

Chapter 29

Beyond Current HIMS: Future Visions and a Roadmap

James Fackler

29.1 Introduction

This chapter explores a future vision for technologically supported healthcare beyond where current health information management systems (HIMS) have taken us. This vision is informed from over three decades of experience as a Pediatric Intensivist working in academic medical centers; as well as being a Medical Informatician engaged for the past 20 years in system design for a major EHR vendor and multiple consulting roles for smaller niche and start-up health information technology (HIT) companies. Based on these experiences and supported by the HIT literature [1] and health policy bodies [2], I believe:

1. The current health care system is not safe,
2. The billions of dollars spent designing, testing, and implementing HIMSs have been spent instantiating the same workflows that created the unsafe current health care system.
3. It is not just unlikely, but rather *impossible*, that current HIMSs can improve the value of care delivered.
4. The primary way HIMS can improve the value of care delivered [3] is by improving the clinical decisions leading to an improvement in patient outcomes.

The cliché that insanity is repeatedly doing the same processes but expecting different results is quite relevant here. Thus, this chapter describes a significant deviation from the United States healthcare industry's current HIMS strategy built as it is upon a few monolithic electronic health record (EHR) vendors with the

J. Fackler, MD
Pediatric Anesthesiology and Critical Care Medicine,
The Johns Hopkins University School of Medicine, Baltimore, MD, USA
e-mail: fackler@jhmi.edu

inherent limitations associated with monolithic size. Instead, what is envisioned here is an altogether different technology support road that has been started by others and importantly, presents solutions built on twenty-first century technologies. Three more predicates will guide this new journey:

1. Current HIMSSs can play a valuable role in the data collection process.
2. The data collected by these HIMSSs must be augmented with data not now routinely captured in HIMSSs.
3. The data collected by these HIMSSs must then be made accessible to vendor-agnostic patient-centric applications.

29.2 The Imperative for Clinical Decision Support Begun by Others...

Two meta-analyses reviews on clinical decision support (CDS) systems, published independently in 2005, concluded that for a CDS to be effective, the system must be automated [4] and must interrupt workflow [5]. Pushing these two points to their logical extensions, for technology to support our clinical decisions, there must be a tight coupling of the clinician and the computer. There must be more than good human computer interfaces; there must be human-computer symbiosis. In 1960 (yes, 55 years ago), JCR Licklider, a psychologist and pioneer in computer science coined the concept of “man-computer symbioses” [6], noting that while none existed in 1960:

The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today. [6]

Licklider’s vision of the tight coupling of man and computers is closer to being realized, but not yet a reality. We may be on the threshold of Licklider’s vision, if the opinions of Kurzweil [6] and others predicting the “singularity” state – the time when there will be no distinction between human and machine – are correct. It would be a fair bet that clinicians in their early 1940s have a good chance of making clinical decisions in a setting of true human-computer symbiosis.

29.3 Decision Support ...

Looking at what’s involved in decision support must begin with a discussion of what decisions clinicians make, how those decisions are made, and how the decisions are best supported. Broad, often excellent, theoretical guidance is published on clinical decisions [7], but again little guidance is available in the literature cataloging the decisions clinicians actually make. A 2010 study of ten faculty pediatric cardiologists found that each physician made close to 160 decisions per day, and of

these, 80 % were made without any basis in published data [9]. The authors further reported that fewer than 3 % of decisions were based on a study relevant to the specific decision [8]. Even less guidance in the form of evidence-based data is available to entrepreneurs and developers of clinical decision support systems in terms of pointing to which decisions *should* be supported.

The best place to start deciding where to start should follow from the work of Daniel Kahneman whose book “Thinking, Fast and Slow” should be required reading for any developer of future clinical decision support applications [9]. Kahneman suggests a dichotomy between System One (fast) thinking and System Two (slow) thinking. System One operates automatically and quickly with little or no effort and no sense of voluntary control. In contrast, System Two allocates attention to the effortful mental activities that demand it, including complex computations or in a clinical context, puzzling through a complicated patient.

At the bedside, a System One decision is often what is euphemistically called “the art of medicine” but is better recognized as intuition. Gary Klein has long studied intuition in experts [10] lauding the expert’s pattern recognition capabilities. Again, pattern recognition is a core cognitive task and senior physicians perform better than do junior physicians [11, 12]. However, given that the vast majority of clinical decisions are based on precious little data and the most senior clinicians do not often make these decisions, it begs the question whether CDS can support the System One, intuition-based pattern recognition decisions. The answer is – yes; but few such solutions, even as prototypes exist [13]. The reader is referred to a recent review on the subject of supervised classifiers that may have applicability to medical diagnosis [14].

The operations of System Two are often associated with the subjective experiences of choice and concentration. The highly diverse operations of System Two have one feature in common: they require attention and are disrupted when attention is drawn away. In the medical context, System Two decisions are those that demand thought, often because a well-established pattern is not recognized and even partial patterns are not obvious or are conflicting. Data may be missing, wrong, and/or are conflicting. Rarely, is there time in a busy clinical environment to engage System Two thinking. While there are strengths and weaknesses associated with System One and System Two thinking, both used at the right times, are crucial for optimal patient care [15]. IT support of the clinical decisions will differ depending on which System is being supported.

At the most fundamental level, CDS should help protect clinicians, and the patients they serve, from cognitive biases. The formal study of cognitive biases was launched in 1974 by Amos Tversky and Daniel Kahneman [16]. Although there are a number of excellent references in the medical literature describing cognitive biases [17–20], a better place for CDS developers to start deciding which apps to build is with a crosswalk of an anti-bias checklist proposed for business decisions [21] into medical decisions. Twelve bias checks are proposed (e.g. check for group-think, check for saliency bias, check for availability bias – which is also called, below, WYSIATI). Whether deciding to build a new manufacturing facility or whether to initiate an invasive surgery for cancer, the decision should seek dissenting

opinions (to avoid groupthink), be certain the recommendation is based on more than the memory of a recent success (to avoid saliency bias) and be certain there isn't a better option not yet considered (to avoid availability bias). CDS developers will do well focusing on as many of the twelve categories of bias as outlined by Kahneman, and his colleagues as is possible.

29.4 Getting the Data in ...

Although it sounds too mundane an issue for a discussion of future HIMS systems, accuracy of the clinical documentation by nurses, physicians and other team members is an under appreciated problem. The cliché of garbage in – garbage out (GIGO) continues to be a major contributor to unsafe systems and renders clinical decision support ineffective, or worse, dangerous. If we are to improve outcomes, we cannot do so if the primary data that clinicians use for their decisions are wrong. Yes, the electronic data are legible, they can be graphed and used in calculations, but unfortunately, all too often the data are erroneous due to omissions, incorrect readings, or disparities between human and medical device readings. Finding evidence for problems with nursing documentation is not hard. One notable study reported on a quality improvement effort in an Italian emergency department and found that triage vital signs were missing in acutely ill trauma victims 10 % of the time even *after* their quality improvement intervention [22]. Mentioning this is not to suggest these adults received poor care – but it does mean as decision support solutions are created with triage vital signs (to, for example, focus the attention of clinicians in a busy environment) that trauma victims might be inappropriately classified by the decision support solution and then these patients might be harmed. Another small survey of trauma resuscitation documentation with a HIMS showed that serial vital signs were not documented a quarter of the time and fully half of the time the Glasgow coma scale and the fluid input-output data were missing [23]. Imagine trying to create a decision support solution for trauma resuscitation without input-output data! However, at this time, HIMSs flow sheets will always be necessary for clinicians to input relevant data that cannot be captured automatically. Level of consciousness, Glasgow coma scale, and capillary refill are extremely important parameters that machines cannot, yet, accurately acquire. The workflow for clinicians (e.g. with voice-based data input) must be augmented and routine use of data error checking routines should be incorporated. HIMSs should prohibit entry of biologically impossible data (e.g. a weight of 874 kg when 87.4 is correct) as well as “implausible” data (e.g. a weight of 87.4 lbs when the patient’s most recent weight was 87.4 kg). These implausible data elements should be based not on population norms for healthy people, but based on that specific, individual patient’s norms and that person’s trends as one would expect to find in a patient-centric system.

Another category of error occurs when other sources of data that are available are just wasted. For example, clinicians periodically do record heart rate – at shorter time intervals when physiological instability demands. Often, EMRs interface with

the physiological monitors and some therapeutic devices to automate data entry into the HIMS system. Doing so requires use of Medical Device Data Systems (MDDS) and these “middle-ware” solutions have well-established regulatory [24] requirements and are commercially available from a number of vendors. MDDS interface with existing bedside medical devices and offer the clinician a time-stamped value for them to verify and then store in the HIMS data tables. Because, for example, the electrocardiogram (ECG) is routinely available as a 240 Hz waveform signal, even verifying and recording the data every 5 min (as is done routinely in by anesthesiologists) means that almost one million data points are lost per patient per hour for just this one signal. That there is value in this single example of wasted data is evidenced by heart rate variability (HRV) analysis and its proposed ability to predict disparate conditions like extubation readiness [25] and subarachnoid hemorrhage [26]. HRV analysis has been shown in a randomized controlled trial to allow early detection of sepsis in low-birth-weight newborns and that with early detection newborn lives are saved [27].

It is important to dig a bit deeper into the challenges of manual data entry biases and the waste of high-fidelity data. A detailed look into the MIMIC II (Multiparameter Intelligent Monitoring in Intensive Care) database work by Hug and Clifford (2007) is in order [28, 29]. Best described on the PhysioNet website (<http://www.physionet.org/>), the MIMIC-II research database has three defining major characteristics: it is publicly and freely available; it encompasses a diverse and very large population of [mostly adult] ICU patients; and it contains high temporal resolution data including lab results, electronic documentation, and bedside monitor trends and waveforms. Developers should also note that because it is built on open source, it allows volunteers to continuously build, refine and share data management and analysis apps. Hug and Clifford (2007) first wanted to determine if the electrocardiogram, systemic arterial blood pressure, and systemic oxygen plethysmography waveform data recorded outside the HIMS (from automated downloads from the physiological monitors) differed from the nursing documented data in the EHR [28]. After developing a filtering algorithm to reject some artifacts, Hug and Clifford found that the monitoring and vital sign data automatically captured as compared to nurse-captured recordings differed not just statistically, but also differed by clinically significant ranges. For example, they found that the least error for each of the four measurements studied occurred on Wednesdays, the highest error rate occurred on Fridays, and errors were most prevalent on the weekend. Diurnal differences were seen as well. And significantly, they detected a significant variation in errors (mean and variance) between data entered by clinicians who logged in anonymously (not now allowed in most mainstream EMRs) compared to those clinicians who were logged in appropriately [28].

In their follow-up paper, Hug and colleagues (2011) analyzed MIMIC II records with both nursing documented data and the automatically captured waveform data not available with the EHR in their 2007 study [29]. For each patient they determined baseline states and then used either the EHR data or the waveform data in an algorithm to predict hypotension. Again in short, the automated data filtered out major artifacts and better predicted episodes of hypotension. Thus, it is not just

interesting that there is variability in the quality of EHR data rather the data quality has potential patient outcome effects. Assuming others confirm these investigator's findings, there are crucial ramifications for restructuring documentation workflows in all areas using continuous vital sign monitors. The takeaway from this research, suggests that clinicians should be viewed as "annotators" of these continuous waveforms rather than arbiter of "truth".

29.5 Federal Drug Administration and Medical Devices Regulation

The above discussion on reliable vital sign capture and recording has comingled two critically distinct regulatory issues. The issue of how to obtain device data (e.g. physiological monitor data and/or therapeutic device data) via an automatic interface for clinician validation and storage within an EHR is regulated as Medical Device Data Systems (MDDS). However, the issue of how to obtain and use the "raw" data from monitors and devices in clinical decision support solutions is not as cleanly described.

In 2011, the Federal Drug Administration (FDA) finalized a rule describing MDDSs stating:

Medical Device Data Systems (MDDS) are hardware or software products that transfer, store, convert formats, and display medical device data. An MDDS does not modify the data or modify the display of the data, and it does not by itself control the functions or parameters of any other medical device. MDDS are not intended to be used for active patient monitoring. [24]

This MDDS rule does not cover the routine use of filtering algorithms and hypotension prediction as illustrated above in the MIMIC II work of Hug and Clifford. The FDA has issued the Food and Drug Administration Safety and Innovation (FDASIA) Health IT Report and delayed a definitive position on the regulation of clinical decision support systems [30]. The report clearly implied a hands-off approach as long as any recommendations passed from a CDS go to a "learned intermediary", meaning any clinician who assumes responsibility for any actions taken [31]. The regulation of an artifact-filtering algorithm used by Hug et al. [29] and any closed-loop CDS solutions will likely be regulated as are current medical devices [31].

Finally, were all these issues not enough as problems in need of solutions in future HIMSSs, actually getting the data acquired remains a challenge. A diatribe asking why consumer electronics are increasing often "plug-and-play" whereas medical devices wallow in a proprietary morass has been done many times by every clinician. As of this writing, there are two well-organized and funded efforts to bring true plug-and-play to medical devices. Most long standing is the Medical Device Plug-and-Play effort spearheaded by Julian Goldman [32, 33]. Work is in part focused on an integrated data environment [33] and dissemination of practical language for health care organizations to use during request for proposals and contracting that would demand vendors support plug-and-play. West Health is also expending substantial effort in the domain of interoperability (see: <http://www.westhealth.org/initiative/our-research>).

So far, the discussion of getting the data in has focused solely on the data collected by clinicians (largely nurses) using standard clinical parameters often facilitated with standard medical equipment. However, there is also a massive untapped trove of data streaming from consumer products. Connected pedometers, scales, pulse oximeters, sphygmomanometers are mainstream consumer devices and are becoming far more common. Location tracking coupled with environmental data and consumer-grade air quality monitors will add previously unavailable data (and might, for example, be useful in an asthma CDS solution). Consumer focused genetic data is also available and although not without problems [34], has broader acceptance than might be guessed [35]. The “Quantified Self” movement has a devoted but still rather small community (see an excellent discussion in the context of a broader vision of the future [36]). Consumers are putting substantial effort into data acquisition, visualization, and analysis. There is some data suggesting the use of consumer device data can improve outcomes [37]. Skepticism should remain high as these early successful reports emerge; the Hawthorne effect can be powerful [38]. However, it seems likely that as sensors improve and become more smoothly incorporated into normal consumer workflow (e.g. by being built into clothing or, maybe, watches) that data quality and availability will improve. Consumer electronics companies, fitness clothing companies and a wide array of startup companies are pushing efforts in this area. HIMSs companies are noticeably absent in this arena, again, making the integration of consumer-generated data with HIMSs data ripe for creation of patient-centric applications that are HIMS vendor agnostic.

Somewhat more established is the use of patient entered data; and “more established” still means efforts less than 10 years old. Consumers have been sharing health stories for millennia if only at the level of chicken soup, garlic necklaces etc. But in early 2006, the social website, *PatientsLikeMe*, opened to the public (<http://www.patientslikeme.com/>). There are hundreds of other such consumer-focused sites but *PatientsLikeMe* has shown some extraordinary successes. From their website, they have about 300,000 members who are recording health data for more than 2,300 conditions and have accumulated 25 million data points [39]. Most extraordinary, if “patientslikeme” is used as a PubMed search term, 46 articles are retrieved. This link between research and social media dynamic is illustrated by this example of lithium and Amyotrophic Lateral Sclerosis patients (ALS). An early (2007) article focused on the story behind an early success understanding on the use of lithium for suppression of symptoms of ALS [40]. In brief, a randomized trial of 44 patients with ALS followed a cohort of 16 who were treated with lithium [41]. No disease progression was reported in this small sized, ALS patient study population when it was published in February 2008. Before then, however, the data had been presented in a conference format and was picked up by the *PatientsLikeMe* social-media community. Working with the investigators from the report, around the time the article was publically released, 116 patients with ALS were already reporting their symptoms within *PatientsLikeMe* while taking doses of lithium much like those reported at the conference. Thereafter, a complete analysis of the *PatientsLikeMe* data was done based on a dataset finalized in February 2010, when 149 patients were eligible for analysis having taken lithium for a part of a year and 78 patients who took lithium for a full year. In short, the analysis showed lithium had no effect on the

progression of ALS (albeit with a low side effect profile) [42]. In an extraordinary waste of resources, a larger randomized controlled trial enrolled patients between 2009 and 2011 and showed the same result as the completely patient entered data from the *PatientsLikeMe* analysis. This randomized clinical trial was published in 2013 [43]. An accompanying editorial [44] simply dismissing the patient entered data analysis should itself be dismissed [45]. To bring this discussion back to point from which it was launched, HIMSS of the future must better acquire accurate data that effectively tells the patient story. But HIMSS must also morph to accommodate this massive consumer-patient data treasure.

29.6 The Future Roadmap Builds on Twenty-First Century Technologies for Vendor-Agnostic Patient-Centric Applications...

If there is a single message readers should take from this chapter it is this: the future of health care information technology is in the dissemination of applications (or “apps”) in a fashion completely analogous to the Android and iOS platform devices. These apps may be as “simple” as gathering data from worn sensors or as “complicated” as combining diagnostic, laboratory, and device data into specific clinical recommendations – even to the point of passing closed-loop instructions to therapeutic devices. Building these apps will require the same entrepreneurial passion and follow through that has gone, and continues to go, into the Angry-Birds-like enterprises. Health care apps, however, must be built with the experience of seasoned clinicians who have the odd combination of out-of-the-box thinking who can entertain the “impossible”, coupled with a 20-something developer partner but with the clinicians’ wisdom to keep the programmer out of clinical trouble.

This concept of the dissemination of apps in a fashion completely analogous to the Android and iOS platform devices for healthcare was first proposed by Ken Mandl and Zak Kohane in an 2009 opinion piece in the *New England Journal of Medicine* [46]. In that article, Mandl and Kohane wrote:

As we seek to design a [HIM] system that will constantly evolve and encourage innovation, we can glean lessons from large-scale information-technology successes in other fields. An essential first lesson is that ideally, system components should be not only interoperable but also substitutable. The Apple iPhone, for example, uses a software platform with a published interface that allows software developers outside Apple to create applications.

Pushing further in a 2012 opinion piece, aptly titled, “Escaping the EHR Trap – The Future of Health IT” [47], Mandl and Kohane urge Health IT vendors to adopt modern technologies wherever possible, and argue that “...” Incentive Programs should not be held hostage to EHRs that reduce...efficiency and strangle innovation. New companies will offer bundled, best-of-breed, interoperable, substitutable technologies...that can be optimized for use in health care improvement. Properly nurtured, these products will rapidly reach the market,




Fig. 29.1 The SMART Platform. Central to the success of the SMART platform is the SMART API that delivers to developers a consistent way to acquire data (from the Container) upon which CDS apps can be built. See text for further details (Reproduced from Mandl et al. [48], with permission of Oxford University Press and the authors)

effectively addressing the goals of ‘meaningful use’, signaling the post-EHR era, and returning to the innovative spirit of EHR pioneers.

What has emerged from Mandl and Kohane’s concepts is the SMART Platform [48] illustrated in Fig. 29.1. A SMART platform enabled HIMS (because, yes, this does describe a post-EHR Health Information Management System) is built on a data container. Writing in 2012, Mandl and colleagues [45] reflected that the SMART data models were still very much a work in progress and limited in scope. The authors further explained that the goal of their data modeling work is not to provide a detailed model for every possible aspect of a patient’s medical history; but rather, to provide highly consistent views for the most common data elements. Because the SMART data models are freely available, this foundational work is accessible to other innovators as well. The SMART model is evolving and now incorporates the Fast Healthcare Interoperability Resources within the Health Level Seven standards organization (or HL7 FHIR pronounced “fire” [49]). As shown in Fig. 29.1, the application program interface (API) also now leverages FHIR. Most remarkable is, within this platform, the wise clinician and 20-something developer dyad can create apps and place them in a public exchange. Local organizations (e.g. hospitals, group practices, payers) can vet applications and then individuals

(clinicians and/or consumers) can further decide which apps to use. Today, SMART enabled applications are being used clinically and in environments beyond those of the original developers [50, 51].

29.7 Using Humans to the Best of Their Ability ...

So now that more and better data is into a vendor-neutral patient-centric platform so the wise clinician/20-something developer dyad can build and disseminate apps, what should they build? According to Licklider, the answer lies in what he called the “man-computer symbiosis” and, therefore, demands an understating of what humans do best. Humans are masters of pattern recognition. In a cognitive task analysis focused on critical care physicians, Fackler and colleagues (2009) identified five broad categories of cognitive activities: pattern recognition; uncertainty management; strategic vs. tactical thinking; team coordination and maintenance of common ground; and, creation and transfer of meaning through stories [35]. Pattern recognition is, however, the prime task after which all other cognitive tasks follow. Additionally, the authors found that while many members of a critical care team used the term ‘pattern’, most physicians could neither define what they meant by ‘patterns’ nor give specific examples of a ‘pattern’. Regardless that clinicians could not be explicit about just what a pattern is, the cognitive task analysis, however, found that pattern recognition did happen in two forms. One pattern was of a complete ‘template’. Asthma is one such complete template based on a minimal history, appearance and breath sounds. A typical template of severe asthma includes the constellation of cues of a patient who is in an upright position, sweaty, speaking in one word answers, exhibiting labored breathing and attentive to his or her own breathing. However, such ‘classic’ complete templates are uncommon.

The second but distinct cognitive task is the real-time merging of pattern fragments (also called ‘packets’) into unique (patient specific) templates. Observed more frequently than identifying a complete template, these packets are recognized as cues that are postulated to be related. It is only through a flexible and dynamic integration of these packets that a complete (or a more complete, but still, partial) template can be created. These templates are context specific. The cue of systolic blood pressure of 80/40 mmHg is quite different in a patient with respiratory failure than in a patient with renal failure, chronic hypertension and altered mental status. Two other cognitive themes from our research [11] are also related and will tie into the decision support discussion below. Critical care clinicians may be uncertain, for example, about missing or possibly erroneous laboratory values. They may be uncertain if a patient’s symptoms do, or do not fit a complete pattern or even partial template. What is often lost in all these discussions, however, is that regardless of this uncertainty, decisions are made, actions are taken, and outcomes may then be equally as uncertain.

Finally, inter-clinician communication is built on pattern recognition but in our study the cognitive theme was identified as creation and use of stories. The term ‘story’ was used explicitly during rounds as senior clinicians often ask, “What’s the patient’s

story?” Reference was also made to the patient’s ‘picture’. Despite differences in terminology, the observational and interview data suggest a common cognitive activity that is closely related to patterns. In both settings, health care teams develop a framework of causal connections and a central theme that tied the various packets of patient data (medical history, test results, etc.) together in a meaningful way.

29.8 Using Machines to the Best of Their Ability ...

So in the man-computer symbiosis, if man is a pattern recognition master, what in a man-computer symbiotic relationship should be the role of computer? In brief, the computer should be a bias-fighter. Pulling again from Kahneman’s book [9], the bias best initially tackled by computers is the “What You See Is All There Is” (or WYSIATI, also called as mentioned above, “the availability bias” [21]). This particular bias is easily understood as its definition is nicely described in its name. It’s equally relevant in what Donald Rumsfeld so famously called “unknown unknowns” or, “you can’t know what you’ve not seen and you don’t even know what you’re missing”. Croskerry (2013) provides an excellent critique of cognitive bias in clinical decision making [52]; and Hough (2013) extends this topic to examine irrationality in decision making throughout our healthcare systems [53]. Again, the reader is referred to the checklist of Kahneman and his colleagues [21].

29.9 Advances in Computer Science and Artificial Intelligent Machines ...

As of this writing (and within my WYSIATI bias) the best potential vendor-agnostic patient-centric decision support solution is exemplified by Watson from IBM. Watson became famous in 2011 when the system crushed the two reigning human champions in *Jeopardy!*. Watson uses a combination of mathematical and computer science techniques applied to massive amounts of unstructured facts. Watson parsed clues of puns and slang and most importantly ranked the confidence of potential answers. Watson meets Licklider’s man-computer symbiosis as it is described on the IBM website (see <http://www.ibm.com/smarterplanet/us/en/ibmwatson/>) as being “a cognitive system that enables a new partnership between people and computers that enhances and scales human expertise.” When Watson approaches a question, “Watson relies on hypothesis generation and evaluation to rapidly parse relevant evidence and evaluate responses from disparate data.” Again, because Watson is handling natural language and most its available data is unstructured text (think textbooks and the medical literature), vast tracks of what to any human is an unknown then become available. Further (and again quoting from the IBM website above), “Through repeated use, Watson literally gets smarter by tracking feedback from its users and learning from both successes and failures. Watson ‘gets smarter’

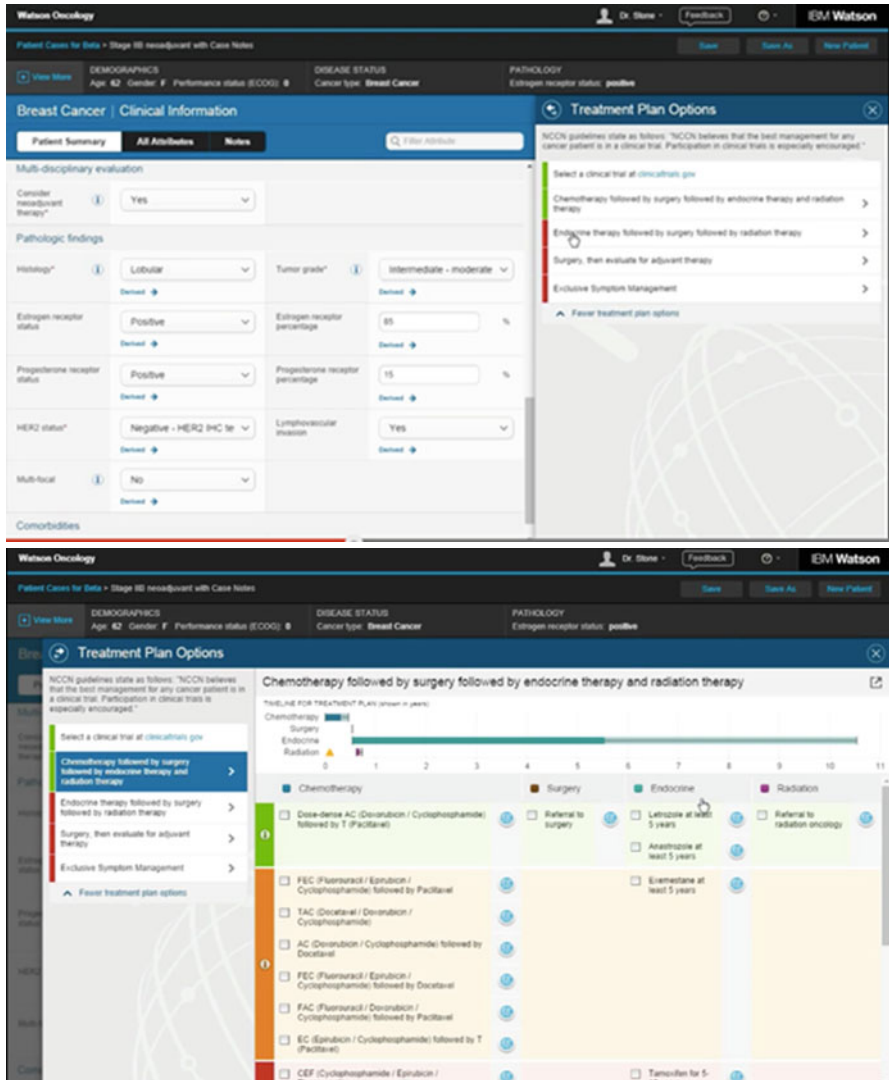


Fig. 29.2 Two screen shots from Watson Oncology. Watson summarizes the EMR data and then suggests the best treatment option based on a diagnosis made solely by the clinicians. Watson is then able to display the available supporting literature including local expert knowledge and patient preferences (Reproduced with permission of IBM)

in three ways: by being taught by its users, by learning from prior interactions, and by being presented with new information.”

The two frames in Fig. 29.2 shows a hypothetical encounter between an expert clinician and Watson for Oncology symbiotic dyad as treatment options are optimized after a diagnosis has been established. Watson for Oncology offers case information, test options, and treatment options. This example shows both the

power of Watson, but also the crucial role of the expert clinician. To be purposely redundant, this must be a symbiotic interaction.

Watson presents the clinical facts as it knows them. These data may come from patient entered data, specific EMR data fields (e.g. age and smoking history) and/or from natural language processed clinical notes. As discussed elsewhere in this chapter these data may be right, wrong, and/or incomplete. The expert clinician must always be cognizant the primary data may be wrong making to avoid the bias of premature closure. Said differently the expert clinician must repeatedly question the primary diagnosis.

In the second frame Watson for Oncology suggests potential treatment plans (along with the hyperlinks to the data and literature supporting the recommendations). Patient preferences can be incorporated into the treatment plan decisions. Yet again, in the symbiotic relationship between the clinician and the computer, it is the human that that will help the patient balance the options.

However, it is important to again emphasize all the work shown in the above example is done to optimize treatment and is not at all focused on the much harder problem of making the correct diagnosis. Diagnostic decision support has long been a focus of clinical informatics [54] but after 40 years of work diagnostic decision support systems remain poorly penetrated [55]. While integration of available data is certainly a problem, even more problematic is the inability of current systems to integrate into the workflow [56]. Said differently, the current diagnostic support systems do not operate in a symbiotic relationship.

Further, although focus on therapy for cancer is laudable and will undoubtedly contribute to adherence to both application of best therapies as well as to patient personalization, there are far more complex problems the approach Watson embodies has the potential to revolutionize. Actually assisting the expert clinician make the diagnosis would yield far more benefit. In only the domain of pediatric critical care, children arrive with a wide array of critical illness that are beyond the full understanding of even the most seasoned clinician [57]. Tests and procedures are done based on precious little data [58]. Diagnostic errors are significant [59].

One need only morph the frames in Fig. 29.2 to imaging the workflow associated with a CDS solution assisting clinicians with diagnostic precision. (And the use of the word “only” in this last sentence is not to trivialize the computer science and engineering necessary.) This CDS solution would first pull data from the EMR and present the available data. Before moving from a case overview workflow to recommendations for testing workflow, many interactions of questions might be necessary. It is the clinician’s role to elicit symptoms and work with the CDS to create as complete a pattern as is possible. As new data becomes available, CDS might assigns new confidences to each potential diagnosis. The CDS solution might “ask” for specific data elements if the solution learns its diagnostic model confidences would be enhanced with additional specific data elements. As more data becomes available, (e.g. family history and a patient’s past medical history) even if no new diagnoses become relevant the confidences the CDS assigns to each diagnosis might fluctuate. With the addition of medication data, side

effects and drug-drug interactions can be added to the problem list. To an experienced clinician even a to rather obvious diagnosis, the CDS has the potential to add diet and life style counseling and/or medication changes to minimize side effects. To be purposefully redundant, in the man-computer symbiosis, will let the clinician be the human touch, the translator of the computer generated confidences, and the overall pattern recognition guru and let the computer make unknown unknowns, known. Again quite cognizant of WYSIATI, there is only one reference to Watson within PubMed [60] and is primarily a descriptive manuscript of work within oncology.

The title of this chapter includes the phrase “a Roadmap”. It is not hyperbole to suggest that opportunities are endless for the wise-clinician/developer dyad to improve patient outcomes. The path should use the SMART platform to build apps that incorporate Watson-like cognitive de-biasing characteristics, and as mentioned in the opening of the chapter, place them within the clinician’s workflow [5]. Not at all parenthetically, as apps are built and evaluated, the outcome analyses must include *patient-centric* measures.

29.10 Data visualization: A Special Instance of Man-Computer Symbiosis

Finally, with good data in both a SMART platform and Watson-consumable forms, data visualization techniques hold a special place in the man-computer symbiosis and decision support efforts of any future HIMS. Again, because humans are pattern recognition masters, presentation of data in a picture is often an effective way

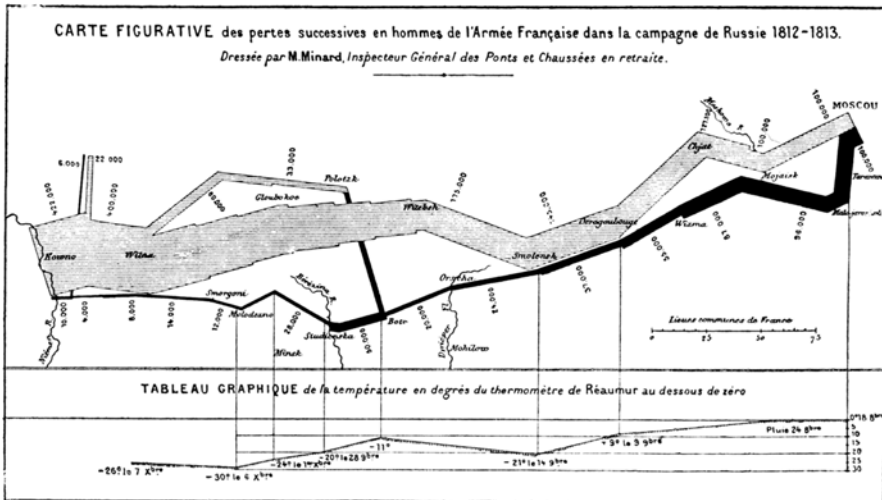


Fig. 29.3 Napoleon’s 1812 Russian Campaign showing devastation of French army as drawn by Charles Minard. Note there are six dimensions of data presented: direction of travel, the global position, time, the soldiers alive and temperature; the latter being responsible for many soldiers’ deaths. Reprinted from Wikimedia Common: <https://commons.wikimedia.org>

for data to be presented to facilitate full or partial template recognition by the clinicians, and also to allow the clinician to see new patterns. Much has been written on data visualization both in the lay press [61] as well as the medical literature. This brief review will highlight only two examples.

Many authors believe the 150-year-old visualization shown in Fig. 29.3 to be the best graphic of all time [62]. The map, drawn by Charles Joseph Minard, shows the losses suffered by Napoleon's army during the Russian campaign of 1812. The size of the top, light-colored band shows the location between the Polish-Russian border and Moscow and the thickness of the band represents the number of soldiers. It is obvious from the diminishing gold bandwidth that the French sustained substantial losses on their march to Moscow. The retreat from Moscow is shown in the black on the bottom, distance is fixed on the “x-axis”, and temperature is added to the graphic in the French troops’ retreat from Moscow. The narrow black line on the bottom right corner captures the devastation of Napoleon’s army at a glance.

Books, too numerous to reference, have been written about data visualization and about medical data visualization (see [63, 64]). The points to make here, in a discussion of the future of HIMSs and a roadmap to the future, are that: (1) there are no current “main-stream” data displays that should be emulated and (2) data visualizations must push not just to present data, but should push into interactive visualizations that allow visual explorations into both patient-centric and population-level data sets with the intent to discover new patient-centric and population-level insights [65]. This latter point will not be explored further here except to point out yet again that the wise-clinician/developer/designer (now) triad must have access to unfettered vendor-agnostic patient-centric data.




Fig. 29.4 Causes of death among the English during the Crimean war as drawn by Florence Nightingale. Reprinted from Wikimedia Common: <https://commons.wikimedia.org>

29.11 Conclusion

I will close this roadmap discussion by highlighting one very and two other rather old visualizations that I believe should serve as “headlights” as the road to the future

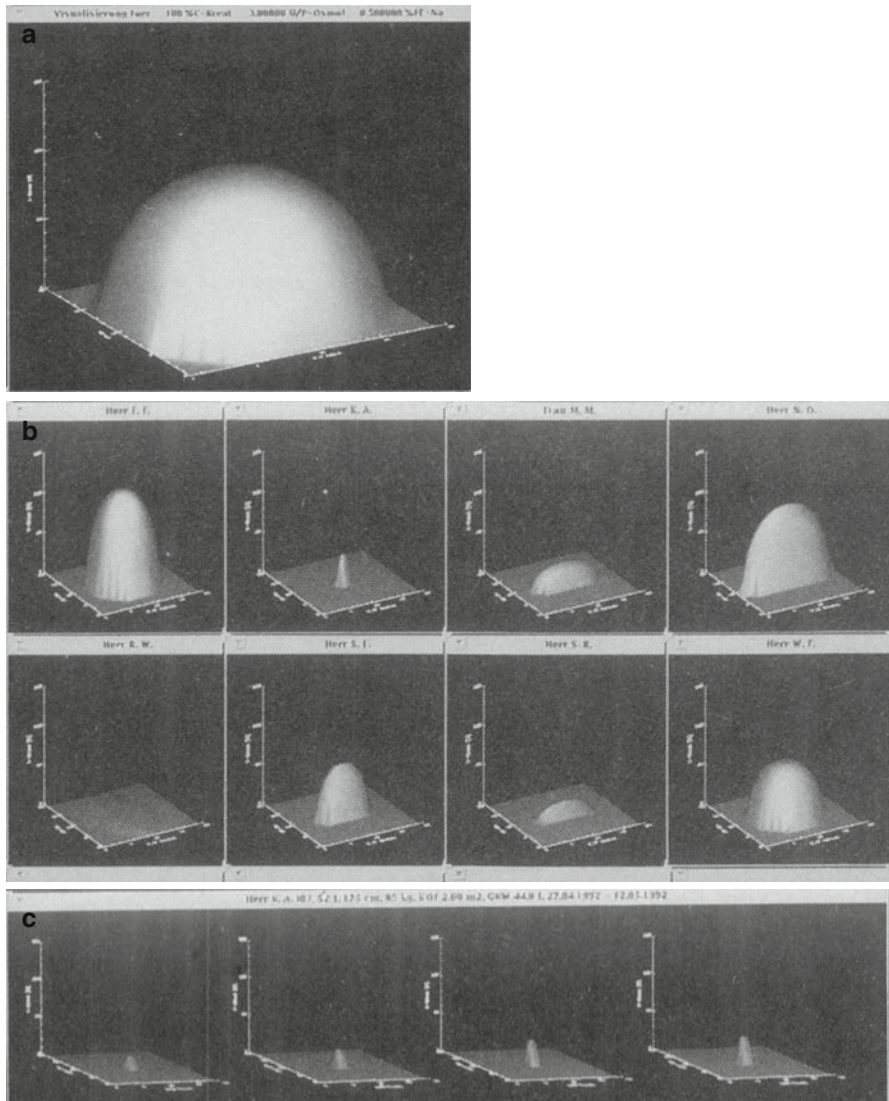
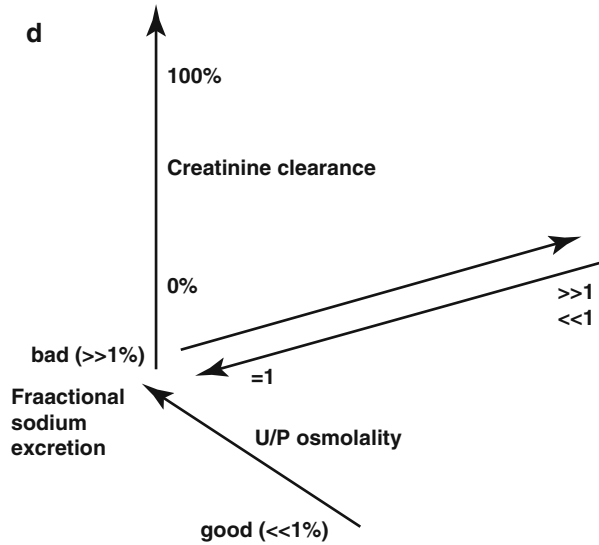


Fig. 29.5 An example of a novel visualization that promotes rapid understanding of complex data (in this case renal function) (Reproduced from Wenkebach et al. [66], with permission of the American Medical Informatics Association, Inc. Frame **a** is normal. Frame **b** shows 8 abnormal patterns. Frame **c** shows a time series of one abnormal pattern. Frame **d** labels the 3 axes)

Fig. 29.5 (continued)



is travelled. Figure 29.4 is now famously about 170 years old from the pen of Florence Nightingale and it still, instantly, tells a story. The distance from the center represents mortality from all causes (in this case during 1 year of the Crimean War). Because the text does not reproduce well here, certainly the reader should ask what is represented by the red, black and blue? The blue wedges represent “Preventable or Mitigable Zymotic Diseases”. The black wedges represent “all other causes”. Only the relatively small wedges actually represent death from “wounds”. Future visualizations should tell so much so “simply”. Enhance the graphic with 2015 available interactions. Allow drill downs into sub-populations such as those dying of cholera, or from a different war, the flu. Drill into the causes of wound mortality to identify patterns so soldier protection can be improved.

In addition, there are two relatively old papers from the medical literature that are worth special mention. First, please look at Fig. 29.5 and assume that the hemisphere drawn in Fig. 29.5a represents normal function. Without knowing anything about the axes, you can then look at the eight patients in Fig. 29.5b and know that none of them are “normal” and that each vary from normal in a different pattern. Finally, in Fig. 29.5c a fourth dimension (time) is added by the sequence of plots for a single patient. Yes, in 2015 this might be animated and additional dimensions may be added with color/shading. But from 1992 when this was published [66], the reader knows that from Fig. 29.5c that function is changing (and either improving or not depending on the direction of time between the pictures). That this was drawn to show renal function is irrelevant (Fig. 29.5d) because the wise-clinician/developer/designer triad, with unfettered vendor-agnostic patient-centric data can create any number of these novel visualizations.

The second paper that the wise-clinician/developer/designer triad with unfettered vendor-agnostic patient-centric data should know about is from Powsner and




Fig. 29.6 A prototype graphic suggested by Powsner and Tufte [67] for the routine display of medical data. To fully appreciate the power of this representation cover the text explanations and realize how much information is transmitted with very few pixels (Reproduced from Powsner and Tufte [67], with permission of Elsevier)

Tufte (1994) also from about 20 years ago [67]. Powsner and Tufte describe design characteristics that can be oversimplified as that can be and in brief, “transmit as much information with as few pixels as is possible.” See Fig. 29.6 as a prototypical representation of serum glucose over time. Note that time on the x-axis is not linear. Much like the example in Fig. 29.5 where pattern recognition is easily supported, here the user quickly knows this particular result that was critically high on hospital admission had not been tested in the previous year, but also that it was normal more than once a year or more before admission. An app designer would do well to heed the design principles these authors espouse. So too, the wise-clinician should encourage the designer to use new visualization techniques.

In conclusion I would like to reproduce the concluding paragraph of the Powsner-Tufte paper and add to it a challenge. Twenty-one years ago, Powsner and Tufte concluded their paper with,

Medical computer systems will soon be able to print a fresh summary for each patient every day. Our proposal for a graphical summary should encourage doctors and nurses to reshape, perhaps re-invent, the medical record *before computer programmers cast institutional convenience into silicon*. Legal and organizational demands for detailed information will not disappear, but these demands need not compromise clinical needs for accessible patient information.

The emphasis is added in the above quote to highlight that Powsner and Tufte saw coming not just the instantiation of 100 year-old, paper-based, data entry and analysis workflows cast into silicon, but also that the codifying these ancient paper-based workflows would compromise accessibility of patient information. Thus the challenge is now even greater, because the roadmap to the future of HIMISs must be disruptive in the every sense of the word [68, 69] The challenge for entrepreneurs plus the wise-clinician/developer dyad or wise-clinician/developer/designer triad is now not merely to undo what the main-stream HIMISs have codified in silicon but to use the data HIMISs record, augment the data as is possible and build apps (including novel visualizations) on the vendor-agnostic patient-centric data. A self-perpetuating cycle must be created as more apps are built and users (again, clinicians and consumers) will clamor for more. As more apps are used they will become more refined. As this cycle spins, there can be optimism that the ultimate goal – improved patient care – will be realized.

References

1. Middleton B, Bloomrosen M, Dente MA, et al. Enhancing patient safety and quality of care by improving the usability of electronic health record systems: recommendations from AMIA. *JAMIA*. 2013;20(e1):e2–8.
2. Institute of Medicine. *Health IT and patient safety: building safer systems for better care*. Washington, DC: The National Academies Press; 2012. p. 1–235.
3. Gordon JE, Leiman JM, Deland EL, Pardes H. Delivering value: provider efforts to improve the quality and reduce the cost of health care. *Annu Rev Med*. 2014;65(1):447–58.
4. Garg AX, Adhikari NKJ, McDonald H, et al. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *JAMA*. 2005;293(10):1223–38.
5. Kawamoto K, Houlihan CA, Balas EA, Lobach DF. Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. *BMJ*. 2005;330(7494):765.
6. Kurzweil R. *Human 2.0*. New Sci. 2005;187(2518):32–7.
7. Fox J, Glasspool D, Patkar V, et al. Delivering clinical decision support services: there is nothing as practical as a good theory. *J Biomed Inform*. 2010;43(5):831–43.
8. Darst JR, Newburger JW, Resch S, Rathod RH, Lock JE. Deciding without data. *Congenit Heart Dis*. 2010;5(4):339–42.
9. Kahneman D. *Thinking, fast and slow*. New York: Macmillan; 2011.
10. Klein G. *Sources of power: how people make decisions*. Cambridge: MIT Press; 1999.
11. Fackler JC, Watts C, Grome A, Miller T, Crandall B, Pronovost P. Critical care physician cognitive task analysis: an exploratory study. *Crit Care*. 2009;13(2):R33.

12. Custer JW, White E, Fackler JC, et al. A qualitative study of expert and team cognition on complex patients in the pediatric intensive care unit. *Pediatr Crit Care Med*. 2012;13(3): 278–84.
13. Shirts BH, Bennett ST, Jackson BR. Using patients like My patient for clinical decision support: institution-specific probability of celiac disease diagnosis using simplified near-neighbor classification. *J Gen Intern Med*. 2013;28(12):1565–72.
14. Amancio DR, Comin CH, Casanova D, et al. A systematic comparison of supervised classifiers. *PLoS ONE*. 2014;9(4):e94137. Shen H-B, ed.
15. Kahneman D, Klein G. Conditions for intuitive expertise. *Am Psychol*. 2009;64(6): 515–51526.
16. Tversky A, Kahneman D. Judgment under uncertainty: heuristics and biases. *Science*. 1974;185(4157):1124–31.
17. Dawson NV, Arkes HR. Systematic errors in medical decision making: judgment limitations. *J Gen Intern Med*. 1987;2(3):183–7.
18. Aberegg SK. Omission bias and decision making in pulmonary and critical care medicine. *Chest*. 2005;128(3):1497–505.
19. Berner ES, Graber ML. Overconfidence as a cause of diagnostic error in medicine. *Am J Med*. 2008;121(5 Suppl):S2–23.
20. Croskerry P, Singhal G, Mamede S. Cognitive debiasing 1: origins of bias and theory of debiasing. *BMJ Qual Saf*. 2013;22 Suppl 2:ii58–64.
21. Kahneman D, Lovallo D, Sibony O. Before you make that big decision. *Harv Bus Rev*. 2011;89(6):50–60. 137.
22. di Martino P, Leoli F, Cinotti F, et al. Improving vital sign documentation at triage: an emergency department quality improvement project. *J Patient Saf*. 2011;7(1):26–9.
23. Bilyeu P, Eastes L. Use of the electronic medical record for trauma resuscitations. *J Trauma Nurs*. 2013;20(3):166–8.
24. Medical Device Data Systems, Medical Image Storage Devices, and Medical Image Communications Devices. 2015. Available at: <http://www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/GeneralHospitalDevicesandSupplies/MedicalDeviceDataSystems/default.htm>. Accessed 22 Mar 2015.
25. Seely AJ, Bravi A, Herry C, et al. Do heart and respiratory rate variability improve prediction of extubation outcomes in critically ill patients? *Crit Care*. 2014;18(2):R65.
26. Park S, Kaffashi F, Loparo KA, Jacono FJ. The use of heart rate variability for the early detection of treatable complications after aneurysmal subarachnoid hemorrhage. *J Clin Monit Comput*. 2013;27(4):385–93.
27. Moorman JR, Waldemar AC, Kattwinkel J, et al. Mortality reduction by heart rate characteristic monitoring in very low birth weight neonates: a randomized trial. *J Pediatr*. 2011;159(6):900–6.e1.
28. Hug CW, Clifford GD. An analysis of the errors in recorded heart rate and blood pressure in the ICU using a complex set of signal quality metrics. *Comput Cardiol*. 2007;34:641–4.
29. Hug CW, Clifford GD, Reisner AT. Clinician blood pressure documentation of stable intensive care patients: an intelligent archiving agent has a higher association with future hypotension. *Crit Care Med*. 2011;39(5):1006–14.
30. FDASIA Health IT Report. 2014:1–34. <http://www.fda.gov/downloads/AboutFDA/CentersOffices/OfficeofMedicalProductsandTobacco/CDRH/CDRHReports/UCM391521.pdf>. Accessed 23 Mar 2015.
31. Karnik K. FDA regulation of clinical decision support software. *J Law Biosci*. 2014;1(2): 202–8.
32. Goldman JM. Solving the interoperability challenge: safe and reliable information exchange requires more from product designers. *IEEE Pulse*. 2014;5(6):37–9.
33. Arney D, Goldman JM, Bhargav-Spantzel A, et al. Simulation of medical device network performance and requirements for an integrated clinical environment. *Biomed Instrum Technol*. 2012;46(4):308–15.
34. Yim S-H, Chung Y-J. Reflections on the US FDA’s warning on direct-to-consumer genetic testing. *Genom Inform*. 2014;12(4):151–5.

35. Carere DA, Couper MP, Crawford SD, et al. Design, methods, and participant characteristics of the Impact of Personal Genomics (PGen) Study, a prospective cohort study of direct-to-consumer personal genomic testing customers. *Genome Med.* 2014;6(12):96.
36. Swan M. Health 2050: the realization of personalized medicine through crowdsourcing, the quantified self, and the participatory biocitizen. *JPM.* 2012;2(4):93–118.
37. Caulfield B, Kaljo I, Donnelly S. Use of a consumer market activity monitoring and feedback device improves exercise capacity and activity levels in COPD. *Conf Proc IEEE Eng Med Biol Soc.* 2014;2014:1765–8.
38. McCambridge J, Witton J, Elbourne DR. Systematic review of the Hawthorne effect: new concepts are needed to study research participation effects. *J Clin Epidemiol.* 2014;67(3):267–77.
39. Available at: <https://www.patientslikeme.com>. Accessed 30 Mar 2015.
40. Frost JH, Massagli MP, Wicks P, Heywood J. How the social Web supports patient experimentation with a new therapy: the demand for patient-controlled and patient-centered informatics. *AMIA Annu Symp Proc.* 2007;6:217–21.
41. Fornai F, Longone P, Cafaro L, et al. Lithium delays progression of amyotrophic lateral sclerosis. *Proc Natl Acad Sci U S A.* 2008;105(6):2052–7.
42. Wicks P, Vaughan TE, Massagli MP, Heywood J. Accelerated clinical discovery using self-reported patient data collected online and a patient-matching algorithm. *Nat Biotechnol.* 2011;29(5):411–4.
43. UKMND-LiCALS Study Group. Lithium in patients with amyotrophic lateral sclerosis (LiCALS). *Lancet Neurol.* 2013;12(4):339–45.
44. Chiò A, Mora G. The final chapter of the ALS lithium saga. *Lancet Neurol.* 2013;12(4):324–5.
45. Wicks P, Vaughan T, Heywood J. Subjects no more: what happens when trial participants realize they hold the power? *BMJ.* 2014;348(jan28 9):g368.
46. Mandl KD, Kohane IS. No small change for the health information economy. *N Engl J Med.* 2009;360(13):1278–81.
47. Mandl KD, Kohane IS. Escaping the EHR trap – the future of health IT. *N Engl J Med.* 2012;366(24):2240–2.
48. Mandl KD, Mandel JC, Murphy SN, et al. The SMART Platform: early experience enabling substitutable applications for electronic health records. *J Am Med Inform Assoc.* 2012;19(4):597–603.
49. Introducing HL7 FHIR®. hlorg. 2014:1–2. <http://hl7.org/implement/standards/fhir/fhir-summary.pdf>. Accessed 29 Mar 2015.
50. Bosl W, Mandel J, Jonikas M, Ramoni RB, Kohane IS, Mandl KD. Scalable decision support at the point of care: a substitutable electronic health record App for monitoring medication adherence. *Interact J Med Res.* 2013;2(2):e13.
51. Klann JG, McCoy AB, Wright A, Wattanasin N, Sittig DF, Murphy SN. Health care transformation through collaboration on open-source informatics projects: integrating a medical applications platform, research data repository, and patient summarization. *Interact J Med Res.* 2013;2(1):e11–8.
52. Croskerry P. From mindless to mindful practice – cognitive bias and clinical decision making. *N Engl J Med.* 2013;368(26):2445–8.
53. Hough D. Irrationality in health care: what behavioral economics reveals about what we do and why. Stanford: Stanford Economics and Finance, Stanford University Press; 2013.
54. Barnett GO, Cimino JJ, Hupp JA, Hoffer EP. DXplain. An evolving diagnostic decision-support system. *JAMA.* 1987;258(1):67–74.
55. Bond WF, Schwartz LM, Weaver KR, Levick D, Giuliano M, Graber ML. Differential diagnosis generators: an evaluation of currently available computer programs. *J Gen Intern Med.* 2011;27(2):213–9.
56. Henderson EJ, Rubin GP. The utility of an online diagnostic decision support system (Isabel) in general practice: a process evaluation. *JRSM Short Rep.* 2013;4(5):31.
57. Fackler JC, Wetzel RC. Critical care for rare diseases. *Pediatr Crit Care Med.* 2001;3(1):89–90.

58. Fackler J, Lehmann HP, Wetzel RC. Critical care for rare diseases (and procedures). *Pediatr Crit Care Med*. 2015;16(3):297–9.
59. Cifra CL, Jones KL, Ascenzi JA, Bhalala US, Bembea MM, Newman-Toker DE, Fackler JC, Miller MR. Diagnostic errors in a PICU. *Pediatr Crit Care Med*. 2015;16(5):468–76.
60. Doyle-Lindrud S. Watson will see you now: a supercomputer to help clinicians make informed treatment decisions. *CJON*. 2015;19(1):31–2.
61. Carey B. Learning to see data. *NY Times*. 2015;164(56,820):SR1.
62. Worth a thousand words. *The economist*. 2007. Available at: <http://www.economist.com/node/10278643>.
63. Wurman RS. *Understanding healthcare*. Newport: Top; 2004.
64. Holzinger A. *Biomedical informatics*. New York: Springer International Publishing; 2014. doi:10.1007/978-3-319-04528-3.
65. Rind A. Interactive information visualization to explore and query electronic health records. *FNT Human Comput Interact*. 2013;5(3):207–98.
66. Wenkebach U, Pollwein B, Finsterer U. Visualization of large datasets in intensive care. *Proc Annu Symp Comput Appl Med Care*. 1992:18–22.
67. Powsner SM, Tufte ER. Graphical summary of patient status. *Lancet*. 1994;344(8919):386–9.
68. Christensen CM, Bohmer R, Kenagy J. Will disruptive innovations cure health care? *Harv Bus Rev*. 2000;78(5):102–12. 199.
69. Hwang J, Christensen CM. Disruptive innovation in health care delivery: a framework for business-model innovation. *Health Aff (Millwood)*. 2007;27(5):1329–35.