The use of sustainable and scalable health care technologies in developing countries

Gari D Clifford¹⁻³

¹Department of Biomedical Informatics, Emory University, Atlanta, GA, USA; ²Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, GA, USA; ³Centre for Affordable Health Care Technology, Kellogg College, University of Oxford, Oxford, UK

Abstract: Although health care technology is progressing, in important ways access is dropping for many communities across the globe. Contributing factors include increasing awareness and options for diagnosis and treatment, and hence expectations, with consequential increases in complexity, training, and costs for health service delivery. Changes in wealth distribution, aging populations, and increasing burdens from chronic conditions also exacerbate these issues. Many attempts at delivering affordable health care technologies into resource-poor regions of the world have had limited impact, often with a focus only on cost or simplicity. However, with close attention to appropriate design considerations, including sustainable business practices and local cultural dynamics, there is the potential to revolutionize the field and deliver health care support technologies that empower, educate, and learn from the user. This article discusses the main concepts in designing relevant and sustainable health care technologies for resource-constrained communities and proposes structural ecosystem changes that may be required to achieve this goal.

Keywords: affordable health care, open source, scalable technology, developing countries, emerging economies

Introduction

This paper is an exploration of a set of possibilities for transforming the current status quo in the ways in which health care technology is used and delivered. These are a collection of observations that the author feels are important and pertinent, and so are inevitably biased. However, it is hoped that this provokes discussion and collaboration. Before delineating the factors to be considered in the design of “appropriate” or sustainable and scalable technology for relatively low-resource environments, it is important to define the term “technology”, particularly in the context of health care. In general, it is a piece of equipment (or software) which enables a user to perform a task more efficiently than without the technology. It is often thought to be something that is electronic in nature, although it can also be a bicycle that has been converted to drive a ventilator or a barrel that has been converted into a water-filtration device, for example. It should be noted, however, that appropriate or sustainable technology does not mean that the technology would be substandard in some way. In fact, given a modern and innovative approach, it could even be superior to existing and more expensive technology and lead to “reverse innovation”. This requires careful thought, and it should be noted that the human right to health includes affordability, accessibility, acceptability, and quality. When designing any tool for resource-constrained environments, it is important to consider several key factors:

Correspondence: Gari D Clifford
Department of Biomedical Informatics,
Emory University, 36 Eagle Row, Atlanta,
GA 30322, USA

1. The needs of the target population (both the users of the tool and the population for which it is intended to improve life).
2. The acceptability of the technology to the population, including behaviors, perceptions, and cultural sensitivities. Will the technology lead to an unintended consequence, such as overprescribing of a dangerous drug? More subtly, it is nontrivial to identify the way a population may react to a given product, potentially viewing low cost or unknown labeled products as inferior.
3. The availability and affordability of supplies and support infrastructure for the new technology. Many current business models for medical equipment rely on the bulk of the profit margin being generated from consumables. If such supplies are manufactured in another country, importation can be slow, expensive, and unreliable or prone to corruption (through informal import taxation).
4. The hidden costs and risks of using the technology. Every new procedure we add to someone's daily life has a direct cost to the recipient (eg, a bus fare to reach the deployment location) and an opportunity cost (eg, in terms of lost earnings).
5. The training needs for the technology to avoid misuse. Many equipment manuals are written for a high literacy audience, and they are often unavailable in nondominant languages.
6. The cost of the technology (both in terms of initial outlay and ongoing costs for supplies and maintenance) relative to alternatives, and the difference in potential impact.
7. The existence of downstream facilities and resources to deal with the output of the new technology. For example, it is no use having a system that can diagnose cancer, if the population thus diagnosed cannot afford or even locate appropriate treatment. The knowledge that the technology provides may even induce stress and exacerbate a condition, or lead to ethical problems as front line health workers are faced with new resource-allocation decisions. Dealing with techno-trash may also be another consideration.
8. The power requirements of the technology. If it is powered by mains electricity, it must be tolerant to daily power outages and power surges. If it is battery powered, the batteries must be locally available and affordable. How to deal with the disposal of the batteries should also be considered.
9. The networked capability of the technology and the ease with which it can be monitored remotely for data uploading, user interaction, and software revisions.
10. After market vigilance and the ability to perform recalls or update the technology.

11. Barriers to entry from national and international regulatory mechanisms and competing products whose distributors may have vested interests in retaining the status quo.
12. The ability for the technology to be self-sustaining and enable a business model to allow competitive use, without creating a colluding oligopoly, or a concentration of power in individuals who may be tempted by corruption.

This article addresses these points with some key examples of appropriate technologies that conform to these designs.

Key design considerations

Needs assessments

Many health care technologies are designed and built to satisfy assumed demands or needs that may not exist. It is true that some technologies are designed first and later find a niche in an unintended market, but this is the reverse design consideration. Rather, a problem or need should first be identified and a solution designed, for which technology may solve some of the issues. If technology can be avoided or minimized in some manner, then it should be. In order to implement a needs assessment of the target communities, extensive surveys and interviews with key stakeholders are required, which have been designed for the specific application in order to be unbiased and noncoercive. An external review of the survey instruments and other documents (such as by an ethical review board) is usually needed. When analyzing the behaviors of the populations, it is important to imagine any unintended consequences. For example, if a new piece of technology is created which can automatically read X-rays on a mobile phone, this may lead to an explosion in the use of X-rays. However, if the population does not have access to well-maintained X-ray machines, or X-rays are ordered inappropriately, then it could lead to significant overexposure to ionizing radiation in the population and an elevated incidence of cancer later in life. Even if the technology is relatively noninvasive, such as in the case of two-dimensional ultrasound, which may seem an attractive solution for screening pregnant women for example, there are sometimes hidden dangers, such as selective gender-based abortions.

Another key issue when considering the needs of the population are the perceived "wants", which may be in conflict with the external assessment of needs. The more an item costs, the more it is generally valued, and in fact can sometimes become a status symbol to indicate to family, friends, and acquaintances that we are successful. When this applies

to medical devices or treatments, however, it can mean that the technology in question can be left on display and may not see the use for which it was intended.

Training

When evaluating the training needed for a given technology, it is important to consider how appropriate it is for the education levels of the users. For example, a two-dimensional ultrasound device requires extensive training and a skilled individual to manipulate the probe and interpret the data it generates. The devices are consequently complicated to use and the output can be highly subjective. A large-scale training program may go some way to addressing this, but the idea of pyramid training has led to mixed results.¹ Not only does information become lost in the transmission (and local cultural practices and beliefs sometimes supplant the training), but if the training is successful, this can sometimes lead to talent migration away from the target population to resource-rich areas that pay more for such talents. Nevertheless, any technology requires some level of training. Ideally, if deployed with existing technologies (such as a mobile phone), the intuitive interface of the existing technology could be used to provide the training. For example, the interface of a smart phone could be used to deliver a simple logic branching guide for decision support. However, it is important that the training is appropriate for the local culture and levels of education, with minimal or no jargon, and matched for minimal literacy. For example, heavy use of pictograms and contextual help are important, since users are disinclined to read instructions. Appropriate automated training and telemedicine infrastructure could lead to an empowering of individuals with previously poor access to knowledge, and enable them to supplement income without actually being drawn away from the communities for which they are being trained (see “Scalability: networking and crowdsourcing” section).

Health care system and behavior evaluation

Once the needs and health-seeking behaviors of the population have been identified, it is important to identify where in the local system, a minimal change can be effected to create the maximum positive effect. This is easy to say of course, and much harder to evaluate when many of the variables are difficult to estimate. Moreover, it may be that there is no need for a change in technology, and a procedural change is required instead, such as retraining. It may even be as simple as improving drug stocks or changing the wording on some packaging. If there is no simple nontechnological solution, then the next step is to create a general product specification

from which several alternative prototypes can be developed and taken to the field to be beta tested.

It is also important to consider both the behavior of the health care worker and the patient in the system. Young et al² point out that behavior change is needed at multiple levels in a product cycle: 1) technology adoption, or willingness to use the technology; 2) engagement, or interest in continuing to use the technology; 3) health-related behavior change, or willingness to modify offline behavior based on relevant health-related information; and 4) health behavior maintenance, or behavior sustainment. Young et al² also note that human factors related to the design and use of mobile health technologies, such as user experience, visual and interface design, usability, and game mechanics affect behavior should be considered. In particular, theoretical approaches, including social, psychological, and behavioral (economics)-based theories, could be used to increase likelihood of health behavior change. They note, however, that little research has explored how these principles can help to inform mobile health design and suggest that this is an important future research direction.

Product design, prototyping, and cocreation

In parallel to beta testing various versions of the prototypes to assess user reactions, it is important to work with the target population to develop the technological solution. Traditionally, focus groups have been used to judge user's reactions to new products. However, it may be ineffective or misleading to use standard approaches to product design when the power structures and cultural dynamics are very different to our own. The very presence of a perceived figure of authority can lead to a user group being very reluctant to put ideas forward, and can be prone to accepting suggestions. It is therefore important to have several initial suggestions, potentially based on local products initially. An excellent example of this type of approach is taken by José Gomez at MIT's DLab Health, where local health care workers are encouraged to use local materials to solve problems, such as modifying a child's toy buzzer found at a local market so that it creates an alert when an intravenous drip bag is empty. Cocreation of the product is important, whereby the target population not only provides input, but also is able to actively participate in the design process. This generates a strong affinity with the product, a sense of ownership, and a deeper understanding of how to use it. Of course, this is not the end of the story, and a medical grade solution with standardized and high quality production has to be eventually created through more standard industrial processes. It is important to create

a professional looking (as well as behaving) product, so that it is not considered inferior to existing systems.

Supply chains

When considering the design and delivery of any technology, it is also important to think through the availability of supply chains, particularly for physical devices, both for initial device delivery, as well as for consumables (such as replacement electrodes), spare parts, and servicing. Equipment from developed countries usually requires proprietary cabling and even specific rechargeable (non-replaceable) batteries. Shipping internationally is not only expensive, but also slow and prone to corruption at the port of entry. Any device should be designed to use standard connectors (such as universal serial bus [USB]). It is often underappreciated that sensors/transducers, connectors, and cabling are often the weak point of any medical device design. USB cables are ubiquitous and low cost, so a natural choice to deliver power and data, although the micro-USB represents a step backward in robustness, with a flimsy female connection. In this respect, although larger, the mini-USB and full-sized USB ports are often preferable, although the cabling is increasingly less common. Care must also be taken to ensure that appropriate electrical isolation is maintained, particularly in the presence of cable splitting or other ad hoc use of the charging ports and there are now low cost USB isolation chips available.

Most markets in developing countries carry mobile phone accessories, compact discs, digital versatile discs, and batteries (of highly varying qualities). However, relying on nonrechargeable batteries is generally a bad design because they are often harvested for other applications or of questionable quality. It is far better to use a rechargeable battery such as those used in mobile phones. A generic design that can take any mobile phone battery is preferable. Harvesting of energy through motion, sunlight, or water flow is another option. However, the expense of such extra power supply devices and the susceptibility to mechanical failure or ad hoc repurposing means that parasitic energy harvesting through direct connection to other powered or charged devices (such as mobile phones) is preferable. This also drives down costs and complexity of manufacture. However, it is still important to develop supply chains for delivering devices. Without a vector for the hardware (or software), it is unlikely that a technology will take off. In recent years, major beverage companies have been joined by the telecom industry in the establishment successful global supply chains. By focusing on these emerging delivery vectors, it may be

possible to find conduits for disruptive technologies, which in time will develop their own supply chains, once a business ecosystem has established. Software also requires a supply chain for delivery, either via physical media, or more commonly via devices such as personal computing devices with Internet connectivity. Although the mobile phone industry is rapidly becoming the obvious vector for software distribution, the cost of connectivity should not be underestimated. Moreover, the proprietary nature of some suppliers can lead to significant add-on costs or barriers. For instance, Apple exerts complete control over which companies may deliver software through the Apple App Store and takes a significant percentage of the purchase price. This control extends to physical hardware too, where a flat rate of \$15 is imposed on every single device sold which needs to connect via the proprietary Apple USB connector. When the health care spending per capita can be as low as \$40 per capita, per annum, \$15 can represent a significant additional cost on the device. Recently, some innovators have designed around this issue by using the headphone connector rather than the USB port (see the iOximeter and Project Hijack, for example), but it provides a much lower power level and restricted bandwidth. Since Android is open source, it offers free access via the USB port, and far greater access to the on-board sensors and monitoring of the device hardware, and is therefore often a preferred platform, particularly in developing countries where Android adoption is significantly greater than iOS usage, particularly in the less wealthy strata of the population.

Scalability: networking and crowdsourcing

It is widely said that there is a plague of pilots in global health care technology, and to some extent this is true – we have many small-scale operations demonstrating the viability of a technology for a specific culture and medical system. However, it is slightly disingenuous to criticize a nascent field like this without addressing the barriers to scaling a technology. Many of the issues as discussed are key factors inhibiting the national or transnational scaling of technology. However, there is a clear opportunity for developing countries and emerging economies to “leap-frog” the antiquated and proprietary approaches to western medical technology, which has led to data silos, inoperability, and a hegemony that dominates the highly regulated and oligopolistic practices of western medicine. One of the key advantages of modern technology is the ability to push and pull information via connected telecommunication equipment, such as the mobile phone (or an embedded GPRS chip). The ability to reach every customer and “touch” every product on an almost daily basis without

significant effort from either the consumer or the supplier of the technology has far-reaching consequences. Imagine a product recall is needed on a million blood pressure devices. If the device is managed via a connected mobile phone app, then the supplier of the product can push a required update to the entire product range in a matter of hours. However, the bidirectionality of the data flow opens up an exciting possibility: the opportunity to add information and data from individual users to an ever-growing database, which can improve the performance of algorithms. Historically, companies have collected such databases without the input of the user, or through focused (and biased) individuals or user groups. However, with cellular-connected devices, a user can be canvassed for their opinion as to the veracity or accuracy of the information coming from the device. If they do not trust it, they can flag it as such, and have it passed to experts or “crowds” of nonexperts to label the data.

This latter innovation is very recent, and almost unheard of in medicine. The idea of consulting more than one doctor’s opinion seems to be the privilege of the rich, where we seek second, or more rarely, third opinions when our diagnosis seems suspect. However, most people do not have the education to know when to request more information, or the resources (time and money) to do so. Moreover, seeking second and third opinions leads to suboptimal and biased aggregations of knowledge. Even in the optimistic scenario that each opinion provides independent additional information, the aggregations of the individual opinions is performed by the patient who is unlikely to weigh them in an unbiased and accurate manner. However, with the advent of mobile and Internet connectivity, there is an opportunity to learn the expertise of multiple individuals in a given context to provide accurate weighing and improved diagnoses. This also can reduce error rates as well as provide automatic grading of a given professional’s talent in a given field.³ Moreover, it can provide piecemeal “crowd sourced” medical diagnoses, and still guarantee a given level of medical quality, and even a professional rating for the individual involved.

An example of such a system for cardiology can be found in Nam et al⁴ and Clifford et al.⁵ In essence, an unskilled worker can take an electrocardiographic recording of a patient in a low-resource environment, using a low-cost device connected to a smartphone.⁶ The phone then processes the information and tells the user if the recording was good enough for interpretation. If not, they are directed on how to take a better recording. (If the user happens to be in a low-connectivity region, another algorithm running on the phone can perform a preliminary screening for prevalent

arrhythmias and queue a referral request or rereview for later. Since the algorithm can provide a confidence estimate on the assessment, the patient can be informed about whether further analysis is needed from the “crowd”, and then send a text message at a later time with the diagnosis and referral). One of the key innovations here is the potential to combine local knowledge across geography and over time. By weighing together multiple opinions of humans (and algorithms) with different skill levels and specializations using a Bayesian framework,³ we can guarantee a given level of accuracy, and integrate out the errors and inaccuracies due to inter- and intrarater variability. Since the data can be pushed back out to other field workers as well as experts (because anyone who is not a white noise generator effectively adds some value to the overall diagnosis), the system can be used to pay (and train) less skilled individuals such as accredited health care activists or traditional birth attendants, etc. This can empower those with poor resources to earn money, self-educate, and gain accreditation and status to improve their socioeconomic status. A schematic of such a system can be seen in Figure 1.

Schemas, ontologies, patient identification, and interoperability

The identification of an individual and attaching it to data from multiple providers, across medical systems and through time has the potential to vastly reduce medical errors and reduce costs and inappropriate treatments. However, creating a unique and portable ID is nontrivial. Although many countries have national ID cards, the poorest and unhealthiest often exist outside such systems, and fraud is not uncommon. Even a fingerprint is nonunique as far as practical technology is concerned. A single fingerprint, even collected in ideal conditions, can lead to false identification 1 time in 10,000. Consider the implications when referencing this against India’s national fingerprint registry! In practice, it is far worse than this, with higher error rates on many low cost fingerprint scanners (such as you find on high-end phones).

Although many universal schemas and ontologies have been proposed to describe data, we have as yet no single “one size fits all” solution. Attempts to design top-down schemas for medical records have usually failed, and an organic bottom-up approach has often been more appropriate.⁷ OpenMRS is a classic example of such a design,⁸ which has been in use for decades in multiple countries around the world. The OHDSI/OMOP (Observational Health Data Sciences and Informatics/Observational Medical Outcomes Partnership) is starting to address this, but the model is mainly

Figure 1 Schematic diagram of a crowd-sourcing framework for medical data.

Notes: A consult can happen either in person with health care worker X, or remotely via a webserver with a health care worker or clinician. The mobile application uploads the data to a centralized server which notifies multiple individuals (both other health care workers with expertise and semitrained data reviewers) that there is medical data to label. As each individual submits their label (via an application or web browser) the webserver aggregates the labels using a Bayesian framework, and estimates the error in the aggregate label or diagnosis. When the error is low enough, a diagnosis is sent to the health care worker or patient with recommendations for treatment or further tests. Each data labeler is paid according to their accuracy.

Abbreviations: SMS, short message service; GPRS, general packet radio service.

developed for relational database-type systems, which are increasingly cumbersome. This does not address the issue of portability of medical data, which is becoming increasingly important for the individual as multiple health care providers deliver services to any given individual. Perhaps the logical solution to this is to have an integrated (and open) personal health record, which the owner (the patient) can choose to share in a limited fashion (both in terms of content and for a specific period of time) with certain health care providers.

The linking of data across databases and through time is another key issue. A patient's history over time is one of the largest factors in their diagnosis. Being able to provide an accurate history over time to any medical provider a patient wishes to see is a key design requirement. Perhaps the only way to address this issue is to enable patients to keep copies of all of their data (both locally on an encrypted partition

of their mobile devices and computers, as well as in the cloud through services such as Microsoft's HealthVault). One of the few open source schemas for doing this is Indivo X. Indivo is one of the children of the Guardian Angel project, a collaboration between Harvard and MIT that envisioned a web-based, automated guardian to manage health data and decisions.⁹

By allowing a patient to collect, store, and share their own medical data in a universal or interoperable schema via a secure method, there is the potential to add other key pieces of data from the individual's environment and behavior. An individual's "exposomes" (the key exposures in their life) such as pollutants, light levels, noise, temperature, humidity, etc, have significant effects on a person's well-being, both mentally and physically. In addition, it is well known that an individual's choices (such as diet, exercise, alcohol

consumption, and social behaviors) can also trigger or lead to poorer health outcomes. Such exposures and behaviors can easily be tracked through the use of social media and mobile phones.¹⁰ This growing field has enormous potential to inform and empower both societies (for public health decisions) and individuals (for modification of behaviors and health choices in day to day living).

In terms of ontologies, the unified medical language system (UMLS) provides a good starting point for describing elements in a schema for later cross mapping. However, there is a nondegeneracy problem in the UMLS (as in any complicated enough schema), which means a specific condition can be described by many different alphanumeric codes (or groups of codes). This makes a generic mapping approach impossible.

When it comes to both data schema and device interoperability, Western medicine suffers from the fallout of early adoption, where the linking of data grew on a system-by-system basis. This is partly because free markets have led to systems that were either not designed to transmit data from one system to another, or proprietary attitudes to data have led to a lack of interoperability. Although data models are emerging as standards, they are often cumbersome to implement in small projects, and so designers choose to ignore them in the prototyping/proof of concept stage, and the eventual implementation of standard data models and schemas becomes rare, due to the pressure to produce and other deadlines. Under “supply chains”, I mentioned standardized cabling (where appropriate). Of course, many people choose to design out cabling and opt for wireless data connections and charging. This introduces other issues, including battery life (both in the short- and long-term), environmental impact of disposing of batteries, and communication security and packet loss. Recently, Bluetooth low energy has improved on both energy consumption and hand shaking to reduce the profound inadequacies of Bluetooth connectivity. There are other ways to transmit data wirelessly of course, such as through the phone’s other sensors, each with their own advantages and disadvantages, including the near-field communication, WiFi, video camera, and microphone. For example, AliveCore chose to use a high pitch sound that the phone’s microphone can pick up. The reason for the design choice is not obvious, but it does work around the need for standard chipsets. However, it can be prone to acoustic and electronic interference (you have to switch off near field communication devices, and other normal sounds like running water appear to degrade the signal quality). However, the use of wireless technology does mean that a specialized locally powered chipset is needed on both the medical device and receiver,

and the user needs to remember to charge the batteries. This is often the Achilles’ heel of medical devices, since the user often forgets to charge or replace batteries in a timely manner, and when the device is needed, no replacement batteries or a convenient sustained power supply is at hand.

Regardless of which communication approach is chosen, standard (and open) interface protocols and data models should be adopted to reduce barriers to adoption and integration, as well as other issues related to legacy devices. It is common practice to keep communication protocols and data closed, because either the prototype was hacked together quickly, or the company wants to create barriers to competitors. However, this is dangerous because it means it is hard for a third party to ensure data integrity, allows the company to price gouge for extra software (which may not be supported later on), and the software is often provided as a compiled library (to protect the secret protocols or ensure the customer has to repurchase over and over again as operating systems evolve and communication standards change). All these factors increase both the price and the risk to the industry and the consumer, generating unnecessary work, slowing down innovation, and directly cutting into the health care budgets of the consumer and provider alike.

Despite this, there are good prospects for linking data on a limited basis. One potential for linking medical records is for leveraging the health information and behaviors concerning our close relatives and friends. We are becoming increasingly aware that not only do we inherit susceptibility to certain diseases from our close relatives through genetics, but also we can adopt or “transmit” unhealthy behaviors from/through our social network.¹¹ Linking our health information to our friends and relatives in a secure and anonymized manner may therefore be an important step in improving health. This then leads to the key issue of privacy and protected health information. Perhaps one of the largest concerns of a user of networked medical technology is the potential for someone to steal their information and use it to the detriment of the owner. Methods employed to abstract key indicators from our health and behavior data and share it with others (as and when we wish) is a key design consideration. However, this leads to an important issue of who owns the data, who can exploit it, and under what circumstances.

Privacy, ownership, and exploitation of medical data

Ownership of medical information is, rather oddly, a topic of debate. While it may seem obvious to anyone that anything recorded from you is owned by you (or entrusted to

your guardian if you are unable to make legal decisions for yourself), many people disagree with this, particularly those who may benefit from doing so. For example, some hospitals argue that they “own” any data they record from you or about you. More surprisingly, commercial manufacturers have argued that they own any data derived from the analysis of a patient’s data by their hardware or software. This is a slippery slope though, because all data is transformed and analyzed in some sense, and the idea of transferred ownership through transformation can easily be extended to claim the manufacturer of the device owns all data recorded by a machine or those data transduced by a machine. This is at odds with the general notion of ownership. For example, you own the rights to your own image in many countries and the publication of it without a grant to release the photo is prohibited. In Germany, for example, it is illegal to publish a photo (through any medium) where a recognizable person is the main subject; when the person is part of a large group (eg, demonstration) or not recognizable this is not forbidden, except if an endorsement of a commercial product is implied. Some countries are a little more liberal, but notably, if a person commissions (ie, pays for) a photo, then they own all the rights to it. It can be reasonably argued, therefore, that the patient is paying for the treatment (through national insurance, private insurance, or directly), went to hospital voluntarily (most of us would probably agree that we have implicitly agreed to be taken, if the need arises, to hospital in an emergency by opting in to a medical insurance scheme), and therefore commissioned the required recordings through a proxy (the doctor who is caring for them). The patient therefore owns their data. However, this does not necessarily dictate who may use the patient’s data.

In some situations, you implicitly grant a third party a royalty-free license to use your data, if, for example, you are in a public place where notices are posted that recordings are being made, and you can reasonably be expected to be on camera. In medicine, it is possible to obtain blanket ethical approval for this in hospitals too, and I would argue essential to push forward research. As long as my data cannot be tied back to my private or public identity outside the hospital (at least without enormous effort [It is well known that it is almost impossible to ensure perfect deidentification of data if external data sources are used]), then I see it as my duty to science to allow researchers to analyze my data. Moreover, it is the duty of the scientific and medical establishment not to waste my data and to use it to improve health care. I would therefore argue that any data recorded through an approved medical system, means that an implicit royalty-free license

has been granted to use the data for nonprofit research, as long as the institution collecting the data can demonstrate they have taken sufficient steps to guard the privacy of the patients in their databases.

This naturally leads to the question of cyber security. With the positive gains in information access, there comes the parallel dangers of data hacking and the release of compromising data, which can be all the more severe in certain cultures. For example, in societies where certain diseases are stigmatized (such as AIDS), reprisals (sometimes violent) can occur. The Health Information Trust Alliance¹² points out those such risks can be minimized by appropriate daily vigilance and rapid response.

Why is all this of interest here? Well it directly pertains to the cost of health care, business models, licensing, and intellectual property (IP). The standard paradigm in medical technology is to think up a “novel” technology, “application”, or “system design”, patent it, and license to a company or build a start-up company around it. The patent is considered IP, but this label is somewhat a misnomer. The idea behind a patent is to grant the inventor a limited exclusive license to make money from it for a period of time. This is the incentive for the inventor to share their idea with the world. As such, the patent is meant to fully explain how to implement their invention. However, over time patents have become enveloped in legal jargon and very few are comprehensible to “someone skilled in the art”. Many patents, when explained in plain language or scientific terms, are in fact obvious to “someone skilled in the art” though. This leads to a rather perverse situation in which obvious (and often overreaching) concepts are patented, and multiple companies infringe on each other’s immense portfolios of patents. A cold war standoff of patent lawyers results and the barriers to entry for small and innovative companies becomes immense. Submarine patents make the situation worse, where individuals or companies throw in speculative patents, without doing the hard work of actually proving the technology or the idea has any potential (or even creating a prototype), and then wait for someone else to do the hard work. They then surface after several years, and use a battery of well-paid lawyers to slice off a pound of flesh. Such a system inhibits innovation and leads to an oligopolistic and stagnant industry.

Regulation, open access, and open source

Regulatory considerations are often a significant issue for technology, and following international standards (such as those from the Institute of Electrical and Electronics Engineers, ISO9000 certification, US Food and Drug

Administration approval, or Conformité Européenne marking) is an established route to international acceptance. However, this can add significant development cost to the technology, and the standards, developed for historically Western markets, do not necessarily fit with the emerging market demands. Overt or poorly enforced laws and regulation can even lead to corruption.¹³ Moreover, the regulations usually only ensure safety, rather than accuracy, and the responsibility for accuracy is off-loaded to the user, who is generally expected to be a trained professional. An alternative approach may be to simply require open sourcing of the technology. This concept usually scares investors and confuses people into thinking that you are giving away the IP though. However, as Richard Stallman and the Free Software Foundation point out, “free” or “open source” licensing means free as in liberated, rather than as in “free beer”. The technology can be used in a not-for-profit situation, but as soon as you wish to profit from it, you must pay a royalty to the owner of the software. Licensing open source software (or providing support for such software), while not common, is still a viable business model, as Red Hat Software has demonstrated.

As Malkin et al point out though, you should not expect a large fee for licensing out IP in developing countries.¹⁴ Perhaps then, it is best not to waste resources filing for IP in the first place. It is entirely possible to file (and be awarded) patents on open source technology, yet still publish open source code for others to evaluate it (for research purposes only). The authors/inventors are still protected from anyone using it for profit without licensing the technology. Moreover, the authors/inventors are entitled to create closed source versions (by changing the license, since they are the owner and free to do so) and create exclusive licensing deals. Perhaps this may obviate the paranoia many people have about IP being stolen in developing countries. In order to increase competition and reduce market prices, it is also important for the application protocol interfaces to be set as open as possible. By creating the lowest cost devices, the companies that sell the devices must compete on service delivery, and so maximize the benefit to the consumer, rather than the profits to the company (at the user's expense).

Sustainability

Sustainability is generally thought to be any system or framework, which does not require external investment (after an initial start-up investment), and that lasts from one generation of people to the next (~20 to 30 years). However, technology evolves much more rapidly than this, so the time period should be thought of as much less, perhaps only 10 years

after the initial funding runs out. In other words, to sustain a technology, there needs to be a motivation for individuals to make a profit from it (otherwise most will seek alternative outlets for their time and resource allocation). In resource-constrained environments, a service delivery model (where the device is free or at cost, but the onward service to support it costs money) or an amortized hire-purchase arrangement is generally more appropriate. A classic example of this is the mobile phone, and this is one of the reasons it has expanded so rapidly.

In fact, the mobile phone is a great example of technology that seems to fill the criteria discussed. However, the rise of the cellphone should not be entirely attributed to just the business model. Illiterate or semiliterate users can operate a phone, informal supply chains for the phone and its services and parts exist, it runs off batteries that are commonly available, and it provides a service that is fundamental and useful (communication), empowering the user. Moreover, the instant connectivity that the mobile phone provides means that key issues around longitudinal medical records can be solved, leapfrogging the antiquated systems in the developed parts of our planet. The opportunity to uniquely identify an individual, track them through their life (via a portable digital identifier such as an email address), and store the data in cloud-based record systems that facilitate interoperability provides opportunities to uncover health and illness trends over the course of an individual's lifetime, and on population levels.

However, in health care it is particularly important to consider how the service is delivered and what value it adds to the consumer's life. To date, very few devices have managed to do so, and most are used for a few months, at best, and then left to become obsolete rapidly. The recent collapse of Zeo, the makers of an electroencephalographic-based sleep staging device, is an example in point. Zeo's CEO admits that their failure was the inability to provide the consumer with “useful and actionable information”, rather than simply “data”. If the consumer does not know what to do with the data to improve their life, then the device is rapidly discarded.

Theoretical frameworks

It is also important to note that there are at least five key theoretical frameworks to describe how theoretical thinking can enrich the study of care coordination. These included the Andersen Behavioral Model, the Donabedian Quality Framework, the Organizational Design Framework, the Relational Coordination Framework, and the related Multilevel Framework.¹⁵ However, the field of resource-constrained health

care is relatively nascent and there is scant evidence for the application of such approaches. Notably, Chib et al¹⁶ conducted a detailed review of 53 mHealth studies addressing one of the three stages of the pathway: inputs, mechanism, or outputs. The main types of interventions studied were related to data collection, consultation between health workers, appointment reminders for health workers and patients, health promotion, medication reminders, health information for patients, and test reminders. The review found very few studies (only six) that have provided any theoretical understanding of adoption and appropriation of technological introduction that produces measurable health outcomes. As a result, there is a lack of a dominant theory in the field, or measures of outputs relevant to making policy decisions.

Entrepreneurship and disruptive innovation

The reader might be wondering where this article is going by now, given that it is in a Journal of Innovation and Entrepreneurship in Health. All of the previous discussions are aimed at describing what I believe are the key elements required in any successful health care technology in emerging markets and developing countries, and where the key barriers and failure points are likely to lie. This in turn has a direct bearing on what types of innovation models are likely to be successful. In Clayton Christensen's book, the *Innovator's Prescription*,¹⁷ he identified three revolutions that are needed for a successful disruptive innovation:

1. Technology enabler: "routinize" previously complicated task.
2. Business model innovation: affordable and convenient.
3. Value network: companies with disruptive mutually reinforcing economic models.

This is well argued, and that this maps well to my own thinking over the last 10 years. His example is to move computing from industrial to personal. Perhaps a more modern recent example is Apple's iPod. At first thought, it is easy to think that Apple revolutionized the phone industry with the iPhone, but in fact Nokia beat them by 2 years. The Nokia N95 (running the open source Symbian OS and sporting a 5 megapixel camera with a Carl Zeiss lens and beautifully designed one-click shutter) emerged in 2005, 2 years before the iPhone. Apple's real revolution was the iPod, which took advantage of the disruption of the music industry by mp3 and peer-to-peer Internet technology (like Napster). It was not even the iPod device itself that drives the revolution (other cheap USB mp3 players had been coming out of China before the iPod arrived). Apple managed to revive

the music industry by offering a revolutionary way to buy music – with a simple one-click approach. I do not think this simplicity can be underestimated, especially in a world where the choice of products is effectively uncountable and loyalty to a reliable and trustworthy service provider can be more important than price. Of course, this is only true as long as the differences in price remain small. Humans are still very price sensitive and high street consumers will often opt for a cheaper medical product (purchased online or from high street pharmacies) that provides unreliable information, over an accurate but higher priced device. The fact is that it is very difficult for nonexperts to assess just what is accurate enough for their own condition. In the next section, I detail a possible framework that can address this issue, empowering the patient and health care worker to become better educated, make better choices, and generate additional income using just their experiences and their mobile phone (or any Internet connection).

A sustainable health care delivery business model to empower patients and health care workers

As indicated earlier, one of the key issues in health care is that of medical uncertainty through subjective judgments. Although humans can outperform automated algorithms in many scenarios (such as image segmentation), there are classes of tasks for which algorithms are better, such as repetitive tasks or those that require constant vigilance and/or fast reactions (eg, implantable cardio defibrillators). Moreover, humans are often influenced by other opinions, or the fear of making a mistake (and thus bias themselves toward low-risk scenarios). This leads to both a bias (due to opinion influences and differing schools of thought/training) and variance (due to changes in vigilance and the inherent unrepeatability of the task). This latter issue, that medical diagnostics generally involves subjective or semisubjective rules of thumb that lead to large intra- and interobserver variation, is often addressed either by implementing systematic check lists, or by seeking a second or third opinion, neither of which address the bias and variance problem correctly. There are several frameworks that have been developed which attempt to address these issues in a systematic fashion, with a notable contribution from Warfield et al¹⁸ who developed a framework for a binary labeling problem. However, few tasks in medicine are binary (even if the "final" decision may be binary – treat or not treat). Recently, we proposed a Bayesian framework that correctly allows for both the bias and variance in a group of annotators for a continuous labeling task.^{3,19}

This has several key effects and advantages over traditional diagnostic routes:

1. The medical label or diagnosis is at least as accurate as the best contribution, and usually more accurate, reducing both bias and variance for any given individual diagnosis. Moreover, no individual or algorithm contributes overwhelmingly to the result, with the “best” diagnosis always being guaranteed.
2. This has the knock-on effect that, if the health care worker (or annotator) knows this, they can be much more confident of adding their opinion to the “crowd”, without worrying that they might be personally responsible for a poor diagnosis. One might argue that this could lead to sloppy behavior. However, since the annotator is paid as a function of their accuracy, poor performers will quickly be disabused from sloppy behavior. By feeding back immediate scores or pay, the annotator has a chance to adjust their approach on the next annotation.
3. We can not only identify poor performers without any gold standard diagnosis, but also identify on what types of data they perform poorly, and provide extra training on these types of data (or choose not to send them such data for annotation).
4. The algorithm can provide confidence intervals on the diagnosis based on the interrater disagreement level, thereby sequentially optimizing the number of requested experts based upon the price a consumer can (or is willing to) pay and the level of accuracy needed.
5. Since the local health care worker (or patient) labels the data immediately, and the remote user only has a limited amount of time to respond to the offer of “employment” on a particular task (because they are in competition with others), this creates an urgency to accelerate diagnoses and allows us to identify how long a diagnosis might take, given the underlying data qualities and the pool of differently skilled “experts” available. If a more urgent diagnosis is needed, then the patient (or attending health care worker) can opt to pay more to increase the incentive to label the data. For example, in a recent experiment using the Amazon Mechanical Turk, we found significant jumps in annotation speed as we changed from 1 cent an annotation to 3 cents.¹⁹
6. Since the user is given immediate feedback, they are effectively receiving training on-the-fly and their skills will develop over time.
7. Algorithms and humans can be combined in an unbiased manner. The general paradigm in medicine is to overread previous results and adjust where necessary, but this leads to a strong bias to accept the first diagnosis.
8. Even in unconnected remote areas, several algorithms can be run on the phone and voted with the human annotation to improve the diagnosis and provide an initial triage at the point of assessment.
9. Since the algorithm can provide confidence intervals on the diagnosis based on the interrater disagreement level, the user can be informed about whether the diagnosis is sufficiently accurate or whether they need to wait for a later update from the cloud-based crowd (via text message for example) once connectivity is restored and the data are synced.
10. At some point, the user will develop enough skill to start flagging diagnoses as suspicious. Flagged data (or those with low confidence in the label/diagnosis) can be pushed forward for further expert analysis or quality control.
11. The database continues to grow over time in both size and quality, with accurate labels being attached to data. Algorithms can be constantly retrained and pushed back to the phone to improve the overall diagnostic accuracy of the system.

Although others have suggested the use of “metered” or pay-as-you-go health care using cloud environments,²⁰ the mechanism of apportioning exact payments based on the exact information each individual can supply to a diagnosis has yet to be suggested in health care. Point #4 might seem highly controversial, since we strive to provide the same quality of care for everyone. However, outside a nationalized health care system (and even assuming no geographic and education-based disparities to affect choice making), it is unrealistic to expect this, and perhaps more efficient in terms of resource allocation for the “consumer”, where they may value things other than disability-adjusted life years or quality-adjusted life years. There are some things that should not be addressed this way of course (such as herd immunity from vaccinations), but that is a matter of regulation and government policy.

Conclusion

Due to the focus of this article, its content provides a speculative look into the future, with which many could argue, and so concrete conclusions are hard to state. However, this article presents a series of ideas that have been demonstrated at various levels from pilots to large-scale randomized clinical trials. If the article describes ideas that do in fact lead to a market revolution, then it will be perceived as prescient, or rather naïve and perhaps amusing otherwise. However, the future is formed by testing multiple competing ideas and

letting natural selection (or deep pockets/government intervention) decide which approach wins out.

In the design discussions, it can be seen that the one key “technology enabler” is the application of diagnostics through a wireless device like the cellphone. By leveraging the input from multiple individuals and aggregating them in an intelligent manner, we can simultaneously increase quality and empower the poorly trained health care workers in the most remote regions of the world. We have developed the necessary technology, but the cultural, social, and economic issues cannot be ignored. They are in fact vital to the success of any technology, and many a product has lost out to inferior technology because of these factors.

This article finishes with a restating of the focus to explore some of the possibilities the author perceives to be promising and exciting, and bounce the ideas off the community in the hope that it resonates and someone wants to take it forward (or collaborate to take the ideas forward). The earlier discussion is by no means exhaustive, and of course there are likely to be competing viewpoints that we need to test in parallel.

Acknowledgments

The author wishes to thank Lisa Stroux for making Figure 1 look much better than my original “artwork”, as well as Rachel Hall-Clifford and the anonymous reviewers for their helpful suggestions. The content is solely the responsibility of the author and does not necessarily represent the official views of any organization.

Disclosure

The author reports no conflicts of interest in this work.

References

1. Bangdiwala SI, Fonn S, Okoye O, Tollman S. Workforce resources for health in developing countries. *Public Health Rev.* 2010;32:296–318.
2. Young SD, Holloway IW, Swendeman D. Incorporating guidelines for use of mobile technologies in health research and practice. *Int Health.* 2014;6(2):79–81.
3. Zhu T, Johnson AEW, Behar J, Clifford GD. Crowd-sourced annotation of ECG signals using contextual information. *Ann Biomed Eng.* 2014; 42(4):871–884.
4. Nam DS, Youn CH, Lee BH, Clifford GD, Healey J. QoS-constrained resource allocation for a grid-based multiple source electrocardiogram application. *Lecture Notes Comput Sci.* 2004;3043:352–359.
5. Clifford GD, Arteta C, Zhu T, et al. A scalable mHealth system for non-communicable disease management. In: *IEEE Global Humanitarian Technology Conference (GHTC'14)*. Silicon Valley, CA: IEEE; 2014.
6. Celi LA, Sarmenta L, Rotberg J, Marcelo A, Clifford G. Mobile care (Moca) for remote diagnosis and screening. *J Health Inform Dev Ctries.* 2009;3(1):17–21.
7. Clifford GD, Blaya J, Hall-Clifford R, Fraser HS. Medical information systems: a foundation for healthcare technologies in developing countries. *Biomed Eng Online.* 2008;7(1):18.
8. Seebregts CJ, Mamlin BW, Biondich PG, et al. Human factors for capacity building: lessons learned from the OpenMRS implementers network. *Yearb Med Inform.* 2010;13–20.
9. Szolovits P, Doyle J, Long WJ, Kohane I, Pauker SG. *Guardian Angel: Health Information Systems*. Technical Report MIT-LCS-TR604. Cambridge, MA: Massachusetts Institute of Technology; 1994. Available from: <http://groups.csail.mit.edu/medg/projects/ga/manifesto/GAtr.html>. Accessed October 14, 2015.
10. Vrijheid M, Slama R, Robinson O, et al. The human early-life exposome (HELIX): project rationale and design. *Environ Health Perspect.* 2014;122(6):535–544.
11. Ackland M, Choi BCK, Puska P. Rethinking the terms non-communicable disease and chronic disease. *J Epidemiol Community Health.* 2003;57:838–839.
12. HITRUST Alliance. Cyber threat intelligence and incident coordination center: protecting the healthcare industry from cyber-attacks. Health Information Trust Alliance (HITRUST), July 2014. Available from: <http://hitrustalliance.net/content/uploads/2014/07/HiTrustC3Datasheet.pdf>. Accessed October 14, 2015.
13. Laffont J. *Regulation and Development, Federico Caffè Lectures*. Cambridge: Cambridge University Press; April 18, 2005.
14. Malkin R, Von Oldenburg Beer K. Diffusion of novel healthcare technologies to resource poor. *Ann Biomed Eng.* 2013;41(9):1841–1850.
15. Van Houdt S, Heyrman J, Vanhaecht K, Sermeus W, De Lepeleire J. An in-depth analysis of theoretical frameworks for the study of care coordination. *Int J Integr Care.* 2013;13:e024.
16. Chib A, van Velthoven MH, Car J. mHealth adoption in low-resource environments: a review of the use of mobile healthcare in developing countries. *J Health Commun.* 2015;20(1):4–34.
17. Christensen CM, Grossman JH, Hwang J. *The Innovator's Prescription: A Disruptive Solution for Health Care*. New York, NY: McGraw-Hill; 2008.
18. Warfield SK, Zou KH, Wells WM. Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation. *IEEE Trans Med Imaging.* 2004;23(7):903–921.
19. Zhu T, McKay JL, Payne A, Ting LH, Clifford GD. Crowdsourced annotation of EMG onset times in healthy individuals and Parkinson disease. Presented at: International Society for Posture and Gait Research (ISPGR) World Congress, June 28–July 2, 2015, Seville, Spain.
20. Sultan N. Making use of cloud computing for healthcare provision: opportunities and challenges. *Int J Inf Manage.* 2014;34(2):177–184.

Innovation and Entrepreneurship in Health

Publish your work in this journal

Innovation and Entrepreneurship in Health is an international, peer reviewed, open access journal publishing original research, reports, reviews and commentaries on innovation and entrepreneurship in health. Special focus will be given to the theory, process, and practice of innovation and entrepreneurship by individuals and organizations

Submit your manuscript here: <http://www.dovepress.com/innovation-and-entrepreneurship-in-health-journal>

Dovepress

within the health care context globally. The manuscript management system is completely online and includes a very quick and fair peer review system, which is all easy to use. Visit <http://www.dovepress.com/testimonials.php> to read real quotes from published authors.