# Chapter 1
# A Brief History and Overview of Health Informatics

## 1.1 Introduction

This first chapter provides a brief history of some, but certainly not all, of the key subdomains within the health informatics field and further explains the potential significance of the FHIR standard that will occupy much of the rest of the book. To do this, the chapter begins with a discussion of early electronic records and clinical decision support tools and then shifts gears to introduce the concept of health information exchange. Later, we discuss interoperability challenges that date back decades and the various ways that existing technologies have been used, sometimes with limited success, to simplify and coordinate the sharing of information among providers. The chapter ends with the premise that widespread adoption of modern web technologies (and FHIR in particular) is transforming health informatics. To help illustrate this the chapter ends with a demonstration FHIR app developed by a team of Georgia Tech students did using these emerging technologies to help predict the onset of a life threatening condition in ICU patients.

One of the things I thought about in preparing the book was the rate at which we are losing the early pioneers in health informatics so I have made a conscientious effort to briefly describe their work wherever it fits into the narrative. In many cases this emphasizes how long-standing many of the key health informatics challenges are, and it enriches the discussion. As I will throughout the book, I provide references or links to supplemental material that should provide more detail and expand your appreciation of the field.

## 1.2 Early Electronic Records and Clinical Decision Support

This book is about health informatics, which encompasses the many applications of computing involved in the *delivery* of healthcare. This is distinct from the related field of bioinformatics which seeks to understand and model the incredibly complex



M. L. Braunstein, *Health Informatics on FHIR: How HL7's New API is Transforming Healthcare*, https://doi.org/10.1007/978-3-319-93414-3_1

biochemical engine within each of our cells. Historically health informatics has therefore been about applying technologies rather than gaining new knowledge. This is changing because of the increasing sophistication of modern analytic and machine learning[1] technologies coupled with the broader availability of digital health data from providers and their patients. It is now clearly the case that health informatics can be the source of new knowledge about how best to diagnose and treat disease in real world patients. Readers may also run into the term 'biomedical informatics' which I, perhaps loosely, describe as the combination of these two technologies into one academic discipline and/or department.

The field has a long history. Since the 1950s when computers could first store and process 'large (by the standards of the day)' amounts of data, software developers have seen healthcare as a fertile domain. Electronic medical record systems (EMRs) date back at least to the early 1960s when Akron Children's Hospital and IBM created HIS, which is claimed to have been the first computer-based clinical hospital information system (Fig. 1.1). I provide a link to an IBM video from the period explaining the system.[2] I believe you will find it quite entertaining.

Ambulatory electronic medical record systems for use outside of hospitals began in the late 1960s. Starting in the 1970s, as you can see in the photo of me back then, I oversaw the development and did some of the programming of one of these early EMRs at the Medical University of South Carolina (MUSC) in Charleston (Fig. 1.2).[3]

Clinical decision support systems were another important early application of informatics to patient care. These systems seek to guide physicians based on scientific evidence as they make diagnostic and therapeutic clinical decisions. They date at least to 1972 when Stanford health informatics pioneer, Dr. Ted Shortliffe, began work on MYCIN, a program to assist physicians to make optimal antibiotic selection to treat infections. We will discuss MYCIN in some detail later on in the book but, a number of years later, Dr. Shortliffe posted a pair of YouTube videos to demonstrate and memorialize the program since, as he said, the computers it could run on were disappearing.[4]

---

[1] Machine Learning is the branch of artificial intelligence focused on the development of algorithms to do things such as classifying or clustering items of interest. In supervised learning input and output data are labeled and an algorithm learns the mapping function from the input data to the output labels by comparing its output to the correct output. In unsupervised learning there is only unlabeled input data and so the model is left to discover and present interesting structures in the data. In medicine, these techniques can use EHR data to classify patients into those that have diabetes and those that do not have it or to cluster diabetic patients into different diabetic subtypes.

In classical machine learning algorithms, the focus is on converting input data into features for supervised and unsupervised learning (feature extraction or feature engineering). There is currently great interest in "deep learning" a subfield that uses massive amounts of labeled data and the power of modern computers to automatically generate useful features and classification models to obtain previously unachievable levels of accuracy.

[2] https://www.youtube.com/watch?v=t-aiKlIc6uk

[3] Braunstein, ML, The Computer in a Family Practice Center: A "Public" Utility for Patient Care, Teaching and Research, *Medical Data Processing*, pp 761–68, Laudet, M, Anderson, J, and Begon, F, Editors. London, Taylor and Francis, 1976.

[4] https://www.youtube.com/watch?v=a65uwr_O7mM https://www.youtube.com/watch?v=ppkg4mQIgXw

**Fig. 1.1** HIS, the first hospital information system. (© 1960 Akron Children's Hospital. All Rights Reserved)

**Fig. 1.2** Both new care models and the role of EMRs in them have both been a subject of interest for many decades. Here in late 1975, Congressman Paul G Rogers (D FL) chair of the House of Representatives Subcommittee on Health and the Environment (middle), visits the very innovative Family Medicine Clinic at the Medical University of SC to see one of the first family medicine residency programs. As part of the visit, Department Chair, Dr. Hiram Curry (left), and I showed him the role of the clinic's very early EMR in scheduling, medical records, billing, pharmacy and patient education. The system also provided support for monitoring patient compliance with their prescriptions, population health and utilization review
*Source*: AAFP Reporter Vol III No 1 January 1976

## 1.3 Health Information Exchange

Sharing of health data also has a long history dating back at least to the 1980s. Historically a sole purpose infrastructure for sharing was envisioned and many were implemented at a health system, community, regional or state basis. While this infrastructure has had many names over the years, it is usually referred to now as a Health Information Exchange (HIE) or a Health Information Network (HIN). A 2011 paper by Gilad Kuperman describes the basic rationale for HIE's as addressing "key healthcare problems that 'siloed' EHRs do not solve". He goes on to give examples of "problems that could only be addressed by interoperability including support for the patient across transitions of care, the ability to perform longitudinal analyses of care, and public-health needs."[5]

Kuperman also lists some of the technical challenges to HIE creation and sustainability: "the need for standards to represent clinical data, the need to identify a patient consistently as they moved among different providers and a framework to assure the patient's privacy." He goes on to say that there are "also questions of who should play a leadership role to address these issues and the kinds of organizational models that could best support interoperability."

Creating workable and sustainable HIEs in the US was one of the goals of the Office of the National Coordinator for Health IT (ONC) the federal agency that created and manages the HITECH program. ONC provided funding to the states for establishing their HIEs and it also funded the Nationwide Health Information Network (NHIN) Prototype Architecture initiative. The NHIN was envisioned as a 'network of networks' connecting the various regional and state health data exchanges. For both technical and political reasons, the NHIN did not require a national patient identifier or a large-scale centralized operation.

ONC then funded the NHIN Trial Implementations project that eventually involved 20 participating organizations. This effort led to CONNECT and open source gateway for secure HIE. CONNECT was so complex and HIE financial resources were so limited that it was not widely adopted.

In 2010, ONC initiated the NHIN Direct project (Direct). We'll discuss this more slowly and in more detail later but, for those who cannot wait, special Direct servers provide email-based capabilities that are tailored to the special needs of health information sharing. Unlike CONNECT, Direct is 'lightweight' and much easier to implement. Typical Direct use cases[6] involve a primary care physician (PCP) sending a patient summary to a specialist as part of a referral, a care provider sending a visit summary to a patient or the transmission of results from a laboratory to an EHR.

Such a highly distributed exchange introduces the trust problem. Trust is essentially knowing for certain who you are dealing with in an electronic network such as the Internet. It is relatively easy within a health system where everyone is employed

[5] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3168299/

[6] A written description of how a user will use a software system to accomplish a task or goal.

and providers are credentialed. When HIE is extended further – to community providers not employed by or formerly affiliated with a health system and particularly to patients in their environment – trust becomes far more challenging.

DirectTrust was created by Dr. David Kibbe to help solve the trust problem for Direct on a national level. The details are beyond the scope of a non-technical book but, given the focus on FHIR, it is interesting that, in August, 2017, FHIR inventor Grahame Grieve and several members of the DirectTrust Policy Committee posted a white paper that suggests that FHIR resources could be 'pushed' (sent) in Direct Messages and DirectTrust certificates could be used with the FHIR RESTful Application Programming Interface (API) to establish trust.[7] We will define APIs later but this development illustrates the possibility of entirely new means of health information exchange that, unlike the earlier efforts, do not require a substantial sole purpose infrastructure. Given HIE's long standing problem of establishing a sustainable business model, that we will turn to next, this could be quite significant.

Beyond the impediments we have already discussed, finding a sustainable business model is arguably the most difficult challenge for HIEs in the US. Since 2001 the eHealth Initiative has surveyed HIEs. It's most recent 2014 survey included 125 of 267 (47%) identified organizations (74 community-based, 25 statewide, and 26 healthcare systems). Nearly 85% were operational but only 36% were financially sustainable with no outside support.[8] These may seem like low percentages given how long HIE has been available but they represent significant progress. The report provides a great deal of other interesting information in an easily read presentation format so I suggest it for readers with a deeper interest in this topic. One key point is that, of the HIEs surveyed, 64 have customers for their services in support of the new payment models such as Accountable Care Organizations (ACOs) and 52 support new care delivery models such as the patient-centered medical home (PCMH). We will discuss these new payment and care delivery models later on, but this clearly illustrates the key role that financial incentives play in informatics adoption.

## 1.4 The Interoperability Challenge

From the 1960s and into the 1980s commercial software was a new industry and the computer systems used in hospitals had limited data storage and processing capabilities. Partially as a result of this, commercial hospital software vendors were new, relatively small organizations that focused on a specific hospital department, such as the clinical laboratory or pharmacy. Each of these so called 'departmental systems' was proprietary and each typically had its own data model. Also, Local Area Networks had yet to appear. This made interoperability, or "data sharing," among

[7] https://www.directtrust.org/wp-content/uploads/2017/08/Direct-FHIR-Whitepaper_vers_1.4_08102017.pdf

[8] https://www.ehidc.org/sites/default/files/resources/files/2014_eHI_Data_Exchange_Survey_Results_Webinar_Slides_0.pdf

these systems, difficult. As the use of these departmentally specific systems grew, their inability to share data became more problematic. For example, when a new patient was admitted to the hospital, the departmental systems had no automated way to know that, to know their demographic information and even what bed they occupied. Some hospitals actually connected a terminal in the nursing unit to each of the systems so that a user could enter information (often the same information) into all of the systems by going from terminal device to terminal device!

In 1987 this problem eventually led to the formation of the global, non-profit healthcare interoperability standards setting organization, Health Level Seven® (HL7). That led to the development of HL7's hugely successful V2 messaging standard for health data. We will only consider data sharing standards from V2 onward. However, there is a rich history of earlier efforts involving some of the pioneers in health informatics. This is chronicled in a white paper those of you interested in the early development of protocols and data sharing networks in healthcare should read.[9]

Despite its age, V2 is still widely utilized today. According to HL7, 95% of US healthcare organizations use one version or another of V2. There are at least a dozen versions listed on the HL7 site. The standard is based on EDI/X12 technology originally developed by the US transportation industry in the 1960s and then adopted by the grocery and manufacturing industries (all these industries had complex supply chains to manage).

Figure 1.3 is an example of an electronic lab test result being reported using V2. We will explore V2 in more detail in a later chapter but the last two (OBX) lines (segments) provide the actual results. In this example, they are the patient's hematocrit (the ratio of the volume of red blood cells to the total volume of blood) of 45% and their erythrocyte (red blood cell) count of 4.94. Both are accompanied by their respective normal ranges. So, are they normal?[10]

```
MSH|^~\&|LABGL1||DMCRES||199812300100||ORU^R01|LABGL1199510221838581|P|2.3
    |||NE|NE
PID|||6910828^Y^C8||Newman^Alfred^E||19720812|M||W|25 Centscheap Ave^^
    Whatmeworry^UT^85201^^P||(555)777-6666|(444)677-7777||M||773789090
OBR||110801^LABGL|387209373^DMCRES|18768-2^CELL COUNTS+DIFFERENTIAL TESTS
    (COMPOSITE)^LN|||199812292128||35^ML|||||||
    IN2973^Schadow^Gunther^^^^MD^UPIN
    ||||||||||^Once||||||CA20837^Spinosa^John^^^^MD^UPIN

OBX||NM|4544-3^HEMATOCRIT (AUTOMATED)^LN||45||39-49
    ||||F|||199812292128||CA20837
OBX||NM|789-8^ERYTHROCYTES COUNT (AUTOMATED)^LN||4.94|10*12/mm3
    |4.30-5.90||||F|||199812292128||CA20837
```

**Fig. 1.3** A laboratory text result is reported as an HL7 V2 message. It is formatted in the EDI/X12 format that is both concise and cryptic. (Courtesy HL7)

[9] http://www.ringholm.com/docs/the_early_history_of_health_level_7_HL7.htm

[10] The answer is yes. The hematocrit of 45 is within the normal range of 39–49 and the erythrocyte count of 4.94 is within the normal range of 4.3–5.9.

This simple example should serve to explain one of the limitations of V2. It was developed at a time when computer memory and storage were both limited and expensive so it is quite intentionally concise, cryptic and, hence, not particularly human readable.

## 1.5 Exciting and Potentially Transformational Times

Earlier in this section you learned that DirectTrust could be used with the FHIR RESTful Application Programming Interface (API). Non-technical readers may not understand what an API is. In non-technical terms it can be understood as a 'contract' that says to software developers that if you send a request from a 'client' computer (e.g., a phone, tablet, notebook or desktop) to a 'server' (the computer where the information is stored) in the specified format you will always get a response in a specified format or initiate a defined action. The key point is that APIs allow developers to access and even update or delete information and other resources on a remote computer without having to understand the technical details of what's going on in the system they are interacting with. As we will learn, this has a significant positive impact on the development of client software. In essence the entire 'app ecosystem' we all use daily on our smartphones is made possible by the phone's APIs.

Any reader who has searched the Internet for information or for a web site to purchase an item did that using an API  Here is an example using Amazon:

https://www.amazon.com/s/ref=nb_sb_noss?url=search-alias%3Daps&field-keywords=size+10+blue+sweater+for+women

You should fairly easily be able to figure out that this hypothetical shopper is looking for a size 10 blue woman's sweater. Try doing this search yourself and you should see a similar string of characters in the field at the top of your browser that displays your query immediately after you click on an icon to initiate the search.

REST is an acronym for **RE**presentational **S**tate **T**ransfer. We mentioned that, using REST, the client and server are independent. The technical term for this is statelessness. Essentially, what this means is that all of the necessary information to handle the request is contained within the request itself. As a result of that, any of the countless thousands of servers that Amazon operates can receive this query and respond to it based entirely on information within the query itself. A moment's reflection should reveal how important that is in managing an environment where queries can arrive at unpredictable times and ask for unpredictable information.

APIs are having a disruptive and transformational impact on many industries. For example, we now take it for granted that we can find and purchase virtually anything with almost no effort. We can conduct most of our banking and financial affairs from anywhere at any time using our phone. We can take a photo of a check and deposit it without visiting the bank or an ATM. This is all done using APIs.

Today, the technologies that support access and sharing of healthcare data are changing rapidly as the number of health-related open APIs and web tools explodes. *The Untapped Potential of APIs in Healthcare*, an article in the Harvard Business Review, begins: “Leaders of most Internet-based businesses have realized the critical importance of using open application programming interfaces (APIs) to expand the reach of their organizations. If the health care industry followed suit, the impact on the quality and cost of care, the patient’s experience, and innovation could be enormous.”[11]

At present, most would agree that HL7’s Fast Healthcare Interoperability Resources or FHIR is the most important healthcare API. FHIR is based on contemporary technologies. Even though FHIR is still under development it is achieving dramatic acceptance and adoption and this may accelerate further now that the R4 version with the first “normative” content was released on 12/27/2018. This book focuses on FHIR because it has the potential to solve many of the interoperability problems that have so long impeded informatics from achieving its potential impact on healthcare quality, cost effectiveness and safety. It does this by offering the promise of far more open and facile access to clinical data (with verified identity and the right permissions). It offers the potential that patients will be able to relatively easily aggregate their health data from whatever providers care for them. They manage and take advantage of that data using apps of their choosing and even make decisions about contributing their data for research or other purposes. In short, we are on the verge of a transformation in healthcare powered by the same technologies and forces that have already transformed most other areas of our lives. We conclude this chapter by looking at an actual example.

**Demonstration FHIR App: Sepsis Watch**

You can see many examples of the transformational potential of an API powered healthcare system by visiting the SMART on FHIR App Gallery.[12] FHIR can be used as the basis for turning electronic medical record systems into open app platforms that are similar in many respects to the app store accessed by the smartphone in your pocket. Much of the attention is on FHIR apps that support physicians in making a more accurate and timely diagnosis and providing more precise, personalized treatment. FHIR also seems poised to solve another longtime challenge in health informatics which is how to give patients access to their complete health record and the tools to use that data to stay well and to more effectively participate in their own care.

[11] https://hbr.org/2015/12/the-untapped-potential-of-health-care-apis

[12] https://apps.smarthealthit.org/

Sepsis is a life threatening generalized organ system dysfunction caused by the body's reaction to a serious infection. Its root cause is the immune systems attempt to deal with an infection by releasing biologic agents into the body that in extreme situations result in generalized inflammation of the vital organs such as the brain, heart and kidneys.[13] Sepsis can be very serious with the risk of death as high as 30% and 50% if it is severe. The most extreme form, septic shock, has death rates of up to 80%.[14]

Sepsis can often be detected early if the right laboratory tests are ordered. This is the goal of the Sepsis Watch FHIR app shown in Fig. 1.4. This app helps the physicians order the right tests in a timely manner. You may wonder where such a tool could be usefully applied.

Sepsis is a particularly common problem in intensive care units (ICU), affecting about 37% of admitted patients and is the leading cause of deaths among patients in non-coronary care intensive care units. The annual cost of hospital care for US patients with sepsis is $14 billion.[15]

Given these statistics, the tele-ICU (eICU) at Emory was felt to be the ideal environment in which to pilot real-time analytics for sepsis.[16] It is also ideally suited to test and refine the various steps involved in the predictive analytics pipeline (the series of steps needed to make a prediction).

Emory Healthcare's eICU currently monitors over 100 patients in five hospitals. The nurses and physicians in the eICU (or 'command center') partner with their counterparts at the bedside to provide an extra layer of monitoring. The eICU provides high-level surveillance using multiscale data acquired from the EHR, near-real time views of physiologic waveforms, and trended data derived from those waveforms. One of the major roles of eICU staff and providers is to respond to alerts and convey concerning trends in data for the bedside providers to act upon. Putting primary responsibility on the eICU to be attentive to and respond to alerts can allow bedside providers to focus their time and resources on patient care for the "right" patient. Furthermore, it can minimize 'alert fatigue,' a common phenomenon where busy clinicians fail to consider advice from information systems.

[13] https://www.ncbi.nlm.nih.gov/pubmed/16424713

[14] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3916382/

[15] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3916382/

[16] Buchman TG, Coopersmith CM, Meissen HW, Grabenkort WR, Bakshi V, Hiddleson CA, Gregg SR. Innovative interdisciplinary strategies to address the intensivist shortage. *Critical care medicine*. 2017 Feb 1;45(2):298–304.

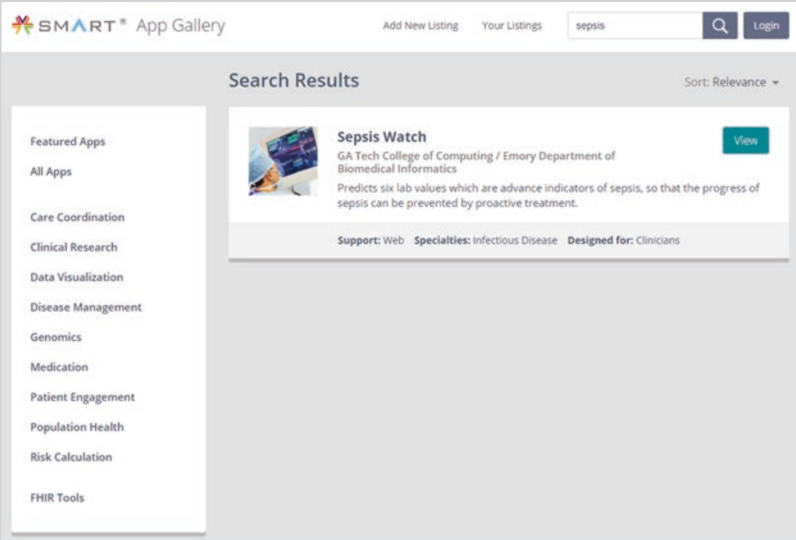


**Fig. 1.4** A video demonstration of the Sepsis Watch app is posted on the SMART on FHIR site. The app was developed by a team of Georgia Tech students taking the author's CS6440 graduate seminar and mentored by Dr. Shemim Namati at Emory School of Medicine's Department of Biomedical Informatics. It predicts six lab values that can be indicators of the onset of sepsis.

## 1.6 A Pivotal Point

The field of health informatics dates back to the early 1960s when computers were first used to store health records. Longstanding interoperability challenges – both technical and non-technical – have impeded the flow of information and progress ever since. We are now at a pivotal point where the building blocks of a more robust and usable technology infrastructure are being developed by HL7 and adopted through the efforts of the government, the health IT and other industries, payers, health care providers, and patients. This book was written in large part to tell that story.