

“20-Percent Doctor Included” & Dr. Algorithm: Speculations and Musings of a Technology Optimist

by Vinod Khosla

khosla ventures

September 2016

TABLE OF CONTENTS

Summary

Introduction

- 01** "The Practice of Medicine" Versus
"The Science of Medicine"
- 02** Healthcare - Innovation Emerging Out of
Complexity
- 03** Replacing 80 Percent of What
Doctors Do?
- 04** Sources, Timing, Incentives, and
Pitfalls of Healthcare Innovation

SUMMARY

01. “The practice of medicine” versus “the science of medicine”

Healthcare today is often the “practice of medicine” rather than the “science of medicine”. In the best cases of practice, good care is delivered, but in the worst cases of the practice of medicine, medical doctors (MD’s as they are called in the US) take moderately educated shots-in-the-dark when it comes to patient care. We should describe improvements in practice over time as how practice improves and why it has gotten better, and how we make it a science! Much of the current practice is driven by conclusions derived from partial information of a patient’s history and current symptoms interacting subjectively with various known and unknown biases of the physician, hospital, and healthcare system as a whole. The future of healthcare should utilize an approach akin to the scientific method, with increased data collection, analysis and experimentation to rapidly improve systems. Physicians could be much more scientific and data-driven with better systems assisting them in the future. With the increasing amount of data and research released every year, it’s hard for the average physician to keep up without technology; the doubling time of medical knowledge is estimated to be less than 3.5 years and accelerating. “Students who graduate in 2020 will experience four doublings in knowledge. What was learned in the first 3 years of medical school will be just 6% of what is known at the end of the decade from 2010 to 2020.”¹ The next generation of medicine will arrive at scientific and data-driven diagnostic and treatment conclusions based on more complete testing of what’s actually going on in a patient’s body. Today all the data in the electronic medical record is only mildly useful in helping with diagnosis. Testing and data collection is limited to what human physicians know how to use when the complex human body could be monitored with orders of magnitude more data

that truly characterizes what the bodies many systems are doing, much like a car today or an America’s Cup boat is very well instrumented.

In the past, the data to make more rigorous and scientific conclusions has simply not been available. And as a result, medical literature is rife with studies about how the practice of medicine does not meet expectations for what would constitute sufficient, correct care. There are plenty of examples that illustrate this, but they tend to share the same themes. (1) Purported experts in their respective fields frequently disagree on the effects of basic procedures instead of agreeing on possibilities, probabilities and potential outcomes. A study on colon cancer experts, for example, showed that there was full distribution across the board (0%-100%) on how valuable colon cancer screening is. (2) Things that are treated as medical fact often end up being completely wrong (yet linger for a while). Prescriptions for antipyretics such as aspirin are typically given to individuals with fever (and have been for over a century). Yet recent studies showcase that prescribing antipyretics to reduce a fever could be significantly more risky than just allowing the fever to run its course (i.e. do nothing)! Whether this particular new fact is true or not, Dr. John Ioannidis has studied this phenomenon extensively, and it holds true throughout medical academia, where it is more likely for a research claim to be false than true.² Even deciding if some condition is a disease or a normal state is sometimes hard. Medicine is complex and hard to find easy answers to but that does not negate the fact that we can do much better than we are doing today.

Misdiagnosis, conflicting diagnoses, and general diagnostic error are also common problems in today’s medical system. These tend to be solvable, but occur for similar reasons as the problems of medical research. Biases, judgment errors, incomplete information, lack of familiarity with

1 Densen, Peter, “Challenges and opportunities facing medical education” *Trans Am Clin Climatol Assoc*, 2011. (<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3116346/>)

2 Ioannidis, John P.A. “Why most published research findings are false” *PLoS Med*, 2005 (<http://www.plosmedicine.org/article/info%3Adoi%2F10.1371%2Fjournal.pmed.0020124>)

the latest research, and other similar factors during patient care leads to billions of dollars in preventable cost, and more importantly, tens of thousands of unnecessary fatalities yearly in the U.S. alone. The difficulty in coordinating care especially in the most difficult patients with multiple conditions and comorbidities impacting one another is a continual burden on the healthcare system, and often a source of poor treatment. A analogy might be a building being assembled by a carpenter, roofer, electrician, plumber, painter, mason, decorator, and landscaper who all have their own, unique plans and designs. Each worker (or their apprentice) comes in and does a bit of work when they have a chance, sometimes leaving a post-it note where others can see it, sometimes just copying a note from the day before and leaving a copy on the stack, sometimes not sharing any information. Each of these professions have their own guild, with their own specialized lingo and preferred way of doing things, often wanting to take as big a share of the job as possible (both for credit and increased profit/billing, but also because of some

mutual mistrust in others abilities), but also just as likely to ignore problems and concerns out their specific domain area. This would be a crazy way to work on a building, but this is how the healthcare system takes care of our selves. This is perhaps why a 2013 study estimated that greater than 400,000 deaths a year were attributable to medical errors in the US and estimated that there were perhaps 4,000,000 events of serious harm attributable to medical errors.³ To put this in perspective, in 2013 there were 37,938 vehicle associated deaths; 41,149 suicides; 16,121 assault/homicides (33,636 total firearm deaths); 6,995 deaths attributed to HIV/AIDS; 584,881 deaths attributed to cancer; and 611,105 deaths attributed to heart problems.⁴ Even when we have the best available science and technology, it is being inappropriately or inadequately applied in ways which have major impact on wellbeing. More integrative patient care across their multiple conditions with software systems with deeply specialized knowledge in multiple specialties and data based insights that may be too complex for any one human to integrate across is one very likely benefit of a data driven system.

Preventability Rationale Standard way	Percentage of Events*
Preventable Events (n=133)	
Error was related to medical judgement, skill, or patient management	58%
Appropriate treatment was provided in a substandard way	46%
The patient's progress was not adequately monitored	38%
The patient's health status was not adequately assessed	23%
Necessary treatment was not provided	17%
Event rarely happens proper precautions and procedures are followed**	14%
Communication between caregivers was poor	8%
Facility's patient safety systems and policies were inadequate or flawed**	3%
Breakdown in hospital environment occurred (equipment failure, etc.)**	2%
Nonpreventable Events (n=155)	
Event occurred despite proper assessment and procedures followed	62%
Patient was highly susceptible to event because of health status	50%
Care provider could not have anticipated event given information available	35%
Patient's diagnosis was unusual or complex, making care difficult	29%
Harm was anticipated but risk considered acceptable given alternatives**	14%

3 James, JT, 2013 http://journals.lww.com/journalpatientsafety/Fulltext/2013/09000/A_New_Evidence_based_Estimate_of_Patient_Harms.2.aspx

4 <http://www.cdc.gov/nchs/fastats/leading-causes-of-death.html>

Study of 780 Medicare patients identified 133 preventable adverse events summarized in the table above.⁵ These problems are what we believe technology will address. Problems such as poor monitoring of patient progress is certainly amenable to technological improvement in the ways we outline. We also believe that some of these “nonpreventable” events will be amenable to improved technology and machine intelligence, such as addressing complex patients, rare events, and improved access to information.

Bringing science and data to medicine will enable us to more rapidly change these debilitating problems. If we are able to collect exponentially more data, as well as collect it continuously, we will have the proper inputs to drive change. Increasing the number of experiments we can run using this data will help us more quickly gain clinical insight and value. The use of data science in particular will help add meaning to all of this collected data, and over time, two very distinct improvements will happen: (i) a better validation of what we accept in medical practice about today’s therapies, prescriptions and procedures; and (ii) the invention of brand new prescriptions, therapies, insights and procedures based on new and more holistic data about a patient. Knowledge (not just text) mining of 50-100 million biomedicine research papers in various journals that constitute human knowledge and research in this area will make the knowledge base of these artificial intelligence software systems more complete and dynamic, including surfacing inconsistencies and special circumstances as such research is reconciled into a more consistent knowledge base. This will be further enhanced by data science based analysis of medical information from individual patients. Together the two dimensions of knowledge graphs (which are likely to be incomplete for a substantial period of time) from research and data science based analysis from patient care data will reinforce each other and accelerate progress and the potential of quality care if the medical establishment doesn’t insert self-interest in the way to slowing it down.

This does not imply that the biological sciences will not be important, as fundamental scientific research in biology will keep improving our

understanding of biological systems and will feed into the complex data science systems and will extend our knowledge base beyond correlations to actual understanding. Biological advances such as the rise of genomic and phenotype knowledge and CRISPR-Cas9 gene editing coupled with technological advances in gene sequencing and gene synthesis will dramatically help speed up experimentation and understanding, and there will be a positive feedback loop so machine learning will be able to guide future experiments. The time period for such a scenario driven by digital technologies could be fifteen years or may take an extra decade or even more but to me, timelines seem to be far less important than the directionality. I also suspect because of reasons related to the “nature of science”, the innovation cycles for biological science-based contributions to medicine will be longer than those for the digital sciences.

02. Healthcare – Innovation emerging out of complexity

The healthcare transition will start incrementally and develop slowly in sophistication, much like a great MD who starts with seven years of medical school and then spends a decade training with the best practitioners by watching, learning and experiencing. Expect many laughing-stock attempts by “toddler computer systems” early in their evolution and learning; they will be the butt of jokes from many writers and doctors. Early printers, typically the dot matrix variety, the toddler generation of computer printers, did not exactly cut it for business correspondence, let alone replace traditional typewriters. But within a few generations the IBM Selectric typewriter was replaced by constantly improving printing technology. Few people in the 1980’s believed that every person would have a personal computer in their home. In 1995, the Internet was not considered as important and companies like AT&T believed that 64 kbps service was all any home would ever need but today Google does not think that a 1000x that number is sufficient. And few believed Google was different than any other search engine in the late 1990’s. And few thought Facebook would take over the world when it was a site for college students in mid 2000’s.

5 Office of the Inspector General, “Adverse Events in Hospitals”, 2010 <http://oig.hhs.gov/oei/reports/OEI-06-09-00090.pdf>

Yet each of these examples showcase that innovation can occur swiftly and dominantly, more like a tsunami than a linear change in the state of the world. In this process, emergent (as of yet unknown) properties will arise from the disruption, leading to a fundamental change to how we will approach healthcare.

Just because a three-year-old child makes some laughable errors does not imply that they will make the same errors as a 21 or 40 year old! Similarly to equate early “toddler digital health systems” to what might eventually be possible is naive. For an imperfect analogy, look at the evolution of cell phones from the clunky to the sublime! I imagine within a few iterations of these systems (here called v0 through v7), we will have a world in where doctors will have much more data fed into their decision-making using the personalized medical equivalent of a Bloomberg financial terminal. Just because a technology does not exist now does not mean it won't exist in the future, and my optimistic speculation is that we are on a path in digital health to create this innovation.

It's just that we won't let the initial toddler systems actually make real decisions while they are learning and growing up in sophistication. They will be in “assist, learn and amplify” mode with new generation of systems being developed every two, three, or four years (a typical development cycle for a sophisticated software system) and with radical improvements in sophistication and capability over seven illustrative generations, much like the cellphone of 1986 (a floor mounted device for your car with heavy handset cords) grew up to be the iPhone of today! This cellphone analogy is one we will return to again multiple times to illustrate how change can happen. The accumulation of data and knowledge graphs will help accelerate these improvements, shifting our current state of “sickcare” to one of “wellcare” where we will have the ability to always think about and understand our health. The transition to the automated science of medicine will likely occur in an organic process of trial and error, starting with initial technologies and ideas that go through multiple iterations and restarts over the coming years.

It's important to note that although many of the

examples used in my discussion are from the United States, one of the most exciting aspects of this revolution in medicine is that it will be able to scale internationally in ways that much of traditional healthcare has been unable to do. The physical technology required is leveraging innovations across technology to produce smaller, more robust devices which can be extremely cheap to produce at scale. Previously prohibitively expensive technologies like genomic sequencing are dropping in cost faster than Moore's law and the ability of software systems to use the genomic, epigenetic, phenotypic, metabolomic, proteomic, and transcriptomic data will expand exponentially. In addition to economies of scale in production, there are key economies of scale in knowledge production, basic sciences discoveries are true everywhere. It is sometimes true that local differences in environment or genetics in a particular population may modify the effectiveness of a particular therapy, but a learning healthcare system will identify those differences much more quickly than the current, unconnected system. Precision, personalized medicine does not need to be expensive medicine; for example, doing the first clinical analysis of a human genome was expensive, but that analysis produced software that significantly speed analysis of the subsequent genomes, and that acceleration is ongoing. Soon it will cost less to do full sequencing than to FedEx the sample. A healthcare system that incorporates a large component of machine learning and is able to leverage mobile devices like mobile phones may enable a developing country to leapfrog the more industrialized nations by not requiring the massive investment in staff and physical infrastructure that developed countries have invested. In much the way that some developing countries have largely leapfrogged the large investment the industrialized world has put into landlines for telecommunications by jumping to mobile phones, a similar process in healthcare is possible. In fact, it may be the only option in many developing countries.

03. Replacing 80-percent of what doctors do?

Technology makes up for human deficiencies and amplifies human strengths – doctors and even other less-trained medical professionals could do so much more than they do now. Today's diagnostic error rate in medical practice

is roughly the equivalent of Google's driverless car having one accident per week; while this would be unacceptable for automated cars, this type of failure rate is permissible in healthcare. In few decades, data will transform diagnostics, to the point where automated systems may displace up to 80-percent of physicians' or healthcare workers' standard work initially. Technological developments will initially AMPLIFY physicians' abilities by arming them with more complete, synthesized and up-to-date research data and knowledge, all of which will lead to better patient outcomes. Software applications (e.g. on a mobile device) will handle many functions that are today handled (inefficiently and sometimes incorrectly) by humans. Computers are much better than people at organizing, recalling, and synthesizing complex information and decisions. This will result in far fewer mistakes and biases than a hot shot MD from Harvard, let alone the average (or median for those statistically inclined) doctor I am most concerned with here.

While it is unclear how this will manifest, patients as consumers will also transform the current role of the physician. With increased sensors and devices that are constantly with us, if needed, the consumer can become the CEO of his or her own health without coordinating with their physician or hospital. This can be particularly empowering in cases where patients have very little access to care. The first prospective randomized trial utilizing multiple smartphone-enabled biosensors was performed recently by Dr. Eric Topol of the Scripps Research Institute in La Jolla, CA⁶. While this pioneering effort did not show reduced healthcare utilization, it showed "evidence of improvement in health self-management which was characterized by a decrease in the propensity to view health status as due to chance factors". Over time, these mobile sensors, devices, and apps can harness their increased data collection and data science sophistication to come up with insights and outperform the average physician in aggregate, though not in every case. While many will choose this path, consumer choice should always be held as a priority for those who want to use our current doctors and healthcare system for care. Let the data accumulate on which systems

perform well in various circumstances.

Common concern about widespread adoption of shared, common baselines of care is that this turns into a cookie-cutter approach that prevents customization and personalization of care. In the scaling up of production systems the opposite is the case. The development of a common foundation or chassis allows resources to be increasingly allocated to customization instead of the baseline. Instead of handcrafting everything from scratch for every individual, individualization will be precision tweaks on top of the baseline standard. Indeed, as we move toward the dream of precision medicine based on a patient's unique, measurable aspects, such as a patient's sequenced genome, it will become impossible for any individual healthcare provider to develop a targeted treatment plan that incorporates an understanding of the complexity of a single body without substantial machine intelligence support, let alone for a panel of patients, each with their own unique genome and life situation. One of the major challenges going forward will be the range of possibilities offered by alternative therapeutic options; however technology can even assist here by helping individuals compare probabilities, compile and reconcile their own goals, values and opportunities, and even perhaps assist in modeling and simulating different outcome options.

I am not suggesting that every physician will change how they practice medicine in 15-25 years, but rather that the thought leaders will be doing so, and the future direction of medicine will be self-evident and the advantages to patient outcomes will be mostly established in well documented studies. We are already seeing this shift happening, with people on the fringe of medicine as well as a few (but growing number of) thought leaders entrenched in medicine taking steps to enable this future. They will gravitate to a world where the best strengths of humans and doctors are harnessed in taking care of patients, while "Dr. Algorithm" systems will work with them to do the bulk of what we know of as diagnostic, monitoring, and prescription work, improving via both automatic feedback mechanisms as well

6 Bloss et al. (2016), A prospective randomized trial examining health care utilization in individuals using multiple smartphone-enabled biosensors. *PeerJ* 4:e1554; DOI 10.7717/peerj.1554

as human scientific input. These medical and technology leaders will show markedly better care and treatment results, and over time, the rest of the world will join the science of medicine.

04. Sources, timing, incentives, and pitfalls of healthcare innovation

The major problems in healthcare are systemic, despite the many doctors who are accomplished, caring, honest and compassionate providers. There are advertent and inadvertent human actions and biases and system incentives that contribute to them. The first problem is that globally, there is a misalignment in incentives, where organizations try to maximize revenue (extra surgeries anyone?) at the expense of optimizing care (just like some car mechanics!). These lead to hidden biases in how we administer care, and has been particularly showcased in the US in its ongoing struggle for large-scale governmental health care reform. The second problem deals with the crawling pace of change in how the AVERAGE doctor operates and gathers new knowledge, even in the presence of a rapid increase of data and knowledge about how to improve care and treatment. Third, there is an incredible increase in the amount and complexity of newly enabled data, vast amounts of research, longitudinal health records, and medical histories. On top of this, new sensors and testing will allow for much more integrative analysis than is currently possible (especially by humans). Utilizing this data will enable much better and more holistic care that will only get progressively better with time, yet are barely prioritized with current healthcare incentives. These lead to my belief that innovation will most likely come from outside the system. Users in more desperate need and lacking traditional healthcare may adapt these innovations first.

As it's actually relatively standard for deep innovation to happen outside of their traditional ecosystems. In most areas this happens from innovators outside the system, acting somewhat naively, failing and then realizing they need some knowledge and collaboration with the system.

Entrepreneurial teams often add domain expertise to their naive "fresh piece of paper" re-invention ideas. Society generally tries to assign more power

to larger entities, like governmental institutions and the Fortune 500 behemoths, but true radical innovation seldom comes from them. Did Walmart reinvent retail or Amazon? Did General Motors perfect the electric car or Tesla (despite its many foibles!)? Did NASA or Lockheed Martin reinvent space launches or SpaceX? Did NBC reinvent media or YouTube? Most importantly did big pharmaceutical companies reinvent biotechnology pharmaceuticals or did Genentech?

If it's outside the system – could innovation come top-down from governments? Typically, growth and innovation tends to be organic for systems that are data-driven and consumer-driven. And typical life cycles of innovation for digital technologies are much shorter than those in tightly regulated healthcare. So even if the U.S. Food and Drug Administration (FDA) or the U.S. government can help by being progressive and helping align incentives, technology (even in its early iterations) will be able to innovate at a faster pace. In fact, one of the biggest risks in slowing medical innovation is slowness or damaging policies by governmental intervention in the face of demonstrated minimal risk. In the worst case, this will bring some forms of technology-driven medical innovation to a halt, but it's also likely that the innovation would just move to more progressive countries that allow for greater experimentation and use of data-driven systems. Technology that helps save costs in a first-order way, as well as technologies that have a strong mobile component can spur this non-US growth in innovation.

There are a lot of improbable sounding possibilities on how data and consumer-driven systems will transform healthcare. Though any particular one is unlikely to become reality, it will be some improbability that will determine the future of health care as it is driven, molded and transformed by digital health technologies. Some improbable scenario today will become tomorrow's reality. I believe over time, we will see a 5x improvement across healthcare: a 5x reduction in doctor work (shifted to data-driven systems), a 5x increase in research (due to the transformation to the "science of medicine"), a 5x lower error rate (particularly in diagnostics), a 5x faster diagnosis (can be on your app), and a 5x cost reduction in care. These are not but rather

“more true than not” speculations. We just have to imagine what might be possible! And we must then have the courage to try and make those possibilities a reality. Healthcare has a radical opportunity to reinvent itself.

INTRODUCTION

Healthcare today often results in suboptimal patient outcomes despite doctors doing the best they can within the current system. Suboptimal outcomes result from the incomplete knowledge and personal biases of today's system and the system being better than it has been. Medicine has historically been approached according to tradition – the experiential evolution of best practices, and a reductionist system of small trials. Optimal treatment outcomes require a healthcare system that is instead primed by holistic, scientifically, probabilistically or other statistically-validated data and conclusions presented to patients as cost/benefit choices. It is time to move beyond the stethoscope, which remains the iconic diagnostic tool for most healthcare professionals worldwide, 200 years after its invention.

Technology will reinvent healthcare as we know it. It is inevitable that, in the future, the majority of physicians' diagnostic, prescription and monitoring, which over time may approach 80-percent of total doctors'/internists' time spent on medicine, will be replaced by smart hardware, software, and testing. This is not to say 80-percent of physicians will be replaced, but rather 80-percent of what they currently do might be replaced so the roles doctors/internists play will likely be different and focused on the human aspects of medical practice such as empathy and ethical choices. Healthcare will become more scientific and more consistent, delivering better-quality care with inexpensive, but orders of magnitude more data-gathering techniques, continual monitoring, more rigorous science and more available and ubiquitous information leading to personalized, precise and consistent (across doctors) insights into a patient. Disease will be measured not by the symptoms it creates but objectively evaluated by the metabolic pathways or physical parts it affects. Many new findings will be outside the reach of most physicians because of the volume of data and the unique holistic insights that data will provide about a patient's very complex condition. Hundreds of thousands or even millions of data points may go into diagnosing a condition and monitoring the progress of a therapy or prescription, well

beyond the capability of any human to adequately consider. Plavex, a blood thinner, is commonly prescribed. How many patents have been genetically tested for whether they metabolize it slowly or rapidly? How many prescribing physicians know it can be done for deciding dosage or if it is the right medication?

This evolution from an entirely human-based healthcare system to an increasingly automated system that enhances human judgment will take time, and there are many ways in which it can happen. Likely the next decade will mostly see systems providing "bionic assist" to physicians and complementing or enhancing their skills. Today's traditional approaches will get better as new approaches, and even new medicine, is invented. As the 80-percent of physician work is replaced over a few decades, the remaining 20-percent will be AMPLIFIED, making them even more effective, and allowing even the average physician or nurse to perform at the level of the very best specialists. Doctors will be able to operate at substantially improved levels of expertise in multiple domains, and they also will be able to handle many more patients. The primary care physician and maybe even the nurse practitioner may be able to operate at the level of six specialists handling six areas of care for one patient with multiple comorbidities in a more coordinated and comprehensive manner without inter-specialist friction. This transition will affect each group of actors in the current system differently. Internal medicine will be transformed to the greatest extent. Procedure-based or interventional medicine may take longer due to the nature of the surgical art, and procedure-volume based incentives. Some constituencies will be affected favorably in some dimensions and worse in others, but the net benefit will be substantially positive for society and individual patients. It is likely that a focus on science, data, and personalization will lead to plenty of unintended benefits that we cannot anticipate today. Nurses will be made much more capable by technology, often replacing the functions only doctors perform today. New medical insights, including ones we cannot yet envision, will be commonplace, and

the practices we follow will be substantially better validated by more rigorous scientific methods. Projects like the Cancer Moonshot will apply rigorous genomic, proteomic and phenotypic tools and within large trials, to optimize the inadequate patient outcomes in oncology practice today. Though medical textbooks won't be "wrong", the current knowledge embodied in them will mostly be replaced by much more precise and advanced methods, techniques, and understandings.

My statements are not forecasts that the hospital burn unit or emergency department will run without any humans on staff (always amazed to watch humans operating in emergency rooms, though it is sad to see patients waiting for attention.) Though the early changes will appear underwhelming and clumsy, in a few decades they will seem obvious, inevitable and well beyond the changes we might envision today. Expect today's expert doctors to think these changes are implausible when they are asked about this possibility, and expect the classic response of "human judgment will not be replaced by technology". To them I say: The nature of technology's exponential curve is non-intuitive for humans; the capabilities of smart technologies in 2040 (hardware, software, tests) are hard to imagine, just as today's smart phones were unimaginable 15 years ago. Even most software experts are unqualified to judge where technology will lead in two decades, let alone doctors who have little familiarity with the rate of progress and possibilities in these areas. The role humans will play in this is hard to define exactly but I suspect strongly that their role in healthcare will change materially. It is possible that a much more cooperative system leveraging the respective

strengths of both humans and technological systems may also evolve, as proposed in the book *Race Against the Machine*. However, the core functions necessary for complex diagnoses, treatment, and monitoring (as a significantly expanded function of healthcare) will more than likely be driven by machine judgment instead of human judgment alone. In fact, as Atul Gawande pointed out, some studies showcase that "our attempt to acknowledge and deal with human complexity [in human ways] causes more mistakes than it prevents."⁷

This transformation will happen in fits and starts along different pathways with many course corrections, steps backward and mistakes as we figure out the best approach. Given the importance of having clarity on what I hypothesize as my forecasts, I want to be clear that they are only directional guesses rather than precise predictions. Further, though many different disciplines will contribute to the innovation in medicine like biological research or new device development, am mostly concerned with the contributions of digital health technologies (smart hardware, software, tests) to medical innovation. This should not be underemphasized, as these contributions, though potentially the most significant, are also the most variable, and hardest to predict in direction, timelines and scope and the ones that will face most resistance from human practitioners and organizations who will likely try and delay them. The rates of progress in internal medicine, procedural medicine, acute care, chronic care, diagnosis etc. will each have its own tortured path to this much better place. The other sciences will continue to contribute much more fundamental insights into human well being.

7 Gawande, Atul. "No Mistake." *The New Yorker*, 1998. (http://www.newyorker.com/archive/1998/03/30/1998_03_30_074_TNY_LIBRARY_000015236)


THE “PRACTICE OF MEDICINE” VERSUS THE “SCIENCE OF MEDICINE”

absolutes Healthcare today is often the “practice of medicine” rather than the “science of medicine”. Diagnoses are partially informed by a variety of medical and non-medical factors, each fraught with problems. Based on historical practice and incremental improvements, it is better than ever but not as good as it can be.

These factors include (i) the patient’s medical history, which is often incomplete and often not considered longitudinally; (ii) the patient’s symptoms, which are either taken poorly, or are poorly communicated by patients; (iii) incentives and promotions driven by the pharmaceutical and medical devices industries; and (iv) the doctor’s partial memory of lessons from medical school, which are laden with cognitive biases on top of potentially being outdated and obsoleted by more recent research, and (v) standard human errors. Doctors are human beings like the rest of us, and they can only keep so many variables in mind when making a decision. Often, if you ask three doctors to look at the same problem, you’ll get three different diagnoses and three different treatment plans. As a patient, how would you feel if a doctor keeps changing his mind about

your disease over time? How would you feel if different doctors say different things about your disease? Today, this happens often. Due to our lack of understanding the specifics of disease mechanisms or changes in our body, we often get the prescription of “rest, eat well and exercise” to a range of maladies. Health data is ‘multiscale’ and comes from a variety of sources (genomic, phenotypic, wearable, lab, biomarkers...) but today’s typical doctor seldom gets and is able to utilize this in a unified way even in those rare times when it is available. The highest probability treatments or prognosis given patient preferences and medical uncertainty is the best path. But often this is hard for humans to achieve or doctors to overcome their biases and to fully include patient preferences in their rushed daily life.

This is not the fault of the doctor or of medicine. To date, the practice of medicine has had very limited information, a range of very complex and sometimes conflicting research, and a large set of for-profit, promotional, and confusing recommendations. It takes the smartest people to get into and out of medical school, and they do the best they can under all the pressures, conflicts,



and realities of the real world of having to make decisions with imperfect information.

The Practice of Medical Research

Let's take a look at modern medicine's view towards fever as an example. For 150 years, doctors have routinely prescribed antipyretics such as aspirin or acetaminophen to help reduce fever. Since it is viewed as an inability of the body to regulate itself, the approach is to aggressively reduce temperatures⁸ in most cases.⁹ In 2005, researchers at the University of Miami, Florida ran a study to test this assumption. Patients were randomly assigned to two groups: (1) the standard treatment of receiving antipyretics if their temperature rose beyond 101.3 degrees Fahrenheit, or (2) if their temperature reached 104 degrees Fahrenheit (considered a high fever). As the trial progressed, seven people who received the standard treatment died, while there was only one death in the group of patients allowed to reach a high fever. At this point, the trial was stopped because the team felt it would be unethical to allow any more patients to receive the standard treatment.¹⁰ When something as basic as fever reduction is a hallmark of the "practice of medicine" and hasn't been tested in over 100 years, we have to ask what else might be practiced due to tradition rather than science? Whether this finding holds up in further study or not it does illustrate the fact that there is a lot more experiential practice in medicine than science.

From hospital to hospital, standard guidelines are not necessarily the same across symptoms and procedures. In fact, even within hospitals, most of what is followed is dependent on the physician rather than probabilistic guidelines that are set. It is amazing that in this day and age each hospital system has to set its own care guidelines? Doesn't science dictate what is best for any given patient?

As the Institute of Medicine says "technological tools, such as decision support tools that can be broadly embedded in electronic health records, hold promise for improving the application of evidence."¹¹ According to one researcher¹² there is good reason to challenge the assumption that every individual practitioner's decision is necessarily correct. Failure of the assumption has immense implications for the quality of care. It implies that the same patient can go to different physicians, be told different things, and receive different care. No doubt some of the differences will not be important. However, some will surely be important—leading to different chances of benefits, different harms, and different costs.

A failure of the assumption also has immense implications for informed consent, expert testimony, consensus development, the concepts of "standard and accepted" or "reasonable and necessary," malpractice, quality assurance programs that are based on statistical norms, and the cost of care. We should have uniform care guidelines throughout the country and care to

8 A. Croskerry, Pat. "The Importance of Cognitive Errors in Diagnosis and Strategies to Minimize Them." *Academic Medicine* 2003: Vol. 78, No. 8 (<http://www.jround.co.uk/error/reading/croskerry1.pdf>).

B. Croskerry, Pat. "From Mindless to Mindful Practice – Cognitive Bias and Clinical Decision Making." *N Engl Med* 2013; 368:2445-2448 June 27, 2013 DOI: 10.1056/NEJMp1303712.

C. Hicks, EP, et al. "Heuristic reasoning and cognitive biases: Are they hindrances to judgments and decision making in orthodontics?" *Am J Orthod Dentofacial Orthop.* 2011 Mar;139(3):297-304. doi: 10.1016/j.ajodo.2010.05.018.

D. Graber, ML, et al. "Diagnostic error in internal medicine." *Arch Intern Med.* 2005 Jul 11;165(13):1493-9.

E. Burgess, D., et al. "Reducing Racial Bias Among Health Care Providers: Lessons from Social-Cognitive Psychology." *J Gen Intern Med.* 2007 June; 22(6): 882–887.

F. Reilly, J., et al. "Teaching about how doctors think: a longitudinal curriculum in cognitive bias and diagnostic error for residents." *BMJ Qual Saf* 2013;22:1044-1050 doi:10.1136/bmjqs-2013-001987.

9 Matthews, Robert, and Clare Wilson. "Fever: friend or foe?." *New Scientist* 207.2771 (2010): 42-45. (<http://www.newscientist.com/article/mg20727711.400-fever-friend-or-foe.html>)

10 Schulman, Carl I., et al. "The Effect of Antipyretic Therapy upon Outcomes in Critically Ill Patients: A Randomized, Prospective Study." 2005: *Surgical Infections* Vol. 6, No. 4, Pg. 369-376.

11 Smith, Mark, et al., eds. *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press, 2013. (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America.aspx>)

12 Eddy DM. The Challenge. *JAMA.* 1990;263(2):287-290. (<http://jama.jamanetwork.com/article.aspx?articleid=380215>)

a patient should not depend on which doctor or which hospital system they happen to be at. Today however 50% of the recommendations made in guidelines are based on expert opinion, case studies, or standards of care rather than to more systematic trials and studies.¹³

As another example, annual health checks or standard physical exams are commonplace ... but why? One study from researchers in Denmark poses questions as to their value.¹⁴ Assumptions that we hold about healthcare need to be challenged in order to get to the best possible outcomes for patients.

Even researchers and purported experts have trouble agreeing as to the best practice of medicine. Let's look at screening for certain types of chronic disease. One study showed that among 17 experts on colon cancer, there is significant disagreement on the proportion of colon cancer deaths prevented by screening (as shown below).¹⁵ If there are such massive distributions of opinions among experts whom should we trust?

Yet even with the knowledge that experts frequently disagree on their estimates on guidelines, we over- utilize them in our official recommendations. In another study by researchers at the Division of Cardiology and Clinical Research Institute at Duke University, it was discovered that among the American College of Cardiology (ACC) and American Heart Association (AHA) guidelines, only 11-percent of recommendations exhibit evidence from multiple randomized trials or meta-analyses (level A evidence, the best kind), whereas 48-percent are based on recommendations based on expert opinion, case studies or standard of care (level C

evidence, the worst kind). To make matters worse, only 19- percent of recommendations in Class I guidelines (the most important recommendations) had level A evidence.¹⁶

One striking example (from researchers at the National Heart and Lung Institute of the UK) of an improper cardiovascular guideline involves the use of β -blockers in patients at risk of cardiac events undergoing non-cardiac surgery. A large meta-analysis of over 10,000 patients shows results that contradict the above guideline by indicating that β -blockers actually increase the risk of stroke in patients getting perfusion therapy. Despite being discredited, the AHA guideline has not been retracted as of 2013.¹⁷

David Eddy, who first started preaching the benefits of evidence-based on medicine back in the 1980's, has said that for a long time, "the concept of requiring evidence as the basis of guidelines wasn't even on the board. Guidelines were set by experts sitting in a room ... reaching consensus, and then going out for lunch."¹⁸ While he now believes in the last decade or so that people are "talking the talk" when it comes to evidence-based guidelines, the rate of adoption is still slow.

Some of our most important clinical guidelines are developed based on expert opinions and case studies as opposed to conclusive research. While clinical practice guidelines are often assumed to be the epitome of evidence-based medicine, it turns out that there is significant lack of evidence in many cases and in extreme examples guidelines may be flat wrong. In some fields, up to 40-50% of the recommendations made in guidelines are based on expert opinion, case

13 Chauhan, Suneet P, et al. "American College of Obstetricians and Gynecologists practice bulletins: an overview." *American journal of obstetrics and gynecology* 194.6 (2006): 1564-1572. (<http://www.ajog.org/article/S0002-9378%2806%2900300-0/>)

14 Krogsbøll LT, et al. "General health checks in adults for reducing morbidity and mortality from disease." *Cochrane Database of Systematic Reviews* 2012, Issue 10. Art. No.: CD009009. DOI: 10.1002/14651858.CD009009.pub2.

15 Eddy, D.M. "Variations in Physical Practice: the Role of Uncertainty." 1984: *Health Affairs* Vol 3, No. 2, Pgs. 74-89.

16 Tricoci, P, et al. "Scientific evidence underlying the ACC/AHA clinical practice guidelines." *National Center for Biotechnology Information* (2009 Feb 25). Retrieved 3 Feb 2014 (<http://www.ncbi.nlm.nih.gov/pubmed/19244190>).

17 Bouri, Sonia, et al. "Meta-analysis of secure randomised controlled trials of β -blockade to prevent perioperative death in non-cardiac surgery." *Heart* (2013). (<http://heart.bmj.com/content/early/2013/07/30/heartjnl-2013-304262>)

18 Eddy, David. "Building a 'Terribly Smart' Doc". [http://www.managedcaremag.com/archives/2008/9/building-- %E2%80%98terribly-smart%E2%80%99-doc-conversation-david-eddy-md-phd](http://www.managedcaremag.com/archives/2008/9/building--%E2%80%98terribly-smart%E2%80%99-doc-conversation-david-eddy-md-phd)

studies, or standards of care rather than on more systematic trials and studies.¹⁹ The Institute of Medicine (IOM) acknowledges that the current state of clinical practice guidelines does not meet the promise and rigor of what should be trusted.²⁰ Dr. John Ioannidis illustrates this situation in a both recent and incredibly important set of guidelines, when the ACC/AHA in late 2013 came out with guidelines for advocating for usage of statins as primary prevention of cardiovascular health risk, “The ACC and the AHA are among the most experienced organizations in medicine that develop guidelines. Their processes are meticulous, including transparent reporting of conflicts. The work behind the guidelines’ development was monumental. References to randomized trials and systematic reviews were continuous (the word “evidence” appears 346 times in the cardiovascular risk assessment report and 522 times in the treatment report alone). Panelists were highly qualified. Statins have been extensively evaluated in numerous randomized clinical trials. The guidelines focused on hard clinical outcomes such as myocardial infarction and stroke. Remaining caveats were explicitly acknowledged in documents covering hundreds of pages. However, this apparently seasoned integration of data and opinion eventually would lead to massive use of statins at the population level; ie, “statinization.” It is uncertain whether this would be one of the greatest achievements or one of the worst disasters of medical history²¹.” The New York Times summed up those guidelines, and the way they got incorporated, as follows²² “...a remarkable and sudden departure from decades of advice on preventing cardiovascular disease. According to the new advice, doctors should not put most people on cholesterol-lowering medications like statins based on cholesterol levels alone. And, despite decades of being urged to do so, patients need not monitor their

cholesterol once they start taking medication. The guidelines do not even set target levels for LDL, the so-called bad cholesterol.” The article astonishingly states “The chairman of the committee that developed the new guidelines, Dr. Neil J. Stone of Northwestern University, said the group was prompted to examine the idea of target LDL levels when two doctors — Dr. Krumholz and Dr. Rodney A. Hayward of the University of Michigan — asked what the evidence was for their efficacy. When the committee looked, Dr. Stone said, they found no evidence. It was generally accepted that lower was better, but no one had shown that an LDL of 90 milligrams per deciliter, for example, was better than 100. And the high doses and multiple drugs many patients were taking to get to target levels raised concerns.” So what is the right way to get an answer? While we have less than high probability conclusions, we use data science applied to EHR practice data across as large a variety of patients as we can. I’d call this practice based evidence (leveraging the best judgments of human doctors today). This evidence will evolve and improve continuously.

This is the fundamental problem with expert opinion instead of statistical analysis! Why did it take two doctors to ask the obvious question and why do we trust experts so much? And how many doctors know that people of Indian origin poorly metabolize statins and require much higher dosage than they are given? Incidentally people of Asian origin achieve similar benefits at lower statin doses than Westerners do.²³ In Japan, the typical dosage prescribed for statins can be roughly half that of the US.²⁴ Yet that’s not considered ‘canon’ amongst physicians in the US, even as more and more studies are starting to examine the substantial effect of genetics and ethnicity should have on statin dosage.

19 Chauhan, Suneet, et al “American College of Obstetricians and Gynecologists practice bulletins: an overview.” American journal of obstetrics and gynecology (2006) (<http://www.ajog.org/article/S0002-9378%2806%2900300-0/>)

20 Institute of Medicine. “Clinical Practice Guidelines We Can Trust” Washington, DC: The National Academies Press, 2011 (http://www.nap.edu/catalog.php?record_id=13058)

21 <http://jama.jamanetwork.com/article.aspx?articleid=1787389>

22 Kolata, “New Cholesterol Advice Startles Even Some Doctors”, New York Times (11/13/2013). (<http://www.nytimes.com/2013/11/14/health/new-cholesterol-advice-startles-even-some-doctors.html>)

23 Liao, James K “Safety and Efficacy of Statins in Asians” Am J Cardiol (2007) (<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2651637/>)

24 Are Statins Equally Effective in Minority Populations? (http://www.medscape.org/viewarticle/581848_2)

Randomized trials throughout all fields of medicine refute initial claims and prior medical literature. And in certain cases, that is a good thing, as the progress of science (and the scientific method) does at times require the refutation of previously held hypotheses. But there is a fundamental difference between research and medicine which derives from science or analysis based on practice data and that which derives from opinion and practice. And medicine should have focus on the former, while likely only resorting to the practice of medicine (i.e. human judgment) when our current understanding is limited enough that proper evidence and guidelines are not clear as to how to treat a patient. And that process in and of itself can be scientific – using a probabilistic approach to determine how to make key decisions in care and research.

So while the jury is still out on whether the trials are actually fully accurate (or if the truth is somewhere in between), the lack of understanding and the constant back-and-forth in the research community shows how flawed the current state of medical knowledge is. As a quick example (beyond the statins case) to bolster this, steroids had been (and still are) widely used to treat people with traumatic brain injury. But the largest randomized trial (with close to 10000 participants) actually found a significant increase in the risk ratio of death with steroids.²⁵ While more meta-analysis needs to be done, the refutation of something so standard in a very serious condition like brain injury leaves cause for concern. And the key in the face of uncertainty is not to continue our old practice but rather to have systems that determine continuously the highest probability risk/reward outcome and to continue to accumulate data and science to refine the uncertainty down. Many of these experiments have been inadvertently run and data to guide "best guesses" exists in patient records. As another example, the medical community is going back and forth on the impact of Niacin in reducing cardiovascular disease (CVD).

In the past, niacin was routinely prescribed for CVD prevention, with multiple (non-randomized) studies confirming that. But recently, two very high-profile randomized trials (AIM-HIGH and HPS2-THRIVE) have strongly refuted those studies and have even claimed the adverse side effects are strong enough to warrant not using niacin in the future.^{26,27} Additional studies have come out which refute the refutation, leading to all kinds of confusion in the medical community!²⁸ (which most physicians won't even be able to incorporate, anyway, due to information overload). The best (if not perfect) answer to many, though not all, such questions lies in more data and better analytics derived initially from electronic medical records and over time from more detailed data including biomarkers, physiologic variables, microbiome and genome data based studies and more research using these data sources.

In other words, most studies we base much of medicine on may not be at least "proven valid". In this context it is not surprising that consumers are constantly whiplashed by the medical or research establishment changing their mind with a study du jour contradicting previous studies. Once again, contradictions aren't bad in and of themselves and are a necessary part of the scientific process. But too many times the quality of research isn't properly examined in these studies and a weighing of benefits vs. harm aren't presented in a meaningful way to either consumers or doctors. And doctors as result can't judge the reliability of research that may be relevant or actionable for their patient's conditions. And while some, such as the IOM report on clinical practice guidelines mentioned above, would advocate for better / more rigorous guidelines to fix the problem, transforming medicine from a practice to a science with vastly increased data and experimentation should eliminate a lot of these problems naturally. Current human-based systems will have a lot of problems trying to cover these deficiencies, even though new v0

25 Alderson, P "Corticosteroids for acute traumatic brain injury (Review) The Cochrane Collaboration, 2009, (<http://www.thecochranelibrary.com/userfiles/ccoch/file/CD000196.pdf>)

26 HPS2-THRIVE. (<http://www.thrivestudy.org/>)

27 Lavigne, Paul and Karas, Richard. "The Current State of Niacin in Cardiovascular Disease Prevention" Journal of American College of Cardiology, 2013 (<http://content.onlinejacc.org/article.aspx?articleid=1486719>)

28 <http://www.mayoclinic.org/diseases-conditions/heart-disease/expert-answers/prevent-heart-attack/faq-20058253>

systems like UpToDate and Epocrates are helping. A study in the New England Journal of Medicine showcased that their medical records indicated that amongst a set of adults, the right and up-to-date care occurred just barely more than half the time (54.9%).²⁹

Experts and human judgment being problematic and non-ideal does not only hold true within medicine. Professor Tetlock³⁰ in his book “Expert Political Judgment: How Good Is It? How Can We Know?” describes a twenty-year study in which 284 experts in many fields were followed when making 28,000 predictions about the future, finding that they were only slightly more accurate than chance, and worse than basic computer algorithms. Using his words “they were only slightly better than dart throwing monkeys” and yet we continue to rely on expert opinion because we have little by way of options. It is better than nothing. But as we shall see we have an opportunity to change this in the next decade or two.

To take this further, Dr. Ioannidis has also done research to show that it is “more likely for a research claim to be false than true.”³¹ His meta-research into biomedical research and clinical studies paint a sobering reality on the effectiveness of the studies we are doing. This compounds problems related to the “practice” of medicine by the non-standard, ad-hoc way those guidelines / recommendations from research are disseminated throughout the medical community. Dr. Ioannidis showed across biomedical studies that many of them claim statistical significance of their work where there is actually a high rate

of false positives. And the majority of studies also utilize private datasets, making it impossible to reproduce the results at hand. Dr. Ioannidis theorizes that it's the various forms of biases (over 235 of them cited across 17 million papers!) that contribute to these results, and has developed prediction models based on some of these biases alone to determine the validity of certain studies.³²

³³ We have enough data stored away in a variety of knowledge bases – both in scientific literature and in hospital EMR systems. We can utilize this data to move towards Dr. Ioannidis' vision of evidence-based practice of medicine, by starting with “practice-based evidence” based on these data sources. By measuring outcomes from the variability already inherent in our medical system today, we can build systems and algorithms that take that information and bootstrap an initial model of global guidelines. And then over time these systems can resolve inconsistencies using both further data and strong medical research, starting with the most harmful of claims and conditions, to build an increasingly accurate model.

But for the time being, we should be wary about a lot of the research that gets disseminated into practice and guidelines of medicine. Research papers: 1) will support that almost anything you eat has some ingredients that are associated with cancer (the same substance may be associated with increased cancer risk in one study, decreased risk in another study), 2) have exaggerated claims on the effect of biomarkers and GWAS-identified loci on disease, 3) showcase plenty of industry (and other) bias in drug target identification, etc ...

29 McGlynn EA, et al. “The quality of health care delivered to adults in the United States” New England Journal of Medicine, 2003 (<http://www.ncbi.nlm.nih.gov/pubmed/12826639>)

30 http://en.wikipedia.org/wiki/Philip_E._Tetlock

31 Ioannidis, John P.A. “Why most published research findings are false” PLoS Med, 2005 (<http://www.plosmedicine.org/article/info%3Adoi%2F10.1371%2Fjournal.pmed.0020124>)

32 Ioannidis, John P.A. “Almost all articles on cancer prognostic markers report statistically significant results.” European Journal of Cancer, 2007 (<http://www.ncbi.nlm.nih.gov/pubmed/17981458>)

33 Alsheikh-Ali, Alawi, Ioannidis, John P.A, et al. “Public Availability of Published Research Data in High-Impact Journals” PLoS One, 2011. (<http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0024357>)

(really, the list goes on).^{34, 35, 36}

As a result, research today can be summed up as follows: “Medical practice has evolved out of centuries of theorizing, personal experiences, bits of evidence, expert consensus, and diverse conflicts and biases, [and] rigorous questioning of [these] long-established practices is difficult.”³⁷

Of course, not all of this research is cited / known as canon throughout the medical community. And the science is certainly difficult, and mistakes should and will be made to progress the field. But the rate of error right now is too high. Dr. Ioannidis’ proposed solutions for more accurate research revolve around increased data sharing, rigorous standards, and a greater focus on meta-analyses and randomized trials. Though valuable and necessary, a lot of this we argue later can be bolstered naturally as a result of technology driving medicine to being a science. Currently, available data in the electronic health record makes more accurate diagnosis and prescription possible. Better biological research and orders of magnitude more physiologic and biomarker data will in the future make this situation substantially better. What happened to the hundreds or thousands of “very similar” (age, demographic, location, genetics, food habits, etc.) patients that have been treated and how effective was each treatment and what does it indicate regarding the probability distribution of disease paths, treatment effectiveness, and disease management? With this “chart” the patient preferences can be effectively considered in simplification qualitative input tools to reduce these insights to patient or doctor preferences as we shall see later! Add to this the knowledge graph extracted from 50-100m research articles and one has the basis to start to

invent a new medicine.

The Practice of Medical Diagnosis

In other areas, moving downstream from research, the practice of medicine is even more noticeable. Doctors frequently disagree on diagnoses. For example, in the field of psychiatry, research has found that psychiatrists using the Diagnostic and Statistical Manual of Mental Disorders (DSM), the standard desk reference for psychiatric diagnoses, have dangerously low diagnostic agreement in many instances. The DSM uses a statistic called the “kappa” to measure the level of agreement between psychiatrists ranging from zero for no agreement and one for complete agreement. In research trials, the DSM V generates a kappa of 0.2 for generalized anxiety disorder and 0.3 for major depressive disorder.

Scientific American described these results for the standard of psychiatric care as “two pitiful kappas”.³⁸ Instead, imagine if we could measure thousands of activities each day and use that data for insights on a patient’s mental state and wellbeing rather than guessing at diagnoses that often change from doctor to doctor. Diagnostics would become much more of a science as we shall see later. Incidentally, in the United States’ fee for service world of healthcare the DSM psychiatry manuals would incorporate symptoms like my leg shake as a mental health “disease” so psychiatrists can bill for it! Often, there are errors of omission where a diagnosis is missed entirely.³⁹ Cases where diagnoses are missed or incorrect could impact up to 10 to 20 percent of cases. It’s also shocking how frequently getting a second opinion changes the diagnosis. If this was really the science of medicine, then

34 Schoenfeld, Jonathan, Ioannidis, John P.A. “Is everything we eat associated with cancer? A systematic cookbook review.” American Journal of Clinical Nutrition, 2012. (<http://ajcn.nutrition.org/content/early/2012/11/27/ajcn.112.047142.full.pdf+html>)

35 Ioannidis, John P.A., Panagiotou, Orestis. “Comparison of Effect Sizes Associated with Biomarkers Reported in Highly Cited Individual Articles and in Subsequent Meta-analyses” JAMA, 2011


36 Prinz, Florian et al “Believe it or not: how much can we rely on published data on potential drug targets?” Nature Reviews, 2011 (<http://www.nature.com/nrd/journal/v10/n9/full/nrd3439-c1.html>)

37 Prasad, Vinay, Cifu, Adam, Ioannidis, John P.A. “Reversals of Established Medical Practices: Evidence to Abandon Ship” JAMA, 2012 (<http://jama.jamanetwork.com/article.aspx?articleid=1104821>)

38 Jabr, Ferris. “Field Tests for Revised Psychiatric Guide Reveal Reliability Problems for 2 Major Diagnoses.” Scientific American (6 May 2012). Retrieved 3 Feb 2014 (<http://blogs.scientificamerican.com/observations/2012/05/06/field-tests-for-revised-psychiatric-guide-reveal-reliability-problems-for-two-major-diagnoses/>).

39 Eddy DM. The Challenge. JAMA. 1990;263(2):287-290. (<http://jama.jamanetwork.com/article.aspx?articleid=380215>)

Value of second opinion



diagnoses should be consistent. One study of a second opinion program in oncology for Partners Healthcare System in Boston resulted in a new plan or a significant change in prior treatment in 90% of cases (although the sample may be biased because the patients were all seeking second opinions).⁴⁰ And in a separate study by the Cleveland Clinic, in 11% of the cases the wrong diagnosis was given, and in 44% of cases enough of a difference was found to suggest moderate to major changes in treatment plan. And an additional 15% of those Cleveland Clinic patients needed additional testing. Is this an acceptable quality of care?

Net-net, patient outcomes are far inferior to and more expensive than what they should be. According to the Institute of Medicine the non-profit health arm of the National Academy of Sciences, 30-percent or \$750 billion is wasted annually in US medical spending.⁴¹ Despite the massive dollar amounts poured into the system, the current benchmarks of performance still aren't good enough. It's trivially easy to find study after study that demonstrates the shortcomings of the practice of medicine. A Johns Hopkins study found that as many as 40,500 patients die in

Intensive Care Units (ICU) in the U.S. each year due to misdiagnosis.⁴² This is equivalent to Google being allowed to introduce a driverless car if it only killed a hundred people a day in car accidents. For framing, the number of fatalities due to ICU misdiagnosis rivals the number of deaths from breast cancer. In the study, the investigators, led by Bradford Winters, examined studies (since 1966) that utilized autopsy to identify diagnostic errors. Class I diagnostic errors, which constitute "missed major diagnoses with potential adverse impact on survival," made up 8% of the patients. These were cases where if the doctors had been aware of the proper diagnosis, therapy and treatment "would have changed." In total, over one in four (28%) of the autopsies had at least one misdiagnosis. Another review of studies in postmortem research suggested that diagnostic errors are implicated in one of every ten patient deaths.⁴³ Yet another study found that system-related factors like poor processes, teamwork and communication were involved in 65-percent of studied diagnostic error cases. Cognitive factors were involved in 75-percent of diagnostic errors with premature closure, which maps to confirmation bias, the practice of persisting with the initial diagnosis

40 Wojcik, Joanne. "Firms seek medical second opinions." *Business Insurance* (4 Dec 2011). Retrieved 3 Feb 2014 (<http://www.businessinsurance.com/article/20111204/NEWS05/312049987>).

41 "Best Care at Lower Cost: The Path to Continuously Learning Health Care in America" Institute of Medicine (6 Sept. 2012). Retrieved 3 Feb. 2014 (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America/Report-Brief.aspx>).

42 Winters, Bradford, et al. "Diagnostic errors in the intensive care unit: a systematic review of autopsy studies." *BMJ Quality and Safety* (21 July 2012). Vol. 21 Issue 11 Retrieved 3 Feb 2014: <http://qualitysafety.bmj.com/content/21/11/894.abstract?sid=647c745e-4b1f-45e4-8883-a700adfc0005>.

43 National Academies of Sciences, Engineering and Medicine, 2015 "Improving diagnosis in health care." Washington, DC: The National Academies Press

and ignoring reasonable alternatives as the most common cause.⁴⁴ Given this, it's not surprising that about 5 percent of adults who seek outpatient care annually suffer a delayed or wrong diagnosis.

Between 44,000 to 98,000 patients die each year from preventable medical errors.⁴⁵ An even when there are not direct errors, knowledge about serious diagnoses very frequently could be found within the medical record. For example, one study created a knowledge base from a diagnostic decision support system to identify high-information clinical findings (HIF's) for a set of selected high-risk diagnoses (HRD's), including myocardial infarction, appendicitis, and a set of carcinomas.⁴⁶ In this study, 25% of the records reviewed contained the HIF's in notes before the HRD was identified. While the study listed a range of reasons why this delay occurs (e.g. non-compliance on the patient side), the relevant take-away is that very frequently, data lives in the medical record that can inform us ahead-of-time about serious problems, and the current medical system does not utilize that information effectively. Another review from the Institute of Medicine indicated that diagnostic errors account for up to 17 percent of adverse events within a hospital.⁴⁷

It is not reasonable to expect a mere human to put all the data in each patients full medical history and to put together every possible combination of symptoms / diagnosis / prescriptions in their head as such "maps of medicine" can be complicated and obtuse. The human body and its diseases are complex, often too complex for unaided understanding by even the smarter humans. The point isn't to critique doctors or medical processes but to point to the complexity that can confound

good decision making except if machines were doing the diagnosis given the magnitude of data, knowledge and possibilities involved.

Some diagnostic errors stem from mistakes in the interpretation of diagnostic tests. For example, pathology, radiology, and the clinical laboratory have error rates of 2% to 5%.⁴⁸ Superimposed on these testing errors are the ubiquitous system-related errors encountered in every health care organization, as well as the cognitive errors indicated above. Diagnostic errors do not occur only in connection with unusual conditions but span the breadth of clinical medicine, from rare disorders to commonplace ones like anemia and asthma.

The Practice of Medical Treatment

Misdiagnosis is only the tip of the iceberg. Compounding the effect of improper diagnosis is miscommunication between doctors and patients. One study found that failure to communicate or document an abnormal test result is unfortunately quite common. In approximately seven-percent of cases, primary care physicians failed to inform or to document informing a patient of an abnormal test result. For one in every 14 tests across the ~550 patients in the meta-study, either the patient was not informed of a clinically significant abnormal test result, or the clinician failed to record reporting the result to the patient. And EMR systems don't automatically make things better (in fact, could exacerbate these mistakes), particularly if there were already poor processes / communication between the physician and patient beforehand.⁴⁹ As many as one-third of hospitalized patients may experience harm or an adverse event, often from preventable errors,

44 Graber, Mark L., Nancy Franklin, and Ruthanna Gordon. "Diagnostic error in internal medicine." *Archives of Internal Medicine* 165.13 (2005): 1493. (<http://archinte.jamanetwork.com/article.aspx?articleid=486642>)

45 Kohn, Linda T., Janet M. Corrigan, and Molla S. Donaldson, eds. *T err is human: building a safer health system*. Vol. 627. National Academies Press, 2000 (<http://www.iom.edu/Reports/1999/to-err-is-human-building-a-safer-health-system.aspx>)

46 Feldman, Mitchell J., et al. "Presence of key findings in the medical record prior to a documented high-risk diagnosis." *Journal of the American Medical Informatics Association* 19.4 (2012): 591-596. (<http://jamia.bmj.com/content/19/4/591>)

47 National Academies of Sciences, Engineering and Medicine, 2015. "Improving diagnosis in health care." Washington, DC: The National Academies Press

48 Graber, Mark L., Robert M. Wachter, and Christine K. Cassel. "Bringing Diagnosis Into the Quality and Safety Equations." *JAMA* 308.12 (2012): 1211-1212. (<http://jama.jamanetwork.com/article.aspx?articleid=1362034>)

49 Casalino, Lawrence P., et al. "Frequency of failure to inform patients of clinically significant outpatient test results." *Archives of internal medicine* 169.12 (2009): 1123. (<http://archinte.jamanetwork.com/article.aspx?articleid=415120>)

according to these studies.⁵⁰ Serious preventable medication errors, although tough to measure, are conservatively estimated to impact over 7 million patient visits, costing over \$15 billion a year in the US.⁵¹ So even if the initial diagnosis might be fine, the set of downstream problems that can result all too (prescription errors just being one example) lead to a much more dysfunctional healthcare system than ideal. An adverse events in hospitals are potentially measured at a rate that is an order of magnitude below the actual number of events. Voluntary reporting can miss up to 90% of adverse events, and these events occur in 1/3 of hospital admissions.⁵² As a result, we see countless examples of hospitals and systems that don't improve patient safety and quality of care over time.⁵³ Many of these errors can be reduced by digital health systems as more and more of the critical decision making is done by the system.

Why this error rate? Clinicians in intensive care units, who care for the sickest patients in a hospital, must manage in the range of 180 activities per patient per day—from replacing intravenous fluids, to administering drugs, to monitoring patients' vital signs. Are errors across multiple healthcare providers taking care of each patient, multiple shifts across which activities must happen and be coordinated and historical and copious information in the patient's health record a surprise?⁵⁴ According to one study, 10% to 20% of cases have delayed, missed, and incorrect diagnosis.⁵⁵

To make matters worse, US Medicare patients

now see an average of seven physicians. Imagine balancing the conflicting recommendations of clinical practice researcher guidelines with as many as seven doctors.⁵⁶ Today, approximately 75 million people in the U.S. are living with multiple chronic conditions and account for an even larger percentage of medical spending. According to a recent study, clinical practice guidelines rarely account for or contain modifications for these comorbid patients making it even more difficult to coordinate care. For instance, a 79-year old woman with osteoporosis, osteoarthritis, type II diabetes, hypertension, chronic obstructive pulmonary disease, and peripheral neuropathy would receive conflicting recommendations on treatment. The guidelines for osteoporosis would recommend that she perform weight-bearing exercise while the guidelines for diabetes would recommend she avoid such exercise.⁵⁷ According to the authors "the use of various clinical practice guidelines developed for single diseases may have adverse effects" in such patients. What if comprehensive software was available to manage and predict the highest probability "best treatments" based on medical data from hundreds of millions of patients? What if the primary care physician could replace the seven specialists to offer more integrated care at lower cost and higher quality because of the help of systems that understand each specialty and help guide the integrative view of generalist physicians? We will postulate that the answers as to the best courses of action are buried in the health records data of the population as a whole and other current and future data and insights derived from such data. A

50 Levinson, Daniel R., and Inspector General. "Adverse events in hospitals: national incidence among Medicare beneficiaries." Department of Health & Human Services (2010). (http://www.mhakeystonecenter.org/documents/january2011/tab_va_nov_2010_hhs_adverse_events_study.pdf)

51 Massachusetts Technology Collaborative (MTC) and NEHI, 2008. "Saving Lives, Saving Money: The Imperative for CPOE in Massachusetts." NEHI (2008). (http://www.nehi.net/writable/publication_files/file/cpoe20808_final.pdf)

52 Classen, David C., et al. "Global trigger tool shows that adverse events in hospitals may be ten times greater than previously measured." *Health Affairs* 30.4 (2011): 581-589. (<http://content.healthaffairs.org/content/30/4/581>)

53 Landrigan, Christopher P., et al. "Temporal trends in rates of patient harm resulting from medical care." *New England Journal of Medicine* 363.22 (2010): 2124-2134 (<http://www.nejm.org/doi/full/10.1056/NEJMsa1004404#t=articleBackground>)

54 Donchin, Y., et al. "A look into the nature and causes of human errors in the intensive care unit." *Quality and Safety in Health Care* 12.2 (2003): 143-147. (<http://qualitysafety.bmj.com/content/12/2/143>)

55 Graber, et al., *JAMA* (2012) <http://jama.jamanetwork.com/article.aspx?articleid=1362034>

56 Pham, Hoangmai H., et al. "Care patterns in Medicare and their implications for pay for performance." *New England Journal of Medicine* 356.11 (2007): 1130-1139. (<http://www.nejm.org/doi/full/10.1056/NEJMsa063979>)

57 Boyd, Cynthia M., et al. "Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases." *JAMA: the journal of the American Medical Association* 294.6 (2005): 716-724 (<http://jama.jamanetwork.com/article.aspx?articleid=201377>).

word of caution in health record data is in order. In the U.S. at least, many such records are principally used for billing purposes and many people speculate that these records are biased towards billable diagnosis.

As various papers elaborate, the use of various clinical practice guidelines (CPG's) developed for single diseases may have adverse effects. Incentives and initiatives such as pay-for-performance can undesirably affect CPG's in caring for older individuals. Specifically, these CPG's often fail to modify recommendations based on comorbidities (which end up being more common in the elderly). So in this hypothetical 79 year old woman with a set of conditions (including COPD, type 2 diabetes, and hypertension), the relevant CPG's would allow for the patient to be prescribed 12 medications with a complicated regimen, with no indication of adverse drug and disease interactions⁵⁸. In the U.S. health care system does not make it particularly easy to treat these individuals, even as incentives are slowly getting better. There are plenty of unanswered questions in dealing with these patients, from how actual adherence rates impact the effects of drugs to incorporation of patient preferences into the guideline-driven prescriptions.⁵⁹ Electronically driven systems will find it easier to encompass this complexity to provide more integrative and personalized recommendations. Approximately 75 million people nationally have 2 or more of these concurrent chronic conditions, so our small (albeit growing) understanding of these interactions is a

huge hole.⁶⁰ As we begin to understand these patients more deeply, they need to become a first-class part of how physicians practice (which will require new data and systems to manage this added complexity in a seamless way).

And how do patients feel about their choices? Fewer than half of patients receive clear information on the benefits and trade-offs of treatments for their condition, and fewer than half are satisfied with their level of control in medical decision making.⁶¹ Algorithmic systems (what I call Dr. Algorithm!) will hand much more power to the patient to ask questions, explore options, choices and consequences in common language making the patient the CEO of their own health!

The Science of Medicine

We currently romanticize the doctor as the all-knowing caring physician that's able to understand our individual state and give us the right prescriptions. And while parts of that is true for the small percentage of us that are well off, this "ideal" physician is not the reality for many if not most of the world's seven billion patients. Yet we can use a scientific process to get significantly closer to this "ideal" physician. Such a system may still need the human touch but the quality of decision-making in diagnosis and treatment should far surpass the current median physician. According to one study the traditional systems for transmitting new knowledge—the ways clinicians are educated, deployed, rewarded, and updated—can no longer

58 Boyd, Cynthia M., et al. "Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases." *JAMA: the journal of the American Medical Association* 294.6 (2005): 716-724. (<http://jama.jamanetwork.com/article.aspx?articleid=201377>)

59 Tinetti, Mary E., Sidney T. Bogardus Jr, and Joseph V. Agostini. "Potential pitfalls of disease-specific guidelines for patients with multiple conditions." *Engl J Med* 351.27 (2004): 2870-2874. (<http://www.ih.gov/nc4/Documents/GuidelinePitfallsNEJM.pdf>)

60 Parekh, Anand K., and Mary B. Barton. "The challenge of multiple comorbidity for the US health care system." *JAMA: the journal of the American Medical Association* 303.13 (2010): 1303-1304. (<http://jama.jamanetwork.com/article.aspx?articleid=185601>)

61 Degner, Lesley F., et al. "Information needs and decisional preferences in women with breast cancer." *JAMA: the journal of the American Medical Association* 277.18 (1997): 1485-1492. (<http://jama.jamanetwork.com/article.aspx?articleid=416136>)

Fagerlin, Angela, et al. "Patients' knowledge about 9 common health conditions: the DECISIONS survey." *Medical Decision Making* 30.5 suppl (2010): 35S-52S. (http://mdm.sagepub.com/content/30/5_suppl/35S)

Olsen, LeighAnne, Robert S. Saunders, and J. Michael McGinnis, eds. *Patients Charting the Course: Citizen Engagement in the Learning Health System: Workshop Summary* National Academies Press, 2011. (<http://www.iom.edu/Reports/2011/Patients-Charting-the-Course-Citizen-Engagement-in-the-Learning-Health-System-Workshop-Summary.aspx>)

Lee, Clara N., et al. "Are patients making high-quality decisions about breast reconstruction after mastectomy?[outcomes article]." *Plastic and reconstructive surgery* 127.1 (2011): 18-26. ()

Sepucha, Karen R., et al. "How does feeling informed relate to being informed? The DECISIONS survey." *Medical Decision Making* 30.5 suppl (2010): 77S-84S. (http://mdm.sagepub.com/content/30/5_suppl/77S)

keep pace with scientific advances.⁶² There clearly are fundamental shortcomings in the older approach to medical education and the transmission and adoption of new knowledge and practices as they come out of medical research but the problems go well beyond that.

In order to achieve innovation, we also must reconsider our outlook on experimentation with regard to medicine. We must increase the rate of experimentation that will have overall benefits even if they may occasionally have measured and understandable negatives. Should we choose to save 1,000 lives if it does some damage to 100 lives in the process? Or should we follow *primum non nocere* and “do no harm”? If we care about overall social good, our choice should be clear even though a few individual patients may not benefit. At the very least, the choice of whether to experiment should be the patient’s (with full information), not the doctor’s, since they are the ones who experience the upside or downside of any course of action. Patients need to be armed with a more complete understanding of the pros and cons for making such medical decisions. Some data indicates that when empowered with full information, patients tend to choose less aggressive therapies than the medical system might choose for them.⁶³ Are we fully considering patient preferences and scientific outcomes data or are we mostly going by “doctor knows best” biases?

It’s unsurprising that approaches to medicine vary so significantly from doctor to doctor as this attitude is built into and reinforced by the ethos of the industry through the Hippocratic Oath 98-percent of medical students swear to some form of this oath upon graduating from medical school, “I will prescribe regimens for the good of my patients according to my ability and my judgment and never do harm to anyone.” Instead of using a mathematical (and therefore objectively

accurate) computation of expected harm versus benefit in light of patient preferences, the oath implicitly allows the doctor to make a subjective recommendation for an aggressive or passive treatment option based on their personal biases. It does not take into account patient preference and risk-reward calculations. It is time that we use math, data, and science to drive decision-making and remove emotional subjectivity and variability (except when it is the patient’s emotional and subjective preferences). Each health care system in the US has its own protocols, but there can only be one “highest probability” answer based on a patient’s condition, history, objectives, and their preferences. Most such care patterns must necessarily be suboptimal given a patient, as there is only one best care path for a patient! Valid question, which we try and answer throughout future sections of the paper, would be “why now”? Diagnostic and treatment options are expanding and changing at an accelerating rate, placing new stresses on clinicians and patients, as well as potentially impacting the effectiveness and efficiency of care delivery. The growing rate of new knowledge generation will lead to change as well. Currently, it would take ~21 hours per day for individual primary care physicians to provide all of the care recommended to meet their patients’ acute, preventive, and chronic disease management needs.⁶⁴ So new systems have to be created to handle that information effectively and offer more considerate care with fewer burdens on the physicians’ time.

Yet still, unfortunately, there is too much of a gap between physician knowledge (even as flawed as it may be) and patient knowledge to really enable patients to be scientifically informed consumers and in proper control of their own health. For example, women who had treatment for breast cancer believed less than half the time “they had achieved their preferred level of control in decision making.”⁶⁵ And another study (although

62 Smith, Mark, et al., eds. *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press, 2013. (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America.aspx>)

63 Smith, Mark, et al., eds. “Best care at lower cost: the path to continuously learning health care in America.” National Academies Press, 2013. (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America.aspx>).

64 Yarnall, Kimberly SH, et al. “Peer Reviewed: Family Physicians as Team Leaders: “Time” to Share the Care.” *Preventing Chronic Disease* 6.2 (2009). (<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2687865/>)

65 Degner, Lesley F., et al. “Information needs and decisional preferences in women with breast cancer.” *JAMA: the journal of the American Medical Association* 277.18 (1997): 1485-1492. (<http://jama.jamanetwork.com/article.aspx?articleid=416136>)

only sampling <100 people) showed that early-stage breast cancer survivors actually knew the answers less than 40% of key questions related to knowledge about breast reconstruction after their mastectomy.⁶⁶ More broadly than that, patients have wide variation in knowledge about key decision-related facts for very common serious surgery, screening, and prescription decisions. For instance, few respondents in one particular study knew the most common side effect of cholesterol drugs or had an appropriate understanding of how much impact medication would have on the risk of heart disease.⁶⁷ To compound matters, patients currently believe they are much more knowledgeable about their health than they actually are, with a negative correlation for information needed for cancer screening decisions and no relationship at all for medication decisions. The one common factor is that consumer's trust in the doctor was associated with feeling extremely well-informed, regardless of the decision type.⁶⁸

And so their actual lack of information leaves plenty of room for improvement. Therefore, when moving medicine from the "practice of medicine" to the "science of medicine," we should take particular care to ensure that the patients involved at the heart of the system are given the ability to bring their own judgments, preferences and biases to help make informed decisions (patient biases should be taken into account!) and be able to ask lot of questions and explore options.

In fact, since 2001, the Institute of Medicine has claimed that much of health care in the United States has lost its focus on the patient. We should move to a world of "patient-anchored care" that is efficient, participatory (for patients), and personalized.⁶⁹ Progress is shifting towards data-driven and consumer-driven medicine, and this shift will happen rapidly, consequentially, and in surprising (yet overall extremely beneficial) ways, giving the patient much more control over their healthcare decisions.

66 Lee, Clara N., et al. "Are patients making high-quality decisions about breast reconstruction after mastectomy?[outcomes article]." *Plastic and reconstructive surgery* 127.1 (2011): 18-26 (<http://www.ncbi.nlm.nih.gov/pubmed/21200195>)

67 Fagerlin, Angela, et al. "Patients' knowledge about 9 common health conditions: the DECISIONS survey." *Medical Decision Making* 30.5 suppl (2010): 35S-52S (http://mdm.sagepub.com/content/30/5_suppl/35S.full.pdf+html)

68 Sepucha, Karen R., et al. "How does feeling informed relate to being informed? The DECISIONS survey." *Medical Decision Making* 30.5 suppl (2010): 77S-84S. (http://mdm.sagepub.com/content/30/5_suppl/77S)

69 Olsen, LeighAnne, Robert S. Saunders, and J. Michael McGinnis, eds. *Patients Charting the Course: Citizen Engagement in the Learning Health System: Workshop Summary* National Academies Press, 2011. (<http://www.iom.edu/Reports/2011/Patients-Charting-the-Course-Citizen-Engagement-in-the-Learning-Health-System-Workshop-Summary.aspx>)

HEALTHCARE – INNOVATION EMERGING OUT OF COMPLEXITY

It is nearly impossible to predict the future of complex systems, even in the short-term, but especially 20 years out into the future. In such systems, innovation typically arises as an emergent property out of thousands of new technologies and ways of doing things, many of which may not even exist at the time prediction gets made. This is much like a tornado “emerging” from winds or waves becoming a tsunami and changing the scale at which effects happen quite non-intuitively. There will be positive interactions between digital health discoveries and more fundamental discoveries and understanding from the biological sciences, other sciences and randomized field trials. In the vast majority of cases, new digital health efforts will hit brick walls and bounce off. In other cases, new ideas will make some initial headway but eventually only become niche solutions in small application areas. The rest, typically very small percentage of everything attempted, will begin to stick in some, sometimes surprising, but fundamental ways, and start to grow exponentially in their impact. These efforts will like be the Google’s, Facebook’s and Twitter’s of the dotcom era. Momentum will build behind them and develop, over time, into a wave of innovation. Steven Johnson best describes this phenomenon of “slow hunches”, whereby the initial versions of ideas struggle and fail, with insight often emerging out of serendipity and leading to new versions that become the right solution as “adjacent possibles” build on each other as obvious next steps to initial ideas.⁷⁰ In these types of systems, complexity begets progress because it creates the right kind of “chaotic soup” environment in which thousands of new ideas can be tried, and every idea is catalyzed by other ideas in this soup of ideas. This is often the nature of innovation.

Healthcare is the perfect example of this type of complex system. One could hardly design a more complicated, chaotic system of human biology, where a patient’s medical state could exist in any of many of different possible combinations of conditions, treatments, and co-morbidities over time. With an ecosystem of so many different actors all motivated by different incentives and optimizing for different objectives, along with evolving scientific understanding, changing government regulations, variations in disease, diagnoses, treatment protocols, insurance coverage, and billing, and personal preferences for patients, doctors, and everyone else, all mixed together in one massive chaotic endeavor. Linear medicine, isolated non-holistic testing of drugs, therapies, and procedures do poorly in this environment though they are better than the “do nothing” or “alternative or traditional medicine” alternative in my opinion. Though better than nothing, they will fail to match the efficacy of “real trials” beyond today’s randomized trials that actual medical practice data will provide. We will have these “post approval” trials that define efficacy for different circumstances instead of one time approval or recall of treatments.

Even with all this chaos, it is important to recognize and pay homage to medicine’s progress over the last fifty years. There has been a massive leap in improving human health, longevity, and curing serious disease. Yet even with this big wave of health improvement so far, this next wave of data-driven care might be substantially larger, resembling a tsunami of change for the better, being more efficacious by being better science, more personalized and precise.

Just as a tsunami, I suspect that fundamental innovation in healthcare is unlikely to happen

⁷⁰ Johnson, Steven. *Where Good Ideas Come From: The Natural History of Innovation* New York: Penguin, 2010

in step functions; it's much more likely to arise as an emergent property out of the chaos and complex system that's really just starting to develop with all the different players, sub-systems, and new innovations interplaying. It will most likely happen very much along lines of trial and error leading to insight and eventual waves of transformation that coalesce into a fundamental reimagining of medical science and practice 10-20+ years out. The fits and starts punctuated equilibrium model of biological evolution may be a good model for digital health progress. Another analogy highlighting the fits and starts of mobile phone evolution (v0 through v7) is explained in detail later. This healthcare transformation starts with the v0 efforts that we've seen over the past and will see in the next few years, which are largely basic, fairly obvious enhancements and digitization of standard healthcare practices that provide some efficiency gains but do little else. The next couple of years or so in healthcare will see what are mostly infant, v0 efforts that have now begun to emerge in earnest and more directly challenge assumptions and "traditional" medical practice by raising critical 1st-order questions about what modern healthcare should look like. People will try a lot of different ideas and approaches with different models of how these new technologies work in isolation. In this stage, we should remember that many of these initial systems will look rudimentary ("toddler systems"). So nearly all of these initial efforts will go through multiple iterations of v1/v2/v3/v4 and more as they bump up against roadblocks and points of failure until they begin to mature and develop real momentum behind them. At different points in time, we'll see things start to click and begin to work more smoothly, gain momentum, and lead to larger changes that build on themselves over time like a wave. Some areas, for example certain aspects of drug discovery, may not get much better and it may still take 10-15 years to develop new pharmaceuticals (even that is hard to state with certainty: see v7 drug discovery described later). But each subsequent version of digital healthcare innovation in a given area will likely only take couple years each to emerge. No one really knows how long this process will take, but I suspect this fundamental reshaping of medicine could take 5-10 innovation cycles in each area, with iteration happening rapidly in some areas and more slowly in others depending on mitigating

factors like regulation, incentives, and other factors. We will later look at how each area will evolve based on technology progress in certain dimensions. But it is worth keeping the cellphone evolution or toddler to expert analogies in mind as we observe the evolution of the ecosystem.

"The Future in Stages"

If one were to look back at medicine in two or three decades, the first iteration with 20/20 hindsight might look like a stage where we see enhancements to traditional medicine with better blood glucose tracking, better software to manage patients, more support tools for today's doctor or health professional, better patient compliance tools, better information coordination, etc. We are already starting to see point innovations like ECG machines that are always available and able to analyze patients ECG's, or digital otoscopes with auto-diagnosis that allow better connection between ear imaging to the digital world with image capture and availability, and handheld or wearable ultrasound imagers connected to mobile devices with cloud-based diagnosis within minutes. These point innovations will make each specialty better with much better heart or ear infection diagnosis and remote monitoring (which doctor wants to see 20 ECG's per patient every month? It is too much work.) because of more frequent and consistent monitoring to manage patient conditions and some new medicine within each point innovation. For example with heart rate monitoring from wearable devices we might start to see heart rate variability more frequently used as an indicator of heart health, or find longer term changes in resting or sleeping heart rate or changes in heart rate variability patterns from hundreds of millions of users across billions of hours of monitoring as a very early indicator of heart conditions, or prescribe antibiotics less because patient with an ear infection can do virtual visits every twelve hours and only get an antibiotics prescription when other alternatives are not working. We might monitor pregnancy more closely with continual inputs from the pregnant mother that are monitored by systems which anticipate problems or have better protocols for diabetic care. Many though not all point innovations will extend the reach of the health care professional and the quality of information they have in a cost effective and likely familiar

way. As systems progress they might do more than is possible today under the supervision of a healthcare professional. Soon new ways of practicing, monitoring, treating, behaviorally modifying, and even diagnosing will start to emerge.

But the biggest contribution to this digitization of medicine will be the accumulation of data. It will accumulate in patient files, patient wearables data, social, other environmental data, and research data. The consumer's health as "wellcare" will start to merge with today's "sickcare" to move towards a broader definition of healthcare, slowed only by the reticence of the establishment and the initial lack of formal medical trials data. The regular emergence of new medical practice may seem surprising today, but will be the norm in the future. Heart monitoring systems will use all this available data to predict heart attacks hours if not days or years in advance. You won't have heart disease or not or "have diabetes" or "not have diabetes". These diseases will be classified in the continuum between the two states. Not that cardiology is "wrong" today but if your heart rate while sleeping is creeping up over two years as measured by your watch or health band will that not be an early telling sign of some change?

The decision support systems of the future will look completely unrecognizable next to today's systems. Today the majority of these systems take pride in being rules-based, such as DXplain used by Mass General, or Isabel's Differential Diagnosis Generator^{71,72}. But with an understanding of all this data that is being accumulated, diagnosis systems in the future will be automated, dynamic sets of algorithms.

Rather than trying to capture human medical knowledge in the form of rules, these "Dr. Algorithms" will evaluate complex data that humans today wouldn't be able to comprehend, and incorporate all that information to come up with the right set of diagnoses and treatment plans. These algorithmic doctors will exist per specialty (i.e. there will be a canonical set of auto-updating algorithms for cardiology) as well

as within internal medicine, and will show up first as optimal care paths (or care guidelines) or as personalized care paths for each individual. More on precision and personalized medicine later.

Today, medical research and our understanding of the human body is flawed. LDL cholesterol is bad. LDL cholesterol does not correlate to heart attacks. Saturated fat is bad. Saturated fat is less bad than carbs.

Saturated fat may be neither bad nor good the litany goes on because of many cognitive biases, leading to (sometimes unintentional) flaws in statistical studies appropriately motivated by financial or "publish or die" motivations in academia. The accumulation of data will start to make such claims more easily (it is never easy though) testable and verifiable. To begin with we will start to see verification (or not) of the claims of the ~ 1500 or so drugs, procedures, treatments that medical practice has at its disposal by re-running the claims through medical records data and real life testing in the complexity of actual patient use patterns. Though it won't always be possible to do this I suspect the vast majority of medical practice will go through such re-validation and many existing treatments will be de-bunked or appropriately narrowed to specific cases (e.g. x helps for disease y if patient has a,b,c conditions). Re-definition of how we diagnose, even classify and treat disease will start. Diabetes may be re-classified as six or twelve different diseases with very different treatment protocols (and by examining EHR data, we have already started to clearly see this split, though without understanding underlying mechanisms yet). We will start to do new things with continuous blood pressure, continuous blood glucose, continuous heart rate, heart rate variability, cardiac output, skin temperature, galvanic skin resistance, voice tonality analysis, metabolic rate, blood oxygen levels, etc as these indicators become inexpensive and routinely available. Other continual measurements like ECG, urine analysis (e.g sodium in urine measured every morning automatically by your toilet seat?), and breath analysis (ketones in breath related to your

71 <http://www.mghlcs.org/projects/dxplain/>

72 <http://www.isabelhealthcare.com/>


metabolic status or diabetes, nitrous oxide in your breath related to your asthma or COPD, or even your lung or breath cancer detected in your breath) will be incorporated. Other periodic measurements and tests such as metabolomics, genomics, epigenetics, imaging studies will contribute to the data pool. These measurements will be tied to a patient's EHR data with detailed physician notes, and will be used alongside fitness bands and mental health monitoring software applications and testable expert opinion from the very best doctors and researchers in order to generate new insights that will fine tune treatment for each disease, often specifically tailored to the patients goals.

Thousands of biomarkers per blood samples and thousands of microbiome species data will add significantly additional dimensions to understanding medicine, though all the data and the complex insights and tradeoffs will be beyond what humans can reasonably comprehend without significant simplification. Machine learning will dramatically augment our current knowledge repository of medicine by incorporating all of this new, previously unrecognized or unrecorded continuous, continual or otherwise unmanageable data, into medicine's practice (or science). This jumble of data or thousands of microbiome

species, hundreds of thousands of physiologic data points (genomic, epigenetic, transcriptomic, metabolomic, microbiome, exposome physiologic, ... sometimes called all the "omics") and wearable data, mental health data points and more will be turned into actionable insights. Historical medicine, developed a hundred years ago, are still primary modes for understanding symptoms today – point measurements of heart rate, blood pressure, and temperature. But so much more is now possible, especially as we find out how to measure and interpret the results (starting by measuring those variables continuously and longitudinally). That new data plus more sophisticated analytics and new forms of data will drive new medicine. It is possible to imagine leveraging the gut microbiome (and companies are already attempting this) coupled with sophisticated machine learning to create a new form of "medical practice" or at least new biotic drugs based on new data. And these will have a range of different advantages from traditional pharmaceuticals or treatments and will often complement those historical tools – each case or condition will be different (and not all will be better!), but already, we are seeing evidence that we can leverage microbiome data to predict diseases earlier, continuously monitor conditions

Humans don't have built-in diagnostics

car




- 400+ on-board sensors
- 750MB/s data processed in google self-driving car

iPhone



- 10+ on-board sensors
- 4 radio

human



- 0 on-board sensors
- annual checkup (maybe)

noninvasively, and develop better therapies for certain conditions (e.g. skin conditions such as better treatment for acne). And this data, coupled with continuous wearable, physiologic data, biomarker data, and data from a range of other sensors and measurements will change medicine.

A few diseases will start to be classified by their “omic” or biomarker imprints. This is already happening for cancer. One company is trying to replace a colonoscopy with a blood test that is derived from data that is capable of measuring 30,000 biomarkers from one blood sample. Of course the research stage may require 30,000 biomarkers from each blood sample but patient disease analysis may only require small subset of these and patient disease monitoring may require an even smaller subset. And yet another company claims they will be able to detect both colon cancer and polyps (precursors to cancer) via a blood test with a different set of biomarkers (focused on small molecules). Each are going through their own trials to verify results, but many of these efforts across many biomarkers and diseases (including beyond cancer eventually) will yield results that in the future will enable us to detect all cancers (and their progression) via a blood test. Cancer progress may be measured by biopsy of tumor DNA or cell free DNA in your bloodstream (cell free blood DNA shows what foreign DNA your body may be naturally attacking), and possibly augmented by the tracking of some subset of the 30,000 biomarkers which may show disease progress or treatment effectiveness. RNA or transcriptome and metabolomics markers are selectively being used in oncology research and advanced treatment though they are still very elementary (what we’d call v0) systems with a very small knowledge and data set. This is the beginning of precision medicine for oncology. Imagine what v7 systems might look like! As computational discovery of drugs helps find drugs likely to work and target the offending biological receptors, monthly monitoring of cancer mutations is becoming possible as mutations happen in real-time as opposed to when the symptoms become visible. Major diseases like diabetes and heart disease will be similarly characterized and this may result in a fragmentation of today’s “diabetes” disease analysis into many different subtypes of diabetes or potentially even completely different

classifications of disease (and the comorbidities associated with it) that all have one common symptom like “poor blood sugar control”. This stage of value derived from biomarker and genomic studies adding value in treating patients will start to accumulate large amounts of this new type of data which we have had little reason to collect prior except in small research studies. We will have economic justification to collect this data that has been heretofore prohibitively expensive. Once we collect such periodic data on millions of patients we will truly start to have the basis to understand the true effect of genomics, epigenetics, blood biomarkers and what each of the thousands of metabolic pathways in our body are doing and how they interact in very complex ways well beyond the ability of even the smartest humans to comprehend. In fact, the Center of Assessment Technology and Continuous Health (CATCH at MIT / Massachusetts General Hospital) was established to discover and apply new ways of quantitatively measuring the human condition in health and disease (with a focus on phenotypes).

The reinvention of medicine into something more precise, consistent, detailed, and comprehensive than that of today’s medical literature will start at that point. Instead of blunt symptom analysis (blood sugar, pain in the chest, dryness in the mouth, tiredness, lethargy, occasional dizziness, or other macro phenomenon) we will start to have very precise indicators of disease states and degree with objective measurements of biomarkers. We will hopefully understand (or at least quantify and statistically infer) the complex network of metabolic pathways the disease is affecting, and what the patient’s body is exactly doing or not doing at the level of physiology, microbiomics, biochemistry and more. Maybe what we today know as symptoms will function as a guide for the tests we give in the future as opposed to defining our disease. Since this is altogether too much data for humans, we will use data science to tell us how to measure the effectiveness of a particular drug for a particular patient’s specific conditions/genomics, environment/ethnicity, as opposed to the status quo of having drug companies develop one drug for all seven billion people on the planet. Massive amounts of data will be reduced to actionable insights to be used by medical and non-medical

professionals initially and eventually by software systems that do more and more of their functions. Though traditional medical literature we use today will often still be correct, possibly the majority of it will be obsoleted by better ways (from data) to diagnose, prescribe, monitor, and manage a patient's condition. Kappa's of 0.2 as in the diagnosis manual for psychiatrists (DSM5) will hopefully become history. The amount of data will be too much for any human to comprehend and yet precisely what we need for trained systems (trained by the world's best research physicians) to provide highly personalized medicine. Though pricing of care for various reasons is very hard to predict, the cost of providing this level of care will be substantially lower than in today's systems (this will be discussed briefly near the end of the paper). We will use cheap compute power and massive data analysis, much like today's advertising systems, which are substantially cheaper and far more precise in targeting advertisement than traditional humans in the advertisement business. It is sad that more compute power is used today to target a human's advertising worth usually a few cents on the Internet than a \$10,000 medical treatment. Some of these efforts are already taking place at a national scale by the NIH— from the Precision Medicine Initiative to the Cancer Moonshot to the BRAIN Initiative. They all center around novel application of large amounts of data to understanding our biology and disease.

When we get to massive amounts of data on hundreds and millions of patients we will have set the stage for a complete re-invention of medicine. Let's take a diversion for a moment. A recent article⁷³ talked about Google's new systems: "Google no longer understands how its 'deep learning' decision-making computer systems have made themselves so good at recognizing things in photos ... Google researchers can no longer explain exactly how the system has learned to spot certain objects, because the programming appears to think independently from its creators, and its complex cognitive processes are inscrutable. This 'thinking' is within an extremely narrow remit, but it is demonstrably effective

and independently verifiable." Such technology development is hard to comprehend or believe but very real. With very large amounts of data this kind of "insight" from the data set in the hands of the very best teams of mathematicians and domain expert scientists will result in completely new ways to practice medicine. Just one example, as we shall see in detail later, of this is a recent tumor pathology study where while pathologists did a good job of reading cancer tissue pathology the system learned to read the same things⁷⁴ But surprisingly, and completely unprompted and without any knowledge of the biology of cancer it discovered new features to look for that human researchers and thousands of pathologists had never thought of. I suspect that in fifteen or twenty years there will be thousands of such new insights comprising what will be the future "practice of medicine" and it will be much more scientific than today's "practice of medicine". I constantly see arguments and naïve "counter proofs" of what computers cannot understand from naïve practitioners who don't understand what is possible today and likely to be possible in two decades from the science of data. Today's doctors are least qualified to opine on where this technology development will lead and what might be possible. Only a small fraction of even data scientists at the leading edge of AI research in system with esoteric names like "deep learning" and "random forest" can comprehend the true potential impact of data science.

Even in this elite set (including those attempting to replicate the full capabilities of the human brain⁷⁵), there is much debate and conjecture about when, how, and in what domains can data science add the most value. These systems, especially currently, do have biases since humans still initially set up the models (with certain known and unknown assumptions that impact the quality of results). And as with all fields, there will always be some things that are more hyped than others (deep learning experts will believe they can change the world of medicine just as cellular biologists believe that). But the results we have seen in the last couple of years have been unprecedented –

73 http://www.theregister.co.uk/2013/11/15/google_thinking_machines/

74 <http://med.stanford.edu/news/all-news/2011/11/stanford-team-trains-computer-to-evaluate-breast-cancer.html>

75 <http://www.bbc.com/news/technology-24710209>

from AlphaGo's efforts to beat the world leader in Go (a much mathematically harder task than chess) to recent efforts bridging textual content with images (such that we can type "show me pictures of flying birds" and automatically those relevant pictures will show up).

My personal opinion is that the recent results (e.g. machines solving the very difficult human task of reCAPTCHA which was considered impossible for machines to do ⁷⁶) prove that progress will be rapid. I suspect technology development in the medical domain will be rapid and relatively easier compared to much more difficult domains like reCAPTCHA and driverless cars or Go playing that Google is attempting. In my estimation building driverless cars is 10x more difficult than replicating those effects with what a doctor does. There are other "too complex for humans" fields already are highly automated like stock trading and autopilots in aircrafts, something that previously was all done by humans. AlphaGo, for example, has been described as having human level of "intuition" when playing the game, given computationally the system could not compute all the possible gameplay options. In fact, this intuition to a degree is superhuman as the most unconventional (and most consequential) moves of AlphaGo were ones human observers didn't foresee ⁷⁷. But humans are now learning from those moves (gaining their own intuition and understanding of the moves) and using it to great effect; Lee Sodol, who lost to AlphaGo, is now winning almost all of his games after the event ⁷⁸.

However in the medical domain the regulatory, profit interests and human elements will slow adoption to a substantially larger degree. Hence timing of such transitions becomes very difficult to prognosticate on and we stay in the domain of speculations and not predictions when it comes to timing.

The evolutionary path for the "re-invention of medicine":

Here's one possible speculative thought exercise

on how the v0 to v7 in healthcare might happen. We use this analogy only imprecisely because the digital world is often measured in versions 1,2,3.. denominated as v0, v1, v2,... . Compared to the very long cycles of significant development (like a new pharmaceutical taking 10-15 years to develop) versions in the digital domain typically take 2-3 years (often much faster, even days and weeks, which result in accelerated rates of experimentation, learning and evolution/progress). I expect digital technologies on medical innovation will dramatically accelerate the rate of change in the medical domain, somewhat slowed by the traditions, regulations and practices in medicine (biology based innovation will continue on its 10-15 year cycle I presume). We suspect the typical version cycle will be 2-4 years for most innovations though it may vary dramatically depending on the domain of medicine and the geography the innovation is applicable to. We use v0-v7 as an approximation of a decade or two or three of innovation as we analogize to the phone market where v0 to v7 went from a floor mounted sewing machine like mobile phone to today's iPhone. This is meant to be a simplistic way for the reader to imagine how large a change can be a result of seemingly small steps, and is not meant to have a causal basis underlying it.

Keep in mind that mobile phones were around for a long time, gradually evolving until one day they seemed more like the "device that used to be a phone" instead of a phone itself. Its primary use evolved from talking to a myriad of other uses. The new medicine I talk about will likely become the "profession that used to be medicine" with complete re-definition of the role of the practitioners, the means of diagnosis, treatment and monitoring. But what exactly and when is nearly impossible to predict. In fact, voice calls now represent a small percentage of the time we spend on these devices today. We have a world that was hard to imagine 15 years ago when even the Motorola Startac (shown above as v5), did not exist.

⁷⁶ <http://www.google.com/recaptcha> <http://en.wikipedia.org/wiki/ReCAPTCHA> <http://singularityhub.com/2013/10/29/tiny-ai-startup-vicarious-says-its-solved-captcha/>

⁷⁷ <http://www.wired.com/2016/05/google-alpha-go-ai/>

⁷⁸ <http://gokifu.com/player/1/Lee%20Sedol>



v0



v1



v2



v3



v4



v5



v6



v7



v8+

So here are our speculations:

The last decade or so in healthcare has seen the rise of several important digital innovations that have brought meaningful efficiencies and, in some cases, improved the quality of care for patients. E- prescriptions, EHRs, new patient management systems and other improvements over the last decade have made the provision and receipt of healthcare more convenient and even yielded some clinical benefit. Patients can now get digital downloads of x-rays to help with care coordination, secure online messaging systems are in use at many hospitals, student clinics,



v0 Medicine

and other care centers to speed communication between doctor and patient, and patients can now book visits to the doctor online instead of via phone after waiting interminably on hold. These services enable and empower consumer choice. As consumers grow to expect that they can instantly book appointments, see reviews ... over time, increasingly, they will grow to expect to have all the information that a doctor has during time of care and to be able to ask questions.

Health information exchanges in the US will further facilitate data use and provide a fuller picture of a patient's health. Health communities are being formed where doctors can electronically answer patients questions or patients can access possibly the largest repository of doctor or patient generated content, creating a valuable medical community. But most important is the digital capture of data. This electronic capture and the necessary exchange of data will be key enabler of future v1-v7 data- science-based products. And the somewhat frustrating and poor implementation of some of these systems like EHR's will improve and get better from their finance and accounting oriented initial implementations to helping with the daily practitioner's work. Though none of this information technology is revolutionary, it enables the seeds of a data science based medical system. Though these mostly administrative systems are more for convenience (or inconvenience for the poorly designed ones) and don't by themselves offer much improvement in health care science or delivery!

These sorts of improvements represent low-hanging fruit in healthcare, with one of the key themes being a (currently rudimentary, but powerful) shift to online and mobile for many healthcare services. Many startups and some larger companies are tackling the remaining obvious challenges in the system as part of their v0 efforts. Here are some illustrative examples.

Some early v0 systems include Genomic Health which uses data discoveries from genomic analysis in a breast cancer test to determine the genetic underpinnings of a patient's cancer. This data helps predict the likelihood of recurrence

in early-stage breast cancer and also the effectiveness of aggressive treatments like chemotherapy, which is prescribed far more often than it needs to be. With the innovations from these companies, doctors are able to individualize treatment options. Other companies and researchers are able to extract signals from blood draws that were previously only possible to gather using biopsies. Doing these "liquid biopsies" for cancers whose cells mutate over the course of the disease give the ability for doctors to personalize care, and was only possible as a function of better digital and genomic sequencing technologies. Another powerful example of a v0 system is in reproductive health – the maternal-blood-based noninvasive prenatal test (NIPT) for fetal chromosomal defects. This test replaces previous invasive procedures like amniocentesis which carry a 1%-2% risk of miscarriage and cost about \$1,500. NIPT started as a \$2,500 test for high-risk pregnancies in 2011 and has quickly become a \$400 test for average/low-risk pregnancies, driven by technology improvements in clinical specificity and test cost. Interestingly, the data analysis from nearly a million women tested by NIPT led to an unexpected discovery – the identification of rare cases of pregnant women with cancer, flagged by the multiple chromosomal defects that originated from their own tumor. Powerful and ethically challenging discovery and yet, an example of how a v0 system performed well beyond the expectation.

Research also gives us a good avenue to see how current v0 systems harness data and algorithms to better medicine. One set of research is from Roger Guimerà and Marta Sales-Pardo, who have developed an algorithm that determine drug-to-drug interactions based on data from previously reported interactions. This is a growing area of concern as 25-percent of adults in the U.S. over the age of 57 take more than five prescriptions simultaneously. Their algorithm does not require any additional pharmacological or biochemical information, it only requires sets of previously reported interactions. The method has proved effective "both in exhaustive pairwise interaction data between small sets of drugs, and in large-scale databases... and [it] can be used efficiently to

discover interactions of new drugs as part of the drug discovery process.⁷⁹ This can enhance and succeed current drug interaction databases like Epocrates.

Also in research, there has been plenty of innovation globally, which has proven the importance of collecting and disseminating data. For example, the Danish National Biobank was established to bring a biobank registry for all residents in Denmark to enable Danish biomedical researchers to harness big data techniques to understand and study the population at scale. Another initiative with similar goals is the Electronic Medical Records and Genomics (eMERGE network established in 2007 in the US (now in Phase II) to bring together genotypic data (DNA repositories) and phenotypic data (via EMR's).

Successes to date of eMERGE have been to use disparate data, standard genomic analysis, and algorithmic techniques to better understand the connection between the genome and phenome. The future goals of eMERGE will be to utilize that data to improve genomic research, and eventually, health outcomes across a diverse patient population.

Using basic data analyses to improve quality of care and decrease mortality rates (without using sophisticated techniques) are also manifestations of v0 uses of data. For example, in Sweden, a comprehensive disease registry for heart attack patients contributed to a 65% reduction in 30-day mortality and a 49% decrease in 1-year mortality from heart attacks.⁸⁰ EMR's, and clinical decision support tools embedded within them, hold promise for improving the application of evidence to medicine.⁸¹ For example, one study showcased improving health outcomes for diabetics by 15% via proper utilization of these EMR's.⁸² These analytics and predictions will only get better over

time though propagation of these new medical insights to daily practice is a significant problem as humans are slow to change their behavior.

The use of data in a clinical research context is not the only v0 shift in healthcare. The shift of healthcare services to online and mobile can occur in a variety of ways. A typical v0 implementation might enable patients to follow medical instructions more easily by digitizing hospital handouts while handling aspects of patient intake, discharge, and procedure preparation on mobile. Whether it's colonoscopy preparation instructions or post-cancer surgery stoma care, or patient intake for your doctor, a system like this could do a lot for the "last mile" in patient care and capture data on real patient behavior in the process. The system can adjust to the level of the patient's capability to understand medical choices or instructions in the patient's native language, not "doctor-speak" though getting it right is a fairly sophisticated technology effort. These types of systems are innovative only in making administration easier, data capture and instruction more timely and accurate and information more complete. They are nurse assistants that make the front office job easier and more leveraged adding substantial economic value. They do save time, reduce cost and improve information completeness and patient experience.

Many startups are taking existing medical practice and making it better, cheaper (though technology has historically increased health care costs because of a variety of factors), more accessible, or some combination of the above. Mobile services can start to provide some of the point functions we'd typically only associate with devices provided at a hospital. One company has a single-lead ECG machine that fits into a smartphone case, and the device has been approved by the US FDA for over-the-counter sale without a prescription.

79 Guimerà, R., et al. "A Network Inference Method for Large-Scale Unsupervised Identification of Novel Drug-Drug Interactions." *PLOS Computational Biology* journal. 5 Dec 2013. DOI: 10.1371/journal.pcbi.1003374.

80 Larsson, Stefan, et al. "Use of 13 disease registries in 5 countries demonstrates the potential to use outcome data to improve health care's value." *Health Affairs* 31.1 (2012): 220-227 (<http://content.healthaffairs.org/content/31/1/220>)

81 Smith, Mark, et al., eds. *Best care at lower cost: the path to continuously learning health care in America*. National Academies Press, 2013. (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America.aspx>)

82 Cebul, Randall D., et al. "Electronic health records and quality of diabetes care." *New England Journal of Medicine* 365.9 (2011): 825-833. (<http://www.nejm.org/doi/full/10.1056/NEJMsa1102519>)



By just placing their fingers on two electrodes embedded into the case, a patient can get a medical- grade, single-lead ECG (with more sophisticated multi-lead devices in process) whenever and wherever they want or need. Soon they will get other patient parameters like respiration rate, blood oxygen, heart rate, blood pressure etc.. The company offers inexpensive ECG reads by a cardiac technician or slightly more expensive reads by cardiologists, but in any case, these services are provided at a fraction of the cost of getting an ECG done in a US hospital. Though it's not as comprehensive as a 12-lead ECG in a hospital, their ECGs are a lot more contextual and immediate. Since you can take an ECG when you are exercising, when you are not feeling well, after you wake up, after you eat, when you are stressed at work, etc. (which your hospital ECG cannot do), you have many more ECG traces and are not limited to a pre-scheduled time when you have your hospital appointment when you may or may not be displaying symptoms. This device and services associated with it lowers costs, makes screening potentially more pervasive (every Walgreens and every village in India could have one and offer free or near-free screening), and lets real cardiac patients constantly monitor themselves five times a day the way diabetics monitor their blood sugar five times a day. Patient engagement has been astronomical so far as well, so you will be able to get this type of monitoring and feedback. Of course, it creates the problem of who will read so many ECGs and that will lead to the need for a v1 solution.

Another company similarly takes functions normally handled at a physician's office, and starts moving them to mobile devices. Another iPhone case can effectively convert your phone to a digital otoscope. In its v0 implementations,

it can function as an ENT assistance device, where mobile phone images can be used for telemedicine, helping reduce the number of times a pediatrician / parents have to send their kid directly to an ENT specialist. Instead, by greatly reducing the friction of delivering an image (or set of images) to the ENT, a parent or pediatrician could use this device to take an image, and have it sent directly to the ENT who can then process it much more quickly (becomes a part of his/her flow). That enough is worthwhile, but this technology can be potentially greatly expanded upon in v1+ to transform ENT care. Once this data is collected in these initial efforts, diagnosis can be done using algorithms alone directly on the images (without requiring a human in the loop). The technology is easily extended to dermatology and over time potentially into ophthalmology. We are already seeing this happen with Google DeepMind (and other researchers') ability to automatically detect diabetic retinopathy from images, though implementation in practices has not yet happened.⁸³ But questions on the economic viability of these early attempts remains.

For some specialties, extensions don't even need to be added to our mobile devices to add value to physicians. Another v0 effort is similar in style to the ECG device in terms of data completeness in real- world contexts and applying it to mental health care. The company offers an app that keeps track of all a patient's activity for mental health patients (under a psychiatrist's supervision) on their mobile phone to give them and their care providers insight about mental health. While a psychiatrist can meet and assess patient once or twice every week or month (if the patient can afford it), this app installed on the mobile phone can constantly, and without any action required by the patient, monitor every activity. Did the patient get out of bed? Did they go to the kitchen or visit their regular restaurant? Are they skipping meals? What about their social activity – Are they texting their friends, emailing their co-workers, or calling their parents the way they usually do? The application gets to understand truly what is normal for patient and, by tracking a 100's of data points a day, detect any deviation from normal, whatever

⁸³ <https://deepmind.com/health/research>

“normal” means for an individual patient. Today, the application reports the status of every patient (red/yellow/green) to their psychiatrist, spots deviations that might be meaningful, and suggests intervention for at-risk patients by alerting their doctor to “check in with the patient”. The mobile app keeps a psychiatrist in closer touch with their patient and lets a nurse more intelligently decide which patients to check in with when he or she starts cold calling her at the beginning of the day. There is too much data already for a human psychiatrist to process at reasonable cost so the company is doing data reduction to insights. That leads to an exciting set of goals for v1 of this system, as we shall see below.

Physiological sensor-rich v0 prototypes we have seen in the wearables space from many startup companies attempt to measure not only daily activities like step counts and heart rate, but also substantially more physiologically nuanced data like continuous blood pressure, hydration, blood oxygenation, cardiac output, blood glucose, and others. While success and adoption remain to be seen, (and early versions are far from accurate, though they will improve and become more comprehensive over time), these will initially be interesting point innovations, some integrated into single devices and others as standalone hardware with both medical and non-medical applications covering fitness training, stress reduction, sleep management, weight reduction, hypertension management, cardiac disease management, diabetes management and others. Beyond that, companies are also using facial recognition and audio signals to detect emotion and mood, which we know have an impact on health. How much more of the human element of care can one provide if one has a direct read of a patient’s nuanced facial messages. These, though possible but not yet being implemented in practice, could be used by the most forward looking doctors as data to add to their diagnosis. These technologies can substantially improve the capabilities of the primary care physicians and provide them with more reliable data about what the patient does outside the doctor’s office. But they fit directly into today’s medical practice with additional and

more accurate information. Most physicians will choose to use this information only marginally until more rigorous multi-year studies based on tens of thousands of patients become available over the course of time. But we need usage for low risk additional input and insight to collect the data so more accurate insights can be derived from it. We need the pioneers who will tolerate the early arrows for this future to happen.

Few companies are trying to develop sensors for breath, food, and other analytes. Every chemical reaction in the human body with a volatile product likely ends up in the human breath. Attempts are being made to find ketones as indicators of ketosis for diabetics or fitness fiends or nitrous oxide as indicators of COPD attacks. Lung cancer and breast cancer detection has been demonstrated through such inexpensive breath sensors and, one day, mass screening could become possible, though first efforts often fail and just become guideposts for how to do it better. Using volatile ketone measurements in place of or in addition to blood sugar to monitor diabetic status may be possible while non-invasive blood sugar efforts continue apace (a very hard problem to crack given the complexity).^{84, 85} Such sensors may develop into diabetes or other disease management tools by helping assess the food one is in-taking and telling us the composition of fats, carbs, and proteins, producing instantaneous food labels (though humans may choose to ignore them, a problem that’s possibly even more challenging to solve). Inexpensive and automatic spectrometers may allow us to assess the constituents like fat in our food or food adulteration with cheaper products. Alongside, daily tests for a variety of biomarkers may make it possible to monitor certain diseases like hypertension, helping with management and providing continual feedback. Over time these data points may help a patient answer questions like “why is my blood pressure good this week but was not great last week?” quantitatively and causally.

An some v0 systems will allow for a patient to ask exactly those questions. IBM’s Watson is an artificially intelligent computer system capable of

84 Am J Physiol Endocrinol Metab 300: E1166–E1175, 2011 University of California Irvine, Institute for Clinical and Translational Science

85 Ueta, Ikuo et al. Journal of Chromatography B, Volume 877, Issue 24, 2009, 2551-2556

answering questions posed in natural language. In February 2013, IBM announced that Watson software system's first commercial application would be for utilization management decisions in lung cancer treatment at Memorial Sloan-Kettering Cancer Center in conjunction with health insurance company WellPoint. IBM Watson's website ⁸⁶ as of June, 201 states "only 20 percent of the knowledge physicians use to diagnose and treat patients today is evidence based. Which means that one in five diagnoses is incorrect or incomplete. Consider that the amount of medical information available is doubling every five years and that much of this data is unstructured. Physicians simply don't have time to read every journal that can help them keep u to date with the latest advances. Given the growing complexity of medical decision making, how can healthcare providers address these problems? Physicians can use Watson to assist in diagnosing and treating patients by having it analyze large amounts of unstructured text and develop hypotheses based o that analysis." Watson has ingested hundreds of thousands of diagnostic reports, millions of pages of medical journal articles alone many hundreds of thousands of patient records for their oncology-focused work. One startup doing a consumer social health management platform is actually taking advantage of Watson's capabilities and has integrated it into their Concierge app. The app will allow customers to ask Watson for personal health recommendations. But Watson is a relatively early and simplistic system in its AI sophistication and much more capable systems are being developed. That bodes well for using knowledge and data promisingly in a decade!

Another v0 effort wants to harness large amounts of previously underutilized data to build a "graph of medicine", connecting symptoms and test results to disease and building an understanding of overlaps and interactions amongst them. Much like IBM's Watson system, their system has ingested millions of articles from medical literature, has curated the graph (essentially an embodiment of knowledge in the medical literature it has ingested) in each specialty, and can detect using existing patient records any potential missed diagnosis. Detecting such missed diagnoses is

obviously better for patients, but it's also good for the care provider and the payer, usually an insurance carrier in the US or a government in many other countries. Such a graph may ask the patient a few extra questions in an automated patient intake app to eliminate a missed diagnosis or offer the doctor something from the latest research paper that may be relevant or actionable in this patient's condition. Drug interaction systems are the v0 version of this and help avoid many an adverse reaction.

Efforts at consistently and continuously monitoring real patient medical records and mining them for insights related to medical treatment efficacy will yield results surprisingly early. Some startups are attempting to do this. The insights will help optimize patient care protocols (guidelines for how patients with certain conditions should be treated) in hospital systems making these care protocols more truly evidence based instead of expert opinion based. And with all the patient data, the company can re-run many medical studies based o real-world patients, their experiences with each treatment, real-world compliance, co-morbidities and more. More so than medical literature and doctor expertise, this kind of analysis may reveal all kinds of errors in current medical practice turning it into more of a science. It will complement the kind of analysis based o medical literature, developing the best science currently possible by reverifying (or not) a treatment's performance in the real world of real patients hence validating or invalidating the original medical studies.

Utilizing data for v0 applications can also harness techniques developed more broadly in data science and mathematics. There is a huge opportunity to dramatically improve patient outcomes by recognizing the value of end-to-end clinical pathways, identifying critical steps, and removing non-productive human biases and variability. One company analyzes a health provider's electronic health records to determine optimal care pathways for common procedures. They apply mathematical techniques to actual clinical data to determine optimal care pathways that produce the best outcomes for prescribed

⁸⁶ http://www.ibm.com/smarterplanet/us/en/ibmwatson/work.html?cmp=usbrb&cm=s&csr=watson.site_20140319&cr=index&ct=usbrb301&cn=sec2menu

patient populations, personalized at the level of the patient and their individual circumstances. As a result, the hospital can spread the adoption of best practices that reduce costs and improve care. In one example, their software was applied to hospital data from gallbladder removals, producing an optimal care pathway that had an average 15% reduction in patient stays and reducing re-admittance rates due to complications, while reducing costs at the hospital.

As another surprising example of v0 innovation, the FDA has approved (as a drug) the first mobile application that has the same effects as a pharmaceutical! This company produces a mobile app that helps lower blood sugar much like the drug Metformin does. It follows protocols for behavior based blood glucose reduction already proven in today's medical literature and captures them in an application that has passed the FDA's standards for a pharmaceutical. This kind of digitization of known medical effects and behavior modification effects to treat disease will be commonplace, often as literal or only somewhat modified versions of known protocols from existing physical world studies. Another company is similarly following a known protocol to reduce weight and showing its effectiveness in reducing diabetic and other associated risks of excess weight by making it mobile based and interactive. The most promising evolution of this in v1 and beyond is the optimization of these protocols based on millions of patients interacting with these systems daily on their mobile devices. The volume of data and hence opportunities to optimize will expand rapidly compared to the very small studies offline methods are usually restricted to for economic reasons. Similarly, meditation has been shown to affect gene expression and it won't be long before someone uses the Harvard study to prove a therapeutic effect.⁸⁷ Multiple small companies are developing blood tests to screen for colon and other cancers, potentially supplementing or even ultimately doing away with colonoscopies. Though the typical doctor only simplistically tracks a few biomarkers

today in your standard blood test, one startup got to its colonoscopy test by evaluating hundreds of thousands of biomarkers in each blood sample. The same techniques will start to be applied to diagnose other diseases from heart condition to diabetes and mental health in their v0 versions. In future versions, sophisticated predictive, monitoring and prognosis prediction capabilities will be incorporated, harnessed by the data flowing in. Monitoring of these biomarkers or patterns of biomarkers and changes in those will be predictive of developing disease.

Professor Larry Smarr utilized a computing techniques and was able to monitor changes in his own blood biomarker patterns over a number of years to detect IBD two years before the doctors could diagnose it, by detecting episodic inflammation peaks via the Complex Reactive Protein (CRP), coupled with data from his gut microbiome and his genome⁸⁸. I suspect these biomarkers or other similar measures like gene expression or the microbiome map will be the primary contributors to a medical diagnosis at some point in the not too distant future and symptoms will only be guideposts for what tests to run. Imagine what will be possible in 15 or 20 years when this technology will be in its later version evolutions!

A example pioneer in the amount of molecular data possible to be collected is Stanford researcher Michael Snyder. The massive collection of genomic, transcriptomic, proteomic, metabolomic and autoantibody data that his lab collected over 14 months and then published in Cell represents an extreme example in the current technological landscape.⁸⁹ The group continues to add to the "Snyderome" with new data streams, and as they and other quantified self pioneers like Larry Smarr's group work both work out the technical details to improve efficiency and make improvements in technology so it gets cheaper and more accessible. These kinds of efforts may lead to the ideal annual physical.

Improvements in both physical technologies

87 Bhasin MK, Dusek JA, et al. "Relaxation Response Induces Temporal Transcriptome Changes in Energy Metabolism, Insulin Secretion and Inflammatory Pathways" PLoS ONE (2013) (<http://www.plosone.org/article/info%3Adoi%2F10.1371%2Fjournal.pone.0062817>)

88 Smarr, Larry. <http://www.slideshare.net/Calit2LS/ucsf-dec-2013final>

89 Chen R, et al. <http://snyderome.stanford.edu/>

and analysis techniques for methods such as mass-spec, sequencing, NMR-spectroscopy, and micro-bead/assays are making highly multiplex measurements more cheap and efficient than single one-off tests using existing methods of clinical chemistry. At the same time, advancements in technologies such as microfluidics (including paper based microfluidics), nanopores, silicon photonics, nanowires, and other pieces of biotechnology are making the sensing devices increasingly cheaper and more portable, even to the point of hand-held disposable peripherals, able to leverage ubiquitous mobile technology (such as the camera, computational power, and connectivity of a smartphone). In v0 level technology we start to move away from molecular diagnostics and lab work being collected only in a clinical setting, to testing in kiosks or mobile care centers, and finally in the home. In a v1/v2 or maybe v3 system there will be home based testing through fingersticks, exhalation, or sampling of other material such a sputum, urine, or feces. The may then be nearly continuous testing through an external system tracking molecular makeup either optically, through microneedle sampling, or outgassing from the tissues from an externally applied device (like a “smart sticker”) or implanted device. In parallel with home systems used to track glucose through fingersticks to implanted continuous glucose monitors and efforts to measure glucose level non-invasively, similar measurement technology will emerge for a vast array of biomarkers. These new data streams will lead to continuously improving understanding of physiology and pathophysiology, both in general and personalization for the individual.

Data-driven health care insurance provider try to isolate the best care providers and exclude the more wasteful providers from the network. One company helps take the best practices and innovations in digital health and brings it to the forefront of health insurance, traditionally an industry known as slow moving. Their success hopefully will help validate the value of data-driven and consumer-driven healthcare, even in the v0 stages, and can accelerate overall progress.

Though we are most familiar with innovation within the Silicon Valley ecosystem, we suspect there is more innovation currently going on

elsewhere too, as hundreds or thousands of entrepreneurs and scientists toil away in secretive, proprietary invention. Each will produce their v0, v1 and beyond instantiation at a different time and in different levels of sophistication going after very different, often specialized applications. A thousand such points of light will bloom and though some will rapidly extinguish, others will permanently start a fire, which may spread and change medicine. Whether the fire is small or large or massive will depend on many factors including in my experience a large dose of luck. I even suspect that the size of the fires will follow the well known Power Law, for those of you who are more mathematically inclined. While the skeptics will focus on the failures with their “told you so” attitude or the triviality of the small fires, the true innovators will recognize that in real life “improbable is not unimportant. We just don’t know which improbable is important”. We just don’t know which fires will spread to be massive transformations and which ones will be smaller fires. Especially in the early stages, that is why we need thousands of fires started in this new domain.

As useful as these sorts of improvements are, the vast majority of the new v0 approaches to healthcare today are primarily enhancements, largely incremental improvements to the existing way of doing things. These v0 systems and even v1 systems will only be marginally useful to existing providers but will start the data engine so useful to AI systems. The v0 mobile was floor mounted in cars and barely even mobile! Moreover, these approaches have had to fit within the current framework of medicine and the limits that come with it. More often than not, they have been geared towards cost control, administration and regulation (like electronic health records), and have been clumsy for patients, doctors, and nurses. Tweaks to existing ways of doing things, though needed and potentially significantly better compared to the current state, will only get us so far. These systems will primarily be better information technology (IT), not real medical invention or innovation. In order to fundamentally transform healthcare for the better, we have to fundamentally transform the rules and ask 1st-order questions. The visible changes are where v1 systems will start to be material though v0 vs. v1 boundaries are admittedly artificial.



v1 Medicine

The v1 stage of healthcare innovation is when we'll start seeing boundaries get pushed in more significant ways. Electronic health records may, for example, actually help patient outcomes more proactively in the next generation. In general v1 and v2 systems will provide "bionic assist" to health care professionals, be they doctors, nurses or others. This bionic assist can resemble systems built on top of EHR's, for example ones that make EHR data digestible and more integrated into a clinician's day- to-day activity (rather than just as a backend system). Just as early versions of driverless car systems may help drivers "stay in the lane" or "stay a safe distance from the car ahead" v1 and v2 systems will use data in intelligent ways beyond what is done today. While v0 systems may be software to help with healthcare workflows, v1 systems will do "intelligent routing," alerting users to errors, cautions, better treatments and more over time.

Patient intake systems of v1 may, based on known literature, actually know additional questions to ask a patient during patient intake to reduce probability of a missed symptom the patient thought was irrelevant to their problem and hence dramatically reduce missed diagnosis rates traditionally stemming from errors like premature closure bias or recency bias of a nurse or doctor. The key here is that the systems will understand the whole body of medical knowledge or will be partnered with systems which understand this knowledge and can make intelligent decisions beyond simple "programmed rules" which are often the characteristic of v0 technologies. The systems will also reduce patient intake time while

improving accuracy and completeness of the data. As medical literature is updated or other practice based evidence emerges, the questions the patient intake system asks will automatically and dynamically change with increasing sophistication as we move from v1 to v2 and beyond. Of course it will also change its questions based on personalization, incorporating factors such as the patients age, gender, ethnicity, weight, emotional state, job stress, past medical history, other unrelated medical conditions, and environmental factors like local flu epidemics, local viruses, local cancers, local infection trends, local disease demographics, and much more. This is real bionic assistance for primary care and other doctors to improve their quality of care and diagnosis while saving time and money. Premature closure or failure to consider all the possibilities, especially the rarer and more serious ones, is one major cause of diagnostic error that can be eliminated.⁹⁰ Other causes such as systematic errors or failure to include symptoms or laboratory tests already in the patient record would also be eliminated by such a combined system, and the patient information-gathering process and the patient's understanding of medical choices could be made much better.⁹¹ These v1 systems will be global as well, and would accomodate patients with poor understanding or comprehension, for example those that may have English as a second language in the US.

Already, thanks to data findings at Kaiser Permanente, using statins within specific protocols for in- hospital stroke patients reduced death rates by 40-percent. This data science based protocol change may be among the most effective treatments for reducing stroke mortality. In yet another example involving statins, data-mining by researchers at Stanford showed that there was an adverse drug reaction (ADR) for those taking a particular type of immunosuppressant commonly prescribed for transplant patients. There are two equally effective immunosuppressant options that doctors prescribe at random: Cyclosporin A and Tacrolimus. Based on data from actual patient

90 Graber, Mark L et al "Diagnostic Error in Internal Medicine" Arch Intern Med (2005) (<http://archinte.jamanetwork.com/article.aspx?articleid=486642>)

91 http://www.hopkinsmedicine.org/news/media/releases/diagnostic_errors_more_common_costly_and_harmful_than_treatment_mistakes and would accomodate patients with poor understanding or comprehension, for example those that may have English as a second language in the US.

health records (EHR data), they were able to show that Tacrolimus is three times safer, meaning patient taking Cyclosporin A in combination with a statin has three times the odds of rhabdomyolysis or serious muscle degeneration. It turns out this drug and all statins are all metabolized (broken down for excretion) by the same enzyme in the liver (CYP34A) so taking both at the same time likely overloads the ability of this enzyme to clear the drugs from the system, and is thus the same as overdosing on statins.⁹² Although practicing physicians may understand the general principles, it is increasingly hard for any person to keep up with all these continually discovered interactions.. Drug-to-drug interactions account for 30-percent of all adverse drug reactions (ADRs), and ADR is the fourth leading cause of death in the US. More and more of these will be surfaced by data analytics and more importantly make it available as insights at the point of care without the doctor having to remember thousands of such factoids and to remember to use them.

Imagine the lives that could be saved with more data-based findings. Over time, the 15,000 or so diseases and the 15,000 or so therapies will be rationalized, retested, personalized and systematized. They will be applied in a more coherent way based on a lot more data on each situation than is available today and more than a human could manage. In future versions of medicine even more data (thousands of biomarkers) will be used to more precisely diagnose, treat, and monitor progress. The cost benefit of each treatment will hopefully be measured in the context of a patient, not for a population as a whole. The above are point discoveries, but hopefully in a decade or two such findings will become the norm and systems will make these “medical discoveries”. A simple blood test, in addition to or without physical, self-reported symptoms may become the standard way to diagnose many diseases, often well before their symptoms are even visible as illustrated by the example of Dr. Smarr’s IBD. The same will likely be true of heart disease where blood biomarkers and/or resting heart rate and/or heart rate variability patterns from your health band may foretell disease instead of waiting for

the first heart attack which is how most cardiac problems are discovered today. Symptoms, like a pain for a bad ligament in your knee, followed by a confirmatory MRI like imaging procedure (and even this will be much easier with new software systems) will still be valuable but much less so. Even a flu coming on may be detected days before you feel it through measurements from a physiology tracking health band.

In v1, data-derived systems and applications will evolve to augment current human care in more and more meaningful ways. The mental health example earlier of too much data for a human psychiatrist to consume is leading to new analysis which has led to, based on a small patient set, discovery of hundreds of micro-behavioral insights that are predictive of patient trajectory and are not in current literature.

Their prediction is near impossible without the data and v1 is about providing very fine-grained behavior monitoring to the psychiatrist beyond the red/yellow/green light of v0 and assisting them in more accurately and precisely diagnosing and treating patients. The new behavioral insights will help the company fine tune its assessment of any patient, detect in the future whether a patient is following their recommended medication or therapy regime, and do monitoring at a level that would be prohibitively expensive (and practically impossible) for a psychiatrist to do. It could detect missed medications, sending alerts related to drug compliance for diseases as serious as bipolar disorder and maniac depression as well as for milder conditions. In a future version it will evolve to fine tune dosage based on these micro-behaviors. Today dosage for a particular patient is usually a random though educated guess followed by a series of iterations. To a forward leaning psychiatrist the “bionic psychiatric assist” this system might provide will allow them to handle more patients at lower costs and provide them with increasing quality of care and more responsive care when the patient needs it. The fine grained and constant knowledge about not only how a patient is doing will be merged with a system’s predictive capability to indicate probabilistically what a patient might do in the

92 Iyer, Srinivasan V., et al. “Mining clinical text for signals of adverse drug-drug interactions.” *Journal of the American Medical Informatics Association* (24 Oct 2013). Retrieved 5 Feb 2014 (<http://jamia.bmj.com/content/early/2013/10/24/amiajnl-2013-001612.abstract>).

future. In contrast, its human counterparts can barely agree on whether a patient has a mental health illness! In v1, it may be able to confirm, change, or fine-tune a diagnosis (or make its own diagnosis in future versions) and, in the future, track the efficacy of a therapy or medication. By collecting vast amounts of per-patient data (today thousands of data points a day) on millions of patients & non-patients, and correlating them with various external data sources (e.g. weather or the stock market or a variety of social or physical factors or data such as genomic and epigenomic information that might affect mental health), companies will keep developing additional insights into human behavior both in wellness and in illness. We hope that well before v7, the span from happiness to mental illness will become a continuum with fine-grained, precise, and consistent measurement based on thousands of behavioral indicators each measured continuously. Hopefully the company will learn how to nudge each patient from their illness state (or their less than well state even if it is not classified as a disease state) towards happiness and wellness. This is the mental health analogue for the goals of Social Physics, which is Sandy Pentland's and MIT Media Lab's effort to harness the power of big data to understand how people and social organizations function via understanding human behavior in those contexts.⁹³ A equal number of new scientific research opportunities will emerge. If your social circle is obese does that increase the probability of you being obese? Beyond the mental, do physical factors like gene expression or microbiome change based on who you hang out with? Biology will start to offer much richer possibilities to better understand "why" as data science surfaces many "whats". The "why's" and "what's" will also blend together as rapid experimentation based on new insights learned from data science can be tested much faster to be biologically researched.

Imagine what v1 of a cellphone based home ECG could be in just a few years for the typical patient. With multiple ECGs/day or hundreds of ECGs/year, coupled with much more contextual information with relevant baselines against which trends can

be tracked, you likely can extract far more insight than a cardiologist working with his traditional 12-lead ECG could even approximate, even if the 12 lead ECG is "better". In fact such an ECG may still have a specialty role and probably will not go away. If hundreds of ECGs were sent to a cardiologist per year per patient, he wouldn't know what to do with them in the v0 version of medicine; but, using the millions of ECGs that startups might accumulate and label with the help of over-reads by cardio technicians and cardiologists, good machine learning algorithms will be able to detect almost everything in a new ECG that humans could detect and more (and likely do it much faster and substantially cheaper).

In fact, we already have indications of moving on to this (somewhat arbitrarily defined) v1 phase. The startup has recently developed an algorithm that identifies atrial fibrillation within an ECG taken from the mobile device and gotten FDA approval for at home auto diagnosed use.⁹⁴ Many more algorithms are in development for other conditions including attempts by some startups to predict heart attacks hours in advance. This is v1 medicine, well beyond what is possible today. This data flow will allow these companies to analyze whether the nature or frequency of the afib can be related to certain circumstances and use that to suggest or influence treatment. At scale, this can represent a fairly fundamental shift in how cardiologists understand, treat, and monitor atrial fibrillation (and eventually, cardiovascular disease in general).

If mental health apps in their v0 form currently have developed hundreds of new behavioral insights or symptoms that are predictive of future patient status, it is possible to imagine many new cardiac insights from the numbers of ECGs, patient records, and other behavioral data from smart wristbands and phones. Constant ECG monitoring will enable better care, and almost every cardiac patient will be able to get uniform quality of care and monitoring at affordable cost. We will have a better understanding of how the patient's heart responds to job stress while they are in the office or to exercise while they are on a

⁹³ Pentland, Sandy. "Social Physics: How Good Ideas Spread" Penguin (2014) (<http://socialphysics.media.mit.edu/>)

⁹⁴ "Silicon Valley in-depth: Tapping medtech's Big Data with AliveCor CEO Euan Thomson" <http://www.massdevice.com/features/silicon-valley-depth-tapping-medtechs-big-data-with-alivecor-ceo-euan-thomson?page=show>

hike or a bike ride or to a large meal or to the flu. How might it change with moderate exercise and what is its true fitness? It is likely that algorithms will in v1 get good enough to match most humans in reading the ECG for many if not all conditions and take this task off the human's plate allowing them to focus on other tasks while providing them a way to assimilate the ECG's outputs and trends into their decision making. We will generally classify such actions expanding what a health professional can do as "bionic assist" much as an auto-pilot removes complexity and expands what a airplane pilot can do, while allowing the human to take over in exceptional circumstances. Most people in the world today discover they have cardiac disease only when they have a heart attack – imagine how much impact we could make on cardiac health by being able to predict when a heart is at-risk through regular, on-mobile monitoring! Possibly by v1 but almost certainly by v5 we will be able to predict with some reasonable percentage of the cases the probability of a heart attack hours and then days and then weeks or years in advance as data accumulates and technology improves. Regular screenings would be possible through free or nearly-free ECGs at every pharmacy or retail store in America, India, or anywhere else because very low-cost devices that could be given away for free. A minimal charge for a screening ECG read by a machine that, even in v1, could then advise you to see a cardiologist for further investigation would be far better than what happens today. Bionic assist for cardiologists will become the norm and especially so where quality cardiac care is limited or unaffordable. Maybe stethoscope sounds accurately correlated to cardiac signal will provide further insight. Some currently known but seldom used variables like heart rate variability (because of difficulty of measuring these constantly) will be continuously measured and integrated into medical decision-making, and other metrics as noted before will be examined for patterns of change and correlation. Much of cardiology will be stuff we know how to do but currently require tests or diagnosis that is not easily or inexpensively accessible. v1 will bleed into v2, v3, v4... with continuously expanding capability and will expand the range of medical

knowledge and disease characterization and monitoring rules. Bionic assist will accumulate usable data on millions of users across multiple years leading to more diagnostic and monitor modalities and a whole new way to practice medicine.

Just as the v1 systems above can dramatically help certain specialties, the new sensors and devices can utilize dozens of different variables to manage core health and behavior. The v1 of physiological sensor- rich prototypes in a variety of form factors like bands, watches, and patches, will measure a number of variables with sufficient accuracy to help manage certain diseases like cardiac disease and hypertension.

System like this could collect thousands of data points per day, which could correlate symptoms with external conditions (such as weather, elevation, and location - does your office always cause an increase in blood pressure; if so it may be time for another job!). Today's car has over 400 sensors. Formula 1 racing has up to 200 data feeds from the sensors in their cars, generating sometimes over 5 GB of data per lap.⁹⁵ These are now vital for being able to compete. Similarly, Larry Ellison's sailing yacht for the America's cup were equipped with 300 sensors collecting over 3000 variables of data every 1/10 of a second.⁹⁶ An even though all these other fields are increasingly using sensors to improve, humans have few to none. For those concerned about their health, this will change (though not everyone will want to subscribe to this brand of medicine just as not everyone uses a smartphone and its many apps). We will start to truly quantify stress, sleep, fitness, hypertension, cardiac output, heart rate and its variability, accumulating billions upon billions of data points across patients, conditions, environments, food, and other habits setting up the data set for evolution v2. Examples would be combining wearable v1 data with mental health v1 data to measure behavioral and physiological parameters, along with environmental and other parameters continuously to be even more fine-grained, precise, and accurate. Even Facebook or Twitter posts may figure into a diagnosis.

⁹⁵ <http://www.zdnet.com/big-data-in-formula-1-racing-7000006919/>

⁹⁶ <http://www.forbes.com/sites/danielfisher/2013/12/23/winging-it-how-larry-ellison-harnessed-big-data-to-win-the-americas-cup/>

Surprisingly, this fine-grained measurement will lead in v2 to much more holistic assessment of a patient and integrative assessment of a patient's mental health.

Computers and data systems (not directly driven by the consumer) will cost effectively do and handle more than our current healthcare system can provide. IBM's Watson or various startups' v1 could be substantially more powerful than its current system or systems like Epocrates (a drug interaction application – drug interactions are a cause of serious medical complications in many of today's hospital patients) by truly understanding much of medical literature and research and knowing how/when/where to apply it.⁹⁷ It is also dynamically updatable as new drug interaction literature comes out. Even currently, studies show that many a doctor use Wikipedia for reference material.⁹⁸ Stefano Bertozzi, dean of public health at Berkeley, AIDS expert, and former Gates Foundation senior fellow, remarked at a recent conference in San Francisco that Wikipedia has been the single greatest advance in public health in the past decade. A graph of medicine can take information from symptoms, diagnoses, and procedures and examine the relationships between each one based on scientific medical data, EHR data, and claims data to create a graph-based representation of all relevant variables in medicine (which can be expanded over time with new nodes and edges). Combine this with the knowledge graph of biomedicine (much like Google's general Knowledge Graph which derives much of its data from the their Knowledge Graph as well as harvesting its own data, ranking its reliability and compiling all results into a database of over 1.6 billion facts collected by machine learning algorithms) we will likely capture from the tens of millions of medical research papers, and we have both research based and practice and data based knowledge.⁹⁹ This medical graph could become the knowledge repository based on medical literature, helping develop guidelines for a treatment protocol available through a mobile or web app to any doctor. This is the next generation

of an interactive medical Wikipedia! Initially, the early versions of this medical graph will just classify diseases, treatments, and procedures as we know it and surface inconsistencies in literature. Then, this same information could be made available to a patient empowering them to ask knowledgeable questions and to assess their options intelligently.

Complementing the medical literature and cross checking it for accuracy is basing its knowledge base on the body of past medical practice captured in patient records. This could enhance internal medicine by being a more comprehensive and unbiased medical detective. It could allow a doctor to ask, "how were patients with similar circumstances, symptoms, complications treated and what were the range of outcomes? What are patient preferences and which of these statistically valid choices matches the patient's preferences?" Though the graph of medicine improves how currently known medicine is practiced, such "complementation" is likely to start to change what we know by adding to or invalidating parts of what we know. It will be the beginning of changing medicine to be much more personalized and precise. Such v1 systems will get substantially better, reading and parsing much more of the available medical literature (and the anecdotal stories around medical care?), continuously updating their body of knowledge and capturing the expertise of the best specialists in each area. Subsequent versions can stay up to date with the thousands of published papers each week and even check if that new research has conflicting conclusions with previous literature. The v0 better care path development and protocol developments will give way in v1 or v2 to much more sophisticated care paths with many more personalized pathways based on both a patient's circumstances, ethnicity, condition, and preferences.

This would be too complicated a protocol set for humans to follow if it was paper based but would be easy in the context of a mobile app.

97 Piatkov, Irina et al. "Drug Interactions, Pharmacogenomics, and Cardiovascular Complication" Drug Discovery (Jan 2013) (<http://www.intechopen.com/books/drug-discovery/drug-interactions-pharmacogenomics-and-cardiovascular-complication>)

98 <http://www.theatlantic.com/health/archive/2014/03/doctors-1-source-for-healthcare-information-wikipedia/284206/>

99 https://en.wikipedia.org/wiki/Knowledge_Vault?oldformat=true

The v1 “bionic assist” systems can go a step further from the current state of the art, by helping physicians identify missed diagnoses seamlessly. Or perhaps by making the latest research, which is relevant to a particular patient, readily available for providers, functioning as contextual memory recall or an extension of their brain. On the patient / consumer side, bionic assist can help bring up good questions for the patient to ask their physician, or provide a second opinion when necessary. In the early going and over time getting more and more comprehensive, surpassing human capability. Early systems will be relatively simplistic but each evolution will get increasingly more sophisticated.

Physiological sensors, many of which will be integrated into medically relevant wearables, will detect atrial fibrillation (AFIB) through continuous inter-beat (RR) interval monitoring and heart rate variability (HRV) monitoring that isn't possible today except in research settings, enhancing current known science and its applicability. The digital health systems will discover previously unknown cases of sleep apnea, poor drug compliance, unknown causes of stress and hypertension, and a host of other indications. But because of vast amounts of data we may discover flu infections hours before the patient feels it because the body has started to respond. Other diseases may show early signs in complex patterns of physiologic data or combinations of physiologic, biomarker and other data. Drug compliance may be monitored based on physiologic parameter shifts. Managing blood sugar may be dependent on current blood sugar, food intake and what body physiology is currently doing and needs to do to avoid certain outcomes (we have seen this proven recently with Eran Segal's research).¹⁰⁰ Every cup of coffee may be automatically recorded because of its impact on the physiological parameters and the affect of coffee consumption on various diseases or states like stress, or mental health or cardiac disease will be truly quantified. Combining physiologic

data with outputs from systems that measure mental health, ECG or other electrical data, and biomarkers from blood sugar to cortisol or a low cost chip to measure all immune system biomarkers (monthly or even weekly for patients with chronic disease) will permit further insights hard to imagine today, hopefully improving patient outcomes as we move to later generations of systems. In fact, Apple's announcement of HealthKit is their way of enabling this reality – moving data collection away from app silos and into a common repository.

Even very simple technologies, such as V0 technologies like glucose monitoring will enable highly customized lifestyle support for improved health and wellbeing. In v1 and beyond there will be increased ability to provide highly customized diets. For example, recent work from the collaborating labs of Eran Segal and Eran Elinav, has shown that each individual has a unique, consistent blood glucose response to different ingested foodstuffs, partly attenuated by the gut microbiome.¹⁰¹ The practical benefit may be the ability to recommend foods that will precisely predict expected blood glucose in a few hours based on any given meal. This could be more key to good blood glucose management than today's pharmaceuticals.

Genetic differences seem to underlie differential responses in blood pressure to sodium intake.¹⁰² Increasingly, individuals have antibodies against common food products, ranging from mild sensitivities to life endangering allergies. Lactose intolerance, the lack of adult expression of lactase is a well known inability to easily digest lactose by a huge segment of the world population. Vitamin D and iron deficiency is widespread even in the industrial world, while conversely the genetic variant of the HFE gene which can lead to iron overload syndrome (especially in men, because they never menstruate) is extremely common in those of Western European ancestry. Common mutations such as those in MTHFR which lower

100 Segal, Eran et al. “Personalized Nutrition by Prediction of Glycemic Responses.” [http://www.cell.com/cell/pdfExtended/S0092-8674\(15\)01481-6](http://www.cell.com/cell/pdfExtended/S0092-8674(15)01481-6)

101 Zeevi D, et al. 2015 <http://www.ncbi.nlm.nih.gov/pubmed/26590418>

102 Wang Y, et al. 2014 <http://www.ncbi.nlm.nih.gov/pubmed/25531090>

enzymatic function and methyfolate levels may motivate pharmacological supplementation.¹⁰³ These are just examples of many of the molecular features that may be used to create customized food/nutrient plans (including customized probiotic and prebiotic supplementation). Although at first this might seem suggest to the imagination, hermetically sealed packets of something disgusting, that doesn't need to be the case. Molecular measurements are leading to new insights into taste preferences. For example, genetic variants in the gene OR6A2 determine whether a person likes cilantro/coriander or if it tastes disgusting to them, methods from biological network analysis are being applied to flavor pairings,¹⁰⁴ part of a future perhaps of food products not only personalized to improve health and wellness, but also to maximize enjoyment. None of this is cost-effective without computerization of interpretation and recommendation.

Companies with digital home health kits (I hope digital first aid kit for the home will appear sooner than later) will enable parents to monitor ear aches from home, using image processing to do initial diagnosis and recommend treatment, perhaps initially supervised and authorized by an ENT. New protocols for antibiotics prescription may be developed based on the ability to more closely monitor ear infections. The systems, based on classifications of hundreds of thousands of images, will likely provide bionic assist to all manner of health professionals via classifying images into not only the most likely diagnosis but also potentially rare but serious conditions. They can also assist by suggesting tests or questions to eliminate those possibilities. Such systems applied to dermatology may discover new diagnosis for skin conditions. Already a company based in the UK has developed a system that is able to detect skin cancer based on the fractal pattern of blood vessels in the image of a skin lesion. Such new ways to analyze images will be discovered often and continually lead to improved versions of systems with each generation having more capability. Today's basic v0 imaging and those auto diagnoses of current known diagnosis in v1,

will increasingly evolve into new, more precise, and sophisticated systems in v2, and beyond that use new medicine not used today!

The range of treatments, effects, and approaches are not limited to what we know, and sometimes, require a fundamental rethink from what we think of as healthcare. v1 technologies will show wide variety from the very serious to the frivolous sounding but effective. It is possible that just like the mobile app being certified as a pharmaceutical drug to treat diabetes and its equivalence to chemical drugs we will see behavioral or versions of most mental disease drugs and treatment protocols. Such "software drugs" are likely to have fewer side effects and hence will often (but not always) have simpler FDA approval paths as biosafety and will be a smaller issue. Whether these are v1 systems with relatively simpler new approaches or more complex problems and need more sophisticated development and hence are v2, v3, v4 systems and need longer development and testing is hard to predict today. These are the unknown unknown's of the medical world. But we speculate that digital drugs like the early ones for diabetes will be common for many disease states, often using patient motivation and behavior change as tools to affect the desired medical goals. Many conditions will start to appear as apps: Sleep order/disorder, stress and stress management, obesity, hypertension and many other apps will be attempted. Some will be more successful than others and many will take multiple revisions to be good enough to be broadly accepted.

Expanding on that further, even though the rate of progress is unknown, it is likely that digital drugs can be as or more effective than pharmaceuticals today to deal with the complexity of the human body and mind. Many of drugs used today are focused on suppressing symptoms rather than fixing the underlying problem of disease. Digital drugs focused on behavior change, coupled with "drugs" that can impact or change things in our microbiome that might be the root cause for certain conditions, will likely do a much better job of addressing disease as complements to

103 Gardner KR, et al. 2014 <http://www.hindawi.com/journals/psychiatry/2014/730956/>

104 Ah YY, et al. 2011 <http://www.nature.com/articles/srep00196>

or replacements for traditional chemical drugs. And these approaches will likely be much easier to personalize as well. By treating the body as a complex system, and recognizing the nuances of interaction between the body's myriad of components and systems and our mind, we can take system-level approaches to combating disease, rather than just focusing on, for example, increasing serotonin receptivity in the brain to combat a wide range of conditions. This in some sense will bring the best of behavioral medicine powered by the ability to understand and communicate at scale because mobile phones and to have conversational AI agents that can nudge us many times a day to addressing disease. Practices (from proponents like Dr. Andrew Weil, Dr. Deepak Chopra and Dr. Dean Ornish) like meditation and mindfulness or versions of wellness medicine will be quantified and applied utilizing the power of the cell phone. Already we are starting to see the gene expression changes driven by meditation as early examples of tools that will be well characterized and applied at scale thru digital devices.

Rich interactions with other systems, with much better validated medical therapies, procedures and protocols, and results of the millions of real-life experiments that doctors run inadvertently will start uncovering new medical relationships. And doctors do run these experiments all the time, as each applies their own belief system within (and sometimes outside) "approved medical practice" because of the way the system is set up today with all its points of unintended

(and often unscientific) variation. All that, along with increasing amounts of wellness care data, traditional medicine data, holistic care, homeopathy care, and other data hopefully will get more completely captured and will inform the best "science of medicine". Patient care protocols across hospital systems and geographies will get better, more consistent, and more personalized. They will be better validated than the opinions of individual medical practitioners, being an aggregate of their collective wisdom tested by statistically valid outcomes data. Within v1, systems may start to objectively score the pathologists and all other doctors on their adherence to known science (and that will be controversial!).

The v1 efforts arise out of entrepreneurs and other innovators within the healthcare system asking the most fundamental 1st-order questions that force a rethink of current "best practices" and a challenging of assumptions. These sorts of questions will give rise to hundreds or thousands of such v1 efforts, each of which will go through multiple iterative versions as they respond to challenges, roadblocks, inefficiencies, and other things in the way. These efforts will collide and adapt, especially as these systems continuously bring more data together, and push more logic and learning into sensors and machine and surface many a contradiction in current practice. In some cases, the system and actors will adjust in response to these efforts, sometimes pushing back and sometimes adopting new approaches while helping to mold them further.



v2, v3, v4...

Typical of a complex system, iterations of what start out as v1 efforts in healthcare will emerge in a chaotic process of trial & error and tweaking in response to points of failure, critique, criticism and imperfection. By necessity, this process will be an organic one, in which both new ideas and the system itself adapt to each other to find modes that start to work and build momentum.

In each of many different areas of medicine, we will see new capabilities emerge every 2-4 years (creating a new generation or version) in the form of these intelligent software or software/hardware systems. In contrast, pharmaceutical discovery will most likely stay in its traditional 10-15 year development & approval cycle, or become at best, half the time because of drug discovery modeling. The nature of medical studies, validation, and regulatory approval will probably have to change to accommodate this new rate of change (or we may place a stranglehold on this new system and delay it substantially). On top of that, there are new software-based companies that attempt to cut this time in half by utilizing machine learning and analytical techniques in the drug discovery process. These efforts can both help reduce the time to develop a drug, but also greatly increase the probability that pharmaceutical companies test drugs with the highest likelihood of success. But our focus is more directly on digital technologies that change how medicine is practiced rather than on software tools for better drug discovery or better robotic surgery. Keep in mind as we get 5+ years out predictability turns to speculation and the further out in years we try and imagine/speculate the more "wild" our speculation gets. This is in contrast to your median physician who in 5-10 years is likely doing something very similar to what he/she is doing today.

As an intertwined relationship between physical technology and a learning healthcare system, we discussed v1 devices such as bands that can be used to monitor heart rate and activity; however as exciting as this technology it may be stepping stone in terms of technology. The rise of cheap commodity antenna/receiver technology prompted by the growth of mobile communications turn the "smart home" into a complex system of interacting EM signals. At first the existing WiFi infrastructure may just be used for simple motion tracking and burglar alarm systems, but advances in data transmission and signal processing will enable low energy microwave and radio waves to passively monitor human activity, not perhaps at the spatial resolution of higher frequency bands such as visible light or X-ray, but certainly at the level to track motion, breathing, and even heart rate and blood flow/perfusion. The ability of EM waves to penetrate into tissue and reflect back may allow passive, continuous observation of changes inside the body. Rarely will one will need to use a stethoscope or blood pressure cuff if such technology can continuously monitor cardiac output and measure volumes of gas exchanged at each breath. This is technology beyond v2 and part of the repeated trends to move to less invasive, more continuous monitoring, more richly featured data collection, provided new features and potential biomarkers and simultaneously offering a more vigilant and convenient system.

In every medical area and specialty, advances are being made; these represent just the first iteration of technology or "point" innovations. Early medical systems will be used in non-critical roles or under physician supervision. They will continue to grow with the help of the top doctors who will be AMPLIFIED by the systems to be able

to understand, recall, deduce and do much more than is possible today by humans alone. The best specialists in every area will make their expertise available to many more patients through digital technologies. So every person can have the best, most researched care — not the overburdened and rushed doctor the average person gets today, especially in the developing world! In fact, the very best doctors, with the skills of the fictional television character, Dr. House, will be an integral part of designing, building and enhancing these advanced systems over time. Data will invalidate some established medical practices, uphold others, and invent new medicine. Just as one company is trying to replace a colonoscopy with a blood test, many other disease characterizations or immune or nervous or circulatory and other system responses will be based on biomarkers or patterns of biomarkers that are not characterized or validated today. As such “tests or monitoring or drug dosing techniques” develop it will spur innovation in measurement technologies to make them painless, low cost, and easily accessible. No more going to a special blood draw location and sending of samples to expensive central labs to get hundreds or thousands of biomarkers one may be interested in tracking, except in rare cases. Already one company has a finger prick test for the doctor’s office (the “system”, if research pans out, will look like a laser printer with a “chip” much like a blood sugar measurement test strip) that measures all immune system related biomarkers in ten minutes with no local expertise needed for the test. Annual physical exams may actually be useful as such blood tests indicate how your body’s thousands of metabolic pathways are trending and what potential future dangers lie ahead. Whether it is future diabetic complications like eye impairment or limb amputation, or early signs of Alzheimer’s or cardiac disease’s that are not yet visible, the signature might often be in the physiologic variables or a biomarker’s complex patterns well in advance. Digital or behavioral solutions can be utilized to slow the progression of disease or change the metabolic pathways. And the guidance for future researchers looking to solve these problems will likely come from insights developed from these techniques and other big data. Many will be calling for more rigorous studies of preliminary conclusions and new insights this data provides, but real confirmation will necessarily take years of tracking a large

population of people following this new approach to medicine. Randomized double blind studies will be an important cog in a much larger wheel. But in order to test these ideas we will need pioneers who apply these ideas and collect vast amounts of data. Many will resist this trend towards unvalidated ideas even if they are “more validated” (and true) than many “human” expert opinions. Experts, especially institutions like the American Medical Association, whose members stand to lose income/power if this scenario happens, will try and slow down this progression if history is a guide under all sorts of excuses.

This isn’t something to fear or regulate against but rather encourage and revel in, because the nature of innovation, including in healthcare, is that real progress comes from complex interactions among independent developments as they build on each other. Keep in mind that without some experimentation and risk innovation slows down and does society more harm than good. Take the current example of driverless cars that is playing out. It is clear that current autopilots are not perfect but even in this imperfect state they will likely save more lives/damage than the deaths/damage their imperfections may cause. Society will clearly be, even in the short run, better off adopting them and putting them on a faster learning curve (only use of this technology in the field will accelerate learning and reducing the flaws). But many, especially those behind or otherwise self interested in delays or otherwise naïve views, will call for delays even if it (like the Hippocratic oath) causes more harm than good. New players combine the best of previous systems with new twists and a complex ecosystem of learning emerges. This new approach will be substantially more effective (and cost-effective). Often directions of study in the more fundamental sciences like the biology and genomics of disease will be prompted by data science observations. The v1 systems will get better and better and over time, especially if they can get sufficient use, becoming more comprehensive and holistic, and will include more and more factors which are too complex for humans to comprehend. Some areas will progress rapidly while others will go slow or even be proven wrong and regress. From the understanding of this complexity, as analyzed and “insighted” by sophisticated machine learning systems new

understandings new science of medicine will start to emerge.

Physiological sensors, many of which integrated into medically relevant wearables, will correlate across diseases and other measurements. Questions we will want answers to will include the likes of: 1) How does each meal affect my illness? 2) Why does my disease feel better today than yesterday? 3) Why was my blood sugar higher last Thursday than the day before? How does it correlate to my blood pressure or what I ate yesterday? 4) How well is my cancer treatment really working this week? The answers will lie, if at all available, in complex correlations between all the variables we have learned to track.

Though many innovations we attempt will fail, enough of them will succeed to change medicine measurably but in ways that follow current medical knowledge around heart rates, blood pressure, CMP blood test panels, and the like initially. Over time, the 15,000 or so diseases and the 15,000 or so therapies will get rationalized, retested, personalized, and systematized & applied in a more consistent and precise way based on much more data for each given situation than is available today and far more comprehensively than a human could provide. The graph of medicine will evolve, extend, and change. Many studies will be upheld and strengthened, while other current medical practices and beliefs will get invalidated or get validated for specific circumstances. Though such medicine is unlikely to be perfect it will be more accurate and precise than we have today. Especially as we integrate the nexus of physiological data, biomarker and internal data, EMR data, genomics data, and these knowledge graphs together.

Professors such as David Sontag are working on how to couple newer forms of machine learning (such as deep learning) with probabilistic inference in order to make new inferences about disease progression that are rooted in scientific research.¹⁰⁵ These types of systems will bring medicine to more of a science than practice, especially as newer data sources are integrated.

Other companies that initially will look more like telemedicine or second opinion services have the intention to build AI systems which can help navigate through various symptoms in a personalized way and operate as both a diagnostic and pharmacological companion.

It is likely that the practice of medicine will not be proven wrong but rather proven not precise and consistent enough. The reinvention of medicine coming from thousands of data science based discoveries and new insights will rewrite medical literature with better diagnoses, monitoring, treatment, counseling, and other recommendations in a decade or two. Today's textbooks likely will become increasingly obsolete in many categories of medical education over time as better insights and methods are discovered. Medical school students slaving over their current textbooks will soon be utilizing apps to help them gain this new knowledge. New and even better systems will be developed. The above are point discoveries, but hopefully, in a decade or two years, such findings will become the norm as systems make these "medical discoveries". Progress through v0/v1/v2/v3/... will be building towards truly personalized medicine based on the most validated and up to date science.

The growing importance of the phenotype as a 'new' data source should not be overlooked. For decades, medicine and disease treatment was symptom-based and diagnostic. Then, there was the "genomics revolution" which we hoped would revolutionize medicine as we sequenced and understood the root causes of every disease from a DNA level. And while genetics and genomics has certainly started to change certain specialties in medicine, "it's becoming increasingly clear that sophisticated – and actionable – understanding of biology and disease requires not only a parts list, but also a nuanced readout of how the parts operate together in the context of a cell, or person. In a word: phenotype."¹⁰⁶ Genomics hasn't lived up to those grand predictions but is adding significant insight to specific areas like cancer. And genomic data is being used further and further to deliver medically relevant information

¹⁰⁵ Sontag, David et al. "Clinical Tagging with Joint Probabilistic Models". <https://arxiv.org/abs/1608.00686>

¹⁰⁶ Shaywitz, David and Ausiello, Dennis. "Decoding Phenotype: The Holy Grail for Today's Medical Scientists" The Atlantic, 2012. (<http://www.theatlantic.com/health/archive/2012/01/decoding-phenotype-the-holy-grail-for-todays-medical-scientists/251797/>)

as the cost of genome sequencing as come down exponentially – allowing companies to allow for “liquid biopsies” eliminating the need for some proportion of very invasive surgeries ... or to allow for mass genetic screening for key breast cancer genes. And there are still missed genomic data coming from the microbiome’s “second genome” and the epigenome as well as other “omics”. One particularly exciting field of research is Professor Steve Quake’s work on single cell RNA sequencing to accurately and quantitatively measure differences in the transcriptomes, giving us the ability to identify and profile of subpopulations of cells within a larger heterogeneous population, something that can’t be done with standard genomic methods currently.¹⁰⁷ Further, your RNA measurements from a blood test can give you a kind of “molecular stethoscope” to track general health issues in various internal organs and yield real-time snapshot of your body. On the 200th anniversary of the stethoscope – the iconic diagnostic device available to 40 million health care professionals worldwide – it is appropriate to launch the era of molecular stethoscope for a more precise diagnosis. It’s hazardous to speculate on what can be possible and especially on the path of change with complex domains, and we will certainly continue to see surprising innovations in helping better understand our health and our diseases.

Others such as Michael Snyder and Stanford are working to build a fully personalized all –omics profile of patients (iPOP). This research is still in its infancy, but Snyder has been able to get a very detailed assessment of participant’s physiological state by collecting their genome, DNA methylome, transcriptome, proteome, metabolome, auto-antibodyome, and their microbiome across multiple locations in the body¹⁰⁸. This comprises billions of individual measurements – that can then be condensed to hundreds or thousands of measurements that can be considered clinically

relevant in the future. Even in this research phase, Snyder’s own –omics profile indicated he had type II diabetes (as well as had a genetic predisposition to it even though there was no family history), and he was subsequently diagnosed with it in more “traditional” ways¹⁰⁹.

That being said, phenotypic data, which are heavily influenced by our environment, can now be electronically captured in increasingly deep (range of phenotypes) and dense (high temporal resolution) ways, and has been shown to give much better indication of a person’s health than other sources (and at minimum, will serve as a necessary complement to the current ‘point’ and genomic data gathered by doctors and researchers). It’s easy to imagine this dense phenotypic data could provide empirically useful segmentation information (i.e. which patients might most likely respond to a particular treatment), but it’s also important to recognize this data might, along with other data, also provide novel, mechanistic insights into the underlying pathophysiology of a particular disease.¹¹⁰ In these speculations we start to mix traditional biological research progress (which is substantial and accelerating too) with more of the digital approaches we are more focused on here. It is illustrative of the fact that clear domains will not exist and each area of progress will reinforce progress in other areas. Even more exciting is the possibility that abilities to find insight from complex data far beyond human capability (the new class of AI systems like DeepMind and beyond) will make these very complex and multi-dimensional data sets even more valuable.

Research groups are currently investigating phenotype to that effect. The US eMERGE initiative was an early of the phenotype. And has spurred a set of other very successful studies, utilizing the EMR to handle phenotyping (PheWAS) which can be mined in future routine clinical practice to

107 <http://www.nature.com/nmeth/journal/v11/n1/full/nmeth.2694.html>

108 Snyder, Michael. “iPOP and its role in participatory medicine” *Genome Medicine*, 2014. (<https://genomemedicine.biomedcentral.com/articles/10.1186/gm512>)


109 Chen, Rui, Snyder, Michael et al. “Personal Omics Profiling Reveals Dynamic Molecular and Medical Phenotypes” *Cell*, 2012. “<http://dx.doi.org/10.1016/j.cell.2012.02.009>”

110 Shaywitz, David. “Can Digital Health Successfully Apply Lessons of Molecular Biomarker Development to Create Useful Measures of Phenotype (Phenomarkers)” *Forbes* (2013). (<http://www.forbes.com/sites/davidshaywitz/2013/11/24/phenomarkers-could-transform-personalized-medicine-can-digital-health-deliver/>)

drive care decisions. In fact, a Penn State study also showcased that PheWAS derived from EMR's accounted for ~66% of genome-wide associations (GWAS) across 14,000+ individuals. The image below showcases PheWAS associations for genes previously associated with glioma and myocardial infarction, along with many others. The current systems mostly just augment genotype, but given widespread use of EMR data mining, may soon be used to find new insights prior to traditional bioinformatics methods. Electronic medical record-based phenotyping, complemented by genomic and other -omic data, has the potential to provide a link between studies that advance the science of medicine (e.g. PheWAS Phenome-wide association studies is a quantitative research technique used by scientists trying to solve the mystery of: What disease associations can we make with a given gene? This is in contrast to GWAS (genome-wide association studies) which asks: What gene is associated with a given disease? Living organisms possess physical and biochemical properties referred to as phenomes)

and studies that advance the practice of medicine. It is worth pointing out that existing EHR data will likely be part of early systems but will be over time be small compared to these other "omic" and other data sources.

Related to the phenotype is the microbiome, the ecological community of microorganisms that are in our body (bacteria, fungi, archaea...). The microbiome is a new source of data that has been yielding new insight as a predictor of disease (just like the sensor data / data from the "point innovation" systems are giving us new but actionable information). For example, the microbiome has been shown to strongly predict children that would eventually go on to have a diagnosis of type 1 diabetes. As the chart below showcases, these children had a significantly reduced level of microbiome complexity shortly after birth. ¹¹¹ The reduced level of complexity and stability in microbia in the gut is now shown to be increasingly important for potentially predicting (or shedding light) into a range of diseases, including



y-axis: microbiome complexity
blue: type 1 diabetes (diagnosed around day a1000) green: control group

¹¹¹ <http://www.diapedia.org/type-1-diabetes-mellitus/microbiome-and-type-1-diabetes>

diabetes and Crohn's disease. Though this is still research, large-scale data, knowledge and insights could lead to large changes in every area of medicine. A lot of microbiome work is very preliminary, but one company has announced a joint collaboration with the Mayo Clinic for microbiome targeted diagnostic testing starting with preterm labor.¹¹²

Most of what eventually happens here and the insights that could emerge will be hard for us to imagine today, but the rise of large amounts of data will allow us to pose questions that currently go unasked. New data analytics systems will handle thousands of columns and hundreds of millions of rows of physiologic and other data, helping us to surface not only the right answers, but suggest what are the right questions to ask in medicine to lead to iterative discovery. Today's hypothesis driven medical investigations are limited by the questions one can hypothesize. Many more smart questions are possible and will surface from data correlations and knowledge graphs! This will be the nature of medical research and will be then packaged into simple apps for deployment with health care professionals who will generally be insulated from such complexity. Hopefully their lives will get simpler than they are today allowing them to focus on the more human elements of care. We will likely find the answers to questions we never knew were important, leading to new medicine and new implications for our health. Complex correlations, designed, hypothesized based on biological sciences, sometimes just intuited, or conjectured or accidentally discovered (probably a very large category of new discovery when data sets get large) will contribute to a rapidly improving science of medicine. This same pattern already exists in other fields. Target, the retail chain, can now probabilistically predict when a woman is pregnant just by looking at how her purchasing habits change over time. Google has become famous for its AB testing on user behavior and, though we may not know why a particular color of button is more likely to lead to a higher click-through-rate on an ad, we know it does and that knowledge is used constantly to affect performance. Prior to

the advent of the data and the resulting insights human advertising executives could not fathom many of these factors as important. A similar pattern will happen in healthcare, as v2 onwards helps us discover non-intuitive relationships that lead to new understanding. Some startups will eventually treat the full patient health record as a search term against 100M record EHR database with many care episodes in each EHR of the medical history of patients, automatically finding the most relevant inadvertent experiments and uncovering the best outcomes. The question we asked earlier "What if comprehensive software was available to manage, analyze and interpret the relationships between diseases?" is more easily answered by such a system.

v2+ of a mental health app may intervene when it detects a patient with declining mental health by sending puppy pictures or asks a parent, sibling, or friend or sibling, or other triggers to call the patient and possibly reduce the probability of a downward spiral at the beginning of a mental health episode. The hundreds of behavioral insights will lead to better, more precise diagnoses than today's diagnostics manual (DSM5) on its perch on the bookshelf, over the next few years, reducing disagreement between multiple psychiatrists evaluating the same patient. v3 or v4 may start to look like a true psychiatrist with more precision and consistency while providing great bionic assist to human care providers, and maybe even Ai to do talk therapy. In v2, v3 or v4 of the cardiologist assist application may look at the temporal evolution of irregularity or other ECG anomalies among the 150 per year ECG's or a cardiac patient. Or the application may find correlations across ECG "features", heart rate variability, blood pressure, hydration or any other physiologic or physical data being monitored and automatically provide diagnosis or suggestions or alerts. It is likely to be predictive of adverse changes or likely cardiac events. Similar progress will happen in smart endocrinologist software that actually starts building custom models of each individuals body to predict blood glucose or hypertension hours in advance the recommend foods, exercises, or other activities to manage it substantially more precisely than is possible today (just as we are seeing in

¹¹² <http://newsnetwork.mayoclinic.org/discussion/mayo-clinic-and-whole-biome-announce-collaboration-joint-development-of-microbiome-diagnostic-testin/>


the aforementioned Eran Segal research). Care plans will be precise and truly customized, auto generated by intelligent software.

That knowledge will in future versions be applied to the special situations of each individual patient in the domain of the 15,000 diseases and the 15,000 diagnostics and therapies available to medicine! Very few human beings will be able to do this as well, but they will be able to substantially enhance the insights gathered by such systems, at least for the next few versions. Man plus machine working together may, at least for a while, do better than either one alone just as even a mediocre chess computer assisting a human can beat the best chess computer today. v2 may re-invent research studies and enhance the quality of medical literature by validating or invalidating or narrowing the scope of each study by using data from orders of magnitude more patients in real circumstances than used in the original study. The current fashion in medicine is to be more “evidence based” because not enough healthcare professionals use such evidence base medicine, instead relying on their personal experiences and often biases. This has been shown to be inferior to the evidence based medicine that has been broadly acknowledged as the right way to deliver care. But if much of the evidence is wrong what are we to do? One can hope to develop “practice based evidence”, i.e. evidence that is based on the millions of practice decisions we have already made in the past and where we know the outcomes. We hope to define what evidence we accept today is right and what is wrong or what other conditions it may be right or wrong under. Usually such complex nuances are often hard to correctly define in small studies. This alone could be a significant change to medical practice and a pathway to highly personalized and top-notch scientific care in v2 and beyond. No doctor will operate alone just as no computer chip designer operates without the help of computers and systems. Design verification is near impossible

without computers along with many other design functions. And the more complex the situation gets the more such systems are needed. Some of the innovation going on is just early experiments and research. As surprising as this sounds, consider the following example: Since 1928, the way breast cancer characteristics are evaluated and categorized has remained largely unchanged. Evaluation is done by hand, under a microscope. Pathologists examine the tumors visually and score them according to a scale first developed eight decades ago. But a new research paper¹¹³ from a machine learning group at Stanford states “Pathologists have been trained to look at and evaluate specific cellular structures of known clinical importance, which get incorporated into the grade of the tumor. However, tumors contain innumerable additional features, whose clinical significance has not previously been evaluated,” said Andrew Beck, MD, a doctoral candidate in biomedical informatics and the paper’s first author. “The computer strips away that bias and looks at thousands of factors to determine which matter most in predicting survival,” said Koller. C-Path, in fact, assesses 6,642 cellular factors, far for more than a human could, especially at reasonable cost. An at some of these kinds of evaluations, incorporations in software systems will be both better and virtually free. Ultimately, the computer yielded results that were a statistically significant improvement over human-based evaluation.

That same research group used machine learning (unsupervised) on continuous monitoring data for newborns in the ICU to predict infant morbidity rates at a much better rate than Apgar, which is the current standard of care assessment for newborns. The figures below showcase how noisy physiological data (top figure) was insighted by the machine learning technique to find signals otherwise unforeseen in the data (bottom figure). These can lead to a potential new way to monitor and treat newborns.

113 <http://stm.sciencemag.org/content/3/108/108ra113> <http://med.stanford.edu/ism/2011/november/computer.html>



Particularly, here, signatures 2 and 5 appeared most likely in sick infants. While signatures 3, 9, and 10 are more common in healthy infants. And as can be seen in the diagram, the “bad” signatures had much lower entropy than normal. But practical encoding of the end user experience will be a better Apgar score for a health professional to use, not the complex waveform we see here. And these are only research findings today still far from clinical practice because a commercial vendor is not implementing them.

An even more recent study has similarly shown that machine learning models were better able to predict the prognosis of non-small cell lung cancer patients better than pathologists were able to. This data was gathered from fully publically available data sets – stained histopathology whole-slide images from The Cancer Genome Atlas and additional images from the Stanford Tissue Microarray Database¹¹⁴. This is the beginning of algorithmic medicine that transcends human capabilities.

Radiology and these examples of pathology in these phases will mostly be driven by advances in computing and machine learning, especially with the interpretation of images. The early stages right now are reading scans to detect diabetic retinopathy; over time, we will see this type of image analysis be used on X-Rays, CT scans, ultrasounds, and MRI’s. Just as ImageNet performance (deep learning systems automatically classifying images in general)

continues to improve and surpass human abilities, so too will systems focused on medical images.

Deep analytics also can enable a first-principles understanding of the drivers of disease. One analytics company for example, applied their solution to analyze a publicly available breast cancer dataset that was already 15 years old. It had been previously analyzed over the course of many years of medical innovation, but not with methods as deep as Topological Data Analysis (TDA). As a result, they found mutations associated with breast cancer patients and validated biomarkers for a pharmaceutical company in a matter of days, instead of months or years of typical biomarker discovery efforts. Such analytics and speed will help pharmaceutical companies and medical research centers access previously inaccessible data for a deeper understanding of the fundamental biology underlying human disease, and lead to better drug candidates. But likely the drug discovery process, though more likely to succeed and somewhat faster and more personalized, will still be slow compared to digital systems.

This group has also worked with Mt. Sinai Hospital to examine a rich dataset of genetic and clinical information for a population of Type I Diabetes patients. Using their software, Mt. Sinai researchers were able to better understand the disease subgroups, biomarkers, and biological pathways of their patient population. As a result, Mt. Sinai’s physicians could give patients

¹¹⁴ Yu, K.-H. et al. Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features. Nat. Commun. 7:12474 doi: 10.1038/ncomms12474 (2016).

the best possible care through more precise diagnosis, more targeted treatments, and more prudent guidance around discharge protocols. This not only improved patient satisfaction, but also reduced complications from one-size-fit-all medicines. Personalized (I use the term, "personalized" for "all omic" medicine to distinguish from precision medicine which often has genomic connotations. Medicine reflects a fundamental shift in understanding patients and disease. With genomic testing costs falling below \$1,000 per patient (and heading to \$100 or less), it is becoming more possible to determine a patient's molecular basis of disease. And with even lower cost sequencing, microbiome distortions or malfunctions, will be a routine part of treatment and diagnosis. As a result, healthcare providers can achieve more precise diagnoses, target therapies more effectively and, in some cases, prevent disease from developing or progressing.

V2/3/4/... drug discovery

The current state of the world has heavy up-front costs and low probability of success at actually having drugs make it all the way through commercial stage (even with the increase in simulation and computer technology over the past 15-20 years).

However, the aforementioned techniques of data-mining and machine learning techniques to clinical records and large molecular datasets can help improve this process dramatically. Large clinical datasets can certainly be used to help improve care delivery such as dosing regimens based on analysis of outcomes, but perhaps more interesting is the ability to discover surprising interactions and novel indications for therapies. The discovery of a side effect may turn out to be the more interesting and useful indication of a medication. For example, Viagra, Rogaine, and Propecia were all originally developed for other purposes: pulmonary hypertension, high blood pressure, and benign prostate hyperplasia, respectively. However, it is the unintended side effects of these medications that have turned

them into widely used blockbusters. Many drug interactions, or between medications and food are only discovered long after medications have been approved and marketed, and only discovered in an accidental, haphazard manner. The fact that grapefruit and grapefruit juice can have powerful interactions with many medications was discovered as an accidental surprise. Data scientists, such as Nicholas Tatonetti at Columbia University are using large datasets to accelerate and systematize the discovery of these kind of interactions and novel indications for medications, such as the fact that when the antidepressant is paroxetine (Paxil) is taken with the cholesterol lowering medication pravastatin, it raises blood glucose, an important finding, as both these medications are very widely used, often in patients with diabetes.¹¹⁵ When data scientists like Prof Tatonetti have access to the clinical datasets, every drug order becomes an experiment.

Far beyond the basic clinical data, computational techniques allows researchers to put together large datasets of molecular level data, including whole genome expression profiles in response to test compounds or in disease pathology, protein interaction and pathway networks, small molecule binding assays, and others. Recent work by researchers like Atul Butte of UCSF and Joel Dudley's group at Mt. Sinai, is showing that you can use data-driven techniques to do proper drug repositioning, which is the discovery of new indications for existing drugs. Drug repositioning has inherent advantages to the standard process of drug discovery because it is much easier to put then through clinical trials since they have already been vetted for safety, dosage, and toxicity (albeit for a different initial purpose). By looking at how the expression of genes across the genome are changed by disease, and by comparing those changes to how gene expression in human tissue samples are effected by a whole library of test molecules, an attempt can be made to find a drug which will work to address the changes in the disease profile to help treat the problem. For example, in one study looking at small cell lung cancer (SCLC), an aggressive subtype of

¹¹⁵ NP Tatonetti, et al. "Detecting drug interactions from adverse-event reports: Interaction between paroxetine and pravastatin increase blood glucose levels" *Clinical Pharmacology & Therapeutics*, 2011 (<http://onlinelibrary.wiley.com/doi/10.1038/clpt.2011.83/abstract>)


lung cancer, the researchers were able to find tricyclic antidepressants and related molecules that induced cell death in SCLC cells, while also inhibiting the growth of other neuroendocrine tumors. The graphic below showcases how the drugs, imipramine and promethazine prevent growth of SCLC tumors in a pre-clinical mouse

model.¹¹⁶ While this drug repurposing is still in an early stage, similar studies by the some of the same researchers have discovered that the epilepsy medication topiramate might be a useful medication for inflammatory bowel disease and that the anti-ulcer medication cimetidine may be another effective therapy for lung cancer.^{117, 118}


116 Nadine S. Jahchan, Joel T. Dudley, Pawel K. Mazur, et al. ----- Neuroendocrine Tumors: Antidepressants as Inhibitors of Small Cell Lung Cancer and Other --- Cancer Discovery (2013) (<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3864571/>)

117 JT Dudley, et al. " Computational Repositioning of the Anticonvulsant Topiramate for Inflammatory Bowel Disease" Science Translational Medicine 1 Aug 2011. (<http://stm.sciencemag.org/content/3/96/96ra76>)

118 Sirota, et al. " Discovery and Preclinical Validation of Drug Indications Using Compendia of Public Gene Expression Data" Science Translational Medicine 1 Aug 2011. (<http://stm.sciencemag.org/content/3/96/96ra77>)

B

Newer approaches use the latest innovations in machine learning such as multi-task learning and deep structural learning for drug discovery. Improvements in our computational modeling of how small molecules like drugs interact with the molecules of the cell at atomic level by researchers like Vijay Pande at Stanford are also expanding our ability to rapidly target diseases with medications. This type of work will likely be a big part of the future of drug discovery. It can start with specific diseases that have stymied traditional methods for curing and treating. But over time, the cost effectiveness of such methods will extend to all diseases, and this will be a core piece to how every pharmaceutical company will find new drugs and therapies (something like the v4 system). Companies are also working on this too – leveraging this new machine learning to find the right medicines or drug targets significantly more efficiently. Machine intelligence is a better tool and human brains or large scale random screening given the high dimensionality or complexity of this problem of finding chemicals that match particular receptors. A v7 drug discovery system will look more radical and novel. Almost all drugs (not just retargeted drugs) will be found using data-driven techniques, with data sources that will make the current genomic information look miniscule. And when those drug go through more protracted clinical trials, the actual delivery of medication will be much more personalized, taking into account genomic, phenotypic, microbiomic, and physiological data in terms of determining dosage and treatment for patients. The majority of the system will look different than it does now, but it will feel obvious

D

when we get to that world. But most important to our digital health thesis these drugs will be both selected (or designed) for an individual's genomics, epigenetics, and proteomics and dosage of course will be tuned based on gene expression and biomarker levels eventually.

And these are just a few out of likely hundreds of such efforts. For example, early data science efforts were more narrowly focused on genetics but bioinformatics is broadening rapidly in scope. Detecting polyps from blood biomarkers instead of colonoscopy is just one example, monitoring their progression or regression based on microbiota or other treatments is a fascinating possibility. Microbiota will certainly be used as drugs as will “digital drugs” that cause behavior change, including foods as drugs. Statin administration protocols at Kaiser discovered by data science becoming the single most powerful “drug” discovered for stroke mortality reduction is another example. The role of stroma in breast cancer pathology is another of hundreds of examples I suspect will emerge. Will some of the thousands of features in a pathology slide along with cell free DNA gene mutations and other biomarkers help specify treatment?

These emerging v2+ systems will not just spit out data, but they'll reveal insights, actionable recommendations, tradeoffs between the set of choices, and more. The decision of what to do may be assisted by a doctor, but the patient will understand too and be far more well-informed (maybe even as much as the doctor). And for

those without the capability to comprehend, systems could dialog at the patient's level of understanding and alternatively assist a doctor or other medical professional to advise them in making more intelligently informed choices based on their personal preferences. Personalization will also be based on our expanding understanding of how individuals respond differently to different medications, and maybe personal lab chips will guide this. Researchers in pharmacogenomics/pharmacogenetics are investigating how variants in an individual's genome modulate how a drug is metabolized and if and when it can actually be effective. Although this is likely to play a much bigger role in v2+ systems, there are actually examples in early v0 systems today. The anticoagulant drug warfarin has what is called a narrow therapeutic window, the amount of drug which helps prevent dangerous clots is not that much less than the amount of drug that prevents effective clotting and can lead to dangerous hemorrhage. Much of that variation in dosage is a result of a handful of known genetic variants. Clinical and genomic data from 4043 patients was used to develop a computational algorithm that returns an optimal dose.¹¹⁹ Similarly, there is aggressive research underway across oncology to develop precision, targeted, multidrug therapies for the different mutations present in individual

tumors; it is already the standard of care to test breast cancer samples for the key proteins that affect response to different medications.

Instead of standardized dosages of pills (25 mg, 50 mg, etc.) given out at the same starting doses to every patient, we can imagine a future of precision, personalized dosages. Basic physiological measurements such as kidney function and the rate that the blood is filtered, which can be estimated through a urine test, can effect how many medications are metabolized and the optimal dose, but rarely plays a role in determining dosing. We can imagine future of personalized precision medicine that takes into account individual genetic variation, basic parameter such as height, weight, kidney function, recent responses to therapy, and other relevant parameters to compute an optimal daily dose of medication for each individual, reducing the chance of dangerous side effects and increasing the likelihood of positive effect. We will be well on our way to the patient being the CEO of their own health, even if they are relatively poorly educated or less knowledgeable about medicine. The patient will know their preferences far better than any doctor could and will use these to pick their personally preferred treatment path, maybe with a healthcare professional's guidance.

119 International Warfarin Pharmacogenetics Consortium, Estimation of the Warfarin Dose with Clinical and Pharmacogenetic Data* NEJM, 2009 (<http://www.nejm.org/doi/full/10.1056/NEJMoa0809329>)



v5 - v7 & beyond

When there are enough of these point innovations, and as they evolve through intermediate versions, they will integrate with each other and start to feel like a tsunami of revolution. Each system will utilize active feedback to continuously improve. The pace of innovation will keep increasing as it has in most digitally driven technologies for a variety of reasons. The systems of approximately a few decades from now (give or take a decade!), which is the time frame I am talking about, will overcome many of the short-term deficiencies of today's technologies, medical professionals and institutions' politics permitting. The medical devices and software systems in two decades might be as different from today's systems as the floor-mounted, multi-pound sewing machine-like car cell phones with bulky handset cords of 1986 are from today's iPhones! New technologies will allow more complex analysis at every level from reading pathology of tumors with many more features than a human could or correlating the causes of stress or atrial fibrillation and suggesting instantaneous remedial measures via mobile notifications. Food may precisely control biomarkers used as measures of disease today.

In v7 of algorithms, we might be surprised! Entirely new relationships may surface, including those far beyond the capability of what we know to look for. What are the relevant features from hundreds of ECGs a year and what shifts in the patterns of the ECG under various condition/activities or at various times are most important for which conditions? Without a corpus of millions of patients and their longitudinal and contextual progression, these correlations are difficult (if not impossible) to detect or confirm. Without low cost high fidelity sensors we cannot apply it to 24x7 ubiquitous and non-intrusive monitoring at home or resolve the various spectral components

or cross-correlate it to heart sounds and their timing or to heart rate variability or body hydration or external variables like temperature, humidity or altitude, or to comorbidities or biomarkers. I would be surprised if we don't discover new ways to characterize cardiac disease and new ways to monitor it closely, unobtrusively, and inexpensively. Eventually, this would lead more importantly to having very fine-grained and highly personalized therapeutic recommendations. Even drug dosing may become very precise based on physiological, genomic, and other phenotypic data. The reading of ECGs automatically and continuously will be routine and many other symptoms of cardiac disease that we cannot correlate to the disease or its treatment today clinically will be better understood. Every patient's facial expressions, analysis of their reactions and emotional state beyond even what the most sensitive humans can read and correlate with mood could be recorded and fed into both their sick care and wellness care. It will be hard for a human to be as holistic as such a system. Reading trends in hundreds of ECG's, correlating it to thousands of minute data points a day, for years, making up millions of data points? What is a mere human being to do? And the cost of these millions of data points would be less than that of getting a single ECG at your hospital. Today's advertising systems though routinely take in billions of data points a day, process them in seconds and inexpensively turn them into programmatic advertising, getting us human victims to click on links or buy stuff we never intended to buy. These systems can predict who may like an entirely new product and target them because of accurate assessments of their psychometric profiles, their inclinations and interests or merely "hacking" their weaknesses. And this happens for ads that are worth a penny! Imagine if one focused this technology and

computing resource to much higher value medical events.

Data itself is not much use till it is reduced to information and then to actionable insights that will start to change the way we diagnose, treat and monitor disease. These insights will be deduced by AI and not humans. Medicine will be personalized by using these new data sources to yield information on both the individual and his/her context. At some point, possible in v5 or beyond, wearable devices will work together to chart every heart beat and measure your blood pressure by the second; data from daily breath & regular urine samples will be measured non-intrusively to help detect what's going on inside your body; activity data will be correlated to other metrics to determine physiological impact; every patient will have their microbiome (or some successor indicators) profiled and the resulting data will help tell us where specific targeted adjustments need to be made; and epigenetics will help us suss out environmental impacts on actual gene expression to figure out how our bodies are really working on a given day. Maybe even an implantable sensor that monitors blood chemistry continuously and with increasing fidelity and comprehensives over time will become commonplace, though biological intervention has a way of delaying deployment. All of this data, and other information we likely haven't envisioned yet, will create baseline models of every single patient and help us detect changes over time that matter. We'll likely find hidden correlations and causal relationships that uncover new knobs to turn to improve an individual patient's health.

The NIH's Precision Medicine Initiative has efforts to bring this personalization to medicine. At a high level, their goals are to advance research by collecting all the data we discuss throughout this paper (biological, genetic, phenotypic, medical records) for ≥ 1 million participants and track over time in order to better understand specifics for various diseases and conditions. Though there will be useful v2+ systems my belief is that at a certain point, in order to have an understanding of all this data which is being collected, we will increasingly

have to rely on systems and algorithms to come up with insights from the data. Or at least to help standardize and administer the learnings in an effective way once the research begins to be implemented – only requiring a doctor in circumstances when the system itself is failing (as is the case with pilots and autopilot systems in aircraft).

Mental health systems, by actually tracking behavior, could ultimately develop a fine-grained understanding of the patient and what makes them better vs. worse and what is effective therapy. This research is already being done on college campuses with positive predictive results.¹²⁰ The application should evolve over time to track the continuum of mental states from various illnesses to various states of wellbeing or happiness and provide constant recommendations to the patient over time to keep them well. Specifically, these algorithms and this information can be used to automatically be your personal guide and therapist. As a college student, they can finely tune whatever micro-therapies or suggestions are needed in order to optimize for grades, well-being, or whatever the student chooses to optimize for. The first step here is showing we can use these systems to predict certain conditions or states of mind; but over time, this will lead to personalized micro and macro interventions that will go beyond "goal setting" into active behavior management. The DSM manual's classifications will be labeled too blunt and even wrong, being supplanted by finer grained and more data science based micro-behavioral scores. It is a level of personalized care and attention that no human being could cost effectively provide. The right foods may be the most powerful drugs, just tuned precisely based on accurate measures of biomarkers, physiologic variables and microbiome. Probiotics and sparingly traditional drugs or biologics may be part of a continually tuned closed loop highly scientific system that actually quantifies the platitudes of practitioners like Mark Hyman, Andrew Weil, and Deepak Chopra. They may really be proven right, but scientifically and quantified and with sufficient qualifications as to when, how and what. The analytics engines

¹²⁰ Campbell, Andrew et al. "StudentLife: assessing mental health, academic performance and behavioral trends".


ACM International Joint Conference on Pervasive and Ubiquitous Computing. <http://dl.acm.org/citation.cfm?id=2632054> of college students using smartphones

will be created by fine-tuned models based on individual variations among millions of people (we probably will not call them patients anymore) with billions or trillions of data points, creating a system that will be much closer to the patient than any doctor could possibly be. It might actually make psychiatry a science with more predictable diagnosis, closer monitoring at a lower cost, and much better disease and wellness management. Today's early research on what the brain is doing and what influences it may turn in to software or software plus activities-based therapy with or without a chemical drug component. And the knowledge will apply to both treating illness and keeping wellness at high levels for all users. All these new methods will rely on massive data collection. Can we get functional MRI's of millions of brains using technology that is more cost effective to make this cost effective? Can we automatically collect behavioral data?

The vast amounts of per patient data on millions of patients and non-patients, and correlations to various other outside data sources like the weather, the economy, the stock market, and all the other external and internal variables that might affect patient well being, will likely take us well beyond monitoring to potentially predicting and treating illness. Already data-mining across large clinical datasets has revealed some surprising

relationships between environment and health. For example, the month of birth seems to be robustly linked to the likelihood of developing a wide range of conditions, including asthma, hypertension, ADHD and nearsightedness.¹²¹ Although it is easy to speculate why this may be the case, basic environmental and behavioral factors that change with season such as the amount and types of pollen in the air, hours and brightness of daylight, average humidity, amount of outdoor play, age at starting school, changes in diet based on food seasonality, and many others may all play a roll at key critical periods of development. As learning systems collect more and more data on the composition of individual diet, changes in the microbiome, patterns of activity and environmental exposure, and how this is attenuated by an individuals unique genomics, we will have a much more holistic picture of human health and wellbeing.

My dream is that a million people have laser printer like device at home that measures a few hundred biomarkers from a pin prick from the patient who monitors themselves every week or more often for a few dollars. Microbiome sampling that is also frequent and with species level detail could add to the dataset. Add all the physiologic variables measured continually or continuously



121 MR Boland, et al. "Birth Month Affects Lifetime Disease Risk: A Phenome-Wide Method" JAMIA, 2015 (<http://jamia.oxfordjournals.org/content/early/2015/06/01/jamia.ocv046>)

and we start to get a data set capable of (speculatively) characterizing what our body is doing. Add known genomic, epigenetic and other data and we have a party. Imagine how well we might be able to characterize the body's systems. Individual and population level knowledge and modeling and predictive capability will emerge (combined with knowledge graphs automatically extracted from research and mostly reconciled for inconsistencies.) We will learn to predict disease before symptoms happen and possible either correct or reduce these deviations from normal body's system behavior. We will be able to monitor progress and effectiveness of treatments, even fine tuning dosing of drugs.

Disease as we know it will also change. An interesting observation was made by the former Chief of Medicine at Massachusetts General Hospital, Dr. Dennis Ausiello who speculated that the term "diabetes" as a disease will disappear in the next decade or two just as the term Dropsy has disappeared. He presumes there are a dozen very distinct diseases that all have a common symptom in "poor blood sugar control" but in fact need very different treatment and management. Some of Joel Dudley's work is already showcasing this conclusion. The image below uses a software tool on diabetic patient genomic and biomarker data, with an emergent visualization that fairly clearly showcases different clusters (implying that all type 2 diabetes is not created equal or are not likely the same disease even if your doctor thinks otherwise – this is a dimensionless map of the space trying to find the topology of data, but emerges seamlessly out of this software). Over time, in v7 and beyond, this type of analysis will be done in real-time, routinely for every patient, so they will be properly mapped to their subgroup, well beyond today's research investigations. And that data will also be tied into the most effective treatments for those subgroups, yielding personalized drugs and therapies across all conditions, driven by math, statistics, and machine learning. Today we largely manage the symptom but as the practice of medicine becomes a true science we will manage the underlying disease as measured by the biological and physical parameters and networks it affects. Diabetes

will likely turn out to be many different diseases with a common system but varied treatments, not the blunt instrument we use today. Research continues to show that data analysis will be able to identify subgroups of diabetes – and based on genomic and phenotypic data (via the EMR), we are starting to understand what these different subgroups are ¹²². Some of the subgroups have marked comorbidities with cancer or neurological diseases, indicating they represent a different type of disease than traditional T2D. Suffice it to say this data is too complex for humans to understand effectively, and we need machine learning systems to make sense of it and enhance our understanding. The body is a complex and intermingled (usually through our blood, nervous and immune system) of thousands of metabolic pathways interacting with each other much like the internet and the national electric grid. If a cluster of power lines failing in Ohio caused a blackout in New York City as happened in 2003, imagine how complex the interactions between systems in our body can be where thousands of metabolic pathways resulting from tens of thousands of genes and the activities of RNA and microbiome based metabolites all get mixed up and communicate through our vascular system. We need to understand and characterize these complex network interactions to improve drugs and treatments, not justified point solutions. I am fascinated by complex systems theory and what it could do for medicine or what the newer AI systems can do for us. It's possible these may turn out to be less biology and more network theory problems. By focusing on detailed information of the phenotypes that are displayed across all diabetics, we will actually be able to understand what these differences are. But first we need to get well beyond the symptoms to the underlying metabolic pathways or network of pathways to understand what is happening. Of course the point of care health care provider will not see anything as complex as the chart above but will in fact be given a patient specific set of recommendations for treatment that will need less knowledge than today's endocrinologist.

Chronic Fatigue: A case study disease:

As the Economist recently described it "CHRONIC-

122 <http://stm.sciencemag.org/content/7/311/311ra174>

FATIGUE SYNDROME, or CFS, which afflicts over 1m people in America and 250,000 in Britain, is certainly chronic and surely fatiguing. But is it truly a syndrome, a set of symptoms reliably associated together and thought to have a single underlying cause—in other words, a definable disease? CFS's symptoms—debilitating exhaustion often accompanied by pain, muscle weakness, sleep problems, “brain fog” and depression—overlap with those of other conditions. These include fibromyalgia (itself the subject of existential doubt), clinical depression, insomnia and other sleep disorders, anemia and diabetes. These overlaps lead some to be skeptical about CFS's syndromic nature. They also mean many people with CFS spend years on an expensive “diagnostic odyssey” to try to find out what is going on. Skepticism about CFS's true nature is reinforced by the number of causes proposed for it. Viruses, bacteria, fungi and other types of parasite have all had the finger pointed at them. So have various chemicals and physical trauma. Evidence that CFS truly does deserve all three elements of its name has accumulated over the years but a definitive diagnostic test has remained elusive. For in a recent Proceedings of the National Academy of Sciences Robert Naviaux of the University of California, San Diego, and his colleagues published evidence that the metabolisms of those diagnosed with CFS are all changing in the same way. Their data suggest it is this cellular response to CFS-triggering traumas, and not the way the response is set in motion, which should define the illness. They also show that this response produces a chemical signal that might be used for diagnosis.”

Dr Naviaux and his team collected and analyzed blood samples from 45 people who had been diagnosed with CFS, and also from 39 controls who were free of any CFS-related symptom. They then trawled through those samples looking at the levels of 612 specific chemicals, known as metabolites, which are produced during the day-to-day operations of living cells. These metabolite profiles, they found, differed clearly and systematically between the patients and the controls. Some 20 metabolic pathways were affected, with most patients having about 40 specific abnormalities. The biggest differences were in levels of sphingolipids, which are involved in intercellular communication, though other molecules played a role as well. These differences

should give clues as to what is happening at a cellular level during CFS. More immediately, a handful of the abnormalities—eight in men and 13 in women—were enough, collectively, to diagnose with greater than 90% accuracy who had the disease. I suspect this is too small a cohort from which to make reliable conclusions (we need 45000 not 45 patient cohorts in complex data domains to avoid spurious correlations – a hilarious translation of these dangers is presented at <http://www.tylervigen.com/spurious-correlations>).

But it is a reasonable hypothesis that by taking a detailed, longitudinal systems biology view on the data and on patients, we will be able to see patterns of relevance that will lead to improved diagnostics, patient cluster identification therapeutic plans, and ultimately prediction and prevention of ME/CFS. Comprehensively measuring facets of human biology related to health, inflammation, immune response and systemic control over time (including response to diet, various treatments) will help fully characterize this and many other diseases. Numerous technologies that have shown promise in helping to understand ME/CFS (bacterial 16S microbiome, clinical immune profiling, gene expression profiling) and new technologies like single cell RNASeq, broad microbiome sequencing and many more hopefully to be invented or made cost effective to apply at large scale) will contribute to the data set needed to do a serious analysis and researchers are attempting this. Profiling these samples within an individual will give us a first look at the diversity of cell populations in ME/CFS and their heterogenous changes from the normal population and within the ME/CFS population over time as their disease is treated or severity changes.

What proof is there that this might work? There are no guarantees but there are indicators: RNA expression studies to date have shown notable differences in pre and post-treatment profiles.

Metabolomic panels used to fully assess up to 1,200 biochemicals in ME/CFS patients using mass spectrometry has already provided guidance on clinical systems malfunctioning in patients and suggestions toward therapeutic approaches. Surveying the gut flora to determine the impact

the microbial constituency has upon the health of the individual is significant. One mechanism underlying these effects involves the ability of the commensals to affect the immune system function through the priming of activated T-cells. The microbial diversity and relative quantity of patients and controls will give us a broad view of immune function (and dysfunction), inflammation and systems biology of ME/CFS patients.


But this case approach is not only applicable to CFS but to many if not most other diseases that are treated with relatively blunt instruments today. The diabetes example discussed previously will likely be dissected with a similar analysis.

Systems Biology:

Prof. Lee Hood of the Systems Biology Institute opines that in getting a disease the first networked perturbation in the body is the important one and simpler to identify. It then leads to a cascade of perturbations in the thousands of metabolic pathways in our body that result in various “symptoms”. He is an advocate of P4 medicine: health care that is predictive, preventive, personalized and participatory.¹²³ Medicine today

is a string of infrequent interventions prompted mainly by symptoms of illness. Hood argues instead for continuous management of health, making full use of whole-genome sequencing and biomarkers to correct disease before it gains a foothold. Hood is embarking on the first big test of his ideas: a nine-month pilot study, dubbed the Hundred Person Wellness Project, in which 10 healthy individuals will be intensively monitored (see ‘An examined life’), offered regular feedback and counseled on lifestyle changes such as shifts in their dietary or sleep habits. The effects of these behavioral changes on their health will, in turn, be tracked using a battery of diagnostic tests. This is an indicator of practices to come. Other efforts include the Google Baseline Project which is focused on collecting all biomarker and genetic information starting with 175 people (later expanding to the thousands) and making similar correlations and predictions about them. The NIH’s Precision Medicine Initiative described earlier is aiming to do something similar, and is working in conjunction with Verily (the team at Google leading the Baseline Project) and leading academic institutions to properly structure their studies and collect data.

¹²³ Gibbs, Wayt W. “Medicine gets up close and personal” *Nature* (11 February 2014) (<http://www.nature.com/news/medicine-gets-up-close-and-personal-1.14702>)



This revolution in data, and our ability to understand it, will lead to real precision medicine, one I call personalized medicine, (which differs from the "Precision Medicine" the NIH and others typically use). Throughout recent history, we would build medical technology that would be easy for humans to use (first and foremost) – starting with the stethoscope two hundred years ago (that we still use in roughly the same way today!). Over time, even though we have added more

and more technology to help aid with diagnosis and treatment, most functional oft-used devices are incremental advances over the previous generation e.g. blood pressure cuffs, heart rate, temperature (all incidentally devices that were available over a century ago because we knew how to measure these variables a century ago; even today these are the most common and first used in diagnosing a patient) EKG machines, etc. But as we shift to a data-driven, scientific world

of medicine where diagnoses and decisions can be driven by machines and algorithms rather than humans, technology we use and measurements we make should not be limited by humans' limited ability to process the data which restricts us to simplistic variables like blood pressure or potassium. In fact, the path of data and data understanding will have the biggest impact on the evolution of healthcare.

The US Precision Medicine Initiative is a good start. Among the scientific opportunities presented by this million-person cohort the initiative intends to assemble is the ability to:

- develop ways to measure risk for a range of diseases based on environmental exposures, genetic factors and interactions between the two;
- identify the causes of individual differences in response to commonly used drugs (commonly referred to as pharmacogenomics);
- discover biological markers that signal increased or decreased risk of developing common diseases;
- use mobile health (mHealth) technologies to correlate activity, physiological measures and environmental exposures with health outcomes;
- develop new disease classifications and relationships;
- empower study participants with data and information to improve their own health; and
- create a platform to enable trials of targeted therapies.

Real precision medicine will go further and ensure that patients can be diagnosed and treated in real-time, with continuous data streams constantly being evaluated to monitor and modulate health. This will go beyond just understanding genomic data (which by and large stays constant), but will take into account a host of biomarkers (measured daily), our microbiome, our voice intonations, our environment, and much more information to help

us make decisions. This data will all be combined in a dynamic knowledge graph derived and continuously updated from medical and scientific research literature, and this network will be key to understanding our biology from a systems level – and going from there to bring applicable, close to real-time changes to our bodies and lifestyle to either treat the root cause of a disease or reach a certain outcome.

There is a lot of overlap between the terms “precision medicine” and “personalized medicine.”¹²⁴ According to the National Research Council, “personalized medicine” is an older term with a meaning similar to “precision medicine.” However, there was concern that the word “personalized” could be misinterpreted to imply that treatments and preventions are being developed uniquely for each individual; in precision medicine, the focus is on identifying which approaches will be effective for which patients based on genetic, environmental, and lifestyle factors. The Council therefore preferred the term “precision medicine” to “personalized medicine.” However, some people still use the two terms interchangeably. Though today the NIH is moving towards precision medicine we may with algorithmic medicine enable personalized precision medicine!

Importantly, as we focus more on a wellness-centric model and move beyond a disease and pathology based model of healthcare, the patient as individual consumer and CEO of their own health and wellbeing will be empowered to decide what changes they need to make, and will have the support of software and tools to help nudge them into making the necessary changes to obtain the outcomes they want. As CEO, patients will need to be given a set of options and preferences, with cost / benefit analysis on the benefits and harms that they can use to make their decisions. Whether it is the dichotomy between wanting more energy, or the desire to be more relaxed, be more social and gregarious or more withdrawn and focused on a private task, our understanding of how all these factors effect an individual will empower patient as consumer to push his or her own physiology and psychology in whichever direction is most desired.

¹²⁴ <https://ghr.nlm.nih.gov/primer/precisionmedicine/precisionvspersonalized>

Being the CEO of your own health goes beyond just wellness though. The accumulation of data that is at the center of this healthcare transformation will drive the patient being able to take control of his or her own health. Just as tools, technologies and algorithms will be built to help clinicians and nurses understand the vast array of complex data they will have to deal with, so too will technologies be built distilling that data for patients. A universal health record will be owned and maintained by the patient, and will have not just traditional EMR data, but also information about his or her genomics, epigenetics, microbiome, and personal data. And as the CEO, the patient will have the ability to let other services analyze or interpret that data for their benefit, giving the cost benefit analysis catered to them so they can make decisions based on the true tradeoffs for their body between various treatments, procedures, drugs and other interventions. These choices which today's patient rarely understands because it's too complex and seldom explained to them will be available to the interested patient in laymen's terms. The ability in plain English to ask infinite questions about the tradeoffs from automated and well-informed algorithmic systems that use conversational AI to dialog with patients will open up true patient choice and control! And incredible patience in responding to a patient's questions at their level!

Although it is hard to envision the world of v7, it is clear most specialties in medicine will (hopefully) change dramatically. The superb medical detectives in Internal Medicine will have huge amounts of help or maybe systems that largely relegate their need to certain types of special cases. A nurse may be able to handle the vast majority of the cases with assistance from Dr. Algorithm (regulations permitting). Imagine this definition of a world where every human's complete health history is online, and histories can be accessed anonymously for data-mining purposes. You visit a doctor with a set of unusual symptoms. Imagine if, at the touch of a button, an algorithm could scan the massive historical

database, identify thousands of patients with similar symptoms, focus on those who are similar to you, then summarize the diagnoses, treatments, and outcomes of those cases. That's the future of medicine! Though, v7 systems won't even need the data, anesthesiology may be automated by v7 of the Sedasys system from Johnson and Johnson. Though that will require even the v0 system being able to be deployed in the face of extensive resistance from the anesthesiology community. As of 2016, J&J has halted sales of its device due to poor sales due to this (mostly unjustified in my view) resistance.¹²⁵ Image processing systems are getting very good at sophisticated identification and classification tasks similar to roughly 80% of the work done in pathology, diagnostic radiology, and dermatology. Oncology could potentially be managed on a daily or weekly basis by fine grained measurement, feedback, monitoring and control recommendations by v7 of oncology systems doing fingerpick liquid biopsies to gauge what the tumor is doing and how your body and its immune system is responding. Robotics may play a larger and larger part in surgery with even micro-robots becoming possible at some point though procedural medicine may or may not progress at the same rate as digital data driven areas.¹²⁶ It is interesting to speculate on what specialties might not be changed by technologies we can imagine today. And there will be technologies that sit beyond our (at least my) ability to imagine today. As Dr. Ausiello has speculated, even the names and definitions of many diseases will change based on more precise characterization.¹²⁷

We discussed the somewhat questionable annual health checkup earlier and the fact that the data is 'multiscale' and comes from a variety of sources (genomic, phenotypic, wearable, lab, biomarkers, microbiome ...) which are seldom considered today. In v7 or potentially well before then, for serious diagnosis or the annual health check ups, data input should include physiome (sensors for individual physiology), exposome (environmental sensors), epigenome, microbiome, metabolome,

125 <http://www.wsj.com/articles/j-j-to-stop-selling-automated-sedation-system-sedasys-1457989723>

126 Forgione, Antonello. "In vivo microrobots for natural orifice transluminal surgery. Current status and future perspectives." *Surgical Oncology*, 2009 (<http://www.unicauca.edu.co/ai/Investigacion/In%20vivo%20microrobots.pdf>)

127 Ausiello, Dennis. <http://radar.oreilly.com/2013/09/ausiello-interview.html>

proteome, transcriptome full body imaging as well as various social graphs. This entered into an increasingly more comprehensive database and query system described in v7 above will at some point in the future be the definitive way to diagnose and treat disease.

As part of the v7 phase, we'll likely also find a host of correlations in healthcare. Though which variables will be exactly important or even relevant is highly speculative, it seems very likely we will discover currently unknown patterns and maybe even have systems biology models of patients to explain particular outcomes and comorbidities. Most likely, such discoveries will initially be data science-based, followed by biology models and hypothesis to help explain what we observe. In my view understanding data science based correlations are important and can happen rapidly but finding science based causality is the ultimate goal and will take substantially more time! Data science insights will often lead to serious biological research topics in order to determine causality and true understanding. It is possible, even likely, that observations and insights from data science will be the trigger for a lot of biology based research which will then enhance understanding of fundamental mechanisms and causality. We are already seeing that happening with the rise of CRISPR-Cas9 and the thousands of researchers that are using the mechanism for genetic experimentation. Feng Zhang, one of the preeminent researchers of CRISPR, recently used data mining to find a new enzyme that in certain instances, may actually perform better than Cas9 for splicing DNA ¹²⁸. Suspect we will see the number of opportunities for scientific research multiply many fold as we apply data science to these problems.

Hopefully, systems biology models will exist for many subsystems from the heart to the endocrine system, successors to Lee Hood's P4 medicine initiative or Google's baseline study, and will help us formulate new basic science hypotheses and disease models. ¹²⁹ Likely, progress will be gated by the rate at which we accumulate data relevant to the various subsystems. Progress on this front

will likely be slower than on the data science-driven insights side, as biology is a fundamentally harder and more complex problem, and maybe more important. Observing the data behavior of these systems will become easier after we develop sensors and measurement techniques for wide use in each subsystem, though the rate of penetration of these measurements is hard to predict. These sensors and the data they provide will look radically different from the data that is currently being pulled out from the EMR or from claims data (what is based on 'current medical practice'). This information will be continuous, dynamic and substantially larger in quantity and eventually even in quality.

As all these new technologies and ways of doing things build and integrate, the billions and billions of data points we accumulate will quantify much of what we currently understand only qualitatively. This new modern healthcare system will show much more reliable causal and correlational information via full longitudinal medical records of hundreds of millions of people over years, likely resulting in entirely new medical insights. Maybe we will not yet understand the full science or have a full systems model of human biology for some time, but many aspects will be modeled and at least characterized partially by the mathematical relationships and correlates we will have discovered, much like what has happened in other areas of technology. The math belying the complex network of interacting metabolic pathway would be a dream but is likely more than two decades away. But in the end complex systems theory or AI systems will likely hold the answer, aided by the micro mechanisms of biology, chemistry, and physics.

In the beginning, these isolated innovations will seem marginal in their impact on established healthcare systems like the US system, but when there are enough of these innovations, they will integrate with each other and start to feel like revolution. The nature of innovation is such that change begins slowly with incremental improvements leading to major shifts over time. The systems of a few decades from now, which is

¹²⁸ <http://www.nature.com/news/alternative-crispr-system-could-improve-genome-editing-1.18432>

¹²⁹ <http://www.wsj.com/articles/google-to-collect-data-to-define-healthy-human-1406246214>

the time frame I am talking about, will overcome many of the short-term deficiencies of today's technologies. By analogy, using today's technology would be like carrying the multi-pound mobile phones from 1986, which were floor mounted cell phones with big handsets and heavy cords

in our pockets rather than iPhones. To imagine the leading edge of healthcare practice in v7 one would have to extrapolate using the following imagery of cell phone evolution:

It is hard to imagine this evolution...



v0



v1



v2



v3



v4



v5



v6



v7



v8+

It is commonly understood in technology innovation that humans generally overestimate changes in the short run and underestimate them in the long run. This is because the nature of the exponential curve is not intuitive for humans. I suspect change in medicine will follow this rule. I want to note again that many approaches (robotic surgery, biological research, chemistry, new device inventions, and others) will contribute to advances in medicine. Just because the focus here has

been on the impact of digital technologies and data science it is not meant to imply that other sciences will not contribute meaningfully. The promise of digital health and its rate of innovation coupled with traditional science innovation should allow for a much more rapid, transformative, and disruptive future. What will be unusual is how large and unexpected the digital health technology change is likely to be.

REPLACING 80-PERCENT OF WHAT DOCTORS DO?

This move towards sensor-based, data-driven healthcare does not mean that we will get rid of human interaction in medicine. In fact, human interaction is a powerful tool, and we should try to better understand the ways human beings positively impact healthcare in order to amplify them even further. For instance, so-called irrational effects like those from placebos and medical ritual prove that human disease is more complicated than just drug-response science. Ultimately, rushed and overloaded MDs' tasks could be done by medical professionals that may not need 10 years of medical school training. Or, when they do have the training, they will be much better caregivers than diagnosticians.... "Even if you have the perfect computer that can tell you what to do, you couldn't expose a severely ill patient to the face of the computer." As Atul Gawande says, "Maybe machines can decide, but only doctors can heal".¹³⁰ Some have argued that in most cases the human body heals itself and doctors are only there to triage unusual cases or address symptoms in the majority of the cases. Also not to be underestimated is the placebo effect that medical care and a doctor or drug can provide. A good system will leverage all these effects and the role of the human and even their training requirements may change as a result.

This new state of medicine will lead to lots of changes with how doctors practice. With exponentially increased data and sensors, insights from that information, and consumers with much more personalized and actionable health information at the tip of their fingers, the majority of the traditional role of the doctor goes away. That's not to say 80% of doctors will go away – it's that systems and consumers (directly) will handle 80% of their current work, so the role of the doctor

will adapt dramatically to hopefully add value in other ways. They will be more focused on the patient (and have to worry less about diagnosis), will be more well versed in complications and comorbidities, and will run hospitals in a different way that will be beneficial from both a cost and care perspective. Care will be more precise and with less judgment. The personalized medicine will include everything from the patient's personal information (physiological, genetic, metabolic, imagine symptoms, etc...) as well as the highest quality hypothesis from the best knowledge graph of medicine we have. The best doctors will learn how to work closely with the v7 systems to provide substantially better care than either the doctors or the technology could provide individually. An expertise will migrate towards empowering health care professionals with far less training to provide higher quality care than is available with even ten years of training today.

The burn unit, surgery unit, or the emergency department of a hospital will have humans but their training requirements may be very different than what is needed today. Medical education will certainly need to change dramatically in yet unpredictable ways and medical schools will need to be reinvented. Even the kinds of candidates selected for medical training will change. The focus in selecting candidates for medical school is more likely to be on EQ than IQ as doctors focus more and more on the human elements of care. We will focus on mirror neurons over grey matter for most practicing (as opposed to research) roles. This also will help with the global shortage of doctors that we currently expect. Beyond diagnosis and treatment, there are many things doctors do that won't be replaced. The point is not that every function of doctors will be

¹³⁰ Gawande, Atul. "No Mistake." *The New Yorker*, 1998. (http://www.newyorker.com/archive/1998/03/30/1998_03_30_074_TNY_LIBRY_000015236)

replaced, but that the assistance of technology will allow for majority to be done in less time, with greater efficacy and skill, but with far less needed traditional medical education. Though surgery will need humans we won't need the super surgeons as robots will do most of the procedure just as auto-pilots do most of the flying today. But we still need pilots. Already driverless cars' self navigating technology is being applied (or attempted – many early attempts fail) to navigate the narrow vessels in human lungs, something humans cannot safely do.

The roles of the individual stakeholders in medicine today will dramatically and differently change as a result of these technological shifts. We are already seeing the start of this change in the US as we get a better understanding of how to care for patients in the long-term. In particular, chronic care will likely become a data driven care protocol, with primary care physicians (PCP's) and nurse practitioners taking an increasingly larger role in managing a patient's health (with patient still being in charge). No longer will the PCP be focused on just diagnosis before referral to a specialist. Upskilling of both the nurse practitioner and the primary care physician will allow them to do what only specialists do today. Instead, they will be in the loop and the patient's primary point of contact to holistically manage their health. The technology for helping do that (and making sense of all the incoming data streams) will become more fundamental to care as well.

Other forms of care outside of chronic care are still important, but their role in the healthcare system and how they will be impacted by technology differs. Procedural or interventional care such as surgeons still requires skill and precision rather than primarily knowledge that can more readily be learned by machine learning systems. Over time, advances in robotics (perhaps coupled with machine learning) may cover more and more of these procedures, but the rate of change will look different from chronic care. On chronic care, ai hope medicine will find cures. But if we don't, managing chronic care will become much better with very strong, frequent feedback and guidance for the patient. .

Similarly, acute care or inpatient care in the hospital will move towards more intelligent, but

automated, monitoring systems. And the skills necessary to help take care of those patients could look more like intelligent nursing – working with those monitoring systems to properly triage and then administer care. This will go a long way in preventing unnecessary complications, infections, or deaths within the hospital – especially as the systems used are able to better deal with complexity (rather than for example setting off too many alarms for a human to properly handle).

In general, the ability for technology to “upskill” workers who are not currently physicians, internists, or surgeons will determine how large an impact this healthcare transformation will have. Upskilling is defined as enabling someone to successfully complete and handle a set of tasks traditionally done by someone with more knowledge than them. In the case of nurse practitioners, in the v0-v7 examples above, we see many examples of how advances in machine learning and automated systems will enable nurse to both help manage patient's chronic care over time (without getting overwhelmed) as well as rely on automated monitoring systems within the hospital to take care of more patients effectively. This upskilling isn't only true of nurses – but we expect PCP's to benefit from this as well – leveraging technology to do the basic tasks of a number of specialists without always needing to refer to them Upskill in general move all healthcare workers towards the top of the skills pyramid.

Consumers and families also benefit from this by having better resources to take control of their health within their own homes. This trend will be an important one as these technological innovations come to market.

Some roles' importance will decline while others will become more prominent. This change in medicine is an empowering one, which should lead to more primary care physicians and nurse practitioners to be able to do a range of medical work previously only handled by specialists. This gives the PCP's and nurses, the ones closest to care, the power and knowledge of these specialists and will be combined with systems that will help them integrate across specialties and comorbidities. This will ensure that the patient's preferences will be best taken into account across

the entire healthcare system.

In each of the many different areas of medicine, we will see new capabilities emerge every 2, 3 or 4 years (a new generation or version 1, 2, 3, 4, etc.) in the form of intelligent software or combined software and hardware systems. v7 will show in a couple of decades in most areas and we may be off by decade but the trends will be obvious. The exact dates are less important and less predictable than the inexorable trend. And though many doctors will keep practicing the way they do today, the leading doctors will have adopted these new methods. Guess which doctor I will seek out if I am a believer in the scientific method?

All the data discussed earlier, millions of data points, will reveal complex correlations between human diseases and wellness states, far beyond the capability of human comprehension and far more quantitative and accurate measurement than possible today. Systems will not just spit out data but also will provide insights, actionable recommendations, tradeoffs and much more at a level of personalization and integrative and holistic comprehensiveness that would be impossible and unaffordable today because the cost of computing is trivial compared to the cost of human attention and the limits on typical human comprehension. The decisions as to what to do may be assisted by a doctor, or likely a humanely trained medical professional (not necessarily a doctor), but also could be made by a patient who, under this scenario, could be as well informed as the doctor. The key to the patient becoming the CEO and decision maker of his or her own health will be such systems that empower the patient against the healthcare system and its many interests. A patient will know when surgery is excessive and unlikely to lead to outcome benefits given their personal preferences. For patients with low health literacy, systems could conduct a dialog at their level of understanding or in their native language or alternatively assist a doctor or other medical professional to advise a patient. I suspect in this area there will be human assist for patients in making decisions.

Today, physicians spend too much time doing

things computers can do, and we should give medical professionals more time for things that uniquely require human involvement. For instance, making inherently subjective decisions that require empathy or a consideration of ethics or providing patients with “warm fuzzies” like comforting kids in pediatric care and providing a friendly ear for lonely patients or their kin. Some of these functions may not need medical school training at all, since they draw on more empathic skills and could actually be done by non-MDs (jobs for UCS film school graduates who can act humane and may have more mirror neurons?). If the human element of care is important, should we not have the most humane humans provide it? The systems I speak of will increasingly allow less trained medical profession to operate at levels of care that are higher than today's fully qualified median doctors, making up for the shortage of doctors and poor care from their overloaded schedules.

The same information will be available to patients allowing them to be empowered to more intelligently engage in and make decisions about their own health. As stated earlier data indicates that when empowered with full information, patients tend to choose less aggressive therapies than the medical system might choose for them.¹³¹ Should we still let doctors prescribe care when consumer preferences may be different or should we empower consumers to be intelligent decision makers? The debate will be very similar to that around driverless cars. Should a system that is better than humans but occasionally makes mistakes, albeit far fewer mistakes than humans be allowed to “practice”? In the driverless car world many knowledgeable people believe a time will come when humans will not be allowed to drive! Will the same issue arise in medicine with doctors not allowed to practice?

It's certainly true that machine learning and data science will not be a panacea. And this is where the role of the human / doctor might become vital and empowering. As mentioned earlier, there is no such thing as the perfect objective function for “becoming healthy and living forever,” at least in the next two decades. So it's might require good human input and understanding to guide the

¹³¹ Smith, Mark, et al., eds. “Best care at lower cost: the path to continuously learning health care in America.” National Academies Press, 2013. (<http://www.iom.edu/Reports/2012/Best-Care-at-Lower-Cost-The-Path-to-Continuously-Learning-Health-Care-in-America.aspx>).

patient or consumer (since the focus will be on wellness) to the outcome they want to achieve. These won't necessarily have to be physicians, but human intervention will be important particularly in the early days of the systems. But then again the possibility exists that no such human assistance is needed and computers will provide the most knowledgeable, sensitive, empathetic assistance.

Human input, apart from just physicians / standard roles in the healthcare ecosystem will influence health more and more as well. As innovations become more consumer-driven, there's a first-class need for great UX and design in order to drive customer adoption. There will be a growing need for humans to help spur this adoption (more so indirectly than directly), because using these new devices systems for improving healthcare will require a strong emotional pull. But remember that Facebook and Twitter have proven that digital systems can be more powerful than humans in driving behavior change and addiction for both good and bad.

We have spent too much time with a broken healthcare system. A shift towards making doctors more holistic and more able to focus on their patients will transform the system. And my technologically optimistic view is that the impetus for that change will be data-driven and consumer-driven systems replacing (and greatly enhancing) the tasks which physicians currently have to do. And what I call Dr. Algorithm will play a starring role in this holistic system.

The advent of Dr. Algorithm and holistic medicine!

One particular example where technology upskilling likely will have a major impact is in clinical decision support. Decision support systems in healthcare are traditionally defined as decision making tools to help clinicians in their workflow, including (but not limited to) computerized alerts, clinical guidelines, order sets, and diagnostic support ¹³². There are a number of systems today that provide decision

support, but the vast majority of them are rules-based and limited in scope. Some examples of these systems for diagnostic support include DXplain and Isabel. DXplain, owned by MGH, is a rule-based knowledge base that has been built up over 25 years. In that time, they emphasize the "240,000 individual data points representing disease / finding relationships." ¹³³ These rules and relationships are used to provide a ranked list of diagnoses when given a set of clinical manifestations. While the premise is useful, it's very hard to have given rules relate to one another, and the system as a whole is fairly rigid, especially as the attributes attributed to each finding or disease have rudimentary ranking between 1 through 5. Similarly, Isabel Healthcare has a differential diagnosis generator that is based on a database that covers diseases and conditions. Even Isabel's "old age" natural language systems, though, still match expert diagnoses 90%+ of the time (but that's still not enough to move fully away from the clinician) ¹³⁴.

Other related systems to help with clinical decision making include event-driven alerts that are often set by simple thresholds, as well as rigid medication adherence mappings to ensure the patient doesn't have an adverse reaction to a newly administered drug. All of these are useful, but only go 5%-10% of the way towards providing comprehensive clinical care and support (therefore requiring doctors and nurses to still do the majority of the work).

These systems at the moment only go a small step of the way towards upskilling nurses and clinicians. True bionic assistance, as described before, will happen as these algorithms become much more advanced and dynamic, and are able to integrate much wider set of data (rather than just disease and symptom understanding). At the point, rather than being clinical decision support tools, these systems will function as "Dr. Algorithm" – with algorithms automatically handling the majority of cases clinicians see. For example, rather than simple threshold-based alerting at the hospital, we will see dynamic alerts

¹³² <https://www.healthit.gov/policy-researchers-implementers/clinical-decision-support-cds>

¹³³ http://dxplain.org/dxplain_faqs.html

¹³⁴ <http://www.isabelhealthcare.com/validation/peer-reviews/accuracy-studies>

and algorithms that will be much more accurate about diagnosing things like sepsis in real-time. And hospitals will move towards much more intelligent order sets, where Dr. Algorithm will be the primary decider in the set of steps / orders and choices presented to patients on a personalized basis, rather than having to rely on physician-by-physician standards that are relied on today. Rather than the knowledge bases which back clinical decision support systems today (and are still by and large limited to experts), the knowledge graphs of the future will be dynamic and will be able to handle vast complexity of data and its reduction to knowledge, insights and actionable choices – such that nurses (or even consumers!) can use those tools and ensure accurate and safe results. These knowledge graphs, which will also be updated with new data continuously from medical and scientific literature along with clinical practice outcomes, get integrated into practice, everyone in the medical system will be upskilled. Dr. Algorithm will exist and thrive across specialties, with different sets of algorithms working on conjunction with one another to come up with accurate diagnoses and care paths for all patients.

The leading edge of medical practitioners will be much more scientific and data-driven in providing patient care once these new tools become available. With the increasing amount of data and research released every year, that's hard to pull off without technology. For example, standard operating procedure involves giving the same drug to millions of people even though we know that each person metabolizes medication at different rates and with different effectiveness. Many of us metabolize aspirin poorly, but the dosage recommendation for all seven billion people on the planet is roughly the same, based on their weight.¹³⁵ Each person should be treated differently, but the average doctor does not provide such personalized care, nor do they have enough time or knowledge to do so. This is a limitation not of doctors but of human beings. Humans cannot till a cornfield as well as tractors either!

And no one actually expects humans to fly a complex airplane as error free as systems except in exceptional circumstances. Airplane autopilots do the vast majority of flying. Even skills that till recently were supposed to require huge human judgment like high-frequency hedge fund stock trading is now substantially done by machines. In fact, across most tasks, research has shown that even relatively simple statistical formulas can outperform human judgment. Gawande cites the late 1980's work of Paul Meehl, David Faust, and Roam Dawes who reviewed more than a hundred studies comparing computers with human judgment in predictions ranging from life expectancy of a patient to survival expectancy of a company. "In virtually all cases, statistical thinking equaled or surpassed human judgment." In the late 80's! Dawes cited the inconsistency and biases of humans, coupled by our inability to accurately consider many factors at once, as leading to the statistical methods outperformance.¹³⁶ It reminds us of the more recent work of Dr. Ioannidis and the Institute of Medicine. Algorithms and data science, refined over decades with much increased computational power, should only increase the computer's gap with human judgment. It is possible that humans and software together will make better decisions at least for a while but eventually the complexity of the systems, with the number of variables and pathways in the body will make human judgment error prone relative to more comprehensive data and knowledge that systems will have.

Complex systems like the body¹³⁷ are not an abstract term but rather one defined by mathematical rigor. We can take all of these data streams and treat them as the inputs to our corpus of medical knowledge. And because we have some (and will continue to accumulate) data on outcomes, we can build on to this knowledge. This will start with simplistic techniques at first, but my belief is over time, techniques more akin to AI complexity research, network theory, and other fields will be used to find the optimal paths forward in enhancing our medical understanding.

¹³⁵ Krasopoulos, G., et al. "Aspirin 'resistance' and risk of cardiovascular morbidity: systematic review and meta-analysis". *BMJ* 2008; 336:195. Retrieved 5 Feb 2014 (<http://www.bmj.com/content/336/7637/195?view=long&pmid=18202034>)

¹³⁶ Dawes, RM, Faust, D, Meehl PE. "Clinical versus actuarial judgment" *Science* (1989) (<http://www.sciencemag.org/content/243/4899/1668.abstract>)

¹³⁷ https://en.wikipedia.org/wiki/Complex_system or <http://www.informatics.indiana.edu/rocha/publications/complex/csm.html>

And applied in specific instances, this won't just include symptoms and other biological data (our genetics, microbiome, etc). These systems will also need to reflect societal needs as well, bringing in cost/benefit analysis into the equation. Many of the decisions we make today and we will make in the future of what to support and not will be driven by societal needs rather than purely scientific ones. But these technological advances will bring science to the table first and foremost when making these decisions.

Healthcare needs to become much more about data-driven deduction, be more precise in diagnostics, prognostication and monitoring, and more consistent across doctors. With more data-driven technology, doctors will have more time to focus on things that computers cannot do yet. Machines can be more comprehensive, more integrative of all the data and much more holistic because of their ability to assimilate a lot more data about all the nuances. Going back to Atul Gawande's article, "the machine, oddly enough, may be holistic medicine's best friend. The professional ethic of making no mistakes has often reduced care to a narrow matter of problem-solving." And our thesis is this narrow problem-solving should be replaced by computers over time. "As expert systems begin to take on more of the technical and cognitive work of medicine, generalist physicians will be in a position to embrace the humanistic dimension of care."¹³⁸

The patient will be empowered with a lot more information and context in order to make decisions in his or her own self-interest instead of being dictated to by the healthcare system with multiple objectives, some of which conflict with the patient's interests. There won't be a magical solution either – because there is not "perfectly healthy" body. The objective function people maximize for will vary by each person and personality, but putting the choice and information with the patient will allow the consumer to make an educated choice (or have a smart default choice) as to what that objective is.

Even the definition of what is disease will evolve and be a personal and/or societal choice. Are low or high thyroid hormones a disease if it is functionally not making a difference or only a minor difference? What if thirty years of drug therapy can increase your risk of vascular calcification and you have high cardiac risk genetics? Which risk does a patient want to take and society want to afford? As a result, much of what physicians do including checkups, testing, diagnostics, prescription, monitoring and behavior modification can be done better with well-designed sensors, active and passive data collection and analytics that drive actionable insights and recommendations without necessarily taking away from the human element of care. There will be a software app for that!

Of course, doctors are supposed to do much more than just measurement. They also are supposed to consume all the data they collect, carefully consider it in the context of the latest medical findings and the patient's history and determine if something is wrong. There is the "eye ball" test that most doctors believe in. Computers also can take on much of the diagnostic and treatment work and even do these functions better than an average (median) doctor, all the while considering more options and making fewer errors. A computer can simultaneously consider hundreds of variable while memory studies show humans can keep seven chunks of information (+/- 2) in their head when making decisions (short term or working memory).¹³⁹ An computers hopefully be able to keep extensive knowledge graphs of the millions of research articles published and be apply to apply them precisely and resolve inconsistencies in knowledge or use medical records data to resolve best paths for multi-morbidity patients with conflicting treatments. In fact, over time, it's likely these computers will operate and feel like human intuition in its abilities to make decisions on the fly. AlphaGo, DeepMind's AI which beat the world's leading Go player, represents the cutting edge in machine learning

¹³⁸ Gawande, Atul. "No Mistake." *The New Yorker*, 1998. (http://www.newyorker.com/archive/1998/03/30/1998_03_30_074_TNY_LIBRY_000015236)

¹³⁹ <https://www.quora.com/How-much-short-term-memory-does-the-brain-have>

today. And AlphaGo has been referred to as having an intuitive feel in the way it made decisions about gameplay.¹⁴⁰ Over time, these “intuitive like” approaches (that in reality are just more and more effective processing of complex data) will look like having one of the world’s best doctors right by your side at all times. The systems would be so good that we would feel guilty (and perhaps in the future even be legally liable) if we made decisions instead of letting this AI system make them.

On the other end, what if you were a heart patient? It’s a simple fact that most doctors could not possibly read and digest all of the latest 5,000 research articles on heart disease or the five hundred ECG’s you might send them or objectively aggregate the results of the best outcomes for that type of patient selected from millions of records of those symptoms and disease. They would find it prohibitive to deduce that your genetics increases the risk of deep vein thrombosis is 500% or that your particular stroke risk can be decreased by certain actions or even foods? They cannot understand your lifestyle fully or understand decades of your longitudinal medical history while they are rushed to make a critical decision. Human beings can understand and integrate simple numbers with qualitative factors but are poor at comprehending hidden patterns in very large and complex data sets. In fact, most of the average doctor’s medical knowledge is from medical school, and cognitive limitations prevent them from remembering the 10,000 to 15,000 diseases humans can get or the 15,000 procedures and treatments at their disposal or the hundreds of relevant medical articles published for the current patient situation they are considering. They can’t be expected to integrate hundreds of biomarkers and historical data on you and your lifestyle and preferences. To practice under this limitation is unfair to the patient if a better option becomes available.

Computers are much better than people at organizing and recalling information. They have larger memories that are less corruptible and don’t

get tired. Moreover, computers unlike people (and doctors) are not subject to cognitive and other biases. The Nobel Prize winning economist, Daniel Kahneman explains in his book, “Thinking Fast and Slow” (a must read for anyone wanting to understand human decision making) that human errors arise from simple rules, otherwise known as heuristics, that people use to form judgments and make decisions. As doctors are overloaded in their work and attempt to use their time efficiently, they use constructs developed over time rather than objective inputs to practice evidence-based medicine. One example is the drug, Avastin, which was approved by the FDA to treat metastatic breast cancer. Adoption was rapid and driven by great marketing, but it was later determined that the drug was not an effective treatment. Ignoring the clinical evidence, oncologists were convinced that Avastin was effective despite scientific evidence to the contrary because of established beliefs and despite the FDA findings to the contrary.¹⁴¹ The recent launch of the Cancer Moonshot project is in recognition of the value and importance of the need for rigorous genomics, proteomic and phenotypic data to deal with the complexity of disease management in current oncology practice. Randomized clinical trials of up to 20,000 cancer patients at all stages of the disease will be completed by 2020, for 20 tumor types and utilizing 60 novel and approved agents. But the data that generates may be even more valuable with proper AI analysis than currently envisioned.

We are seeing examples now of how today’s sophisticated AI systems are able to look back at previous research and find results and understanding we hadn’t thought possible. After looking at data written off as useless 20 years ago on spinal cord injury, one system was able to use topological data analysis to determine high blood pressure’s detrimental effects on long term recovery from injury.¹⁴² This data was all there in the past, but without retrospective data mining techniques, would never have been discovered. As AI gets more and more sophisticated over time,

140 <http://www.wired.com/2016/05/google-alpha-go-ai/>

141 <http://ari.aynrand.org/issues/government-and-business/regulations/The-Avastin-Travesty>

142 “The Latest Medical Breakthrough In Spinal Cord Injuries Was Made by a Computer Program” <http://www.fastcoexist.com/3052282/the-latest-medical-breakthrough-in-spinal-cord-injuries-was-made-by-a-computer-program>

these data sets will continue to get more valuable in ways we can't yet imagine or understand.

Similarly, let's look back at the breast cancer pathology example discussed earlier. Since 1928, the way breast cancer characteristics are evaluated and categorized remains largely unchanged. It is done by hand, under a microscope. Pathologists examine the tumors visually and score them according to a scale that was first developed eight decades ago. Yet a machine-learning group at Stanford¹⁴³ states their computational model yielded results that were a statistically significant improvement over human-based evaluation. Over time, these tasks currently handled by pathologists should move to computational systems that are able to process thousands of features and related biomarkers and gene mutations more accurately. The pathologists and radiologists roles are likely to be substantially replaced though certain specialty functions may remain.

Moreover, the computers discovered that the characteristics of the cancer cells and the surrounding cells, known as the stroma, were both important in predicting patient survival. Could a human pathologist consider 6,642 cellular factors, something that would be routine for a v7 system to do? Even the smartest human pathology researchers did not come up with the extra features and yet very nascent technology systems did. Can humans keep up with these new tools or computational power or should they just learn to leverage them or let them take over certain functions? How much more would such systems discover if we had vast amounts of data and could humans learn to apply all the new learning accurately or precisely?

The authors of this study state that "Through machine learning, we are coming to think of cancer more holistically, as a complex system rather than as a bunch of bad cells in a tumor.... The computers are pointing us to what is significant, not the other way around." said Matt van de Rijn MD, PhD, a professor of pathology and co-author of the study. "If we can teach computers to look at a tumor tissue sample and predict

survival, why not train them to predict from the same sample which courses of treatment or drugs a given patient might respond to best? Or even to look at samples of non-malignant cells to predict whether these benign tissues will turn cancerous," asks Koller. "This is personalized medicine." But Van de Rijn does not see computers replacing pathologists, possibly because he came from an MD background. It is hard to predict where the limits of such systems are but as a betting man I would bet on the systems over humans to provide the best pathology care in a few decades. To be fair there are many patients that will prefer the human pathologist just as today people prefer homeopathic or natural medicines as personal preferences. We will at least be able to offer both choices and will let the patient decide.

Computers can remember more complex information more quickly and completely resulting in far fewer mistakes and biases than a hot shot MD from Harvard. Contrary to popular opinion, computers also are better at integrating and balancing considerations of patient symptoms, history, demeanor, environmental factors and population management guidelines than the average physician. Besides, who wants to be treated by an average or below-average physician? Remember that by definition, 50- percent of MDs are below average (median) if the system/ automated alternative is reliably and repeatably better! Not only that, but computers also have much lower error rates. Already, thousands of data points can be collected on a cell phone everyday for months, if not years. Shouldn't we take advantage of that when it comes to our health even if humans may not be great at integrating such data into their thinking?

The next generation of medicine will arrive at scientific and data-driven diagnostic and treatment conclusions based on probabilities and real testing of what's actually going on in a patient's body. This kind of approach will be used in the diagnosis and monitoring of many diseases and will be much more personalized to the patient and automated in its operation than physicians currently can provide. Systems will utilize more complex models of interactions within the human

143 Beck, A., et al. "Systematic Analysis of Breast Cancer Morphology Uncovers Stromal Features Associated with Survival." *Sci Transl Med* 9 November 2011: Vol. 3, Issue 108, p. 108ra113 *Sci. Transl. Med.* DOI: 10.1126/scitranslmed.3002564.

body and more sensor data than a human doctor could comprehend in order to suggest diagnosis and the most appropriate therapy. A doctor looking at your immune system may be able to look at a few immune markers, but a system with hundreds or thousands of biomarkers from millions of patients could detect complex interaction patterns between thousands of biomarkers and their relationship to patient data. And these markers will apply to all cells and molecules in the body – from your DNA, RNA, small molecules, proteins, gut bacteria, and more – and this data will be obtained relatively simply, with either a few drops of blood or noninvasively (e.g. with saliva) taken at home. Will humans be able to comprehend a thousand microbiome species and other data points interacting in complex ways? Among the 15,000 diseases and 15,000 therapies and the thousands of metabolic pathways in the human body that, though worth understanding through biological science research and network research of interacting metabolic pathways and machine learning systems, will reduce and personalize the net recommendations in an individual patient's context, tuned to the patient's explicit preferences. These systems also could monitor continuously the effect of many drugs. Thousands of baseline and disease multi-omic (genomic, metabolomics, microbiomic and other) data points, more integrative history and demeanor will go into each diagnosis, all hard for today's human doctor to process accurately. Data science will be key to this and in turn, it will reduce costs, reduce physician workloads and improve patient care. In my opinion, over the next 15-20 years, data science might do more to improve healthcare than all of the biological sciences combined.

In the future, doctors assisted or substituted by systems will be able to tailor their explanations to the health literacy level of patients using common-language terms and adapting to each individual's sophistication level using computerized dialog managers. These computerized managers will be patient, unlike your typical "doctor in a hurry" with the common and unfortunate case overload. This matters; according to the Institute of Medicine, nearly 100 million U.S. adults have limited health literacy skills. Improving people's

understanding of their health will positively affect their health outcomes! Machines can be infinitely patient with repetition, can record and replay information at later times, even quiz patients to check understanding, and if they have monitoring of patient physiological response (voice, heart rate, etc.) modern techniques in affective computing may enable them to tune their responses to the patient's current emotional state. And don't be confused by early clumsy attempts at dialog by computers! At Boston University School of Medicine's pilot study of the automated "virtual nurse system," 74-percent of hospital patients said they preferred receiving their discharge instructions from the virtual nurse, rather than their human doctors or nurses.¹⁴⁴ Imagine v7 systems that manage these conversations much better and may even be multi-lingual with the ability to adjust dynamically to the level of healthcare understanding of each patient. Or maybe a very different type of medical professional than today's doctor, assisted by machine intelligence, can play a role here?

An example from chess is worth noting. In the book *Race Against the Machine*,¹⁴⁵ the authors cover an intriguing fact. While chess computers today can routinely beat humans at playing the very complex game of chess, even a mediocre chess program assisted by a human chess player can beat the best chess computer. Will medical computers follow the same path and will the best combination be that of human doctor and machine doctor working together? Chess has certain mathematical and computing characteristics that make this combination of human plus computer doing better than either alone a high probability. It is not clear what characteristics the domain of medicine will have. Today's AlphaGo computer from Google is better than any human Go player but given its unique "non-brute force computing approach", but can it be even better than the machine alone with human assistance? It is hard to predict. But clearly playing against AlphaGo has made Lee Sedol better. Previously he was winning 75% of his matches against humans but since playing and losing to AlphaGo, and studying his techniques, he has increased his win percentage against other

144 "The Re-Engineered Hospital Discharge Program to Decrease Rehospitalization." Boston University School of Medicine. Retrieved 5 Feb 2014 (<http://www.bu.edu/fammed/projectred/publications/RED%20Fact%20Sheet%202-7-09%20v2.pdf>).

humans to 90%.

Doctors will struggle to keep up, but they will increasingly rely on these tools to make decisions. Over time, they will increase their reliance on technology for triage, diagnosis and decision-making. Eventually, we will need fewer doctors, or we will make up for the anticipated shortage of doctors in both the developed and developing world. Regardless, every patient will receive better and more accessible care than what is available today. Diagnosis and treatment planning will be done by a computer used in concert with the empathetic support from medical personnel, selected and trained more for their caring personalities than for their diagnostic abilities. No brilliant diagnostician with bad manners, like the TV character Dr. House, will be needed for direct patient contact. Instead, Dr. House's best contribution will be to serve as the teacher for the new "Dr. Algorithm", which we will use to provide diagnoses. The most humane humans like nurse practitioners or other medical professionals will provide the care. And each patient may receive more human contact and care than they do today. But the most brilliant diagnosis and prescription will likely be done by AI systems with training help (during their development) from the best and most brilliant clinicians in each area of specialty. As Yogi Berra said: It's difficult to make predictions, especially about the future!

We have spent too much time with a broken healthcare system. A shift towards Dr. Algorithm as doctor will make it more holistic and more able to focus on their patients will transform the system. And my technologically optimistic view is that this cannot happen cost effectively without Dr. Algorithm.

Change by Specialty: The easy & the hard speculations!

What does it take to get to these v7 systems in healthcare?

It is worth emphasizing that each field within medicine will progress and adopt technology at differing rates. Some specialties may progress relatively rapidly (e.g. sleep medicine), some much slower (e.g. slowed down by incumbents or hard technical problems), some will get much more

precise (e.g. endocrinology), and for some it's even harder than usual to speculate on the timing of impact (e.g. robotic and VR surgery).

Rather than make strict claims about the buckets each specialty fits into, it is valuable to understand the mechanisms and axes by which technological progress will gain steam in healthcare. It's speculation to pretend we know exactly how the changes will happen, but these axes can be used to both evaluate and help spur technological change across disciplines.

Medicine to date has been limited by factors human doctors can use or understand; as a result, our willingness to collect complex data in the practice of medicine has been limited because we had no good uses for this data since humans could not use it (how does a human evaluate fractal patterns for skin cancer or 1000 bacterial species in the microbiome for relationships to depression or Alzheimer's?). We are seeing a number of approaches to change that, including

U Finding excuses to gather more data

- Getting medical records directly from consumers
- Collecting sensor data from other devices e.g. mobile phone
- Collect data for one simple purpose but integrate it into a larger set for more complex analysis, a common pattern in entrepreneurial companies

U Understanding data in different ways

- Using algorithmic systems and complex math to extract patterns from data e.g. subtypes of diabetes or probability of a cardiac condition in the next few days
- Leveraging new sources of data that humans today couldn't understand or properly integrate into practice e.g. microbiome markers
- Use variety of approaches to understand complexity – statistical machine learning, deep learning, topological data analysis, and even more algorithmic techniques to

be discovered

U Building (and growing) a knowledge graph

- Use research publications, textbooks and medical records to build our current state of evidence and knowledge
- Continue to build o top of these systems to allow for greater intelligence and precision over time

U Discovering new medicine

- Computational drug discovery will lead to many novel insights and faster paths to drug targets
- Novel therapeutics that leverage these different systems (e.g. the microbiome) will become more common

U Integrating all variables into one systematic approach

- Humans can only use a few factors in our decision making, but software systems can use thousands or more to make decisions
- Today humans would have a hard time integrating genomic data (predispositions), epigenetic or gene expression data, hundreds of metabolic variables, microbiome data, medical record data from the EMR, as well as personal and geographic history
- would expect v7 versions of the current Google Brain or DeepMind systems, adapted to medicine, will be able to better navigate across all these variables than humans can to make or at least heavily inform our decisions about care

Each of these axes will overlap; in conjunction, we will use these methodologies (and new ones we haven't yet discovered) to progress medicine forward and make it a mostly automated, highly

levered, scientific practice that gets more accurate over time. The below examples all illustrate how these different axes interplay to move towards a 20% doctor future.

1. Allergy and Immunology

80% currently:

A allergist / immunologist currently spends the majority of their time focused on diagnosing and alleviating allergic or adverse conditions of the human body to external substances ¹⁴⁵. It's estimated up to one fifth of the Western Hemisphere suffers from some type of allergy, yet at the moment, we still do not have a definitive idea of why allergies develop (although we do understand the immune reaction afterward) ¹⁴⁶.

Transitions:

We are starting to see early efforts trying to get new understanding of our immune system. Many of these efforts will attempt to pool together data and generate new knowledge from those databases, such as Harvard's machine learning data repository for immunology:

<http://bio.dfci.harvard.edu/DFRMLI/>.

We are also starting to see new research from domains previously not studied extensively showing actual prevention of allergies, in particular with the microbiome in mice:

<http://news.sciencemag.org/biology/2014/08/gut-microbe-stops-food-allergies>

These efforts, coupled with increased levels of funding to tackle the root cause and prevention of allergies, should lead the path to better understanding mechanisms and potentially finding breakthroughs in immunology research ¹⁴⁷. These approaches are likely to be unconventional, coinciding with the unconventional set of people leading these funding efforts e.g.

20% future:

It will take some of these scientific breakthroughs to come to fruition, but when that happens, immunology will fully change to being about

¹⁴⁵ <http://www.aaaai.org/professional-education-and-training/careers-in-a-i.aspx>

¹⁴⁶ <http://www.webmd.com/allergies/understanding-allergies-basics>

¹⁴⁷ <http://www.bbc.com/future/story/20150409-why-do-we-have-allergies>

understanding individual's specific conditions, and being able to cure or at least severely limit allergic reactions. Although it is likely that some aspect of the allergic reaction is beneficial to the human body, a better understanding of the underlying mechanisms will enable us to fine tune drugs or treatments to limit the harmful side effects. An allergist, rather than dealing with allergies after the fact and on a somewhat broad way (as is done today), will be data-driven and have more variation and granularity in how they are able to treat patients. And even further on, an allergist's job will be about fully preventing allergies. In fact, over time, the primary care physician will handle much of this prevention, and they will leverage automated knowledge (from software) to handle people's allergies just as they do with their variety of choices today.

Much of the insight that will lead to this future will come from the axes defined above:

- 1) data science or biological research enhancing our biological understanding of immune pathways
- 2) new medicine i.e. data science yielding new, unexpected insights about allergies just through researchers data mining health records
- 3) knowledge graph i.e. having vast repository of knowledge about allergic reactions and immunology.

In the future it is likely that the practicing allergist's knowledge will be more complete, more current and probably more in depth and current in a software system that makes it all available to a primary care provider. Continuous environmental and patients metabolic and microbiome data may let software be predictive of potential events and allow even a nurse to respond quickly to intervention needs because of the knowledge available in such systems. The software may be able to access detailed genomic profiles, microbiome profiles and other personal information in recommending courses of action.

2. Anesthesiology

80% currently:

Current practice for anesthesiologists centers around the administration of anesthetics, and the monitoring of vital signs during surgery. Then post-surgery, they are responsible for ensuring you are safely moved from anesthesia ¹⁴⁸. Due to previous advances in technology, medications, and knowledge, risk in routine anesthesiology is considered quite low currently (mortality rates less than 1:100,000)

Technology does play a vital role in current day anesthesia, primarily from the side of decision support tools that integrate with the anesthesia information management systems (AIMS) to give quasi real-time guidance to the care team throughout and after the surgical procedure ¹⁴⁹. Yet, still, these are rudimentary efforts that don't take place at the center of the anesthesiologists workflow for the most part. And the monitoring tools, even post-operation, are still quite cumbersome and are frequently standalone machines.

Transitions:

Given how rudimentary the current systems are, how standard general anesthesiology practice has become, and how expensive it is ... we will start to see automation and machine learning taking on more of the anesthesiologist's role. We are already seeing this with a major player, Johnson & Johnson, putting Sedasys out to start replacing anesthesiologists. And as would be expected, anesthesiologists, recognizing the threat of this automated system, pushed back, to the point where J&J has stopped selling the device.

<http://blogs.wsj.com/corporate-intelligence/2013/09/26/jjs-sedasys-puts-challenge-to-anesthesiologists/>

<http://www.wsj.com/articles/j-j-to-stop-selling-automated-sedation-system-sedasys-1457989723>

In this case, it seems over time more than likely that Sedasys-like systems will be able to cover 80%+ of anesthesiology cases. While Sedasys

148 <https://www.asahq.org/WhenSecondsCount/physiciananesthesiologistrole.aspx>

149 http://www.anesthesiologynews.com/ViewArticle.aspx?d=Educational+Reviews&d_id=161&i=June+2015&i_id=1194&a_id=32822

right now is targeted to somewhat simplistic screenings, researchers at University of British Columbia are showing they could have automated systems administer anesthesia for complicated brain and heart surgeries, even for children ¹⁵⁰.

Even if these systems get their start “bionically assisting” anesthesiologists, providing the default settings to administer, over time, given the range of cases we are already seeing being handled by automated systems in this v0 stage, these systems will be more than capable of performing the administration and monitoring of drugs during and after surgery.

20% future:

Because these systems are already in place and are showing that a machine can, in some cases, perform better (and certainly cheaper), than an anesthesiologist ... it is not that farfetched to imagine in the future, the majority of anesthetic functions will be handled by machine-learning driven automated systems. While there will be fits and starts, and it's more likely that simple procedures will be handled first (they will yield the most cost savings and be less risky), over time, we won't need to worry about a person administering our drugs during surgery; the machine will handle it for us. And then after surgery, through better non-invasive monitoring devices and our smartphones, we'll have a more pleasant experience as we recover. However, progress in technology can and will continue to be hindered by incumbents fighting for their current incentives – such as today's anesthesiologists who aggressively fight adoption of these automated devices (like Sedasys)

The key contribution of the anesthesiologist is the knowledge repository in their head based on education and experience both in recommending treatment and responding to unexpected events. All of this key contribution, in greater depth and currency, will reside in software systems capable of monitoring many more variables and find more complex interaction patterns than a human could (think of the complex patterns and future pathways in a game that AlphaGo was able to intuit applied to the patients' body's reaction to

anesthesia and other procedures/medications).

3. Cardiology

80% currently:

Cardiologists manage and focus on a lot of aspects related to heart health and disease. Their primary functions include: diagnosis of cardiac conditions (with data generally coming from a variety of tests taken during a visit), managing of care, and performing some non-surgical procedures (including stents, catheterization, pacemaker setup)

Transitions:

The best example (v0) of showcasing the transition from 80% cardiologist to 20% cardiologist is with the mobile ECG. A few years ago, these ECG's could only do one lead, and were somewhat simply attached as a case to your mobile device.

Now, the vision of what a mobile ECG can do has dramatically expanded and will soon offer six lead ECG's in a mobile phone form factor or single lead ECG's that are in your Apple watch strap, always available. But more importantly, for the first time, due to machine learning over millions of ECG's collected, there are algorithms approved by the FDA which can automatically detect atrial fibrillation, as well as what a “normal” ECG looks like. Their accuracy rates are better than the average cardiologist and patients no longer need to wonder at home or at party if they are having a cardiac AFIB episode.

New data collection has also shown even more surprising connections – where continuous wearable devices that are on the market today that only measure heart rate and heart rate variability can yield indication of when your heart rhythm may be abnormal (and therefore, signal when you should take an ECG). This type of science would have been unheard of just a few years ago, and has the impact to change the way we practice cardiology. It is far more responsive than either calling an ambulance or waiting till the weekend is over to schedule an appointment! And algorithms well beyond AFIB detection will show

¹⁵⁰ https://www.washingtonpost.com/business/economy/new-machine-could-one-day-replace-anesthesiologists/2015/05/11/92e8a42c-f424-11e4-b2f3-af5479e6bbdd_story.html

up to diagnose the myriad of cardiac conditions. Most characteristics of an ECG will be measured accurately in this mobile, patient controlled form factor with the occasional (<20%) of the cases requiring more extensive and expensive ECG's.

20% future:

The v7 version shows a completely different world of cardiology, where 90%+ of cardiac conditions are cost effectively monitored directly from your mobile device. Cardiologists as we know them today will be used for people who prefer them, for routine non-surgical procedures, or are those who continue to do both fundamental and data science based research. Procedural cardiology like surgery will also get more and more automated but likely more driven by robotics progress than by data science.

But the core detection, monitoring, and managing of all cardiac conditions (not just afib) will become routinely done over time via these new technologies. The shape of the ECG waveform may even detect your potassium levels and other biomarkers (potassium is key to heart function) all in a mobile device. At home use of data (such as potassium level or ST depression or the other things today's ECG is used for and many more indicators than we understand today will be measured and more importantly processed by Deep Mind like software to make recommendations directly to the patient or their care provider. They will require not just the proper algorithms and data science expertise, but also the right user experience and design to easily connect with the patient without requiring continued human interfaces. Alternatively, these systems and devices will also connect to the electronic health record, enabling nurses and other upskilled members of the care team, who don't have to be cardiologists, to walk through diagnoses with the patient. All this will be backed up by much more complex data analysis than is possible for humans to do for each patient today and the vast and current knowledge of cardiology captured by intelligent knowledge software based on all the published research and the correlations systems find in treatment effectiveness and cautions from more advanced EMR and other data. Decision making will be far more nuanced and personalized than macro guidelines used today which seldom use detailed data or genomic

data. Many more features of ECG waveforms and cardiac imaging will be used then are used today. Adding your median cardiologist to this mix is more likely to result in errors than benefits, except at the leading edge of cardiology research where human cardiologists like possibly continue to play a major role for some time.

4. Dermatology

80% currently:

Dermatologists are able to diagnose and treat many conditions of the skin – though the traditional process requires the patient going in to see the dermatologist, who is then able to examine the patient, diagnose, and then determine the various treatment options.

Transitions:

Similar to cardiology, we are starting to see devices that look like simple mobile phone attachments that can examine conditions of the skin. Because so much of dermatology is on the visual layer, large sets of the current technologies used to both diagnose and treat conditions are different forms of lighting. We are seeing those technologies move from being only in-office to portable. A simple phone camera is often used to transmit skin images to the physician today but the critical feature of such systems is they are helping with data collection on which machine learning systems can operate and learn from using the diagnosis (often captured in the EMR) that the human dermatologists is making today. This, will be the data feedstock that artificial intelligence systems need to learn to diagnose the images in addition to just transmitting them. They then start to use the knowledge in dermatology research publications to automatically recommend further tests or call in other software specialist systems (is this an indication of adrenal function and what might confirm or disprove that hypothesis?). Such software systems have fewer biases and more complete "memory" of all the possible conditions. They will likely also be able to read many more features that the human eye might not recognize in an image (e.g is the pattern fractal and might that increase the probability that this is a skin cancer conditions?)

20% future:

Telemedicine is already showing use for dermatology and skin conditions, creating an easier more cost-effective process for dealing with or seeking treatment or diagnosis. Over time, more and more research utilizing a range of factors about your biomarkers and your microbiome will give us a much better understanding of dermatological conditions (what pattern of microbes increase psoriasis and what other implications might that have for the patients health?), changing the shape of what tools and technologies we need to detect and treat them. We are also seeing companies that are leveraging new business models to send personalized skin care treatments directly to the consumer, over time integrating more data to provide better results.

The axes of technological change will include

- 1) Finding excuses to gather more data – whether it be the direct to consumer route, or using novel forms of data gathering (e.g. iphone attachments) to collect more information
- 2) Integrating all variables into one systematic approach – leveraging our understanding of biomarkers, microbiome, and visual processing algorithms to understand our skin better
- 3) Discovering new medicine – the above will yield novel therapeutics and over the counter medication that will be both personalized and effective

The v7 version of this, as a result, won't look anything like the way we have to practice dermatology today given current limitations, but rather will be a more proactive and real-time way of keeping our skin healthy. Such systems will use many more image features, biomarker, genomic and microbiome and other data to diagnose, treat and monitor progress on dermatological conditions. And each patient, based on hundreds of factors and very current knowledge bases, will receive highly personalized treatment.

5. Diabetes

80% currently:

While endocrinologists do focus on more than just diabetes (all hormones / anything produced by the endocrine system), diabetes is one of the most prevalent diseases that they help treat and manage.

However, diabetes management and treatment right now is incredibly rudimentary: diabetes is an epidemic tied to many comorbidities, we focus on crude invasive measurements that are cumbersome for patients, we haven't yet granularly separated out different subtypes of diabetes even though it's clear they are there. But that takes more data than humans can handle in their head. But without some fundamental changes in both the practice of medicine and the underlying technologies and data science behind understanding and treating diabetes, 80% of managing diabetes will still look roughly the same.

Transitions:

Since diabetes in many cases is caused by lifestyle choices, one fundamental shift that needs to occur must be on the management side. Many have tried before, but new advances in technology and data analysis are making the prospect of noninvasive glucose monitoring more and more likely which in turn will generate orders of magnitude more data than we have today on what causes blood sugar swings. This coupled with applications on the smartphone which are already being developed will enable real-time management of diabetes that an entire care team and family can share and understand which again will increase the data available to machine learning systems.

Secondly, our scientific understanding of diabetes is changing due to our ability to use algorithms to navigate the complexity of biological data. Science Translational Medicine published a study showing how complex topological data analysis was able to identify different subgroups of type 2 diabetes patients based on a variety of different clinical and genomic data sources¹⁵¹. In just this initial research, they showed different subtypes that corresponded pretty distinctly to different comorbidities associated with that subtype of

¹⁵¹ <http://www.ayasdi.com/blog/news/identification-of-type-2-diabetes-subgroups-through-topological-data-analysis-of-patient-similarity/>

diabetes. This understanding will help treatments and management significantly going forward. Food, microbiome, genomics, transcriptomics, continuous blood sugar measures will all help decrease symptomatic health.

20% future:

Like cardiology and many of these other specialties, we can envision a world where the continuous detection and monitoring of glucose is driven by the mobile device, and there are applications and algorithms that help us manage diabetes. If the above transitions occur, software and data science will continuously and rapidly make it seamless for the patient to understand, manage and control their specific diabetes (whatever it will be called then) in real time. As Denny Ausiello has said, over time, “diabetes will go the way of dropsy, an obsolete term and we will recognize the myriad of diseases today classified as diabetes.” But that will take much more data. Given that endocrinology tries to characterize the behavior of thousands of metabolic pathways in our body that all interact with each other through the circulatory (or vascular) system with thousands of biomarkers as signaling systems, and each biomarker does multiple things (serotonin modulates depression and immune response) it is hard to imagine great endocrinology care without being able to understand and remember all the possible interactions in this complex network. Only large intelligent software systems will be able to monitor all these systems and interpret all the data we might have available. V7 systems will reduce this complex data and related knowledge graphs (including downstream implications like if thyroid is hypoactive will treatment effect vascular calcification in twenty years for this patient and how does genomics effect the optimal decision?) and reduces it to the optimal actionable recommendations.

Humans couldn't possible handle all this complexity without significant errors. Infact weekly or daily (for at risk patients) monitoring of these metabolic pathways and hormonal behavior through hundreds of biomarkers could become routine through a laser printer like device at home or in your local pharmacy!

6. Emergency Medicine

80% currently:

Emergency medicine by definition requires a hodgepodge of different medical and triaging skills, where everyone in the department needs to be highly effective and alert at all times. Given the complexities this causes, this can lead to inefficiencies, misdiagnoses, overtreatment, and other errors in the ER. Many efforts are underway to help streamline these processes, but there are a few underlying transitions that should shift emergency medicine for the better in the longer term.

Transitions:

We are seeing a few technology changes in the near term that are having an impact on emergency departments. Firstly, the rise of telemedicine (similar to the rise in urgent care facilities) helps enable much better routing of patients into the ER. Over time, we'd also expect many things to be treated outside of the ER, and telemedicine coupled with diagnostic advances from your mobile device will help jumpstart that shift. Secondly, the rise of better analytics and data tools in the hospital will lead to much better handoffs in the ER. Almost certainly much better triage will be done by software systems and not nurses or doctors in 80% of the cases ¹⁵².. Over the long term, good software will lead to upskilling the nurses in the ER to handle more and more tasks while increasing the quality of care.

Constant monitoring using simple wearable devices in the ER will allow real time decisions and continual reevaluation of both triage as well as monitoring of patients and alarms, alerts and treatments.

20% future:

Emergency medicine will need to exist for a while (and is the hardest to speculate on replacing), and optimize for similar skill sets (quick thinkers, generalists, etc) that they do now. However, computer systems should increasingly handle the load and complexity of triaging, monitoring and managing patient flow (including handoffs), leading to a much more streamlined process for the emergency care team when needing

¹⁵² http://www.digitalhealth.net/digital_patient/47751/babylon-health-launches-new-ai-triage-tool

to treat patients and the application of much deeper knowledge skills through upskilling of the care providers. This will improve all the common frustrations listed above with emergency departments, while maintaining high quality and efficacy bar.

7. General Surgery

80% currently:

Surgery is still a field driven by human specialists who are experts with their hands first and foremost.

Transitions:

Intuitive Surgical paved the way for robotic surgery, but it will still much more “robotics” than driven by software. Given how long it took for Intuitive to gain market share, technological advancements in surgery will continue to happen in fits and starts, coupled with movements in other specialties, especially anesthesiology.

20% future:

We are starting to see much more nimble systems controlled by software that can perform multiple surgeries, effectively the next generation of Intuitive Surgical. As these machines and systems gain adoption, in the long run, the majority of surgery will be robotic and controlled by algorithms guiding movement and detection. One promising example of these are systems that are derivatives of self-driving cars being used to navigate the much more predictable environment of the human body for robotic surgery. Couple these self navigation technologies to navigate the body and image and physiologic and other data pattern recognition (as illustrated in anesthesiology above), one might see unpredictable and hopefully exciting advances in surgery and procedural medicine. But given the dimensionality of what goes on in a complex surgery, especially in emergency medicine, it is hard to speculate on replacing 80% of the human role.

8. Oncology

80% currently:

We have made a lot of progress using data as well as focusing research efforts to better understand

and segment various forms of cancer. Both positive progress with precision medicine as well as negative results have not moved fast enough (with cancer still being one of the largest killers in the US).

Transitions:

The arrival of Cologuard, a stool DNA test that can detect colon cancer, is a big step towards helping both the diagnosis of cancer, and the reduced need for a general screening of colonoscopies. That being said, Cologuard is not nearly as effective as colonoscopies at detecting polyps, which can be precursors to

developing cancer over time. But the market need is clear, and these companies are paving the way to changing gastroenterology.

Transitions:

On both sides of the coin, we are starting to see initial efforts to rapidly change and enhance our ability to diagnose cancers, as well as more effectively treat cancers once they have already occurred.

Some companies are drastically bringing down the cost of genetic testing (and commoditization of that industry will open up much higher value analysis applications). Others are using sophisticated signal to noise techniques to make “liquid biopsies” feasible, thus enabling increased understanding and personalized treatment of cancers without requiring a biopsy. Others are going even further and are trying to use a blood test to diagnose colon cancer (as well as polyps that are key precursors to it).

These approaches will all continue to advance, and we’ll continue to use new forms of data to further enhance our understanding and treatment of cancers.

20% future:

There’s a lot of room to grow from there. Much of actually curing cancer will be about either understanding the biological mechanisms, or being able to do exponentially more experiments to see what works and what doesn’t. But in the future, we will think of oncologists as controlled experimenters enhanced by data science. And while not proposing that we will necessarily cure

cancer in the medium term ... we will be able to tackle them piece by piece, with the many of our innovations driven initially by efforts by machine learning.

The future of colon cancer diagnosis (and many others) will be driven by biomarkers from a simple blood test. Data science will help extract the right biomarkers (small molecules, proteins, etc) that highly correlate with the presence of cancer precursors that can then be removed (in the case of polyps) or alleviated.

Today, the most complex cases of cancer have a tumor board, where multiple human experts will all examine a case, and share their opinions in diagnosis and treatment path. In the future, this will be a virtual tumor board, where different “AI experts” will be absorbing different sources of data, interpreting, and giving responses in order to advise the oncologist on the right path of treatment. And these systems will each get better over time.

9. Psychiatry

80% currently:

“In the future, when we think of the private sector and health research, we may be thinking of Apple more than Lilly and Pfizer,” Tom Insel, former head of the National Institute of Mental Health who recently joined Alphabet to work on the intersection of mental health and technology¹⁵³.

It’s clear we are in a mental health crisis, and the field of psychiatry has not kept up with other medical fields in progressing towards being scientific and data-driven. On top of that, even the institutions that should be best equipped for managing mental health patients are running into severe trouble (both in terms of bandwidth and ability to care for those patients).

As psychiatry takes the leap and moves towards being more data-driven, we will see the role of the psychiatrist change substantially, and the ways we measure and evaluate “disease” vs. health, and

therefore manage treatment will be different as well.

Transitions:

As mentioned above, when the leading mental health researcher leaves the national institute for Google Life Sciences, that’s making a statement for where the field is going. Top researchers are all pushing for the disbandment of the DSM (the bible for psychiatry discussed earlier), and focus more on biological validity of particular conditions¹⁵⁴. This will make psychiatry both data-driven and science-driven, which will change the practice and the standards in the field completely.

On top of that, we are seeing research and products live that indicate we can diagnose and monitor mental health conditions using newly available data from our smartphone, our recorded voice and its intonations, and other sources that haven’t been included in any current standards of psychiatry. As in cardiology, psychiatry calls for 24x7 monitoring of patient which a human can do so we must rely on sensors like your smartphones¹⁵⁵. One system measures passively on your cell phone all movement (how often did the patient get out of bed? Go to the kitchen?), behaviors (like frequency or response time to email, Facebook, Twitter and messaging interactions), and hundreds of data points a day and has been able to identify hundreds of micro behaviors that are not in psychiatric literature. These are then used to predict and identify mental health episodes like bipolar behavior or maniac depressive episode. These are much more likely to be combined with the biomarker studies to create diverse objective measurements of data to help diagnose, manage, and eventually treat mental health episodes.

20% future:

The majority of a psychiatrist’s job will be mobile monitoring, intervening, and administering mobile therapeutics (hopefully more often than pharma products which will have greater side effects). It will be a long road to get here, but given how much progress has been made in diagnosing mental

¹⁵³ <http://www.technologyreview.com/news/541446/why-americas-top-mental-health-researcher-joined-alphabet/>

¹⁵⁴ <http://www.technologyreview.com/view/514571/nimh-will-drop-widely-used-psychiatry-manual/>

¹⁵⁵ <https://ginger.io/in-the-news/> <https://www.computer.org/csdl/mags/pc/2012/04/mpc2012040036-abs.html>, <http://hd.media.mit.edu/tech-reports/TR-670.pdf>

health episodes using smartphone app, it's only a matter of time before those apps will also be able to alleviate those conditions. And given the real-time nature of these mobile applications, and their ability to learn given increased, personalized data, our understanding of mental health should exponentially increase in a short period of time. suspect we will see digital drugs that modify the trajectory of episodes (will puppy pictures or a call from you sibling change the direction of a depression episode if provided early before it gets serious given one will be able to identify the onset of these episodes very early?)

Perhaps even more important, research is showing that our mental state is a continuum – so everyone, not just those with psychiatric conditions, will be able to benefit from these monitoring and therapeutic tools. Everyone will be able to move more to the side of wellness and mental health in a low cost, scalable fashion.

We are starting to see conversational AI systems that might use guidelines from psychiatric research to make 24x7 talk therapy available to all patients through AI bots? Or will they atleast assist a therapist dramatically lighten their workload by suggesting questions and answers and be comprehensive in their knowledge of psychiatry?

10. Radiology & Pathology

80% currently:

Radiologists use a variety of medical images for diagnostic purposes. This task is inherently a computational one, and is being done by humans primarily for historical reasons. Unfortunately, even in the 1990's, it was clear that technical progress has outpaced human ability to be error free and without variation in reading images ¹⁵⁶.

Transitions:

Two changes are occurring simultaneously that is changing the practice of radiology (and pathology though different will also face very similar progression, though it has additional complications). 1) we are seeing databases of medical images opening up for researchers and commercial entities to develop algorithms over

them. Hospitals are looking for ways to reduce costs, and algorithmic identification and diagnosis is orders of magnitude cheaper than requiring a radiologist. 2) we are seeing those algorithms bring us new understanding of diagnosis, and over time they will prove to be more accurate than humans. Newer data science techniques are able to detect both common and rare conditions as long as the available training data is there.

20% future:

In particular with radiology, it is inevitable the majority of radiology will be algorithmically driven. In certain subspecialties, there may still be people augmenting the computer systems. And we will always need people for new algorithm development, and a further understanding of measurement, especially as newer or more portable imaging techniques come out (such as portable ultrasounds). But over time, much more than 80% of radiology will be completely displaced by active learning computer vision systems. Unlike say emergency medicine where the reduction in the role of humans is much harder to speculate on, in radiology it feels much more a matter of time than speculation.

11. Sleep Medicine

80% currently:

Sleep medicine is focused on treating either sleeping disorders (including sleep apnea or insomnia) or abnormal events that happen during sleep (such as sleep walking). We are increasingly seeing the wide range of effects our sleep quality has on all areas of our lives, whether it be performance at work, athletics, or elsewhere. However, the field of sleep medicine still suffers from (at least) two fundamental issues: (1) their focus on sleep studies puts patients in a very controlled situation, and studies typically have very small sample size prior to being published. While these are still useful (in particular for diagnosing individual conditions), they don't collect nearly enough data to take advantage of many of the new data science techniques that can help us understand sleep better. (2) the devices and mechanisms by which to traditionally treat sleep disorders are extremely cumbersome.

¹⁵⁶ <http://www.ncbi.nlm.nih.gov/pubmed/9536897>

For example, even though CPAP is extremely effective for sleep apnea, a large proportion of the population still won't use CPAP devices due to their inconvenience.

Transitions:

We are already seeing a massive change in the way we think about and analyze sleep, and it happened orthogonally of traditional sleep medicine. As wearable trackers became more and more popular, they were able to gather the most comprehensive data on primary sleep patterns. While they weren't capturing the same dataset as what one would capture in a sleep clinic, because they were able to capture that data on many orders of magnitude of scale greater than even all sleep clinics combined, data scientists have been able to deduce habits, trends, and insights into our sleep that were previously unknown. Though early systems are laughable it is hard to see why they won't get substantially better by v3 or v5! Along with newer technologies that substantially reduce the invasiveness of monitoring many more things (continuous HR, respiration, environmental factors or pollutants....) and do it seamlessly, sleep medicine will surely be data driven with closed, less intrusive loops to treatment inside the patients home.

These trackers also had the additional benefit of bringing sleep science and knowledge front and center to the average consumer. The increased data collection will make the discipline much more scientific and large scale in its applicability because of the low cost of providing much better and responsive care. It is likely that the whole discipline of sleep medicine will change radically, driven by digital health technologies. These, combined with other smart devices such as smart lighting to help with maintaining proper circadian rhythm, will lead to a better awareness and probably better sleeping patterns amongst consumers. They will also help drive adoption of more consumer-friendly CPAP devices.

20% future:

The norm for a sleep study in the future will be able to use thousands of patients, via ResearchKit or equivalent applications. These combined with more open data access will lead to a rapid enhancement of our understanding of the science of sleep. We will see applications beyond our current set of sleep medicine, as well as different therapeutic mechanisms (including meditation apps, direct brain stimulation). The 80% of what sleep medicine is now will go towards data science and better device manufacturing.

04

SOURCE, TIMING, INCENTIVES, AND PITFALLS OF HEALTHCARE INNOVATION

I have tried to focus on a single dimension of healthcare innovation, that which is driven by digital health technologies or healthcare innovation driven by computing and software. To clarify, there also will be other sources of innovation, and in particular, the biological sciences will play a large part. That said, if I was forced to pick (a choice I prefer not to make), I personally believe that computational, software and digital health technologies have more potential to drive innovation in applied medical care and even be guides for new biological research over the next one to two decades than the biological sciences (although the latter's contribution will be substantial nevertheless). In fact, hybrid efforts where both technologies combine to substantially improve our healthcare toolbox may be the most promising!

Healthcare technology to date has been made subservient to the current system in order to maximize profits and more often than not has led to increased healthcare costs, but this will begin to trend towards less costly and more easily accessible mobile technologies (warning: the cost of technology does not always relate to price charged for it, so we will need to be vigilant on the way this impacts the actual price people pay since the potential for substantially lower costs will exist). This trend can be hijacked, and vigilance and competition will be essential. Less bureaucracy will increase competition and help realize more of the potential for cheaper, more accessible and better technology. Because traditional medical drugs, procedures and protocols have higher potential for negative consequences the regulatory bodies have necessarily needed to be more bureaucratic and slower moving. Often, but not always, digital health technologies I suspect will offer fewer

and less material negative consequences and hence easier paths to approval. The U.S. Food and Drug Administration (FDA) can help by being progressive, but it also is possible that the innovation first will happen in parts of the world that have very limited access to healthcare today, such as the uninsured in the U.S. or the many people in India who don't live within geographic proximity to a doctor or medical facilities. If that happens in underprivileged locations, such studies will require that rigorous ethical standards be met, and the focus will have to be on the net healthcare benefit to the communities.

These new technologies, often starting from outside the system or on its periphery, will allow us to provide care to those who don't have it now and will prevent simple medical issues from getting worse before they are addressed (hopefully with help from the U.S. FDA, which has been treading carefully (as it should) into these waters by approving an ECG machine as an over the counter device without a prescription and approving iPhone software applications as a pharmaceutical drug! A more nuanced approach to risks and benefits for new approaches and insights will need to be taken by the regulators to encourage experimentation that is not easily accepted in the healthcare system (except inadvertently and by doctor's individual biases running many more experiments than regulators would knowingly allow). Food & behavior could actually become drugs in measurable/quantifiable ways, not just a "Mark Hyman" evangelicalism of if you eat $x + y + z$, then ... Of course regulatory bodies all over the world could delay the advent of this radical vision of medicine by a decade or two by either slowing down v0 and v1 innovations as these won't be able to prove efficacy early or by delaying the more radical forms such as implanted

nanobots or extensive genetic engineering of human cells, as a natural human reaction to change.

Where will this innovation come from? In most areas, this happens from innovators outside the system who act somewhat naively, failing and then realizing they need some knowledge and collaboration with the system. Entrepreneurial teams often add domain expertise to their naïve “fresh piece of paper” re-invention ideas. Mainstream players tend to address lower risk but more incremental innovation in mainstream markets, while often entrepreneurs try and address smaller but exponentially growing markets (not all entrepreneurs go after re-invention but the ones driving the change with which we are concerned do). These are not absolutes but rather “more true than not” rules. Society generally tries to assign more power to larger entities, like governmental institutions and the Fortune 50 behemoths, but true radical innovation seldom comes from them. As a reminder of ground covered earlier, did Walmart reinvent retail or Amazon? Did General Motors reinvent electric cars or Tesla? Did SpaceX reinvent space launches or NASA and Lockheed Martin? Did Google invent social networks or Facebook? Did NBC reinvent media or YouTube? What did Google know about media? Most importantly did big pharmaceutical companies reinvent biotechnology pharmaceuticals or did Genentech? The recent book about Elon Musk and the progress of Tesla and SpaceX and the education, failures and evolution of their naivety is a good role model of how things may develop.¹⁵⁷

Let’s take the birth of Genentech as an example. It only came to exist because big pharma was too conservative to innovate in the 1980’s on drugs using new technology, now known as biotechnology (Bob Swanson, an associate at a venture firm, co-founded Genentech from cutting-edge university research). Eventually, larger companies did later help scale biotech to what it is today. As another example, even though Toyota and Volkswagen did not lead the electric car revolution (although Toyota did less risky version with the Prius hybrid), I suspect that they will help electric car technologies like Tesla scale and become mainstream. They also are following

Google’s lead in the driverless car, but are unlikely to introduce them in radically new configurations. Instead, they mostly focus on their use as “safety” technologies that help you avoid drifting outside your lane or keep a safe distance from the car in front of you. They have little interest in innovating around the idea of car sharing, where a car can come to you anytime you want it, and therefore can be shared by a dozen people, since it would reduce the total overall demand for cars. It is to be noted that given the pressure around developing driverless cars from the naïve players like Google, Uber and Tesla, many traditional players are being forced to respond and follow. Something similar may happen in the medical domain. For consumers, driverless cars with a software driver will be safer and more importantly less expensive than a car that one owns. Further it will enable cost effective Uber like services that will be cheaper than owning a car for most people. Whether that occurs or not in medicine will likely depend upon which entrepreneur wills this change to happen. It is possible for health care to get much cheaper and much better for everyone except maybe the health companies and some medical professionals. If health care costs are half the percentage of our GDP that it is in the US today, some people and organizations will get hurt and they will fight these changes. The collateral benefit will be the elimination of human errors and limitations and the introduction of much more accessible, better, scientific and personalized medicine and wellness.

Since innovation is unlikely to come from established companies, then where will all come from? Some believe we have to work within the constraints of the medical establishment in order to advance innovation, but I disagree. Again the story of Tesla working outside the established automotive industry is a good role model. Some will follow reluctantly and some will try to lead, but most organizations in traditional healthcare will fight this trend towards reduced costs, because it also reduces profits. Any reduction in health care costs is a reduction in some healthcare providers revenue and often, profit. This is one reason why it is outsiders who will most likely disrupt the system. By comparison, landline phone calling

¹⁵⁷ Vance, Ashlee. “Elon Musk: Tesla, SpaceX, and the Quest for a Fantastic Future”.

rates from companies like AT&T didn't decline until mobile operators changed the standard for what a phone call should cost. Remember how expensive long distance calling was not very long ago? Innovation seldom happens from the inside because existing incentives are usually set up to discourage disruption. By comparison, many other industries have a storied history of having innovation come from outside the system. What did I know about computing when I started Sun Microsystems? I did not even have a computer science degree and had never visited an IBM or DEC facility. To be fair, outside innovators often work with those inside the system to mutual advantage and often drive changes that are then scaled by the more innovative insiders.

Although institutions will individually adopt some point innovations, it is unlikely that we will see large-scale adoption by fee-for-service organizations unless it makes them more money (Proton Beam accelerators, anyone?).¹⁵⁸ An the accountable care organizations in the US or various national health systems will focus on cost minimization! Some external forces like accountable care organizations (a wholly U.S. phenomenon) will grease the adoption of such innovations, because their small organizational footprint requires them to more intensely manage patient costs. Other uninsured consumers who have few alternatives also will try these initially risky and less well established approaches. It is possible such systems first get adopted in countries with severe shortages of medical resources like India. Further, the timing of such adoption and when the knee of the adoption curve happens is hard to predict and may depend on extraneous forces like the FDA, good or bad publicity and other events. I suspect self diagnosis has increased dramatically worldwide, but under the visible radar. The latest findings from the Pew Research Center's Internet and American Life Project showed that 35- percent of

U.S. adults have gone online to self-diagnose a medical condition and 59-percent look up health information online.¹⁵⁹ v0 efforts, like WebMD and Google health queries are relatively poor beginnings.

By contrast, the medical system is invested in doing things the same way. If a hospital could cure a patient in half the time, would they be willing to cut their business in half? Would a hospital work to reduce infections? Today, hospitals try and increase revenue per bed per night. If a patient gets bed ulcer, it means \$40,000 in extra revenue for the hospital.¹⁶⁰ How hard are they likely to work to eliminate the scourge? This is a phenomenon that affects the U.S. and the developing world, which both operate under a private for-profit model; incentives for change and innovation in most large government run organizations also are limited. Some non-U.S. organizations are attempting to run early experimentation with these efforts. The risk to very large organizations (like the National Health Service in the UK) to jump headlong into such new efforts and the potential backlash would not be good policy management. But they can help a lot by trying it at the edges of the traditional system.

While many things may increase patient's quality of care, if they are not financially beneficial or beneficial to the central authority in centrally provided health systems, they will not be optimized for. This was especially true in the fee-for-service world, where studies have shown that clinicians reimbursed for each service tend to recommend more visits and services than clinicians reimbursed under other payment methods, sometimes increasing visits between 11% and 61% (depending on the specialty)¹⁶¹. Psychiatrists have the same problem with incentives. They do not get paid to cure patients; they get paid to endlessly treat them over many expensive therapy sessions, which does affect

158 RAND Corporation. Redirecting Innovation in U.S. Health Care – Medical Arms Race (http://www.rand.org/content/dam/rand/pubs/research_reports/RR300/RR308/RAND_RR308.pdf)

159 Fox, Susannah and Duggan, Maeve. "Health Online 2013". Pew Research Internet Project. 15 Jan 2013 (<http://www.pewinternet.org/2013/01/15/health-online-2013/>).

160 Halvorson, George. "Don't Let Health Care Bankrupt America" 2013, p59

161 Helmchen, Lorens A., and Anthony T. Lo Sasso. "How sensitive is physician performance to alternative compensation schedules? Evidence from a large network of primary care clinics." *Health Economics* 19.11 (2010): 1300-1317. (<http://onlinelibrary.wiley.com/doi/10.1002/hec.1551/abstract>).

their behavior (although, many psychiatrists certainly focus on patient welfare). From the therapists' perspective, would you rather have a steady stream of repeat patients that fill your hours or have to attract new patients? The former has minimal patient acquisition costs, and a continual revenue stream. A cured patient is just like a terminated account at a cable TV subscription service. The excess surgeries some surgeons do are well documented. A USA Today study found that unnecessary surgeries may account for 10 to 20-percent of operations depending on the speciality.¹⁶² For instance, with spinal fusions, the study found that 10-percent of the surgeries paid for by Medicare in 2011 were unnecessary because "there was no medical basis for them or because doctors did not follow standards of care by exploring non-surgical treatments. In just one year and for a specific kind of spinal injury, the sum total cost savings would have been \$157 million. Another study published by the Journal of the American Medical Association in 2011 showed that in 22.5-percent of pacemaker-related surgeries were unnecessary.¹⁶³ That wasn't good for the patient or the payer but it was good for the doctor. It's clear that not all doctors or hospitals and other institutions behave the same way. There is a distribution of behaviors from very ethical patient centric care to profit maximizers with a disregard for the best interests of the patient. But when incentives are misaligned the overall statistics tend to be misaligned with patient welfare choices.

Another area of misaligned incentives is end-of-life care. Today, 25-percent of U.S. Medicare dollars or more than \$125 billion is spent on services for five-percent of beneficiaries in their last year of life.¹⁶⁴ In many instances and especially with terminal illness, the additional treatment does not increase quality of life. This disproportionate funding presents many challenging questions:

should we provide the best care possible or the best care we can afford, and how should budgets be allocated? By being in denial of our mortality, we are driving up the cost of healthcare, and those who stand to financially benefit within the medical system are not incentivized to make any changes. These choices need to be made by society as a whole rather than by individual physicians. When society bears the cost, we need consistency in spending decisions. In the cases where individuals bear the cost, then excess or premium care can be determined by the payer.

Expecting the medical establishment to do anything different is like expecting them to reduce their own profits. Medical device manufacturers, like those that build and sell huge scanning systems, don't want to cannibalize sales of their expensive equipment by providing cheaper, more accessible monitoring devices like a \$99 ECG monitor. The traditional players will lobby, goad, pay and even intimidate doctors and regulators to reject these new devices, falsely claiming that they aren't as good if they provide only 90-percent of the functionality (but at 5-percent of the cost) in many cases.

For instance, many anesthesiologists are strongly opposing the use of the Sedasys¹⁶⁵ device, which would automate sedation during colonoscopy for many patients and remove anesthesiologists from the room and reduce their income.¹⁶⁶ According to a Johnson & Johnson (JNJ)-sponsored study by RAND, over \$1 billion is spent each year sedating patients undergoing otherwise painful colonoscopies. An anesthesiologist's involvement typically adds \$600 to \$2,000 to the procedure's cost, according to a research letter published online by JAMA Internal Medicine. In fact, health-care number crunchers have frequently targeted anesthesiologists for savings. Their median annual salary of \$286,000 is ninth among

¹⁶² <http://www.usatoday.com/story/news/nation/2013/06/18/unnecessary-surgery-usa-today-investigation/2435009/>

¹⁶³ Eisler, Paul and Hansen, Barbara. "Doctors perform thousands of unnecessary surgeries." USA Today. 2013 June 2 (<http://www.usatoday.com/story/news/nation/2013/06/18/unnecessary-surgery-usa-today-investigation/2435009/>).

¹⁶⁴ Wang, Penelope. "Cutting the high cost of end-of-life care." Fortune Magazine (12 Dec 2012). Retrieved 4 Feb 201 (<http://money.cnn.com/2012/12/11/pf/end-of-life-care-duplicate-2.moneymag/>).

¹⁶⁵ Frankel, TC "New machine could one day replace anesthesiologists", The Washington Post, https://www.washingtonpost.com/business/economy/new-machine-could-one-day-replace-anesthesiologists/2015/05/11/92e8a42c-f424-11e4-b2f3-af5479e6bbdd_story.html

¹⁶⁶ Rockoff, "Robots vs. Anesthesiologists", The Wall Street Journal (2013).(<http://online.wsj.com/news/articles/SB10001424052702303983904579093252573814132>)

all physicians and third among non-surgeons surveyed by PayScale.com, a salary data and software firm. In 2009, \$1.3 billion was spent on 12.5 million of such gastroenterology probes, according to a RAND Corp. study paid for by JNJ's Ethicon Endo-Surgery unit, the maker of Sedasys.¹⁶⁷ The study also suggested that \$1.1 billion of the spending was for low-risk patients who didn't need it. And this is US- only numbers. As of 2016, J&J has stopped sales of its Sedasys system due to this pushback.

Another example is when hospitals compete with one another for business; they will market high-cost technologies to patients. There's no evidence that the new technology — for instance, proton beam therapy — is actually better than traditional treatment options, but hospitals want the branding of expensive new “exclusive” equipment that attracts more patients. In yet another example, let's look at the ongoing fight between doctors and nurse practitioners. In the first half of 2013 alone, the California Medical Association, the lobbying group representing doctors, spent \$1.2 million fighting scope-of- practice¹⁶⁸ bills. In order to keep their exclusivity and reduce competition, doctors are fighting to prevent nurse practitioners from prescribing medication, for example antibiotics, and treating patients independently.¹⁶⁹

From the pharmaceutical companies' perspective, they push marginally different drugs instead of generic solutions that may actually be better for patients, because they want to have exclusive drugs that have less price competition. In actuality, they want to increase the number of drug subscribers and generate recurring revenue for as long as possible. They would rather sell a cholesterol-lowering drug than encourage healthier eating habits, which would reduce their own profits. If they permanently reduced cholesterol, they would lose customers!

Another good example is the cardiovascular “polypill”. The polypill is intended to reduce blood pressure and cholesterol, both of which lead to

cardiovascular disease. It combines multiple different medications that are often taken separately and each of which have received FDA approval. By combining these medications into one pill, it creates an easy and affordable alternative that epidemiological models suggest would have significant public health benefits. However, U.S. regulators have not yet deemed the polypill safe and efficacious enough for approval for lack of appropriate trials. Further, due to the high cost of securing FDA approval of new drugs, major pharma companies have little incentive to develop the polypill and do the trials to get it approved, simply because the profit margin isn't high enough even though the health benefits could be significant, this has likely stalled development of the polypill.

The misaligned incentives with pharmaceutical companies go even deeper with drug trials. Companies want drugs or devices approved with small patient populations, because their priority is getting to market as quickly as possible with the lowest budget possible. One cannot blame them for doing this given the immense cost of an FDA trial, as they are not in the money losing business. When testing a new treatment, a small number of people with just a few variables in trials means a higher probability of approval and being approved for routine clinical use faster. For example, a drug is tested in a conventional medical trial with patients who are between 60 and 70 years old and have reached a certain progression or stage in a disease. The trial demonstrates the drug is effective in 65-percent of patients tested, which may be enough for it to be approved for routine clinical use. Now that the drug is widely available, it's still unclear for whom it will be most effective, since they used such a narrow focus during the trial phase ignoring most of the population. As indicated in the first section, this excludes accounting for even the most basic demographic factors like ethnicity and other age ranges, not to mention comorbidities, where 80% of chronic care costs are spent.

167 “An Analysis of the Labor Markets for Anesthesiology”, RAND Corporation (2010). (http://www.rand.org/content/dam/rand/pubs/technical_reports/2010/RAND_TR688.pdf)

168 <http://www.ncsl.org/research/health/scope-of-practice-overview.aspx>

169 Mason, Melanie. “State bill to boost use of nurse practitioners goes no where”. Los Angeles Times (01 Sept 2013). Retrieved 5 Feb 2014 (<http://articles.latimes.com/2013/sep/01/local/la-me-healthcare-20130902>).

Economists at the University of Chicago and MIT have done a detailed analysis which shows that pharmaceutical research is skewed against research and development of treatments that prevent long term health problems and extend life by large amounts because of the expense (both in terms of investment but also in terms of patent protection and lifetime economic value of such potential therapeutics).¹⁷⁰ In short, for conditions like cancer, the economic incentives are all to develop therapies which extend the end of life by a short amount instead of developing therapies which prevent disease in the first place. This is observed across the spectrum of therapeutics. It is typically far more financially rewarding to treat the chronic symptoms of disease than try to cure or prevent it.

The most fundamental tenant of science is that it is objectively based on facts; the idea that research is not reliable and clinical trials cannot be debunked. As mentioned before, the work of Professor John Ioannidis at Stanford University shows the opposite to be true. He has instead shown that published research findings are more likely to be wrong than right due to sample sizes that aren't large enough, effect sizes and statistical significance that are too small, other interests and prejudices that are present, and improper treatment and pre-selection of tested relationships.¹⁷¹ Correlated effects can complicate findings and lead to incorrect attribution of causal relationships where there are none, e.g. variations due to patients' ethnicities, the presence of co-morbidities and the influence of environmental factors that aren't properly accounted for — all of these factors can lead medical research astray.

But this misleading data does not exclusively reside in the domain of for-profit institutions. One example is from hematology and oncology researchers at Amgen who attempted to replicate 53 "landmark studies" related to cancer, but they were only able to reproduce the results in six.¹⁷² This meant in only ~10% of cases were the results of the previous cancer studies reproducible

enough to meet the quality criteria to drive further cancer drug development research. The errors were in this case driven by the "publish or perish" incentives in the academic world or sometimes by lack of full understanding of statistical validity among biological science researchers. This creates a high level of bias in the data and puts potential patients at risk, but the offending parties are not incentivized to change the system because they get their product to market quickly or meet their publication goals. do hope a good knowledge graph will surface many of these inconsistencies and resolve many and highlight the rest.

To be fair, there are many great doctors, researchers and many ethical organizations and people. The point is that the incentives in healthcare make innovation from within unlikely. Fortunately, it doesn't matter if the establishment tries to do this or not, because these changes will happen regardless. Game theory and game mechanics will need to be thought through to align incentives of all the players to get the right technologically possible solutions to be adopted more quickly. Otherwise, we will see large delays, active obstruction and lots of fear, uncertainty and doubt spread by the various parties.

We must realign incentives with patients as the central focus rather than the system. This change may start at the periphery with the people who need it the most — the millions of uninsured people in the U.S. or the hundreds of millions of people in India without access to any doctors. There aren't enough rural doctors in India and few of them have access to the New England Journal of Medicine, CT scanner or even reliable electricity, but most potential patients and healthcare providers have cell phones. Like the Internet, healthcare will likely shift to a consumer driven system (or change will be substantially delayed and slowed), which will allow for less money to be spent on healthcare costs. It will allow us to provide care to those who don't have it now. It will help avoid errors and provide basic services to those who cannot afford full healthcare

170 Budish E, et al. 2015 <http://www.nature.com/articles/srep00196>

171 Ioannidis, John PA. "Why most published research findings are false." *PLoS medicine* 2.8 (2005): e124. (<http://www.plosmedicine.org/article/info%3Adoi%2F10.1371%2Fjournal.pmed.0020124>)

172 Begley, C. Glenn, et al. "Drug Development: Raise standards for preclinical cancer research." *Nature* 483, 531– 53 (29 March 2012) doi:10.1038/483531a.

services. It will prevent simple things from getting worse before being addressed.

There has been much ado in various blogs about how Silicon Valley and outsiders don't understand healthcare and therefore should not or cannot try to understand and innovate it. As I explained above, and granting the rare exception, it is hard for insiders to innovate within a system, at least when it comes to radical innovation. That's not to say that Silicon Valley and other outsiders won't leverage the system and have partners and doctors from inside the system helping them. Many others work with doctors in testing and studying their ideas and technologies. Most startups we are funding have physicians on their team and collaborate with other healthcare partners who are inclined towards innovation and experimentation. There are many progressive members in today's healthcare system driving to make these changes happen.

Some comments on costs & technology/practice dissemination in healthcare are warranted as it is a material driver of what does or does not happen. The goal here is to describe one potential vision of where medicine might go rather than on economics, but just like incentives, economics does drive a lot of what technology does or does not get adopted. Incentives or game mechanisms to increase competition and encourage lower cost digital technology will become necessary and policy will encourage or discourage this and hence determine the rate of adoption and pricing of these new approaches to medicine. The pharmaceutical mobile app has unfortunately been priced as high as at \$180/mo (varies by plan) when it did not cost as much as a pharma drug to develop. We must look at ways to encourage dissemination of technology that is closer to Internet models than traditional healthcare models. We need creativity of business models. Personalized & genomic data in healthcare can increase or decrease cancer treatment costs for example. We will need to focus on "accessibility" of digital health technology broadly instead of the traditional focus on maximizing profits, which has often been a hallmark of technology in healthcare. We will also need to encourage data openness even at some risk (e.g. datadonors.org) & better anonymization

technology.

To go a bit further, we adapt some of George Halvorson's work to understanding how cost impacts the healthcare system, and why technology and an entrepreneurial mindset can go a long way in alleviating the current concerns (specific to the US, but the lessons are more broadly applicable).¹⁷³ Halvorson was CEO and Chairman of Kaiser Permanente. His distilled premise is that the way to drive down costs in the US healthcare system is by moving healthcare to a state where the incentives and goals are to improve people's care and well being. When the system is financially incentivized to achieve those goals (i.e. that is aligned with profit maximization for the individual entities), healthcare costs will rapidly and dramatically come down. The best hospitals and healthcare systems in the US operate this way (just as the top doctors do), but it is not true of the typical situation.

Rather than discuss all the specifics Halvorson mentions to reduce costs, we can focus in on three interrelated themes that greatly impact care costs right now: lack of proper coordinated care, lack of incentives / business models aligned to improve patient care, and the lack of connection between improving healthcare in every day life or for chronic disease care versus healthcare while you are in the hospital for an acute need or for a procedure. Digital health technologies can help improve each of these costs if applied correctly.

The lack of coordinated care and dysfunction in the care delivery system has a lot of, now obvious, side effects. For example, something seemingly as simple as delivering prescription drugs leads to more than million prescription drug mistakes in the US year. And this is one of the more simplistic examples of where a natural coordinated care system would help ensure that patients were getting the proper prescriptions. The early v0 systems can help ensure that many different entities in the hospital are reaching similar conclusions for their care, ensuring coordination. But even simpler than that are just the v1 efforts of EMR's that are trying to both incorporate a range of relevant data sources as well as provide a UX that is useful and intuitive

¹⁷³ Halvorson, George. "Don't Let Health Care Bankrupt America" 2013

for people in the healthcare system. These front end systems will nudge people to do the right thing and help prevent simple mistakes across complicated coordinated care, and early versions of technological innovation is helping.

Related, coordinated care is so rare across the healthcare system because individual entities (hospitals, clinics, doctors) are not incentivized to care for these patients. One of the major benefits of the Affordable Care Act (PPACA) is the push for more Accountable Care Organizations (ACO's) that helps fix those incentives from a regulatory standpoint (an important step, but not one to bank the future on). But the current state of incentives is dismal. Business models currently allow for doctors to effectively not care about improving their care outcomes, and many times, poor care outcomes actually benefit hospitals. For example, in the past pressure ulcers from patient immobility generated a lot of cash flow for hospitals. "A five thousand dollar patient can become a fifty thousand dollar patient if the ulcer for that patient is diagnosed slowly and if the treatment for the patient is delayed."¹⁷⁴ Although most payers, including Medicare/Medicaid no longer reimburse healthcare providers for bedsores that appear under their care, and make hospitals financially responsible for paying for treatments for many problems that results from gross examples of negligence in care,¹⁷⁵ the fee for service system is one that pays for care activities and utilization of healthcare resources instead of patient health and wellness. Just informing and incentivizing providers to care about alleviating this outcome leads to improved results (from a care standpoint), as the dramatic increase in attention to pressure ulcers as a result of changes in reimbursement demonstrates. While incentives many times are functions of non- technological systems, as the data-driven diagnostic systems evolve, they can be focused on maximizing an objective function of improved care. Systems deployed at a hospital could give suggestions that ensured not extra fees and revenues, but rather proper care for the patients at the bedside. And if the patient has access to these "second opinion" tools, they might be able to determine their increased likelihood of a

problem (like a pressure ulcer) before the hospital (who still might not be incentivized) and call for treatment.

Similarly, at the moment, patients / consumers currently have very little data to work with when navigating the healthcare system. Not only can they rarely evaluate the performance of given hospitals or physicians (although this is changing), but they aren't focused on how actions in their every day lives are vitally important to their overall care (well-being and costs). As Halvorson notes, building systems and support networks to encourage everyone in the US to walk for 30 minutes a day will cut incidences of diabetes by half, risks of colon cancer in males by half, as well as decreased rates of depression, stroke, and a slew of other serious medical conditions. But this home activity does not generate revenue for hospitals. If the healthcare system as we know it didn't change, but consumer-driven technologies (focused on the home and the internet) somehow enabled this outcome of walking, it would be the largest contributor to the average person's well-being (by far), compared to advances in medical, biological, and even data science! Digital technology can make known good behavior easier to follow as Omada and Welldoc are showing, but we are still very early and crude in implementing and incorporating those into our healthcare system. We will need all of these systems (consumer-driven, data-driven, hospital-focused, etc ...) to focus on improving wellcare before we get the transformation outlined in this thesis.

These are a just a set of many things that might affect the cost, accessibility and rate of adoption of these new digital technologies. I mention them because they have large implications but these issues are not the primary subjects of this thesis or vision for medicine.

The reality is that healthcare has to move in this direction in order to make it affordable for everyone. There are many arguments and challenging questions in the blogs about this point of view. Some have answers, many reflecting naiveté in understanding how technology

174 Halvorson, George. "Don't Let Health Care Bankrupt America" 2013. Page 60

175 https://www.cms.gov/medicare/medicare-fee-for-service-payment/hospitalacqcond/hospital-acquired_conditions.html

incrementally changes and evolves, and many questions and criticisms truly don't have answers yet. The early versions of IBM's Watson computer did abysmally in competing against humans in the game show Jeopardy. But within a few years, the system beat the best humans at a game judged too difficult for computers! The same happened with chess-playing computers until they beat the world champion at chess! And today in 2016 AlphaGo, DeepMind and other similar technologies make IBM's Watson and similar systems look like so yesteryear!

Some reporters ask good doctors with no knowledge of machine learning technology if this will happen. Those answers I find largely irrelevant except to the extent they point to capabilities computers don't have today. How qualified is a doctor (who doesn't understand the rate of change from IBM's Watson Jeopardy computer to AlphaGo and DeepMind or the changes in the level of sophistication) in predicting what software technology would do in two decades? The shortcomings a doctor may point to help define what computers need to get good at! Just because an answer doesn't exist today does not mean that it won't be found or that we won't find workarounds. Some things will come as tradeoffs to make healthcare more affordable. My guesstimates will be wrong on many counts as new technologies and approaches, and sometimes-unforeseen problems, emerge. Many comments have come from good and passionate doctors (there are plenty of them around!) and bloggers who have health insurance and can already afford good care. I personally worry about the bottom half of doctors globally who are too rushed, too overburdened, too mercenary or too out-of-date with their education, especially in the developing world. And I worry about the patients who have little access to doctors or modern healthcare.

Entrepreneurs can come at these challenges from the outside or inside the system and inject new insight. They can ask naïve questions that get at the heart of assumptions that may be both pervasive and unperceived. They can jump in as often happens by underestimating the complexity of the problem and then be forced to solve the complexity to survive. Necessity is often the mother of invention in entrepreneurial world and

naivety an important tool to help get started when more knowledgeable people might be discouraged from starting because they understand the true complexity of the problem. This was amply illustrated in the book on Elon Musk's efforts at Tesla and SpaceX. Once in trouble, entrepreneurs can leverage the many insiders at the right time to provide real understanding of medicine. They can build smart computers to be objective cost minimizers WHILE being care optimizers. Domain expertise has a place, and the smartest doctors aren't outraged at this idea (just the ones with knee-jerk reactions). People always react against technological progress, and many don't have the imagination to see how the world is changing (when street lights were first introduced, Yale students objected to them as detailed in the book "When Old Technologies Were New"). But there will be many good doctors willing to assist in this transition. Eric Topol (author of "The Creative Destruction of Medicine" and "The Patient Will See You Now") and Dr. Daniel Kraft, have called for a data-driven approach to healthcare and are examples of insiders who think like outsiders. There's no question that many naïve innovators from outside the system, maybe even 90-percent of them, will attempt this change and fail. But a few of these outsiders will succeed and change the system. The fact that the vast majority of dotcom companies failed during the "dotcom bust" in the late nineties did not prevent the ones that survived and those that emerged, like Google, Facebook, Twitter and many others, from improving on the failed efforts and from changing the system. The innovators will get the appropriate help from insiders and leverage their expertise, and there will be many good doctors willing to assist in this transition.

This evolution from an entirely human-based to an increasingly automated healthcare system will take time, and there are many ways in which it can happen, but with a little bit of luck, it won't take as long as people think. The move will happen in fits and starts along different pathways with many course corrections, steps backward and mistakes as we figure out the best approach forward. It's impossible to predict how this will ultimately happen, and while I may not predict the pathway, it does not mean it will not eventually happen. It may be the case that all significant efforts will have to be catalyzed by outsiders. The

healthcare system might actually start responding to these threats from the inside and change as a result. Maybe we will start seeing disruption at the fringes along slippery but shallow slopes. The transition could start as a hundred small changes in different areas of medicine and in different ways and end with an overhaul of healthcare that takes place over a couple decades. During all this, many or most in this effort will fail, but a few will succeed and change the world. For lack of sufficient technological progress, institutional drag or backlash, we might see delays of a decade or more. For those of us who support entrepreneurs and companies that help create this change, most investments will be losses but more money will be made than lost through the few successes. None of us knows for sure how this space will turn out, but there's a huge opportunity for technologists, entrepreneurs and other forward-thinkers to reduce healthcare expenditures and improve patient care at the very same time.

There are a lot of improbable sounding possibilities for how digital health technology may impact healthcare. Though any particular one is unlikely to become a reality, it will be some improbability that will determine the future of health care as it is driven, molded and transformed by digital health technologies. Some improbable scenario today will become tomorrow's reality. These are not absolutes but rather "more true than not" speculations. We just have to imagine what might be possible!

Imagine the Impossible

In 1985, it was hard to imagine a personal computer (PC) in every home. In 1990, it was hard to imagine grandmother using email. In 1995, it was hard to imagine a pervasive consumer Internet. In 2000, it was hard to imagine Google being ubiquitous. In 2005, it was hard to imagine that the mobile phone would be universal or that it would function primarily as something other than a device to talk to other people. In 2010, it was hard to imagine Facebook's true impact. Maybe in 2015, it is be hard to imagine what digital health will do to the global healthcare system. By 2025 or 2035, what will digital health technologies do to medicine?

While I will certainly be wrong in my specific

predictions or timelines of technology evolution, strongly suspect that I will be directionally right. Again, these are not absolutes but rather "more likely to be more-true-than-not speculations.

Change is constant. Optimists (like me) are usually wrong about the degree of change in the short run, but even they underestimate the exponential change and improvement over the longer term as growth begets technology. The experts who scoffed at the personal computer as a toy; the Internet as minor compared to AT&T's global telecommunication vision (hence AT&T was sold to Cingular wireless for a song for its failure to adapt and only the brand really remains) ; Google as yet another search engine; Facebook as no more than a teenage phenomenon, will similarly scoff at, dismiss or minimize the impact of digital health technologies. Regardless, the future will likely happen, and it will be driven by even more surreal and naïve visions of entrepreneurial energy and passion for a vision.

Imagine what will be possible in 15-30 years or so when this technology will be in its fifth or even tenth evolution! Even if every doctor and patient has not adopted these technologies yet, the thought leaders and early adopters will, and the future of medicine will be obvious. And it will be altogether too complex for today's doctors to handle. Their roles will change though it is hard to postulate "how"! I suspect what they do today will change by 80% or more, but new unforeseeable roles may emerge as happens often with technological progress.

Imagine a future where medical professionals are mobile, since they no longer need to be tied to physical hospitals that house expensive equipment. While hospitals won't go away entirely, they may become less common or smaller in scale since inexpensive sensors generating lots of data can supplant some of the things hospitals offer today. The patient monitoring and assist functions of hospitals may change in response to technological development. By comparison, many parts of the developing world skipped the need for landline telephones and went straight to the proliferation of mobile phones. This could be a way for developing nations to leapfrog developed nations in medical care.

Imagine how technology also will impact health insurance companies. Health insurance could give people financial rewards for improving their health and further making the consumer the CEO of his own health. The healthcare consumer may choose to be part of this trend or not!

Over time, more actionable data will improve patient health while reducing costs and improving the quality of life for billions of people. As the population gets healthier, the conversation will ultimately shift from disease management to health optimization. Healthier people can work longer, which is important for the ageing population in developed countries in order to grow and maintain GDP. There are many second order effects as technology increases the health of the population. The trend is in our favor as technology constantly becomes cheaper and more accessible. Recent trends in healthcare technology notwithstanding, digital technologies do increase competition because of their short innovation cycle compared to the long cycle for biology-based technologies in healthcare.

Possibly the single most interesting aspect of this technology evolution may be to make the consumer the CEO of his or her own health. This goes much further and is much more intriguing than a rudimentary linear extrapolation of the Quantified Self movement, which would imply an exponential increase in the number of sensors and inputs will lead to a much better health outcome. Smart software and hardware systems may offer consumers much more sophisticated advice and choices (sometimes in conjunction with new types of care providers). Making the consumer the CEO of his/her own health will result as a function of incentives and easily actionable “wellcare” information.

My speculation is that this technology and system adaptation will allow for a 5x5 improvement across healthcare. There will be a 5x reduction in doctor work, enabling doctors to focus on patient care, research, or new things not yet imagined. There will be a 5x increase in research,

as medicine shifts to being more data-driven and all patients’ data is online (in a secure, anonymized way) to be used for studies. There will be a 5x lower error rate since the majority of conditions will have computer systems guiding them towards the correct results by default. There will be 5x faster diagnosis due to the advent of consumer-driven systems (diagnosis from your home) as well as fast and sophisticated hardware and software systems replacing the current back-and-forth and communication between all the various entities of the healthcare system (not just your physician, but your nurses, PCP’s, insurers, etc...). And lastly, speculate that healthcare will be 5x cheaper, although this will require a very conscious effort in ensuring we do not promote technological systems that just add to and increase the cost of care.

The net result of all the increased data, sensors, and medical insights discussed above will manifest itself globally when your average consumer’s (person’s) default choice will be interfacing with healthcare systems (wearables, doctors, hospitals, v7 technologies) in a way that benefits their long-term health. They will have all the data they need, but won’t have to analyze it (unless they want to!) ... the preferred action will organically emerge. Technology will not force these actions those but instead will nudge by informing patients of their choices and resulting consequences (which is why great design and great understanding of human motivation and behavior change is vital). So as CEO, the consumer will continue to have full choice over what they do, allowing for continued experimentation that will then give additional data to feed in our algorithms and systems. This process will be self-reinforcing, leading to better healthcare over time. And the consumer will drive healthcare innovation (with data compounding the positive impact of this new world).

And finally, relying on systems will not absolve us of responsibility for the difficult choices society faces around ethical choices!