"20% doctor included": speculations & musings of a technology optimist

by Vinod Khosla

In recent talks at the University of California, San Francisco and other venues, I laid out my thesis on how mobile devices, big data, and artificial intelligence will disrupt healthcare. I believe it is inevitable that the majority of physicians' diagnostic and prescription work (physicians do other work too) will be replaced with smart hardware and software, and healthcare will become better, more consistent across physicians, and more scientific. The remaining 20% of physicians' work will be AMPLIFIED, giving them better capability. Given the importance of having clarity on what I hypothesized (my forecasts are directional guesses rather than precise predictions) about how healthcare might change, I'd like to make clarify my views and respond to some of the recent commentary.

Let me start with a summary:

1. Healthcare today is often really the "practice of medicine" rather than the "science of medicine". In the worst cases of the practice of medicine, doctors just take moderately educated shots in the dark when it comes to patient care. Physicians should be much more scientific and data-driven. That's hard for the average physician to pull off without technology, because of the increasing amount of data and research released every year. Next-generation medicine will be the scientific arrival at diagnostic and treatment conclusions and real testing of what's actually going on in your body. And, it will be much more personalized than your physician can provide. Data science will be key to this.

2. Technology makes up for human deficiencies and amplifies strengths – MDs and even other less trained medical professionals can do much more than they do now. By 2025 more data-driven, automated healthcare will displace up to 80% of physicians' diagnostic and prescription work. It will AMPLIFY physicians by arming them with more complete, synthesized, and up-to-date research data, all leading to better patient outcomes. Computers are much better than people at organizing and recalling information. They have larger and less corruptible memories, remember more complex information much more quickly and completely, and make far fewer mistakes than a hot shot MD from Harvard. Contrary to popular opinion, they're also better at integrating and balancing considerations of patient symptoms, history, demeanor, environmental factors, and population management guidelines than the average physician.

3. The healthcare transition will start incrementally and develop slowly in sophistication, much like a great MD who starts with seven years of med school and then spends a decade training with the best practitioners by watching, learning, and experiencing. Expect many laughing-stock attempts by "toddler MD computer systems" early in their evolution and learning. These systems will grow with the help of the top MDs and AMPLIFY them so everyone can have the best, most researched care, not the average, overburdened, and rushed doctor the average person gets today! In fact, the very best MDs will be an integral part of designing and building these advanced systems.

4. The systems of 10-15 years from now (which is 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 then would be like our carrying the multi-pound phones from 1987 (they were floor mounted cell phones with big handsets and heavy cords) in our pockets rather than iPhones. There will be point innovations that seem immaterial, but, when there are enough of them, they will integrate with each other and start to feel like a revolution. In the meantime, expect these early systems and tools to be the butt of jokes from many a writer or MD. Early printers, typically "dot matrix", did not exactly cut it for business correspondence, let alone replace traditional means. Early medical systems will be used in non-critical roles or under physician supervision. Eventually, this shift in how healthcare is delivered will allow for less money to be spent on capital equipment, cutting health care costs. It will allow us to provide care to those who don't have it now. And, it will prevent simple things from getting worse before being addressed.

5. The human element of care, provided by humane humans, not rushed, overloaded MDs, will still be around. These people may have MDs anyway, but they won't need ten or twenty years of medical school training. Or, when they have the training, they will become much better diagnosticians and caregivers. Beyond diagnosis and treatment, there are many things doctors do that won't be replaced.

6. The problem in healthcare is not with doctors, many of whom are accomplished, caring, honest, and compassionate providers. The problem is the incredible increase in complexity of the newly enabled data (some extrinsic), vast amounts of research, longitudinal health records, and histories without the self-reporting inaccuracies of the patient that allows for the much more integrative analysis that is now possible. The problem is also the misalignment of incentives in medicine, where organizations try to maximize revenue (extra surgeries anyone?) at the expense of optimizing care. This is why innovation will most likely happen from the outside. It is also important to realize that I refer to the "AVERAGE GLOBAL doctor" or healthcare concern, not every doctor or company. The standards of performance in some parts of the world and in some parts of this country are very different than those in the best metropolitan hospitals in the United States. I also note that 50% of MD's are below average though every doctor most people know is above average!

Practice vs. science

One of my principal concerns about healthcare today is that it's often really the "practice of medicine" rather than the "science of medicine". Take modern medicine's view towards fever as an example. For 150 years, doctors have routinely prescribed antipyretics (aspirin, acetaminophen, etc.) to help reduce fever, since fever was viewed as an inability of the body to regulate itself and therefore needed to be reduced aggressively in all cases. But in 2005, researchers at the University of Miami, Florida, ran a study of 82 intensive care patients, for whom protection from high temperatures was traditionally thought to be important. The patients were randomly assigned to receive antipyretics either if their temperature rose beyond 101.3°F (the "standard treatment") or only if their temperature reached 104°F. As the trial progressed, seven people getting the standard treatment died, while there was only one death in the group of patients allowed to have a higher fever. At this point, the trial was stopped because the team felt it would be unethical to allow any more patients to get the *standard treatment*. So when something

as basic as fever reduction is a hallmark of the "practice of medicine" and hasn't been challenged for 100+ years, we have to ask "what else might be practiced due to tradition rather than science?"

In the worst cases of the practice of medicine, the "average" doctors just take moderately educated shots in the dark when it comes to patient care. A diagnosis is partially informed by the patient's medical history (but often not really), partially informed by symptoms (but patients aren't very good at communicating what's really going on), and mostly informed by pharma advertising and the doctor's half-remembered lessons from medical school (which are laden with cognitive biases, recency biases, and other very human errors, besides potentially having been obsoleted by more recent research). Many times, 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 do you feel when a doctor keeps changing his mind about your disease over time? How do you feel if different doctors say different things about your disease? Today, this happens often. In some areas, psychiatry for example, doctors frequently disagree on diagnoses. 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. The DSM V uses a statistic called the "kappa" to measure the level of agreement between psychiatrists (ranging from 0 for no agreement and 1 for complete agreement). In research trials, the DSM V, which is set to be published in May 2013, 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". And often, there are errors of omission where a diagnosis is just missed entirely.

The net effect of all this is patient outcomes that are far inferior to and more expensive than what they should be. The current benchmarks of performance aren't good enough, and 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 an ICU in the US each year due to misdiagnosis, rivaling the number of deaths from breast cancer. Yet another study found that 'system-related factors', e.g. poor processes, teamwork, and communication, were involved in 65% of studied diagnostic error cases. 'Cognitive factors' were involved in 75%, with 'premature closure' (sticking with the initial diagnosis and ignoring reasonable alternatives, or, more fancifully termed, the "confirmation bias") as the most common cause. These types of diagnostic errors also add to rising healthcare expenditures, costing $300,000 per malpractice claim.

Physicians should be much more scientific and data-driven in providing patient care. That's hard to pull off without technology, because of the increasing amount of data and research released every year. For example, standard operating procedure involves giving the same drug to millions of people even though we know that each patient metabolizes medication at different rates and with different effectiveness. Many of us are even resistant to aspirin. Each person should be treated differently, but the average doctor can't handle the information required to do that. Nor does he have enough time or knowledge to do it. Healthcare needs to become much more about data-driven deduction and less about trial-and-error. Next-generation medicine will be the scientific arrival at diagnostic and treatment conclusions based on probabilities and real testing of what's actually going on in your body. And, it will be much

more personalized than your physician can provide. Systems will utilize more complex models of interactions within the human body and much more sensor data than a human MD could comprehend to suggest diagnosis. Thousands of baseline and disease mulit-ohmic (genomic, metabolomics, microbiomic, and other) data points, more integrative history, and demeanor will go into each diagnosis. Ever-improving dialog manager systems will help make data capture and exploration from patients more accurate and comprehensive. Data science will be key to this. In the end, this would reduce costs, reduce physician workloads, and improve patient care. Doctors will also be able to tailor their explanations to the health literacy level of patients using common-language terms and adapting the sophistication level using computerized dialog managers. These computerized managers will be patient, unlike your typical "doctor in a hurry" with the usual unfortunate case overload. This matters because, according to the Institute of Medicine, nearly 100M US adults have "limited health literacy skills" that are most likely to affect their health outcomes!

Replacing 80% of what doctors do?

At the heart of my view is that much of what physicians do (checkups, testing, diagnosis, prescription, behavior modification, etc.) can be done better by well-designed sensors, passive and active data collection, and analytics, without necessarily taking away from the human element of care (especially in the 2020's decade when this technology will be in its fifth or tenth evolution)! Let's take a brief look at a standard physical as an example. During a physical, a patient will get measured in multiple ways – from weight and blood pressure to pulse and respiration. A nurse, doctor, or other caregiver will spend 15–30 minutes to run through all these different routines. Every single one of these measurements could be done in real-time by the patient in more representative environments at home if he or she just had the right sensors hooked up wirelessly to a mobile device. All that data could be measured and transmitted to the doctor's office in less time and for less money than it takes to gas up the car and get to the hospital in the first place (leave aside sitting in the waiting room). ZocDoc* allows patients to check in for an appointment and provide basic information ahead of time. This type of form could easily be expanded to include vitals. And our "dialog manager" could ask many follow-on questions and probe other symptom possibilities, providing a more complete patient record WHILE reducing the amount of time the doctor has to spend with the patient, if any!). No doubt, the best doctors will likely do this better than our dialog manager, even for the next one or two decades. But, the system will be an immediate dramatic improvement compared to the average hurried and overloaded doctor (or worse, the developing world "no medical-school doctor" living within 50 miles of a rural patient)!

Of course, doctors aren't supposed to just measure. They're supposed to consume all that data, carefully consider it in context of the latest medical findings and the patient's history, and figure out if something's wrong. Computers can take on much of diagnosis and treatment work, and even do these functions better than an average doctor could (while considering more options and making fewer errors). What if you're a heart patient? It's a simple fact that most doctors couldn't possibly read and digest all of the latest 5,000 research articles on heart disease. In fact, most of the average doctor's

medical knowledge is from when they were in medical school, and cognitive limitations prevent them from remembering the 10,000+ diseases humans can get.

Computers are much better than people at organizing and recalling information. They have larger and less corruptible memories, remember more complex information much more quickly and completely, and make far fewer mistakes than a hotshot MD from Harvard. Contrary to popular opinion, they're also 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 the average or below-average physician? Remember, 50% of MDs are below-average! Not only that, computers have much lower error rates. Shouldn't we take advantage of that when it comes to our health?!

Technology makes up for human deficiencies and amplifies strengths – MDs and even other less trained medical professionals can do much more. Eventually, computers *will* replace 80% of what doctors do but amplify the doctor's capabilities, reducing misdiagnosis, in what they do by arming them with more complete, synthesized, and up-to-date data, all leading to better patient outcomes. Physicians spend too much time doing things computers can do, and we should give them more time for things that uniquely require human involvement, like providing patients "warm & fuzzies", comforting kids in pediatric care, making inherently subjective decisions that require empathy or a consideration of ethics, and providing a friendly ear for lonely patients. Some of these functions may not need medical school training at all, but rather draw on more empathic skills and could actually be done by non-MDs. Lifecom, an AI diagnostics engine company, showed in clinical trials that medical assistants using a knowledge engine were 91% accurate in diagnosis *without using labs, imaging, or exams*. Another clinical study by the same company demonstrated that greater than 75% of cases can be safely triaged to be treated by RNs, with the remainder handled by doctors. Another study at MassGen found that 25% of the time, a medical record for patients who wound up with 'high risk diagnoses' had 'high information clinical findings' before a physician eventually made the diagnosis — in other words, there was a significant delay that might have been avoided had a clinical decision support system been used to parse the notes!

Initially, many doctors will be against this transition and won't support it, but new technologies will make the receptive doctors much better at their job – quicker, more accurate, and more fact-based. There is a tremendous opportunity here in the influx of data that has never before been available. At first, computers will just help in decision support, starting by leveraging guidance from the very best doctors. Eventually (sometime in the next 10-15 years), computers will become better diagnosticians than your average doctor. Once we have a large enough dataset on different people in different situations, along with an addressable database of research studies, we will be amazed at how much better computers can do over today's patient outcomes. We'll be able to identify patterns and interactions among various areas of physiology in ways that weren't possible before, making the link between changes in one area of the body causing symptoms in another area. Doctors will struggle to keep up, but then will increasingly rely on these tools to make decisions. Over time, they will increase their reliance on technology for triage, diagnosis, and decision-making, so that we'll need fewer doctors and every patient will receive the best care. 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, a la "Dr. House", will be needed in direct patient contact. He can best serve as the trainer for the new "Dr. Algorithm", which we'll use to provide the diagnosis, while the most humane humans , (nurse practitioners or other medical professionals) will provide the care.

Eventually, computers will model and track your entire health state. They'll read your mood by analyzing your facial expression; gauge your social activity through the number of emails you sent, calls you made, or things you tweeted; track your mobility and activity from your GPS as Ginger.io* does in assisting mental health patients or Jawbone* UP band; and monitor your vitals through your food intake, galvanic skin response, heart rate, and skin temperature, among other things. No doctor could be this integrative. There are already numerous startups (and others soon to be launched) that plan to collect health data in a frictionless, easy way in order to create better baseline systems models of the body for patients. Others will do predictive analytics on that information and head off problems before they arise. Still more will suggest lifestyle approaches to improve the way people live. Some of this change is already happening, but this quiet rumbling of data-driven diagnostics will become an avalanche in the future, playing out similarly to the explosion of cellphones. Imagine today's systems, improved 10X by a decade or two of evolution and competition. Nobody expected cell phones to take over India in 2000, but now in 2012, few people remember that cell phones weren't expected to be universal (AT&T even killed off their mobile business in the 80s because McKinsey told them the total US market would be less than a million devices by 2020).

Device- and data-driven healthcare will also extend the reach of medicine. It will change public health, especially in the developing world. A country like India needs ten times as many doctors to serve everyone well, and that's not affordable. Most doctors there don't have access to the latest, expensive research journals and couldn't assimilate all the information contained in them even if they had the time, patience, and inclination to read them. A mobile phone could provide the needed testing and diagnosis to the remotest villages and at very affordable costs. This type of care is just not possible using today's medical school graduates, who tend to be clustered around cities.

Systems will start as clumsy toddlers and develop to maturity and efficiency!

Don't expect ace diagnosis systems overnight. They may start as seemingly minor point innovations or as clumsy-sounding systems not ready for prime time.

Imagine using a device like the AliveCor* iPhone case to take an ECG after every workout. What about every time you feel lightheaded or numb? What about every single morning, just like diabetics who measure their blood sugar multiple times a day? Now what if you could get an ECG case for free and do measurements for less than $1/test? If you're a heart patient, this device and others like it would capture a lot more information than your annual or semiannual ECG check at the doctor's office. Not only that, the office check will cost hundreds to even thousands of dollars, and what's more, you

probably wouldn't be exhibiting any symptoms during the in-person visit anyway if your condition is intermittent. What if you instead sent 500 ECGs to your doctor over the course of a year for less than it costs to get one ECG done in the hospital? What do you think the average physician would do with all that valuable data? He or she would have no clue what to do with it, which is why we'd need software to "auto-diagnose" the ECG. Today, most heart disease is identified only after patients have heart attacks. But imagine having preventative cardiac care, with every at-risk patient having an ECG every morning for less than a buck. After being trained to identify abnormalities, machine-learning software could *predict* episodes indicated by the ten or so ECGs out of that set of 500 that the cardiologist should pay attention to (before getting to a point computers can do an EKG read themselves for pennies), simultaneously making his or her job easier AND more effective. We could discover most heart disease well before a heart attack or stroke and address it at a fraction of the cost of care that would be needed following such a trauma. But we need a decade of data to be really good at it.

Many dermatology appointments could be handled by CellScope* which produces low-cost iPhone attachments for imaging skin moles, rashes, ear infections, and (in the future) your retina or throat. The resulting images, taken at home, could be processed by sophisticated algorithms running in the cloud to detect patterns that warrant closer inspection (e.g. SkinVision uses the fractal nature of patterns in a skin image to determine more accurately than most general physicians whether you have skin cancer). You might get a diagnosis a lot faster than it takes to get a doctor's appointment and then take your child to the clinic. And follow-up 'visits' could happen every six or twelve hours! A device like the Eyenetra* could give you an eye test and fit you for eyeglasses at little cost or hassle. Technology could handle diagnoses, lab orders, and writing prescriptions, asking for human assistance or input only when necessary.

Every metabolic process that has a volatile byproduct causes changes in your breath. Adamant* is a very risky startup that's attempting to produce a chip that can detect hundreds of gases in your breath. If you're asthmatic, it can measure the level of nitrous oxide and predict whether you're at risk of having an attack. If you're diabetic and have ketones in your breath, it can detect that too and tell you that your body is undergoing ketosis. It can detect if you have lung cancer and even tell you what *type* of lung cancer. It will even detect whether you're burning fat or sugars during exercise, because each results in different component concentrations in your breath. This little chip will do all this inexpensively, for far less than a big, expensive CT scanner that'll just tell you that you have a nodule in your lungs but can't tell you what kind of lung cancer you have. Eventually, you won't have a doctor stare at you and tell you that you look well; instead, your doctor will be able to look at the levels of hundreds of compounds in your breath and *know* whether you're well.

Speaking of looking well, a startup named Ginger.io* determines patients' mental health based on a variety of metrics. It can monitor your rate of emailing, tweeting, texting, and calling to gauge your social activity. Using motion sensors and your phone's GPS, it can even know if you're hiding in your bedroom, eating in your kitchen, or just staying in bed. By watching for changes in your behavior, it can tell how you're doing far better than a psychiatrist could possibly determine and actually calls your psychiatrist if you're in the danger zone of an episode. For example, detecting a behavioral pattern

change that's indicative of bipolar disorder could help us prevent shooting sprees of the type we've seen recently.

There are many other startups doing innovative things in healthcare. Proteus is helping address drug non-compliance, one of the biggest problems in medicine. They've designed a clever system that combines a pill sensor, a body patch, and a mobile app. They attach a tiny, ingestible sensor to pills that gets activated by stomach acids. When a patient takes the pill, the sensor sends a signal to the body patch, which then relays the signal to the app. This system will allow caregivers to remotely monitor patient adherence by individual, time-stamped pill consumption events. This is a far better solution than having your doctor base a diagnosis on the one-time blood test done in the clinic. Empatica uses sensors on patients' wrists to measure bio-signals that correlate with emotion. Imagine having a continuous and, more importantly, accurate and objective measurement of your emotions for a month instead of only the latest, biased description that you might give (or forget to give) to your hurried physician. Several companies are also improving remote monitoring and diagnosis in the clinical setting. AirStrip Technologies provides real-time vital signs to physicians' mobile devices. Sotera Wireless provides a battery-powered mobile device for monitoring vital signs. Agile Diagnosis and Lifecom are improving clinical decision-making by providing decision trees with probability-based outcomes for physicians at the point of care. These and other startups are forcing us to rethink healthcare from diagnostics to treatment.

This is only the beginning of the many generations of improvement that are likely to happen in the next two decades, but we have already begun to see meaningful impacts on healthcare. Studies have demonstrated the ability of computerized clinical decision support systems to lower diagnostic errors of omission significantly, directly countering the 'premature closure' cognitive bias. Isabel is a differential diagnosis tool and, according to a Stony Book study, matched the diagnoses of experienced clinicians in 74% of complex cases. The system improved to a 95% match after a more rigorous entry of patient data. Even IBM's fancy Watson computer, after beating all humans at the very human intelligence-based task of playing *Jeopardy*, is now turning its attention to medical diagnosis. It can process natural language questions and is fast at parsing high volumes of medical information, reading and understanding 200 million pages of text in 3 seconds.

Right now, Isabel and Watson require physicians to ask follow-up questions of the patient, a point of inefficiency and potential cognitive bias introduction. But, Lifecom has the medical text-parsing, knowledge-base generation, and runtime diagnostic capabilities of Watson, while also being able to automatically propose a follow-on question or test to eliminate candidates from the differential diagnosis list. This accelerates the diagnostic process and lessens biases introduced by the human element of question-framing. Watson also relies on machine-learned, purely statistical relationships among symptoms, findings, and causes, whereas a system like Lifecom adds to that ability the use of a medical ontology, the web of physiological, anatomical, and other concepts, as well as their interrelationships. It's too early to tell which approach will work better in the long term, but the point is that these systems will continue to evolve and, eventually, won't need physicians as intermediaries.

In the beginning, these point innovations will seem immaterial, but, when there are enough of them, they will integrate with each other and start to feel like a revolution. The medical devices and software systems of 2020 will be as different from today's computers as the car floor-mounted, multi-pound cell phones with bulky handset cords of 1986 are from today's iPhones!

Digital first-aid kits

The confluence of all these different sensing, monitoring, and communication technologies will naturally lead to the creation of 'digital first aid kits' that cost < $100 and can used at home between doctor visits. These kits will include mobile apps to help patients determine how serious a new medical problem might be, as well as monitoring devices that can track and analyze blood pressure, A1c levels, ECG waveforms, blood oxygen levels, skin/ear/ENT conditions, social interaction, and other indicators of how well patients are managing their diabetes, asthma, depression, and other chronic conditions. Clinical staff will help train patients on these devices and apps, and, using software, they'll help them interpret health trend data and provide recommendations during return visits.

Healthcare service stations of the future

The emergence of healthcare service stations, found in your local pharmacy or supermarket, will be another exciting phenomenon resulting from the proliferation of devices and data. These walk-in stations will be convenient — 75% of the US population lives within a 5-minute drive of a Walgreens. Right now, you can go to most of these stores and get some basic advice, receive a flu shot, or have your eyes checked, but not much else. In the near future, new technologies will enable many more services, providing high-quality care at low cost. In these centers, pharmacists and nurse practitioners will take patient histories, do simple exams, and prescribe medications or follow-on tests, aided by computerized diagnostics and telemedicine well before 2025.

Patients will be able to walk into these advanced healthcare stations without appointments and give their medical updates to personally-chosen avatars on private screens. These avatars will serve as healthcare concierges, elucidating symptoms, wirelessly uploading data collected by sensors on the patient's phone, suggesting tests (to be done by the patient using easy-to-use tools on-site or at home), and giving prescriptions through a clinical support system, whose ability to diagnose and recommend effective treatments will have been validated against the best general practitioners and specialists. Software will help RNs diagnose illnesses, recommend effective treatment options, and refer patients to specialists when necessary. In ten years, it's likely that genetic testing will be routine, extending diagnostic capability and allowing treatment selection based on genotype. Nurse practitioners will perform exams and take image scans using inexpensive, disposable tools (some of which might be part of the digital first-aid kit). Images will be analyzed in real-time by diagnostic software and transmitted to an on-call specialist, who pulls up patient information, stored in the cloud, and connects with the NP and patient using telepresence systems. Assessment and treatment apps will be prescribed, downloaded,

and installed on patients' phones, just as easily as prescriptions for cold medicine. And avatars will be accessible while patients are at home, providing real continuity of care. Ultimately, with ubiquitous walk-in clinics in the retail setting driven by powerful decision-support software, the distribution of healthcare professionals will change in parallel with the distribution channel. The pyramid will flatten, with many more RN-level professionals trained to effectively leverage software and interact with patients, and many fewer expensive MDs specialized to handle what will become the long tail of care.

The human element

Some of the critics of more automated healthcare argue that medicine isn't just about inputting symptoms and receiving a diagnosis; it's about building personal relationships of trust between providers and patients. The move towards sensor-based, data-driven healthcare doesn't mean that we'll get rid of human interaction. Serotonin administered by comforting humans and the placebo effect are only a few of many ways that complex interactions help recovery, and we shouldn't lose these tools (instead, let's understand them better, quantify them, and amplify them). Providing good bedside manner, giving comfort, and answering certain types of questions can often be handled better by a person than a machine, but you don't need a medical degree to do that (except in some specialty areas like surgery or research). Nurses, nurse practitioners, social workers, and other types of less expensive, non-MD caregivers could do this just as well as doctors (if not better) and spend more time providing personal, compassionate care. For a while, surgery and other procedures may require human doctors with greater knowledge than a nurse practitioner. The point is that not every function doctors do will be replaced, but the majority will be done in less time and with greater efficacy over the next two decades. At the other end, some surgical procedures requiring extreme precision (e.g. tumor irradiation) are best handled by surgical robots or robots and humans working in concert.

Consider hospital discharge for a moment. You might think this is a function that could (or should) only be performed by a human being. Interestingly, a Boston Medical Center study showed that patients preferred receiving discharge instructions from a computer instead of a human and appreciated the amount of time and information provided by the computer. Additionally, patients with lower health literacy reported a significantly greater bond with the computer compared to patients with higher health literacy, supporting the notion that computers can provide adequate emotional support in certain circumstances.

I'm not advocating the removal of the human front-end to patient care. I'm arguing that we should focus on building robust back-end sensor technology and diagnostics through sophisticated machine learning and artificial intelligence operating on clinical and non-clinical data in greater volumes than humans can handle, in order to create a much more comprehensive understanding of patients. What's more, these combined hardware/software systems will do better at follow-ups than over-burdened, time-constrained doctors. They'll recognize and flag exceptions in ECGs, respiration, movement, or anything else for caregivers to respond to in more targeted ways, actually improving provider-patient interactions and health outcomes.

A transition to automation has already happened in several other areas where we once thought human judgment was required. When you're on board that cross-country flight from New York to Los Angeles, most of the flying is being done by auto-pilot, not by a human (though a human is there for certain emergency situations). Investing was long considered a unique bastion of human judgment. Algorithmic trading now drives the vast majority of volume in the stock markets, with computers long ago replacing gesturing traders wearing funny-looking coats in the stock pits. Cars have been parking themselves and avoiding accidents using lane-assist for some time now, and Google already has a self-driving car that's had zero accidents driving more than 300,000 miles on normal streets (a much harder problem than automating many physician functions). Most humans couldn't drive that much without at least hitting a curb or two. Within a few decades, vehicles with steering wheels will become as quaint a concept as hand-cranked engines. David Cope at University of Santa Cruz has even developed software that has matched original Bach and a music professor composing in the Bach tradition in the music's "Bachness" to the chagrin of many.The same mental shift about human involvement and its gradual replacement by computers will also happen in healthcare. This would create a much more comprehensive understanding of patients and actually improve provider-patient interactions and health outcomes with more personalized treatment. In the end, physicians will be able to spend more, higher quality time with each patient, as if they had only 300 to manage, rather than 3,000 (though he will be able to manage 10,000+ patients with computer assist)! Caregivers will actually have MORE time to spend talking to their patients, making sure they understand, socializing care, and finding out the harder-to-measure pieces of information from patients because they will be spending a lot less time gathering data and referring to old notes. And they will be able to handle many more patients, reducing costs in the system.

The source of healthcare innovation

Where will all this innovation in healthcare come from? Some believe we have to work within the constraints of the medical establishment in order to advance it. I disagree. 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 reduces profits. That's why the system will most likely be disrupted by outsiders. Land-line phone call rates didn't decline until mobile operators changed the rules of 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, and doctors and hospitals are all invested in doing things the same way. If a hospital could cure you in half the time, would they be willing to cut their business in half? Pharma companies push marginally different drugs instead of generic solutions that may actually be better for patients because they don't necessarily want to cure you; they want you to be a drug subscriber and generate recurring revenue for as long as possible. They'd rather sell you a cholesterol-lowering drug than encourage you to eat healthier and reduce their own profits. If they permanently reduced your cholesterol, they would lose a customer! Psychiatrists in general have the same problem with incentives. They don't get paid to cure you; they get paid to treat you over multiple expensive therapy sessions that never end. What would you rather do as a therapist: have a steady stream of repeat patients that fill your hours or have

to attract new patients? The former has minimal patient acquisition costs. 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 $29 * ECG machine. The traditional players will lobby/goad/pay/intimidate doctors and regulators to reject these new devices, emptily claiming that they aren't "as good" and provide only 90% of the functionality (but at 5% of the cost). Expecting the medical establishment to do anything different is like expecting them to reduce their own profits. To be fair, these are generalizations and there are many great doctors and many ethical organizations and people. The point is that the incentives in healthcare make innovation from within extremely unlikely. Fortunately, it doesn't matter if the establishment tries to do this or not, because it will happen regardless. And it may start at the periphery, e.g. with the 40 million uninsured people in this country or the hundreds of millions of people in India with no access to any doctor. There aren't enough rural doctors in India and few of them have access to the New England Journal of Medicine or a CT scanner or even reliable electricity. But, most potential patients have cell phones. This shift in how healthcare is delivered will allow for less money to be spent on capital equipment, cutting health care 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 services. And, it will prevent simple things from getting worse before being addressed.

There has been much ado in the blogs about how Silicon Valley and outsiders don't understand healthcare and hence should not or cannot try to understand and innovate in it. As I explained above, and granting the rare exceptions, it is hard for insiders to innovate within a system, at least when it comes to radical innovation. That's not to say these Silicon Valley and other outsider "innovators" won't leverage the system or have partners and doctors from inside helping them. Lifecom's CEO is a trauma surgeon and teaches surgery and critical care on faculty. One of Proteus' founders is an MD, as is one of the founders of AirStrip Technologies. IBM Watson is working with . Many others work with in testing and studying their ideas and technologies. Most startups we are funding have MDs on their team and collaborate with other healthcare partners.

The reality is that healthcare has to move in this direction in order to make it affordable to everyone. There are many arguments and challenging questions in the blogs about this point of view. Some have answers, many reflecting naivety in understanding how technology increments and evolves, and many questions and criticisms don't have answers. But, just because an answer doesn't exist today does not mean it 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 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.

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 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 *can* have 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. But, there will be many good doctors willing to assist in this transition. Eric Topol (author of "The Creative Destruction of Medicine") 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% of them, will attempt this change and fail. But, a few of these outsiders will succeed and change the system. They 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 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. 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'll 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, ending 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 those of us who support entrepreneurs and companies that help create this change, most investments will be lost 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.

* A Khosla Ventures investment