Data-Driven, Automated Healthcare: Its Promise, Challenges, and Regulatory Playing Field

Christina Ren
Stanford University

Abstract

The future of data-driven, automated healthcare holds both promise and uncertainty. The regulatory playing field for this emerging industry has yet to pan out – to date, the US Food and Drug Administration (FDA) still considers health information technology (HIT) software as a medical device. Thus far, the FDA does not scrutinize HIT software as much as traditional medical devices. Nonetheless, the FDA has been criticized for using an outdated regulatory framework to govern modern health technologies. As such technologies with the potential to affect hundreds of millions lives become increasingly prevalent and salient, will the FDA be able to 1) handle the onslaught of new companies and regulatory burden 2) mitigate risks to public health while promoting innovation? This paper will discuss the potential of data-driven, automated healthcare, explain current regulations, explore newly proposed legislations, and finally provide recommendations for a smarter regulatory environment for AI-driven healthcare. These recommendations suggest on focusing on field testing and post-market safety instead of approval, providing regulatory understanding towards learning pains in medical AI industry, changing quality metrics to include diagnostic metrics, encouraging patients to embrace medical AI, and avoiding mandate and incentive policies within regulation.

It is inevitable that, in the future, the majority of physicians' diagnostic, prescription and monitoring, which over time may approach 80-percent of total doctor time spent on medicine, will be replaced by smart hardware, software, and testing.

– Vinod Khosla (Khosla, 2014, p. 1)

Such a bold claim was made by co-founder of Sun Microsystems and prominent venture capitalist, Vinod Khosla. Since 2012, he has predicted that data-driven, automated systems can, should, and will replace what physicians currently do. While this may irk, challenge, or even amuse many who find such comments dubious, Khosla's claims should not be dismissed. The staggering growth in computing power and combinatorial power of innovation can turn improbable outcomes – including artificial intelligence (AI) in medicine – into tomorrow's reality.

We have seen incredible advances in automation across industries like aviation, and finance. In healthcare, companies like Tencent and Lumiata are striving towards smart diagnostics and assisted clinical decisions. Using data points from journal articles, public data sets, physician notes, and patient-gathered information, Lumiata claims to have built the world's first Medical Graph that seeks to "mimic multi-dimensional human reasoning." Xiaoyi, a Chinese robot featuring AI, recently became the first robot to pass the national medical licensing examination and is set to assist physicians with diagnosis and treatment this year (Si 2017).

This paper will not focus on a cost benefit analysis of such a future (though it will touch upon its potential). Instead, given this foreseeable transition towards a future that turns vast amounts of data into personalized, real-time knowledge, the central question becomes: how will all of this be regulated?

This paper will first discuss the potential of data-driven, automated healthcare and provide examples of what it currently is and may become (Section I). Then, the paper will explain current regulations and the rationale for its current limitations (Section II), explore newly proposed legislations (Section III), and finally provide recommendations for implementing a smarter regulatory environment for AI-driven healthcare (Section IV).

I. Data-Driven, Automated Healthcare

Current practice of medicine is error-ridden

Even with extensive training, there is often substantial variability in the conclusions made by practitioners for the same patient, or for patients with similar conditions. In one study, cardiologists were given the same patient information and half recommended cardiac surgery while the other half did not. Two years later, 40% of cardiologists disagreed with their previous assessments even though they were given the exact same data (Eddy, 1990). For pathologists reviewing a patient's tissue sample, agreement in diagnosing breast cancer cases was around 75%, but for

diagnostically challenging cases, could be as low as 48% (Elmore *et al.*, 2015).

The lack of consensus is not surprising considering how much information needs to be reviewed. A pathologist needs to scrutinize several slides per patient, with each slide containing over ten gigapixels. That is a lot of data to cover in a limited amount of time.

These mistakes can be costly, however. Diagnostic errors contribute to over 40,000 deaths per year. Autopsy studies have shown that in 8% of patients, the diagnostic error was serious enough that it may have caused or directly contributed to the patient's death. Had doctors been aware of the proper diagnosis, treatment likely would have been different (Winters *et al.*, 2012).

Applications of medical AI to address these problems

AI may help address some of healthcare's pressing issues. By using vast amounts of data and algorithms, technologies can identify patterns that were either previously unidentified or labor intensive and/or error-prone for humans to generate. Promising applications range from disease detection, data generation and infrastructure, and health management.

Today, risk for cardiovascular disease is predicted using standard medical guidelines from the American Heart Association (AHA) based on eight risk factors including age, blood pressure, and cholesterol levels. A recent study found that machine learning programs created their own criteria that correctly predicted 7.6% more cardiovascular events than the standard method with fewer false positives (Weng *et al.*, 2017). Improved prediction rates lead to preventive treatment for those at risk. In a test sample of around 83,000, that translates to an additional 355 patients whose lives could have been saved. Many of the risks factors identified by the algorithm are not included in the AHA guidelines, such as severe mental illness and taking oral corticosteroids. Another study found that models for retinal images are now also able to predict various cardiovascular risk factors such as age, smoking status, blood pressure, and major adverse cardiac events like heart attacks based on physical characteristics of the eye (Poplin *et al.*, 2018).

In cancer detection, Google released its research on tumor identification through pathology images last year. Its automated approach detected 92% of the tumors, compared to a pathologist's 73% sensitivity (Liu *et al.*, 2017).

Work on healthcare AI spans more than just interesting research and training models. In China, technology giants like Tencent and Alibaba have already pushed several AI-driven medical products into hospitals. In 2017, Tencent launched the AI Medical Innovation System (AIMIS, or miying in Chinese) which is used for speedy cancer and lung disease detection in more than 100 hospitals across China (Wee and Mozur, 2018). Reported accuracy rates for this technology are 90% for early

diagnosis of esophageal cancer, 97% for diabetic retinopathy, and 95% for lunch sarcoidosis.

Tencent is also looking to use technology to improve access to quality healthcare in rural areas. WeDoctor is a platform that trains village doctors to use equipment and assist with diagnostics. It also automatically uploads medical records and has remote consultation features to help build smart clinics in rural regions.

Of course, medicine and human biology is complex; AI cannot solve every problem. Marty Kohn, former chief medical scientist at IBM working on IBM Watson Health recognizes there has been more hype than impact: "in certain niches, AI is here and has been for years. But it's not happening at scale. And it hasn't yet helped large numbers of patients." (Duncan 2017)

Automated systems are also somewhat of a black box, and even the researchers themselves don't always know why some algorithms work and others don't. Efficient regulatory frameworks are needed to guard against unintended consequences while allowing for innovation in this space.

II. Current Regulations

Under the Federal, Food, Drug, and Cosmetic Act, HIT software and diagnostic tests are considered medical devices. By definition, a medical device is an "instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article… intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease" (Center for Devices and Radiological Health). In 1976, the Medical Device Amendments required a "reasonable assurance of safety and effectiveness" based on valid scientific evidence that "adequately demonstrates the absence of unreasonable risk of illness or injury." Over 35 years have passed, and the regulations for medical devices have remained mostly unchanged. Because of the Federal, Food, Drug, and Cosmetic Act, the FDA currently mandates that medical software comply with regulations that apply to more traditional medical devices. Thus, understanding how medical devices are regulated is the closest thing to understanding how data-driven, automated healthcare systems are currently regulated.

General Overview

Medical device establishments are required to register and list their devices with the FDA. The FDA classifies devices as class I, II, or III based on the safety of the device as determined by the agency. At minimum, all manufacturers for medical devices are required to report any "device-related deaths and serious injuries, and malfunctions that may… result in death or serious injury" (Shuren, 2010). Pre-market approval (PMA) and post-market surveillance of medium to high-risk devices (class II and III) may also apply. The PMA process can be a lengthy one and very costly; it typically involves data from clinical trials proving safety

and effectiveness of the device, as well as manufacturing information on the device.

510(k) Clearance Pathway

The FDA allows moderate risk medical devices to come to market through a mechanism called a 510(k) clearance. If a medical device is "substantially equivalent" to a device already on the market ("predicate devices"), then the manufacturer(s) can essentially bypass providing clinical data for their own medical device. Instead, they can use existing information from predicate devices. They are also not subject to as stringent of post-market oversight as devices on the regular PMA track.

De novo 510(k) process

If the FDA rejects the 510(k) submission and establishes that the device is "not substantially equivalent" to the predicate device the manufacturers referenced, the device is automatically classified as a Class III device, the highest risk classification. To bump the classification down to Class I or II and avoid submitting a PMA, manufacturers can apply through the de novo 510(k) process (Center for Devices and Radiological Health). The de novo pathway is typically for devices that represent moderate to low risk (class I or II) and is a novel device with no predicate devices. While more troublesome and lengthy than the 510(k) clearance (especially if following a rejected 510(k) clearance), the de novo pathway is comparatively better than submitting a PMA. Unlike the PMA, the de novo pathway does not require extensive clinical trial data to support its safety and effectiveness, only that data that demonstrates the general and support controls "support a classification of Class I or Class II". In February 2015 for example, 23andMe approved a carrier screening test for Bloom Syndrome through the de novo pathway. ("Evaluation of Automatic Class III…")

Expedited Access Pathway (EAP)

The FDA has been piloting a program in recent years that would accelerate medical device approvals for devices that treat or diagnosis "life threatening or irreversibly debilitating diseases or conditions." The idea is to reduce PMA requirements while increasing post-market requirements to provide patients more timely access to new medical devices (Center for Devices and Radiological Health). Applications received through EAP will receive priority review and the FDA must determine device classification within 120 days.

Limitations of Medical Device Regulations on AI-Driven Healthcare

It is clear that the language used by the FDA for medical devices was written for more traditional devices in mind, such as "surgical lasers, wheelchairs, sutures, pacemakers, vascular grafts, intraocular lenses, and orthopedic pins." Yet because there is no policy specifically for medical software that analyzes data to help make clinical decisions, such software would still classify as medical devices.

Jeffery Shuren, Director of FDA's Center for Devices and Radiological Health stated that, "to date, FDA has largely refrained from enforcing our regulatory requirements with respect to health information technology devices" (Shuren 2010). For HIT vendors, registering, obtaining premarket approval, and reporting adverse events is voluntary. But clinical decision-making software will likely span beyond what is currently acknowledged as "HIT". It will involve integrating analysis of published papers, past outcomes for thousands of patients with similar medical profiles, and (perhaps most contentious of all in terms of public safety) patient-specific information including risk factors, whole genome sequencing, meta-genomic analysis of the gut microbiome, biomarkers from blood samples, patient-gathered data from wearables, and much more.

Already, a few companies have run into regulatory friction with the FDA. On November 2013, FDA ordered 23andMe to take its direct-to-consumer genetic tests off the market due to "serious concerns… [raised] if test results are not adequately understood by patients or if incorrect test results are reported" (Department of Health and Human Services). The FDA stated that 23andMe blatantly ignored and did not provide additional information necessary to complete de novo 510(k) approval even after several months of submission and negotiation. During that time instead, 23andMe focused its efforts on new marketing campaigns to expand consumer base without market authorization from the FDA. While 23andMe has recently gotten approval on a limited number of genetic tests including carrier screening tests and certain genetic risks like select BRCA mutations, the approach the FDA has taken thus far – that is, approve each genetic test individually – is not promising for clinical decision-making software that aims to serve as a comprehensive diagnostic tool for a spectrum of different human disease conditions.

In an industry that will become increasingly more convergent, the FDA will find it increasingly challenging to fit technologies neatly into a single regulatory pathway or center. We have already begun to see this in novel therapies—stents (medical devices) with a drug component, stem cells (biologics) as delivery systems for drugs, and smart drugs that allow for continuous monitoring in the body (medical devices). Currently the Office of Combination Products determines which center leads the product review, but this process can significantly delay the review process (Lamar, 2015). Forcing a hybrid product into one of three categories as the primary label is like shoving a square peg on a round hole – eventually the agency will run into some major backlog.

In addition, current expedited access pathways for medical device approval are reserved for life threatening or irreversibly debilitating diseases only. For automated systems that could potentially encompass a wide variety of human states of health—including routine problems or complex chronic conditions—the risk-benefit profile may not be as stark. This means an accurate pathway for earlier diagnosis of cancer, stroke,

and heart attacks would be considered something that would benefit a life threatening or irreversibly debilitating disease. However, a diagnosis of a child's ear infection, mild allergic reaction to an unidentifiable source, or response to a new treatment plan for type II diabetes would not be considered. This is quite lopsided, considering AI systems will likely deal with more routine cases as opposed to critical conditions. Such expedited pathways would thus likely be ill suited for clinical decision-making software.

Confusion within the Industry

Existing traditional regulatory structures are ill suited for the iterative design and validation of data-driven, automated medical systems. In the past year, the FDA has openly recognized this and has outlined several areas to be addressed, including possibly removing clinical decision support software from their jurisdiction all together ("Digital Health Innovation…") However, these efforts are still in the early stages in the form of draft guidelines and pilot programs that have yet to be implemented. Present day companies remain in limbo, forced to navigate ambiguous regulatory waters.

In the absence of defined regulatory pathways, companies and legal consults are expected to exercise judgment how likely their product will be subject to regulations, and if so, the extent of regulation. The FDA has been taking a more hands-off approach compared to its approach in 2013 (when it ordered 23andMe to recall its product), but much remains uncertain.

III. Proposed Legislations

21st Century Cures Act sparks regulatory overhaul

The clinical decision-making industry is still in its infancy, but some of the key players have already taken regulatory matters in their own hands. IBM has led the way in influencing how the FDA should deal with technologies like Watson Health. In 2014 IBM spent $5 million lobbying in Washington, deploying two lobbying firms and its own lobbyists to have clinical decision support regulation be included in a bill called the 21st Century Cures Act (Edney, 2015). Republican chairman of the House Energy and Commerce Committee Fred Upton spearheaded the bill, designed to "accelerate the discovery, development, and delivery" of medical therapies (H.R.6, 114th Cong.). The law was enacted by the US Congress in December 2016. Admittedly, the act is quite the Christmas tree bill, because there is a section that specifically deals with software: Sensible Oversight for Technology Which Advances Regulatory Efficiency (SOFTWARE) act.

The SOFTWARE act defines health software in a very broad sense. Health software includes potential uses such as administrative and operational support as well as "use to analyze information to provide patient-specific recommended options to consider in the prevention,

diagnosis, treatment, cure, or mitigation of a particular disease or condition" (H.R.6, 114th Cong, p. 221). Within a year and a half of passing the law, the Secretary of Health and Human Service will have consulted "external stakeholders" before issuing regulations, administrative orders, and guidance. The Secretary may implement a new framework for the regulation of software that overrides previous regulations. Issues to be addressed include classification, standards for development, validation, verification, labeling requirements, and post-marketing requirements for software. Perhaps most importantly, health software would be excluded from the definition of a medical device (H.R.6, 114th Cong, p. 225). The FDA will retain "primary jurisdiction" in regulation health software.

The bill highlights several trends in the modern regulatory framework of automated healthcare. First, the bill substantially broadened its definition of health software. Previous versions made a distinction between medical and health software. Health and medical software were similar in function, except medical software provides physicians with recommendations of treatment or course of action "without the need for such professionals to perform additional interpretation." Medical software would be more closely regulated by the FDA, while health software would not. By merging the two definitions into one, the government is essentially recognizing health software will likely be a combination of the two definitions.

Second, the bill encourages progress in the field—it seeks to offer flexibility to an embryonic, evolving market while protecting patients from harm. The senators who wrote this recognize that something should be done, so they are mandating the FDA to come up with some new standards of approval.

FDA's response: Digital Health Innovation Action Plan

In response to the enactment of the 21st Century Cures Act, the FDA has released the Digital Health Innovation Action Plan in 2017. The goal is to redesign policies to reflect the needs of digital health technologies. Action items include issuing draft guidelines for how clinical decision support software will be regulated, how products with both software and medical device functions will be categorized, and whether software changes need to submit a 501(k) clearance pathway (traditionally used for devices that have similar devices already approved.) The agency has also been working with the International Medical Device Regulators Forum (IMDRF) to create new risk categorization frameworks for software as a medical device ("Digital Health Innovation…")

One of the biggest components of the Digital Health Innovation Action Plan is the Digital Health Software Precertification (Pre-Cert) Pilot Program. The FDA is looking to streamline health technology oversight by preapproving software developers instead of their product(s). That way, developers can market lower risk software as a medical device without

premarket review of individual products. Developers are evaluated based on "patient safety, product quality, clinical responsibility, cybersecurity responsibility, and proactive culture" ("Software Precertification Program…") In lieu of stringent pre-market review, products are evaluated by patient feedback and real world performance data. The FDA has selected nine companies – including Apple, Samsung, and Verily –to participate in the pilot program launching by the end of 2018. There are plans to expand this program to more organizations in 2019.

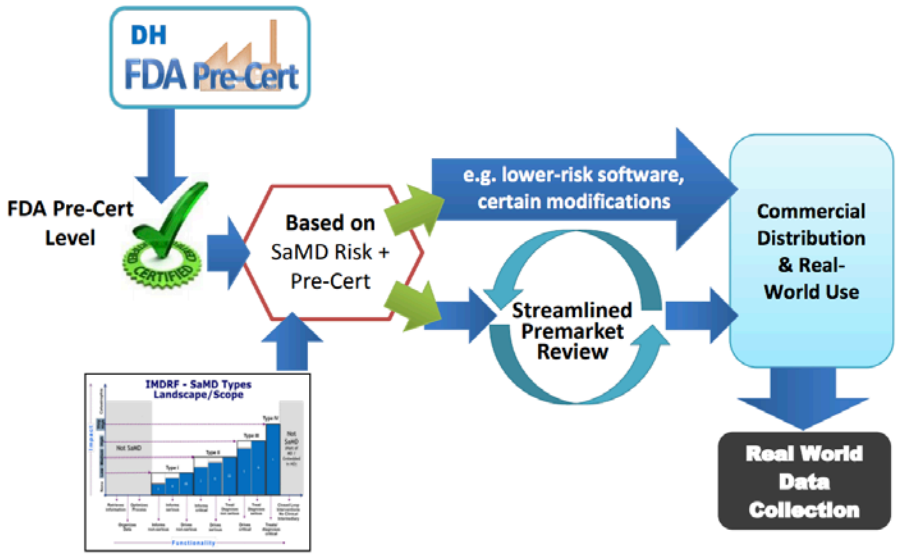

Figure 1. High level concept of the reimagined approach using FDA Pre-Cert for Software (Image credit: FDA)

IV. Moving Forward

Regulation in this field is in a bit of a catch-22. Developers want to know how exactly they will be regulated now and in the future. In order to impose effective regulations and minimize confusion, the government should offer specific, up-to-date guidelines instead of outdated regulations. In fact, the industry may actually suffer in the long run if it is regulated inappropriately. The latest Pre-Cert Program is a good start in ensuring safety without inhibiting developers' ability to innovate.

Thus far, the FDA has taken the "enforcement discretion" approach, allowing companies to do their own risk assessment on themselves. This is not necessarily bad; being overly prescriptive so early on could create a bigger mess in the long-term. Thus, this section will not offer detailed recommended solutions but instead discuss high-level recommendations, foreseeable challenges, and how smart regulations will play an important role in shaping data-driven, automated healthcare systems.

A. Focus on field testing and post-market safety instead of premarket approval

Premarket approval will likely neither be very feasible nor useful when it comes to dealing with medical AI systems. Let's say to validate improved

patient outcomes 200 people are enlisted in a study—half of them receives care with clinical decision-making support and the other does not. By the time this expensive and time-consuming clinical trial ends, so many things will have changed – the system updated its database, revised its algorithm, etc. Unlike drugs, biologics, and medical devices, the system is not a static, finished product.

A smarter solution is not to focus on extensive premarket approval but instead run performance on simulations of patient cases before market release and then carefully monitor outcomes post-release. For general diagnosis and treatment systems, the FDA can run a series of 1,000 hypothetical patients through the system and compare patient outcomes to the same fictional patients diagnosed with typical physicians. If the health software produced equivalent or better patient outcomes compared to the physicians', the software can be released in the market. Afterwards, post-market review will evaluate how the software is doing with real patients. Example questions in post-market review could be: How do doctors use this technology and is it improving their diagnostic accuracy? Have there been instances in which patient health and safety were compromised and if so, was that independent from professional interpretation? What are areas that could be improved? The documentation of treatment and patient outcomes in real world situations should be ultimately what decides if the health software is safe and effective.

B. Put AI-driven healthcare on assist, learn, and amplify mode

As the medical AI movement progresses, there will undeniably be roadblocks and challenges. Many possibilities of data usage will be pursued, and many will most likely fail. But the technologies that do succeed will determine the future of healthcare driven by technology. While the industry is still in its infancy, there will undeniably be cases with less than optimal results.

The earliest versions of "Dr. Algorithm" will be, as Vinod Kholsa puts it, "toddler computer systems" in training (Kholsa, 2014, p. 3). IBM Watson Health and other similar systems will likely generate some laughable outcomes and be the butt of jokes for many. It will be "clumsy and underwhelming". But like the evolution of mobile phone technology, each new version will be increasingly more sophisticated.

So how should we deal with this for early patient users? To mitigate potential risk, there will be a period of time when clinical decision-making software will strictly be in "assist, learn and amplify" mode. Until it has been sufficiently tested in a wide variety of real scenarios, the software should be accompanied by professional guidance, supervision, and interpretation. While health software can certainly give recommendations and point out things that physicians may have otherwise missed, initially physicians will still call the shots and be responsible for the patients.

C. Change the way hospitals report quality measures and emphasize diagnostic accuracy

In 1993, MIT Professor Erik coined the term "productivity paradox" to describe how productivity often appears to decline after the introduction of information technology (Brynjolfsson, 1993). One of the reasons is that adaptive forces have not yet caught up to the technology. The other reason is that productivity is not being measured the right way. Music streaming services like Spotify, for example, allow greater accessibility to music. But if "productivity" were measured through record sales, you'd think people are listening to less music now. The same applies to clinical decision-making software—current publically reported quality measures do not often report on diagnostic accuracy. There is therefore no incentive to use tools that could aid with improved diagnosis. And for practitioners adopting new systems, they might not see improved outcomes using the current quality measures.

This paradox is important as we begin to adopt medical AI systems. If outcome measurements do not demonstrate that such systems are beneficial, then future attempts to expand such technologies might be jeopardized because of premature conclusions made under current metrics.

D. Encourage physicians to embrace assistance

This recommendation is less regulation-focused and argues about changing society's perception of a physician's role. Even if data supports the conclusion that clinical decision-making software improves patient outcomes, cultural acceptance and cognitive bias of physicians are potential roadblocks to successful integration. Doctors have traditionally held the position as gatekeepers of medical knowledge. The adage "the doctor knows best" puts healthcare professionals as the authoritative forefront of diagnosis and treatment. For those that pride themselves as highly competent doctors, they may view using tools as a cop out. Robert Wachter, MD., author of *The Digital Doctor* and prominent academic physician at UCSF, admits that he had felt the need to keep such a façade as a resident back in the 1980s: "more than once I found myself stumped about a diagnosis or treatment, but I was too embarrassed to admit that to my patient. At times I even told little lies—" Excuse me, my pager just went off"— before I left the room to look something up" (Wachter, 2015, p. 108.) While this mentality is starting to change amongst the younger generation of physicians, it's important for professionals, teaching institutions, and society to change this perception.

E. Avoid mandates and incentive policies at all costs

Instead of restrictive policies, another side of the extreme would be to issue mandates and incentive policies. This may seem strange and counterintuitive now, as most people are concerned about the safety of medical software. But hypothetically in the future, provided that medical software has proven itself to be safe and effective, the government may think it's a good idea to have every physician adopt AI-drive assists to make care more consistent and personalized. This is faulty reasoning and can produce disastrous results. The government's decision to require

physicians to adopt electronic health records in their practice through HITECH is a good example of what not to do. While HITECH pushed many to transition to electronic health records, it has also been criticized for forcing physicians to adopt clunky, early versions of EHR software and disrupting physician workflow. Like many other industries transitioning to automated systems, it's best to let medical AI run its course in the market.

The early days of clinical decision-making software will be on a demand basis. Patients who seek more technology enabled medical care will choose doctors who use such supportive and patient-geared technologies. Patients who prefer a more traditional means of interaction with a doctor will not. Unlike the transition from paper to electronic medical records, the type of care patients wish to receive should, at least initially, be their preference.

A similar situation will occur on the side of healthcare professionals. Most will be skeptical, even oppose it. When I once asked a room full of medical students from Asia if they think healthcare will be dominated by machine intelligence or remain more or less the same thirty years down the road, they overwhelmingly were on the latter side of the spectrum. But early adopters will be among the first to change their practice.

V. Concluding Remarks

The future of data-driven, automated healthcare holds both promise and uncertainty. It is promising because of its ability to minimize labor intensive, error-prone aspects of medicine and drive new health insights. It is uncertain because the regulatory playing field for this emerging industry has yet to pan out. Current regulations under medical devices are ill suited for medical AI, so new regulatory standards will need be in place. Minimal regulations should be implemented to allow for innovation, though performance testing and post-market surveillance should be done to ensure safety for patients. Additional considerations to ensure a smooth, effective adoption of medical AI will involve more than just reasonable regulations but instead, mobilizing and engaging people to embrace a future with a more data-driven, automated approach to healthcare.

To many, Vinod Kholsa's prediction conjures a science fiction reality of a cold, apathetic robot doctor. Human interaction, physician intuition and touch are instead replaced with sensors, automated voices, and screens. This is an understandable instinctual reaction, but also a limiting and false notion of how medical AI will likely be integrated in healthcare systems. Physicians will not be obsolete, because the human element of care cannot be replaced. As Atul Gawande puts it, "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. Maybe machines can decide, but only doctors can heal." (Gawande 1998)

References

Alper, Brian S. et al. (2004). How Much Effort Is Needed to Keep up with the Literature Relevant for Primary Care? *Journal of the Medical Library Association* 92.4: 429–437. Print.

Carson, Erin. (2015) IBM Watson Health Computes a Pair of New Solutions to Improve Healthcare Data and Security *TechRepublic*. Retrieved from https://www.techrepublic.com/article/ibm-watson-health-computes-a-pair-of-new-solutions-to-improve-healthcare-data-and-security/

Center for Devices and Radiological Health. (n.d.). Evaluation of Automatic Class III Designation (De Novo) Summaries. Retrieved from https://www.fda.gov/AboutFDA/CentersOffices/OfficeofMedicalProductsandTobacco/CDRH/CDRHTransparency/ucm232269.htm

Center for Devices and Radiological Health. (n.d.). What is a medical device? Retrieved from https://www.fda.gov/AboutFDA/Transparency/Basics/ucm211822.htm

Department of Health and Human Services. (2013, November 22). Personal Genome Service (PGS) [Letter to Ann Wojcicki]. Retrieved from https://www.fda.gov/ICECI/EnforcementActions/WarningLetters/2013/ucm376296.htm

Developing a Software Precertification Program: A Working Model (2018) Retrieved from https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/DigitalHealthPreCertProgram/UCM605685.pdf

Digital Health Innovation Action Plan. (2017) Retrieved from https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf

Duncan, David. (2017) Can AI Keep You Healthy? *MIT Technology Review*. Retrieved from https://www.technologyreview.com/s/608987/how-ai-will-keep-you-healthy/

E. Brynjolfsson. (1993). The Productivity Paradox of Information Technology. *Communications of the ACM* 36: 66– 77.

Eddy, David. (1990). The Challenge. *JAMA.* 1990;263(2):287-290.

Edney, Anna. (2015). This Medical Supercomputer Isn't a Pacemaker, IBM Tells Congress. *Bloomberg,* Retrieved from www.bloomberg.com/news/articles/2015-01-29/this-medical-supercomputer-isn-t-a-pacemaker-ibm-tells-congress>.

Elmore, Joann, *et al.* (2015) Diagnostic Concordance Among Pathologists Interpreting Breast Biopsy Specimens. *JAMA.* 2015;313(11):1122–1132.

Evaluation of Automatic Class III Designation for the 23andMe Personal Genome

Service Carrier Screening Test for Bloom Syndrome. (2015, February). Retrieved from https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN140044.pdf

Eyers, Sally et al. (2010). The Effect on Mortality of Antipyretics in the Treatment of Influenza Infection: Systematic Review and Meta-Analyis. *Journal of the Royal Society of Medicine* 103.10 (2010): 403–411.

Graber ML;Mathew A. (2008) Performance of a web-based clinical diagnosis support system for internists. *J Gen Intern Med.* 23(suppl 1): 37-40

Gawande, Atul. (1998). No Mistake. *The New Yorker*. Retrieved from https://www.newyorker.com/magazine/1998/03/30/no-mistake

Hernandez, Daniela. (2014). Artificial Intelligence Is Now Telling Doctors How to Treat You. Wired Magazine. Retrieved from www.wired.com/2014/06/ai healthcare

H. R. 6, 114th Cong. (2015) 21st Century Cures Act. (enacted). Retrieved from https://www.congress.gov/114/bills/hr6/BILLS-114hr6rfs.pdf

Khosla, Vinod. (2014). 20-Percent Doctor Included: Speculations and Musings of a Technology Optimist." 1–51. Print.

Lamar, A. and Burr, R. (2015). Innovation for Healthier Americans. Retrieved from https://www.help.senate.gov/imo/media/Innovation_for_Healthier_Americans.pdf

Liu, Yun, *et al.* (2017) Detecting Cancer Metastases on Gigapixel Pathology Images. *Google AI.*

McGlynn EA, *et al.* (2003). The quality of health care delivered to adults in the United States. *New England Journal of Medicine* 348:2635-2645.

Poplin, Ryan, *et al.* (2018). Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering* 2, 158–164.

Raths, David. (2015). Regulators Can't Keep Pace with Health-Care Technology Innovation. *Government Technology*. Retrieved from http://www.govtech.com/health/Digital-Health-Dilemma-Regulators-Struggle-to-Keep-Pace-with-Health-Care-Technology-Innovation.html

Shuren, Jeffrey. (2010). Testimony of Jeffrey Shuren, Director of FDA's Center for Devices and Radiological Health. Retrieved from http://www.cchfreedom.org/pr/Health%20IT%20Deaths%20-%20FDA%20jeffrey%20Shuren.pdf

Si, Ma & Yu, Cheng. (2017). Chinese robot becomes world's first machine to pass medical exam. *China Daily.* Retrieved from

http://www.chinadaily.com.cn/bizchina/tech/2017-11/10/content_34362656.htm

Thompson & M. Jason Brooke. (2011). FDA's Approach to Clinical Decision Support Software: A Brief Summary. *FierceHealthIT*. Retrieved from http://www.fiercehealthit.com/special-reports/fdas-approach-clinical-decision-support-software-brief-summary

Wachter, Robert (2015). The Digital Doctor: Hope, Hype, and Harm at the Dawn of Medicine's Computer Age. *McGraw-Hill Education*.

Weng, Stephen et al. (2017) Can machine-learning improve cardiovascular risk prediction using routine clinical data? PLOS ONE 12(4): e0174944.

Wee, Sui-Lee & Mozur, Paul. (2018) Amazon Wants to Disrupt Health Care in America. In China, Tech Giants Already Have. *NY Times*. Retrieved from https://www.nytimes.com/2018/01/31/technology/amazon-china-health-care-ai.html

Winters B et al. (2012). Diagnostic errors in the intensive care unit: a systematic review of autopsy studies. *BMJ Qual Saf*. 2012; 21: 894-902.