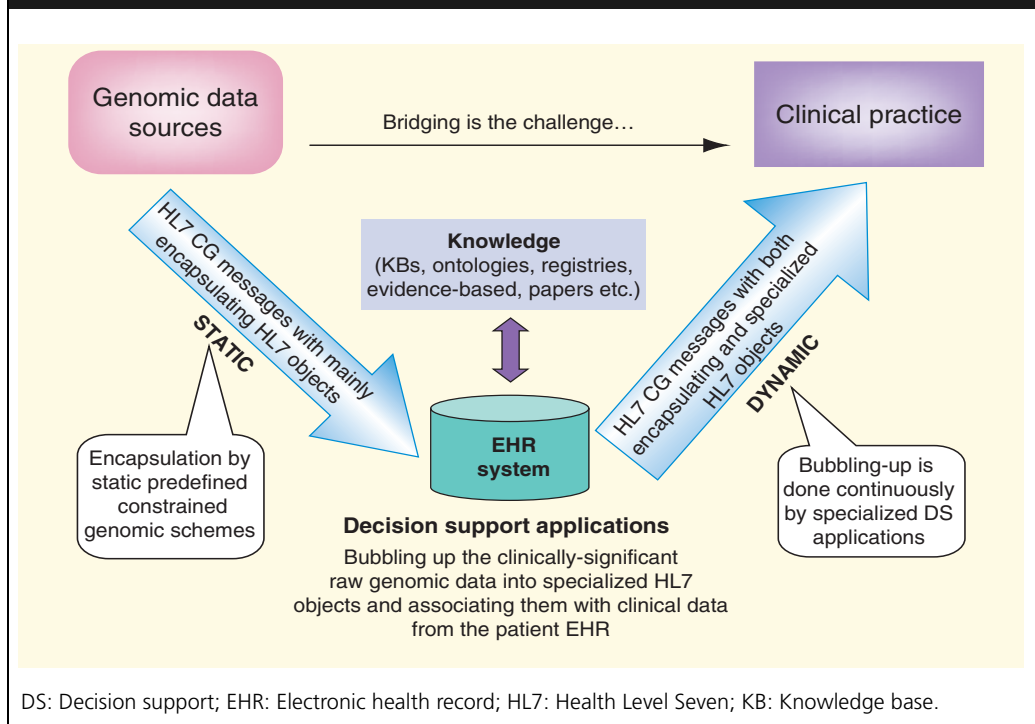
bubbling-up process are held in other HL7 objects in the Genotype model that can also be associated with clinical data in the patient's medical records (represented in the Genotype model as the 'clinical phenotype' classes).

These static and dynamic phases lead to a gradual 'distillation' of the raw genomic data in the context of diagnosis and treatment provided to a specific patient, while holding parts of the raw data within the HL7 objects so that they can be parsed again when, for example, new knowledge becomes available. The complete raw genomic data will be accessible only 'by reference', possibly using the Life Sciences Identifier (LSID – a new OMG specification [112]).

This situation is analogous to raw imaging data. The average referring physician does not examine the dozens of CT images taken in a study. Instead, he or she looks at the already designated regions of interest (ROIs) or just reads the radiologist's report. Similarly, many clinicians might not be able to grasp the raw genomic data represented by the bioinformatics markup. However, it could be helpful to have parts of that data in the HL7 instances for purposes such as evidence tracing and secondary or future processing. The HL7 classes should be seen as representing the digest of the raw genomic data that is most pertinent to the

healthcare practice itself. There is room here for applications that might parse the bioinformatics markup, intelligently populate the HL7 classes, and subsequently associate them with the appropriate clinical data. The HL7 classes in the Genotype model have the advantage of being better tied with the other HL7 classes in the patient EHR (for example, in a problem/allergy list). As a result, these classes are more capable of linking individual genomic data to the clinical data of that individual.

Various bioinformatics markups could be recognized so that there will not necessarily be one winning format. The issue should not be which bioinformatics model is the best fit with a healthcare reference information model such as the HL7 RIM? Instead, we should focus on how we could develop mechanisms to digest genomic data in various representations and link them to the RIM for the benefit of personalized medicine.

Figure 3 shows a conceptual workflow where the above coexistence takes place and is executed step-wise. First, the static phase takes place and HL7 messages are sent with merely encapsulating objects carrying raw genomic data. In the second phase, these messages are enriched with 'bubbled-up' objects that are used by the end-user application at the point of care.

Highlights

- The adoption of electronic health records (EHRs) is crucial for the success of personalized medicine.
- Integration of clinical data with genomic data is a challenge, partly due to ontological gaps between these two distinct information worlds.
- Incorporation of high-resolution individual genomic data (for example, full sequencing) could pinpoint the integration of clinical genomic data to the precise individual condition, beyond the use of generic genetic profiles.
- Longitudinal and cross-institutional EHRs could increase quality of care and patient safety, and provide an adequate container for personal genomic data.
- The sustainability model of lifetime EHRs is a major issue. A model of multiple competing Independent Health Records Banks is proposed as a comprehensive solution to the problem of fragmentation in healthcare, mainly in light of ethical and medical considerations.
- EHR functional and informational models are being developed by international Standards Developing Organizations.
- In parallel, standards for correlated clinical–genomics patient data are being developed by the Health Level 7 Clinical Genomics Special Interest Group, enabling the realization of the 'encapsulate & bubble-up' paradigm. This process allows the encapsulation of raw genomic data, while bubbling-up the most clinically-significant items (e.g., single nucleotide polymorphisms) and associating them with known phenotypes (scientific knowledge) or observed clinical phenotypes residing in the EHR.
- True semantic incorporation of clinical genomic standardized data into an EHR could be achieved by using the new generation of medical informatics standards, which are based on a central Reference Information Model (RIM), from which all specifications are derived.

For example, if the raw data is the full sequencing of certain genes, known mutations could be identified. However, in addition, the most clinically significant SNPs could be 'bubbled-up' from the raw sequencing data and associated with clinical phenotypes based on the current knowledge. This annotated and enriched message could then be presented to the genetic counselor at the point of care. Savel and co-workers [10] used the Genotype model to represent C-reactive protein, pentraxin-related (CRP) data in the National Health and Nutrition Examination Survey (NHANES) project of the USA Centers for disease Control and Prevention (CDC) [113]. A sample XML was crafted to demonstrate the usability of the 'encapsulate & bubble-up' approach.

Conclusion

The implications of EHR standardization and sustainability for personalized medicine are enormous. Patient-centric, longitudinal, and cross-institutional EHR will be the major carrier of a complete patient health history, including genomic data. While posing many privacy concerns [114], it is clear that only such a complete history can assure the highest quality of care, and

in particular, the best utilization of genomic testing results in clinical decision making, realizing the vision of personalized medicine. However, the fragmentation in the healthcare information world is a great barrier to this end. Even if EHR standards are already in place, it might still be difficult to get a complete EHR due to the possible unavailability of several of the data sources where the patient's medical records actually reside.

In terms of informatics, medical informatics and bioinformatics have thus far evolved separately, and therefore rely on different information models. In addition, only small portions of the raw genomic throughput are currently relevant to clinical practice, however, it is foreseen that a growing amount of this data will become significant in the future. The integration of raw genomic data into operational clinical information systems so that the data becomes useful at the point of care is another great challenge in this field. The 'encapsulate & bubble-up' approach developed by the HL7 Clinical Genomics SIG has been described as a possible way to address this challenge.

Outlook

The field of EHR will evolve in the next 5–10 years and attempt to address the fundamental problem of individual data fragmentation in the world of healthcare information. In the next 5 years, it is foreseen that regional EHR registries [11] will enable the federation of patient data on the fly based on metadata and pointers stored in regional registries or even national registries [12]. However, the inevitable solution of IHRBs briefly described in this paper will emerge in the next 5–10 years, as it becomes apparent that federation of patient data is not a safe process in a world of increasing mobility of people, and as many of the data sources might not be available (for example, small clinics that went out of business) when the data is needed at the point of care. In addition, it will be difficult to intelligently integrate genomic data into a 'virtual' patient EHR. A persistent EHR residing in an IHRB will be a better choice as the integration and summarization processes will be running continuously with the utilization of the most accurate and up to date knowledge available.

Acknowledgment

Bibliography

Papers of special note have been highlighted as either of interest (•) or of considerable interest (••) to readers.

1. Ruano G: *Quo vadis* personalized medicine? *Personalized Med.* 1(1), 1–7 (2004).

2. Davis RL, Khoury MJ: The journey to personalized medicine. *Personalized Med.* 2(1), 1–4 (2005).

3. Vazquez ES: Personalized therapy: an interdisciplinary challenge. *Personalized Med.* 1(1), 127–130 (2004).

4. Lawrence D: *From Chaos to Care: The Promise of Team-Based Medicine.* Perseus Publishing, Cambridge, MA, USA (2002).

5. Kohn LT, Corrigan JM, Donaldson MS (Eds): *To Err Is Human.* Institute of Medicine, Washington, DC, USA (1999).

6. Kohn LT, Corrigan JM, Donaldson MS (Eds): *Crossing the Quality Chasm.* Institute of Medicine, Washington, DC, USA (2001).

7. Salvatore JG (Ed.): Balancing medical innovation and patient confidentiality. *The Pfizer Journal.* 5(3), 1–39 (2001).

8. Shabo A, Vortman P, Robson B: Who's afraid of lifetime electronic medical records? Proceedings of the Towards Electronic Health Records Conference (TEHRE 2001). London, UK (2001).

•• **This paper describes in detail the vision of Independent Health Records Banks and discusses why it might be the inevitable sustainability model of lifetime EHRs.**

9. Shabo A: Structuring the medical narrative in patient records – a further step towards a multi-accessible EHR. Synopsis of the patient records section. In: *IMIA Yearbook of Medical Informatics 2004: Towards Clinical Bioinformatics.* Haux R, Kulikowski E (Eds), Schattauer, Stuttgart, Germany, 317–320 (2004).

10. Savel T, Lin B, Shabo A, McQuillan GM: The future of transmitting human genomic data in the public health information network (PHIN): Using a prototype Health Level 7 shared genotype refined message information model (HL7 R-MIM) and the extensible markup language (XML). The 3rd Annual PHIN (Public Health Information Network) Conference, Atlanta, Georgia, USA (2005).

•• **This paper presents the first experimental use of the HL7 Clinical Genomics Specifications within the NHANES project of the USA CDC.**

11. Karp S: Santa Barbara County Care Data Exchange. Cornerstones for Electronic Healthcare. Washington DC, USA, (2004).

•   **Santa Barbara is one of the first regions in the USA to prototype with regional electronic health records.**

12. Yasnoff WA, Humphreys BL, Overhage JM *et al.*: A consensus action agenda for achieving the national health information infrastructure. *J. Am. Med. Inform. Assoc.* 11(4), 332–338 (2004).

•   **A comprehensive review of the considerations leading to the need for a National Health Information Infrastructure.**

Websites

101 http://www.healthcareinformatics.com/newsclips/newsclips03_3_05.htm
Dietrich L, Rothschild A, Ball MJ: Patient-centered clinical documentation – Problem-driven electronic health records can help defragment healthcare. *Healthcare Informatics Online.* Accessed April 23, 2005.

102 http://www.hl7.org
HL7 (Health Level Seven) – An ANSI-Accredited Standards Developing Organization. Accessed April 23, 2005.

•   **The HL7 site is the source of the HL7 Standard Specifications, including the EHR Functional Model and the Clinical Genomics specifications. Some resources are restricted to members only.**

103 http://www.centc251.org/
Health Informatics – Electronic healthcare record communication – Part 1: Extended architecture. ENV13606–1, Committee European Normalisation, CEN/TC 251 Health Informatics Technical Committee, 2000. Accessed April 23, 2005.

104 http://www.openehr.org/
OpenEHR. An international not-for-profit foundation, working towards interoperable, life-long electronic health records, proven in practice and understanding the social, clinical and technical challenges of electronic records for health care in the information society. Develops open source code for EHR interoperability. Accessed April 23, 2005.

105 http://www.cdisc.org/
Clinical Data Interchange Standards Consortium (CDISC). Accessed April 23, 2005.

106 http://secure.cihi.ca/cihiweb/en/downloads/infostand_ihisd_isowg1_mtg_denoct_contextdraft.pdf
Health Informatics – Electronic Health Record Definition, Scope, and Context. ISO Draft Technical Report 20514, Second Draft, August 2003. Accessed April 23, 2005.

107 http://www.hl7.org
Dickinson G, Fischetti L, Heard S: *HL7 EHR System Functional Model – Draft Standard for Trial Use.* July, 2004. Accessed April 23, 2005.

108 http://www.hl7.org
HL7 Clinical Genomics Special Interest Group.

109 http://www.mged.org/Workgroups/MAGE/mage.html
MAGE-ML Microarray and Gene Expression Markup Language. A standard for the representation of microarray expression data. Accessed April 23, 2005.

110 http://www.bsml.org
Bioinformatic Sequence Markup Language (BSML). Accessed April 23, 2005

111 http://www.snomed.org/
Systematized Nomenclature of Medicine (SNOMED). Originally developed by the College of American Pathologists (CAP), and has been recently licensed to the US government. Accessed April 23, 2005.

112 http://xml.coverpages.org/lsid.html
The Object Management Group (OMG) Life Sciences Identifiers (LSID) specifications are persistent, location-independent, resource identifiers for uniquely naming biologically significant resources including, but not limited to, individual genes or proteins, or data objects that encode information about them. Accessed April 23, 2005.

113 http://www.cdc.gov/nchs/nhanes.htm
The Division of Health and Nutrition Examination Surveys (DHANES) is one of the survey divisions at the CDC's National Center for Health Statistics (NCHS). Accessed April 23, 2005.

114 http://www.nap.edu/books/0309056977/html/index.html
US Committee on Maintaining Privacy and Security in Health Care: For the Record: Protecting Electronic Health Information, part one. Accessed April 23, 2005.