

Toward a science of learning systems: a research agenda for the high-functioning Learning Health System

RECEIVED 10 May 2014
 REVISED 18 August 2014
 ACCEPTED 23 September 2014
 PUBLISHED ONLINE FIRST 23 October 2014



Charles Friedman¹, Joshua Rubin¹, Jeffrey Brown², Melinda Buntin³, Milton Corn⁴, Lynn Etheredge⁵, Carl Gunter⁶, Mark Musen⁷, Richard Platt², William Stead³, Kevin Sullivan⁸, Douglas Van Houweling¹

ABSTRACT

Objective The capability to share data, and harness its potential to generate knowledge rapidly and inform decisions, can have transformative effects that improve health. The infrastructure to achieve this goal at scale—marrying technology, process, and policy—is commonly referred to as the Learning Health System (LHS). Achieving an LHS raises numerous scientific challenges.

Materials and methods The National Science Foundation convened an invitational workshop to identify the fundamental scientific and engineering research challenges to achieving a national-scale LHS. The workshop was planned by a 12-member committee and ultimately engaged 45 prominent researchers spanning multiple disciplines over 2 days in Washington, DC on 11–12 April 2013.

Results The workshop participants collectively identified 106 research questions organized around four system-level requirements that a high-functioning LHS must satisfy. The workshop participants also identified a new cross-disciplinary integrative science of cyber-social ecosystems that will be required to address these challenges.

Conclusions The intellectual merit and potential broad impacts of the innovations that will be driven by investments in an LHS are of great potential significance. The specific research questions that emerged from the workshop, alongside the potential for diverse communities to assemble to address them through a ‘new science of learning systems’, create an important agenda for informatics and related disciplines.

Key words: learning health system, population health, system science, research agenda

BACKGROUND AND SIGNIFICANCE

Vision of a Learning Health System

The capability to share data—and harness its potential to generate knowledge rapidly and inform decisions—can have transformative effects on complex systems that produce goods and provide services. Within specific sectors of the economy, individual organizations and collaborating groups have leveraged their data to increase productivity, gain competitive advantage, and revolutionize business models.^{1–3} Domains ranging from transportation to agriculture, to environmental sciences, to supply chains, to military services and intelligence have demonstrated the capacity and need to derive significant benefits from leveraging data, systems, and human interconnectedness on ever-increasing scales.⁴

However, attempts to employ these types of approaches to realize transformative impacts for the most challenging societal

problems—including health and education—have not enjoyed similar success.⁵ While the health sector can point to examples of large clinical research networks, as well as increasing adoption of electronic health record (EHR) systems and other information technologies, this sector has not undergone the type of IT-enabled transformation visible across other industries.⁶

One widely-held vision for realizing the transformations previously described, at significant scale and scope in the health domain, is generally known as the Learning Health System (LHS). A series of publications from the Institute of Medicine (IOM)^{7,8} has described the LHS and documented the need for it. The IOM defines the LHS as a vision for an integrated health system “... in which progress in science, informatics, and care culture align to generate new knowledge as an ongoing, natural by-product of the care experience, and seamlessly refine and deliver best practices for continuous improvement in

Correspondence to Dr Charles P Friedman, Department of Learning Health Sciences, University of Michigan, 1111 E. Catherine St., Ann Arbor, MI 48109, USA; cpfried@umich.edu

©The Author 2014. Published by Oxford University Press on behalf of the American Medical Informatics Association.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com
 For numbered affiliations see end of article.

health and healthcare.”⁷ Though articulated in various forms, the underlying concept is straightforward: harness the power of data and analytics to learn from every patient, and feed the knowledge of “what works best” back to clinicians, public health professionals, patients, and other stakeholders to create cycles of continuous improvement.

Consistent with the characteristics of a continuously learning system as articulated by the IOM in 2012,⁹ the LHS can also be thought of as common infrastructure, governance, and incentive structures coupled with shared values that enable such a health system and culture. The LHS also promises to empower a new era of personalized medicine.¹⁰ Realizing the promise of personalized medicine requires a key component of what the LHS enables—“knowing what works, understanding why it works, learning for whom it works, and applying that knowledge to address patient needs. . . .”¹¹

The Office of the National Coordinator for Health Information Technology (ONC) set achievement of rapid learning as the pinnacle goal of its 5-year strategic plan (2011–2015).¹² In 2014, ONC identified the national-scale LHS as its 10-year strategic goal and stated that, “This ‘LHS’ should also enable lower health care costs, improved population health, truly empower consumers, and drive innovation.”¹³

The LHS imperative

The imperative to achieve an LHS is anchored in a number of realities of the current health system that render its status quo unsustainable.¹⁴ The USA spends roughly 18% of its gross domestic product (GDP) on healthcare, and, by some projections, could almost double that share by 2040.¹⁵ Healthcare spending per capita by the USA greatly exceeds that of other ‘first world’ countries,¹⁶ while underperforming these nations on many indices of population health status. As one study notes, “While the U.S. performs well on breast and colorectal cancer survival rates, it has among the highest rates of potentially preventable deaths from asthma and amputations due to diabetes, and rates that are no better than average for in-hospital deaths from heart attack and stroke”.¹⁷ Other studies suggest that as much as \$750 billion annually, approximately 25% of total expenditures, is literally wasted on administrative inefficiencies, fraud, abuse, and treatments that do not work.^{8,18} Even with such spending, it has long been recognized that patients do not receive over 45% of recommended care¹⁹ and quality is improving slowly.²⁰ Such urgency itself underpins a series of LHS-related IOM reports²¹ dating back to the original ‘The Learning Healthcare System’ vision report of 2007.²²

There is a growing recognition of the US healthcare system’s inability to routinely study its own behaviour; an LHS would provide such capability, and would significantly address many of the current challenges faced by the system. For example, researchers required 4 years to recognize that a commonly prescribed drug was leading to tens of thousands of deaths.²³ In the presence of an LHS infrastructure to support routine surveillance leveraging timely data from a population the size of the USA, this safety signal could have been detected in approximately 2 months.²⁴ Researchers who developed and tested

analytic methodologies for learning from EHR data, utilizing the UK’s General Practice Research Database (GPRD), highlighted the importance of learning at large scale. These researchers concluded that, “Ideally future databases should be much larger than GPRD, which includes about eight million patients. On the basis of our work to date, we estimate that 40–50 million patients are needed for the breadth of future studies we can envisage.”²⁵

The growing recognition of the importance of an LHS is seen in an increasing number of publications, amplifying the calls from the IOM.²¹ For example, a 2013 editorial in the *New England Journal of Medicine* cites the LHS as one of three essential steps needed to reduce healthcare’s ‘GDP Footprint’,¹⁴ while another has challenged all academic health centers to become LHSs, so they can be drivers of needed change in health research and health care delivery.²⁶ Further illustrating the emerging national imperative, the entire July 2014 issue of *Health Affairs* was devoted to the potential of ‘big data’ to transform healthcare and health,²⁷ including an extensive discussion of a new rapid-learning agenda.²⁸

Overall, realizing the LHS on a national scale by harmonizing and synergizing multiple efforts has taken on many features of a “Big Hairy Audacious Goal”,²⁹ which “changes the time frame and simultaneously creates a sense of urgency. . . .”³⁰

Pathway and progress toward an LHS

A foundation for the LHS is emerging as the nation’s healthcare system becomes increasingly digital.³¹ This is due in part to an estimated over \$30 billion federal investment to promote meaningful use of health information technology through the HITECH Act provisions of the American Recovery and Reinvestment Act of 2009.³² Averaging several indices, the nation’s health system is approximately 35% digital now but is expected to be 80% digital by 2019.³³ Interoperability requirements built into HITECH will make significant but imperfect progress toward standardizing data stored in EHRs.

Other foundational elements of a national LHS include organizations that, within their own boundaries, have become LHSs. While several organizations have made significant progress toward learning status, Kaiser Permanente, Geisinger Health System, and the US Veterans Administration stand out as examples of institutions that are already harnessing the power of data to improve the health of the populations they serve.^{34,35} In addition, federated learning networks among independent organizations have formed to share data and harvest the benefits of learning from it, with the HMO Research Network, the Vaccine Safety Datalink, the Food and Drug Administration (FDA) Mini-Sentinel, and CancerLinQ^{7,36} as key examples. These developments both document the benefits to be derived from a nationally-scaled LHS and suggest that a national-scale LHS is technically and organizationally achievable. A recent paper examining four emerging data networks concludes: “these four programs [are] examples of the first stage in the development of a shared national big-data resource that leverages the investments of many agencies and organizations for the benefit of multiple networks and users”.³⁷

From these developments, it is evident that achievement of a national-scale LHS will not be the work of a single organization, stakeholder group, or governmental entity. Rather, it is anticipated that the LHS will require active participation of and cooperation among multiple and diverse stakeholders, nationwide and ultimately globally.³⁸ While there is increasing national focus in the USA, the LHS is increasingly a global phenomenon. There is increasing collaboration among international and multinational initiatives including TRANSFoRM (which “aims to develop the technology that facilitates a learning healthcare system” for the European Union),³⁹ with explicit recognition that key LHS enablers such as standards should be developed in ways conducive to global harmonization.⁴⁰

The LHS research agenda

The LHS is a consummate challenge for researchers across a range of scientific disciplines. Progress at each stage of development will require the best possible answers to an enormously wide range of questions.

The scope of these challenges was first suggested at a 2010 IOM workshop that focused on the digital infrastructure for the LHS. Participants in this workshop identified the LHS as an ultra-large scale, ultra-complex, cyber-socio-technical system,⁷ recognizing that at a national and ultimately global scale, the LHS is a massive system of systems. That same year, in a separate study addressing similar issues, a group of researchers wrote: “The U.S. urgently needs a major initiative to develop software and systems engineering foundations for a national-scale health information network (NHIN). The NHIN will be an ultra-large-scale (ULS) system. An ULS systems perspective therefore must guide these activities.”⁴¹

It follows that achieving the LHS vision requires more than merely addressing difficult engineering challenges; profound and difficult socio-technical challenges will need to be addressed by bringing together a wide-range of disciplines. Moreover, the LHS is not a ‘build out’ project. Along the way, many challenges—including government policies, competition among healthcare providers, rivalries among technology vendors and developers, concerns related to balancing privacy and security against the value created by secure sharing of data—will need to be addressed. The LHS will need to be sustainable through a sound business and governance model. The complex trans-disciplinary work ahead may build upon emerging specializations of the past decade, including “implementation and integration sciences. . . the theory and methods necessary to tackle complex societal issues and problems”.⁴²

METHODS: THE WORKSHOP

In this spirit, the National Science Foundation convened a 2-day invitational workshop to identify the fundamental scientific and engineering research challenges to achieving a national-scale LHS. The workshop was planned by a 12-member committee and ultimately engaged 45 prominent researchers spanning multiple disciplines over two full days in Washington, DC on 11–12 April 2013. While fully recognizing that the LHS is ultimately a global challenge and that the findings would

have potential to generalize across national boundaries, support from a US national agency directed a national focus for the workshop.

The participants were challenged to view the LHS as a ‘system after next’.⁴³ As such, the LHS can succeed only by creating novel combinations of role, process and technology. This must occur by working back from the future, not by figuring out how to fix the various problems with a current system that fails to learn rapidly, routinely, and at scale.

The research questions, that were the focal product of the workshop, developed through two rounds of breakout group deliberations. In each round, the participants divided into four groups with different composition. In the first round, each group examined one of four use cases describing the operation of a high-functioning LHS at scale. The use cases separately addressed LHS operation in support of health care quality, population health, personal health, and biomedical research and illuminated what learning could mean—and how it could ultimately improve health—in each of these contexts. Each group identified the research challenges that must be addressed to realize a very high-functioning LHS capable of executing its assigned use case. The LHS use cases may be found at <http://healthinformatics.umich.edu/lhs/nsfworkshop>.

In the second round, each group focused on one of four broad system-level requirements for the high-functioning LHS:

1. An LHS trusted and valued by all stakeholders
2. An economically sustainable and governable LHS
3. An adaptable, self-improving, stable, certifiable, and responsive LHS
4. An LHS capable of engendering a virtuous cycle of health improvement.

Each group identified the research challenges that must be addressed to realize a very high-functioning LHS capable of meeting the assigned requirement.

As part of the plenary session closing the workshop, each participant wrote the research challenge question stemming from the 2-day meeting that he/she felt was most important and that he/she could envision himself/herself being personally involved in addressing.

Following the workshop, the planning committee compiled, refined, and organized the research challenges into the form of questions. The questions were circulated back to all participants for review and comment, which resulted in further refinement of the questions. A modification of the four system-level requirements used for the second round of the workshop’s group deliberations formed the basis for organizing the key findings. Ultimately, 106 questions were generated and they were organized into 4 categories and 19 subcategories. In reviewing the report, the 19 subcategories themselves (box 1) may paint the most comprehensive picture of the breadth and nature of the research challenge questions generated through the workshop methodology. The challenges and questions that follow form a broad scientific agenda for realizing a high-functioning LHS. Further, in addition to the granular level of findings

Box 1: Learning Health System (LHS) research challenge questions: categories and subcategories

1. System requirement 1. An LHS trusted and valued by all stakeholders
 - Defining, measuring, and tracking confidence and trust
 - Building confidence and trust in the data inputs
 - Building confidence and trust in the process
 - Generating value while building confidence and trust in the outputs
 - Building confidence and trust in the system as a whole
2. System requirement 2. An economically sustainable and governable LHS
 - Private sector incentives and markets
 - The role of the public sector in the LHS
 - Conceptualizing value in the context of the LHS
 - The LHS and the healthcare/health system as a whole
3. System requirement 3. An adaptable, self-improving, stable, certifiable, and responsive LHS
 - An adaptable LHS
 - A self-improving LHS (that builds trust in the system)
 - A certifiable LHS
 - A stable LHS (through trust and confidence)
 - A responsive LHS
4. System requirement 4. An LHS capable of engendering a virtuous cycle of health improvement
 - Rapidly creating knowledge that engenders learning
 - Communication within the LHS: What mechanisms will enable communication of methods used and results obtained, in actionable forms, to all stakeholders with interest in the results?
 - Building a smart system: How can the LHS become smart enough?
 - Learning about learning, research about research
 - Key questions specific to health improvement

of the research challenge questions themselves, there was a second level of transcendent findings, illuminating themes interwoven into a new science of learning systems anchored in the socio-technical challenges identified.

RESULTS: THE RESEARCH CHALLENGES AND QUESTIONS

We provide below an illustrative sample of the research challenge questions generated as a formal result of the workshop. This illustrative sample was selected from among the full set of

questions (n = 106). The representative sample below includes at least one question from all subcategories, containing a mix of high-level and specific questions, with an intent of conveying the perspectives deriving from the diverse disciplines represented. A full listing is available in the report itself (<http://healthinformatics.umich.edu/lhs/nsfworkshop>).

Requirement 1: an LHS trusted and valued by all stakeholders

- What methods will be needed to measure confidence, trust, and trustworthiness?
- What approaches will promote monitoring of confidence and trust at varying levels of scale, and thus enable improvement of the LHS over time?
- Can the potential of a dataset to generate valid knowledge (fitness for use) be computed from the dataset itself? What metadata are required?
- How can we analyze motives and incentives for security breaches and attacks on the LHS, and best structure governance and security policies and mechanisms to guard against them? What is the relationship between the real and perceived security of the LHS and stakeholders' levels of confidence and trust in it?
- What value metrics will assess the magnitude of improvement in health outcomes and cost-efficiency, as well as social and behavioral impacts, associated with LHS activities?
- How can the LHS be designed, engineered, and operated as a self-defending and self-repairing system for purposes of protecting individual and institutional privacy and the integrity of data knowledge against malicious attack and accidental disclosure?

Requirement 2: an economically sustainable and governable LHS

- What value is the LHS uniquely capable of producing, for which stakeholders, and how can this value be identified, measured, and incentivized? How can the science of networks, markets, game theory, and mechanism design inform the design and operation of the LHS?
- How can predictive models and empirical studies address what ingredients, essential to standing up and sustaining the LHS, have no private rationale for funding or are unlikely to be funded privately, and therefore should be considered for catalysis stemming from public funding? How can lessons from other rapid learning systems and ultra-large scale systems in other sectors be applied to thinking about incentives and markets vis-à-vis the LHS?
- How can predictive models and empirical studies address the relationship between incentive structures

and data sharing and other key activities essential to the function of the LHS that will advance the public interest?

Requirement 3: an adaptable, self-improving, stable, certifiable, and responsive LHS

- What features of the LHS, as a socio-technical ecosystem, will best position it to incorporate innovation in data sources that include real-time monitors, sensors, and devices intended to augment bodily functions and directly improve a person's health?
- How do we combine data generated in learning-from-doing and from observational data with careful experiments to be sure knowledge is scientifically solid and valid for stated purposes, but also as widely useful as possible?
- What analysis must be conducted, what evidence must be gathered, and what arguments must be constructed and maintained over time, about the development, design, and operation of the system and about the environment in which it operates to enable system certification for safety, performance, and other critical system properties?
- How do we make data sufficiently self-describing so that, for example, the system might be able to identify, without human intervention, data that are most relevant to addressing a question?
- How can we understand and develop mechanisms to balance the tradeoff between speed (rapidity) and accuracy?

Requirement 4: an LHS capable of engendering a virtuous cycle of health improvement

- What existing analytical and inferential methods serve the needs of the LHS, and what needs of the LHS require new methods? In particular, in the context of an LHS at scale, what methods are needed to address bias, confounding, propagation of erroneous information, spurious correlations, and other potential sources of mis-learning in the LHS?
- What rigorous methods for data description and new data quality metrics will capture the semantics that enable rapid learning from (big) heterogeneous data streams? What approaches will describe fitness for use within the context of a specific purpose? For example, what are the semantics for characterizing data and identifying and describing bias?
- How do we develop ways to communicate generated results, and surrounding uncertainties, to others who may wish to replicate (or build upon) the work done, as well as to the general public?
- How can the LHS become smart enough to detect attempts to answer a question that is not answerable with the resources in the system?

- What are the best approaches to measuring and understanding the predictive value of models resulting from big data?
- What new methods are needed to determine the value of knowledge generated by the LHS and resulting actions?
- To support health care quality improvement, how can the LHS develop a complex multidimensional model of a person's future health status and communicate it to them or their trusted healthcare provider to promote their health and wellbeing?

DISCUSSION: THE LHS AND AN EMERGING SCIENCE OF CYBER-SOCIAL LEARNING SYSTEMS

The challenges and specific research questions—organized into four areas corresponding to requirements a high-functioning LHS must meet—constitute one level of expression of the workshop's findings. At a higher level, the findings in their totality and their multiple interactions generate a broader perspective, suggesting that rising to the challenge of the LHS may require a novel emergent science of large-scale learning systems best seen as an evolution from the science of information systems, through a science of cyber-physical systems, and ultimately to a science of cyber-physical-social ecosystems. While these concepts were explicitly discussed, the workshop did not describe this new science in detail, but rather was suggestive of it, with the LHS providing the real world, imperative, driving problem that brings this science into sharper focus. A second workshop product was therefore the suggestion of this new perspective, the applicability of which may transcend the health domain that was the original focus of this workshop.

Sciences have been developed and refined to study information systems, systems responsible for the digital representation, processing, storage, transmission, and use of information enabled by advances in computer science and engineering, as well as other fields. A next generation science of cyber-physical systems has also emerged, driven by advances in sensors, the ongoing miniaturization of digital computers, the wireless networking of edge devices to massive cloud computing systems, computer control of robotic machinery, and machine learning and artificial intelligence. After considering the challenges associated with realizing a national-scale LHS, the workshop participants came to recognize cyber-social ecosystems as a next step in this evolution.

A national-scale LHS will have to be understood and designed as such a cyber-social ecosystem: a large-scale, decentralized, *human-intensive*, cyber-catalyzed and cyber-supported information processing system. The system as a whole—not just the digital infrastructure, but also networks of people and institutions—will have to be understood not just as users of a technological infrastructure, but also as parts of the information system itself. So, for example, trust in the digital infrastructure is necessary to induce privacy-sensitive citizens to agree to permit information flows that are critical enablers of computations essential to the beneficial operation of the

system. Institutions with commercial interests in the data that they generate will have to be incentivized or otherwise motivated to transmit such data under certain conditions.

From this perspective, the LHS cyber-infrastructure must be understood as a platform that empowers multiple and diverse stakeholders to individually and collectively drive innovation across the healthcare ecosystem. By serving as a platform that enables *data as a service* to facilitate and incentivize data sharing and data use, the LHS will underpin a host of unforeseeable innovations in data-, knowledge-, and evidence-driven healthcare, bio-surveillance in the public interest, and health-related research and development. Just as no one involved in the creation of the Internet in the 1960s could have reliably predicted the ultimate emergence of Google, Amazon, Facebook, or Twitter, no one today can predict what the LHS will catalyze in the decades to come. Yet, just as some original visionaries saw the emergence of personal computing and an “intergalactic computer network” (that grew into the world wide web)⁴⁴ ahead of their times, establishing a robust foundational framework for the LHS, with many elements articulated in the aforementioned IOM reports, will enable analogous types of emergent future possibilities, growth, and evolution. If it is properly designed, the LHS can be expected to have profoundly catalytic effects on future innovation for health, healthcare, and the economy.

Achieving a learning system will require discovery of methods to loosely couple the parts of this complex ecosystem in ways that enable them to align with one another over time. Just as the Internet easily survives the loss of component systems due to loose coupling, and just as it fosters the emergence of completely unforeseen functions, so an architecture enabling loose but influential coupling in the health ecosystem will allow each component of the system to participate in the evolving network while maintaining its own autonomy, pursuing its own interests, and optimizing its own performance, while also being able to rely on trusted other parts of the system doing their parts in concert.

In sum, the findings of this workshop identify a host of deeply important, difficult, and unresolved research issues at the intersection of computer science, software and systems engineering, the behavioral and social sciences,⁴⁵ healthcare, clinical medicine,⁴⁶ research and statistical methodologies,²⁵ and public health. The authors and the workshop participants recognized that addressing the deep scientific questions related to the LHS—many of which will be anchored in the new envisioned science of cyber-social ecosystems—will require not only collaboration among multiple and diverse stakeholder types, but will also require research, learning, and innovation to occur at the intersection of multiple and diverse disciplines. No one discipline alone has the tools or knowledge requisite to attack this multifaceted sociotechnical challenge. It follows that neither the research community currently organized around the National Science Foundation (NSF) alone, nor around the National Institutes of Health (NIH) alone, has the requisite intellectual capabilities to empower this emerging scientific perspective comprehensively.

The LHS requires a new and significant crossing of capabilities that today are present only in these largely disconnected communities. Realizing this goal will require the NSF to continue to expand its embrace of health and healthcare as a driving domain for fundamental research, as it will require the NIH to expand its embrace of computing as a transformative power for biomedical research and healthcare innovation and delivery. The potential benefits are manifold. If the LHS is successful, there will be important lessons for how to leverage large amounts of real-world data, mechanisms for learning from such data, feedback components aimed at mobilizing lessons learned to inform decisions and actions, and cyber-social ecosystems—entailing networks of computers, machines, people, and organizations—to improve performance and bring about transformation in many other sectors outside of health.⁴⁷ Therefore, additional mechanisms, of which this workshop was one starting point, and resources to bring together the largely disconnected research communities, will be needed to underpin an interdisciplinary, multi-stakeholder journey to realize such promise.

The specific research questions that emerged from the workshop, alongside the potential for diverse communities to assemble to address them through a “new science of learning systems”, were deeply exciting to the workshop participants.

CONTRIBUTORS

The authors have participated to varying extents in the preparation of this manuscript, but all have met the four authorship criteria listed below: substantial contributions to the conception or design of the work; or the acquisition, analysis, or interpretation of data for the work; drafting the work or revising it critically for important intellectual content; final approval of the version to be published; agreement to be accountable for all aspects of the work in ensuring that questions related to the accuracy or integrity of any part of the work are appropriately investigated and resolved.

FUNDING

This work was supported by National Science Foundation grant number IIS -1249354.

COMPETING INTERESTS

Authors CF, JR, and JB are members of the interim steering committee for the Learning Health Community, a grass-roots, not-for-profit entity seeking to promote a national scale learning health system. Author MM serves on the scientific advisory boards of Apixio, HealthFidelity, and Kyron companies which have granted him stock options.

PROVENANCE AND PEER REVIEW

Not commissioned; externally peer reviewed.

DATA SHARING STATEMENT

A set of documents which reports the findings of the workshop in much greater detail, including all 106 research questions, can be accessed (<http://healthinformatics.umich.edu/lhs/nsfworkshop>).

REFERENCES

1. Ayres I. Super crunchers: why thinking-by-numbers is the new way to be smart. New York, NY: Bantam Books, 2007: 1–260.
2. Davenport T, Harris J. Competing on analytics: the new science of winning. Boston, MA: Harvard Business School Publishing Corporation, 2007:1–218.
3. Tapscott D, Williams A. Wikinomics: how mass collaboration changes everything, expanded edition. New York, NY: The Penguin Group, 2008:1–351.
4. Northrop L. Presentation: “Ultra-Large-Scale Systems: Scale Changes Everything”. Software Engineering Institute. Carnegie Mellon University. SMART Event. 6 March 2008. http://resources.sei.cmu.edu/asset_files/presentation/2008_017_001_23733.pdf
5. See e.g.: Cutler, D. Searching for Health Care’s Entrepreneurial Spirit. Harvard Business Review. HBR Blog Network. 1 March 2011. <http://blogs.hbr.org/2011/03/searching-for-health-cares-ent/>
6. Clark A, Hanna K, Parrish G. Still thinking research: strategies to advance the use of electronic health records to bridge patient care and research. Washington, DC.: FasterCures, 2011:1–44.
7. Institute of Medicine (IOM). Digital infrastructure for the learning health system: the foundation for continuous improvement in health and health care. In: Grossman C, Powers B, McGinnis JM, eds. *The foundation for continuous improvement in health and health care workshop series summary*. National Academies Press, 2011: 1–311.
8. Smith M, Saunders R, Stuckhardt L, et al. In: Smith M, Saunders R, Stuckhardt L, McGinnis J. eds. *Best care at lower cost: the path to continuously learning health care in America*. 1st edn. Washington, DC: National Academies Press, 2012:1–450. <http://www.iom.edu/bestcare>
9. Institute of Medicine (IOM). Table of Characteristics of a Continuously Learning Health Care System. 2012. <http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America/Table.aspx>
10. Hamburg M, Collins F. The path to personalized medicine. *N Engl J Med* 2010;363:301–4. <http://www.nejm.org/doi/pdf/10.1056/NEJMp1006304>
11. Personalized Medicine Coalition. *The case for precision medicine*. 4th edn. Washington, DC. 2014. http://www.personalizedmedicinecoalition.org/Userfiles/PMC-Corporate/file/pmc_case_for_personalized_medicine.pdf
12. Office of the National Coordinator for Health Information Technology (ONC). Federal Health IT Strategic Plan: 2011–2015. 2011. <http://www.healthit.gov/sites/default/files/utility/final-federal-health-it-strategic-plan-0911.pdf>
13. Office of the National Coordinator for Health Information Technology (ONC). Connecting Health and Care for the Nation: A 10-Year Vision to Achieve an Interoperable Health IT Infrastructure. 2014. <http://healthit.gov/sites/default/files/ONC10yearInteroperabilityConceptPaper.pdf>
14. Milstein A. Code red and blue—safely limiting health care’s GDP footprint. *N Engl J Med* 2012;368:1–3.
15. Executive Office of the President, Council of Economic Advisers. *The economic case for health reform*. Washington, DC: The White House. 2009. <http://www.whitehouse.gov/administration/eop/cea/TheEconomicCaseforHealthCareReform>
16. OECD. *Health at a glance 2013: OECD indicators*. OECD Publishing, 2013. <http://www.oecd.org/els/health-systems/Health-at-a-Glance-2013.pdf>
17. Mahon M, Weymouth J. U.S. Spends far more for healthcare than 12 industrialized nations, but quality varies. New York: The Commonwealth Fund, 2012. http://www.commonwealthfund.org/~media/Files/News/News%20Releases/2012/Apr/SquiresOECDRelease_FINAL_5312wolink.pdf
18. See also: Berwick D, Hackbarth M. Eliminating waste in US health care. *JAMA*, 2012;307:1513–6. <http://news.medicin.duke.edu/wp-content/uploads/2013/08/Eliminating-Waste-in-US-Healthcare-Berwick.pdf>
19. McGlynn E, Asch SM, Adams J, et al. The quality of health care delivered to adults in the United States. *N Engl J Med* 2003;348:2635–45. <http://www.nejm.org/doi/full/10.1056/NEJMsa022615>
20. National Healthcare Quality Report 2012. Rockville, MD, 2013.
21. (IOM) Medicine, I. of. Roundtable on Value & Science-Driven Health Care. The Learning Health System Series. 2012. <http://iom.edu/~media/Files/Activity%20Files/Quality/VSRT/Core%20Documents/LearningHealthSystem.pdf>
22. Institute of Medicine. The Learning Healthcare System. National Academies Press, 2007. <http://www.iom.edu/Reports/2007/The-Learning-Healthcare-System-Workshop-Summary.aspx>
23. Rubin R. “How did Vioxx debacle happen?”. *USA Today* 12 October 2004. http://usatoday30.usatoday.com/news/health/2004-10-12-vioxx-cover_x.htm
24. Norton L. Creating infrastructure for big data in the public and private sectors. Presented at ECRI Institute Annual Conference: Data BIG and Small: What Healthcare Decision Makers are Using Now; November 2013. https://www.ecri.org/Video/2013_TA_Conf/4-Session-1-Norton.mov
25. Tannen R, Weiner M, Xie D. Use of primary care electronic medical record database in drug efficacy research on cardiovascular outcomes: comparison of database and randomised controlled trial findings. *BMJ* 2009;338:b81. http://www.bmj.com/highwire/filestream/377160/field_highwire_article_pdf/0/bmj.b81
26. Dzau VJ, Cho A, ElLaissi W, et al. Transforming academic health centers for an uncertain future. *N Engl J Med* 2013; 369:991–3.
27. Health Affairs. Using big data to transform care. July 2014: Volume 33, Issue 7. <http://content.healthaffairs.org/content/33/7.toc>
28. Etheredge LM. Rapid learning: a breakthrough agenda. *Health Aff* 2014;33:1155–62. <http://content.healthaffairs.org/content/33/7/1155.full.pdf+html>

29. See e.g.: Collins J, Porras J. Built To Last: Successful Habits of Visionary Companies. 1994.
30. Buchanan L. How to Achieve Big, Hairy, Audacious Goals: Leadership expert Jim Collins explains what drives some entrepreneurs to relentlessly pursue bold ideas—and succeed where others have failed. Inc. November 1, 2012. <http://www.inc.com/leigh-buchanan/big-ideas/jim-collins-big-hairy-audacious-goals.html>
31. Halamka JD. *The Rise of Electronic Medicine*. MIT Technology Review. 1 September 2011. <http://www.technologyreview.com/news/425298/the-rise-of-electronic-medicine/>
32. O’Harlow RJr. *The Machinery Behind Health-Care Reform*. *The Washington Post. Business* 16 May 2009. <http://www.washingtonpost.com/wp-dyn/content/article/2009/05/15/AR2009051503667.html>
33. Blavin FE, Buntin MB. Forecasting the use of electronic health records: an expert opinion approach. Centers for Medicare & Medicaid Services, Office of Information Products & Data Analytics. Medicare & Medicaid Research Review 2013: Volume 3, Number 2 (unpublished). http://www.cms.gov/mmrr/Downloads/MMRR2013_003_02_A02.pdf (unpublished)
34. Technology, P. C. of A. on S. and. Report to the President realizing the full potential of health information technology to improve healthcare for Americans: the path forward. Washington, DC, 2010:1–108.
35. Etheredge LM. A rapid-learning health system. *Health Aff*, 2007;26:w107–18.
36. “CancerLinQ.” ASCO.org. American Society of Clinical Oncology, n.d. Web. 06 Nov. 2013. <http://www.asco.org/quality-guidelines/cancerlinq>
37. Curtis L, Brown J, Platt R. Four health data networks illustrate the potential for a shared national multipurpose big-data network. *Health Aff* 2014;33:1178–86. <http://content.healthaffairs.org/content/33/7/1178.full.pdf+html>
38. Summit D. at the M. 2012L. H. S. (2012). Core Values Underlying a National-Scale Person-Centered Continuous Learning Health System (LHS) (pp. 1–2). <http://www.learninghealth.org/corevalues/>
39. The TRANSFoRm Project EU. Translational Research and Patient Safety in Europe. Funding from the European Union’s Seventh Framework Programme for research, technological development and demonstration under grant agreement no 247787 [TRANSFoRm]. <http://www.transformproject.eu/>
40. See e.g., CDISC. CDISC launches learning health community standards initiative. Essential Standards to Enable Learning (ESTEL). Austin, TX, 2013. <http://www.cdisc.org/lhs>
41. Sullivan K, Knaus W, Marks R. An Ultra-Large-Scale Systems Approach to National-Scale Health Information Systems. Future of Software Engineering Research. FoSER 2010, November 7–8, 2010, Santa Fe, NM., 2010 ACM 978-1-4503-0427-6/10/11. NITRD, 2011. http://www.nitrd.gov/nitrdgroups/images/b/ba/An_Ultra-Large-Scale_Systems_Approach_to_National-Scale_Health_Information_Systems_-p365.pdf
42. See e.g.: Bammer G. “Integration and Implementation Sciences: Building a New Specialization”. The Resilience Alliance. *Ecol Soc* 2005;10:6. <http://www.ecologyandsociety.org/vol10/iss2/art6/ES-2005-1360.pdf>
43. Gorry GA. Envisioning the system after next. *EDUCOM Rev*. 1997:30–3.
44. Internet Hall of Fame. Internet Hall of Fame Pioneers, Inductees, Posthumous Recipient: J.C.R. Licklider. <http://www.internethalloffame.org/inductees/jcr-licklider>
45. See e.g.: Cutler D. Where are the health care entrepreneurs? The failure of organizational innovation in health care. National Bureau of Economic Research. NBER Working Paper No. 16030. Issued in May 2010. <http://www.nber.org/papers/w16030>
46. National Research Council. *Toward precision medicine: building a knowledge network for biomedical research and a new taxonomy of disease*. Washington, DC: The National Academies Press, 2011.
47. Sullivan K. Cyber-Social Ecosystems. Kevin J. Sullivan—Blog. 21 December 2013. <https://www.kevinsullivan.com/blog/%5B%20node%3Acreated%3Acustom%3AY%5D/12/21/cyber-social-ecosystems>

AUTHOR AFFILIATIONS

¹University of Michigan, Ann Arbor, Michigan, USA

²Harvard Pilgrim Health Care Institute, Boston, Massachusetts, USA

³Vanderbilt University, Nashville, Tennessee, USA

⁴National Library of Medicine, Bethesda, Maryland, USA

⁵George Washington University, Washington, District of Columbia, USA

⁶University of Illinois, Urbana-Champaign, Illinois, USA

⁷Stanford University, Stanford, California, USA

⁸University of Virginia, Charlottesville, Virginia, USA